
Coordinated Double Machine Learning

Nitai Fingerhut¹ Matteo Sesia² Yaniv Romano¹

Abstract

Double machine learning is a statistical method for leveraging complex *black-box* models to construct approximately unbiased treatment effect estimates given observational data with high-dimensional covariates, under the assumption of a partially linear model. The idea is to first fit on a subset of the samples two non-linear predictive models, one for the continuous outcome of interest and one for the observed treatment, and then to estimate a linear coefficient for the treatment using the remaining samples through a simple *orthogonalized* regression. While this methodology is flexible and can accommodate arbitrary predictive models, typically trained independently of one another, this paper argues that a carefully *coordinated learning* algorithm for deep neural networks may reduce the estimation bias. The improved empirical performance of the proposed method is demonstrated through numerical experiments on both simulated and real data.

1. Introduction

A fundamental problem in data science is to estimate from high-dimensional observational data the effect of one variable, the *treatment*, on an outcome of interest, accounting for all covariates. For example, one may want to understand the effect of temperature on atmospheric pollution using measurements from weather stations. Alternatively, in a retrospective medical study one may want to estimate the relation between drug dosage and changes in blood pressure, accounting for differences in patients characteristics and clinical histories. These questions have traditionally been addressed through multivariate parametric models. However, it is difficult to obtain exact inferences if the data include many covariates. Further, simple modelling assump-

tions may not be appropriate if the outcome depends on the covariates in an unknown and possibly complicated way.

Sophisticated machine learning (ML) algorithms, such as deep neural networks and random forests, are playing increasingly relevant roles in estimation applications, even though they were originally designed for prediction tasks. Although generally difficult to analyze theoretically and often hard to interpret, ML models are sometimes appealing for estimation because they can flexibly account for complex relations between the covariates and the outcome. For example, the output of ML models can be translated into treatment effect estimates using an approach known as *double machine learning* (DML) (Chernozhukov et al., 2018). DML enjoys desirable theoretical properties—it leads to approximately unbiased estimates of the treatment effect—and it is becoming widely applied in many fields, including public health (Chernozhukov et al., 2021; Torrats-Espinosa, 2021), finance (Feng et al., 2020), and economics (Knaus, 2020; Semenova et al., 2017).

The goal of this paper is to improve the estimation performance of DML by adapting the learning algorithms used to fit its two underlying predictive models—one tasked with predicting the outcome and the other with predicting the treatment given the covariates. Concretely, a novel loss function is introduced to jointly fit the two ML models as to minimize the final DML estimation bias, and a cross-validation technique is developed to tune the relevant hyperparameters. The key idea is to coordinate the two predictive models so that the impacts of their respective errors approximately cancel out, in contrast with the typical practice of training the two models separately. Concretely, this idea is explored in this paper within the context of deep neural networks.

1.1. Problem Statement and Technical Background

Consider n data points (X_i, D_i, Y_i) drawn i.i.d. from some distribution assumed to follow a partially linear model (Robinson, 1988):

$$Y_i = g(X_i) + D_i\theta_i + U_i, \quad D_i = m(X_i) + V_i. \quad (1)$$

Above, the outcome $Y_i \in \mathbb{R}$ may depend on the covariates $X_i \in \mathbb{R}^d$ through the unknown function g . The treatment $D_i \in \mathbb{R}$ may depend on X_i through the unknown function m , and it affects Y_i linearly. The inferential target is the

¹Departments of Electrical and Computer Engineering and of Computer Science, Technion, Israel ²Department of Data Sciences and Operations, University of Southern California, CA, USA. Correspondence to: Nitai Fingerhut <nitai@campus.technion.ac.il>.

average treatment effect $\theta := \mathbb{E}[\theta_i] \in \mathbb{R}$, while m and g are *nuisance* functions, in the sense that their estimation is not of direct interest. Note that θ_i is assumed to be independent of X_i . The noise terms U_i and V_i have mean zero and are independent of one another as well as of X_i . In the absence of unmeasured confounders, and with the understanding that the values of X and D are determined prior to that of Y , one may think of interpreting θ as a causal parameter. However, as these assumptions may not always hold in practice, the methodology presented in this paper will not allow us to make rigorous claims of causal inference from real data.

Double Machine Learning. In DML (Chernozhukov et al., 2018), the observations are randomly divided into two disjoint subsets, $\mathcal{I}_1, \mathcal{I}_2 \subseteq \{1, \dots, n\}$. The data indexed by \mathcal{I}_1 are utilized to train two ML models. The goal of one model is to predict D given X , and its regression function is denoted by $\hat{m}(X)$. The goal of the other model is to predict Y given X or, equivalently, to compute an estimate $\hat{\ell}(X)$ of

$$\ell(X) := \mathbb{E}[Y | X] = g(X) + \theta m(X). \quad (2)$$

This is the Robinson-style estimator (Robinson, 1988), but it is not the only way to implement DML. An alternative version of DML is based on an approximation of $g(X)$ instead of $\ell(X)$, although $\ell(X)$ is easier to reconstruct from the data using a fully non-parametric prediction model for $Y | X$. Nonetheless, the methods presented in this paper could be extended to the version of DML based on $g(X)$. The second data subset, indexed by \mathcal{I}_2 , is used to compute residuals \hat{U}_i and \hat{V}_i , for all $i \in \mathcal{I}_2$, defined respectively as:

$$\hat{U}_i := Y_i - \hat{\ell}(X_i), \quad \hat{V}_i := D_i - \hat{m}(X_i). \quad (3)$$

Finally, the DML treatment effect estimate is obtained as:

$$\hat{\theta} := \left(\sum_{i \in \mathcal{I}_2} \hat{V}_i^2 \right)^{-1} \sum_{i \in \mathcal{I}_2} \hat{V}_i \hat{U}_i. \quad (4)$$

This strategy is summarized in Algorithm 1. Its strength is that, under certain assumptions it leads to an approximately unbiased $\hat{\theta}$; i.e., $\mathbb{E}[\hat{\theta}] \approx \theta$. In particular, it mitigates through data splitting the possible bias due to estimation errors in \hat{m} and $\hat{\ell}$, which are typically impossible to recover exactly.

Algorithm 1 DML

Input: data $\{(X_i \in \mathbb{R}^d, D_i \in \mathbb{R}, Y_i \in \mathbb{R})\}_{i=1}^n$.
Split $\{1, \dots, n\}$ into two disjoint subsets: $\mathcal{I}_1, \mathcal{I}_2$.
Train \hat{m} and $\hat{\ell}$ on \mathcal{I}_1 to estimate m and ℓ in (1) and (2).
Estimate θ in (1) with $\hat{\theta}$ in (4) using the data in \mathcal{I}_2 .
Output: treatment effect estimate $\hat{\theta}$.

One disadvantage of data splitting is that it reduces the number of observations effectively available for training the predictive models. To overcome this issue, Chernozhukov

et al. (2018) suggested *cross-fitting*: the entire DML procedure can be repeated twice (or even more times) swapping the roles of \mathcal{I}_1 and \mathcal{I}_2 (or repeatedly dividing $\{1, \dots, n\}$ into different disjoint subsets), and then all the resulting estimates $\hat{\theta}$ are averaged. For simplicity, cross-fitting will not be applied here because it concerns an issue that is orthogonal to the focus of this paper; however, it could similarly improve further the performance of the proposed method.

Preview of the proposed method. While DML is relatively robust in theory (Chernozhukov et al., 2018), it is not guaranteed to produce accurate estimates of θ in finite samples. In fact it can be very biased if the black-box regression models \hat{m} and $\hat{\ell}$ do not accurately approximate m and ℓ in (1). Unfortunately, this issue is not uncommon in practice if the data are high-dimensional or if the true m and ℓ have a complicated form. This is exemplified in Figure 1, which visualizes the results of semi-synthetic numerical experiments based on financial data borrowed from Chernozhukov & Hansen (2004); see Section 4 for more details. Here, the sample size is 2000 and DML is applied using two deep neural networks to estimate m and ℓ . The experiments are repeated 250 times with different realizations of the randomness in the semi-synthetic data to show the distribution of the treatment effect estimates. The results demonstrate that $\hat{\theta}$ is biased, regardless of whether $\theta = 0$ or $\theta > 0$.




Figure 1. Distribution of DML effect estimates on semi-synthetic data. Red: standard DML. Blue: proposed coordinated DML. Both methods are applied over 5000 experiments with a sample size of 2000, using deep neural networks as base predictive models. The outcome is a simulated measure of net financial assets, and the treatment variable measures 401(k) eligibility. The dotted vertical lines indicate the homogeneous treatment effect (left: 0, right: 10).

Figure 1 also previews the performance of our novel method, which differs from the standard DML insofar as it is based on carefully trained neural networks designed to estimate m and ℓ in such a way as to approximately minimize the downstream bias of $\hat{\theta}$. This is in contrast with the traditional approach of estimating m and ℓ one-by-one, with off-the-shelf ML algorithms, in such a way as to separately maximize the predictive accuracy for D and Y , respectively.

1.2. Related work

This work is most closely related with [Rostami et al. \(2021\)](#), which proposed using a multi-task predictive model to estimate m and ℓ simultaneously. Their idea is to share the neural network architecture for the tasks of predicting $D | X$ and $Y | X$, so that each model can borrow strength from the other, but they do not consider a new loss function designed to optimize the estimation performance of DML. Instead, they focus solely on predictive accuracy while training.

[Mackey et al. \(2018\)](#) suggested reducing DML bias by replacing the *first-order* orthogonalized estimator in (4) with a *higher-order* estimator. By contrast, we focus on learning m and ℓ , not on changing the downstream estimation of θ . Future work may explore combining the two approaches.

Numerous other extensions of DML have been proposed to move beyond the partially linear regression model in (1) ([Bach et al., 2020](#); [Chang, 2020](#); [Colangelo & Lee, 2020](#); [Kallus & Uehara, 2020](#); [Klaassen et al., 2018; 2021](#); [Lewis & Syrgkanis, 2020](#); [Narita et al., 2020](#); [Semenova & Chernozhukov, 2021](#); [Kurz, 2021](#)). The idea developed in this paper could also be relevant in many of those settings, but such extensions are not considered here.

2. Bias in Double Machine Learning

Before introducing the proposed method, we begin by discussing the sources of DML bias in finite samples. The following analysis is adapted from [Chernozhukov et al. \(2018\)](#), and it is presented here with slightly different but essentially equivalent assumptions and a simplified notation to highlight the idea motivating our contribution. Without loss of generality, we will assume in the theoretical analyses that the total sample size is $2n$ and $|\mathcal{I}_1| = |\mathcal{I}_2| = n$, which simplifies the notation. Mathematical proofs are in [Appendix A](#).

2.1. Theoretical Analysis

Assumption 2.1 (i.i.d. data). The data $\{(X_i, D_i, Y_i)\}_{i=1}^{2n+1}$ are i.i.d. and follow the partially linear model in (1).

Assumption 2.2 (boundedness). The noise terms U, V and the estimation errors $\Delta\hat{m}(X) := m(X) - \hat{m}(X)$ and $\Delta\hat{\ell}(X) := \ell(X) - \hat{\ell}(X)$ of the two nuisance functions in (1) are bounded almost surely. The individual treatment effects are also bounded almost surely. (These technical assumptions are convenient but could be relaxed.)

Theorem 2.3. *Under Assumptions 2.1–2.2, for n large enough, with probability at least $1 - O(n^{-1})$ the DML treatment effect estimate satisfies*

$$\hat{\theta} - \theta = \mathcal{B}_{\text{DML}} + O\left(\sqrt{(\log n)/n}\right),$$

where \mathcal{B}_{DML} is defined as:

$$\mathcal{B}_{\text{DML}} := \frac{\mathbb{E}\left[\Delta\hat{m}(X)\Delta\hat{\ell}(X) \mid \mathcal{I}_1\right] - \theta\mathbb{E}\left[(\Delta\hat{m}(X))^2 \mid \mathcal{I}_1\right]}{\text{Var}[V] + \mathbb{E}\left[(\Delta\hat{m}(X))^2 \mid \mathcal{I}_1\right]}. \quad (5)$$

Above, the conditioning on \mathcal{I}_1 represents conditioning on the data indexed by \mathcal{I}_1 . In words, [Theorem 2.3](#) implies $\hat{\theta}$ will be close to θ , in the large- n limit, if $\text{Var}[V] > 0$ and \hat{m} estimates m accurately. In general, however, $\hat{\theta}$ may not necessarily be close to θ , or even unbiased, if $\Delta\hat{m}$ is large. [Appendix B](#) provides a practical demonstration of this phenomenon and validated empirically [Theorem 2.3](#).

It is worth mentioning that, in the alternative version of DML based on an ML model for g instead of one for ℓ ([Chernozhukov et al., 2018](#)), the analogous \mathcal{B}_{DML} term in (5) is simpler and only depends on $\mathbb{E}[\Delta\hat{m}(X)\Delta\hat{\ell}(X) \mid \mathcal{I}_1]$, not on $\mathbb{E}[(\Delta\hat{m}(X))^2 \mid \mathcal{I}_1]$. Unfortunately, it is unclear how to accurately estimate the function g in practice without also estimating θ . Thus, it would feel a little “circular” to express $\hat{\theta} - \theta$ in terms of something implicitly depending on the accuracy of an estimate of θ . By contrast, the result in (5) is more informative because it connects the DML estimation error to the predictive accuracy of two straightforward predictive models for $Y | X$ and $D | X$. In any case, the methodology proposed in the next section is still applicable even if DML is implemented using \hat{g} instead of $\hat{\ell}$.

2.2. Conditions for Asymptotic Consistency

[Theorem 2.3](#) tells us that the quality of $\hat{\theta}$ as an estimate of θ depends on the quality of \hat{m} and $\hat{\ell}$ as approximations of m and ℓ , respectively. The following simple corollary clarifies the connection between the rates of convergence of \hat{m} and $\hat{\ell}$ and that of $\hat{\theta}$, in the limit of large sample sizes.

Assumption 2.4 (consistency of \hat{m}). There exist sequences $\eta_n^m \rightarrow 0$ and $\rho_n^m \rightarrow 0$, as $n \rightarrow \infty$, such that

$$\mathbb{P}\left[\mathbb{E}\left[(\Delta\hat{m}(X))^2 \mid \mathcal{I}_1\right] \leq \eta_n^m\right] \geq 1 - \rho_n^m.$$

Assumption 2.5 (convergence of $\hat{\ell}$). There exist a decreasing sequence η_n^ℓ (not necessarily going to 0) and a sequence $\rho_n^\ell \rightarrow 0$, as $n \rightarrow \infty$, such that

$$\mathbb{P}\left[\mathbb{E}\left[(\Delta\hat{\ell}(X))^2 \mid \mathcal{I}_1\right] \leq \eta_n^\ell\right] \geq 1 - \rho_n^\ell.$$

Note that [Assumption 2.4](#) (resp. [2.5](#)) is weaker than requiring $\hat{m}(X)$ to converge to $m(X)$ in L2 norm (resp. $\hat{\ell}(X)$ to $\ell(X)$) and, intuitively, it says that $\Delta\hat{m}$ (resp. $\Delta\hat{\ell}$) converges with high probability to 0 as $1/\sqrt{\eta_n^m}$ (resp. $1/\sqrt{\eta_n^\ell}$). Under these assumptions, [Theorem 2.3](#) implies the following.

Corollary 2.6. *Under Assumptions 2.1–2.5 and if $\text{Var}[V] > 0$, for n large enough, with probability at least $1 - O(1/n + \rho_n^m + \rho_n^\ell)$ the DML $\hat{\theta}$ satisfies*

$$\left| \hat{\theta} - \theta \right| \leq O \left(\sqrt{(\log n)/n} + \eta_n^m + \sqrt{\eta_n^\ell \eta_n^m} \right). \quad (6)$$

In words, Corollary 2.6 states that $\hat{\theta}$ will converge to θ roughly at rate $n^{-1/2}$ (a common standard of statistical efficiency, known as \sqrt{n} -consistency), if either $\eta_n^m \leq O(n^{-1})$, or $\eta_n^m \leq O(n^{-1/2})$ and $\eta_n^\ell \leq O(n^{-1/2})$. In the first case, $\hat{\ell}$ does not need to converge to ℓ but \hat{m} must be \sqrt{n} -consistent, which is a rather strong assumption given that the true m may have a complicated form. In the second case, it suffices for \hat{m} to be $\sqrt[4]{n}$ -consistent, as long as $\hat{\ell}$ also converges to ℓ at the same rate. The latter situation provides the main motivation for the DML methodology. In fact, thanks to data splitting and orthogonalization, relatively slow rates of convergence for \hat{m} and $\hat{\ell}$ lead to \sqrt{n} -consistency for $\hat{\theta}$.

The novelty of this work begins by noting that there exists a third way for $\hat{\theta}$ to achieve \sqrt{n} -consistency. The following additional corollary of Theorem 2.3 lays out the main idea.

Corollary 2.7. *Under Assumptions 2.1–2.4 and if $\text{Var}[V] > 0$, for n large enough, with probability at least $1 - O(1/n + \rho_n^m)$ the DML $\hat{\theta}$ satisfies*

$$\left| \hat{\theta} - \theta - \frac{\mathbb{E} \left[\Delta \hat{m}(X) \Delta \hat{\ell}(X) \mid \mathcal{I}_1 \right]}{\text{Var}[V]} \right| \leq O(\tilde{\eta}_n), \quad (7)$$

where $\tilde{\eta}_n = \sqrt{(\log n)/n} + \eta_n^m$.

Corollary 2.7 implies $\hat{\theta}$ can be \sqrt{n} -consistent even if \hat{m} is only $n^{1/4}$ -consistent and $\hat{\ell}$ does not converge to ℓ , as long as the estimation errors of \hat{m} and $\hat{\ell}$ are *uncorrelated*. In general, $\Delta \hat{m}$ and $\Delta \hat{\ell}$ are two arbitrary functions of the same random variable X , and therefore they may be correlated. However, there can exist functions $\Delta \hat{m}, \Delta \hat{\ell}$ such that $\mathbb{E}[(\Delta \hat{m}(X))^2 \mid \mathcal{I}_1] > 0$, $\mathbb{E}[(\Delta \hat{\ell}(X))^2 \mid \mathcal{I}_1] > 0$ while $\mathbb{E}[\Delta \hat{m}(X) \Delta \hat{\ell}(X) \mid \mathcal{I}_1] = 0$, and that would be a particularly desirable situation for the estimation of θ because it would relax the asymptotic convergence assumptions required for \sqrt{n} consistency.

The above result suggests that, if the goal is to estimate θ with DML, one should try to fit ML models \hat{m} and $\hat{\ell}$ that, in addition to being accurate in their respective prediction tasks, also have approximately uncorrelated errors. Inspired by this observation, a new loss function and a corresponding coordinated training algorithm designed to obtain DML estimates with lower bias are developed next.

3. Coordinated Double Machine Learning

3.1. Loss Function

Unfortunately, it is not possible to train a joint ML model for \hat{m} and $\hat{\ell}$ to directly minimize the \mathcal{B}_{DML} estimation error term in (5) because that involves several unobservable quantities. One issue is that $\Delta \hat{m}$ and $\Delta \hat{\ell}$ are not visible; however, one can observe the residuals defined in (3): $\hat{V}_i = \Delta \hat{m} + V_i$ and $\hat{U}_i = \Delta \hat{\ell} + \theta V_i + U_i$, respectively. Another issue is that it is unclear how to estimate the second term in the numerator of (5) because it depends on the unknown θ . The simplest solution would be to take a step back from the specific form of \mathcal{B}_{DML} and to minimize a loss function of the form

$$\mathcal{L}_0 := \frac{\alpha}{|\mathcal{I}_2|} \sum_{i \in \mathcal{I}_2} \hat{V}_i^2 + \frac{\beta}{|\mathcal{I}_2|} \sum_{i \in \mathcal{I}_2} \hat{U}_i^2,$$

with the scaling hyperparameters α and β discussed in the next section. This is essentially the approach of Rostami et al. (2021). Here, we propose to add a regularization term

$$\mathcal{L}_b := \left| \frac{1}{|\mathcal{I}_2|} \sum_{i \in \mathcal{I}_2} \hat{V}_i \hat{U}_i \right|, \quad (8)$$

so that the final loss function takes the form

$$\mathcal{L} := \mathcal{L}_0 + \gamma \mathcal{L}_b, \quad (9)$$

where γ is an additional hyperparameter that will be tuned via cross-validation, as explained in the next section. In the special case of $\gamma = 0$, the loss in (9) reduces to the (weighted) sum of the mean-squared-error losses for \hat{m} and $\hat{\ell}$, which are typically minimized separately in DML. The motivation for the new regularization term is two-fold. First, \mathcal{L}_b is equivalent to \mathcal{B}_{DML} if $\theta = 0$. This case corresponds to a *null hypothesis* of no treatment effect, and unbiasedness under the null can be especially useful to obtain a reliable hypothesis test based on the DML estimator (Chernozhukov et al., 2018). Second, note that for n large,

$$\mathcal{L}_b \approx \left| \mathbb{E} \left[\Delta \hat{m}(X) \Delta \hat{\ell}(X) \right] + \theta \text{Var}[V] \right|.$$

Thus, a small $\Delta \hat{m}(X)$ due to asymptotic consistency of \hat{m} for large sample sizes would imply that $\mathbb{E}[\Delta \hat{m}(X) \Delta \hat{\ell}(X)]$ is (much) greater in absolute value than both $\theta \text{Var}[V]$ and $\theta \mathbb{E}[(\Delta \hat{m}(X))^2 \mid \mathcal{I}_1]$, provided that $\Delta \hat{\ell}(X)$ does not converge to 0 as quickly. Then, in this limit,

$$\mathcal{B}_{\text{DML}} \approx \frac{\mathbb{E} \left[\Delta \hat{m}(X) \Delta \hat{\ell}(X) \mid \mathcal{I}_1 \right]}{\text{Var}[V] + \mathbb{E} \left[(\Delta \hat{m}(X))^2 \mid \mathcal{I}_1 \right]},$$

and therefore minimizing \mathcal{L} may reduce \mathcal{B}_{DML} more effectively than simply targeting \mathcal{L}_0 . Of course, it may not always be the case that ℓ is more difficult to estimate than m and the

above conditions are satisfied; therefore, the proposed approach may sometimes provide no improvement compared to standard DML. In fact, on the one hand, if the regularization hyperparameter γ is too small, there may be little difference between the model learnt with the proposed approach and that learnt by fitting \hat{m} and $\hat{\ell}$ separately. On the other hand, if γ is too large, minimizing (9) may lead to \hat{m} and $\hat{\ell}$ corresponding to relatively high $\mathbb{E}[(\Delta\hat{m}(X))^2 | \mathcal{I}_1]$ and $\mathbb{E}[(\Delta\hat{\ell}(X))^2 | \mathcal{I}_1]$, hindering the estimation of θ . This trade-off motivates the following adaptive tuning strategy that prevents the proposed method from doing much harm.

3.2. Training and Hyperparameter Tuning

After randomly splitting the available observations as usual into two disjoint subsets, $\mathcal{I}_1, \mathcal{I}_2$, the initial step of the proposed coordinated DML (C-DML) procedure, summarized in Algorithm 2, is to estimate the hyperparameters α, β, γ in the loss function (9). This is achieved by further splitting \mathcal{I}_2 into two disjoint subsets, $\mathcal{I}_{2,1}, \mathcal{I}_{2,2}$, and applying the standard DML procedure (Algorithm 1) to the data in $\mathcal{I}_1 \cup \mathcal{I}_{2,1}$, obtaining predictive models \hat{m}_0 and $\hat{\ell}_0$ based on \mathcal{I}_1 , as well as an estimate $\hat{\theta}_0$ of θ based on $\mathcal{I}_{2,1}$. The scaling hyperparameters α and β in (8) are set to be the inverse of the mean squared prediction error corresponding to \hat{m}_0 and $\hat{\ell}_0$, respectively, evaluated on the hold-out data in \mathcal{I}_2 . The purpose of this choice is to ensure that the joint loss function in (8) weights the predictive performances of the learnt model for the two learning tasks evenly, accounting for their possibly different intrinsic difficulties. The regularization penalty is scaled similarly, dividing it by the absolute value of the empirical covariance of the DML residuals. The value of the regularization hyperparameter γ is tuned with the following cross-validation procedure. First, the function \hat{g}_0 is defined as $\hat{g}_0(X) := \hat{\ell}_0(X) - \hat{\theta}_0\hat{m}_0(X)$. Second, joint models for \hat{m} and $\hat{\ell}$ are repeatedly trained on the data in $\mathcal{I}_1 \cup \mathcal{I}_{2,1}$ to minimize the proposed regularized loss in function (9) over a grid of possible values for γ , by applying Algorithm 3 for each γ . Note that this algorithm also outputs a preliminary DML treatment effect estimate $\hat{\theta}_1(\gamma)$ evaluated on the data in $\mathcal{I}_{2,1}$ and based on the \hat{m} and $\hat{\ell}$ learnt from the data in \mathcal{I}_1 . For each $\hat{\theta}_1(\gamma)$, the mean squared difference between Y and $\hat{g}_0(X) + D\hat{\theta}_1(\gamma)$ is evaluated on the hold-out data in $\mathcal{I}_{2,2}$. Then, $\hat{\gamma}$ is selected as the value of γ minimizing the above mean squared prediction error. Finally, the joint training procedure described in Algorithm 3 is applied to the full data set using the tuned hyperparameters $\alpha, \beta, \hat{\gamma}$.

3.3. Demonstration of Hyperparameter Tuning

Figure 2 gives a practical demonstration of the hyperparameter tuning procedure described in Algorithm 2. A synthetic data set is generated by sampling covariates $X \in \mathbb{R}^{10}$ from a standard Gaussian auto-regressive process of order one,

Algorithm 2 C-DML

Input: data $\{(X_i \in \mathbb{R}^d, D_i \in \mathbb{R}, Y_i \in \mathbb{R})\}_{i=1}^n$, list of regularization hyperparameters $\tilde{\Gamma}$.

Split $\{1, \dots, n\}$ into two disjoint subsets: $\mathcal{I}_1, \mathcal{I}_2$.

Split \mathcal{I}_2 into two disjoint subsets: $\mathcal{I}_{2,1}, \mathcal{I}_{2,2}$.

Apply Algorithm 1 to the data in $\mathcal{I}_1 \cup \mathcal{I}_{2,1}$, obtaining \hat{m}_0 and $\hat{\ell}_0$ based on \mathcal{I}_1 , and an estimate $\hat{\theta}_0$ based on $\mathcal{I}_{2,1}$.

Set $\alpha = \left(|\mathcal{I}_{2,1}|^{-1} \sum_{i \in \mathcal{I}_{2,1}} \hat{V}_i^2\right)^{-1}$,
 $\beta = \left(|\mathcal{I}_{2,1}|^{-1} \sum_{i \in \mathcal{I}_{2,1}} \hat{U}_i^2\right)^{-1}$.

Scale $\Gamma = \left(|\mathcal{I}_{2,1}|^{-1} \left| \sum_{i \in \mathcal{I}_{2,1}} \hat{V}_i \hat{U}_i \right| \right)^{-1} \cdot \tilde{\Gamma}$.

Define the function $\hat{g}_0(X) := \hat{\ell}_0(X) - \hat{\theta}_0\hat{m}_0(X)$.

for $\gamma \in \Gamma$ **do**

Apply Algorithm 3 with $(\mathcal{I}_1, \mathcal{I}_{2,1}, \alpha, \beta, \gamma)$, obtaining $\hat{\theta}_1(\gamma)$.

Evaluate on $\mathcal{I}_{2,2}$ the mean squared error ϕ_γ :

$$\phi_\gamma := \frac{1}{|\mathcal{I}_{2,2}|} \sum_{i \in \mathcal{I}_{2,2}} \left(Y_i - \hat{g}_0(X_i) - D_i \hat{\theta}_1(\gamma) \right)^2.$$

end for

Set $\hat{\gamma} := \arg \min_{\gamma \in \Gamma} \phi_\gamma$.

Apply Algorithm 3 with $(\mathcal{I}_1, \mathcal{I}_2, \alpha, \beta, \hat{\gamma})$, obtaining $\hat{\theta}(\hat{\gamma})$.

Output: a treatment effect estimate $\hat{\theta}(\hat{\gamma})$.

Algorithm 3 C-DML with fixed $(\mathcal{I}_1, \mathcal{I}_2, \alpha, \beta, \gamma)$

Input: subsets: $\mathcal{I}_1, \mathcal{I}_2$, hyper-parameter α, β, γ .

Input: data $\{(X_i \in \mathbb{R}^d, D_i \in \mathbb{R}, Y_i \in \mathbb{R})\}_{i=1}^n$, hyper-parameters α, β .

Train simultaneously \hat{m} and $\hat{\ell}$ on \mathcal{I}_1 to estimate m and ℓ in (1) and (2), by minimizing \mathcal{L} .

Estimate θ in (1) with $\hat{\theta}$ in (4) using the data in \mathcal{I}_2 .

Output: a treatment effect estimate $\hat{\theta}$.

with correlation parameter $\rho = 0.8$. The samples are divided into two disjoint groups; precisely, each is assigned to either the *majority* (containing $\approx 80\%$ of the observations) or the *minority* (containing $\approx 20\%$ of the observations) group, by comparing the value of the first feature, X_0 , to a suitable fixed threshold. The variables D and Y are generated from a partially linear model (1), with:

$$m(X) = \begin{cases} x_1 + 10x_3 + 5x_6, & \text{if } X \in \text{majority,} \\ 10x_1 + x_3 + 5x_6, & \text{if } X \in \text{minority,} \end{cases} \quad (10)$$

and

$$g(X) = \begin{cases} x_0 + 10x_2 + 5x_5, & \text{if } X \in \text{majority,} \\ 10x_0 + x_2 + 5x_5, & \text{if } X \in \text{minority,} \end{cases} \quad (11)$$

and the treatment effects are homogeneous with $\theta = 1$. The variables U and V are independent standard Gaussian noise.




Figure 2. Demonstration of our cross validation procedure described in Algorithm 2 (blue). Top: final treatment effect estimation bias as a function of the hyperparameter γ . Bottom: log prediction error ϕ_γ evaluated on hold-out data $\mathcal{I}_{2,2}$ as a function of γ . Algorithm 2 adaptively chooses the value of γ leading to the smallest hold-out prediction error (vertical line).

3.4. Demonstration of C-DML Bias Reduction

This section provides an empirical demonstration of the theoretical arguments made in Sections 2 and 3.1 to justify the proposed C-DML procedure as a lower-bias alternative to standard DML. A synthetic set is generated with a setup similar to that of 3.3, but with different nuisance functions to emphasize the versatility of the proposed method:

$$m(X) = \begin{cases} \text{ReLU}(\frac{1}{2}x_1^2 + x_3^3 + x_5), & \text{if } X \in \text{majority}, \\ \text{ReLU}(-\frac{5}{2}x_1^2 + x_4 + x_9), & \text{if } X \in \text{minority}, \end{cases}$$

$$g(X) = x_9 + |x_2| + \frac{1}{2}e^{x_4 + x_5}.$$

Above, $\text{ReLU}(x) := x$ if $x > 0$ and $\text{ReLU}(x) := 0$ otherwise. The noise variables U and V are independent Gaussian with zero mean and variances σ_U^2 and 1, respectively. The noise variance σ_U^2 will be varied as a control parameter.

Figure 3 compares several performance measures for the standard DML and the proposed C-DML, as a function of σ_U^2 . The statistics are averaged over 500 experiments based on independent samples with 2000 observations each. The results show C-DML leads to more accurate estimates of θ than regular DML if σ_U^2 is large, while the two methods perform similarly when σ_U^2 is small—the standard DML has slightly lower bias but higher variance under the latter regime. This performance gap can be understood by looking at other metrics, such as the absolute average of $\Delta\hat{m}(X) \cdot \Delta\hat{\ell}(X)$; this increases with σ_U^2 but stays lower for C-DML compared to DML, consistently with the theory in Section 2. Further, as σ_U^2 grows, $(\Delta\hat{\ell}(X))^2$ becomes larger for both DML and C-DML, consistently with the increased difficulty of predicting $Y | X$. At the same time,




Figure 3. Different performance metrics for DML and C-DML on synthetic data, as a function of the noise variance in the distribution of $Y | X$. From the top: estimation bias; mean squared estimation error; absolute covariance of $\Delta\hat{m}$ and $\Delta\hat{\ell}$; mean squared error for \hat{m} ; mean squared error for $\hat{\ell}$.

$(\Delta\hat{m}(X))^2$ does not increase with DML, as that method decouples the prediction tasks for $Y | X$ and $D | X$, but it increases for C-DML. In fact, the novel loss function in (9), with the adaptive hyperparameter tuning procedure described in Section 3.2, allows C-DML to sacrifice some of the predictive accuracy of \hat{m} and $\hat{\ell}$ in return for partially cancelling out $\Delta\hat{m}(X)\Delta\hat{\ell}(X)$, which leads to lower bias.

4. Numerical Experiments

4.1. Setup

In this section, the estimation performance of C-DML (Algorithm 2) is compared to that of standard DML (Algorithm 1) on synthetic and semi-synthetic data. In each experiment, the data are divided into three disjoint subsets, namely \mathcal{I}_1 (containing 50% of the observations), $\mathcal{I}_{2,1}$ (containing 25% of the observations), and $\mathcal{I}_{2,2}$ (containing 25% of the observations). For standard DML, the predictive models are trained on \mathcal{I}_1 and $\hat{\theta}$ is evaluated on $\mathcal{I}_2 = \mathcal{I}_{2,1} \cup \mathcal{I}_{2,2}$; the same approach is followed for C-DML, but in that case the

$\mathcal{I}_{2,1}$ and $\mathcal{I}_{2,2}$ are also utilized to tune the hyperparameters, as discussed in Section 3. For both methods, the predictive models are deep neural networks as described in Appendix D.1.1. Early stopping is applied for both methods to avoid overfitting, using predictive performance evaluated on the hold-out data in $\mathcal{I}_{2,1}$ as the stopping criterion.

4.2. Synthetic Data

Synthetic data are generated from the same model with homogeneous treatment effects as in Section 3.3, with ρ varying from 0.1 to 0.9. The nuisance functions (10) and (11) depend on different but consecutive features of the vector X . Therefore, these functions are approximately independent of one another if ρ is close to 0, in which case the new regularization penalty in (8) will tend to have little effect. By contrast, if ρ is close to 1, $m(X)$ and $g(X)$ become more correlated and so do $\Delta\hat{m}$ and $\Delta\hat{\ell}$, allowing the new regularization to play a more important role on the accuracy of $\hat{\theta}$. As an additional benchmark, standard DML is also applied using random forest predictive models instead of deep neural networks; see Appendix D for details.

Figure 4 compares the performances of the three methods as a function of ρ , averaging over 100 independent experiments. The results show that C-DML tends to yield to more accurate estimates $\hat{\theta}$, both when $\theta = 0$ and when $\theta = 1$. As anticipated, the improvement is more noticeable when ρ is large and $\Delta\hat{m}$ tends to become correlated with $\Delta\hat{\ell}$.

4.3. Semi-synthetic Data

To assess the performance of the proposed method in a realistic scenario we carry out experiments on semi-synthetic data, where the covariates are obtained from a real-world study but the outcome is generated from a partially linear model under which the true treatment effects are known exactly. In particular, we borrow a data set from Chernozhukov & Hansen (2004) which has also been used by Chernozhukov et al. (2018). In this study, the goal is to predict the effect on net financial assets of an individual’s 401(k) eligibility. The covariates $X \in \mathbb{R}^{13}$ measure various personal and financial details, and some of them are correlated with one other.

The real data are transformed into semi-synthetic data with known treatment effect with the following pre-processing. First, one features is picked to serve as the treatment variable D , and the remaining ones are stored in $X^{\text{semi}} = (X \setminus D) \in \mathbb{R}^{d-1}$. Second, the observations are split into disjoint subsets. Third, a random forest is trained to predict Y given X^{semi} , using the first subset of observations. Fourth, the function $g_{\text{RF}}(X^{\text{semi}})$ is defined as the output of the fitted random forest, to be evaluated on the second subset of the data. Finally, the semi-synthetic data set is constructed as:

$$\{(X_i^{\text{semi}} \in \mathbb{R}^{d-1}, D_i \in \mathbb{R}, Y_i^{\text{semi}} \in \mathbb{R})\}_i,$$




Figure 4. Distribution of treatment estimation errors obtained by applying the proposed C-DML (blue) and standard DML (red) on 100 independent experiments with synthetic data, using deep neural network predictive models. The horizontal axis represents the correlation between different covariates generated from a Gaussian autoregressive process, which controls the correlation of the nuisance functions. The green boxes represent the performance of standard DML applied with random forest predictive models.

where i indices only the observations in the second subset, and Y_i^{semi} follows the partial linear model from (1):

$$Y_i^{\text{semi}} = g_{\text{RF}}(X_i^{\text{semi}}) + D_i\theta_i + U_i.$$

Above, U_i is independent standard Gaussian noise, and θ_i is the treatment effect for individual i , which we may fix to be either homogeneous (as assumed so far) or heterogeneous.

Figure 1 reports on the results obtained by applying DML and C-DML to the above data in the case of a homogeneous treatment effect θ . Figure 5 reports on the results obtained in the case of heterogeneous treatment effects. In the latter, θ_i is sampled independently of everything else from a normal distribution with mean θ and variance one; two different values of θ are considered: $\theta = 0$ and $\theta = 10$. In both settings, the results show that the C-DML estimates suffer from lower bias compared to standard DML, especially when $\theta = 0$, which is consistent with the theory in Section 2. The results of additional numerical experiments with semi-synthetic data are reported in Appendix C. Those experiments involve Beijing air quality data (Zhang et al., 2017), Facebook blog feedback data (Buza, 2014), and other data from the DoubleML Python package (Bach et al., 2022).

5. Real Data Application

The proposed method is applied to study the relation between temperature and ozone concentration using the Bei-




Figure 5. DML estimates on semi-synthetic data with heterogeneous effects. The dotted vertical lines indicate the true average treatment effect (left: 0, right: 10). Other details are as in Figure 1.

jing air quality data set (Zhang et al., 2017). Although ozone generally increases with temperature, both quantities can be affected by other atmospheric variables through possibly complicated non-linear relations (Camalier et al., 2007). Therefore, we model the influence of temperature (D) and other covariates (X) on ozone (Y) with a partially linear model (1), aiming to estimate θ . As the ground truth is unknown, estimation accuracy cannot be directly measured. However, it is possible to quantify and compare the performance of different methods in terms of their self-consistency, by inspecting the results obtained from the analysis of increasingly large subsets of the data. Concretely, seven subsets of observations with sizes ranging from 1,000 to 10,000 are constructed by randomly sampling from the 420,768 observations in full data set. On each data subset, C-DML and DML are applied as in the previous semi-synthetic experiments, and a 95% confidence interval for θ is computed via the bootstrap (Bach et al., 2022). Note that a non-parametric bootstrap based on 200 samples with replacement is utilized instead of the faster multiplier bootstrap recommended by Bach et al. (2022), due to the relatively small size of the data sets considered here.

Figure 6 visualizes the bootstrap confidence intervals for θ obtained with each method as a function of the sample size. While the true value of θ is unknown, one would expect the confidence intervals to overlap with one another if the estimates are unbiased and the bootstrap reliably quantifies their standard errors. No evidence of the contrary can be seen for C-DML, whose estimate of $\hat{\theta} \approx 2.4$ is relatively stable for all sample sizes. By contrast, the standard DML estimates are less stable and their confidence intervals obtained with different sample sizes are more likely to be disjoint, suggesting lower estimation accuracy.

6. Conclusion

The ability to compute approximately unbiased treatment effect estimates starting from any ML predictive models undoubtedly is a desirable strength of DML (Chernozhukov et al., 2018). Yet the estimation bias in finite samples is affected not only by the accuracy of the underlying predic-




Figure 6. Effect of temperature on ozone concentration, estimated from Beijing air quality data (Zhang et al., 2017) with C-DML or DML. Estimates are shown as a function of the number of samples utilized. Error bars indicate 95% bootstrap confidence intervals.

tive models, but also on the covariance of their error terms, which is a quantity that cannot be easily controlled with standard single-task training techniques. Given that the DML bias can be large in practice, depending in unforeseeable ways on the data at hand as well as on the particulars of the predictive models employed, it may be wise for practitioners to consider carefully how to choose and train the ML models utilized with this methodology. This paper provides a principled solution based on joint learning of the two underlying predictive models with a regularized multi-task loss function that explicitly attempts to minimize the estimation errors. This approach has the downside of involving additional data splits and more expensive computations, but the experiments carried out in this study suggest that such cost may be well-justified by higher accuracy, especially if the null hypothesis of zero average effects is true.

There are several directions for future research. For example, one could investigate whether our adaptive hyperparameter tuning algorithm can be improved, especially to facilitate the analysis of very large data sets for which the current implementation may be computationally prohibitive. Further, we have focused for convenience on deep neural networks, but the methodology could be extended to other predictive models, such as random forests. Finally, this work could be generalized to deal with a binary treatment or outcome.

Software Availability

A Python implementation of the methods described in this paper is available from <https://github.com/nitaifingerhut/C-DML.git>, along with tutorials and code to reproduce the experiments.

Acknowledgements

N.F. and Y.R. were supported by the Israel Science Foundation (grant 729/21). Y.R. also thanks the Career Advancement Fellowship, Technion, for providing research support.

References

- Bach, P., Klaassen, S., Kueck, J., and Spindler, M. Uniform inference in high-dimensional generalized additive models. *preprint arXiv:2004.01623*, 2020.
- Bach, P., Chernozhukov, V., Kurz, M. S., and Spindler, M. DoubleML – An object-oriented implementation of double machine learning in Python. *J. Mach. Learn. Res.*, 23(53):1–6, 2022.
- Bach, P. S. *Applications in High-Dimensional Econometrics*. PhD thesis, Staats- und Universitätsbibliothek Hamburg Carl von Ossietzky, 2021.
- Buza, K. Feedback prediction for blogs. In *Data analysis, machine learning and knowledge discovery*, pp. 145–152. Springer, 2014.
- Camalier, L., Cox, W., and Dolwick, P. The effects of meteorology on ozone in urban areas and their use in assessing ozone trends. *Atmos. Environ.*, 41(33):7127–7137, 2007.
- Chang, N.-C. Double/debiased machine learning for difference-in-differences models. *Econom. J.*, 23(2):177–191, 2020.
- Chernozhukov, V. and Hansen, C. The effects of 401 (k) participation on the wealth distribution: an instrumental quantile regression analysis. *Rev. Econ. Stat.*, 86(3):735–751, 2004.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters. *Econom. J.*, 21(1):C1–C68, 01 2018.
- Chernozhukov, V., Kasahara, H., and Schrimpf, P. Causal impact of masks, policies, behavior on early covid-19 pandemic in the us. *J. Econom.*, 220(1):23–62, 2021.
- Colangelo, K. and Lee, Y.-Y. Double debiased machine learning nonparametric inference with continuous treatments. *preprint arXiv:2004.03036*, 2020.
- Feng, G., Giglio, S., and Xiu, D. Taming the factor zoo: A test of new factors. *J. Finance*, 75(3):1327–1370, 2020.
- Kallus, N. and Uehara, M. Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *J. Mach. Learn. Res.*, 21:167–1, 2020.
- Klaassen, S., Kück, J., Spindler, M., and Chernozhukov, V. Uniform inference in high-dimensional Gaussian graphical models. *preprint arXiv:1808.10532*, 2018.
- Klaassen, S., Kueck, J., and Spindler, M. Transformation models in high dimensions. *J. Bus. Econ. Stat.*, pp. 1–11, 2021.
- Knaus, M. C. Double machine learning based program evaluation under unconfoundedness. *preprint arXiv:2003.03191*, 2020.
- Kurz, M. S. Distributed double machine learning with a serverless architecture. In *Companion ACM/SPEC Int. Conf. Perform. Eng.*, pp. 27–33, 2021.
- Lewis, G. and Syrgkanis, V. Double/debiased machine learning for dynamic treatment effects. *preprint arXiv:2002.07285*, 2020.
- Mackey, L., Syrgkanis, V., and Zadik, I. Orthogonal machine learning: Power and limitations. In *Int. Conf. Mach. Learn.*, pp. 3375–3383. PMLR, 2018.
- Narita, Y., Yasui, S., and Yata, K. Off-policy bandit and reinforcement learning. *preprint arXiv:2002.08536*, 2020.
- Robinson, P. M. Root-n-consistent semiparametric regression. *Econometrica*, pp. 931–954, 1988.
- Rostami, M., Saarela, O., and Escobar, M. The bias-variance tradeoff of doubly robust estimator with targeted ℓ_1 regularized neural networks predictions. *preprint arXiv:2108.00990*, 2021.
- Semenova, V. and Chernozhukov, V. Debiased machine learning of conditional average treatment effects and other causal functions. *Econom. J.*, 24(2):264–289, 2021.
- Semenova, V., Goldman, M., Chernozhukov, V., and Taddy, M. Estimation and inference on heterogeneous treatment effects in high-dimensional dynamic panels. *preprint arXiv:1712.09988*, 2017.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, 2014.
- Torrats-Espinosa, G. Using machine learning to estimate the effect of racial segregation on covid-19 mortality in the united states. *Proc. Natl. Acad. Sci. U.S.A.*, 118(7), 2021.
- Zhang, S., Guo, B., Dong, A., He, J., Xu, Z., and Chen, S. X. Cautionary tales on air-quality improvement in beijing. *Proc. R. Soc. A*, 473(2205):20170457, 2017.

A. Mathematical Proofs

A.1. Proof of Theorem 2.3

We begin by restating the assumptions more formally.

Assumption A.1 (i.i.d. data). The data $\{(X_i, Y_i)\}_{i=1}^{2n+1}$ are i.i.d. and follow the partially linear model in (1). In particular, the noise terms U and V are independent of X , as well as of one another, and have mean zero.

Assumption A.2 (boundedness). There exist positive constants $C_1, C_2, C_3, C_4, C_\theta$ such that $-C_1 < U_i < C_1$ and $-C_2 < V_i < C_2$ a.s. in (1), $-C_3 < \Delta\hat{m}(X_i) < C_3$, $-C_4 < \Delta\hat{\ell}(X_i) < C_4$, and $-C_\theta \leq \theta \leq C_\theta$.

Proof. Recall that the DML treatment effect estimate is defined in (4) as:

$$\hat{\theta} = \left(\frac{1}{n} \sum_{i \in \mathcal{I}_2} \hat{V}_i^2 \right)^{-1} \frac{1}{n} \sum_{i \in \mathcal{I}_2} \hat{V}_i \left(Y_i - \hat{\ell}(X_i) \right),$$

where we set $|\mathcal{I}_2| = n$ to simplify the notation (assuming the total number of samples is $2n$), without loss of generality.

Let's start by looking at the numerator above:

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \hat{V}_i \left(Y_i - \hat{\ell}(X_i) \right) \\ &= \frac{1}{n} \sum_{i=1}^n (\Delta\hat{m}(X_i) + V_i) \left(g(X_i) + m(X_i)\theta_i + V_i\theta_i + U_i - \hat{\ell}(X_i) \right) \\ &= \frac{1}{n} \sum_{i=1}^n (\Delta\hat{m}(X_i) + V_i) \left(\Delta\hat{\ell}(X_i) + V_i\theta_i + U_i \right) \\ &= \frac{1}{n} \sum_{i=1}^n \Delta\hat{m}(X_i)\Delta\hat{\ell}(X_i) + \frac{1}{n} \sum_{i=1}^n \Delta\hat{m}(X_i)(U_i + V_i\theta_i) + \frac{1}{n} \sum_{i=1}^n V_i\Delta\hat{\ell}(X_i) + \frac{1}{n} \sum_{i=1}^n V_i^2\theta_i + \frac{1}{n} \sum_{i=1}^n V_iU_i \\ &= \underbrace{\frac{1}{n} \sum_{i=1}^n \theta_i V_i^2 + \frac{1}{n} \sum_{i=1}^n \Delta\hat{m}(X_i)\Delta\hat{\ell}(X_i)}_{\triangleq \psi_1} + \underbrace{\frac{1}{n} \sum_{i=1}^n (U_i + V_i\theta_i) \Delta\hat{m}(X_i) + \frac{1}{n} \sum_{i=1}^n V_i\Delta\hat{\ell}(X_i) + \frac{1}{n} \sum_{i=1}^n V_iU_i}_{\triangleq \psi_2}. \end{aligned}$$

Hoeffding's inequality yields that, with probability at least $1 - O(n^{-1})$,

$$|\psi_1 - \mathbb{E}[\psi_1 | \mathcal{I}_1]| \leq (C_\theta C_2^2 + C_3 C_4) \sqrt{\frac{2 \log n}{n}},$$

where

$$\mathbb{E}[\psi_1 | \mathcal{I}_1] = \theta \text{Var}[V] + \mathbb{E}[\Delta\hat{m}(X)\Delta\hat{\ell}(X) | \mathcal{I}_1].$$

As $\mathbb{E}[\psi_2 | \mathcal{I}_1] = 0$, Hoeffding's inequality similarly also yields that, with probability at least $1 - O(n^{-1})$,

$$|\psi_2| \leq (C_1 C_3 + C_\theta C_2 C_3 + C_2 C_4 + C_1 C_2) \sqrt{\frac{2 \log n}{n}}.$$

Therefore, with probability at least $1 - O(n^{-1})$,

$$\left| \frac{1}{n} \sum_{i=1}^n \hat{V}_i \left(Y_i - \hat{\ell}(X_i) \right) - \left(\theta \text{Var}[V] + \mathbb{E}[\Delta\hat{m}(X)\Delta\hat{\ell}(X) | \mathcal{I}_1] \right) \right| \leq O\left(\sqrt{\frac{\log n}{n}} \right).$$

Let us now look at the denominator in the definition of $\hat{\theta}$.

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \hat{V}_i^2 &= \frac{1}{n} \sum_{i=1}^n (\Delta \hat{m}_0(X_i) + V_i)^2 \\ &= \underbrace{\frac{1}{n} \sum_{i=1}^n V_i^2 + \frac{1}{n} \sum_{i=1}^n (\Delta \hat{m}_0(X_i))^2}_{\triangleq \psi_3} + 2 \underbrace{\frac{1}{n} \sum_{i=1}^n V_i \Delta \hat{m}_0(X_i)}_{\triangleq \psi_4}. \end{aligned}$$

Applying Hoeffding's inequality as above gives us that, with probability at least $1 - O(n^{-1})$,

$$\left| \frac{1}{n} \sum_{i=1}^n \hat{V}_i^2 - (\text{Var}[V] + \mathbb{E}[\Delta \hat{m}^2(X) | \mathcal{I}_1]) \right| \leq O\left(\sqrt{\frac{\log n}{n}}\right).$$

Combining the above results with a first-order Taylor expansion of $1/(1+x)$ for $x \approx 0$, one obtains that, with probability at least $1 - O(n^{-1})$,

$$\begin{aligned} O\left(\sqrt{\frac{\log n}{n}}\right) &\geq \left| \hat{\theta} - \frac{\theta \text{Var}[V] + \mathbb{E}[\Delta \hat{m}(X) \Delta \hat{\ell}(X) | \mathcal{I}_1]}{\text{Var}[V] + \mathbb{E}[\Delta \hat{m}^2(X) | \mathcal{I}_1]} \right| \\ &= \left| \hat{\theta} - \theta - \frac{\mathbb{E}[\Delta \hat{m}(X) \Delta \hat{\ell}(X) | \mathcal{I}_1] - \theta \mathbb{E}[\Delta \hat{m}^2(X) | \mathcal{I}_1]}{\text{Var}[V] + \mathbb{E}[\Delta \hat{m}^2(X) | \mathcal{I}_1]} \right|. \end{aligned}$$

□

B. Empirical Bias Analysis

To verify empirically the analysis in Section 2, the following experiments are conducted. Synthetic data are generated from a partially linear model (1), as in Section 3.3, with a sample size of 2000. The standard DML is applied with two deep neural network estimators as described in Section D.1.1, and θ is estimated as detailed in Algorithm 1. Figure 7 compares the empirical and theoretical DML bias over 10000 independent experiments, confirming the results from Section 2. Figure 8 reports analogous results obtained after adding independent Gaussian noise with mean zero and variance σ_ℓ^2 to the DML predictive model $\hat{\ell}(X)$. As σ_ℓ^2 grows and the predictive models becomes less accurate, the estimation error increases, but the relation between empirical and theoretical bias continues to hold.

Figure 7. Scatter plot of the empirical ($\Delta\hat{\theta} = \hat{\theta} - \theta$) vs. theoretical (\mathcal{B}_{DML}) DML bias in 10000 independent experiments with synthetic data. The treatment effects are homogeneous with $\theta = 1$. The three ticks on each axis denote the minimum, mean, and maximum values. The mean squared error reported on top reflects the distance between the theoretical and empirical bias.





Figure 8. Scatter plot of the empirical ($\Delta\hat{\theta} = \hat{\theta} - \theta$) vs. theoretical (\mathcal{B}_{DML}) DML bias, as in Figure 7. The DML predictive model is perturbed by adding independent Gaussian noise: $\hat{\ell}(X) + \mathcal{N}(0.0, \sigma_\ell^2)$. Left: $\sigma_\ell = 1$; center: $\sigma_\ell = 10$; right: $\sigma_\ell = 25$.

C. Additional Results From Numerical Experiments

Synthetic data are obtained from several real-world data sets as explained in Section 4.3. In these additional experiments, the treatment effects are homogeneous, and their mean θ is varied as a control parameter. The predictive models utilized in these experiments are deep neural networks, as explained in Appendix D.1.1. The experiments are repeated 5000 times.

BEIJING AIR QUALITY

The *Beijing air quality* data set, presented in Zhang et al. (2017), details the air pollution in various train stations in Beijing. The covariates $X \in \mathbb{R}^{14}$ correspond to several measurements, such as temperature, air pressure, date, and time. The semi-synthetic data are obtained considering temperature as the treatment, and aiming to predict the concentration ($[ug/m^3]$) of PM2.5 in the air.




Figure 9. DML and C-DML treatment effect estimates for the *Beijing air quality* semi-synthetic data set, with a sample size of 2000. Other details are as in Figure 1.

FACEBOOK BLOG FEEDBACK

The *Facebook blog feedback* data set of Buza (2014) originates from Facebook blog posts. The covariates $X \in \mathbb{R}^{279}$ detail attributes of the posts, such as time in the air (with respect to some baseline), day of posting, and keywords in the text. The target is to predict the number of comments in the first 24 hours after posting the blog. The post length (measured in words) is utilized as treatment, aiming to predict the popularity of the post. Due to the large amount of features, a larger sample size of 4000 is used in these experiments. The results are presented in Figure 10.

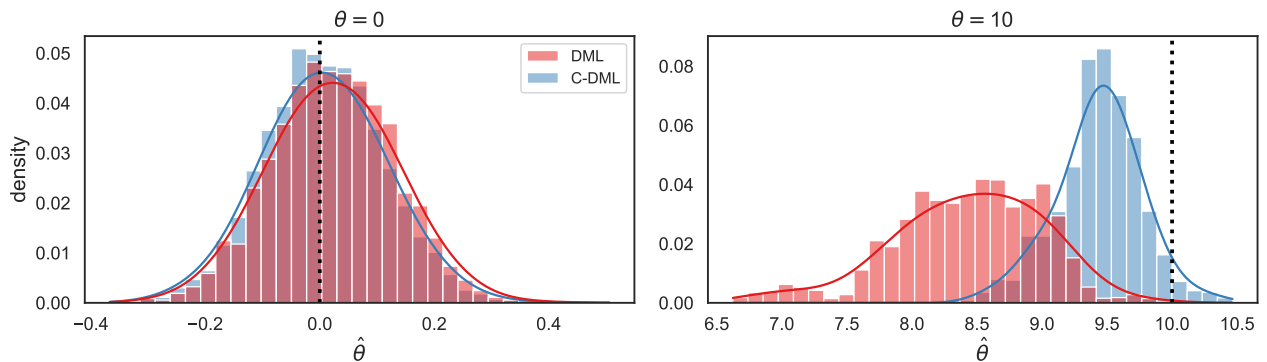


Figure 10. DML and C-DML treatment effect estimates for the *Facebook blog feedback* data set, with a sample size of 4000. Other details are as in Figure 1.

CCDDHNR2018

The *CCDDHNR2018* data set of [Bach \(2021\)](#) is integrated into `DoubleML` Python package of [Bach et al. \(2022\)](#), and can be used to generate semi-synthetic data following the PLR model from (1) with adjustable parameters, such as the number d of covariates. Specifically, two different values of d are considered: $d = 10$ and $d = 50$. The results are in Figure 11. When $\theta = 0$, both methods perform well, although C-DML has slightly lower variance. When $\theta = 10$, the relative performance depends on d : DML has lower bias (although higher variance) with $d = 10$, but C-DML clearly outperforms with $d = 50$.

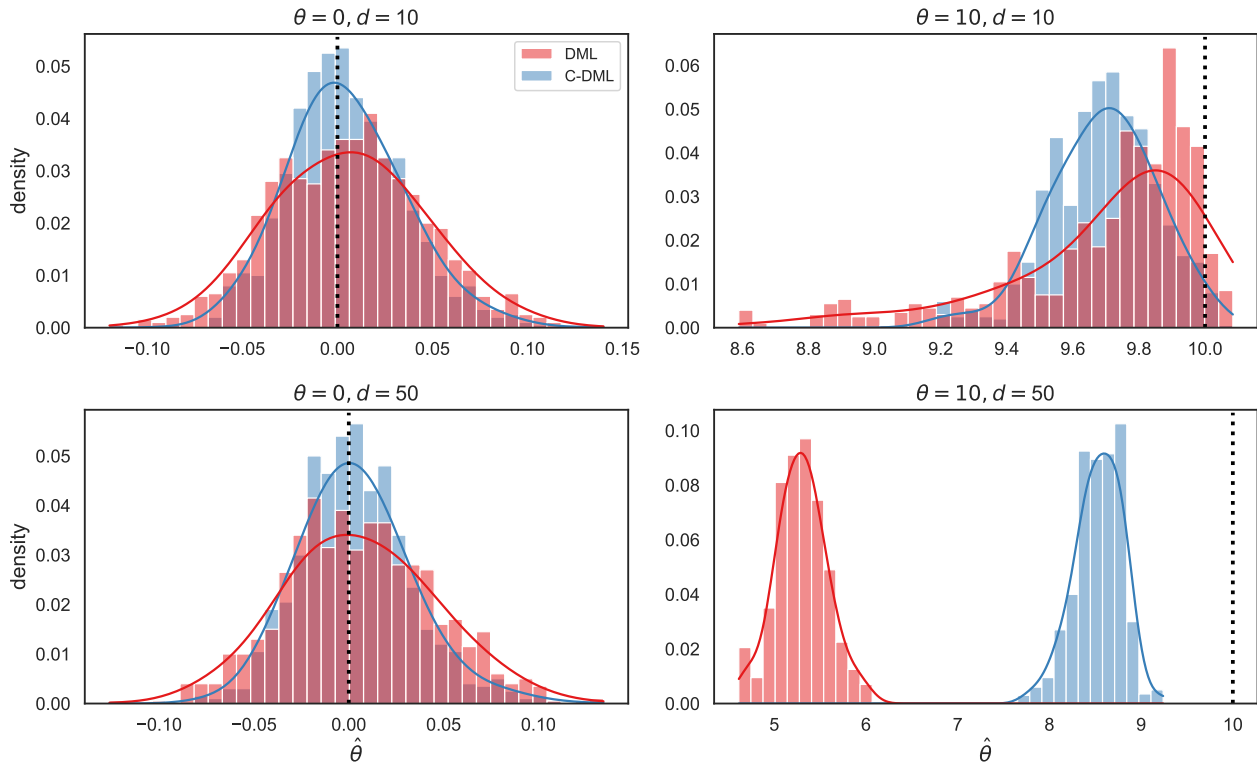


Figure 11. DML and C-DML treatment effect estimates for the *CCDDHNR2018* data set, with a sample size of 2000. Top: data with $d = 10$ covariates. Bottom: data with $d = 50$ covariates. Other details are as in Figure 1.

D. Further Information About the Models and Training Schemes

D.1. Implementation Details

D.1.1. NEURAL NETWORKS

The deep neural networks utilized in this paper have 3 layers. Given a d -dimensional input, the 3 layers have width equal to $(\lfloor d/2 \rfloor, \lfloor d/4 \rfloor, 1)$ respectively. In the case of C-DML, we used the exact same architecture to form $\hat{\nu}(X)$, $\hat{\ell}(X)$, and coordinated their training with respect to the loss in (9). We make this choice to have a fair comparison between DML and C-DML, such that the number of parameters for both methods is the same. (Future work may suggest a shared architecture, where $\hat{\nu}(X)$, $\hat{\ell}(X)$ share weights except the last layer that provides two outputs instead of one.) Learning is carried out via gradient descent, with a gradient step for all samples at once. The learning rate is fixed to 0.01, clipping gradients with norms larger than 3. Early stopping is utilized to avoid overfitting; the number of epochs (capped at 2000) is tuned by evaluating the loss function on a hold-out data set.

The experiments of Figure 4 are based on a more expressive deep learning network to provide a fairer comparison with the random forest model in Section D.1.2. This neural network differs from the previous one in the number of layers and in their width. The number of layers here is 5, and their width are $(d, d, d, d, 1)$, respectively. Dropout regularization (Srivastava et al., 2014) is applied between the second and third layer, with the probability of retaining a hidden unit equal to 0.1.

D.1.2. RANDOM FORESTS

Random forest regression models are implemented using the Python package `sklearn`. The default hyper-parameters are utilized, except the number of trees in the forest and the maximal depth, both of which are set equal to 20.