
Interval Estimation of Individual-Level Causal Effects Under Unobserved Confounding

Nathan Kallus*

Xiaojie Mao*

Angela Zhou*

Cornell University and Cornell Tech

Abstract

We study the problem of learning conditional average treatment effects (CATE) from observational data with unobserved confounders. The CATE function maps baseline covariates to individual causal effect predictions and is key for personalized assessments. Recent work has focused on how to learn CATE under unconfoundedness, *i.e.*, when there are no unobserved confounders. Since CATE may not be identified when unconfoundedness is violated, we develop a functional interval estimator that predicts bounds on the individual causal effects under realistic violations of unconfoundedness. Our estimator takes the form of a weighted kernel estimator with weights that vary adversarially. We prove that our estimator is sharp in that it converges exactly to the tightest bounds possible on CATE when there may be unobserved confounders. Further, we prove that personalized decision rules derived from our estimator achieve optimal minimax regret asymptotically. We assess our approach in a simulation study as well as demonstrate its application by comparing conclusions from a real observational study and clinical trial.

1 Introduction

Learning individual-level (conditional-average) causal effects from observational data is a key question for determining personalized treatments in medicine or in assessing policy impacts in the social sciences. Many recent advances have been made for the important

question of estimating conditional average treatment effects (CATE), which is a function mapping baseline covariates to individual causal effect predictions [1, 3, 10, 13, 19, 26, 34, 38]. However, all of these approaches need to assume *unconfoundedness*, or that the potential outcomes are conditionally independent of treatments, given observed covariates. That is, that all possible confounders have been observed and are controlled for.

While unconfoundedness may hold *by design* in ideal settings like randomized controlled trials, the assumption is almost always *invalid to some degree* in any real observational study and, to make things worse, the assumption is inherently *unverifiable*. For example, passively collected healthcare databases often lack part of the critical clinical information that may drive both doctors' and patients' treatment choices, *e.g.*, subjective assessments of condition severity or personal lifestyle factors. The expansion and linkage of observational and administrative datasets does afford greater opportunities to observe important factors that influence selection into treatment, but some hidden factors will always remain and there is no way to prove otherwise. When unconfoundedness does not hold, one might find causal effects in observational data where there is actually no real effect or vice versa, which in turn may lead to real decisions that dangerously rely on false conclusions and may introduce unnecessary harm or risk [17].

Therefore, sensitivity analysis of causal estimates to realistic violations of unconfoundedness is crucial for both credible interpretation of any findings and reliable decision making. Traditional sensitivity analysis and modern extensions focus on bounding feasible values of *average treatment effects* (ATE) or the corresponding *p*-values for the hypothesis of zero effect, assuming some violation of unconfoundedness [7, 31, 40].

However, ATE is of limited use for individual-level assessments and personalization. For such applications, it is crucial to study the *heterogeneity* of effects as covariates vary by estimating the CATE. Given an in-

*Authors in alphabetical order.

dividual with certain baseline covariates, the sign of the CATE for their specific values determines the best course of action for the individual. Furthermore, learning CATE also allows one to generalize causal conclusions drawn from one population to another population [11].

In this paper, we develop new methodology and theory for learning bounds on the CATE function from observational data that may be subject to some confounding. We propose a functional interval estimator for CATE that is derived from a weighted kernel regression, where the weights vary adversarially per a standard sensitivity model that specifies how big the potential impact of hidden confounders might be on selection. The proposed estimators admit efficient computation by a sorting and line search procedure. We show that our estimator is *sharp* in that it converges pointwise to the tightest possible set of identifiable CATE functions – the set of CATE functions that are consistent with both the population-level observational-data generating process and the assumed sensitivity model. That is, our interval function is asymptotically neither too wide (too conservative) nor too narrow (too optimistic). We show that the minimax-regret of the derived personalized decision rules converges to the best possible under the assumed sensitivity model. We assess the success of our approach in considering individual-level effects in a simulation study, and demonstrate its application on real data, a likely-confounded observational study and companion clinical trial for hormone replacement therapy.

2 Related work

Learning CATE. Studying heterogeneous treatment effects by learning a functional form for CATE under the assumption of unconfoundedness is a common approach [1, 3, 10, 13, 19, 26, 34, 38]. Under unconfoundedness, CATE is given by the difference of two identifiable regressions and the above work study how to appropriately tailor supervised learning algorithms specifically to such a task. In particular, [1] consider estimating CATE under unconfoundedness using kernel regression on a transformation of the outcome given by inverse propensity weighting (IPW). Our bounds arise from adversarially-weighted kernel regression estimators, but in order to ensure sharpness, the estimators we use reweight the sample rather than the outcome.

Sensitivity analysis and partial identification. Sensitivity analysis in causal inference considers how the potential presence of unobserved confounders might affect a conclusion made under the assumption

of unconfoundedness. It originated in a thought experiment on the effects of smoking in lung cancer where it is argued that unobservable confounding effects must be unrealistically large to refute observational evidence [6]. In our approach, we use the marginal sensitivity model (MSM) introduced by [37], which bounds the potential impact of unobserved confounding on selection into treatment. Specifically, it bounds the ratio between the propensity for treatment when accounting only for observables and when accounting also for unobservables. This model is closely related to the Rosenbaum sensitivity model [31], which bounds the ratio between the propensity for treatment between any two realization of unobservables. While the Rosenbaum sensitivity model is traditionally used in conjunction with matching, [39] recently proposed an empirical risk minimization approach to estimate the CATE sensitivity bounds based on sieve estimation methods and further provide ATE bounds. See [40] for more on the relationship between the Rosenbaum sensitivity model and the MSM considered in this paper. [37, 40] consider the sensitivity of ATE estimates under the MSM but not sharpness. [2] consider a related problem of bounding a population average under observations with unknown but bounded sampling probabilities and prove sharpness assuming discrete outcomes. Instead of relying on a sensitivity model, [23] consider sharp partial identification of ATE under no or weak assumptions such as monotone response and [22] consider corresponding minimax-regret policy choice with discrete covariates. [24] consider sup-norm bounds on propensity differences and show sharp partial identification of bounds for CATE and ATE by integrating partially identified bounds on the conditional quantile treatment effect.

Personalized decision making. Optimal personalized decision rules are given by thresholding CATE if known and hence a natural approach to learning such policies is to threshold CATE estimates [30]. This of course breaks down if CATE is not estimable. Our paper derives the appropriate extension to partially identified CATE and shows that decision rules derived from our estimates in fact achieve optimal minimax regret. Under unconfoundedness, recent work has also studied directly learning a *structured* decision policy from observational data since best-in-class CATE estimates (*e.g.*, best linear prediction) need not necessarily lead to best-in-class policies (*e.g.*, best linear policy) [4, 8, 14–16, 36, 41]. Recently, [17] studied the problem of finding such structured policies that are also robust to possible confounding under a similar sensitivity model. However, this approach produces only a policy and not a CATE estimate, which itself is an important object for decision support as one would like to interpret the policy relative to predicted effects and

understand the *magnitude* of effect and *uncertainties* in its estimation. Causal effect estimates are of course also important for influencing other conclusions, directing further the study of causal mechanisms, and measuring conclusions against domain knowledge.

3 Problem Set-up

We assume that the observational data consists of triples of random variables $\{(X_i, T_i, Y_i) : i = 1, \dots, n\}$, comprising of covariates $X_i \in \mathcal{X} \subseteq \mathbb{R}^d$, assigned treatments $T_i \in \{0, 1\}$, and real-valued outcomes $Y_i \in \mathcal{Y} \subseteq \mathbb{R}$. Using the Neyman-Rubin potential outcome framework, we let $Y_i(0), Y_i(1) \in \mathcal{Y}$ denote the potential outcomes of each treatment. We let the observed outcome be the potential outcome of the assigned treatment, $Y_i = Y_i(T_i)$, encapsulating non-interference and consistency assumptions, also known as SUTVA [33]. Moreover, $(X_i, T_i, Y_i(0), Y_i(1))$ are i.i.d draws from a population $(X, T, Y(0), Y(1))$.

We are interested in the CATE function:

$$\tau(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x]$$

If X contained all confounders, then we could identify CATE by *controlling* for it in each treatment group:

$$\tilde{\tau}(x) = \mathbb{E}[Y \mid T = 1, X = x] - \mathbb{E}[Y \mid T = 0, X = x]. \quad (1)$$

If unconfoundedness held in that potential outcomes are independent of assigned treatment given X , then it's immediate that $\tilde{\tau}$ would equal τ . However, in practice there will almost always exist *unobserved confounders* not included in X , *i.e.*, unconfoundedness with respect to X is *violated*. That is, in general, we may have that

$$(Y(0), Y(1)) \not\perp T \mid X.$$

In such general settings, $\tau(x) \neq \tilde{\tau}(x)$ and indeed $\tau(x)$ may *not* be estimated from the observed data even with an infinite sample size [31].

For $t \in \{0, 1\}$, let $e_t(x) = \mathbb{P}(T = t \mid X = x)$ be the *nominal* propensity for treatment given only the observed variables and $e_t(x, y) = \mathbb{P}(T = t \mid X = x, Y(t) = y)$ be the *complete* propensity accounting for all confounders. In this paper, we use the following sensitivity model to quantify the extent of violation of the unconfoundedness with respect to the observed covariates X . The model measures the degree of confounding in terms of the odds ratios of the nominal and complete propensities [37, 40].

Definition 1 (Marginal Sensitivity Model). *There exists $\Gamma \geq 1$ such that, for any $t \in \{0, 1\}, x \in \mathcal{X}, y \in \mathcal{Y}$,*

$$\frac{1}{\Gamma} \leq \frac{(1 - e_t(x))e_t(x, y)}{e_t(x)(1 - e_t(x, y))} \leq \Gamma. \quad (2)$$

Taking logs, eq. (2) can be seen as bounding the absolute difference between the logits the nominal propensity and the complete propensity by $\log \Gamma$. When unconfoundedness with respect to X holds, we have that $e_t(x) = e_t(x, y)$ and (2) holds with $\Gamma = 1$. When $\Gamma = 2$, for example, then the true odds ratio for an individual to be treated may be as much as double or as little as half of what it is actually observed to be given only X . As Γ increasingly deviates from 1, we allow for greater unobserved confounding.

4 An Interval Estimator for CATE

4.1 Population estimands

We start by characterizing the population estimands we are after, the population-level upper and lower bounds on CATE. As discussed above, without unconfoundedness, there is no single CATE function that can be point identified by the data. Under the MSM with a given Γ , we can conceive of the set of identified CATE functions as consisting of all the functions that are consistent with both the population of observational data (X, T, Y) and the MSM. All such functions are observationally equivalent in that they cannot be distinguished from one another on the basis of observational data alone. This defines a particular interval function that maps covariates to the lower and upper bounds of this set and this is the function we wish to estimate.

For $t \in \{0, 1\}$ and $x \in \mathcal{X}$, let $f_t(y \mid x)$ denote the density of the distribution $\mathbb{P}(T = t, Y(t) \leq y \mid X = x) = \mathbb{P}(T = t, Y \leq y \mid X = x)$. Note that these distributions are identifiable based on the observed data as they only involve observable quantities. Further, define $\mu_t(x) = \mathbb{E}[Y(t) \mid X = x]$, which is *not* identifiable from data, such that $\tau(x) = \mu_1(x) - \mu_0(x)$, and note that

$$\mu_t(x) = \mu_t(w_t; x) = \frac{\int y w_t(y \mid x) f_t(y \mid x) dy}{\int w_t(y \mid x) f_t(y \mid x) dy} \quad (3)$$

where $w_t(y \mid x) = 1/e_t(x, y)$, which too is unidentifiable. Eq. (3) is useful as it decomposes $\mu_t(x)$ cleanly into its identifiable ($f_t(y \mid x)$) and unidentifiable ($w_t(y \mid x)$) components. Based on the MSM (2), we can determine the uncertainty set that includes all possible values of $w_t(y \mid x) = 1/e_t(x, y)$ that agree with the model, *i.e.*, violate unconfoundedness by no more than Γ :

$$\mathcal{W}_t(x; \Gamma) = \{w_t(\cdot \mid x) : w_t(y \mid x) \in [\alpha_t(x; \Gamma), \beta_t(x; \Gamma)] \forall y\} \quad (4)$$

where $\alpha_t(x; \Gamma) = 1/(\Gamma e_t(x)) + 1 - 1/\Gamma$, $\beta_t(x; \Gamma) = \Gamma/e_t(x) + 1 - \Gamma$.

These are exactly the $w_t(\cdot | x)$ functions that agree with both the known nominal propensities $e_t(x)$ and the MSM in eq. (2). Eq. (4) is derived directly from eq. (2) by simple algebraic manipulation. We define the population CATE bounds under the MSM correspondingly.

Definition 2 (CATE Identified Set Under MSM). *The population bounds under the MSM with parameter Γ for the expected potential outcomes are*

$$\bar{\mu}_t(x; \Gamma) = \sup_{w_t \in \mathcal{W}_t(x; \Gamma)} \mu_t(w_t; x), \quad (5)$$

$$\underline{\mu}_t(x; \Gamma) = \inf_{w_t \in \mathcal{W}_t(x; \Gamma)} \mu_t(w_t; x), \quad (6)$$

and the population bounds for CATE are

$$\bar{\tau}(x; \Gamma) = \bar{\mu}_1(x; \Gamma) - \underline{\mu}_0(x; \Gamma), \quad (7)$$

$$\underline{\tau}(x; \Gamma) = \underline{\mu}_1(x; \Gamma) - \bar{\mu}_0(x; \Gamma). \quad (8)$$

Therefore, the target function we are interested in learning is the map from x to identifiable CATE intervals:

$$\mathcal{T}(x; \Gamma) = [\underline{\tau}(x; \Gamma), \bar{\tau}(x; \Gamma)].$$

4.2 The functional interval estimator

We next develop our functional interval estimator for $\mathcal{T}(x; \Gamma)$. Toward this end, we consider the following kernel-regression-based estimator for $\mu_t(x)$ based on the *unknown* weights $\mathbf{W}_t^* = (W_{ti}^*)_{i=1}^n$ according to the complete propensity score, $W_{ti}^* = 1/e_t(X_i, Y_i(t))$:

$$\hat{\mu}_t(\mathbf{W}_t^*; x) = \frac{\sum_{i=1}^n \mathbb{I}(T_i = t) \mathbf{K}\left(\frac{X_i - x}{h}\right) W_{ti}^* Y_i}{\sum_{i=1}^n \mathbb{I}(T_i = t) \mathbf{K}\left(\frac{X_i - x}{h}\right) W_{ti}^*}, \quad (9)$$

where $\mathbf{K}\left(\frac{X_i - x}{h}\right) = K\left(\frac{X_{i,1} - x_1}{h}\right) \times \dots \times K\left(\frac{X_{i,d} - x_d}{h}\right)$, $K(\cdot)$ is a univariate kernel function, and $h > 0$ is a bandwidth. In particular, all we require of K is that it is bounded and $\int uK(u)du = 0$, $\int u^2K(u)du < \infty$ (see Thm. 1). For example, we can use the Gaussian kernel $K(u) = \exp(-u^2/2)$ or uniform kernel $K(u) = \mathbb{I}(|u| \leq \frac{1}{2})$. If we knew the true weights \mathbf{W}_t^* , then basic results on non-parametric regression and inverse-probability weighting would immediately give that $\hat{\mu}_t(\mathbf{W}_t^*; x) \rightarrow \mu_t(x)$ as $n \rightarrow \infty$ if $h \rightarrow 0$ and $nh^d \rightarrow \infty$ [27], i.e., the estimator, eq. (9), would be consistent when the complete propensity scores are known.

However, the estimator in eq. (9) is an infeasible one in practice because $1/W_{ti}^* = e_t(X_i, Y_i(t))$ is unknown and cannot be estimated from any amount of observed data. Instead, we bracket the range of feasible weights and consider how large or small eq. (9) might be. For $t \in \{0, 1\}$ and $\Gamma \geq 1$, we define

$$\widehat{\mathcal{W}}_t = \{\mathbf{W}_t : \alpha_t(X_i; \Gamma) \leq W_{ti} \leq \beta_t(X_i; \Gamma) \forall i\} \subseteq \mathbb{R}^n,$$

where $\alpha_t(\cdot)$ and $\beta_t(\cdot)$ are defined in (4). Our interval CATE estimator is

$$\widehat{\mathcal{T}}(x; \Gamma) = [\hat{\underline{\tau}}(x), \hat{\bar{\tau}}(x)], \quad \text{where} \quad (10)$$

$$\hat{\bar{\tau}}(x) = \hat{\bar{\mu}}_1(x) - \hat{\underline{\mu}}_0(x), \quad \hat{\underline{\tau}}(x) = \hat{\underline{\mu}}_1(x) - \hat{\bar{\mu}}_0(x), \quad (11)$$

$$\hat{\bar{\mu}}_t(x; \Gamma) = \sup_{\mathbf{W}_t \in \widehat{\mathcal{W}}_t} \hat{\mu}_t(\mathbf{W}; x), \quad (12)$$

$$\hat{\underline{\mu}}_t(x; \Gamma) = \inf_{\mathbf{W}_t \in \widehat{\mathcal{W}}_t} \hat{\mu}_t(\mathbf{W}; x). \quad (13)$$

Note that $\alpha_t(x), \beta_t(x)$ depend on $e_t(x)$. Since we mainly focus on dealing with unobserved confounding, we assume that we know the nominal propensity scores $e_t(x)$ for simplicity as it is in fact identifiable.

4.3 Computing the interval estimator

Our interval estimator is defined as an optimization problem over n weight variables. We can simplify this problem by characterizing its solution. Using optimization duality, Lemma 2 in appendix shows that, in the solution, each weight variable realizes its bounds (upper or lower) and that weights are monotone when sorted in increasing Y_i value. This means that one need only search for the inflection point. As summarized in the following proposition, this means that the solution is given by a simple discrete line search to optimize a unimodal function, after sorting.

Proposition 1. *Suppose that we reorder the data so that $Y_1 \leq Y_2 \leq \dots \leq Y_n$. Define the following terms for $k \in \{1, \dots, n\}$, $x \in \mathcal{X}$, and $\Gamma \geq 1$:*

$$\bar{\lambda}(k; x, \Gamma) = \frac{\sum_{i \leq k} \tilde{\alpha}_i^K(t, x; \Gamma) Y_i + \sum_{i \geq k+1} \tilde{\beta}_i^K(t, x; \Gamma) Y_i}{\sum_{i \leq k} \tilde{\alpha}_i^K(t, x; \Gamma) + \sum_{i \geq k+1} \tilde{\beta}_i^K(t, x; \Gamma)},$$

$$\underline{\lambda}(k; x, \Gamma) = \frac{\sum_{i \leq k} \tilde{\beta}_i^K(t, x; \Gamma) Y_i + \sum_{i \geq k+1} \tilde{\alpha}_i^K(t, x; \Gamma) Y_i}{\sum_{i \leq k} \tilde{\beta}_i^K(t, x; \Gamma) + \sum_{i \geq k+1} \tilde{\alpha}_i^K(t, x; \Gamma)},$$

where

$$\tilde{\alpha}_i^K(t, x; \Gamma) = \mathbb{I}[T_i = t] \alpha_t(X_i; \Gamma) \mathbf{K}\left(\frac{X_i - x}{h}\right),$$

$$\tilde{\beta}_i^K(t, x; \Gamma) = \mathbb{I}[T_i = t] \beta_t(X_i; \Gamma) \mathbf{K}\left(\frac{X_i - x}{h}\right).$$

Then we have that

$$\hat{\bar{\mu}}_t(x) = \bar{\lambda}(k^H(x, \Gamma); x, \Gamma), \quad \hat{\underline{\mu}}_t(x) = \underline{\lambda}(k^L(x, \Gamma); x, \Gamma),$$

where

$$k^H(x, \Gamma) = \inf\{k = 1, \dots, n : \bar{\lambda}(k; x, \Gamma) \geq \bar{\lambda}(k+1; x, \Gamma)\},$$

$$k^L(x, \Gamma) = \inf\{k = 1, \dots, n : \underline{\lambda}(k; x, \Gamma) \leq \underline{\lambda}(k+1; x, \Gamma)\}.$$

4.4 Sharpness guarantees

We next establish that our interval estimator is *sharp*, i.e., it converges to the identifiable set of CATE values.

That is to say, as a robust estimator that accounts for possible confounding, our interval is neither too wide nor too narrow – asymptotically, it matches exactly what can be hoped to be learned from any amount of observational data. The result is based on a new *uniform* convergence result that we prove for the weight-parametrized kernel regression estimator, $\hat{\mu}_t(W; x)$, to the weight-parametrized estimand $\mu_t(w_t; x)$. Although the uniform convergence may not hold in general, it holds when restricting to monotone weights, which is where we leverage our characterization of the optimal solution to eqs. (13) and (12) as well as a similar result characterizing the population version in eqs. (5) and (6) using semi-infinite optimization duality [35].

Theorem 1. *Suppose that*

- i. K is bounded, $\int K(u) < \infty$, $\int uK(u)du = 0$, and $\int u^2K(u)du < \infty$,
- ii. $n \rightarrow \infty$, $h \rightarrow 0$ and $nh^{2d} \rightarrow \infty$,
- iii. Y is a bounded random variable,
- iv. $e_t(x)$ and $f_t(y | x)$ are twice continuously differentiable with respect to x for any fixed $y \in \mathcal{Y}$ with bounded first and second derivatives,¹
- v. $e_t(x, y)$ is bounded away from 0 and 1 uniformly over $x \in \mathcal{X}$, $y \in \mathcal{Y}$, $t \in \{0, 1\}$.

Then, for $t \in \{0, 1\}$,

$$\begin{aligned} \hat{\bar{\mu}}_t(x) &\xrightarrow{P} \bar{\mu}_t(x), & \hat{\underline{\mu}}_t(x) &\xrightarrow{P} \underline{\mu}_t(x), \\ \hat{\bar{\tau}}(x) &\xrightarrow{P} \bar{\tau}(x), & \hat{\underline{\tau}}(x) &\xrightarrow{P} \underline{\tau}(x). \end{aligned}$$

In words, Theorem 1 states that under fairly general assumptions, if the bandwidth h is appropriately chosen, our bounds for both conditional average outcomes and CATE are pointwise consistent and hence sharp.

4.5 Personalized Decisions from Interval Estimates and Minimax Regret Guarantees

We next consider how our interval CATE estimate can be used for personalized treatment decisions and prove that the resulting decisions rules are asymptotically minimax optimal. Let us assume that the outcomes Y_i correspond to *losses* so that *lower outcomes are better*. Then, if the CATE were known, given an individual with covariates x , clearly the optimal treatment decision is $t = 1$ if $\tau(x) < 0$ and $t = 0$ if $\tau(x) > 0$ (and either if $\tau(x) = 0$). In other words, $\pi(x) = \mathbb{I}(\tau(x) < 0)$

minimizes the risk

$$\begin{aligned} V(\pi; \tau) &= \mathbb{E}[\pi(X)Y(1) + (1 - \pi(X))Y(0)] \\ &= \mathbb{E}[Y(1)] + \mathbb{E}[\pi(X)\tau(X)] \end{aligned}$$

over $\pi : \mathcal{X} \rightarrow \{0, 1\}$. If $\tau(x)$ can be point-estimated, an obvious approach to making personalized decisions is to threshold an estimator of it. If the estimator is consistent, this will lead to asymptotically optimal risk.

This, however, breaks down when CATE is unidentifiable and we only have an interval estimate. In particular, there is no single identifiable value of $V(\pi; \tau)$. Instead, we focus on the worst case regret relative to a default $\pi_0(x)$ given by the MSM:

$$\bar{R}_{\pi_0}(\pi; \Gamma) = \sup_{\tau(x) \in \mathcal{T}(x; \Gamma)} \forall x \in \mathcal{X} (V(\pi; \tau) - V(\pi_0; \tau))$$

The default represents the decision that would have been taken in the absence of any of our observational data. For example, in the medical domain, if there is not enough clinical trial evidence to support treatment approval, then the default may be to not treat, $\pi_0(x) = 0$. At the population level, the uniformly best possible policy we can hope for is the minimax regret policy:

$$\pi^*(\cdot; \Gamma) \in \operatorname{argmin}_{\pi: \mathcal{X} \rightarrow \{0, 1\}} \bar{R}_{\pi_0}(\pi; \Gamma) \quad (14)$$

Proposition 2. *The following is a solution to eq. (14):*

$$\pi^*(x; \Gamma) = \mathbb{I}(\bar{\tau}(x; \Gamma) \leq 0) + \pi_0(x)\mathbb{I}(\underline{\tau}(x; \Gamma) < 0 < \bar{\tau}(x; \Gamma))$$

This minimax-optimal policy always treats when $\bar{\tau}(x; \Gamma) \leq 0$ and never treats when $\underline{\tau}(x; \Gamma) \geq 0$ because in those cases the best choice is unambiguous. Whenever the bounds contain 0, the best we could hope for is 0 regret, which we can always achieve by mimicking π_0 .

Next, we prove that if we approximate the true minimax-optimal policy, by plugging in our own interval CATE estimates in place of the population estimands, then we will achieve optimal minimax regret asymptotically.

Theorem 2. *Let*

$$\hat{\pi}(x; \Gamma) = \mathbb{I}(\hat{\bar{\tau}}(x; \Gamma) \leq 0) + \pi_0(x)\mathbb{I}(\hat{\underline{\tau}}(x; \Gamma) < 0 \leq \hat{\bar{\tau}}(x; \Gamma)).$$

Then, under the assumptions of Theorem 1,

$$\bar{R}_{\pi_0}(\hat{\pi}(\cdot; \Gamma); \Gamma) \xrightarrow{P} \min_{\pi: \mathcal{X} \rightarrow \{0, 1\}} \bar{R}_{\pi_0}(\pi; \Gamma) = \bar{R}_{\pi_0}(\pi^*(\cdot; \Gamma); \Gamma)$$

4.6 Extension: interval estimates for the partial conditional average treatment effect

In subsections 4.1 - 4.5, we consider CATE conditioned on all observed confounders X . However, in many applications, we may be interested in heterogeneity of

¹Note that we can also use $\alpha_t(x; \Gamma)$ and $\beta_t(x; \Gamma)$ as the bounds on W_{ti} in the definition of our estimators in eqs. (12) and (13). In this case, we don't need derivative assumptions on $e_t(\cdot)$. In practice, using x or X_i leads to similar results.

treatment effect in only a few variables, conditioning on only a subset of variables X_S , with $S \subset \{1, \dots, n\}$ as the corresponding index set. For example, in medical applications, fewer rather than more variables are preferred in a personalized decision rule due to cost and interpretability considerations [30]. Therefore, we consider estimation of the *partial* conditional average treatment effect (PCATE):

$$\tau(x_S) = \mu_1(x_S) - \mu_0(x_S), \quad (15)$$

where $\mu_t(x_S) = \mathbb{E}[Y(t) \mid X_S = x_S]$ for $t \in \{0, 1\}$.

Analogously, we define $f_t(y, x_{S^c} \mid x_S)$ as the joint conditional density function of $(T = t, Y(t), X_{S^c})$ given $X_S = x_S$, where $S^c = \{1, \dots, n\} \setminus S$ denotes the complement of S . We further define the following for a weight functional $w_t^P(\cdot, \cdot \mid x_S)$:

$$\mu_t(w_t^P; x_S) = \frac{\iint y w_t^P(x_{S^c}, y \mid x_S) f_t(y, x_{S^c} \mid x_S) dy dx_{S^c}}{\iint w_t^P(x_{S^c}, y \mid x_S) f_t(y, x_{S^c} \mid x_S) dy dx_{S^c}} \quad (16)$$

Since $\mu_t(w_t^P; x_S) = \mu_t(x_S)$ when $w_t^P(x_{S^c}, y \mid x_S) = \frac{1}{e_t(x_S, x_{S^c}, y)}$, we define the population interval estimands for PCATE under the MSM as follows.

Definition 3 (PCATE Identified Set Under MSM). *The population bounds under the MSM with parameter Γ for the partial expected potential outcomes and PCATE are*

$$\bar{\mu}_t(x_S; \Gamma) = \sup_{w_t^P \in \mathcal{W}_t^P(x_S; \Gamma)} \mu(w_t^P; x_S), \quad (17)$$

$$\underline{\mu}_t(x_S; \Gamma) = \inf_{w_t^P \in \mathcal{W}_t^P(x_S; \Gamma)} \mu(w_t^P; x_S), \quad (18)$$

$$\bar{\tau}(x_S; \Gamma) = \bar{\mu}_1(x_S; \Gamma) - \underline{\mu}_0(x_S; \Gamma), \quad (19)$$

$$\underline{\tau}(x_S; \Gamma) = \underline{\mu}_1(x_S; \Gamma) - \bar{\mu}_0(x_S; \Gamma). \quad (20)$$

where $\mu(w_t^P; x_S)$ is defined in (16), and

$$\mathcal{W}_t^P(x_S; \Gamma) = \left\{ w_t^P : w_t^P(x_{S^c}, y \mid x_S) \in [\alpha_t(x_S, x_{S^c}; \Gamma), \beta_t(x_S, x_{S^c}; \Gamma)], \forall x_{S^c} \in \mathcal{X}_{S^c}, \forall y \in \mathcal{Y} \right\}.$$

We can extend our interval estimators to the above:

$$\hat{\mu}_t(\mathbf{W}_t; x_S) = \frac{\sum_{i=1}^n \mathbb{I}(T_i = t) \mathbf{K}\left(\frac{X_{i,S} - x_S}{h}\right) W_{ti} Y_i}{\sum_{i=1}^n \mathbb{I}(T_i = t) \mathbf{K}\left(\frac{X_{i,S} - x_S}{h}\right) W_{ti}}, \quad (21)$$

$$\hat{\bar{\mu}}_t(x_S; \Gamma) = \sup_{\mathbf{W}_t \in \widehat{\mathcal{W}}_t} \hat{\mu}_t(\mathbf{W}_t; x_S), \quad (22)$$

$$\hat{\underline{\mu}}_t(x_S; \Gamma) = \inf_{\mathbf{W}_t \in \widehat{\mathcal{W}}_t} \hat{\mu}_t(\mathbf{W}_t; x_S) \quad (23)$$

These PCATE interval estimators use the partial covariates X_S in the kernel function but the complete

observed covariates X in the nominal propensity score (in $\alpha_t(\cdot)$ and $\beta_t(\cdot)$),² compared to CATE interval estimators, eqs. (9) and (11), that use X in both the kernel function and nominal propensity score. In appendix section D, we prove appropriate analogues of Theorems 1 and 2 for our PCATE interval estimators under analogous assumptions. In this way, we can use the complete observed covariates X in the nominal propensity scores to adjust for confounding as much as possible, so that unobserved confounding is minimal, while only estimating heterogeneity in a subset of interesting covariates.

4.7 Practical Considerations

Boundary bias: For bounded space \mathcal{X} , kernel-regression-based estimators may have high bias at points x near the boundaries. This can be alleviated by truncating kernels at the boundary and replacing any kernel term of the form $\mathbf{K}\left(\frac{X_i - x}{h}\right)$ by a re-normalized version $\mathbf{K}\left(\frac{X_i - x}{h}\right) / \int_{x' \in \mathcal{X}} \mathbf{K}\left(\frac{X_i - x'}{h}\right) dx'$ so that all kernel terms have the same integral over the bounded \mathcal{X} [9, 18]. We take this approach in our experiments.

Propensity score estimation: Although our theoretical results in section 4 assume that the nominal propensity score $e_t(x)$ is known, these results still hold if we use a consistent estimator for it. Indeed, $e_t(x)$ is identifiable. If parametric estimators are used for propensity score estimation, *e.g.*, linear logistic regression, model misspecification error may occur. In this case, we can interpret the marginal sensitivity model as the log odds ratio bound between the complete propensity score and the best parametric approximation of the nominal propensity score. Consequently, the resulting CATE sensitivity bounds also incorporate model misspecification uncertainty. See [40] for more details on marginal sensitivity model for parametric propensity score.

Selection of the sensitivity parameter Γ . The parameter Γ bounds the magnitude of the effects of unobserved confounders on selection, which is usually unknown. [12] suggest calibrating the assumed effect of the unobserved confounders to the effect of observed covariates. For example, we can compute the effect of omitting each observed covariate on the log odds ratio of the propensity score and use domain knowledge to assess plausible ranges of Γ to determine if we could have omitted a variable that could have as large an effect as the observed one.

²We could also use $\alpha_t(x_S, X_{i,S^c}; \Gamma)$ and $\beta_t(x_S, X_{i,S^c}; \Gamma)$ here. See also footnote 1.




Figure 1: Bounds on CATE for differing values of Γ (dashed), compared to original confounded kernel regression (purple dashed) and true CATE (black).

5 Experiments

Simulated Data. We first consider an one-dimensional example illustrating the effects of unobserved confounding on conditional average treatment effect estimation. We generate a binary unobserved confounder $u \sim \text{Bern}(1/2)$ (independent of all else), and covariate $X \sim \text{Unif}[-2, 2]$. We fix the nominal propensity score as $e(x) = \sigma(0.75x + 0.5)$. For the sake of demonstration, we fix an underlying “true” Γ^* value and set the complete propensity scores as $e(x, u) = \frac{u}{\alpha_t(x; \Gamma^*)} + \frac{1-u}{\beta_t(x; \Gamma^*)}$ and sample $T \sim \text{Bern}(e(x, u))$. This makes the complete propensities achieve the extremal MSM bounds corresponding to Γ^* , with u controlling which bound we reach. We choose an outcome model to yield a nonlinear CATE, with linear confounding terms and noise randomly generated as $\epsilon \sim N(0, 1)$:

$$Y(t) = (2t - 1)X + (2t - 1) - 2 \sin(2(2t - 1)X) - 2(2u - 1)(1 + 0.5X) + \epsilon$$

When we learn the confounded effect estimate, $\tilde{\tau}(X)$, from data as in eq. (1), we incur a confounding term ($\tilde{\tau}(x) - \tau(x)$) that grows in magnitude with positive x :

$$2(2 + x)(\Pr[u = 1 \mid \frac{X}{\tilde{\tau}} = 1] - \Pr[u = 1 \mid \frac{X}{\tilde{\tau}} = 0]). \quad (24)$$

In Fig. 1, we compute the bounds using our estimators, eqs. (12) and (13), for varying choices of Γ on a dataset with $n = 2000$ where $\log \Gamma^* = 1$. We use a Gaussian kernel with bandwidths chosen by leave-one-out cross-validation for the task of unweighted regression in each treatment arm. The bounds are centered at the confounded kernel regression estimate of CATE (purple long-dashed line).

By construction, the deviation of the confounded CATE from the true CATE, eq. (24), is greater for more positive x ; our approach learns bounds whose

widths reflect the appropriate “size” of confounding at each x . While the confounded estimation suggests a large region, $x \in [0, 1.25]$, where treatment $\pi(x) = 1$ is optimal, the true CATE suggests that treatment at many of these x values is harmful, and our bounds suggest that the benefit of treatment is ambiguous in the confounded-optimal region.

In Section E.1 of the Appendix, we include a similar example to illustrate how the bounds would change for assessing partial CATE: additional confounding dimensions tend to increase the outcome variation, so the bounds widen.

In Table 1, we compare the *true* policy values, $V(\pi; \tau)$, achieved by the decision rules derived from our interval CATE estimates, following Section 4.5 and letting $\pi_0(x) = 0$ (never treat). We consider 20 Monte Carlo replications for each setting of Γ^* and report 95% confidence intervals. Any omitted confidence interval is smaller than ± 0.01 . Note that on the diagonal of Table 1, we assess a policy with a “well-specified” Γ equal to Γ^* , which achieves the best risk for the corresponding data generating process. The case of $n = 5000$ essentially gives the population-level optimal minimax regret. Finally, for comparison, we include the policy values of both the thresholding policy based on the confounded CATE learned by IPW-weighted kernel regression using nominal propensities ($\hat{\tau}(x)$) and the truly optimal policy based on the true (and unknowable) τ . The policy value of the confounded policy suffers in comparison to the policies from our estimated bounds. Specifying an overly conservative Γ achieves similar risk in this setting, while underspecifying Γ compared to the true Γ^* incurs greater loss.

Assessment on Real-World Data: Hormone Replacement Therapy. To illustrate the impacts of unobserved confounding, we consider a case study of a parallel clinical trial and large observational study from the Women’s Health Initiative [28]. Hormone replacement therapy (HRT) was the treatment of interest: previous observational correlations suggested protective effects for onset of chronic (including cardiovascular) disease. While the clinical trial was halted early due to dramatically increased incidence of heart attacks, the observational study evidence actually suggested preventive effects, prompting further study to reconcile these conflicting findings based on unobserved confounding in the observational study [20, 29, 32]. Follow-up studies suggest benefits of HRT for younger women [5].

We consider a simple example of learning an optimal treatment assignment policy based on age to reduce endline systolic blood pressure, which serves as a proxy

Table 1: Policy risk for various policies under data generating processes with different Γ^* (lower is better)

$\log(\Gamma^*)$	$n = 1000$			$n = 5000$			$n = 1000$	$n = \infty$
	$\hat{\pi}(x; e^{0.5})$	$\hat{\pi}(x; e^1)$	$\hat{\pi}(x; e^{1.5})$	$\hat{\pi}(x; e^{0.5})$	$\hat{\pi}(x; e^1)$	$\hat{\pi}(x; e^{1.5})$	$\mathbb{I}(\hat{\tau}(x) < 0)$	$\mathbb{I}(\tau(x) < 0)$
0.5	-1.65 ± 0.01	-1.64 ± 0.01	-1.63 ± 0.00	-1.68	-1.63	-1.63	-1.60 ± 0.00	-1.68
1	-1.58 ± 0.02	-1.64 ± 0.02	-1.64 ± 0.01	-1.62	-1.67	-1.63	-1.48 ± 0.04	-1.68
1.5	-1.51 ± 0.02	-1.60 ± 0.02	-1.63 ± 0.02	-1.52	-1.63	-1.67	-1.36 ± 0.04	-1.68




Figure 2: CATE estimated from unconfounded data (CT) vs. confounded observational study data (OS): confounding leads to opposite conclusions.

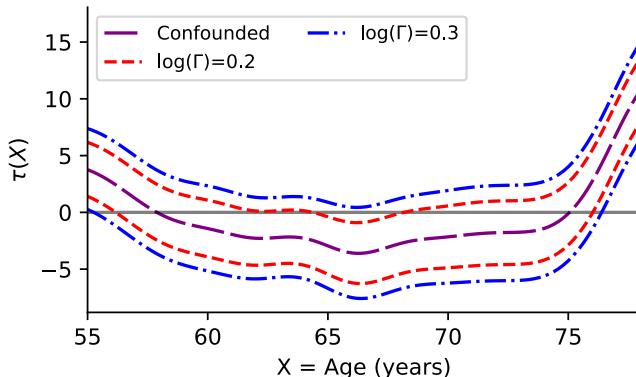


Figure 3: Bounds on CATE estimated from the WHI observational study. The observational study is highly sensitive to unobserved confounding.

outcome for protective effects against cardiovascular disease. Thus we consider learning the PCATE for $S = \{\text{age}\}$ while controlling for *all* observed baseline variables. The observed covariates are 30-dimensional (after binary encodings of categorical variables) and include factors such as demographics, smoking habits, cardiovascular health history, and other comorbidities (*e.g.*, diabetes and myocardial infection). We restrict attention to a complete-case subset of the clinical-trial data ($n = 14266$), and a subset of the observational study ($n = 2657$).

For comparing findings from the clinical trial and ob-

servational study, Fig. 2 plots estimates of the partial conditional average treatment effect on systolic blood pressure over age by a difference of LOESS regressions. In the clinical trial (CT, orange) we used a simple LOESS regression on age for each treatment condition, and in the observational study (OS, blue) we used IPW-weighted regression with propensities estimated using all observed baseline variables. A negative CATE suggests that HRT reduces systolic blood pressure and might have protective effects against cardiovascular disease. The clinical trial CATE is not statistically significantly different from zero for ages above 67. The observational CATE crucially displays the *opposite* trend, suggesting that treatment is statistically significantly beneficial for women of ages 62-73. We display 90% confidence intervals obtained from a 95% confidence interval for individual regressions within each treatment arm.

In Fig. 3, we apply our method to estimate bounds on $\tau(X_{\text{age}})$. We estimate propensity scores using logistic regression. Our bounds suggest that the estimated CATE from the observational study is highly sensitive to potential unobserved confounding: for sensitivity parameter as low as $\log(\Gamma) = 0.2$ ($\Gamma = 1.22$), we see that the sensitivity bounds contain 0 for nearly all individuals such that we would likely prefer to less intervention. To interpret this value of Γ , we compute the distribution of instance-wise $\Gamma_{i,j}$ parameters between the propensity estimated from all covariates for individual i , and the propensity estimated under dropping each covariate dimension j : $\Gamma_{i,j} = \frac{(1-e_{t_i}(X_i))e_{t_i}(X_{i,-j})}{e_{t_i}(X_i)(1-e_{t_i}(X_{i,-j}))}$. The maximal such Γ value is observed by dropping the indicator for 1-4 cigarettes smoked per day, which leads to a maximal $\Gamma = 1.17$ value.

6 Conclusion

We developed a functional interval estimator that provides bounds on individual-level causal effects under realistic violations of unconfoundedness. Our estimators, which we prove are sharp for the tightest bounds possible, use a weighted kernel estimator with weights that vary adversarially over an uncertainty set consistent with a sensitivity model.

References

- [1] J. Abrevaya, Y.-C. Hsu, and R. P. Lieli. Estimating conditional average treatment effects. *Journal of Business & Economic Statistics*, 33(4):485–505, 2015.
- [2] P. Aronow and D. Lee. Interval estimation of population means under unknown but bounded probabilities of sample selection. *Biometrika*, 2012.
- [3] S. Athey and G. Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.
- [4] S. Athey and S. Wager. Efficient policy learning. *arXiv preprint arXiv:1702.02896*, 2017.
- [5] S. H. Bakour and J. Williamson. Latest evidence on using hormone replacement therapy in the menopause. *The Obstetrician & Gynaecologist*, 17(1):20–28, 2015.
- [6] J. Cornfield, W. Haenszel, E. C. Hammond, A. M. Lilienfeld, M. B. Shimkin, and E. L. Wynder. Smoking and lung cancer: recent evidence and a discussion of some questions. *Journal of the National Cancer Institute*, 22(1):173–203, 1959.
- [7] V. Dorie, M. Harada, N. B. Carnegie, and J. Hill. A flexible, interpretable framework for assessing sensitivity to unmeasured confounding. *Statistics in Medicine*, 2016.
- [8] M. Dudik, D. Erhan, J. Langford, and L. Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 2014.
- [9] T. Gasser and H.-G. Müller. Kernel estimation of regression functions. In *Smoothing techniques for curve estimation*, pages 23–68. Springer, 1979.
- [10] D. P. Green and H. L. Kern. Modeling heterogeneous treatment effects in large-scale experiments using bayesian additive regression trees. In *The annual summer meeting of the society of political methodology*, 2010.
- [11] E. Hartman, R. Grieve, R. Ramsahai, and J. S. Sekhon. From sate to patt: combining experimental with observational studies to estimate population treatment effects. *JR Stat. Soc. Ser. A Stat. Soc. (forthcoming)*. doi, 10:1111, 2015.
- [12] J. Y. Hsu and D. S. Small. Calibrating sensitivity analyses to observed covariates in observational studies. *Biometrics*, 69(4):803–811, 2013.
- [13] F. D. Johansson, N. Kallus, U. Shalit, and D. Sonntag. Learning weighted representations for generalization across designs. *on ArXiv*, 2018.
- [14] N. Kallus. Recursive partitioning for personalization using observational data. In *International Conference on Machine Learning (ICML)*, pages 1789–1798, 2017.
- [15] N. Kallus. Balanced policy evaluation and learning. *To appear in Proceedings of Neural Information Processing Systems*, 2018.
- [16] N. Kallus and A. Zhou. Policy evaluation and optimization with continuous treatments. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1243–1251, 2018.
- [17] N. Kallus and A. Zhou. Confounding-robust policy improvement. *To appear in Proceedings of Neural Information Processing Systems*, 2018.
- [18] S. Kheireddine. *On Boundary Correction in Kernel Estimation*. PhD thesis, Université Mohamed Khider-Biskra, 2016.
- [19] S. R. Künzel, J. S. Sekhon, P. J. Bickel, and B. Yu. Meta-learners for estimating heterogeneous treatment effects using machine learning. *arXiv preprint arXiv:1706.03461*, 2017.
- [20] D. A. Lawlor, G. D. Smith, and S. Ebrahim. Commentary: The hormone replacement-coronary heart disease conundrum: is this the death of observational epidemiology? *International Journal of Epidemiology*, 2004.
- [21] B. K. Lee, J. Lessler, and E. A. Stuart. Improving propensity score weighting using machine learning. *Statistics in medicine*, 29(3):337–346, 2010.
- [22] C. Manski. *Social Choice with Partial Knowledge of Treatment Response*. The Econometric Institute Lectures, 2005.
- [23] C. F. Manski. *Partial identification of probability distributions*. Springer Science & Business Media, 2003.
- [24] M. Masten and A. Poirier. Identification of treatment effects under conditional partial independence. *Econometrica*, 2018.
- [25] D. F. McCaffrey, G. Ridgeway, and A. R. Morral. Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychological methods*, 9(4):403, 2004.
- [26] X. Nie and S. Wager. Learning objectives for treatment effect estimation. *arXiv preprint arXiv:1712.04912*, 2017.
- [27] A. Pagan and A. Ullah. *Nonparametric econometrics*. Cambridge university press, 1999.
- [28] A. T. Pedersen and B. Ottesen. Issues to debate on the women’s health initiative (whi) study. epidemiology or randomized clinical trials-time out for hormone replacement therapy studies? *Human Reproduction*, 2003.

- [29] R. L. Prentice, M. Pettinger, and G. L. Anderson. Statistical issues arising in the women’s health initiative. *Biometrics*, 2005.
- [30] M. Qian and S. A. Murphy. Performance guarantees for individualized treatment rules. *Annals of statistics*, 39(2):1180, 2011.
- [31] P. Rosenbaum. *Observational Studies*. Springer Series in Statistics, 2002.
- [32] J. E. Rossouw, J. E. Manson, A. M. Kaunitz, and G. L. Anderson. Lessons learned from the women’s health initiative trials of menopausal hormone therapy. *Obstetrics & Gynecology.*, 2013.
- [33] D. B. Rubin. Comments on “randomization analysis of experimental data: The fisher randomization test comment”. *Journal of the American Statistical Association*, 75(371):591–593, 1980.
- [34] U. Shalit, F. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- [35] A. Shapiro. *Semi-Infinite Programming*, chapter On Duality Theory of Conic Linear Problems, pages 135–165. 2001.
- [36] A. Swaminathan and T. Joachims. Counterfactual risk minimization. *Journal of Machine Learning Research*, 2015.
- [37] Z. Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 2012.
- [38] S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, (just-accepted), 2017.
- [39] S. Yadlowsky, H. Namkoong, S. Basu, J. Duchi, and L. Tian. Bounds on the conditional and average treatment effect in the presence of unobserved confounders. *arXiv preprint arXiv:1808.09521*, 2018.
- [40] Q. Zhao, D. S. Small, and B. B. Bhattacharya. Sensitivity analysis for inverse probability weighting estimators via the percentile bootstrap. *ArXiv*, 2017.
- [41] X. Zhou, N. Mayer-Hamblett, U. Khan, and M. R. Kosorok. Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, 112(517):169–187, 2017.