COVARIATE BALANCING FOR ESTIMATING CONTINUOUS TREATMENT EFFECTS USING SUPREMUM OF INTEGRAL PROBABILITY METRICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Estimating causal effects from observational data requires controlling for confounding. Weighting approaches seek to remove dependence between treatment and confounders by reweighting the units. We propose Covariate Balancing using Supremum of Integral Probability Metrics (CBSupIPM), a weighting method for continuous treatments. CBSupIPM learns weights by minimizing the largest discrepancy, over treatment levels, between the weighted conditional covariate distribution and the population covariate distribution. We establish uniform convergence of the weighted Nadaraya–Watson estimator of the average dose–response function (ADRF) based on CBSupIPM weights. When an ideal weight function is sufficiently smooth and the true response surface belongs to a suitable class, the estimator matches the nonparametric convergence rate of the corresponding regression estimator in an unconfounded setting, up to logarithmic factors. Also, uniform consistency is guaranteed even when the true response surface belongs to a broader class, such as uniformly Lipschitz continuous functions in the covariates. For the numerical experiments, we construct a realistic semi-synthetic dataset using the National Medical Expenditure Survey (NMES). We also use repeated simulations with synthetic data to assess method performance more reliably and examine how it changes as different data-generating factors vary. In these numerical experiments, CBSupIPM shows the most favorable empirical performance in terms of ADRF estimation error among the comparators in most settings.

1 INTRODUCTION

In observational studies, treatment is not assigned at random. Consequently, observed outcome differences across treatment levels can reflect differences in pretreatment covariates as well as treatment effects. Weighting seeks to reduce confounding by balancing the distributions of observed pretreatment covariates across treatment groups (Li et al., 2018).

In this paper, we consider continuous treatments such as dosage, smoking exposure, and pollution. Two adjustment strategies for assessing the causal effect of continuout treatments are regression adjustment and weighting. Regression adjustment models the expected outcome given covariates and treatment, then averages over covariates to estimate the average dose–response function (ADRF) (Singh et al., 2024; Schwab et al., 2020; Park et al., 2022; Cai et al., 2023). Misspecifying that model, however, can bias the resulting estimate.

Weighting approximately removes treatment–covariate dependence by reweighting units. A common approach constructs weights based on the generalized propensity score (GPS), the conditional treatment density given covariates (Robins et al., 2000; Fong et al., 2018). Since the weights are proportional to the inverse of the conditional density of the treatment, small density regions could produce unnecessary large weights. In addition, model misspecification can leave residual confounding.

These concerns motivate approaches that construct weights by directly targeting covariate balance. Without modeling the density of the treatment, nonparametric covariate balancing generalized propensity score (npCBGPS) (Fong et al., 2018) and entropy balancing for continuous treatments

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

(EBCT) (Vegetabile et al., 2021; Tübbicke, 2022) construct weights so that the reweighted data satisfy certain moment constraints on interactions between treatment and covariates. Among feasible weights, npCBGPS selects those maximizing the empirical likelihood, whereas EBCT selects those minimizing the Kullback entropy metric.

Instead of moment conditions, distributional approaches try to make the entire distribution of the treatment and covariates become independent. Examples include the kernel optimal orthogonality weighting (KOOW), covariate balance via discrepancy minimization (CBDM), and distance covariance optimal weights (DCOW). These approaches minimizes a discrepancy between the weighted joint distribution of covariates and treatment with the product of their marginal distributions.

In this paper, we propose Covariate Balancing using Supremum of Integral Probability Metrics (CBSupIPM), which minimizes the largest discrepancy across treatment levels between the weighted conditional covariate distribution and the marginal covariate distribution. Our contributions are as follows. First, we use weighted population distributions to formulate the weighting objective. We decompose the error of the weighted ADRF estimator to clarify how covariate balance affects the ADRF estimation. Second, we establish uniform convergence rates over trimmed regions that expand to the interior of the treatment support. Under the stated conditions, sufficient smoothness of the response surface and balancing score, and suitable choices of network architecture yield an $h_n^2 + (nh_n)^{-1/2}$ convergence rate up to logarithmic factors. This matches the polynomial order of the pointwise convergence rate established for penalized DCOW (Huling et al., 2024). Third, the uniform error bounds yield corresponding rate bounds for estimating the maximum ADRF value over the same treatment regions. Finally, CBSupIPM achieves lower errors in estimating both the ADRF and its maximum value than the competing methods.

2 PRELIMINARIES

2.1 NOTATION AND ASSUMPTIONS

Let $\boldsymbol{X} \in \mathcal{X} \subseteq \mathbb{R}^{d_X}$, $T \in \mathcal{T} \subseteq \mathbb{R}$, and $Y^t \in \mathcal{Y} \subseteq \mathbb{R}$ be pretreatment covariates, treatment, and the potential outcome at treatment level t, respectively. All random elements are defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and $\mathbb{E}$ denotes expectation with respect to $\mathbb{P}$.

For random elements U and V, let $\mathbb{P}_U := \mathbb{P} \circ U^{-1}$, denote its Lebesgue density by p_U when it exists, and write $\mathbb{P}_{U|V=v}$ for the conditional distribution of U given $V = v$. For any integrable f,

$$\mathbb{E}_{U|V=v}[f(U,V)] := \int f(u,v)\, \mathbb{P}_{U|V=v}(du) = \mathbb{E}[f(U,V) \mid V = v].$$

We observe i.i.d. samples $\{(\boldsymbol{X}_i, T_i, Y_i)\}_{i=1}^n$ from the observed data distribution $\mathbb{P}_{\boldsymbol{X},T,Y}$, where $Y = Y^T$. We define the average dose-response function (ADRF)

$$\mu(t) = \mathbb{E}[Y^t], \quad t \in \mathcal{T}.$$

To estimate ADRF based on observed data, as usual, we assume the following two causal conditions:

Assumption 2.1 (Weak unconfoundedness)**.** For every $t \in \mathcal{T}$,

$$Y^t \perp\!\!\!\perp T \mid \boldsymbol{X}$$

Weak unconfoundedness means that, conditional on covariates, treatment assignment gives no additional information about potential outcome at each treatment level. In other words, there are no unobserved confounders.

Assumption 2.2 (Overlap)**.** For $\mathbb{P}_{\boldsymbol{X}}$-a.e. $\boldsymbol{x} \in \mathcal{X}$,

$$\mathbb{P}_T \ll \mathbb{P}_{T|\boldsymbol{X}=\boldsymbol{x}}$$

The symbol $\ll$ denotes absolute continuity. Overlap means that treatment sets with positive marginal probability must also have positive conditional probability for almost every covariate value.

2.2 WEIGHTED NADARAYA–WATSON ADRF ESTIMATOR

We define the true response surface by $f^*(\boldsymbol{x}, t) := \mathbb{E}[Y^t \mid \boldsymbol{X} = \boldsymbol{x}]$. Under the above assumptions, f^* is identified from the observed data distribution $\mathbb{P}_{\boldsymbol{X},T,Y}$ through

$$f^*(\boldsymbol{x}, t) = \mathbb{E}[Y^t \mid \boldsymbol{X} = \boldsymbol{x}] = \mathbb{E}[Y \mid \boldsymbol{X} = \boldsymbol{x}, T = t]. \tag{1}$$

Thus, ADRF is identified from $\mathbb{P}_{\boldsymbol{X},T,Y}$ through

$$\mu(t) = \int f^*(\boldsymbol{x}, t)\, \mathbb{P}_{\boldsymbol{X}}(d\boldsymbol{x}).$$

For given $W : \mathcal{X} \times \mathcal{T} \to [0, \infty)$ with $0 < \mathbb{E}[W(\boldsymbol{X}, t) \mid T = t] < \infty$, we define the weighted conditional covariate law $\mathbb{P}^W_{\boldsymbol{X}|T=t}$ by

$$\mathbb{P}^W_{\boldsymbol{X}|T=t}(d\boldsymbol{x}) := \frac{W(\boldsymbol{x}, t)\mathbb{P}_{\boldsymbol{X}|T=t}(d\boldsymbol{x})}{\mathbb{E}[W(\boldsymbol{X}, t) \mid T = t]}. \tag{2}$$

Similarly, define the corresponding weighted conditional outcome law by

$$\mathbb{P}^W_{Y|T=t}(dy) := \int_{\mathcal{X}} \mathbb{P}_{Y|\boldsymbol{X}=\boldsymbol{x},T=t}(dy)\, \mathbb{P}^W_{\boldsymbol{X}|T=t}(d\boldsymbol{x}). \tag{3}$$

Definition 2.1 (Exact balancing weight function)**.** We call W an *exact balancing weight function* if

$$\mathbb{P}^W_{\boldsymbol{X}|T=t} = \mathbb{P}_{\boldsymbol{X}}, \qquad \forall t \in \mathcal{T}. \tag{4}$$

Under the overlap condition, $W^*(\boldsymbol{x}, t) = \dfrac{d\mathbb{P}_T}{d\mathbb{P}_{T|\boldsymbol{X}=\boldsymbol{x}}}(t)$ becomes an exact balancing weight function. When treatment densities exist, we have $W^*(\boldsymbol{x}, t) = p_T(t)/p_{T|\boldsymbol{X}=\boldsymbol{x}}(t)$.

Under $\mathbb{E}[W(\boldsymbol{X}, t) \mid T = t] = 1$, the exact balancing weight function W is unique almost everywhere.

For the exact balancing weight function W, weak unconfoundedness gives

$$\mathbb{P}_{Y^t}(dy) = \mathbb{P}^W_{Y|T=t}(dy). \tag{5}$$

Define

$$\mu_W(t) := \frac{\mathbb{E}_{(\boldsymbol{X},Y)|T=t}[W(\boldsymbol{X}, t)Y]}{\mathbb{E}_{\boldsymbol{X}|T=t}[W(\boldsymbol{X}, t)]}. \tag{6}$$

$$= \frac{\mathbb{E}_{(\boldsymbol{X},Y)|T=t}[W(\boldsymbol{X}, T)Y]}{\mathbb{E}_{\boldsymbol{X}|T=t}[W(\boldsymbol{X}, T)]}. \tag{7}$$

Then, $\mu(t) = \mu_W(t)$ for an exact balancing weight function W. The proof is attached in Appendix B.1.

Under $\mathbb{E}[W(\boldsymbol{X}, t) \mid T = t] = 1$, the exact balancing weight function W is unique almost everywhere. Thus, given an estimator $\widehat{W}$ of the exact balancing weight function, we can construct the corresponding weighted Nadaraya–Watson(NW) ADRF estimator as

$$\widehat{\mu}_{\widehat{W}}(t) = \frac{\sum_{i=1}^n K_h(T_i - t)\widehat{W}(\boldsymbol{X}_i, t)Y_i}{\sum_{i=1}^n K_h(T_i - t)\widehat{W}(\boldsymbol{X}_i, t)}, \tag{8}$$

where $K_h(u) = h^{-1}K(u/h)$, K is a kernel, and $h > 0$ is the bandwidth. Note that $\mu_W(t)$ is the population target of the $\widehat{\mu}_W$.

However, most existing weighting methods estimate weight values at the observed units, rather than the balancing weight function W itself. For a given weighting method A, let $\widehat{w}_i^A$ denote the weight assigned to unit i. Within our formulation, we interpret these weights as estimates of $W(\boldsymbol{X}_i, T_i)/\sum_i W(\boldsymbol{X}_i, T_i)$ of the exact balancing weight function W. Substituting these weights into kernel estimates of the conditional expectations in Eq. (7) yields

$$\widehat{\mu}_A(t) = \frac{\sum_{i=1}^n K_h(T_i - t)\widehat{w}_i^A Y_i}{\sum_{i=1}^n K_h(T_i - t)\widehat{w}_i^A}. \tag{9}$$

We use this estimator in our numerical experiments. The construction of $(\widehat{w}_i^A)_{i=1}^n$ for various methods are detailed in Appendix A.1.

2.3 Review of Weighting Methods

We summarize the shared principles in existing weighting methods; see Appendix A.1 for detailed explanations.

Weighting based on the GPS estimates $p_{T|\boldsymbol{X}}(t \mid \boldsymbol{x})$ and calculate the sample weights as $\widehat{w}_i = \widehat{p}_T(T_i)/\widehat{p}_{T|\boldsymbol{X}}(T_i \mid \boldsymbol{X}_i)$ (Robins et al., 2000; Hirano & Imbens, 2004). The covariate balancing generalized propensity score (CBGPS) method uses the moment balance equation for estimating $p_{T|\boldsymbol{X}}(t \mid \boldsymbol{x})$ (Fong et al., 2018).

Instead of estimating $p_{T|\boldsymbol{X}}(t \mid \boldsymbol{x})$, npCBGPS and EBCT estimate the sample weights directly by maximizing the empirical likelihood or the relative entopry under certain moment balance constraints (Fong et al., 2018; Vegetabile et al., 2021; Tübbicke, 2022). Distributional approaches aim to find a weight vector $\boldsymbol{w}$ such that

$$\widehat{\mathbb{P}}^{\boldsymbol{w}}_{\boldsymbol{X},T} \approx \widehat{\mathbb{P}}_{\boldsymbol{X}} \otimes \widehat{\mathbb{P}}_T,$$

Examples are KOOW (Kallus & Santacatterina, 2019), CBDM (Martinet, 2020) and DCOW (Huling et al., 2024).

Our approach is to find a weight vector such that $\widehat{\mathbb{P}}^{\boldsymbol{w}}_{\boldsymbol{X}|T=t} \approx \widehat{\mathbb{P}}_{\boldsymbol{X}}$ for all t. An advantage is to be able to control the error between $\widehat{\mathbb{P}}^{\boldsymbol{w}}_{\boldsymbol{X}|T=t}$ and $\widehat{\mathbb{P}}_{\boldsymbol{X}}$ uniformly in t, whihc results in the uniform consistency of the corresponding ADRF estimator.

3 Covariate Balancing using Supremum of IPMs

3.1 SupIPM risk and bias control

For two distributions $\mathbb{P}, \mathbb{Q}$ and a discriminator class $\mathcal{F}$, we define $\mathrm{IPM}_{\mathcal{F}}$ as

$$\mathrm{IPM}_{\mathcal{F}}(\mathbb{P}, \mathbb{Q}) = \sup_{f \in \mathcal{F}} \left| \int f \, d\mathbb{P} - \int f \, d\mathbb{Q} \right|. \tag{10}$$

Thus a small IPM means that the two distributions have similar expectations for every $f \in \mathcal{F}$.

CBSupIPM aims to choose the weight function W which minimizes the essential supremum of IPM over the interior of the treatment support:

$$\mathcal{R}_{\mathcal{F}}(W) = \sup_{t \in \mathrm{int}(\mathcal{T})} \mathrm{IPM}_{\mathcal{F}} \left(\mathbb{P}^{W}_{\boldsymbol{X}|T=t}, \mathbb{P}_{\boldsymbol{X}} \right). \tag{11}$$

The following proposition shows that this population risk controls the discrepancy between the population target of the weighted NW ADRF estimator associated with W and the ADRF.

Proposition 3.1. *Suppose that there exists a constant $C_{f^*} > 0$ such that $f^*(\cdot, t)/C_{f^*} \in \mathcal{F}$ for almost every $t \in \mathrm{int}(\mathcal{T})$. Then,*

$$\|\mu_W - \mu\|_{L_\infty(\mathrm{int}(\mathcal{T}))} \leq C_{f^*} \sup_{t \in \mathrm{int}(\mathcal{T})} \mathrm{IPM}_{\mathcal{F}} \left(\mathbb{P}^{W}_{\boldsymbol{X}|T=t}, \mathbb{P}_{\boldsymbol{X}} \right) = C_{f^*} \mathcal{R}_{\mathcal{F}}(W)$$

Appendix B.2 proves this bound.

We parameterize a positive normalized weight function through a score s as follows:

$$W_s(\boldsymbol{x}, t) = \frac{\exp\{s(\boldsymbol{x}, t)\}}{\mathbb{E}[\exp\{s(\boldsymbol{X}, t)\} \mid T = t]}. \tag{12}$$

We model the score as an element of a given sieve $\mathcal{S}_n$; a class of neural networks whose size may grow with the sample size.

Let

$$\widehat{\mathbb{P}}^{W_s}_{\boldsymbol{X}|T=t} = \sum_{i=1}^{n} A_i^s(t) \delta_{\boldsymbol{X}_i}, \qquad \widehat{\mathbb{P}}_{\boldsymbol{X}} = \frac{1}{n} \sum_{i=1}^{n} \delta_{\boldsymbol{X}_i}.$$

The empirical SupIPM risk is defined as

$$\widehat{R}_{\mathcal{F}}(W_s) := \sup_{t\in\mathcal{T}_{\epsilon_n}} \mathrm{IPM}_{\mathcal{F}}\left(\widehat{\mathbb{P}}^{W_s}_{\boldsymbol{X}|T=t}, \widehat{\mathbb{P}}_{\boldsymbol{X}}\right). \tag{13}$$

We estimate the CBSupIPM score by:

$$\widehat{s}^{\mathrm{CBSupIPM}} \in \operatorname*{argmin}_{\boldsymbol{s}\in\mathcal{S}_n} \widehat{\mathcal{R}}_{\mathcal{F}}(W_s) = \operatorname*{argmin}_{\boldsymbol{s}\in\mathcal{S}_n} \sup_{t\in\mathcal{T}_{\epsilon_n}} \mathrm{IPM}_{\mathcal{F}}\left(\widehat{\mathbb{P}}^{W_s}_{\boldsymbol{X}|T=t}, \widehat{\mathbb{P}}_{\boldsymbol{X}}\right).$$

3.2 ReLU and sigmoid discriminator classes

The IPM depends on the discriminator class $\mathcal{F}$. We consider the class $\mathcal{F}_{\mathrm{ReLU}}$ for the discriminator class: Let $\mathbb{B}^{dx} = \{z : \|z\|_2 \le 1\}$ and $\mathbb{S}^{dx-1} = \{z : \|z\|_2 = 1\}$. Following Park et al. (2025), define the ReLU discriminator on the unit ball:

$$\mathcal{G}_{\mathrm{ReLU}} = \left\{g_{\gamma,b} : \mathbb{B}^{dx} \to [0,\infty) : g_{\gamma,b}(z) = (\gamma^\top z + b)_+,\ \gamma \in \mathbb{S}^{dx-1},\ b \in [-1,1]\right\}. \tag{14}$$

To apply this class for $\boldsymbol{x} \in [-1,1]^{dx}$, we use the transformation

$$\psi(\boldsymbol{x}) = \begin{cases} \boldsymbol{x}\dfrac{\|\boldsymbol{x}\|_\infty}{\|\boldsymbol{x}\|_2}, & \boldsymbol{x} \neq 0, \\ 0, & \boldsymbol{x} = 0 \end{cases} \tag{15}$$

which maps $[-1,1]^{dx}$ to $\mathbb{B}^{dx}$. The resulting function class space is

$$\mathcal{F}_{\mathrm{ReLU}} = \{g_{\gamma,b} \circ \psi : g_{\gamma,b} \in \mathcal{G}_{\mathrm{ReLU}}\} \tag{16}$$

$$= \{f_{\gamma,b} : [-1,1]^{dx} \to [0,\infty),\ f_{\gamma,b}(\boldsymbol{x}) = \max(\gamma^\top \psi(\boldsymbol{x}) + b,\, 0)\}. \tag{17}$$

3.3 Computation of SupIPM Objective

Define the local weights

$$A_i^s(t) = \frac{K_{h_n}(T_i - t)\exp\{s(\boldsymbol{X}_i, t)\}}{\sum_{j=1}^n K_{h_n}(T_j - t)\exp\{s(\boldsymbol{X}_j, t)\}}. \tag{18}$$

Then, the empirical risk can be calculated by

$$\sup_{t\in\mathcal{T}_{\epsilon_n}} \mathrm{IPM}_{\mathcal{F}}\left(\widehat{\mathbb{P}}^{W_s}_{\boldsymbol{X}|T=t}, \widehat{\mathbb{P}}_{\boldsymbol{X}}\right) = \sup_{t\in\mathcal{T}_{\epsilon_n}} \sup_{f\in\mathcal{F}} \left|\sum_{i=1}^n A_i^s(t) f(\boldsymbol{X}_i) - \frac{1}{n}\sum_{i=1}^n f(\boldsymbol{X}_i)\right|.$$

Thus, for the discriminator class $\mathcal{F} = \mathcal{F}_{\mathrm{ReLU}}$, the resulting problem has the form

$$\min_{s\in\mathcal{S}_n} \max_{t\in\mathcal{G}_n} \sup_{f\in\mathcal{F}} \left|\sum_{i=1}^n A_i^s(t) f(\boldsymbol{X}_i) - \frac{1}{n}\sum_{i=1}^n f(\boldsymbol{X}_i)\right|.$$

This minimax problem is solved by iterative gradient updates.

4 Uniform Consistency of CBSupIPM estimator

Using $A_i^s(t)$ from (18), write the corresponding NW estimator as $\widehat{\mu}^s(t) = \sum_i A_i^s(t) Y_i$ for μ. The region $\mathcal{T}_{\epsilon_n}$ excludes a boundary strip of width ϵ_n, shrinking as n grows. S_n is the neural network with ReLU activation function. Let $d = d_X + 1$; L_n, H_n are the depth and width of $\mathcal{S}_n$. The conditions in Appendix B.4 require Hölder smoothness of the true response surface, treatment smoothness, controlled noise, regular kernel and treatment density, and Hölder smoothness of the exactly balancing score s^*, together with bounds on network size and bandwidth.

Theorem 4.1 (Uniform convergence of ADRF estimator under Hölder smooth response surfaces)**.** *Under Assumption B.1 –B.8, up to logarithmic factors, for* $q_\alpha = \min\left\{1, \frac{2\alpha}{dx+3}\right\}$,

$$\left\|\widehat{\mu}^{\widehat{s}_n} - \mu\right\|_{L_\infty(\mathcal{T}_{\epsilon_n})} = O_P\left[\left\{h_n^2 + (nh_n)^{-1/2} + H_n^{-\beta/d}\right\}^{q_\alpha} + \sqrt{\frac{h_n L_n H_n^2}{n}}\right].$$

Remark : Taking $h_n \asymp n^{-1/5}$, $\epsilon_n \asymp (\log n)^{-1}$, $L_n \asymp \log n$, $H_n \asymp n^{3d/\{5(d+\beta)\}}$, gives, when $\alpha \geq (d_X + 3)/2$ and $\beta \geq 2d$,

$$\left\| \widehat{\mu}^{\widehat{s}_n} - \mu \right\|_{L_\infty(\mathcal{T}_{\epsilon_n})} = O_P(n^{-2/5}) \tag{19}$$

up to logarithmic factors. Theorem 3.7 of Huling et al. (2024) states, under its moment and regularity conditions, that at a fixed treatment level t_0,

$$\widehat{\mu}_{\mathrm{DCOW}}(t_0) - \mu(t_0) = O_P\{(nh_n)^{-1/2} + h_n^2\}.$$

When the true response surface is Hölder smooth, the CBSupIPM bound holds uniformly on every compact subset of $\mathrm{int}(\mathcal{T})$, with the same convergence rate, up to logarithmic factors, as the DCOW pointwise convergence rate

Uniform convergence also controls the estimation error of the maximum ADRF value, since

$$\left| \sup_{t \in \mathcal{T}_{\epsilon_n}} \widehat{\mu}^{\widehat{s}_n}(t) - \sup_{t \in \mathcal{T}_{\epsilon_n}} \mu(t) \right| \leq \left\| \widehat{\mu}^{\widehat{s}_n} - \mu \right\|_{L^\infty(\mathcal{T}_{\epsilon_n})}.$$

In fact, our numerical experiments demonstrate the effectiveness of CBSupIPM in estimating the maximum ADRF value.

5 Experiments

We evaluate the empirical performance of CBSupIPM using semi-synthetic data based on NMES data (Bryer, 2025) and synthetic data.

5.1 Synthetic Experiments

For the synthetic study, we consider linear and nonlinear designs based on Kang & Schafer (2007). We conducted 50 replications for each data-generating process setting to ensure the stability of the performance of the ADRF estimator. The data generating mechanisms and corresponding ADRFs are specified in Appendix C. The linear and nonlinear designs each use 2000 training and 2000 independent test observations per repetition. The baseline is $(d_X, k_T, k_Y, \pi_0, b_T) = (5, 5, 5, 0.2, 0.2)$, where k_T and k_Y count covariates relevant to treatment and outcome generation, $\pi_0 = \Pr(T = 0)$, and b_T sets the treatment coefficient scale. We conduct sensitivity analyses varying (d_X, k_T, k_Y) as one factor, π_0, or b_T, holding the other factors fixed. The results are shown in Table 9 -14.

5.2 Semi-Synthetic Experiments

For the semi-synthetic study, we use data from the 1987 National Medical Expenditure Survey (NMES), available as `nmes` in the R package `TriMatch`. Treatment T is cumulative smoking exposure measured in pack-years, and Y denotes total medical expenditure. We use the transformed outcome $Y^* = \log(1+Y)$. After excluding observations with missing values in treatment, outcome, or the eight selected covariates, the sample contains 19,352 individuals, including 9,804 with $T = 0$. For NMES, we retain the observed covariates and treatment values and generate potential outcomes using a fitted outcome regression and randomly permuted, centered residuals. See Appendix C.2 for details of NMES data generation.

5.3 Baselines and Evaluation

ADRF estimation over the treatment distribution. We compare CBSupIPM with GPS, DCOW, and EBCT in a semi-synthetic NMES study and linear and nonlinear synthetic studies.

For replication $r \in \{1, \ldots, R\}$ and weighting method A, let μ_r denote the true ADRF, and let $\widehat{\mu}_{r,A}$ denote its estimator using method A. We draw $t_{r,j} \sim \mathbb{P}_T,\ j = 1, \ldots, m$. To reduce distortion from ADRF extrapolation, we restrict evaluation grid to

$$\mathcal{T}_{\epsilon,r} = \{t_{r,j}\}_{j=1}^m \cap \{t : \epsilon < F_{T,r}(t) < 1 - \epsilon\}.$$

Table 1: ADRF estimation over the treatment distribution

	Linear			Nonlinear			NMES		
Method	RMSE	MAE	SUP	RMSE	MAE	SUP	RMSE	MAE	SUP
Equal	0.1655 (0.0074)	0.1299 (0.0056)	0.3668 (0.0143)	0.2037 (0.0081)	0.1587 (0.0075)	0.4650 (0.0188)	0.3553 (0.0118)	0.2785 (0.0103)	0.6983 (0.0237)
GPS	0.1786 (0.0290)	0.1199 (0.0062)	0.3850 (0.0440)	0.1694 (0.0062)	0.1307 (0.0052)	0.4201 (0.0173)	0.4028 (0.0192)	0.3061 (0.0130)	0.8738 (0.0572)
npCBGPS	0.1466 (0.0050)	0.1162 (0.0041)	0.3431 (0.0110)	0.1986 (0.0094)	0.1528 (0.0080)	0.4326 (0.0186)	0.3963 (0.0151)	0.3158 (0.0137)	0.7638 (0.0283)
EBCT	0.1445 (0.0050)	0.1149 (0.0042)	0.3341 (0.0101)	0.1921 (0.0090)	0.1472 (0.0077)	0.4281 (0.0186)	0.3700 (0.0137)	0.2940 (0.0125)	0.7071 (0.0251)
DCOW	0.1414 (0.0048)	0.1112 (0.0038)	0.3366 (0.0116)	0.1821 (0.0080)	0.1399 (0.0064)	0.4259 (0.0173)	0.3467 (0.0169)	0.2636 (0.0128)	0.7348 (0.0352)
CBSupIPM	**0.1314** (0.0043)	**0.1040** (0.0032)	**0.3138** (0.0107)	**0.1630** (0.0078)	**0.1259** (0.0061)	**0.3853** (0.0162)	**0.2960** (0.0120)	**0.2335** (0.0102)	**0.5851** (0.0226)
Oracle	0.1634 (0.0165)	0.1189 (0.0052)	0.3699 (0.0325)	0.1531 (0.0049)	0.1193 (0.0040)	0.3668 (0.0116)	–	–	–

We compare the ADRF estimation performance of the weighting methods using the following metrics.

$$\mathrm{RMSE}_A = \left[\frac{1}{R} \sum_{r=1}^{R} \frac{1}{|\mathcal{T}_{\epsilon,r}|} \sum_{t \in \mathcal{T}_{\epsilon,r}} \{\widehat{\mu}_{r,A}(t) - \mu_r(t)\}^2 \right]^{1/2},$$

$$\mathrm{MAE}_A = \frac{1}{R} \sum_{r=1}^{R} \frac{1}{|\mathcal{T}_{\epsilon,r}|} \sum_{t \in \mathcal{T}_{\epsilon,r}} |\widehat{\mu}_{r,A}(t) - \mu_r(t)|,$$

$$\mathrm{SUP}_A = \frac{1}{R} \sum_{r=1}^{R} \left\{ \sup_{t \in \mathcal{T}_{\epsilon,r}} |\widehat{\mu}_{r,A}(t) - \mu_r(t)| \right\},$$

Here, we use $R = 50$, $m = 2000$, and $\epsilon = 0.05$, and $F_{T,r}$ is the treatment CDF (empirical CDF for semi-synthetic data). The results are presented in Table 1.

Standard errors are in parentheses. The smallest mean in each column is bold. Equal uses the sample weights $\widehat{w}_i^{\mathrm{Equal}} = 1/n$ and Oracle uses the sample weights

$$\widehat{w}_i^{\mathrm{Oracle}} = \frac{W^*(\boldsymbol{X}_i, T_i)}{\sum_j W^*(\boldsymbol{X}_j, T_j)}$$

for $W^*(\boldsymbol{x}, t) = \dfrac{d\mathbb{P}_T}{d\mathbb{P}_{T|\boldsymbol{X}=\boldsymbol{x}}}(t)$ Oracle is unavailable for NMES because its treatment assignment density is unknown. Results use 50 replications. Smaller values are better. Parentheses are Monte Carlo standard errors. Bold marks the smallest value among completed non-oracle methods.

In the linear setting, CBSupIPM outperforms oracle weighting. A related finding for binary treatments is that using a nonparametric estimate of the propensity score can improve asymptotic efficiency relative to using the true propensity score (Hirano et al., 2003).

Estimation of the maximum ADRF value and the maximizing treatment. For replication r, let

$$M_r = \sup_{t \in \mathcal{T}} \mu_r(t), \qquad t_r^\star \in \arg\max_{t \in \mathcal{T}} \mu_r(t)$$

denote the maximum ADRF value and a maximizing treatment over $\mathcal{T}$, respectively. We partition the sample into K disjoint folds $I_{r,1}, \ldots, I_{r,K}$. For method A, let $\widehat{\mu}_{r,A}^{(k)}(t)$ be the weighted NW estimator computed using observations in $I_{r,k}$. Each fold selects a candidate treatment on the common evaluation grid $\mathcal{T}_{\epsilon,r}$:

$$\widehat{t}_{r,A}^{(k)} \in \arg\max_{t \in \mathcal{T}_{\epsilon,r}} \widehat{\mu}_{r,A}^{(k)}(t), \qquad k = 1, \ldots, K. \tag{20}$$

Table 2: Estimation of the maximum ADRF and its maximizer

	Linear (M)		Nonlinear (M)		Linear (t)		Nonlinear (t)	
Method	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
Equal	0.2530 (0.0201)	0.2181 (0.0183)	0.2601 (0.0229)	0.2134 (0.0213)	0.2044 (0.0285)	0.1484 (0.0201)	0.2499 (0.0243)	0.2024 (0.0209)
GPS	0.2544 (0.0244)	0.2100 (0.0205)	0.2539 (0.0223)	0.2169 (0.0189)	0.1612 (0.0166)	0.1215 (0.0151)	0.2070 (0.0249)	0.1593 (0.0189)
npCBGPS	0.2466 (0.0214)	0.2095 (0.0186)	0.2529 (0.0244)	0.1997 (0.0222)	0.1720 (0.0161)	0.1341 (0.0154)	0.3017 (0.0402)	0.2219 (0.0292)
EBCT	0.2443 (0.0213)	0.2059 (0.0188)	0.2407 (0.0225)	0.1942 (0.0203)	0.1613 (0.0162)	0.1244 (0.0147)	0.2436 (0.0259)	0.1871 (0.0223)
DCOW	0.2455 (0.0273)	0.2011 (0.0201)	0.2607 (0.0253)	0.2106 (0.0220)	0.1573 (0.0125)	0.1316 (0.0123)	0.2269 (0.0238)	0.1780 (0.0201)
CBSupIPM	**0.2045** (0.0208)	**0.1720** (0.0158)	**0.2011** (0.0188)	**0.1688** (0.0156)	**0.1433** (0.0175)	**0.1044** (0.0140)	**0.1926** (0.0194)	**0.1528** (0.0167)
Oracle	0.2501 (0.0221)	0.2096 (0.0195)	0.2381 (0.0227)	0.2001 (0.0184)	0.1920 (0.0235)	0.1422 (0.0184)	0.2127 (0.0270)	0.1598 (0.0201)

For each evaluation fold v, we evaluate the $K-1$ candidates selected by the other folds and retain the largest estimated value. Averaging these values across the K evaluation folds gives

$$\widehat{M}_{r,A} = \frac{1}{K}\sum_{v=1}^{K} \max_{k\in\{1,\dots,K\}\setminus\{v\}} \widehat{\mu}_{r,A}^{(v)}\left(\widehat{t}_{r,A}^{(k)}\right).$$

We estimate $t_r^\star$ by $\bar{t}_{r,A} = \dfrac{1}{K}\sum_{k=1}^{K}\widehat{t}_{r,A}^{(k)}$.

We compare the performance of the weighting methods in estimating the maximum ADRF value and the maximizing treatment using the following metrics.

$$\mathrm{RMSE}_A^{(M)} = \left[\frac{1}{R}\sum_{r=1}^{R}\left\{\widehat{M}_{r,A} - M_r\right\}^2\right]^{1/2},$$

$$\mathrm{MAE}_A^{(M)} = \frac{1}{R}\sum_{r=1}^{R}\left|\widehat{M}_{r,A} - M_r\right|,$$

$$\mathrm{RMSE}_A^{(t)} = \left[\frac{1}{R}\sum_{r=1}^{R}\left\{\bar{t}_{r,A} - t_r^\star\right\}^2\right]^{1/2},$$

$$\mathrm{MAE}_A^{(t)} = \frac{1}{R}\sum_{r=1}^{R}\left|\bar{t}_{r,A} - t_r^\star\right|.$$

The results are presented in Table 2.

Also, we examine how closely each method estimates the maximum value and its maximizer of ADRF. Figure 1 presents the results over 50 replications in the linear setting, respectively.

Each panel shows 50 points, one for each replication. The point for replication r and method A has coordinate $\left(\bar{t}_{r,A}, \widehat{M}_{r,A}\right)$. The red curve is the true ADRF, and the red point marks $(t_r^\star, M_r)$ which is the fixed point over r. All methods tend to overestimate the maximum ADRF value, but CBSupIPM shows the smallest average overestimation. Figure for the nonlinear setting is presented in Figure 2.

6 Discussion

Contribution. We connect covariate balancing to ADRF estimation through an error decomposition that motivates the CBSupIPM objective. We establish uniform convergence rates that also

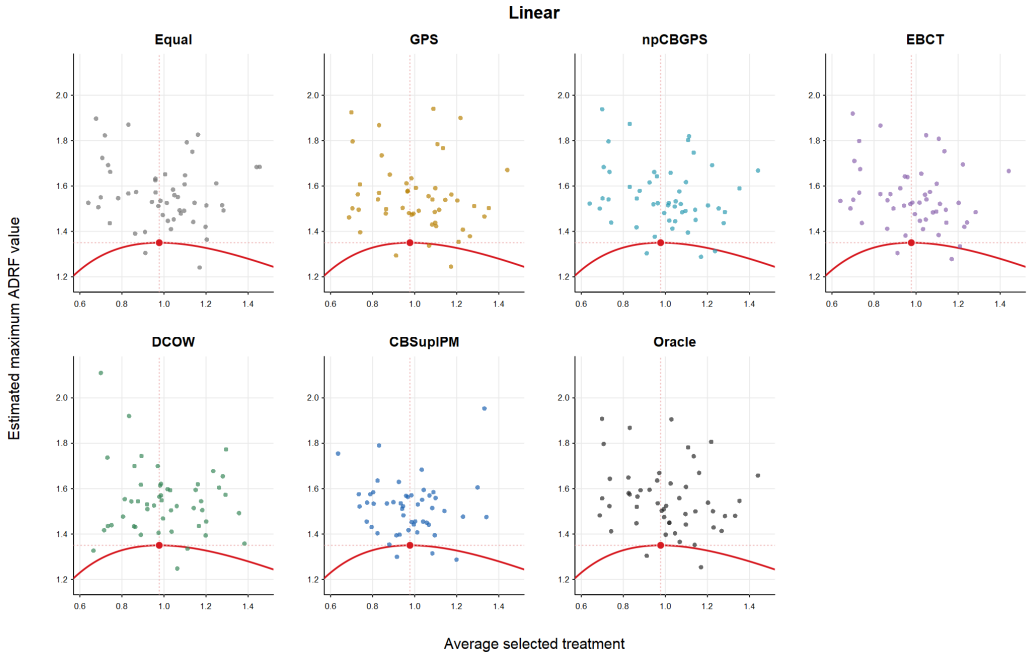


Figure 1: Maximum ADRF value estimation in the linear setting

control estimation of the maximum ADRF value. Our numerical results demonstrate the effectiveness of CBSupIPM for both estimation tasks.

Limitation. Although unconfoundedness and overlap are widely used, they are strong assumptions for causal identification and are difficult to verify from observed data alone. Training and evaluation can also require substantial computation time. Our asymptotic analysis is limited to the weighted NW estimator. Establishing corresponding asymptotic results for CBSupIPM combined with local linear estimation or doubly robust methods is an important direction for future research.

AI assistance. We used large language models(LLMs) to assist with English writing and editing, identify relevant literature, and provide feedback on the manuscript. We read and revised all text generated by LLMs. We verified the references and information through searches. We read the feedback and revised the manuscript ourselves. We also used these tools to help develop experimental code and interpret experimental results. We also used these tools to format manuscript content, including tables and figures, in LaTeX. We checked the formatting by inspecting the compiled output. We take responsibility for the final content, including the code, numerical results, and interpretations.

REFERENCES

Jason Bryer. *TriMatch: Propensity Score Matching of Non-Binary Treatments*, 2025. URL `https://jbryer.github.io/TriMatch/`. R package version 1.0.1.

Hengrui Cai, Chengchun Shi, Rui Song, and Wenbin Lu. Jump interval-learning for individualized decision making with continuous treatments. *Journal of Machine Learning Research*, 24(140): 1–92, 2023.

Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Gaussian approximation of suprema of empirical processes. *The Annals of Statistics*, 42(4):1564–1597, 2014.

Christian Fong, Chad Hazlett, and Kosuke Imai. Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *The Annals of Applied Statistics*, 12(1):156–177, 2018. doi: 10.1214/17-AOAS1101.

Keisuke Hirano and Guido W. Imbens. The propensity score with continuous treatments. In Andrew Gelman and Xiao-Li Meng (eds.), *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives: An Essential Journey with Donald Rubin's Statistical Family*, Wiley Series in Probability and Statistics, chapter 7, pp. 73–84. John Wiley & Sons, Ltd., 2004. doi: 10.1002/0470090456.ch7.

Keisuke Hirano, Guido W. Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003. doi: 10.1111/1468-0262.00442.

Jared D. Huling, Noah Greifer, and Guanhua Chen. Independence weights for causal inference with continuous treatments. *Journal of the American Statistical Association*, 119(546):1657–1670, 2024. doi: 10.1080/01621459.2023.2213485.

Kosuke Imai and David A Van Dyk. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467):854–866, 2004.

Nathan Kallus and Michele Santacatterina. Kernel optimal orthogonality weighting: A balancing approach to estimating effects of continuous treatments. *arXiv preprint arXiv:1910.11972*, 2019.

Joseph D. Y. Kang and Joseph L. Schafer. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science*, 22(4):523–539, 2007. doi: 10.1214/07-STS227.

Fan Li, Kari Lock Morgan, and Alan M Zaslavsky. Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, 113(521):390–400, 2018.

Guillaume Martinet. A balancing weight framework for estimating the causal effect of general treatments. *arXiv preprint arXiv:2002.11276*, 2020.

Hyung Park, Eva Petkova, Thaddeus Tarpey, and R. Todd Ogden. A single-index model with a surface-link for optimizing individualized dose rules. *Journal of Computational and Graphical Statistics*, 31(2):553–562, 2022. doi: 10.1080/10618600.2021.1923521.

Yuha Park, Kunwoong Kim, Insung Kong, and Yongdai Kim. Relu integral probability metric and its applications. *arXiv preprint arXiv:2504.18897*, 2025.

James M. Robins, Miguel A. Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000. doi: 10.1097/00001648-200009000-00011.

Johannes Schmidt-Hieber. Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.

Patrick Schwab, Lorenz Linhardt, Stefan Bauer, Joachim M. Buhmann, and Walter Karlen. Learning counterfactual representations for estimating individual dose-response curves. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(4):5612–5619, 2020. doi: 10.1609/aaai.v34i04.6014.

Rahul Singh, Liyuan Xu, and Arthur Gretton. Kernel methods for causal functions: Dose, heterogeneous and incremental response curves. *Biometrika*, 111(2):497–516, 2024. doi: 10.1093/biomet/asad042.

Stefan Tübbicke. Entropy balancing for continuous treatments. *Journal of Econometric Methods*, 11(1):71–89, 2022. doi: 10.1515/jem-2021-0002.

Aad W Van Der Vaart and Jon A Wellner. Weak convergence. In *Weak convergence and empirical processes: with applications to statistics*, pp. 16–28. Springer, 1996.

Brian G. Vegetabile, Beth Ann Griffin, Donna L. Coffman, Matthew Cefalu, Michael W. Robbins, and Daniel F. McCaffrey. Nonparametric estimation of population average dose-response curves using entropy balancing weights for continuous exposures. *Health Services and Outcomes Research Methodology*, 21(1):69–110, 2021. doi: 10.1007/s10742-020-00236-2.