

Towards optimal doubly robust estimation of heterogeneous causal effects

Edward H. Kennedy

Department of Statistics & Data Science, Carnegie Mellon University
e-mail: edward@stat.cmu.edu

Abstract: Heterogeneous effect estimation is crucial in causal inference, with applications across medicine and social science. Many methods for estimating conditional average treatment effects (CATEs) have been proposed, but there are gaps in understanding if and when such methods are optimal. This is especially true when the CATE has nontrivial structure (e.g., smoothness or sparsity). Our work contributes in several ways. First, we study a two-stage doubly robust CATE estimator and give a generic error bound, which yields rates faster than much of the literature. We apply the bound to derive error rates in smooth nonparametric models, and give sufficient conditions for oracle efficiency. Along the way we give a general error bound for regression with estimated outcomes; this is the second main contribution. The third contribution is aimed at understanding the fundamental statistical limits of CATE estimation. To that end, we propose and study a local polynomial adaptation of double-residual regression. We show that this estimator can be oracle efficient under even weaker conditions, and we conjecture that they are minimal in a minimax sense. We go on to give error bounds in the non-trivial regime where oracle rates cannot be achieved. Some finite-sample properties are explored with simulations.

MSC2020 subject classifications: 62G08, 62H12, 62N01.

Keywords and phrases: Conditional effects, influence function, minimax rate, nonparametric regression.

Received May 2022.

Contents

1	Introduction	3009
2	Setup & illustration	3010
	2.1 Notation	3010
	2.2 Simple motivating illustration	3011
3	General pseudo-outcome regression	3013
4	DR-Learner	3017
	4.1 Previous work	3017
	4.2 Construction & analysis	3018
	4.3 Examples & illustrations	3021
	4.4 Simulation experiments	3022

arXiv: [2004.14497](https://arxiv.org/abs/2004.14497)

5	Local polynomial R-learner	3023
5.1	Construction	3024
5.2	Main error bound & oracle results	3025
5.3	Faster rates with known covariate density	3029
6	Discussion	3030
A	Proofs	3032
A.1	Proof of Theorem 1	3032
A.2	Proof of Proposition 1	3033
A.3	Proof of Proposition 2	3033
A.4	Proof of Theorem 2	3033
A.5	Proof of Corollary 1	3034
A.6	Proof of Theorem 3	3034
A.6.1	Term (A.4)	3035
A.6.2	Term (A.5)	3037
A.6.3	Term (A.6)	3040
A.6.4	Combining bounds	3042
B	R code	3044
	Acknowledgments	3046
	Funding	3046
	References	3046

1. Introduction

Heterogeneous effect estimation plays a crucial role in causal inference, with applications across medicine and social science, e.g., improving understanding of variation, and informing policy or optimizing treatment decisions. The most common target parameter in this setup is the conditional average treatment effect (CATE) function, $\mathbb{E}(Y^1 - Y^0 \mid X = x)$, which measures the expected difference in outcomes had those with covariates $X = x$ been treated versus not. The CATE is typically identified under standard causal assumptions (including no unmeasured confounding) as the difference between two regression functions, $\tau(x) \equiv \mathbb{E}(Y \mid X = x, A = 1) - \mathbb{E}(Y \mid X = x, A = 0)$.

Important early methods for estimating the CATE often employed semiparametric models, for example partially linear models assuming $\tau(x)$ to be constant [39, 34], or structural nested models in which $\tau(x)$ followed some known parametric form [32, 51, 48, 53]. These approaches reflect a commonly held conviction that the CATE may be more structured and simple than the rest of the data-generating process (with the zero treatment effect case an obvious example). This is similar in spirit to the common belief that interaction terms are more often zero than “main effects”.

Recent years have seen a move towards more flexible estimators of $\tau(x)$. The first formal nonparametric model where the CATE has its own complexity separate from regression functions seems to be Example 4 of Robins et al. [35]. van der Laan [48] (Section 4.2) proposed an important model-free “meta-algorithm” for estimating the CATE, a variant of which we study in this paper. The last

5–10 years has seen even more emphasis on nonparametrics and incorporating machine learning [16, 21, 3, 44, 43, 30, 55, 25, 15, 20]. The present work is in a similar vein, focusing on (i) providing more flexible CATE estimators with stronger theoretical guarantees, and (ii) pushing forward our understanding of optimality and the fundamental limits of CATE estimation. At the start of each of the Sections 3, 4, and 5, we detail related work and describe how our results fit.

After describing the setup and presenting a simple motivating illustration in Section 2, we go on to present our three main contributions: (i) a model-free oracle inequality for regression with estimated pseudo-outcomes (given in Section 3); (ii) general conditions for oracle efficiency of a doubly robust estimator we term the DR-Learner, with applications to specific regression methods under smoothness conditions (in Section 4); and (iii) a more refined analysis of a specialized estimator we call the lp-R-Learner, showing that faster rates can be achieved in the non-oracle regime, and giving a partial answer towards understanding the fundamental limits of CATE estimation (in Section 5).

2. Setup & illustration

We assume access to an iid sample of observations of $Z_i = (X_i, A_i, Y_i)$, where $X \in \mathbb{R}^d$ are covariates, $A \in \{0, 1\}$ is a binary treatment or exposure, and $Y \in \mathbb{R}$ an outcome of interest. The distribution of Z is indexed by the covariate distribution and nuisance functions:

$$\begin{aligned}\pi(x) &= \mathbb{P}(A = 1 \mid X = x) \\ \mu_a(x) &= \mathbb{E}(Y \mid X = x, A = a) \\ \eta(x) &= \mathbb{E}(Y \mid X = x)\end{aligned}$$

Our goal is estimation of the difference in regression functions under treatment versus control $\tau(x) \equiv \mu_1(x) - \mu_0(x)$. Under standard causal assumptions of no unmeasured confounding, consistency, and positivity or overlap ($\epsilon \leq \pi \leq 1 - \epsilon$ w.p.1, which we assume throughout), the function $\tau(x)$ also equals $\mathbb{E}(Y^1 - Y^0 \mid X = x)$, where Y^a is the counterfactual outcome under $A = a$. We refer to $\mu_1(x) - \mu_0(x)$ as the CATE, noting that our results hold regardless of whether the causal assumptions do. The average treatment effect (ATE) is given by $\mathbb{E}\{\tau(X)\}$.

2.1. Notation

We use $\mathbb{P}_n(f) = \mathbb{P}_n\{f(Z)\} = \frac{1}{n} \sum_i f(Z_i)$ as shorthand for sample averages. When $x \in \mathbb{R}^d$ we let $\|x\|^2 = \sum_j x_j^2$ denote the usual (squared) Euclidean norm, and for generic (possibly random) functions f we let $\|f\|^2 = \int f(z)^2 d\mathbb{P}(z)$ denote the (squared) $L_2(\mathbb{P})$ norm. We use the notation $a \lesssim b$ to mean $a \leq Cb$ for some universal constant C , and $a \asymp b$ to mean $cb \leq a \leq Cb$ so that $a \lesssim b$ and $b \lesssim a$. We let $a_n \sim b_n$ mean $a_n/b_n \rightarrow 1$ as $n \rightarrow \infty$.

At various points we refer to s -smooth functions, which we define as those contained in the Hölder class $\mathcal{H}(s)$, a canonical function class in nonparametric regression, density estimation, and functional estimation. Intuitively, it contains smooth functions that are close to their $\lfloor s \rfloor$ -order Taylor approximations. More precisely, $\mathcal{H}(s)$ is the set of functions $f : \mathcal{X} \rightarrow \mathbb{R}$ that are $\lfloor s \rfloor$ -times continuously differentiable with partial derivatives up to order $\lfloor s \rfloor$ bounded, and for which

$$|D^m f(x) - D^m f(x')| \lesssim \|x - x'\|^{s - \lfloor s \rfloor}$$

for all x, x' and $m = (m_1, \dots, m_d)$ such that $\sum_j m_j = \lfloor s \rfloor$, where $D^m = \frac{\partial^{\lfloor s \rfloor}}{\partial x_1^{m_1} \dots \partial x_d^{m_d}}$ is the multivariate partial derivative operator.

2.2. Simple motivating illustration

Consider a simple data-generating process where the covariates X are uniform on $[-1, 1]$,

$$\pi(x) = 0.5 + 0.4 \operatorname{sign}(x)$$

and $\mu_1(x) = \mu_0(x)$ are equal to the piecewise polynomial function defined on page 10 of Györfi et al. [18], given by

$$\mu_a(x) = \begin{cases} (x+2)^2/2 & \text{if } -1 \leq x \leq -0.5, \\ x/2 + 0.875 & \text{if } -0.5 \leq x < 0, \\ -5(x-0.2)^2 + 1.075 & \text{if } 0 < x \leq 0.5, \\ x + 0.125 & \text{if } 0.5 \leq x < 1, \end{cases}$$

which is illustrated in Figure 1. Figure 1 also shows $n = 1000$ simulated data points from this data-generating process, approximately half of which are treated (shown on the left panel) and the other half untreated (shown on the right). Also shown are estimates of the corresponding μ_1 and μ_0 functions, using default tuning parameters with the `smoothing.spline` function in base R.

An interesting but likely common phenomenon occurs in this simple example. The individual regression functions are non-smooth, and difficult to estimate well on their own; this is especially true in the region where there are fewer treated individuals. Thus the estimate $\hat{\mu}_1$ tends to oversmooth on the left, where there are more untreated individuals; in contrast, the estimate $\hat{\mu}_0$ tends to *undersmooth* on the right, where there are more treated individuals. This means a naive plug-in estimator of the CATE that simply takes the difference $\hat{\mu}_1 - \hat{\mu}_0$ will be a poor and overly complex estimator of the true difference, which is not only a constant but zero.

In contrast, suppose for simplicity that the propensity scores π were known. Then a regression of the inverse-probability-weighted (IPW) pseudo-outcome $\frac{(A-\pi)Y}{\pi(1-\pi)}$ would, up to constants, behave just as an oracle estimator that had access to the actual counterfactual difference $Y^1 - Y^0$, since the conditional




FIG 1. Plot of simulated data where the regression functions μ_1 and μ_0 individually are complex and difficult to estimate, but their difference is simply constant and equal to zero. Thus a naive plug-in estimator of the CATE will be overly complex, yielding large errors.




FIG 2. Estimated CATE curves in a simple simulated example. The plug-in method inherits the large errors from estimating the individual regression functions, which are complex and non-smooth. The IPW and doubly robust methods adapt to the smoothness of the CATE, which is constant in this example, with the doubly robust methods more efficient.

mean of this pseudo-outcome is exactly $\tau(x)$. Figure 2 shows results from this procedure, as well as two other more efficient and doubly robust versions described in subsequent sections, again all using default tuning parameter choices from `smoothing.spline`. For these simulated data, the doubly robust estimators are much more efficient than the IPW estimator, and do a much better job of adapting to the correct underlying simplicity of the true τ .

The results shown in Figure 2 are typical for this data-generating process: across 500 simulations, the IPW and doubly robust estimators gave smaller inte-

grated squared bias across X by a factor of 10–100, respectively, and the doubly robust estimators improved on the integrated variance of the IPW estimator by nearly a factor of 20.

Remark 1. If the CATE has no additional structure (e.g., smoothness, sparsity) beyond that of the individual regression functions μ_a , then a simple plug-in estimator $\hat{\tau}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x)$ would be sufficient. However, in general the complexities of μ_a and τ may differ substantially, and in practice we often expect the CATE to be much more structured, as described in the Introduction.

In the following sections we study the error of procedures like those illustrated above, giving model-free guarantees for practical use, as well as a more theoretical study of the fundamental limits of CATE estimation in a large non-parametric model.

3. General pseudo-outcome regression

In this section we lay out a framework for characterizing the error of general two-stage regression estimators that regress estimated “pseudo-outcomes” on a covariate vector. We define a new notion of stability (which we prove holds for generic linear smoothers, for example), and then show how this stability can be combined with a small-bias condition to yield oracle efficiency. In addition to laying a foundation for the analysis of a doubly robust CATE estimator in the next section, the result should also be of independent interest in other problems involving regression with estimated or imputed outcomes [13, 1, 40, 41, 23, 6].

With a few exceptions, previous work on regression with estimated outcomes has not appeared to exploit sample splitting, and has largely focused on particular pseudo-outcomes, and particular regression estimators (in both stages); in contrast we lean on sample splitting in order to obtain a more general result that is agnostic about the methods used, beyond a basic stability condition. Prominent examples of previous work include Ai and Chen [1], Rubin and van der Laan [40], and Foster and Syrgkanis [15], all of which gave results for pseudo-outcomes of a general form. Ai and Chen [1] and Rubin and van der Laan [40] did not use sample splitting and so limited their attention to particular estimators. Ai and Chen [1] used sieves, and focused more on a finite-dimensional component appearing in the pseudo-outcome. Rubin and van der Laan [40] considered least squares, penalized methods, and linear smoothers in the second stage, while being agnostic about the first-stage regression. However, their error bounds do not allow one to exploit double robustness in the pseudo-outcome, which is a crucial advantage in practice and which our result does allow.

Our results are closest in spirit to Foster and Syrgkanis [15], who study generic empirical risk minimization when the loss involves complex nuisance functions. However there are some important distinctions to be made. Most importantly, Foster and Syrgkanis [15] assume the loss satisfies a Neyman orthogonality property in order to obtain squared error rates; in contrast, our bound does not rely on this structure, but can exploit it when it holds, as shown in the following section. Crucially, our bound will also be seen to yield doubly robust errors,

whereas the orthogonality-based results of Foster and Syrgkanis [15] yield errors that are second-order but not doubly robust. Double robustness is essential when different nuisance components are estimated with different errors. Lastly, Foster and Syrgkanis [15] pursue global error bounds, while we consider pointwise error, which can be beneficial for developing inferential tools. These and some other differences are discussed further in the next section.

In what follows we first define a new notion of estimator stability, which can be viewed as a form of stochastic equicontinuity for nonparametric regression (as opposed to averaging).

Definition 1 (Stability). Suppose $D^n = (Z_{01}, \dots, Z_{0n})$ and $Z^n = (Z_1, \dots, Z_n)$ are independent training and test samples, respectively, with covariate $X_i \subset Z_i$. Let:

1. $\hat{f}(z) = \hat{f}(z; D^n)$ be an estimate of a function $f(z)$ using the training data D^n ,
2. $\hat{b}(x) = \hat{b}(x; D^n) \equiv \mathbb{E}\{\hat{f}(Z) - f(Z) \mid D^n, X = x\}$ the conditional bias of the estimator \hat{f} ,
3. $\hat{\mathbb{E}}_n(Y \mid X = x)$ denote a generic regression estimator that regresses outcomes (Y_1, \dots, Y_n) on covariates (X_1, \dots, X_n) in the test sample Z^n .

Then the regression estimator $\hat{\mathbb{E}}_n$ is defined as *stable* at $X = x$ (with respect to a distance metric d) if

$$\frac{\hat{\mathbb{E}}_n\{\hat{f}(Z) \mid X = x\} - \hat{\mathbb{E}}_n\{f(Z) \mid X = x\} - \hat{\mathbb{E}}_n\{\hat{b}(X) \mid X = x\}}{\sqrt{\mathbb{E}(\{\hat{\mathbb{E}}_n\{f(Z) \mid X = x\} - \mathbb{E}\{f(Z) \mid X = x\}\}^2)}} \xrightarrow{p} 0 \quad (3.1)$$

whenever $d(\hat{f}, f) \xrightarrow{p} 0$.

Sample splitting plays an important role here: the regression procedure as defined estimates the pseudo-outcome on a separate sample, independent from the one used in the second-stage regression via $\hat{\mathbb{E}}_n$. Examples of such procedures can be found in subsequent sections, e.g., as illustrated in Figures 3 and 5. The main role of sample splitting is that it allows for informative error analysis while being agnostic about the first- and second-stage methods.

Remark 2. With iid data, one can always obtain separate independent samples by randomly splitting the data in half (or in folds); further, to regain full sample size efficiency one can always swap the samples, repeat the procedure, and average the results, popularly called cross-fitting and used for example by Bickel and Ritov [7], Robins et al. [35], Zheng and van der Laan [56], and Chernozhukov et al. [10]. In this paper, to simplify notation we always analyze a single split procedure, with the understanding that extending to an analysis of an average across independent splits is straightforward.

As mentioned above, Definition 1 can be viewed as a generalization of the classic condition that

$$\frac{(\mathbb{P}_n - \mathbb{P})(\hat{f} - f)}{1/\sqrt{n}} \xrightarrow{p} 0$$

when $\|\widehat{f} - f\| \xrightarrow{P} 0$ (cf. Andrews [2], Lemma 19.24 of van der Vaart [52], Lemma 2 of Zheng and van der Laan [56], term $\mathcal{I}_{3,k}$ of Chernozhukov et al. [10], Lemma 2 of Kennedy et al. [24], etc.). More specifically, Definition 1 generalizes from settings involving averages \mathbb{P}_n to generic regressions $\widehat{\mathbb{E}}_n$, where slower than root-n rates can appear in the denominator scaling. We also note that the stability in Definition 1 is local or pointwise, in the sense that it is defined at a particular $X = x$ (because pointwise rates are the focus of this paper). More global variants of stability are also possible; for example, after this work appeared on arxiv, Rambachan et al. [31] used an L_2 version of stability (their Assumption B.1 and Proposition B.1) to study L_2 rates (their Lemma B.1).

The next result shows that generic linear smoothers are stable, with respect to a natural weighted L_2 distance. This and all other proofs are given in Appendix A. Linear smoothers are a fundamental class of regression estimators and include, for example, linear regression, series methods, local polynomials, nearest neighbor matching, smoothing splines, kernel ridge regression, and some versions of random forests [5, 42, 54]. We suspect other kinds of estimators that are not linear smoothers are also stable in the sense of Definition 1, but leave this to future work.

Theorem 1. *Linear smoothers $\widehat{\mathbb{E}}_n\{f(Z) \mid X = x\} = \sum_i w_i(x; X^n)\widehat{f}(Z_i)$ are stable in the sense of Definition 1, with respect to distance*

$$d(\widehat{f}, f) = \|\widehat{f} - f\|_{w^2} \equiv \sum_{i=1}^n \left\{ \frac{w_i(x; X^n)^2}{\sum_j w_j(x; X^n)^2} \right\} \int \{\widehat{f}(z) - f(z)\}^2 d\mathbb{P}(z \mid X_i),$$

whenever $1/\|\sigma\|_{w^2} = O_{\mathbb{P}}(1)$ for $\sigma(x)^2 = \text{var}\{f(Z) \mid X = x\}$.

We note that Theorem 1 recovers results similar to Lemma 2 of Kennedy et al. [24] for averages, i.e., when the weights all equal $w_i(x; X^n) = n^{-1}$, but also covers the much larger class of all linear smoothers. The norm $\|\cdot\|_{w^2}$ defining the relevant distance $d(\widehat{f}, f)$ for linear smoothers is a natural weighted and conditional version of an $L_2(\mathbb{P})$ norm. Namely, it is a weighted average of the conditional $L_2(\mathbb{P})$ norm (given X), weighted more towards the point $X = x$ depending on how localized the weights $w_i(x; X^n)$ are.

The next result shows how stability and consistency of \widehat{f} yield a rate of convergence result, relative to an oracle estimator that regresses the true unknown $f(Z)$ on X , and that a further small-bias condition yields oracle efficiency (i.e., asymptotic equivalence to the oracle estimator). This lays out a recipe for deriving convergence rates and conditions for oracle efficiency for general pseudo-outcome regression problems, which we will use in the CATE setup in the next section.

Proposition 1. *Under the same setup from Definition 1, also define:*

- i. $m(x) = \mathbb{E}\{f(Z) \mid X = x\}$ the conditional expectation of $f(Z)$ given X ,
- ii. $\widehat{m}(x) = \widehat{\mathbb{E}}_n\{\widehat{f}(Z) \mid X = x\}$ the regression of $\widehat{f}(Z)$ on X in the test samples,
- iii. $\widetilde{m}(x) = \widehat{\mathbb{E}}_n\{f(Z) \mid X = x\}$ the corresponding oracle regression of $f(Z)$ on X ,

and the oracle mean squared error $R_n^*(x)^2 = \mathbb{E}[\{\tilde{m}(x) - m(x)\}^2]$. If:

1. the regression estimator $\widehat{\mathbb{E}}_n$ is stable with respect to distance d , and
2. $d(\widehat{f}, f) \xrightarrow{P} 0$,

then

$$\widehat{m}(x) - \tilde{m}(x) = \widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} + o_{\mathbb{P}}(R_n^*(x)).$$

Therefore if it further holds that

3. $\widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} = o_{\mathbb{P}}(R_n^*(x))$,

then \widehat{m} is oracle efficient, i.e., asymptotically equivalent to the oracle estimator \tilde{m} in the sense that

$$\frac{\widehat{m}(x) - \tilde{m}(x)}{R_n^*(x)} \xrightarrow{P} 0.$$

We now go on to discuss each component of Proposition 1. The oracle risk $R_n^*(x)$ will typically be known (at least up to constants) based on structural assumptions on the target regression function $m(x)$ as well as the form of the estimator $\widehat{\mathbb{E}}_n$. For example, if m is s -smooth and $\widehat{\mathbb{E}}_n$ is an appropriate minimax optimal estimator (e.g., based on series or local polynomial regression) then $R_n^*(x) \asymp n^{-1/(2+d/s)}$. Similarly, if m is s -sparse and $\widehat{\mathbb{E}}_n$ is an appropriate minimax optimal estimator (e.g., lasso) then $R_n^*(x) \asymp \sqrt{s \log d/n}$, under some assumptions. The stability condition was described previously; as mentioned there, it holds for linear smoothers and we expect it to hold more generally. The consistency condition $d(\widehat{f}, f) \xrightarrow{P} 0$ is mild since it does not impose any particular rate requirement, and would generally follow from standard regression results (depending on the form of \widehat{f}); although consistency at any rate is sufficient for asymptotic negligibility, finite-sample behavior could depend on the rate at which $d(\widehat{f}, f)$ converges to zero.

Under just stability of $\widehat{\mathbb{E}}_n$ and consistency of \widehat{f} , one obtains the rate of convergence result

$$\widehat{m}(x) - m(x) = \tilde{m}(x) - m(x) + \widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} + o_{\mathbb{P}}(R_n^*(x)).$$

The bias term $\widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\}$ (recalling $\widehat{b}(x) = \mathbb{E}\{\widehat{f}(Z) - f(Z) \mid D^n, X = x\}$) is the crucial piece that will generally determine whether the estimator \widehat{m} is oracle efficient; if not, the rate would follow from this bias term. The bias $\widehat{b}(x)$ can in be determined on a case-by-case basis, depending on the estimator \widehat{f} . For simple plug-in estimators of f , the bias would typically have a first-order dependence on the estimation error in whatever nuisance functions appear in f . However, in some cases one will be able to construct influence-function-based estimators of f that have only a second-order dependence on nuisance estimation error. The availability of such estimators will depend on the form of m , but would be expected to exist when m is a structured combination of regressions or densities (e.g., a dose-response curve, CATE, counterfactual density, etc.). We refer to Section 5.3 of Kennedy [22] for some related general discussion of using influence functions for non-pathwise differentiable quantities like m .

Note $\widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\}$ is actually an estimate of the regression of the bias term $\widehat{b}(X)$ at $X = x$; thus it can be viewed as a smoothed bias, which may be more or less averaged or localized depending on the form of $\widehat{\mathbb{E}}_n$. Since the form of $\widehat{b}(x)$ will often be well-understood (e.g., derived analytically based on the form of \widehat{f}), in the following proposition we relate it to its smoothed analog $\widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\}$ in the linear smoother case.

Proposition 2. *If $\widehat{b}(x) = \widehat{b}_1(x)\widehat{b}_2(x)$, and $\widehat{\mathbb{E}}_n$ is a linear smoother with*

$$\sum_i |w_i(x; X^n)| = O_{\mathbb{P}}(c_n),$$

then

$$\widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} = O_{\mathbb{P}}(c_n \|\widehat{b}_1\|_{w,p} \|\widehat{b}_2\|_{w,q})$$

for the norm $\|f\|_{w,p} = [\sum_i \{ \frac{|w_i(x; X^n)|}{\sum_j |w_j(x; X^n)|} \} |f(X_i)|^p]^{1/p}$ and $1/p + 1/q = 1$ ($p, q > 1$).

Proposition 2 shows that the smoothed bias can be expressed in terms of natural weighted norms of the components of the bias $\widehat{b}(x)$ itself. For many linear smoothers $\sum_i |w_i(x; X^n)| \leq C$ with probability one, so that $c_n = 1$ (e.g., this is a condition in the celebrated theorem of Stone [45] guaranteeing weak universal consistency of linear smoothers). For series estimators with $w_i(x; X^n) = \frac{1}{n} \rho(x)^\top \widehat{Q}_{hx}^{-1} \rho(X_i)$ where ρ is some basis of dimension k and $\widehat{Q}_{hx} = \mathbb{P}_n(\rho \rho^\top)$, a simple bound can yield $c_n = \sup_x \sqrt{\rho(x)^\top \rho(x)}$ (which is of order \sqrt{k} for many series). We suspect the dependence on k can be avoided with more careful analysis, or when working with $L_2(\mathbb{P})$ norms.

4. DR-Learner

In this section we analyze a two-stage doubly robust estimator we refer to as the DR-Learner, following the naming scheme from Nie and Wager [30] and Künzel et al. [25]. After detailing previous work, we describe the algorithm in detail, and then give model-agnostic error bounds which apply for arbitrary first-stage estimators, and as long as the second-stage estimator is stable in the sense of Definition 1. We go on to apply the error bound in multiple nonparametric models incorporating smoothness or sparsity structure, and then explore the performance of the DR-Learner in simulations.

4.1. Previous work

Variants of the DR-Learner have been used before, though often tied to particular estimators and not incorporating sample splitting, which can allow for model-agnostic error bounds and reduced bias. van der Laan [48] (Section 4.2) appears to be the first to propose the general DR-Learner approach, i.e., flexible regressions of the pseudo-outcome (4.1) below on covariates. Specifically, van der

Laan [48] (along with van der Laan [49] and Luedtke and van der Laan [27]) advocates regressing the pseudo-outcome on $V \subseteq X$ to construct candidate CATE estimators, and then selecting among them with a tailored cross-validation approach [50]. The main distinction with our work is they did not give specific error guarantees for the CATE.

To the best of our knowledge, previous papers that do give specific error rates for the DR-Learner either employ stronger nuisance estimation conditions than we show are required, or else do not allow the CATE to be smoother or more structured than the individual regression functions μ_a . Lee et al. [26] studied a local-linear version of the DR-Learner, but assumed the first-stage nuisance error was negligible. Semenova and Chernozhukov [43] and Zimmert and Lechner [57] studied series and local-constant variants, respectively, but required conditions on nuisance estimation that are as restrictive as for the ATE. Fan et al. [14] also studied a local-constant variant, but did not consider the case where the CATE is smoother than the regression functions.

Our results are closest to Foster and Syrgkanis [15], who also considered the DR-Learner and gave an oracle inequality for generic empirical risk minimization when the loss involves complex nuisance functions. However there are some important distinctions to be made. First and most importantly, our error bound for the DR-Learner is doubly robust, involving a simple second-order product of nuisance errors. In contrast, the results of Foster and Syrgkanis [15] yield errors that are second-order but not doubly robust, instead involving L_4 errors of all nuisance components. This means our error bound is tighter when the propensity score and outcome regressions are estimated at different rates. Second, we focus on local estimation error at a point, whereas Foster and Syrgkanis [15] consider global error (e.g., integrated MSE); in this sense our work is complementary. A third important distinction is that our results can be used to justify the validity of inferential procedures, such as confidence intervals and hypothesis testing, whereas Foster and Syrgkanis [15] focus on global rates of convergence.

There is also a related literature on regression with estimated pseudo-outcomes, similar in spirit to Foster and Syrgkanis [15], though giving somewhat more specialized results. For example, Ai and Chen [1] and Rubin and van der Laan [40] both studied regression with general pseudo-outcomes, though they focused on particular estimators, and did not give doubly robust error bounds. Ai and Chen [1] restricted their attention to sieves, and focused more on a finite-dimensional component appearing in the pseudo-outcome. Rubin and van der Laan [40] considered least squares, penalized methods, and linear smoothers in the second stage, but their error bounds are not doubly robust.

Künzel et al. [25] considered various method-agnostic “meta-learners”, but not the DR-Learner; further, none of their methods are doubly robust, and so in general would inherit larger error rates from the underlying regression estimators.

4.2. Construction & analysis

The algorithm below describes our proposed construction of the DR-Learner.

Algorithm 1 (DR-Learner). Let (D_1^n, D_2^n) denote two independent samples of n observations of $Z_i = (X_i, A_i, Y_i)$.

Step 1. Nuisance training:

- (a) Construct estimates $\hat{\pi}$ of the propensity scores π using D_1^n .
- (b) Construct estimates $(\hat{\mu}_0, \hat{\mu}_1)$ of the regression functions (μ_0, μ_1) using D_1^n .

Step 2. Pseudo-outcome regression: Construct the pseudo-outcome

$$\hat{\varphi}(Z) = \frac{A - \hat{\pi}(X)}{\hat{\pi}(X)\{1 - \hat{\pi}(X)\}} \{Y - \hat{\mu}_A(X)\} + \hat{\mu}_1(X) - \hat{\mu}_0(X) \quad (4.1)$$

and regress it on covariates X in the test sample D_2^n , yielding

$$\hat{\tau}_{dr}(x) = \hat{\mathbb{E}}_n \{ \hat{\varphi}(Z) \mid X = x \}. \quad (4.2)$$

Step 3. Cross-fitting (optional): Repeat Step 1–2, swapping the roles of D_1^n and D_2^n so that D_2^n is used for nuisance training and D_1^n as the test sample. Use the average of the resulting two estimators as a final estimate of τ . K -fold variants are also possible.

Figure 3 gives a schematic illustrating the DR-Learner construction.




FIG 3. Schematic illustrating the DR-Learner approach. In the first stage, the nuisance functions $\hat{\pi}$ and $(\hat{\mu}_0, \hat{\mu}_1)$ are estimated from training sample D_1^n . In the second stage, these estimates are used to construct an estimate of the pseudo-outcome $\hat{\varphi}$, which is then regressed on X using the test sample D_2^n .

Remark 3. The DR-Learner approach is motivated by the fact that (4.1) is the (uncentered) efficient influence function for the ATE [33, 19]; this drives many of its favorable properties. The intuition is that, to efficiently estimate the ATE, the standard doubly robust estimator averages the pseudo-outcome $\hat{\varphi}$, so to estimate the CATE the DR-Learner regresses $\hat{\varphi}$ on covariates. For a review of influence functions and semiparametric theory we refer to van der Laan and Robins [51], Tsiatis [46], and Kennedy [22].

Next we present the main result of this section, which gives error bounds for the DR-Learner procedure (relative to an oracle) for arbitrary first-stage estimators, as long as the second-stage estimator is stable in the sense of Definition 1.

Theorem 2. *Let $\hat{\tau}_{dr}(x)$ denote the DR-Learner estimator detailed in Algorithm 1, which regresses the estimated pseudo-outcome $\hat{\varphi}(Z)$ on covariates X . Assume:*

1. *The second-stage regression estimator $\hat{\mathbb{E}}_n$ is stable with respect to distance d ,*
2. *$d(\hat{\varphi}, \varphi) \xrightarrow{P} 0$,*

Let $\tilde{\tau}(x) = \hat{\mathbb{E}}_n\{\varphi(Z) \mid X = x\}$ denote an oracle estimator that regresses the true pseudo-outcome $\varphi(Z)$ on X , and denote its risk by $R_n^(x)^2 = \mathbb{E}\{[\tilde{\tau}(x) - \tau(x)]^2\}$. Then*

$$\hat{\tau}_{dr}(x) - \tilde{\tau}(x) = \hat{\mathbb{E}}_n\{\hat{b}(X) \mid X = x\} + o_{\mathbb{P}}(R_n^*(x))$$

for

$$\hat{b}(x) = \sum_{a=0}^1 \frac{\{\hat{\pi}(x) - \pi(x)\}\{\hat{\mu}_a(x) - \mu_a(x)\}}{a\hat{\pi}(x) + (1-a)(1 - \hat{\pi}(x))}$$

and $\hat{\tau}_{dr}(x)$ is oracle efficient in the sense of Proposition 1 if $\hat{\mathbb{E}}_n\{\hat{b}(X) \mid X = x\} = o_{\mathbb{P}}(R_n^(x))$.*

The bound on the DR-Learner error given in Theorem 2 shows that it can only deviate from the oracle error by at most a (smoothed) product of errors in the propensity score and regression estimators, thus allowing faster rates for estimating the CATE even when the nuisance estimates converge at slower rates. Importantly the result is agnostic about the methods used, and requires no special tuning or undersmoothing.

Theorem 2 gives a smaller risk bound compared to earlier work. Semenova and Chernozhukov [43] and Zimmert and Lechner [57] gave a remainder error larger than the product of the nuisance mean squared errors, with oracle efficiency requiring this product to shrink faster than $1/nd$ and $1/(n/h)$ for series and kernel-based second stage regressions, respectively (for h a shrinking bandwidth). Fan et al. [14] assumed the CATE was only as smooth as the individual regression functions, giving a larger oracle risk. The results from Foster and Syrgkanis [15] are closest to ours, but as mentioned earlier their error bounds are not doubly robust, and instead involve L_4 errors of all nuisance components.

Remark 4. Several important oracle inequalities for cross-validated selection of estimators exist in the literature [50, 27, 12]. These are relevant for cross-validated CATE estimation, but are conceptually different from our Theorem 2; for example, they use a different oracle. Namely, in cross-validated selection the oracle is the best performing among a group of learners, whereas in our setup the oracle is the specified learner $\hat{\mathbb{E}}_n$ when it is given access to the true pseudo-outcome $\varphi(Z)$.

Remark 5. In addition to showing double robustness and giving conditions for oracle efficiency, Theorem 2 also has important implications for inference.

Namely, if the DR-Learner is oracle efficient so that $\hat{\tau}(x) - \tilde{\tau}(x) = o_{\mathbb{P}}(R_n^*(x))$, then whenever an inferential result is available for the oracle estimator $\tilde{\tau}(x)$ (which is just a standard regression of oracle pseudo-outcomes on covariates), this would also apply to $\hat{\tau}(x)$, immediately allowing for the construction of confidence intervals, tests, etc. For example, if one used a local polynomial estimator for the second-stage regression $\hat{\mathbb{E}}_n$, then under standard conditions [28] $\tilde{\tau}(x)$ would be asymptotically normal, and so $\hat{\tau}(x)$ would be as well, by virtue of the oracle equivalence. Thus standard confidence intervals could be constructed, treating the estimated pseudo-outcomes as one would usual observed outcomes in nonparametric regression. This is a benefit of the asymptotic equivalence to the oracle in Theorem 2. In contrast, the results from Foster and Syrgkanis [15], for example, instead give conditions under which the global risk of a DR-Learner is within a multiplicative constant of that of the oracle, which is not as immediately useful for inference.

4.3. Examples & illustrations

An important feature of the oracle result in Theorem 2 is that it is essentially model-free: beyond a mild consistency assumption, it only requires that the second-stage regression estimator satisfies the stability condition in Definition 1. In the following corollaries, we illustrate the flexibility of this result by applying it in settings where the nuisance functions and CATE are smooth or sparse (i.e., in settings where local polynomial, series, lasso, or random forest estimators would work well). Similar results could be obtained in generic models with known bounds on mean squared error rates.

Corollary 1. *Suppose the assumptions of Theorem 2 hold. Further assume:*

1. $\hat{\mathbb{E}}_n$ is a minimax optimal linear smoother with $\sum_i |w_i(x; X^n)| = O_{\mathbb{P}}(1)$.
2. The propensity score π is α -smooth, and $\|\hat{\pi} - \pi\|_{w,2} = O_{\mathbb{P}}(n^{-1/(2+d/\alpha)})$.
3. The regressions μ_a are β -smooth, and $\|\hat{\mu}_a - \mu_a\|_{w,2} = O_{\mathbb{P}}(n^{-1/(2+d/\beta)})$.
4. The CATE τ is γ -smooth.

Then

$$\hat{\tau}_{dr}(x) - \tau(x) = O_{\mathbb{P}}\left(n^{\frac{-1}{2+d/\gamma}} + n^{-(\frac{1}{2+d/\alpha} + \frac{1}{2+d/\beta})}\right)$$

and the DR-Learner is oracle efficient if

$$\sqrt{\alpha\beta} \geq \frac{d/2}{\sqrt{1 + \frac{d}{\gamma}(1 + \frac{d}{2\bar{s}})}}, \tag{4.3}$$

where $\bar{s} = (\frac{\alpha^{-1} + \beta^{-1}}{2})^{-1}$ is the harmonic mean of (α, β) .

Corollary (1) illustrates how the DR-Learner can adapt to smoothness in the CATE even when the propensity score and regression functions may be less smooth, and gives sufficient conditions for achieving the oracle rate depending on nuisance smoothness and dimension d .

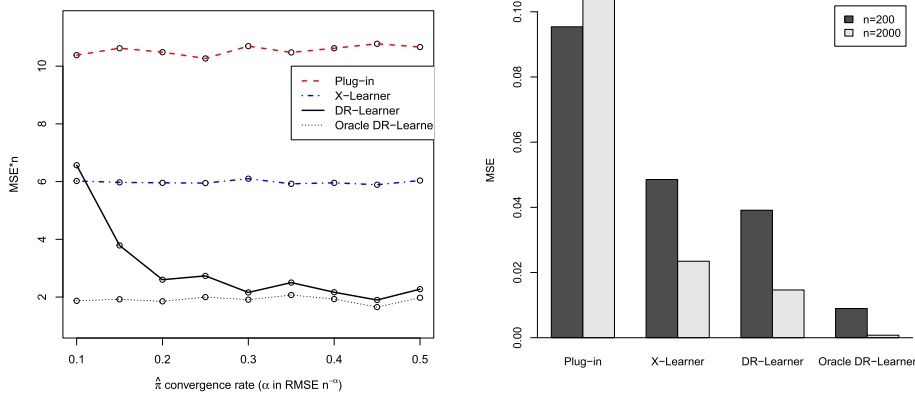
It is instructive to compare the sufficient condition in (4.3) to the analogous condition for root- n consistency of a standard doubly robust estimator of the average treatment effect, which is $\sqrt{\alpha\beta} \geq d/2$ (cf. Equation 25 of Robins et al. [36]). First, as the CATE smoothness gets larger, the sufficient condition (4.3) for the CATE approaches that for the ATE, as should be expected (intuitively, an infinitely smooth CATE should be nearly as easy to estimate as the ATE). Second, the term in the denominator of (4.3), dividing $d/2$, can be interpreted as a “lowered bar” for optimal estimation, due to the fact that the oracle rate $n^{\frac{-1}{2+d/\gamma}}$ is slower than root- n . This phenomenon was also noted in the dose-response estimation problem by Kennedy et al. [23], and is in contrast with the error bounds given in Nie and Wager [30] and Zimmert and Lechner [57], which required ATE-like conditions for oracle efficiency of CATE estimators, via $n^{-1/4}$ or faster rates on the nuisance estimators (note that if $\alpha = \beta$ then $\sqrt{\alpha\beta} \geq d/2$ means the nuisance error rate is faster than $n^{-1/4}$). In the next section we will show how the sufficient condition (4.3) can even be improved upon.

Similar results could be obtained for other models. For example, suppose the propensity score and regression functions are α - and β -sparse, respectively, and estimated at rates $\alpha \log d/2$ and $\beta \log d/n$, and similarly for the CATE with γ -sparsity. The corresponding condition for oracle efficiency would be $\alpha\beta \log^2 d/n^2 \leq \gamma \log d/n$, which is $\sqrt{\alpha\beta} \leq \gamma n/\log d$. However we note that we have not proved second-stage estimators attaining these rates satisfy the stability condition of Definition 1.

4.4. Simulation experiments

In this section we study some finite-sample properties via simulations (R code is available in Appendix B). We use four methods for CATE estimation: a plug-in that estimates the regression functions μ_0 and μ_1 and takes the difference (called the T-Learner by Künzel et al. [25]), the X-Learner from Künzel et al. [25], the DR-Learner from Section 4, and an oracle DR-Learner that uses the true pseudo-outcome in the second-stage regression.

First we use the piecewise polynomial model from the motivating example in Section 2.2, with outcome and 2nd-stage regressions fit with `smoothing.spline` in R. Figure 4a shows the mean squared error for the four CATE methods at $n = 2000$ (based on 500 simulations with MSE averaged over 500 independent test samples), across a range of convergence rates for the propensity score estimator $\hat{\pi}$. To control the convergence rate we constructed this estimator as $\hat{\pi} = \text{expit}\{\text{logit}(\pi) + \epsilon_n\}$, where $\epsilon_n \sim N(n^{-\alpha}, n^{-2\alpha})$ so that $\text{RMSE}(\hat{\pi}) \sim n^{-\alpha}$ (the estimator was truncated to lie within $[0.01, 0.99]$ to ensure positivity holds). The results show that the plug-in estimator inherits the large error in estimating the individual regression functions, while the DR-Learner achieves much smaller errors and adapts to the smoothness of the CATE. The X-Learner has MSE in between the two. Consistent with Theorem 2, the MSE of the DR-Learner approaches that of the oracle as the propensity score estimation error decreases (i.e., as the convergence rate gets faster).



(a) Piecewise polynomial model from 2.2 (b) High-dimensional model

FIG 4. Simulation results.

Next we consider a high-dimensional logistic model where $X = (X_1, \dots, X_d) \sim N(0, I_d)$, and $\text{logit}\{\pi(x)\} = \frac{1}{2\sqrt{\alpha}} \sum_{j=1}^{\alpha} x_j$, and $\text{logit}\{\mu_a(x)\} = \frac{1}{\sqrt{\beta}} \sum_{j=1}^{\beta} x_j$, so the propensity score and individual regressions are α and β sparse, respectively, while the CATE is zero. The normalization of the coefficients ensures $\pi(X) \in [0.2, 0.8]$ with high probability and similarly for μ_a . We use standard cross-validation-tuned lasso for all model-fitting, i.e., to estimate the propensity scores, regression functions, and second-stage fits. Figure 4b shows mean squared errors for the four CATE methods when $d = 500$ and $\alpha = \beta = 50$, for $n = 200$ and $n = 2000$, across 100 simulations (median of mean square errors is reported due to high skewness). As expected from Theorem 2, the DR-Learner is closer to the oracle than the plug-in or X-Learner, though this high-dimensional setup yields larger estimation error than the previous simulation model. The relative performance of the DR-Learner seems to improve with sample size.

5. Local polynomial R-learner

The previous section gave general sufficient conditions under which the DR-Learner attains the oracle error rate of an estimator with direct access to the difference $Y^1 - Y^0$, showing that this rate can in fact be achieved whenever the product of nuisance errors is of smaller order. This raises the crucial question of what happens when this product is not sufficiently small: in such regimes, is there any hope at still attaining the oracle error rate?

The current section provides a first answer to this question, with a more refined analysis of a different estimator. Specifically, we analyze a double-sample-split local polynomial adaptation of the R-Learner [30, 39], which we call the lp-R-Learner for short. The R-Learner of Nie and Wager [30] is a nonparametric

RKHS regression-based extension of the double-residual regression method of Robinson [39]. A nonparametric series-based version of the R-Learner was also proposed in Example 4 of Robins et al. [35], though assuming known propensity scores and not incorporating outcome regression. Chernozhukov et al. [9] studied a lasso version of the R-learner. Some important distinctions between our results and those in previous work include the following: (i) our estimator is built from local polynomials, and incorporates a specialized form of sample splitting inspired by Newey and Robins [29] for bias reduction, (ii) our sufficient conditions for attaining oracle efficiency are substantially weaker than the $n^{-1/4}$ rates in Nie and Wager [30] and Chernozhukov et al. [9], and (iii) we give specific rates of convergence outside the oracle regime.

We first describe the lp-R-Learner in detail, then give the main error bound result, which holds under a Hölder-smooth model, and is valid for a wide variety of tuning parameter choices. Following that, we optimize the bound with specific tuning parameter choices, under different sets of conditions, and discuss the resulting rates.

5.1. Construction

The algorithm below describes the lp-R-Learner construction.

Algorithm 2 (lp-R-Learner). *Let $(D_{1a}^n, D_{1b}^n, D_2^n)$ denote three independent samples of n observations of $Z_i = (X_i, A_i, Y_i)$.*

Let $b : \mathbb{R}^d \mapsto \mathbb{R}^p$ denote the vector of basis functions consisting of all powers of each covariate, up to order $\lceil \gamma \rceil$, and all interactions up to degree $\lceil \gamma \rceil$ polynomials (cf. Masry [28]). Let $K_{hx}(X) = \frac{1}{h^d} K(\frac{X-x}{h})$ for $K : \mathbb{R}^d \mapsto \mathbb{R}$ a bounded kernel function with support $[-1, 1]^d$, and h a bandwidth parameter.

Step 1. Nuisance training:

- (a) *Using D_{1a}^n , construct estimates $\hat{\pi}_a$ of the propensity scores π .*
- (b) *Using D_{1b}^n , construct estimates $\hat{\eta}$ of the regression function $\eta = \pi\mu_1 + (1 - \pi)\mu_0$, and estimates $\hat{\pi}_b$ of the propensity scores π .*

Step 2. Localized double-residual regression:

Define $\hat{\tau}_r(x)$ as the fitted value from a kernel-weighted least-squares regression (in the test sample D_2^n) of outcome residual $(Y - \hat{\eta})$ on basis terms b scaled by the treatment residual $(A - \hat{\pi}_b)$, with weights $(\frac{A - \hat{\pi}_a}{A - \hat{\pi}_b})K_{hx}$. Thus $\hat{\tau}_r(x) = b(0)^\top \hat{\theta}$ for $\hat{\theta}$ the minimizer over $\theta \in \mathbb{R}^p$ of

$$\mathbb{P}_n \left(K_{hx}(X) \left\{ \frac{A - \hat{\pi}_a(X)}{A - \hat{\pi}_b(X)} \right\} \left[\{Y - \hat{\eta}(X)\} - \theta^\top b(X - x) \{A - \hat{\pi}_b(X)\} \right]^2 \right). \quad (5.1)$$

Step 3. Cross-fitting (optional): Repeat Step 1–2 twice, first using (D_{1b}^n, D_2^n) for nuisance training and D_{1a}^n as the test sample, and then using (D_{1a}^n, D_2^n)

for training and D_{1b}^n as the test sample. Use the average of the resulting three estimators of τ as the final estimator $\hat{\tau}_r$.

Remark 6. The kernel weights in the second step regression need to be multiplied by the ratio $(A - \hat{\pi}_a)/(A - \hat{\pi}_b)$ in order to ensure the independence of relevant products of nuisance estimators (i.e., that they are built from separate samples D_{1a}^n and D_{1b}^n). This allows for multiplicative biases and thus faster rates due to undersmoothing, first introduced by Newey and Robins [29] but for \sqrt{n} -estimable functionals. In other words, this ensures the bias of the lp-R-Learner equals a product of biases of the nuisance estimators; other, different ratios used in this kernel weight would therefore generally not work in the same way.

Figure 5 gives a schematic illustrating the lp-R-Learner construction.




FIG 5. Schematic illustrating the lp-R-Learner approach. In the first stage, the nuisance functions $\hat{\pi}_a$ and $(\hat{\pi}_b, \hat{\eta})$ are estimated from training samples D_{1a}^n and D_{1b}^n , respectively. In the second stage, these estimates are used in a kernel-weighted least squares regression of residuals $(Y - \hat{\eta})$ on residual-scaled basis terms $(A - \hat{\pi}_b)b$, with weights $(\frac{A - \hat{\pi}_a}{A - \hat{\pi}_b})K_{hx}$.

5.2. Main error bound \mathcal{E} oracle results

Before giving the main error bound in this section, we first present the following condition on the nuisance estimators that we use in the analysis. At a high level, this condition requires the nuisance estimators to be linear smoothers with particular bias and variance bounds.

Condition 1. The nuisance estimators $(\hat{\pi}_a, \hat{\pi}_b, \hat{\eta})$ are linear smoothers of the form

$$\hat{\pi}_j(x) = \sum_{i \in D_{1j}^n} w_{i\alpha}(x; X_{1j}^n) A_i \tag{1a}$$

$$\widehat{\eta}(x) = \sum_{i \in D_{1b}^n} w_{i\beta}(x; X_{1b}^n) Y_i$$

with weights $w_i(x; X_{1b}^n)$ depending on tuning parameter k , which are localized in the sense that

$$w_i(x; X_{1b}^n) = 0 \text{ whenever } \|X_i - x\| > 1/k^{1/d} \quad (1b)$$

and which satisfy the conditional bias and variance bounds

$$\begin{aligned} |\mathbb{E}\{\widehat{\pi}_j(x) \mid X_{1j}^n\} - \pi(x)| &\lesssim k^{-\alpha/d} & \text{var}\{\widehat{\pi}_j(x) \mid X_{1j}^n\} &\lesssim k/n \\ |\mathbb{E}\{\widehat{\eta}(x) \mid X_{1b}^n\} - \eta(x)| &\lesssim k^{-\beta/d} & \text{var}\{\widehat{\eta}(x) \mid X_{1b}^n\} &\lesssim k/n. \end{aligned} \quad (1c)$$

Conditions (1a)–(1c) are relatively standard. Many popular estimators take the form given in (1a), as discussed just before Theorem 1. Condition (1b) holds for several prominent linear smoothers. For example, it holds for series estimators built from k basis terms, if properly localized, and for standard kernel or local polynomial estimators when taking the bandwidth parameter as $h \sim k^{-d}$, as shown for example in Proposition 1.13 of Tsybakov [47].

Condition (1c) also has been shown to hold for series and local polynomial estimators, for example, when the underlying regression function is appropriately smooth. In particular, under standard conditions, the bias part of (1c) would hold for these methods when the propensity score π is α -smooth and the regression function η is β -smooth; we again refer to Belloni et al. [4] and Tsybakov [47] for a review of related results. For (1c) to hold uniformly over all $x \in \mathcal{X}$ would typically mean these bounds would only hold up to log factors; however our result will only require the bias to be controlled locally, near the point at which the CATE is to be estimated.

The next result gives error bounds on the lp-R-Learner from Algorithm 2, under Hölder smoothness conditions.

Theorem 3. *Let $\widehat{\tau}_r(x)$ denote the lp-R-Learner estimator detailed in Algorithm 2. Assume:*

1. *The estimator $\widehat{\eta}$ and observations Z are bounded, and X has density bounded above.*
2. *The estimators $\widehat{\pi}_j$ satisfy $\mathbb{P}\{\epsilon \leq \widehat{\pi}_j(x) \leq 1 - \epsilon\} = 1$ for some $\epsilon > 0$.*
3. *The eigenvalues of the sample Gram matrices \widehat{Q}_{hx} and \widetilde{Q}_{hx} defined in (A.7) are bounded away from zero in probability.*
4. *The nuisance estimators $(\widehat{\pi}_a, \widehat{\pi}_b, \widehat{\eta})$ satisfy Condition 1, with the bias and variance bounds holding for all x' such that $\|x' - x\| \leq h$.*

Let $s = \frac{\alpha+\beta}{2}$ denote the average smoothness of the propensity score and regression function. Then, if the CATE $\tau(x)$ is γ -smooth and $k^{-(\alpha \wedge \beta)/d} \lesssim k/n$, we have

$$\widehat{\tau}_r(x) - \tau(x) = O_{\mathbb{P}}\left(h^\gamma + k^{-2s/d} + k^{-2\alpha/d} + \frac{1}{\sqrt{nh^d}}\left(1 + \frac{k}{n}\right)\right).$$

Before detailing the result and implications of Theorem 3, we first discuss the assumptions. The first part of Assumption 1 is mostly to simplify presentation, and could be weakened at the cost of added complexity; the second part ensuring X has bounded density is more crucial, but still mild. Assumption 2 is standard in the causal literature, and in theory could be guaranteed by simply thresholding propensity score estimates; however, extreme propensity values (i.e., positivity violations) are an important issue in practice, especially in the nonparametric and high-dimensional setup [11]. Assumption 3 is relatively standard (see, e.g., Assumption (LP1) of Tsybakov [47]) but would restrict how n and h scale; when d is fixed, standard bandwidth choices should suffice. Assumption 4 is arguably most crucial, and does the most work in the proof (together with the specialized sample splitting); however, as detailed in the discussion of Condition 1 prior to the theorem statement, Assumption 4 uses standard conditions commonly found in the nonparametric regression literature.

Now we give some discussion and interpretation of the (in-probability) error bound of Theorem 3. The first three terms are the bias, and the last two are the variance (on the standard deviation scale). The bias has three components. The first h^γ bias term comes from the bias of an oracle estimator with access to the true propensity score and regression function, and matches the bias of an oracle with direct access to the difference $Y^1 - Y^0$. The other two bias terms come from nuisance estimation: the first $k^{-2s/d}$ term is the product of the biases of the propensity score and regression estimators, whereas the second $k^{-2\alpha/d}$ term is the squared bias of the propensity score estimator. If the propensity score is at least as smooth as the regression function, then the first $k^{-2s/d}$ term will dominate. Some heuristic intuition about why these specific bias terms arise is as follows: by virtue of its least squares construction, the lp-R-Learner can be viewed as a product of the inverse of an “ $X^T X$ -like” term involving products of $\hat{\pi}_a$ and $\hat{\pi}_b$, and an “ $X^T Y$ -like” term involving products of $\hat{\pi}_a$ and $\hat{\eta}$.

The variance has two components. As with the bias, the first $1/\sqrt{nh^d}$ term is the standard deviation of an oracle estimator with access to the true nuisance functions. The second term is a product of this oracle standard error with k/n , which is the additional variance coming from nuisance estimation (in fact k/n is the product of the standard deviations of the nuisance estimators). In standard regimes the tuning parameter k would have to be chosen of smaller order than n (e.g., $k \log k/n \rightarrow 0$ as in Belloni et al. [4] and elsewhere) in order for Condition 1 to hold, making the variance contribution from nuisance estimation asymptotically negligible. This last point will be discussed further shortly.

Now we give conditions under which the oracle rate is achieved by the lp-R-Learner, which we conjecture are not only sufficient but also necessary conditions.

Corollary 2. *Suppose the assumptions of Theorem 3 hold. Further assume $\alpha \geq \beta$ so the propensity score is smoother than the regression function, and take*

1. $h \sim n^{-1/(2\gamma+d)}$, and
2. $k \sim n/\log^2 n$ so that $k \log k/n \rightarrow 0$.

Then, up to log factors,

$$\widehat{\tau}_r(x) - \tau(x) = O_{\mathbb{P}}(n^{-\gamma/(2\gamma+d)} + n^{-2s/d})$$

and the oracle rate is achieved if the average nuisance smoothness satisfies $s \geq \frac{d/4}{1+d/2\gamma}$.

Corollary 2 shows that an undersmoothed lp-R-Learner can be oracle efficient under weaker conditions than the error bound given in Theorem 2 would indicate for the DR-Learner.

Remark 7. Note taking $k \sim n/\log^2 n$ amounts to undersmoothing as much as possible: it drives down nuisance bias, letting the variance k/n go to zero only very slowly, since the contribution of the latter is asymptotically negligible for CATE estimation as long as $k \lesssim n$. In general, undersmoothing can often require specific knowledge of smoothness parameters, and data-driven approaches are often lacking. However, it is worth noting that in Corollary 2 the choice of k only depends on the sample size, and not on the underlying nuisance smoothness, for example, making it more feasible to implement. For example, one could set k and then use cross-validation or other approaches to select h [50, 6]; however we leave a formal study of tuning parameter selection to future work.

The result in Corollary 2 is reminiscent of a similar phenomenon in marginal ATE estimation, where Robins et al. [37] showed that the condition $s \geq d/4$ is necessary and sufficient for the existence of root- n consistent estimators of the ATE functional $\mathbb{E}\{\mathbb{E}(Y | X, A = 1)\}$. Our result thus shows that $s \geq d/4$ is also sufficient for oracle efficient estimation of the CATE, but now the oracle rate is $n^{-\gamma/(2\gamma+d)}$ rather than root- n , and so there is in fact a weaker bar for oracle efficiency (namely $s \geq \frac{d/4}{1+d/2\gamma}$), depending on the smoothness γ . At one extreme, when the CATE is infinitely smooth, the condition we give for oracle efficiency of the lp-R-Learner recovers the usual $s \geq d/4$ condition for the ATE. At the other extreme, when the CATE is non-smooth, oracle efficiency can be much easier to achieve (e.g., if the CATE is only $\gamma = 1$ -smooth, it is only required that $s \geq 1/2$ even for arbitrarily large dimension d). We conjecture that the condition $s \geq \frac{d/4}{1+d/2\gamma}$ is not only sufficient but may also be necessary for oracle efficiency in the above Hölder model, making the proposed lp-R-Learner minimax optimal in this regime when $\alpha \geq \beta$; however, we leave a proof of this to future work.

When $s < \frac{d/4}{1+d/2\gamma}$ and the oracle rate is not achieved, the rate $n^{-2s/d}$ is slower than the usual functional estimation rate $n^{-4s/(4s+d)}$, which is minimax optimal for the ATE if the covariate density is smooth enough [37], and also occurs for simpler functionals like the expected density [7, 8]. To illustrate this gap, if $s = d/8$ then the rate in Corollary 2 is $n^{-1/4}$ while the usual functional minimax rate is $n^{-1/3}$.

Remark 8. There is a trade-off between the DR-Learner and lp-R-Learner. In short, the DR-Learner provides somewhat more general guarantees, and is computationally more straightforward, while the lp-R-Learner can achieve faster

rates in somewhat more specialized Holder smoothness models, but is also somewhat more difficult to compute in practice, and requires undersmoothing.

5.3. Faster rates with known covariate density

Here we briefly consider how the lp-R-Learner rates from Corollary 2 can be improved by exploiting structure in the covariate density, as in Robins et al. [35, 38]. This is somewhat paradoxical since the CATE itself does not depend on the covariate density.

Some intuition for a possible rate improvement is as follows. If the density of the covariates X was known, then one could construct nuisance estimators $(\hat{\pi}_a, \hat{\pi}_b, \hat{\eta})$ satisfying Condition 1 without any restriction on the tuning parameter k , such as the $k \log k/n \rightarrow 0$ restriction employed in Corollary 2. For example, a series estimator $\hat{\eta}$ could satisfy Condition 1 for any choice of k , if it took the form

$$\hat{\eta}(x) = b(x)^T [\mathbb{E}\{b(X)b(X)^T\}]^{-1} \mathbb{P}_n\{b(X)Y\} \tag{5.2}$$

where b is a vector of appropriate basis functions (different from those used in the lp-R-Learner construction), and similarly for $\hat{\pi}$. Intuitively this is because restrictions like $k \log k/n \rightarrow 0$ are only required to ensure the inverse of the sample Gram matrix $\mathbb{P}_n(bb^T)$ converges to a limit (in operator norm), whereas if the density of X was known, then one could simply compute the population Gram matrix $\mathbb{E}(bb^T)$ exactly.

We conjecture however that the $n^{-2s/d}$ rate from Corollary 2 may not be improvable in the non-random fixed X case. A fixed design setup was recently considered by Gao and Han [17], though in a somewhat specialized model where the propensity scores have zero smoothness. In fact when $\alpha = 0$ our rate matches their lower bound in the non-oracle regime.

The next result gives rates for the lp-R-Learner when there are no restrictions on the choice of nuisance tuning parameter k .

Corollary 3. *Suppose the assumptions of Theorem 3 hold, and in particular suppose Assumption 4 holds without any restrictions on k (e.g., as if the density of X is known). Then if $\alpha \geq \beta$ the convergence rate in Theorem 3 is optimized by taking $h \sim n^{\frac{-3s}{2s\gamma+(s+\gamma)d}}$ and $k \sim n^{\frac{3\gamma d/2}{2s\gamma+(s+\gamma)d}}$ which gives*

$$\hat{\tau}(x) - \tau(x) = O_{\mathbb{P}}\left(n^{-\gamma/(2\gamma+d)} + n^{\frac{-3s}{2s+d(1+s/\gamma)}}\right),$$

and so the oracle rate is achieved if $s \geq \frac{d/4}{1+d/2\gamma}$.

Remark 9. When the density of the covariates X is unknown but smooth and estimable at fast enough rates, we expect results similar to those in Corollary 3 to still hold. This phenomenon also occurs for the ATE functional [35, 38]. However, here the analysis of the lp-R-Learner becomes much more complicated, so we leave this to future work.

When the nuisance tuning parameter k is unrestricted, one needs to balance all three dominant terms from Theorem 3: the two bias terms h^γ and $k^{-2s/d}$, as well as the variance term $\frac{k/n}{\sqrt{nh^d}}$. Notice the “elbow” at $s \geq \frac{d/4}{1+d/2\gamma}$ occurs here even when the density is known, perhaps again suggesting that this condition for oracle efficiency may be both sufficient and necessary. Whether the rate $n^{\frac{-3s}{2s+d(1+s/\gamma)}}$ is minimax optimal or not in the non-oracle regime is unknown; we will pursue this in future work.

Interestingly, the rate $n^{\frac{-3s}{2s+d(1+s/\gamma)}}$ is slightly slower than the usual functional estimation rate $n^{-4s/(4s+d)}$. For example, when $\gamma \rightarrow \infty$ and $s = d/8$, the CATE rate from Corollary 3 is $n^{-3/10}$ whereas the classic functional estimation rate is $n^{-1/3}$. Therefore there do appear to be benefits of exploiting structure in the covariate density, but whether the gap between the improved rate of Corollary 3 and the usual functional rate can be closed is unclear. An interesting more philosophical question is whether structure in the covariate density should be exploited for CATE estimation in practical settings, since the CATE itself does not depend on the covariate density.

6. Discussion

In this paper we studied the problem of estimating conditional treatment effects, giving new results that apply to a wider variety of methods and that better exploit when the CATE itself is structured, compared to the current state-of-the-art. Sections 3 and 4 were more practically oriented, giving model-free yet informative error bounds for general regression with estimated outcomes, and the DR-Learner method for CATE estimation, along with examples from smooth and sparse models, illustrating the flexibility of Proposition 1 and Theorem 2. In contrast, Section 5 was more theoretically oriented, aiming instead at understanding the fundamental statistical limits of CATE estimation, i.e., the smallest possible achievable error. We derived upper bounds on this error with a specially constructed (and tuned) estimator called the lp-R-Learner. Namely we showed that in a Hölder model, oracle efficiency is possible under weaker conditions on the nuisance smoothness than indicated from the DR-Learner error bound given in Theorem 2 – which itself is an improvement over conditions given in previous work.

Figure 6 summarizes the various error rates in this paper graphically, in an illustrative setup where the dimension is $d = 20$, CATE smoothness is $\gamma = 2d$, and nuisance smoothness $\alpha = \beta = s$ varies from 0–10. This shows the various gaps between methods/rates, in an interesting regime where the CATE is relatively smooth compared to the nuisance functions.

Our work raises some interesting open questions, some of the more immediate of which we list here for reference:

1. Can the error bounds in Proposition 1 and Theorem 2 be improved without committing to particular first- or second-stage methods?
2. What rates can be achieved by specialized sample splitting and tuning of a DR-Learner, rather than an R-Learner?

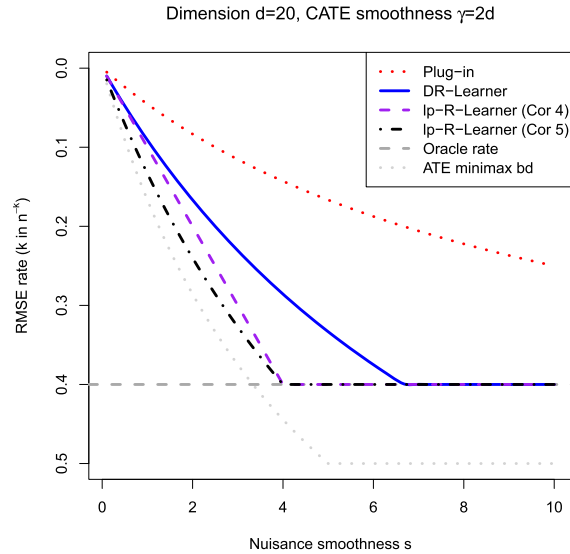


FIG 6. Illustration of error bounds in this work, as a function of nuisance smoothness $s = \alpha = \beta$. The dotted gray line represents the classical minimax lower bound for functional estimation, which is \sqrt{n} when $s \geq d/4$. The dashed gray line is the oracle rate that would be achieved by a minimax optimal regression using $Y^1 - Y^0$. The red line is the bound for a plug-in estimator, which is just the rate for estimating the individual regression functions. The blue line is the rate for the DR-Learner given in Theorem 2, which matches the oracle under conditions given there. The purple line is the rate achieved by the lp-R-Learner when tuned as in Corollary 2, which matches the oracle when $s \geq \frac{d/4}{1+d/2\gamma}$. The black line is the improved rate achieved by the lp-R-Learner when tuned as in Corollary 3, e.g., with known covariate density.

3. What are the analogous results of Corollaries 1 for other function classes?
4. Can the error bounds in Theorem 3 be obtained with other methods?
 A natural alternative is a higher-order influence function approach [35, 38], which could avoid the $\alpha \geq \beta$ condition, perhaps at the cost of some extra complexity.

We have obtained partial answers to some of these questions, but most remain unanswered and left for future work. One of the deepest open questions is whether the rates given in Corollaries 2 and 3 are minimax optimal or not (in the fixed and random design setups, respectively).

Appendix A: Proofs

A.1. Proof of Theorem 1

Here we use the notation of Proposition 1. Letting $T_n = \widehat{m}(x) - \widetilde{m}(x) - \widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\}$ denote the numerator of the left-hand side of (3.1), we will show that

$$T_n = O_{\mathbb{P}}\left(\frac{\|\widehat{f} - f\|_{w^2}}{\|\sigma\|_{w^2}} R_n^*(x)\right)$$

which yields the result when $1/\|\sigma\|_{w^2} = O_{\mathbb{P}}(1)$. First note that for linear smoothers we have

$$T_n = \widehat{\mathbb{E}}_n\{\widehat{f}(Z) - f(Z) - \widehat{b}(X) \mid X = x\} = \sum_{i=1}^n w_i(x; X^n) \{\widehat{f}(Z_i) - f(Z_i) - \widehat{b}(X_i)\}$$

and this term has mean zero since

$$\mathbb{E}\{\widehat{f}(Z_i) - f(Z_i) - \widehat{b}(X_i) \mid D^n, X^n\} = \mathbb{E}\{\widehat{f}(Z_i) - f(Z_i) - \widehat{b}(X_i) \mid D^n, X_i\} = 0$$

by definition of \widehat{b} and iterated expectation. Therefore

$$\begin{aligned} \mathbb{E}(T_n^2 \mid D^n, X^n) &= \text{var}\left[\sum_{i=1}^n w_i(x; X^n) \{\widehat{f}(Z_i) - f(Z_i) - \widehat{b}(X_i)\} \mid D^n, X^n\right] \\ &= \sum_{i=1}^n w_i(x; X^n)^2 \text{var}\{\widehat{f}(Z_i) - f(Z_i) \mid D^n, X_i\} \leq \|\widehat{f} - f\|_{w^2}^2 \sum_{i=1}^n w_i(x; X^n)^2 \end{aligned} \tag{A.1}$$

where the second line follows since $\widehat{f}(Z_i) - f(Z_i)$ are independent given the training data, and the third since $\text{var}(\widehat{f} - f \mid D^n, X) \leq \mathbb{E}\{(\widehat{f} - f)^2 \mid D^n, X\}$ and by definition of $\|\cdot\|_{w^2}$. Further note that $R_n^*(x)^2$ equals

$$\begin{aligned} \mathbb{E}\{[\widetilde{m}(x) - m(x)]^2\} &= \mathbb{E}\left(\left[\sum_{i=1}^n w_i(x; X^n) \{f(Z_i) - m(X_i)\} \right. \right. \\ &\quad \left. \left. + \sum_{i=1}^n w_i(x; X^n) m(X_i) - m(x)\right]^2\right) \\ &= \mathbb{E}\left\{\sum_{i=1}^n w_i(x; X^n)^2 \sigma(X_i)^2\right\} + \mathbb{E}\left\{\left[\sum_{i=1}^n w_i(x; X^n) m(X_i) - m(x)\right]^2\right\} \\ &\geq \mathbb{E}\left\{\|\sigma\|_{w^2}^2 \sum_{i=1}^n w_i(x; X^n)^2\right\} \end{aligned} \tag{A.2}$$

(A.3)

where the second line follows from iterated expectation and independence of the samples, and the third by definition of $\|\cdot\|_{w^2}$ (and since the squared bias term from the previous line is non-negative). Therefore

$$\begin{aligned} \mathbb{P}\left\{\frac{\|\sigma\|_{w^2}|T_n|}{\|\widehat{f}-f\|_{w^2}R_n^*(x)}\geq t\right\} &= \mathbb{E}\left[\mathbb{P}\left\{\frac{\|\sigma\|_{w^2}|T_n|}{\|\widehat{f}-f\|_{w^2}R_n^*(x)}\geq t\mid D^n, X^n\right\}\right] \\ &\leq \left(\frac{1}{t^2R_n^*(x)^2}\right)\mathbb{E}\left\{\|\sigma\|_{w^2}^2\mathbb{E}\left(\frac{T_n^2}{\|\widehat{f}-f\|_{w^2}^2}\mid D^n, X^n\right)\right\} \\ &\leq \left(\frac{1}{t^2R_n^*(x)^2}\right)\mathbb{E}\left\{\|\sigma\|_{w^2}^2\sum_{i=1}^nw_i(x;X^n)^2\right\}\leq\frac{1}{t^2} \end{aligned}$$

where the second line follows by Markov’s inequality, the third from the bound in (A.1) and iterated expectation, and the last from the bound in (A.3). Therefore the result follows since we can always pick $t^2 = 1/\epsilon$ to ensure the above probability is no more than any ϵ .

A.2. Proof of Proposition 1

Stability and consistency together imply

$$\widehat{m}(x) - \widetilde{m}(x) = \widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} + o_{\mathbb{P}}(R_n^*(x))$$

by definition. Therefore if $\widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} = o_{\mathbb{P}}(R_n^*(x))$ the result follows.

A.3. Proof of Proposition 2

We have

$$\begin{aligned} \widehat{\mathbb{E}}_n\{\widehat{b}(X) \mid X = x\} &= \sum_{i=1}^nw_i(x;X^n)\widehat{b}_1(X_i)\widehat{b}_2(X_i) \\ &\leq \left\{\sum_{i=1}^n|w_i(x;X^n)|\|\widehat{b}_1(X_i)\|^p\right\}^{1/p}\left\{\sum_{i=1}^n|w_i(x;X^n)|\|\widehat{b}_2(X_i)\|^q\right\}^{1/q} \\ &= \left(\sum_{i=1}^n|w_i(x;X^n)|\right)\|\widehat{b}_1\|_{w,p}\|\widehat{b}_2\|_{w,q} \end{aligned}$$

where the second line follows by Holder’s inequality, and the last by definition of the norms.

A.4. Proof of Theorem 2

This follows from Proposition 1, letting $f(Z) = \varphi(Z)$ and noting that

$$\widehat{b}(x) = \left\{\frac{\pi(x)}{\widehat{\pi}(x)} - 1\right\}\{\mu_1(x) - \widehat{\mu}_1(x)\} - \left\{\frac{1 - \pi(x)}{1 - \widehat{\pi}(x)} - 1\right\}\{\mu_0(x) - \widehat{\mu}_0(x)\}$$

$$= \sum_{j=0}^1 \left\{ \frac{\widehat{\pi}(x) - \pi(x)}{j\widehat{\pi}(x) + (1-j)(1-\widehat{\pi}(x))} \right\} \{\widehat{\mu}_j(x) - \mu_j(x)\} \equiv \widehat{b}_0(x) + \widehat{b}_1(x)$$

by iterated expectation.

A.5. Proof of Corollary 1

The rate result follows since

$$\begin{aligned} \widehat{m}(x) - m(x) &= \widehat{m}(x) - \widetilde{m}(x) + \widetilde{m}(x) - m(x) \\ &= O_{\mathbb{P}}\left(n^{-\left(\frac{1}{2+d/\alpha} + \frac{1}{2+d/\beta}\right)}\right) + \widetilde{m}(x) - m(x) \\ &= O_{\mathbb{P}}\left(n^{-\left(\frac{1}{2+d/\alpha} + \frac{1}{2+d/\beta}\right)}\right) + O_{\mathbb{P}}\left(n^{\frac{-1}{2+d/\gamma}}\right) \end{aligned}$$

where the second equality follows from Proposition 2 together with Assumptions 1–3, and the third since $\widehat{\mathbb{E}}_n$ is minimax optimal and the CATE is γ -smooth. For the oracle efficiency condition, note that

$$\begin{aligned} n^{-\left(\frac{1}{2+d/\alpha} + \frac{1}{2+d/\beta}\right)} &\leq n^{-\frac{1}{2+d/\gamma}} \\ \iff \frac{1}{2+d/\alpha} + \frac{1}{2+d/\beta} &\geq \frac{1}{2+d/\gamma} \\ \iff 4 + d/\alpha + d/\beta &\geq \frac{4 + 2d/\alpha + 2d/\beta + d^2/\alpha\beta}{2+d/\gamma} \\ \iff 4 + (4 + d/\alpha + d/\beta)d/\gamma &\geq d^2/\alpha\beta \\ \iff \alpha\beta &\geq \frac{d^2}{4 + (4 + d/\alpha + d/\beta)d/\gamma} = \frac{d^2/4}{1 + (1 + d/4\alpha + d/4\beta)d/\gamma} \end{aligned}$$

which yields the result.

A.6. Proof of Theorem 3

To simplify notation, in this subsection we largely omit function arguments, for example writing $\widehat{\tau} = \widehat{\tau}(x)$, $b = b(X - x)$, $K_{hx} = K_{hx}(X)$, etc. We also define $b_0 = b(0)$ and

$$\begin{aligned} \widehat{\phi} &= \widehat{\phi}(Z) = (A - \widehat{\pi}_a)(Y - \widehat{\eta}) & \phi &= \phi(Z) = (A - \pi)(Y - \eta) \\ \widehat{\varphi}_a &= \widehat{\varphi}_a(Z) = (A - \widehat{\pi}_a)(A - \widehat{\pi}_b) & \nu &= \nu(X) = \pi(1 - \pi). \end{aligned}$$

We let X^n denote all covariates in training/test samples $(D_{1a}^n, D_{1b}^n, D_2^n)$, and we let $D_1^n = (D_{1a}^n, D_{1b}^n)$ denote the training data.

First note that by definition

$$\widehat{\tau} = b_0^T \mathbb{P}_n(bK_{hx}\widehat{\varphi}_a b^T)^{-1} \mathbb{P}_n(bK_{hx}\widehat{\phi}).$$

Thus we begin with the central decomposition

$$\begin{aligned} \widehat{\tau} - \tau &= b_0^\top \mathbb{P}_n (bK_{hx} \nu b^\top)^{-1} \mathbb{P}_n (bK_{hx} \phi) - \tau \\ &\quad + b_0^\top \mathbb{P}_n (bK_{hx} \nu b^\top)^{-1} \mathbb{P}_n \{bK_{hx} (\widehat{\phi} - \phi)\} \\ &\quad + b_0^\top \{ \mathbb{P}_n (bK_{hx} \widehat{\varphi}_a b^\top)^{-1} - \mathbb{P}_n (bK_{hx} \nu b^\top)^{-1} \} \mathbb{P}_n (bK_{hx} \widehat{\phi}) \\ &\equiv b_0^\top \widetilde{Q}_{hx}^{-1} \mathbb{P}_n (bK_{hx} \phi) - \tau \end{aligned} \tag{A.4}$$

$$+ b_0^\top \widetilde{Q}_{hx}^{-1} \mathbb{P}_n \{bK_{hx} (\widehat{\phi} - \phi)\} \tag{A.5}$$

$$+ b_0^\top (\widehat{Q}_{hx}^{-1} - \widetilde{Q}_{hx}^{-1}) \mathbb{P}_n (bK_{hx} \widehat{\phi}) \tag{A.6}$$

where we define

$$\widetilde{Q}_{hx} \equiv \mathbb{P}_n \{b(X - x)K_{hx}(X)\nu(X)b(X - x)^\top\} \tag{A.7}$$

$$\widehat{Q}_{hx} \equiv \mathbb{P}_n \{b(X - x)K_{hx}(X)\widehat{\varphi}_a(Z)b(X - x)^\top\}.$$

The general approach in this proof is to use conditional error bounds for each term in the decomposition above, from which one arrives at bounds in probability (cf. Lemma 6.1 of Chernozhukov et al. [10]).

The term on the right-hand side of (A.4) is the difference between τ and an oracle version of the lp-R-Learner that has access to the true nuisance functions and so is built from ν and ϕ ; we will show that it attains the same order as the oracle rate $n^{-\gamma/(2\gamma+d)}$. The term (A.5) captures the error from estimating $(\widehat{\pi}_b, \widehat{\eta})$ in $\widehat{\phi}$; we will show its order is the product of the biases from estimating $(\widehat{\pi}_b, \widehat{\eta})$ plus a typically smaller variance term. The term (A.6) captures the error from estimating $(\widehat{\pi}_a, \widehat{\pi}_b)$ in $\widehat{\phi}_a$; we will show it behaves similarly to (A.5), except involving the product of the biases of $(\widehat{\pi}_a, \widehat{\pi}_b)$. In regimes where the oracle rate is not achievable and the propensity score π is smoother than the regression function μ , we will show that the term (A.5) dominates.

A.6.1. Term (A.4)

The analysis of the term in (A.4) follows that of a standard local polynomial estimator of pseudo-outcome ϕ/ν on X with a special choice of kernel. This is because we can write

$$b_0^\top \mathbb{P}_n (bK_{hx} \nu b^\top)^{-1} \mathbb{P}_n (bK_{hx} \phi) = \sum_{i=1}^n w_i(x; X^n) \frac{\phi(Z_i)}{\nu(X_i)}$$

where the weights are

$$w_i(x; X^n) \equiv \frac{1}{n} b(0)^\top \widetilde{Q}_{hx}^{-1} b(X_i - x)K_{hx}(X_i)\nu(X_i). \tag{A.8}$$

Thus the oracle estimator in (A.4) is a local polynomial estimator of the regression of ϕ/ν on X using scaled kernel function $K_{hx}\nu$, which has the same support as K_{hx} and has a smaller upper bound since $\nu \leq 1/4$.

The above implies that the weights w_i satisfy analogues of Proposition 1.12 and Lemma 1.3 in Tsybakov [47], i.e., that they reproduce polynomials in X up to degree $\lceil \gamma \rceil$, and have the localizing properties given in the following lemma.

Lemma 1. Assume (i) $|K(u)| \leq K_{\max} \mathbb{1}(\|u\| \leq 1)$ and (ii) $\sup_x \|b(x)\| \leq B$, and define

$$\lambda_n \equiv \|\tilde{Q}_{hx}^{-1}\| \quad \text{and} \quad \xi_n \equiv \mathbb{P}_n(\|X - x\| \leq h)/h^d.$$

Then the weights in (A.8) satisfy:

1. $\max_i |w_i(x; X^n)| \lesssim \lambda_n/nh^d$.
2. $\sum_{i=1}^n |w_i(x; X^n)| \lesssim \lambda_n \xi_n$.
3. $w_i(x; X^n) = 0$ when $\|X_i - x\| > h$.

Proof. Property (3) follows from Assumption (i). Properties (1) and (2) also follow immediately after noting that

$$|w_i(x; X^n)| \leq \frac{K_{\max} B}{4nh^d} \|\tilde{Q}_{hx}^{-1}\| \mathbb{1}(\|X_i - x\| \leq h)$$

using the submultiplicative property of the operator norm, with Assumptions (i)–(ii) and the facts that $\|b(0)\| = 1$ and $\nu = \pi(1 - \pi) \leq 1/4$. \square

Note also that $\mathbb{E}(\phi/\nu \mid X) = \tau$ since

$$\mathbb{E}\{\phi(Z) \mid X\} = \mathbb{E}[\{A - \pi(X)\}\{Y - \eta(X)\} \mid X] = \pi(X)\{1 - \pi(X)\}\tau(X)$$

by iterated expectation. Therefore by the same logic as in Proposition 1.13 of Tsybakov [47], using Lemma 1 with the Hölder condition on τ , we have

$$\begin{aligned} \mathbb{E}\{b_0^T \mathbb{P}_n(bK_{hx}\nu b^T)^{-1} \mathbb{P}_n(bK_{hx}\phi) - \tau \mid X^n\} &= \sum_{i=1}^n w_i(x; X^n) \{\tau(X_i) - \tau(x)\} \\ &= \sum_{i=1}^n w_i(x; X^n) \left(\left[\sum_{0 < |j| \leq \lceil \gamma \rceil - 1} \frac{D^j \tau(x)}{j!} (x_i - x)^j \right. \right. \\ &\quad \left. \left. + \sum_{|k| = \lceil \gamma \rceil} \frac{\lceil \gamma \rceil}{k!} \int_0^1 (1-t)^{\lceil \gamma \rceil - 1} \{D^k \tau(x + t(x_i - x)) - D^k \tau(x)\} dt (x_i - x)^k \right] \right) \\ &\lesssim \sum_{i=1}^n |w_i(x; X^n)| \|X_i - x\|^\gamma \mathbb{1}(\|X_i - x\| \leq h) \lesssim h^\gamma \lambda_n \xi_n \end{aligned}$$

where the second line uses a multivariate Taylor approximation of τ at x evaluated at X_i , and the third line follows by the polynomial-reproducing property of w_i , the Hölder assumption on τ , and Property 3 of Lemma 1, with Property 2 of Lemma 1.

For the variance we have

$$\text{var}\{b_0^T \mathbb{P}_n(bK_{hx}\nu b^T)^{-1} \mathbb{P}_n(bK_{hx}\phi) \mid X^n\} = \sum_{i=1}^n w_i(x; X^n)^2 \text{var}(\phi \mid X = X_i)$$

$$\lesssim \max_i |w_i(x; X^n)| \sum_{i=1}^n |w_i(x; X^n)| \lesssim \frac{\lambda_n^2 \xi_n}{nh^d}$$

since the variance of ϕ is bounded, and using Properties 1–2 of Lemma 1. Therefore term (A.4) satisfies

$$\mathbb{E}[\{b_0^\top \tilde{Q}_{hx}^{-1} \mathbb{P}_n(bK_{hx}\phi) - \tau\}^2 \mid X^n] \lesssim \left(h^{2\gamma} \xi_n + \frac{1}{nh^d}\right) \lambda_n^2 \xi_n. \tag{A.9}$$

Since $\xi_n = O_{\mathbb{P}}(1)$ and $\lambda_n = O_{\mathbb{P}}(1)$ by assumption, the above rate will end up matching the classical $n^{-\gamma/(2\gamma+d)}$ rate, when balancing bias and variance by taking $h \sim n^{-1/(2\gamma+d)}$.

A.6.2. Term (A.5)

Now we bound the conditional mean and variance of term (A.5). First note we have

$$\begin{aligned} \mathbb{E}(\hat{\phi} - \phi \mid D^n, X^n) &= \mathbb{E}\{(A - \hat{\pi}_a)(Y - \hat{\eta}) - (A - \pi)(Y - \eta) \mid D^n, X^n\} \\ &= \mathbb{E}\{(A - \pi + \pi - \hat{\pi}_a)(Y - \eta + \eta - \hat{\eta}) - (A - \pi)(Y - \eta) \mid D^n, X^n\} \\ &= \{\hat{\pi}_a(X_i) - \pi(X_i)\} \{\hat{\eta}(X_i) - \eta(X_i)\} \equiv \hat{R}_2(X_i) \end{aligned} \tag{A.10}$$

by iterated expectation. Let

$$\mathcal{B}_n(x; \hat{f}) = \mathbb{E}\{\hat{f}(x) \mid X^n\} - f(x)$$

denote the pointwise conditional bias of a generic estimator \hat{f} of f at x .

Then the conditional mean of term (A.5) is

$$\begin{aligned} \mathbb{P}_n\{bK_{hx}(\hat{\phi} - \phi)\} \mid X^n &= \sum_{i=1}^n w_i(x; X^n) \nu(X_i)^{-1} \mathbb{E}(\hat{\phi}_i - \phi_i \mid X^n) \\ &= \sum_{i=1}^n w_i(x; X^n) \nu(X_i)^{-1} \mathbb{E}\{\hat{R}_2(X_i) \mid X^n\} \\ &= \sum_{i=1}^n w_i(x; X^n) \nu(X_i)^{-1} \mathcal{B}_n(X_i; \hat{\pi}_a) \mathcal{B}_n(X_i; \hat{\eta}) \\ &\lesssim \sum_{i=1}^n \frac{\lambda_n}{nh^d} |\mathcal{B}_n(X_i; \hat{\pi}_a) \mathcal{B}_n(X_i; \hat{\eta})| \mathbb{1}(\|X_i - x\| \leq h) \\ &\lesssim \sup_{\|x' - x\| \leq h} |\mathcal{B}_n(x'; \hat{\pi}_a)| |\mathcal{B}_n(x'; \hat{\eta})| \lambda_n \xi_n \lesssim k^{-\alpha/d} k^{-\beta/d} (\lambda_n \xi_n) \end{aligned} \tag{A.11}$$

where the second line follows by iterated expectation, the third since $\hat{\pi}_a \perp \hat{\eta}$, the fourth and fifth by Properties 1–3 of Lemma 1 and since $\nu \geq \epsilon(1 - \epsilon)$, and

the last by Condition (1c) via Assumption 4. For the conditional variance we have the decomposition

$$\text{var} [b_0^T \mathbb{P}_n (bK_{hx} \nu b^T)^{-1} \mathbb{P}_n \{ bK_{hx} (\hat{\phi} - \phi) \} \mid X^n] \tag{A.12}$$

$$= \text{var} \left[\sum_{i=1}^n w_i(x; X^n) \left\{ \frac{\hat{\phi}(Z_i) - \phi(Z_i)}{\nu(X_i)} \right\} \mid X^n \right]$$

$$= \mathbb{E} \left(\text{var} \left[\sum_{i=1}^n w_i(x; X^n) \left\{ \frac{\hat{\phi}(Z_i) - \phi(Z_i)}{\nu(X_i)} \right\} \mid D^n, X^n \right] \mid X^n \right) \tag{A.13}$$

$$+ \text{var} \left(\mathbb{E} \left[\sum_{i=1}^n w_i(x; X^n) \left\{ \frac{\hat{\phi}(Z_i) - \phi(Z_i)}{\nu(X_i)} \right\} \mid D^n, X^n \right] \mid X^n \right) \tag{A.14}$$

For the term in (A.13) note that

$$\text{var} \left[\sum_{i=1}^n \left\{ \frac{w_i(x; X^n)}{\nu(X_i)} \right\} (\hat{\phi}_i - \phi_i) \mid D^n, X^n \right]$$

$$= \sum_{i=1}^n \left\{ \frac{w_i(x; X^n)}{\nu(X_i)} \right\}^2 \text{var} (\hat{\phi}_i - \phi_i \mid D^n, X^n)$$

$$\lesssim \frac{1}{\epsilon(1-\epsilon)} \max_i |w_i(x; X^n)| \sum_{i=1}^n |w_i(x; X^n)| \lesssim \frac{\lambda_n^2 \xi_n}{nh^d}$$

where the second line used that $\text{var}(\hat{\phi} - \phi \mid D^n, X^n)$ and $1/\nu$ are bounded, and the last Properties 1–2 of Lemma 1.

The second term (A.14) equals

$$\text{var} \left[\sum_{i=1}^n \left\{ \frac{w_i(x; X^n)}{\nu(X_i)} \right\} \hat{R}_2(X_i) \mid X^n \right] = \sum_{i=1}^n \left\{ \frac{w_i(x; X^n)}{\nu(X_i)} \right\}^2 \text{var} \{ \hat{R}_2(X_i) \mid X^n \} \tag{A.15}$$

$$+ \sum_{i \neq j} \left\{ \frac{w_i(x; X^n) w_j(x; X^n)}{\nu(X_i) \nu(X_j)} \right\} \text{cov} \{ \hat{R}_2(X_i), \hat{R}_2(X_j) \mid X^n \} \tag{A.16}$$

For the first term (A.15) above we have

$$\text{var} \{ \hat{R}_2(X_i) \mid X^n \} = \text{var} [\{ \hat{\pi}(X_i) - \pi(X_i) \} \{ \hat{\eta}(X_i) - \eta(X_i) \} \mid X^n]$$

$$= \text{var} \{ \hat{\pi}(X_i) \mid X^n \} \text{var} \{ \hat{\eta}(X_i) \mid X^n \}$$

$$+ \text{var} \{ \hat{\pi}(X_i) \mid X^n \} [\mathbb{E} \{ \hat{\eta}(X_i) - \eta(X_i) \mid X^n \}]^2$$

$$+ \text{var} \{ \hat{\eta}(X_i) \mid X^n \} [\mathbb{E} \{ \hat{\pi}(X_i) - \pi(X_i) \mid X^n \}]^2$$

$$\lesssim \left(\frac{k}{n} \right)^2 + \frac{k}{n} (k^{-2\alpha/d} + k^{-2\beta/d}) \tag{A.17}$$

where in the second line we use the fact that $\text{var}(V_1V_2) = \text{var}(V_1)\text{var}(V_2) + \text{var}(V_1)\mathbb{E}(V_2)^2 + \text{var}(V_2)\mathbb{E}(V_1)^2$ if $V_1 \perp\!\!\!\perp V_2$, and the third the bias and variance from Conditions 1b-1c. (A slightly worse bound could have used $\text{var}(V_1V_2) \leq \mathbb{E}(V_1^2)\mathbb{E}(V_2^2)$ in the second line, which would give the expression in the third line plus a $k^{-2(\alpha+\beta)/d}$ term, which is smaller order in the undersmoothing regime.) Therefore when $\max(k^{-2\alpha/d}, k^{-2\beta/d}) = k^{-2(\alpha \wedge \beta)/d} \lesssim k/n$, the first term in (A.15) satisfies

$$\sum_{i=1}^n \left\{ \frac{w_i(x; X^n)}{\nu(X_i)} \right\}^2 \text{var}\{\widehat{R}_2(X_i) \mid X^n\} \lesssim \frac{\lambda_n^2 \xi_n}{nh^d} \left(\frac{k}{n}\right)^2$$

using Properties 1-2 of Lemma 1.

For the second term in (A.16) we have

$$\begin{aligned} & \text{cov}\{\widehat{R}_2(X_i), \widehat{R}_2(X_j) \mid X^n\} \\ &= \text{cov}\{\widehat{\pi}(X_i), \widehat{\pi}(X_j) \mid X^n\} \mathbb{E}[\{\widehat{\eta}(X_i) - \eta(X_i)\}\{\widehat{\eta}(X_j) - \eta(X_j)\} \mid X^n] \\ & \quad + \text{cov}\{\widehat{\eta}(X_i), \widehat{\eta}(X_j) \mid X^n\} \mathbb{E}[\{\widehat{\pi}(X_i) - \eta(X_i)\}\{\widehat{\pi}(X_j) - \eta(X_j)\} \mid X^n] \\ & \quad + \text{cov}\{\widehat{\pi}(X_i), \widehat{\pi}(X_j) \mid X^n\} \text{cov}\{\widehat{\eta}(X_i), \widehat{\eta}(X_j) \mid X^n\} \\ & \lesssim |\text{cov}\{\widehat{\pi}(X_i), \widehat{\pi}(X_j) \mid X^n\}| \left(k^{-2\beta/d} + \frac{k}{n}\right) \end{aligned} \tag{A.18}$$

$$\begin{aligned} & \quad + |\text{cov}\{\widehat{\eta}(X_i), \widehat{\eta}(X_j) \mid X^n\}| \left(k^{-2\alpha/d} + \frac{k}{n}\right) \\ & \quad + |\text{cov}\{\widehat{\pi}(X_i), \widehat{\pi}(X_j) \mid X^n\} \text{cov}\{\widehat{\eta}(X_i), \widehat{\eta}(X_j) \mid X^n\}| \end{aligned} \tag{A.19}$$

where the second line uses the fact that $\text{cov}(A_iB_i, A_jB_j) = \text{cov}(A_i, A_j)\mathbb{E}(B_iB_j) + \text{cov}(B_i, B_j)\mathbb{E}(A_iA_j) + \text{cov}(A_i, A_j)\text{cov}(B_i, B_j)$ whenever $(A_i, A_j) \perp\!\!\!\perp (B_i, B_j)$, and the third Cauchy-Schwarz. Now

$$\begin{aligned} & \text{cov}\{\widehat{\pi}(X_i), \widehat{\pi}(X_j) \mid X^n\} \tag{A.20} \\ &= \text{cov}\left\{ \sum_{\ell=1}^n w_{\ell\alpha}(X_i; X^n) A_{\ell}, \sum_{\ell=1}^n w_{\ell\alpha}(X_j; X^n) A_{\ell} \mid X^n \right\} \mathbb{1}(\|X_i - X_j\| \leq 2k^{-1/d}) \\ &\leq \sqrt{\text{var}\{\widehat{\pi}(X_i) \mid X^n\}} \sqrt{\text{var}\{\widehat{\pi}(X_j) \mid X^n\}} \mathbb{1}(\|X_i - X_j\| \leq 2k^{-1/d}) \\ &\lesssim \left(\frac{k}{n}\right) \mathbb{1}(\|X_i - X_j\| \leq 2k^{-1/d}) \end{aligned} \tag{A.21}$$

where the first line uses Property 3 of Lemma 1, which implies $\widehat{\pi}(X_i)$ and $\widehat{\pi}(X_j)$ are built from different observations if X_i and X_j are far enough apart, the second Cauchy-Schwarz, and the third Condition (1b). Therefore when $k^{-2(\alpha \wedge \beta)/d} \lesssim k/n$, and defining

$$\omega_{ni} \equiv \frac{k}{n-1} \sum_{j \neq i} \mathbb{1}(\|X_i - X_j\| \leq 2k^{-1/d})$$

$$\xi'_n \equiv \frac{1}{nh^d} \sum_{i=1}^n \mathbb{1}(\|X_i - x\| \leq h) \omega_{ni}$$

the second term in (A.16) satisfies

$$\begin{aligned} & \sum_{i \neq j} \left\{ \frac{w_i(x; X^n) w_j(x; X^n)}{\nu(X_i) \nu(X_j)} \right\} \text{cov} \{ \widehat{R}_2(X_i), \widehat{R}_2(X_j) \mid X^n \} \\ & \lesssim \left(\frac{k}{n} \right)^2 \sum_{i \neq j} \left\{ \frac{|w_i(x; X^n) w_j(x; X^n)|}{\nu(X_i) \nu(X_j)} \right\} \mathbb{1}(\|X_i - X_j\| \leq 2k^{-1/d}) \\ & \lesssim \left(\frac{\lambda_n}{nh^d} \right)^2 \left(\frac{k}{n} \right)^2 \sum_{i \neq j} \mathbb{1}(\|X_i - x\| \leq h) \mathbb{1}(\|X_i - X_j\| \leq 2k^{-1/d}) \\ & = \xi'_n \left(\frac{\lambda_n^2}{nh^d} \right) \left(\frac{k}{n} \right) \end{aligned}$$

where the first inequality follows from (A.19) and (A.21), the second by Properties 1 and 3 of Lemma 1, and the last by definition.

Therefore combining bounds on the four terms (A.11), (A.13), (A.15), and (A.16), we have that the term (A.5) satisfies

$$\begin{aligned} & \mathbb{E} \left([b_0^\top \widetilde{Q}_{hx}^{-1} \mathbb{P}_n \{ bK_{hx}(\widehat{\phi} - \phi) \}]^2 \mid X^n \right) \\ & \lesssim k^{-4s/d} (\lambda_n \xi_n)^2 + \frac{\lambda_n^2}{nh^d} \left(\xi_n + \left(\frac{k}{n} \right) \left(\xi'_n + \xi_n \frac{k}{n} \right) \right) \end{aligned} \quad (\text{A.22})$$

A.6.3. Term (A.6)

Note term (A.6) equals

$$\begin{aligned} & b_0^\top \{ \mathbb{P}_n (bK_{hx} \widehat{\varphi}_a b^\top)^{-1} - \mathbb{P}_n (bK_{hx} \nu b^\top)^{-1} \} \mathbb{P}_n (bK_{hx} \widehat{\phi}) \\ & = b_0^\top (\widehat{Q}_{hx}^{-1} - \widetilde{Q}_{hx}^{-1}) \mathbb{P}_n (bK_{hx} \widehat{\phi}) \\ & = \{ b_0^\top \widetilde{Q}_{hx}^{-1} (\widetilde{Q}_{hx} - \widehat{Q}_{hx}) \} \{ \widehat{Q}_{hx}^{-1} \mathbb{P}_n (bK_{hx} \widehat{\phi}) \} \end{aligned} \quad (\text{A.23})$$

For the first term in the product in (A.23) we have

$$b_0^\top \widetilde{Q}_{hx}^{-1} (\widehat{Q}_{hx} - \widetilde{Q}_{hx}) = \sum_{i=1}^n \left\{ \frac{w_i(x; X^n)}{\nu(X_i)} \right\} \{ \widehat{\varphi}_a(Z_i) - \nu(X_i) \} b(X_i - x)^\top$$

which is a finite-dimensional vector with elements

$$\sum_{i=1}^n \left\{ \frac{b_\ell(X_i - x)}{\nu(X_i)} \right\} w_i(x; X^n) \{ \widehat{\varphi}_a(Z_i) - \nu(X_i) \}$$

which we can tackle with similar logic as for term (A.5). First note

$$\mathbb{E}(\widehat{\varphi}_a - \nu \mid D^n, X^n) = \mathbb{E}\{ (A - \widehat{\pi}_a)(A - \widehat{\pi}_b) - (A - \pi)^2 \mid D^n, X^n \}$$

$$\begin{aligned}
 &= \mathbb{E}\{(A - \pi + \pi - \hat{\pi}_a)(A - \pi + \pi - \hat{\pi}_b) - (A - \pi)^2 \mid D^n, X^n\} \\
 &= \{\hat{\pi}_a(X_i) - \pi(X_i)\}\{\hat{\pi}_b(X_i) - \pi(X_i)\} \equiv \widehat{R}_{2\pi}(X_i) \tag{A.24}
 \end{aligned}$$

by iterated expectation. Defining $\mathcal{B}_n(x; f)$ as in the previous subsection, the conditional means of the elements of (A.23) (omitting subscripts ℓ) are

$$\begin{aligned}
 &\mathbb{E}\{b_0^\top \widetilde{Q}_{hx}^{-1}(\widehat{Q}_{hx} - \widetilde{Q}_{hx}) \mid X^n\}_\ell \\
 &= \sum_{i=1}^n \left\{ \frac{b(X_i - x)}{\nu(X_i)} \right\} w_i(x; X^n) \mathbb{E}\{\widehat{\varphi}_a(Z_i) - \nu(X_i) \mid X^n\} \\
 &= \sum_{i=1}^n \left\{ \frac{b(X_i - x)}{\nu(X_i)} \right\} w_i(x; X^n) \mathbb{E}\{\widehat{R}_{2\pi}(X_i) \mid X^n\} \\
 &= \sum_{i=1}^n \left\{ \frac{b(X_i - x)}{\nu(X_i)} \right\} w_i(x; X^n) \mathcal{B}_n(X_i; \hat{\pi}_a) \mathcal{B}_n(X_i; \hat{\pi}_b) \\
 &\lesssim \sum_{i=1}^n \frac{\lambda_n}{nh^d} |\mathcal{B}_n(X_i; \hat{\pi}_a) \mathcal{B}_n(X_i; \hat{\pi}_b)| \mathbb{1}(\|X_i - x\| \leq h) \\
 &\lesssim \sup_{\|x' - x\| \leq h} |\mathcal{B}_n(x'; \hat{\pi}_a)| |\mathcal{B}_n(x'; \hat{\pi}_b)| (\lambda_n \xi_n) \lesssim k^{-2\alpha/d} (\lambda_n \xi_n)
 \end{aligned}$$

where the second line follows by iterated expectation, the third since $\hat{\pi}_a \perp\!\!\!\perp \hat{\pi}_b$, the fourth and fifth by Properties 1–3 of Lemma 1 and since $\nu \geq \epsilon(1 - \epsilon)$, and the last by Condition (1c) via Assumption 4.

The analysis of the conditional variance follows exactly the same logic as for term (A.5), and is of the same order. For the second term in the product in (A.23) we have

$$\|\widehat{Q}_{hx}^{-1} \mathbb{P}_n(bK_{hx}\widehat{\phi})\| \lesssim \widehat{\lambda}_n \|\mathbb{P}_n(bK_{hx}\widehat{\phi})\|$$

and note

$$\begin{aligned}
 &\mathbb{E}\{\|\mathbb{P}_n(bK_{hx}\widehat{\phi})\|^2 \mid X^n\} \\
 &= \sum_j \mathbb{E}\{\mathbb{P}_n(b_j K_{hx}\widehat{\phi})^2 \mid X^n\} \\
 &= \sum_j [\mathbb{E}\{\mathbb{P}_n(b_j K_{hx}\widehat{\phi}) \mid X^n\}^2 + \text{var}\{\mathbb{P}_n(b_j K_{hx}\widehat{\phi}) \mid X^n\}].
 \end{aligned}$$

Therefore for b_ℓ the ℓ -th component of $b(X_i - x)$

$$\mathbb{E}\{\mathbb{P}_n(b_\ell K_{hx}\widehat{\phi}) \mid X^n\} = \mathbb{P}_n\{b_\ell K_{hx} \mathbb{E}(\widehat{\phi}_i \mid X^n)\} \lesssim \frac{1}{nh^d} \sum_{i=1}^n \mathbb{1}(\|X_i - x\| \leq h) = \xi_n$$

using boundedness of the kernel K and observations Z . For the variance we have

$$\text{var}\{\mathbb{P}_n(b_\ell K_{hx}\widehat{\phi}) \mid X^n\} = \mathbb{E}[\text{var}\{\mathbb{P}_n(b_\ell K_{hx}\widehat{\phi}) \mid D^n, X^n\} \mid X^n] \tag{A.25}$$

$$+ \text{var}[\mathbb{E}\{\mathbb{P}_n(b_\ell K_{hx}\widehat{\phi}) \mid D^n, X^n\} \mid X^n] \quad (\text{A.26})$$

For the term in (A.25) note

$$\begin{aligned} & \text{var}\{\mathbb{P}_n(b_\ell K_{hx}\widehat{\phi}) \mid D^n, X^n\} \\ &= \left(\frac{1}{nh^d}\right)^2 \sum_{i=1}^n b_\ell(X_i - x)^2 K\left(\frac{X_i - x}{h}\right)^2 \text{var}(\widehat{\phi}_i \mid D^n, X^n) \\ &\lesssim \left(\frac{1}{nh^d}\right)^2 \sum_{i=1}^n \mathbb{1}(\|X_i - x\| \leq h) = \frac{\xi_n}{nh^d} \end{aligned}$$

using the boundedness of the kernel, X , and $\widehat{\phi}$. For the term in (A.26)

$$\begin{aligned} & \text{var}[\mathbb{E}\{\mathbb{P}_n(b_\ell K_{hx}\widehat{\phi}) \mid D^n, X^n\} \mid X^n] = \text{var}[\mathbb{P}_n\{b_\ell K_{hx}\mathbb{E}(\widehat{\phi} \mid D^n, X^n)\} \mid X^n] \\ &= \text{var}\{\mathbb{P}_n\{b_\ell K_{hx}\widehat{R}_2(X_i) \mid X^n\} \lesssim \frac{1}{nh^d} \left(\xi_n + \left(\frac{k}{n}\right) \left(\xi'_n + \xi_n \frac{k}{n}\right)\right) \end{aligned}$$

where the second line follows from the definition of $\widehat{R}_2(X_i)$ and since $\mathbb{E}(\phi \mid X_i)$ is constant given X^n . The rest follows the same logic as for the term in (A.14), noting that λ_n terms do not appear since here there is only a kernel weight, and so no corresponding $\|\widetilde{Q}_{hx}^{-1}\|$ term.

Therefore the square of the term (A.6) has conditional expectation bounded above by a constant multiple of

$$\begin{aligned} & \left\{ k^{-4\alpha/d} (\lambda_n \xi_n)^2 + \frac{\lambda_n^2}{nh^d} \left(\xi_n + \left(\frac{k}{n}\right) \left(\xi'_n + \xi_n \frac{k}{n} \right) \right) \right\} \\ & \times \widehat{\lambda}_n \left\{ \xi_n^2 + \frac{1}{nh^d} \left(\xi_n + \left(\frac{k}{n}\right) \left(\xi'_n + \xi_n \frac{k}{n} \right) \right) \right\}. \quad (\text{A.27}) \end{aligned}$$

A.6.4. Combining bounds

Now we use the following lemma (or equivalently Lemma 6.1 of Chernozhukov et al. [10]) to deduce unconditional convergence from the previously derived conditional bounds.

Lemma 2. *Suppose a random variable Z_n satisfies*

$$|\mathbb{E}(Z_n \mid X^n)| \lesssim b_n \text{ and } \text{var}(Z_n \mid X^n) \lesssim s_n^2$$

for some $b_n = b(X^n)$ and $s_n^2 = s(X^n)^2$. Then $Z_n = O_{\mathbb{P}}(b_n + s_n)$.

Proof. Note for some M_ϵ depending on any $\epsilon > 0$ we have

$$\begin{aligned} \mathbb{P}\left(\frac{|Z_n|}{b_n + s_n} \geq M_\epsilon\right) &= \mathbb{E}\left\{\mathbb{P}\left(\frac{|Z_n|}{b_n + s_n} \geq M_\epsilon \mid X^n\right)\right\} \\ &\leq \mathbb{E}\left\{\frac{\mathbb{E}(Z_n^2 \mid X^n)}{(b_n + s_n)^2 M_\epsilon^2}\right\} \end{aligned}$$

$$\leq \mathbb{E} \left\{ \frac{\mathbb{E}(Z_n | X^n)^2 + \text{var}(Z_n | X^n)}{(b_n^2 + s_n^2)M_\epsilon^2} \right\} \leq \frac{C}{M_\epsilon^2}$$

where the first equality uses iterated expectation, the second Markov’s inequality, the third the fact that $(b_n + s_n)^2 \geq b_n^2 + s_n^2$, and the last the bounds on $|\mathbb{E}(Z_n | X^n)|$ and $\text{var}(Z_n | X^n)$. The result follows taking $M_\epsilon = \sqrt{C/\epsilon}$. \square

Recall our conditional bounds in (A.9), (A.22), and (A.27) involve the quantities $(\lambda_n, \widehat{\lambda}_n, \xi_n, \xi'_n)$. The quantities λ_n and $\widehat{\lambda}_n$ are $O_{\mathbb{P}}(1)$ by Assumption 3, and in the following lemmas we show that ξ_n and ξ'_n are as well, as long as the covariate density is bounded (Assumption 1).

Lemma 3. *Assume X has a density bounded above by some $C < \infty$. Then*

$$\xi_n \equiv \mathbb{P}_n(\|X - x\| \leq h)/h^d = O_{\mathbb{P}}(1).$$

Proof. First note that

$$\mathbb{E}(\xi_n) = \frac{1}{h^d} \int \mathbb{1}(\|t - x\| \leq h) p(t) dt \leq \frac{C}{h^d} \frac{\pi^{d/2} h^d}{\Gamma(1 + d/2)} \leq 5.3C$$

using the bound on the density and the volume of a d -ball of radius h . Therefore for any $\epsilon > 0$ we have

$$\mathbb{P}(\xi_n \geq 5.3C/\epsilon) \leq \epsilon$$

by Markov’s inequality, which yields the result. \square

Lemma 4. *Assume X has a density bounded above by some $C < \infty$. Define*

$$\omega_{ni} \equiv \frac{k}{n-1} \sum_{j \neq i} \mathbb{1}(\|X_i - X_j\| \leq 2/k^{1/d}).$$

Then

$$\xi'_n \equiv \frac{1}{nh^d} \sum_{i=1}^n \mathbb{1}(\|X_i - x\| \leq h)\omega_{ni} = O_{\mathbb{P}}(1)$$

Proof. For any i we have

$$\mathbb{E}(\omega_{ni} | X_i) = k \int \mathbb{1}(\|X_i - t\| \leq 2/k^{1/d}) p(t) dt \leq 2^d 5.3C$$

just as in Lemma 3. Therefore

$$\begin{aligned} \mathbb{E} \left\{ \frac{1}{nh^d} \sum_{i=1}^n \mathbb{1}(\|X_i - x\| \leq h)\omega_{ni} \right\} &= \mathbb{E} \left\{ \frac{1}{nh^d} \sum_{i=1}^n \mathbb{1}(\|X_i - x\| \leq h)\mathbb{E}(\omega_{ni} | X_i) \right\} \\ &\leq 2^d 5.3C \mathbb{E}(\xi_n) \leq 2^d (5.3C)^2 \end{aligned}$$

by iterated expectation and Lemma 3. Therefore again the result follows by Markov’s inequality. \square

Therefore, combining the bounds in (A.9), (A.22), and (A.27), together with the facts that $(\lambda_n, \hat{\lambda}_n, \xi_n, \xi'_n)$ are all $O_{\mathbb{P}}(1)$ due to Assumptions 1 and 3, we have the unconditional convergence results

$$\begin{aligned} b_0^{\top} \widehat{Q}_{hx}^{-1} \mathbb{P}_n(bK_{hx}\phi) - \tau &= O_{\mathbb{P}}\left(h^{\gamma} + \frac{1}{\sqrt{nh^d}}\right) \\ b_0^{\top} \mathbb{P}_n(bK_{hx}\nu b^{\top})^{-1} \mathbb{P}_n\{bK_{hx}(\widehat{\phi} - \phi)\} &= O_{\mathbb{P}}\left(k^{-2s/d} + \frac{1}{\sqrt{nh^d}}\left(1 + \sqrt{\frac{k}{n}} + \frac{k}{n}\right)\right) \\ b_0^{\top} \widetilde{Q}_{hx}^{-1}(\widetilde{Q}_{hx} - \widehat{Q}_{hx})\widehat{Q}_{hx}^{-1} \mathbb{P}_n(bK_{hx}\widehat{\phi}) &= O_{\mathbb{P}}\left(k^{-2\alpha/d} + \frac{1}{\sqrt{nh^d}}\left(1 + \sqrt{\frac{k}{n}} + \frac{k}{n}\right)\right). \end{aligned}$$

Note we can discard the $\sqrt{k/n}$ term since, if $k \leq n$ then the constant term 1 on the far right side dominates, whereas if $k \geq n$ then $k/n \geq \sqrt{k/n}$ and the k/n term dominates.

Appendix B: R code

Piecewise polynomial model from motivating example in Section 2.2:

```
set.seed(1234)
expit <- function(x){ exp(x)/(1+exp(x)) }
logit <- function(x){ log(x/(1-x)) }
n <- 4*2000; nsim <- 500; rateseq <- seq(0.1,0.5,by=0.05)
res2 <- NULL

for (rate in rateseq){

  res <- data.frame(matrix(nrow=nsim,ncol=4))
  colnames(res) <- c("plugin","x1","drl","oracle.drl")

  for (i in 1:nsim){

    ## simulate data
    s <- sort(rep(1:4,n/4)); x <- (runif(n,-1,1))
    ps <- 0.1 + 0.8*(x>0)
    mu0 <- (x <= -.5)*0.5*(x+2)^2 + (x/2+0.875)*(x>-1/2 & x<0) +
      (x>0 & x<.5)*(-5*(x-0.2)^2 +1.075) + (x>.5)*(x+0.125)
    mu1 <- mu0; tau <- 0
    a <- rbinom(n,1,ps)
    y <- a*mu1 + (1-a)*mu0 + rnorm(n,sd=(.2-.1*cos(2*pi*x)))

    ## estimate nuisance functions
    pihat <- expit( logit(ps) + rnorm(n,mean=1/(n/4)^rate,sd=1/(n/4)^rate))
    mu1hat <- predict(smooth.spline(x[a==1 & s==2],y[a==1 & s==2]),x)$y
    mu0hat <- predict(smooth.spline(x[a==0 & s==2],y[a==0 & s==2]),x)$y

    ## construct estimators
    plugin <- mu1hat-mu0hat
    x1 <- predict(smooth.spline(x[a==1 & s==3],(y-mu0hat)[a==1 & s==3]),x)$y
    x0 <- predict(smooth.spline(x[a==0 & s==3],(mu1hat-y)[a==0 & s==3]),x)$y
    x1 <- pihat*x0 + (1-pihat)*x1
```

```

pseudo <- ((a-pihat)/(pihat*(1-pihat)))*(y-a*mu1hat-(1-a)*mu0hat) + mu1hat-mu0hat
drl <- predict(smooth.spline(x[s==3],pseudo[s==3]),x)$y
pseudo.or <- ((a-ps)/(ps*(1-ps)))*(y-a*mu1-(1-a)*mu0) + mu1-mu0
oracle.drl <- predict(smooth.spline(x[s==3],pseudo.or[s==3]),x)$y

## save MSEs
res$plugin[i] <- (n/4)*mean((plugin-tau)[s==4]^2)
res$x1[i] <- (n/4)*mean((x1-tau)[s==4]^2)
res$drl[i] <- (n/4)*mean((drl-tau)[s==4]^2)
res$oracle.drl[i] <- (n/4)*mean((oracle.drl-tau)[s==4]^2)

}

res2 <- rbind(res2, c(rate, apply(res,2,mean)))
}

```

High-dimensional model example:

```

set.seed(1234)
library(glmnet)
expit <- function(x){ exp(x)/(1+exp(x)) }
logit <- function(x){ log(x/(1-x)) }

nsim <- 100; res <- data.frame(matrix(nrow=nsim,ncol=4))
colnames(res) <- c("plugin","x1","drl","oracle.drl")

n <- 4*2000; d <- 500; alpha <- d/10; beta <- alpha

for (i in 1:nsim){

## simulate data
s <- sort(rep(1:4,n/4)); x <- matrix(rnorm(n*d), n, d)
mu0 <- expit(as.numeric(x %*% rep(c(1, 0), c(beta,d-beta)))/sqrt(beta/1))
ps <- expit(as.numeric(x %*% rep(c(1, 0), c(alpha,d-alpha)))/sqrt(alpha/0.25))
tau <- 0; mu1 <- mu0 + tau
a <- rbinom(n,1,ps); y <- rbinom(n,1,a*mu1+(1-a)*mu0)

## estimate nuisance functions
pihat <- predict(cv.glmnet(x[s==1,],a[s==1], family="binomial"),newx=x,
  type="response", s="lambda.min")
mu0hat <- predict(cv.glmnet(x[a==0 & s==2,],y[a==0 & s==2], family="binomial"),
  newx=x, type="response", s="lambda.min")
mu1hat <- predict(cv.glmnet(x[a==1 & s==2,],y[a==1 & s==2], family="binomial"),
  newx=x, type="response", s="lambda.min")

## construct estimators
plugin <- mu1hat-mu0hat
x1 <- predict(cv.glmnet(x[a==1 & s==3,],(y-mu0hat)[a==1 & s==3]),newx=x)
x0 <- predict(cv.glmnet(x[a==0 & s==3,],(mu1hat-y)[a==0 & s==3]),newx=x)
x1 <- pihat*x0 + (1-pihat)*x1

pseudo <- ((a-pihat)/(pihat*(1-pihat)))*(y-a*mu1hat-(1-a)*mu0hat) + mu1hat-mu0hat
drl <- predict(cv.glmnet(x[s==3,],pseudo[s==3]),newx=x)
pseudo.or <- ((a-ps)/(ps*(1-ps)))*(y-a*mu1-(1-a)*mu0) + mu1-mu0
oracle.drl <- predict(cv.glmnet(x[s==3,],pseudo.or[s==3]),newx=x)

## save MSEs
res$plugin[i] <- (n/4)*mean((plugin-tau)[s==4]^2)

```

```

res$xml[i] <- (n/4)*mean((x1-tau)[s==4]^2)
res$dr1[i] <- (n/4)*mean((dr1-tau)[s==4]^2)
res$oracle.dr1[i] <- (n/4)*mean((oracle.dr1-tau)[s==4]^2)
}

```

Acknowledgments

The author thanks Sivaraman Balakrishnan, Matteo Bonvini, Aaron Fisher, Virginia Fisher, Jamie Robins, and Larry Wasserman for very helpful discussions.

Funding

This research was supported by NSF DMS Grant 1810979, NSF CAREER Award 2047444, and NIH R01 Grant LM013361-01A1.

References

- [1] C. Ai and X. Chen. Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843, 2003. [MR2015420](#)
- [2] D. W. Andrews. Asymptotics for semiparametric econometric models via stochastic equicontinuity. *Econometrica*, 62(1):43–72, 1994. [MR1258665](#)
- [3] S. Athey and G. Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016. [MR3531135](#)
- [4] A. Belloni, V. Chernozhukov, D. Chetverikov, and K. Kato. Some new asymptotic theory for least squares series: pointwise and uniform results. *Journal of Econometrics*, 186(2):345–366, 2015. [MR3343791](#)
- [5] G. Biau. Analysis of a random forests model. *The Journal of Machine Learning Research*, 13(1):1063–1095, 2012. [MR2930634](#)
- [6] A. F. Bibaut and M. J. van der Laan. Data-adaptive smoothing for optimal-rate estimation of possibly non-regular parameters. *arXiv preprint arXiv:1706.07408*, 2017.
- [7] P. J. Bickel and Y. Ritov. Estimating integrated squared density derivatives: sharp best order of convergence estimates. *Sankhyā*, 381–393, 1988. [MR1065550](#)
- [8] L. Birgé and P. Massart. Estimation of integral functionals of a density. *The Annals of Statistics*, 23(1):11–29, 1995. [MR1331653](#)
- [9] V. Chernozhukov, M. Goldman, V. Semenova, and M. Taddy. Orthogonal machine learning for demand estimation: High dimensional causal inference in dynamic panels. *arXiv preprint arXiv:1712.09988*, 2017.
- [10] V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. M. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. [MR3769544](#)

- [11] A. D'Amour, P. Ding, A. Feller, L. Lei, and J. Sekhon. Overlap in observational studies with high-dimensional covariates. *arXiv preprint arXiv:1711.02582*, 2017. [MR4215042](#)
- [12] I. Díaz, O. Savenkov, and K. Ballman. Targeted learning ensembles for optimal individualized treatment rules with time-to-event outcomes. *Biometrika*, 105(3):723–738, 2018. [MR3842895](#)
- [13] J. Fan and I. Gijbels. Censored regression: local linear approximations and their applications. *Journal of the American Statistical Association*, 89(426):560–570, 1994. [MR1294083](#)
- [14] Q. Fan, Y.-C. Hsu, R. P. Lieli, and Y. Zhang. Estimation of conditional average treatment effects with high-dimensional data. *arXiv preprint arXiv:1908.02399*, 2019. [MR4356575](#)
- [15] D. J. Foster and V. Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019. [MR4630373](#)
- [16] J. C. Foster, J. M. Taylor, and S. J. Ruberg. Subgroup identification from randomized clinical trial data. *Statistics in Medicine*, 30(24):2867–2880, 2011. [MR2844689](#)
- [17] Z. Gao and Y. Han. Minimax optimal nonparametric estimation of heterogeneous treatment effects. *arXiv preprint arXiv:2002.06471*, 2020.
- [18] L. Györfi, M. Kohler, A. Krzykacz, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, 2002. [MR1920390](#)
- [19] J. Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 66(2):315–331, 1998. [MR1612242](#)
- [20] P. R. Hahn, J. S. Murray, and C. M. Carvalho. Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects. *Bayesian Analysis*, 2020. [MR4154846](#)
- [21] K. Imai and M. Ratkovic. Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443–470, 2013. [MR3086426](#)
- [22] E. H. Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *arXiv preprint arXiv:2203.06469*, 2022.
- [23] E. H. Kennedy, Z. Ma, M. D. McHugh, and D. S. Small. Nonparametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B*, 79(4):1229–1245, 2017. [MR3689316](#)
- [24] E. H. Kennedy, S. Balakrishnan, and M. G'Sell. Sharp instruments for classifying compliers and generalizing causal effects. *The Annals of Statistics*, 48(4):2008–2030, 2020. [MR4134784](#)
- [25] S. R. Künnzel, J. S. Sekhon, P. J. Bickel, and B. Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019. [MR4051203](#)
- [26] S. Lee, R. Okui, and Y.-J. Whang. Doubly robust uniform confidence band for the conditional average treatment effect function. *Journal of Applied Econometrics*, 32(7):1207–1225, 2017. [MR3734484](#)
- [27] A. R. Luedtke and M. J. van der Laan. Super-learning of an opti-

- mal dynamic treatment rule. *The International Journal of Biostatistics*, 12(1):305–332, 2016. [MR3505699](#)
- [28] E. Masry. Multivariate regression estimation local polynomial fitting for time series. *Stochastic Processes and their Applications*, 65(1):81–101, 1996. [MR1422881](#)
- [29] W. K. Newey and J. M. Robins. Cross-fitting and fast remainder rates for semiparametric estimation. *arXiv preprint [arXiv:1801.09138](#)*, 2018.
- [30] X. Nie and S. Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 2021. [MR4259133](#)
- [31] A. Rambachan, A. Coston, and E. H. Kennedy. Counterfactual risk assessments under unmeasured confounding. *arXiv preprint [arXiv:2212.09844](#)*, 2022.
- [32] J. M. Robins. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics – Theory and Methods*, 23(8):2379–2412, 1994. [MR1293185](#)
- [33] J. M. Robins and A. Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995. [MR1325119](#)
- [34] J. M. Robins, S. D. Mark, and W. K. Newey. Estimating exposure effects by modelling the expectation of exposure conditional on confounders. *Biometrics*, 48(2):479–495, 1992. [MR1173493](#)
- [35] J. M. Robins, L. Li, E. J. Tchetgen Tchetgen, and A. W. van der Vaart. Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and Statistics: Essays in Honor of David A. Freedman*, pages 335–421, 2008. [MR2459958](#)
- [36] J. M. Robins, L. Li, E. J. Tchetgen Tchetgen, and A. W. van der Vaart. Quadratic semiparametric von mises calculus. *Metrika*, 69(2-3):227–247, 2009. [MR2481922](#)
- [37] J. M. Robins, E. J. Tchetgen Tchetgen, L. Li, and A. W. van der Vaart. Semiparametric minimax rates. *Electronic Journal of Statistics*, 3:1305–1321, 2009. [MR2566189](#)
- [38] J. M. Robins, L. Li, R. Mukherjee, E. J. Tchetgen Tchetgen, and A. W. van der Vaart. Minimax estimation of a functional on a structured high dimensional model. *The Annals of Statistics*, 45(5):1951–1987, 2017. [MR3718158](#)
- [39] P. M. Robinson. Root-n-consistent semiparametric regression. *Econometrica*, 56(4):931–954, 1988. [MR0951762](#)
- [40] D. B. Rubin and M. J. van der Laan. A general imputation methodology for nonparametric regression with censored data. *UC Berkeley Division of Biostatistics Working Paper Series*, 194, 2005.
- [41] D. B. Rubin and M. J. van der Laan. Doubly robust censoring unbiased transformations. *UC Berkeley Division of Biostatistics Working Paper Series*, 208, 2006. [MR2306842](#)
- [42] E. Scornet. Random forests and kernel methods. *IEEE Transactions on Information Theory*, 62(3):1485–1500, 2016. [MR3472261](#)
- [43] V. Semenova and V. Chernozhukov. Estimation and inference about conditional average treatment effect and other structural functions. *arXiv*, 2017.

- [44] U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, Volume 70, pages 3076–3085. JMLR.org, 2017.
- [45] C. J. Stone. Consistent nonparametric regression. *The Annals of Statistics*, 595–620, 1977. [MR0443204](#)
- [46] A. A. Tsiatis. *Semiparametric Theory and Missing Data*. New York: Springer, 2006. [MR2233926](#)
- [47] A. B. Tsybakov. *Introduction to Nonparametric Estimation*. New York: Springer, 2009. [MR2724359](#)
- [48] M. J. van der Laan. Statistical inference for variable importance. *The International Journal of Biostatistics*, 188, 2005. [MR2275897](#)
- [49] M. J. van der Laan. Targeted learning of an optimal dynamic treatment, and statistical inference for its mean outcome. *UC Berkeley Division of Biostatistics Working Paper Series*, 317:1–90, 2013.
- [50] M. J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. *UC Berkeley Division of Biostatistics Working Paper Series*, Paper 130, 2003. [MR0902239](#)
- [51] M. J. van der Laan and J. M. Robins. *Unified Methods for Censored Longitudinal Data and Causality*. New York: Springer, 2003. [MR1958123](#)
- [52] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge: Cambridge University Press, 2000. [MR1652247](#)
- [53] S. Vansteelandt and M. M. Joffe. Structural nested models and g-estimation: the partially realized promise. *Statistical Science*, 29(4):707–731, 2014. [MR3300367](#)
- [54] I. Verdinelli and L. Wasserman. Forest guided smoothing. *arXiv preprint [arXiv:2103.05092](#)*, 2021.
- [55] S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. [MR3862353](#)
- [56] W. Zheng and M. J. van der Laan. Asymptotic theory for cross-validated targeted maximum likelihood estimation. *UC Berkeley Division of Biostatistics Working Paper Series*, Paper 273:1–58, 2010.
- [57] M. Zimmert and M. Lechner. Nonparametric estimation of causal heterogeneity under high-dimensional confounding. *arXiv preprint [arXiv:1908.08779](#)*, 2019.