# Minimax Optimal Nonparametric Estimation of Heterogeneous Treatment Effects 

Zijun Gao<br>Department of Statistics<br>Stanford University<br>Email: zijungao@stanford.edu

Yanjun Han<br>Department of Electrical Engineering<br>Stanford Univeristy<br>Email: yjhan@stanford.edu


#### Abstract

A central goal of causal inference is to detect and estimate the treatment effects of a given treatment or intervention on an outcome variable of interest, where a member known as the heterogeneous treatment effect (HTE) is of growing popularity in recent practical applications such as the personalized medicine. In this paper, we model the HTE as a smooth nonparametric difference between two less smooth baseline functions, and determine the tight statistical limits of the nonparametric HTE estimation as a function of the covariate geometry. In particular, a two-stage nearest-neighbor-based estimator throwing away observations with poor matching quality is near minimax optimal. We also establish the tight dependence on the density ratio without the usual assumption that the covariate densities are bounded away from zero, where a key step is to employ a novel maximal inequality which could be of independent interest.


## 1 Introduction and Main Results

Causal inference aims to draw a causal relationship between some treatment and target responses. Nowadays, personalized medicine and huge available data make heterogeneous treatment effect (HTE) estimation meaningful and possible. While there are various practical approaches of estimating the HTE [AI16, $\mathrm{PQJ}^{+} 17, \mathrm{WA18}, \mathrm{KSBY19]}$, some important theoretical questions remain unanswered.

In this paper, we consider the Neyman-Rubin potential outcome model [Rub74] for the treatment effect. Assume for simplicity that there are $n$ individuals in the treatment group and the control group, respectively, where the generalizations to different group sizes are straightforward. For each individual $i \in[n]$ in the control group, we observe a vector of covariates $X_{i}^{0} \in \mathbb{R}^{d}$ and her potential outcome $Y_{i}^{0} \in \mathbb{R}$ for not receiving the treatment. Similarly, for individual $i$ in the treatment group, the covariates $X_{i}^{1}$ and the potential outcome $Y_{i}^{1}$ under the treatment are observed. We assume the following model for the potential outcomes: for any $i \in[n]$,

$$
\begin{equation*}
Y_{i}^{0}=\mu_{0}\left(X_{i}^{0}\right)+\varepsilon_{i}^{0}, \quad Y_{i}^{1}=\mu_{1}\left(X_{i}^{1}\right)+\varepsilon_{i}^{1} \tag{1}
\end{equation*}
$$

where $\mu_{0}, \mu_{1}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ are the baseline functions for the control and treatment groups, respectively, and $\varepsilon_{i}^{0}, \varepsilon_{i}^{1}$ are modeling errors. The heterogeneous treatment effect $\tau$ is defined to be the difference of the baseline functions:

$$
\begin{equation*}
\tau(x) \triangleq \mu_{1}(x)-\mu_{0}(x) \tag{2}
\end{equation*}
$$

In other words, the treatment effect $\tau(x)$ is the expected change in the outcomes after an individual with covariate $x$ receives the treatment, which is usually heterogeneous as $\tau(x)$ typically varies in $x$. The target is to find an estimator $\hat{\tau}$ which comes close to the true HTE $\tau$ under the $L_{2}$ norm of functions, based on the control and treatment observations $\left\{\left(X_{i}^{0}, Y_{i}^{0}\right)\right\}_{i \in[n]},\left\{\left(X_{i}^{1}, Y_{i}^{1}\right)\right\}_{i \in[n]}$.

In HTE estimation, modeling of the baseline functions $\mu_{0}, \mu_{1}$ or the HTE $\tau$ plays an important role. In practice, the treatment effect $\tau$ is typically easier to estimate than the baseline functions $\mu_{0}, \mu_{1}$, as ideally the treatment effect depends solely on the single treatment. In this paper, we assume that both the baseline and treatment effect functions are nonparametric functions, with an additional constraint that $\tau$ is smoother than $\mu_{0}$ :
Assumption 1 (Baseline and HTE functions). The baseline $\mu_{0}$ and the treatment effect $\tau$ belong to $d$-dimensional Hölder balls with smoothness parameters $\beta_{\mu} \leq \beta_{\tau}$, respectively.

Recall the following definition of Hölder balls.
Definition 1 (Hölder ball). The Hölder ball $\mathcal{H}_{d}(\beta)$ with dimension $d$ and smoothness parameter $\beta \geq 0$ is the collection of functions $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ supported on $[0,1]^{d}$ with

$$
\left|\frac{\partial^{s} f}{\partial x_{1}^{\beta_{1}} \cdots \partial x_{d}^{\beta_{d}}}(x)-\frac{\partial^{s} f}{\partial x_{1}^{\beta_{1}} \cdots \partial x_{d}^{\beta_{d}}}(y)\right| \leq L\|x-y\|_{2}^{\alpha}
$$

for all $x, y \in \mathbb{R}^{d}$ and all multi-indices $\left(\beta_{1}, \cdots, \beta_{d}\right) \in \mathbb{N}^{d}$ with $\sum_{i=1}^{d} \beta_{i}=s$, where $\beta=s+\alpha$ with $s \in \mathbb{N}, \alpha \in(0,1]$. Throughout we assume that the radius $L$ is a fixed positive constant and omit the dependence on $L$.

Assumption 1 imposes no structural assumptions on $\mu_{0}$ and $\tau$ except for the smoothness, and assumes that the $\operatorname{HTE} \tau$ is smoother and thus easier to estimate than the baseline $\mu_{0}$. This typically holds as intuitively fewer factors contribute to the HTE than baselines, and semiparametric models usually assume simple forms of HTEs such as constant or linear [Ken16, CCD $^{+} 18$ ], where Assumption 1 could be treated as a nonparametric counterpart. Standard results in nonparametric estimation reveal that if one naïvely estimates the baselines $\mu_{0}$ and $\mu_{1}$ separately based on (1), then each function can be estimated within accuracy $\Theta\left(n^{-\beta_{\mu} /\left(2 \beta_{\mu}+d\right)}\right)$, and so is the difference $\tau$ in (2). However, if the covariates in the control and treatment groups match perfectly, i.e. $X_{i}^{0}=X_{i}^{1}$ for all $i \in[n]$, then differencing (1) gives an improved estimation accuracy $\Theta\left(n^{-\beta_{\tau} /\left(2 \beta_{\tau}+d\right)}\right)$ for the HTE $\tau$. In general, the estimation performance of HTE is an interpolation between the above extremes and depends heavily on the quality of the covariate matching. To model such qualities, we have the following assumption on the covariates $X_{i}^{0}, X_{i}^{1}$.

Assumption 2 (Covariates). The covariates are generated under fixed or random designs below:

- Fixed design: the covariates are generated from the following fixed grid

$$
\left\{X_{i}^{0}\right\}_{i \in[n]}=\left\{0, \frac{1}{m}, \cdots, \frac{m-1}{m}\right\}^{d}, \quad\left\{X_{i}^{1}\right\}_{i \in[n]}=\left\{0, \frac{1}{m}, \cdots, \frac{m-1}{m}\right\}^{d}+\Delta,
$$

for some vector $\Delta \in \mathbb{R}^{d}$ with $\|\Delta\|_{\infty} \leq 1 /(2 m)$, and $m \triangleq n^{1 / d}$ is assumed to be an integer.

- Random design: the covariates are i.i.d. sampled from unknown densities $g_{0}, g_{1}$ on $[0,1]^{d}$ :

$$
X_{1}^{0}, X_{2}^{0}, \cdots, X_{n}^{0} \stackrel{\text { i.i.d. }}{\sim} g_{0}, \quad X_{1}^{1}, X_{2}^{1}, \cdots, X_{n}^{1} \stackrel{\text { i.i.d. }}{\sim} g_{1},
$$

where there is a bounded likelihood ratio $\kappa^{-1} \leq g_{0}(x) / g_{1}(x) \leq \kappa$ on the densities.
Under the fixed design, the covariates in both groups are evenly spaced grids in $[0,1]^{d}$, with a shift $\Delta$ quantifying the matching distance between the control and treatment groups. The fixed design is not very practical, but the analysis will provide important insights for the HTE estimation. The random design model is more meaningful and realistic without any matching parameter, and the density ratio $g_{0} /\left(g_{0}+g_{1}\right)$ corresponds to the propensity score, a key quantity in causal inference. We assume that $\kappa^{-1} \leq g_{0} / g_{1} \leq \kappa$ everywhere, as it is usually necessary to have a propensity score bounded away from 0 and 1 . Besides the bounded likelihood ratio, we remark that it is not assumed that the densities $g_{0}, g_{1}$ are bounded away from zero.
Finally it remains to model the noises $\varepsilon_{i}^{0}, \varepsilon_{i}^{1}$, and we assume the following mild conditions.
Assumption 3 (Noise). The noises $\varepsilon_{i}^{0}, \varepsilon_{i}^{1}$ are mutually independent, zero-mean and of variance $\sigma^{2}$.

Based on the above assumptions, the target of this paper is to characterize the minimax risks of the nonparametric HTE estimation under both fixed and random designs. Specifically, we are interested in the following minimax $L_{1}$ risk

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\mathrm{fixed}}(\Delta) \triangleq \inf _{\hat{\tau}} \sup _{\mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right), \tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right)} \mathbb{E}_{\mu_{0}, \tau}\left[\|\hat{\tau}-\tau\|_{1}\right]
$$

for fixed design with matching parameter $\Delta \in \mathbb{R}^{d}$, where the infimum is taken over all possible estimators $\hat{\tau}$ based on the observations $\left\{\left(X_{i}^{0}, Y_{i}^{0}\right)\right\}_{i \in[n]},\left\{\left(X_{i}^{1}, Y_{i}^{1}\right)\right\}_{i \in[n]}$. As for the minimax risk for random designs, we include the dependence on the density ratio $\kappa$ and use the $L_{1}$ norm based on the density $g_{0}$ (as we do not assume that the densities are bounded away from zero):

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\mathrm{random}}(\kappa) \triangleq \inf _{\hat{\tau}} \sup _{\mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right), \tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right), g_{0}, g_{1}} \mathbb{E}_{\mu_{0}, \tau}\left[\|\hat{\tau}-\tau\|_{L_{1}\left(g_{0}\right)}\right]
$$

The main aim of this paper is to characterize the tight minimax rates for the above quantities, and in particular, how they are determined by the covariate geometry and the full set of parameters $(n, \kappa, \sigma)$ of interest. In addition, we aim to extract useful practical insights (instead of fully practical algorithms) based on the minimax optimal estimation procedures established in theory. Moreover, we mostly focus on the special nature of the HTE estimation problem instead of general nonparametric estimation, and therefore we elaborate less on broad issues of nonparametric statistics such as adaptation/hyperparameter estimation and refer interested readers to known literature (see Section 5).

Our first result is the characterization of the minimax rate for HTE estimation under fixed designs.
Theorem 1 (Fixed Design). Under Assumptions $1-3$

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\mathrm{fixed}}(\Delta) \asymp n^{-\frac{\beta_{\mu}}{d}}\left(n^{\frac{1}{d}}\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}+\left(\frac{\sigma^{2}}{n}\right)^{\frac{\beta_{\tau}}{2 \beta_{\tau}+d}}
$$

Theorem 1 shows that, as the covariate matching quality improves (i.e. $\|\Delta\|_{\infty}$ shrinks), the estimation error of the HTE decreases from $n^{-\beta_{\mu} / d}$, which is slightly better than the estimation error $n^{-\beta_{\mu} /\left(2 \beta_{\mu}+d\right)}$ for the baselines, to the optimal estimation error $\left(\sigma^{2} / n\right)^{\beta_{\tau} /\left(2 \beta_{\tau}+d\right)}$ when the learner has $n$ direct samples from $\tau$. We also remark that the covariate matching quality is determined by the $\ell_{\infty}$ norm of $\Delta$, and the matching bias does not depend on the noise level $\sigma$.
The minimax rate for HTE estimation under random designs exhibits more interesting behaviors.
Theorem 2 (Random Design). Under Assumptions 1 • 3 if $\beta_{\tau} \leq 1$ and $\kappa \leq n$, then

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\mathrm{random}}(\kappa)=\widetilde{\Theta}\left(\left(\frac{\kappa}{n^{2}}\right)^{\frac{1}{d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}}+\left(\frac{\kappa \sigma^{2}}{n^{2}}\right)^{\frac{1}{2+d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}}+\left(\frac{\kappa \sigma^{2}}{n}\right)^{\frac{\beta_{\tau}}{2 \beta_{\tau}+d}}\right)
$$

Theorem 2 provides the first minimax analysis of the HTE estimation with tight dependence on all parameters $(n, \kappa, \sigma)$ without assuming densities bounded away from zero. Interestingly, Theorem 2 shows that there are three regimes of the minimax rate of the HTE estimation under random designs, and in particular, there is an intermediate regime where neither the matching bias nor the estimation of $\tau$ dominates. Moreover, when it comes to the dependence on the density ratio $\kappa$, the effective sample size is $n / \sqrt{\kappa}$ in the first two regimes, while it becomes $n / \kappa$ in the last regime.

### 1.1 Related work

The nonparametric regression of a single function has a long history and is a well-studied topic in nonparametric statistics, with various estimators including kernel smoothing [Ros56], local polynomial fitting [Sto77, FGG $^{+}$97], splines [DBDBM ${ }^{+}$78, GS93], wavelets [DJKP95], nearest neighbors [Fix51, Sto77], and so on. We refer to the excellent books [Nem00, GKKW06, Tsy09, Joh11, BD15] for an overview. However, estimating the difference of two nonparametric functions remains largely underexplored.

In the history of causal inference, the average treatment effect (ATE) has long been the focus of research [IR15]. Recently, an increasing number of works seek to estimate the HTE, and there are two major threads of approaches: parametric (semi-parametric) estimation and nonparametric estimation.

In parametric estimation, a variety of parametric assumptions such as linearity are imposed on the baseline functions and HTE. The problem then reduces to classic parametric estimation and many methods are applicable, e.g. [TAGT14, $\left.\mathrm{IR}^{+} 13\right]$. In this paradigm, classic parametric rate is obtained under proper regularity assumptions. To further deal with observational study and bias due to the curse of high-dimensionality, $\left[\mathrm{CCD}^{+} 18\right]$ develops a double/debiased procedure allowing a less accurate estimation of the mean function without sacrificing the parametric rate.

In nonparametric estimation, there are a large number of practical approaches proposed, such as splinebased method [Che07], matching-based method [XBJ12] and tree-based method [AI16, PQJ ${ }^{+}$17, WA18]. However, relatively fewer work study the statistical limit of nonparametric estimation of HTE. One work [NW17] argued that a crude estimation of baselines suffices for the optimal estimation of HTE, while their approach did not take into account the impact of the covariate geometry. Some works AS18, KSBY19] studied the minimax risk with smoothness assumptions on the baseline functions, but they did not directly model the smoother HTE function and thus arrived at a relatively easier claim that the difficulty is determined by the less smooth baseline. In contrast, we assume a smoother HTE and provide additional insights on how to construct pseudo-observations and discarding poor-quality data based on the covariate geometry. Hence, to the best of our knowledge, this paper establishes the first minimax rates with tight dependence on both the covariate geometry, and the parameters $\kappa$ related to propensity scores as well as $\sigma$ related to noise levels.

We also comment that in nonparametric regression problems with random design, a uniform lower bound on the density is usually necessary in most past works [BCLZ02, GKKW06]. To overcome this difficulty, there is a recent line of research showing that to estimate nonparametric functionals, Hardy-Littlewood-type maximal inequalities are often useful to deal with densities close to zero for both kernel methods [HJWW17] and nearest neighbors [JGH18]. Therefore, this work is a continuation of the above thread with a novel maximal inequality on another different problem.

### 1.2 Notations

Let $\mathbb{R}, \mathbb{N}$ be the set of all real numbers and non-negative integers, respectively. For $p \in[1, \infty]$, we denote by $\|\cdot\|_{p}$ both the $\ell_{p}$ norm for vectors and the $L_{p}$ norm for functions. Let $[n] \triangleq$ $\{1,2, \cdots, n\}$, and $a \wedge b \triangleq \min \{a, b\}$. We call $K: \mathbb{R}^{d} \rightarrow \mathbb{R}$ a kernel of order $k$ if $\int K(x) d x=1$ and $\int\left(v^{\top} x\right)^{\ell} K(x) d x=0$ for any $v \in \mathbb{R}^{d}$ and $1 \leq \ell \leq k$. For non-negative sequences $\left\{a_{n}\right\}$ and $\left\{b_{n}\right\}$, we write $a_{n} \asymp b_{n}$, or $a_{n}=\Theta\left(b_{n}\right)$, to denote that $c b_{n} \leq a_{n} \leq C b_{n}$ for some absolute constants $c, C>0$ independent of $n$. The notation $a_{n}=\widetilde{\Theta}\left(b_{n}\right)$ means that the above inequality holds within multiplicative polylogarithmic factors in $n$.

### 1.3 Organization

In Section 2, we construct the minimax optimal HTE estimator via a combination of kernel methods and covariate matching under the simple fixed design setting. For the random design, a two-stage nearest-neighbor based algorithm is proposed in Section3 to trade off two types of biases and the variance. The efficacy of the proposed estimator is shown via numerical experiments in Section 4 Some limitations and future works are discussed in Section 5, and all proofs are relegated to the Appendices.

## 2 Fixed Design

This section is devoted to the minimax rate of HTE estimation under fixed design. To construct the estimator, we first form pseudo-observations of the outcomes in the treatment group on the covariates in the control group based on covariate matching, and then apply the classic kernel estimator based on perfectly matched pseudo-observations. We also sketch the ideas of the minimax lower bound and show that both the matching bias and the estimation error are inevitable.

### 2.1 Estimator Construction

Recall that in the perfect matching case, i.e. $\Delta=0$, the HTE function $\tau$ can be estimated by classic nonparametric estimators after taking the outcome difference $Y_{i}^{1}-Y_{i}^{0}$ in (1). This basic idea will be

```
Algorithm 1 Estimator Construction under Fixed Design
    Input: control observations \(\left\{\left(X_{i}^{0}, Y_{i}^{0}\right)\right\}_{i \in[n]}\), treatment observations \(\left\{\left(X_{i}^{1}, Y_{i}^{1}\right)\right\}_{i \in[n]}\), dimension-
    ality \(d\), kernel \(K\) of order \(\left\lfloor\beta_{\tau}\right\rfloor\), bandwidth \(h>0\), query \(x_{0} \in[0,1]^{d}\).
    Output: HTE estimator \(\hat{\tau}\left(x_{0}\right)\) at point \(x_{0}\).
    for \(x \in\left\{X_{i}^{0}\right\}_{i \in[n]}\) do
        Choose \(t=\left\lfloor\beta_{\mu}\right\rfloor+1\), and find neighboring treatment covariates \(\left\{x_{i}\right\}_{i \in[t]^{d}}\);
        Compute the weights \(\left\{w_{i}\right\}_{i \in[t]^{d}}\) according to (3);
        Compute the pseudo-observation \(\hat{Y}^{1}(x)\) according to 4).
    end for
    Output the final estimator \(\hat{\tau}\left(x_{0}\right)\) according to (5).
```

used for general fixed designs: first, for each $i \in[n]$, we form pseudo-observations $\hat{Y}_{i}{ }^{1}$ with target mean $\mu_{1}\left(X_{i}^{0}\right)$, where the observed covariate $X_{i}^{1}$ in the treatment group moves to the nearest covariate $X_{i}^{0}$ in the control group. Next, we apply the kernel estimator in the perfect matching case to the pseudo-difference $\hat{Y}_{i}^{1}-Y_{i}^{0}$. The covariate matching step also makes use of a suitable kernel method, where the matching performance depends on the smoothness of the baseline function $\mu_{0}$.
Specifically, the pseudo-observations $\hat{Y}_{i}^{1}$ are constructed as follows. Let $t \triangleq\left\lfloor\beta_{\mu}\right\rfloor+1$, and fix a covariate $x$ in the set of control covariates. For each $j \in[d]$, let $x_{1, j}, \cdots, x_{t, j}$ be the $t$ closest grid points in $\left\{\Delta_{j}, 1 / m+\Delta_{j}, \cdots 1-1 / m+\Delta_{j}\right\}$ to $x_{j}$, the $j$-th coordinate of $x$. Moreover, for each coordinate $j \in[d]$, we also compute the following weights $w_{1, j}, \cdots, w_{t, j}$ such that

$$
\begin{equation*}
\sum_{i=1}^{t} w_{i, j}\left(x_{i, j}-x_{j}\right)^{\ell}=\mathbb{1}(\ell=0), \quad \forall 0 \leq \ell \leq\left\lfloor\beta_{\mu}\right\rfloor \tag{3}
\end{equation*}
$$

To form the pseudo-observation of the treatment outcome $\hat{Y}^{1}(x)$ on the covariate $x$, we use the treatment observations $Y^{1}\left(x_{i}\right)$ for all multi-indices $i=\left(i_{1}, \cdots, i_{d}\right) \in[t]^{d}$ and $x_{i} \triangleq\left(x_{i, 1}, x_{i, 2}, \cdots, x_{i, d}\right)$, with $w_{i}=w_{i, j}$ being the weight of the covariate $x_{i}$. Note that each $x_{i}$ belongs to the covariate grid of the treatment group and therefore $Y^{1}\left(x_{i}\right)$ is observable. Finally, given the above covariates $x_{i}$ and weights $w_{i}$, the pseudo-observation of the treatment outcome at the control covariate $x$ is

$$
\begin{equation*}
\hat{Y}^{1}(x)=\sum_{i \in[t]^{d}} w_{i} Y^{1}\left(x_{i}\right) \tag{4}
\end{equation*}
$$

In approximation theory, this is an interpolation estimate of $\mu_{1}(x)$ based on the function values $\left\{\mu_{1}\left(x_{i}\right)\right\}_{i \in[t]^{d}}$ in the neighborhood of $x$.
Given the pseudo-observations, the final HTE estimator $\hat{\tau}\left(x_{0}\right)$ for any $x_{0} \in[0,1]^{d}$ is defined as

$$
\begin{equation*}
\hat{\tau}\left(x_{0}\right) \triangleq \frac{\sum_{i=1}^{n} K\left(\left(X_{i}^{0}-x_{0}\right) / h\right)\left(\hat{Y}^{1}\left(X_{i}^{0}\right)-Y_{i}^{0}\right)}{\sum_{i=1}^{n} K\left(\left(X_{i}^{0}-x_{0}\right) / h\right)} \tag{5}
\end{equation*}
$$

i.e. the Nadaraya-Watson estimator [Nad64, Wat64] applied to the pseudo-differences, where $K$ is any kernel of order $\left\lfloor\beta_{\tau}\right\rfloor$, and $h \in(0,1)$ is a suitable bandwidth. The complete description of the estimator construction is displayed in Algorithm 1.

The performance of the above estimator is summarized in the following theorem.
Theorem 3. There exists a constant $C>0$ independent of $(n, h)$ such that

$$
\sup _{\mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right), \tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right)} \mathbb{E}_{\mu_{0}, \tau}\left[\|\hat{\tau}-\tau\|_{1}\right] \leq C\left(n^{-\frac{\beta_{\mu}}{d}}\left(n^{\frac{1}{d}}\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}+h^{\beta_{\tau}}+\frac{\sigma}{\sqrt{n h^{d}}}\right) .
$$

In particular, the estimator $\hat{\tau}$ in (5) achieves the upper bound in Theorem 1 for $h=\Theta\left(n^{-1 /\left(2 \beta_{\tau}+d\right)}\right)$.
We sketch the proof idea of Theorem 3 here. Using the definition of pseudo-observations $\hat{Y}^{1}$, for each $i \in[n]$ we have $\hat{Y}^{1}\left(X_{i}^{0}\right)-Y_{i}^{0}=\tau\left(X_{i}^{0}\right)+b_{\mu}\left(X_{i}^{0}\right)+\tilde{\varepsilon}_{i}$, where $\tau\left(X_{i}^{0}\right)$ is the target treatment effect at the point $X_{i}^{0}, b_{\mu}(x)=\sum_{i \in[t]^{d}} w_{i} \mu_{1}\left(x_{i}\right)-\mu_{1}(x)$ is the matching bias incurred by the linear interpolation, and $\tilde{\varepsilon}_{i}$ is a linear combination of the error terms. Then the first term of Theorem 3 is an upper bound of the matching bias, while the other terms come from the traditional bias-variance tradeoff in nonparametric estimation. We leave the full proof to the appendix.

### 2.2 Minimax Lower Bound

In this section, we show that the above HTE estimator is minimax optimal via the following minimax lower bound.

Theorem 4. There exists a constant $c>0$ independent of $n$ such that for any HTE estimator $\hat{\tau}$, it holds that under Gaussian noise,

$$
\sup _{\mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right), \tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right)} \mathbb{E}_{\mu_{0}, \tau}\left[\|\hat{\tau}-\tau\|_{1}\right] \geq c\left(n^{-\frac{\beta_{\mu}}{d}}\left(n^{\frac{1}{d}}\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}+\left(\frac{\sigma^{2}}{n}\right)^{\frac{\beta_{\tau}}{2 \beta_{\tau}+d}}\right)
$$

As for the proof of Theorem 4, note that the estimation error $\left(\sigma^{2} / n\right)^{\beta_{\tau} /\left(2 \beta_{\tau}+d\right)}$ is optimal in the nonparametric estimation of $\tau$ even if there is a perfect covariate matching. Hence, it remains to prove the first term, i.e. the matching bias. The proof is based on Le Cam's two-point method [Yu97] and a functional optimization problem. Consider two scenarios $\left(\mu_{0}, \tau\right)$ and $\left(\mu_{0}^{\prime}, \tau^{\prime}\right)$ with $\sigma=0$ and $\mu_{0}^{\prime}(x) \equiv 0, \tau^{\prime}(x) \equiv 0$, where $\left(\mu_{0}(x), \tau(x)\right)$ is the solution to the following optimization problem:

$$
\begin{array}{ll}
\operatorname{maximize} & \|\tau\|_{1} \\
\text { subject to } & \mu_{0}\left(X_{i}^{0}\right)=0, \quad \mu_{0}\left(X_{i}^{1}\right)+\tau\left(X_{i}^{1}\right)=0, \quad i=1, \cdots, n \\
& \mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right), \tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right) \tag{6}
\end{array}
$$

Note that when $\left(\mu_{0}, \tau\right)$ is a feasible solution to (6), under both scenarios $\left(\mu_{0}, \tau\right)$ and $\left(\mu_{0}^{\prime}, \tau^{\prime}\right)$ all outcomes are identically zero in both groups. Hence, these scenarios are completely indistinguishable, and it remains to show that the objective value of (6) is at least the first error term. We defer the complete proofs to the appendix.

## 3 Random Design

The random design assumption is more practical and complicated than the fixed design, as the nearest matching distances of different individuals are typically different without the regular grid structure. In particular, under the random design some covariates have better matching quality than others, in the sense that the minimum pairwise distance between the control and treatment covariates is $\Theta\left(n^{-2 / d}\right)$ while the typical distance is as large as $\Theta\left(n^{-1 / d}\right)$. In this section, we propose a two-stage nearest-neighbor estimator for HTE under the random design, and show that the minimax estimation error exhibits three different behaviors.

### 3.1 Estimator Construction

In the previous section, we construct pseudo-observations at each control covariate using the same treatment grid geometry in the neighborhood of that covariate, thanks to the fixed grid design assumption. However, the local geometries in the random design may differ spatially. For example, for $n$ i.i.d. uniformly distributed random variables on $[0,1]$, the typical spacing between spatially adjacent values is $\Theta\left(n^{-1}\right)$, while the minimum spacing is of the order $\Theta\left(n^{-2}\right)$ with high probability. Therefore, the key insight behind the new estimator construction is to look for the best geometry in each local neighborhood. Specifically, we propose the following two-stage estimator: at the first stage we collect $m_{1}$ nearest control covariates of the query point, while at the second stage we find the 1-nearest-neighbor treatment covariates for all above control samples and only pick $m_{2} \leq m_{1}$ of them with desirably small nearest neighbor distances. Note that the nearest-neighbor estimator is a better candidate than the kernel-based estimator here, as the nearest-neighbor distance can adapt to different density regimes and result in adaptive bandwidths.
The detailed estimator construction is displayed in Algorithm 2. Note that the first stage is similar in spirit to the kernel estimator, where we simply collect $m_{1}$ control samples in the small neighborhood and aim to take the (weighted) average. However, the second stage is very different in the sense that it asks to throw away samples with poor covariate matching quality to all treatment covariates and only keeps $m_{2}$ pairs of samples. Finally, a simple average is taken to the difference of the selected $m_{1}$ pairs, which is sufficient for smoothness $\beta_{\tau} \leq 1$. Although Algorithm 2 can be applied to higher smoothness in principle, we will discuss the fundamental challenges of $\beta_{\tau}>1$ in Section 5

```
Algorithm 2 Estimator Construction under Random Design
    Input: control observations \(\left\{\left(X_{i}^{0}, Y_{i}^{0}\right)\right\}_{i \in[n]}\), treatment observations \(\left\{\left(X_{i}^{1}, Y_{i}^{1}\right)\right\}_{i \in[n]}\), dimension-
    ality \(d\), number of neighbors \(m_{1}, m_{2}\) with \(m_{1} \geq \kappa m_{2}\), query point \(x_{0} \in[0,1]^{d}\).
    Output: HTE estimator \(\hat{\tau}\left(x_{0}\right)\) at query point \(x_{0}\).
    Find \(m_{1}\) nearest control covariates \(\left\{X_{i}^{0}\right\}_{i \in I_{1}}\) in the Euclidean distance to \(x_{0}\), with \(\left|I_{1}\right|=m_{1}\);
    for \(i \in I_{1}\) do
        Compute the minimum distance \(d_{i}=\min _{j \in[n]}\left\|X_{i}^{0}-X_{j}^{1}\right\|_{2}\) to the treatment covariate;
        Let \(j(i) \in[n]\) be any minimizer to the above minimization program.
    end for
    Find the smallest \(m_{2}\) elements of \(\left(d_{i}\right)_{i \in I_{1}}\) and denote the index set by \(I_{2}\), with \(\left|I_{2}\right|=m_{2}\);
    Return the final estimate \(\hat{\tau}\left(x_{0}\right)=m_{2}^{-1} \sum_{i \in I_{2}}\left(Y_{j(i)}^{1}-Y_{i}^{0}\right)\).
```

The optimal choices of $\left(m_{1}, m_{2}\right)$ depend on the density ratio $\kappa$ and noise level $\sigma$ as follows:

$$
\left(m_{1}, m_{2}\right)=\widetilde{\Theta} \begin{cases}\left(\kappa^{\frac{\beta_{\mu}}{\beta \tau+\beta_{\mu}}} n^{\frac{\beta_{\tau}-\beta_{\mu}}{\beta_{\tau}+\beta_{\mu}}}, 1\right)_{d \beta_{\mu}} & \text { if } \sigma \leq \sigma_{1} \\ \left(n \cdot\left(\kappa \sigma^{2} / n^{2}\right)^{\frac{2 \beta_{\mu} \beta_{\tau}}{2 \beta_{\mu} \beta_{\tau}+d\left(\beta_{\mu}+\beta_{\tau}\right)}},\left(n^{2} / \kappa\right)^{\frac{2 d\left(\beta_{\mu}+\beta_{\tau}\right)}{2 \beta_{\mu} \beta_{\tau}+d\left(\beta_{\mu}+\beta_{\tau}\right)}} \sigma^{\frac{2 d \beta_{\mu} \beta_{\tau}+d\left(\beta_{\mu}+\beta_{\tau}\right)}{2 \beta}}\right) & \text { if } \sigma_{1}<\sigma \leq \sigma_{2} \\ \left(n^{\frac{2 \beta_{\tau}}{2 \beta_{\tau}+d}}\left(\sigma^{2} \kappa\right)^{\frac{d}{2 \beta_{\tau}+d}},(n / \kappa)^{\frac{2 \beta_{\tau}}{2 \beta_{\tau}+d}} \sigma^{\frac{2 d}{2 \beta_{\tau}+d}}\right) & \text { if } \sigma>\sigma_{2}\end{cases}
$$

with noise thresholds ( $\sigma_{1}, \sigma_{2}$ ) given by

$$
\sigma_{1}=\left(\kappa / n^{2}\right)^{1 / d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}, \quad \sigma_{2}=n^{\left(\beta_{\tau}-\beta_{\mu}\right) /\left(2 \beta_{\tau}\right)-\beta_{\mu} / d} / \sqrt{\kappa} .
$$

The rationale behind the above choices lies in the following theorem.
Theorem 5. Given Assumptions $1 / 3$ and $\beta_{\mu} \leq \beta_{\tau} \leq 1$, fix any integer parameters $\left(m_{1}, m_{2}\right)$ with $m_{1} \geq \kappa m_{2}$ and $m_{2}=\Omega(\log n)$, the following upper bound holds for the estimator $\hat{\tau}$ in Algorithm 2 .

$$
\sup _{\mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right), \tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right)} \mathbb{E}_{\mu_{0}, \tau}\left[\|\hat{\tau}-\tau\|_{L_{1}\left(g_{0}\right)}\right] \leq \operatorname{polylog}(n) \cdot\left(\left(\frac{\kappa m_{2}}{n m_{1}}\right)^{\frac{\beta_{\mu}}{d}}+\left(\frac{m_{1}}{n}\right)^{\frac{\beta_{\tau}}{d}}+\frac{\sigma}{\sqrt{m_{2}}}\right)
$$

where $\operatorname{poly} \log (n)$ hides poly-logarithmic factors in $n$. In particular, if the parameters $\left(m_{1}, m_{2}\right)$ are chosen as above, the estimator $\hat{\tau}$ achieves the upper bound of Theorem 2 .

We explain the originations of the three types of errors in Theorem 5. The first error comes from the covariate matching bias, where it is given by the average of the $m_{2}$ smallest 1-nearest-neighbor distances between control-treatment pairs at the second stage. The second error comes from the bias incurred at the first stage, where the key quantity of interest is the distance between a random query point $x_{0} \sim g_{0}$ to its $m_{1}$-nearest-neighbor. Finally, the third error comes from the observational noise, where it is expected that the noise magnitude drops off by a factor of $\sqrt{m_{2}}$ after averaging. Hence, the optimal choice of the parameters $\left(m_{1}, m_{2}\right)$ aims to balance the above types of errors, with an additional constraint $\kappa m_{2} \leq m_{1}$ to ensure that no single treatment covariate is matched to too many control covariates and therefore lead to a small variance.
The full proof of Theorem 5 is postponed to the appendix, but we remark that an appropriate maximal inequality is the key to handling general densities $g_{0}, g_{1}$ which may be close to zero. Specifically, let $f$ be any density function on $[0,1]^{d}$, and define the following minimal function of $f$ as

$$
\begin{equation*}
m[f](x) \triangleq \inf _{0<r \leq \sqrt{d}} \frac{1}{\operatorname{Vol}_{d}(B(x ; r))} \int_{B(x ; r)} f(t) d t \tag{7}
\end{equation*}
$$

where $B(x ; r)$ denotes the ball centered at $x$ with radius $r, \operatorname{Vol}_{d}$ is the $d$-dimensional volume, and $f(t)=0$ whenever $t \notin[0,1]^{d}$. It is clear that $m[f](x) \leq f(x)$ for almost all $x \in[0,1]^{d}$, while the following lemma summarizes a key property of $m[f]$ which is very useful in the proof of Theorem 5 and could be of independent interest.
Lemma 1. For every $\lambda>0$, there exists a constant $C_{d}>0$ depending only on $d$ such that

$$
\int_{[0,1]^{d}} f(x) e^{-\lambda \cdot m[f](x)} d x \leq \begin{cases}\exp \left(-\lambda / e C_{d}\right) & \text { if } \lambda<e C_{d} \\ C_{d} / \lambda & \text { if } \lambda \geq e C_{d}\end{cases}
$$

### 3.2 Minimax Lower Bound

The following minimax lower bound complements Theorem 5 and shows the near minimax optimality (up to logarithmic factors) of the two-stage estimator in Algorithm 2
Theorem 6. Assume that $\beta_{\mu} \leq \beta_{\tau} \leq 1$. Then under Gaussian noises and appropriate densities $g_{0}, g_{1}$ satisfying Assumption [2] the minimax risk under the random design satisfies

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\text {random }}(\kappa) \geq \frac{1}{\operatorname{polylog}(n)}\left(\left(\frac{\kappa}{n^{2}}\right)^{\frac{1}{d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}}+\left(\frac{\kappa \sigma^{2}}{n^{2}}\right)^{\frac{1}{2+d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}}+\left(\frac{\kappa \sigma^{2}}{n}\right)^{\frac{\beta \tau}{2 \beta_{\tau}+d}}\right)
$$

The proof of Theorem6 is very involved and postponed to the appendix. Similar to the fixed design, the last error term comes from the classical nonparametric estimation of $\tau$, and establishing the first error term requires to solve a linear program in the same spirit to (6). However, as a regular grid is no longer available, constructing a near-optimal solution to becomes much more challenging. In addition, proving the second error term requires a hybrid approach of testing multiple hypotheses and constructing an exponential number of feasible solutions to (6). To this end, the following novel representation of the Hölder continuity condition is very useful throughout the lower bound argument. Lemma 2. Let $\beta \leq 1$, and $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n}$ be fixed covariate-value pairs on $[0,1]^{d} \times \mathbb{R}$. Then there exists some function $f \in \mathcal{H}_{d}(\beta)$ (with radius $L$ ) with $f\left(x_{i}\right)=y_{i}$ for all $i \in[n]$ if and only if for any $i, j \in[n]$, we have $\left|y_{i}-y_{j}\right| \leq L\left\|x_{i}-x_{j}\right\|_{2}^{\beta}$.

## 4 Numerical Experiments

We illustrate the efficacy of the proposed estimator in Algorithm 2 via some numerical experiments. Specifically, we aim to show that the two main ingredients of Algorithm 2, i.e. constructing pseudoobservations based on covariate matching and discarding observations with poor matching quality, are key to improved HTE estimation. We compare our estimator (which we call selected matching) with the following three estimators: the full matching estimator which never discards samples (i.e. $m_{2}=m_{1}$ always holds in Algorithm 2], the $k N N$ differencing and kernel differencing estimators which apply separate $k$-NN or kernel estimates to both baselines and then take the difference. Note that the idea of differencing-based estimators were essentially used in [AS18, KSBY19].

The experimental setting is as follows. We choose the parameter values $n \in\{50,100, \cdots, 1,000\}$, $d, \kappa \in\{1, \cdots, 10\}$, and $\sigma=2 / \sqrt{n}$. Denote by $\phi_{\mu, \sigma^{2}}(x)$ the pdf of the normal random variable $\mathcal{N}\left(\mu, \sigma^{2}\right)$, the baseline function and the HTE are chosen to be

$$
\mu_{0}(x)=2 \phi_{0.1,0.15^{2}}(x)+2.5 \phi_{0.4,0.05^{2}}(x)+4 \phi_{0.8,0.1^{2}}(x), \quad \tau(x)=\phi_{0,1}(2 x-1)
$$

for $d=1$ and $x \in[0,1]$, and for general dimensions, the variable $x$ is replaced by $\sqrt{d}(\bar{x}-0.5)+0.5$. An illustration of $\mu_{0}$ and $\tau$ for $d=1$ is plotted in Figure 1 , where $\tau$ is smoother and $\mu_{0}$ is more volatile. For each given $(n, d, \kappa, \sigma)$, we generate $n$ control covariates $X_{1}^{0}, \cdots, X_{n}^{0}$ following the i.i.d. density $g_{0}(x)=2 \kappa /(\kappa+1)$ when $x_{1} \in[0,1 / 2]$ and $g_{0}(x)=2 /(\kappa+1)$ when $x_{1} \in(1 / 2,1]$. Similarly, the treatment covariates $X_{1}^{1}, \cdots, X_{n}^{1}$ are i.i.d. generated following the density $g_{1}(x)=2-g_{0}(x)$, and the responses $Y_{i}^{0}, Y_{i}^{1}$ are defined in (1) with i.i.d. $\mathcal{N}\left(0, \sigma^{2}\right)$ noises. The algorithm parameters are determined by the optimal bias-variance tradeoffs in theory. Finally, the performance of HTE estimation is measured via the root mean squared error (RMSE) averaged over 100 simulations. The source codes are available at https://github.com/Mathegineer/NonparametricHTE.
The experimental results are displayed in Figures 1 and 2 , and we have the following observations. First, the differencing based estimators typically perform worse than the matching based estimators, especially at the places where the baseline is more volatile. Therefore, the differencing suffers more from the less smooth baselines. Second, although the full matching estimator behaves similarly as selected matching when $d=\kappa=1$, the performance of full matching deteriorates significantly if either $d$ or $\kappa$ increases, and selected matching is relatively more stable. This observation supports our intuition that the pairs with poor covariate matching quality should be discarded.

## 5 Further Discussions

A major concern of this work is to capture higher order smoothness under the random design, which we remark is very challenging from both sides of achevability and lower bounds. For example,


Figure 1: The baseline function $\mu_{0}, \operatorname{HTE} \tau$, as well as the estimates.


Figure 2: The RMSE as a function of $n, \kappa$ or $d$ with default parameters $(n, \kappa, d)=(1000,1,1)$.
there are a number of open questions related to capturing higher order smoothness in nonparametric statistics, e.g. the adaptation to small densities $\left[\mathrm{PR}^{+} 16\right]$ despite the known asymptotic equivalence [Nus96], in the estimation of conditional variance [RR14], non-smooth functionals [HJWW17], the analysis of nearest-neighbor methods [JGH18], local goodness-of-fit tests [BW19]. Specializing to our setting, to capture the smoothness $\beta_{\mu}>1$ and reduce the matching bias, locally there need to be at least two treatment-control pairs mutually of small distance. However, instead of the $\Theta\left(n^{-2}\right)$ minimum distance for $n$ uniform random variables in $[0,1]$, the counterpart for two pairs becomes at least $\Theta\left(n^{-4 / 3}\right)$. Hence, we expect that there will be a phase transition from $\beta_{\mu} \leq 1$ to $\beta_{\mu}>1$ in the estimation performance, and it is an outstanding open question to characterize this transition.
As for the lower bound, the main difficulty of the extension to $\beta \geq 1$ lies in Lemma 2 To the best of our knowledge, there is no simple and efficient criterion for the existence of a Hölder $\beta$-smooth function with specified values on general prescribed points. To illustrate the difficulty, assume that $\beta$ is an integer and consider the following natural generalization of the condition in Lemma 2 for any $i_{0}, i_{1}, \cdots, i_{\beta} \in[n]$, it holds that $\beta!\cdot\left|f\left[x_{i_{0}}, x_{i_{1}}, \cdots, x_{i_{\beta}}\right]\right| \leq L$, where $f\left[x_{0}, x_{1}, \cdots, x_{m}\right]$ is the divided difference of $f$ on points $x_{0}, x_{1}, \cdots, x_{m}$ [dB05]. Standard approximation theory shows that this condition is necessary, and it reduces to Lemma 2 for $\beta=1$. However, it is not sufficient when $\beta \geq 2$. For a counterexample, the value specification $f(-3)=f(-1)=1, f(1)=f(3)=0$ satisfies the divided difference condition with $\beta=2, L=1 / 4$, but no choice of $f(0)$ satisfies both $|f[-3,-1,0]| \leq 1 / 8$ and $|f[0,1,3]| \leq 1 / 8$.
We also point out some other extensions. First, the current estimators rely on the knowledge of many parameters which are unlikely to be known in practice. For the smoothness parameters and the density ratio parameter, their impact factors through the bias and it is possible to apply Lepski's trick [LMS97] or data-driven parameters [FG95] to achieve full or partial adaptation. For other parameters such as the noise level, pre-estimation procedures such as [DJKP95] are expected to be helpful. Second, in our work, we study the minimax risk after conditioning on both random realizations of the covariates and group assignments, meaning that the worst-case HTE can depend on both the above realizations. In contrast, NW17] only conditions on the random covariates and considers an expected risk taken with respect to the randomness in the group assignments, where the worst-case HTE cannot depend on the group assignments. The minimax risk in the latter setting remains unknown, and it is an outstanding question to compare the minimax rates in these settings.

## Broader Impact

This work mainly provides theoretical tools and bounds for the HTE estimation in causal inference, as well as potentially useful practical insights such as the two-stage nearest-neighbors and throwing away observations with poor covariate matching quality. This special form of nonparametric estimation problems could be a useful addition to the literature in nonparametric statistics, and theorists and practitioners working on causal inference may potentially benefit from this work.

## References

[AI16] Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. Proceedings of the National Academy of Sciences, 113(27):7353-7360, 2016.
[AS18] Ahmed Alaa and Mihaela Schaar. Limits of estimating heterogeneous treatment effects: Guidelines for practical algorithm design. In International Conference on Machine Learning, pages 129-138, 2018.
[BCLZ02] Lawrence D Brown, T Tony Cai, Mark G Low, and Cun-Hui Zhang. Asymptotic equivalence theory for nonparametric regression with random design. The Annals of statistics, 30(3):688-707, 2002.
[BD15] Gérard Biau and Luc Devroye. Lectures on the nearest neighbor method. Springer, 2015.
[BW19] Sivaraman Balakrishnan and Larry Wasserman. Hypothesis testing for densities and high-dimensional multinomials: Sharp local minimax rates. The Annals of Statistics, 47(4):1893-1927, 2019.
$\left[C C D{ }^{+}\right.$18] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.
[Che07] Xiaohong Chen. Large sample sieve estimation of semi-nonparametric models. Handbook of econometrics, 6:5549-5632, 2007.
[dB05] Carl de Boor. Divided differences. Surveys in Approximation Theory, 1:46-69, 2005.
[DBDBM $\left.{ }^{+} 78\right]$ Carl De Boor, Carl De Boor, Etats-Unis Mathématicien, Carl De Boor, and Carl De Boor. A practical guide to splines, volume 27. springer-verlag New York, 1978.
[DJKP95] David L Donoho, Iain M Johnstone, Gérard Kerkyacharian, and Dominique Picard. Wavelet shrinkage: asymptopia? Journal of the Royal Statistical Society: Series B (Methodological), 57(2):301-337, 1995.
[FG95] Jianqing Fan and Irene Gijbels. Data-driven bandwidth selection in local polynomial fitting: variable bandwidth and spatial adaptation. Journal of the Royal Statistical Society: Series B (Methodological), 57(2):371-394, 1995.
[ $\mathrm{FGG}^{+}$97] Jianqing Fan, Theo Gasser, Irène Gijbels, Michael Brockmann, and Joachim Engel. Local polynomial regression: optimal kernels and asymptotic minimax efficiency. Annals of the Institute of Statistical Mathematics, 49(1):79-99, 1997.
[Fix51] Evelyn Fix. Discriminatory analysis: nonparametric discrimination, consistency properties. USAF school of Aviation Medicine, 1951.
[GKKW06] László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. A distribution-free theory of nonparametric regression. Springer Science \& Business Media, 2006.
[GKOV17] Weihao Gao, Sreeram Kannan, Sewoong Oh, and Pramod Viswanath. Estimating mutual information for discrete-continuous mixtures. In Advances in neural information processing systems, pages 5986-5997, 2017.
[GS93] Peter J Green and Bernard W Silverman. Nonparametric regression and generalized linear models: a roughness penalty approach. Chapman and Hall/CRC, 1993.
[HJWW17] Yanjun Han, Jiantao Jiao, Tsachy Weissman, and Yihong Wu. Optimal rates of entropy estimation over lipschitz balls. arXiv preprint arXiv:1711.02141, 2017.
[ IR ${ }^{+}$13] Kosuke Imai, Marc Ratkovic, et al. Estimating treatment effect heterogeneity in randomized program evaluation. The Annals of Applied Statistics, 7(1):443-470, 2013.
[IR15] Guido W Imbens and Donald B Rubin. Causal inference in statistics, social, and biomedical sciences. Cambridge University Press, 2015.
[JGH18] Jiantao Jiao, Weihao Gao, and Yanjun Han. The nearest neighbor information estimator is adaptively near minimax rate-optimal. In Advances in neural information processing systems, pages 3156-3167, 2018.
[Joh11] Iain M Johnstone. Gaussian estimation: Sequence and wavelet models. Under contract to Cambridge University Press, 2011.
[Ken16] Edward H Kennedy. Semiparametric theory and empirical processes in causal inference. In Statistical causal inferences and their applications in public health research, pages 141-167. Springer, 2016.
[KSBY19] Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. Proceedings of the National Academy of Sciences, 116(10):4156-4165, 2019.
[LMS97] Oleg V Lepski, Enno Mammen, and Vladimir G Spokoiny. Optimal spatial adaptation to inhomogeneous smoothness: an approach based on kernel estimates with variable bandwidth selectors. The Annals of Statistics, pages 929-947, 1997.
[MU05] Michael Mitzenmacher and Eli Upfal. Probability and computing: Randomized algorithms and probabilistic analysis. Cambridge University Press, 2005.
[Nad64] Elizbar A Nadaraya. On estimating regression. Theory of Probability \& Its Applications, 9(1):141-142, 1964.
[Nem00] Arkadi Nemirovski. Topics in non-parametric statistics. Ecole d?Eté de Probabilités de Saint-Flour, 28:85, 2000.
[Nus96] Michael Nussbaum. Asymptotic equivalence of density estimation and gaussian white noise. The Annals of Statistics, pages 2399-2430, 1996.
[NW17] Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. arXiv preprint arXiv:1712.04912, 2017.
[PQJ ${ }^{+}$17] Scott Powers, Junyang Qian, Kenneth Jung, Alejandro Schuler, Nigam H Shah, Trevor Hastie, and Robert Tibshirani. Some methods for heterogeneous treatment effect estimation in high-dimensions. arXiv preprint arXiv:1707.00102, 2017.
[ $\mathrm{PR}^{+}$16] Tim Patschkowski, Angelika Rohde, et al. Adaptation to lowest density regions with application to support recovery. The Annals of Statistics, 44(1):255-287, 2016.
[Ros56] Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. The Annals of Mathematical Statistics, pages 832-837, 1956.
[RR14] Thomas S Richardson and Andrea Rotnitzky. Causal etiology of the research of james m. robins. Statistical Science, pages 459-484, 2014.
[Rub74] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of educational Psychology, 66(5):688, 1974.
[Sto77] Charles J Stone. Consistent nonparametric regression. The annals of statistics, pages 595-620, 1977.
[TAGT14] Lu Tian, Ash A. Alizadeh, Andrew J. Gentles, and Robert Tibshirani. A simple method for estimating interactions between a treatment and a large number of covariates. Journal of the American Statistical Association, 109(508):1517-1532, 2014. PMID: 25729117.
[Tsy09] Alexandre B Tsybakov. Introduction to nonparametric estimation. Springer, 2009.
[WA18] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, 113(523):1228-1242, 2018.
[Wat64] Geoffrey S Watson. Smooth regression analysis. Sankhyā: The Indian Journal of Statistics, Series A, pages 359-372, 1964.
[XBJ12] Yu Xie, Jennie E Brand, and Ben Jann. Estimating heterogeneous treatment effects with observational data. Sociological methodology, 42(1):314-347, 2012.
[Yu97] Bin Yu. Assouad, fano, and le cam. In Festschrift for Lucien Le Cam, pages 423-435. Springer, 1997.

## A Auxiliary Lemmas

Lemma 3. Fix any $t \in \mathbb{N}, m \geq 1$ and $0 \leq \Delta \leq 1 /(2 m)$. For any $i \in[m]$, let $x_{0}, \cdots, x_{t}$ be the $(t+1)$ nearest neighbors of $(i-1) / m$ in $\{\Delta, 1 / m+\Delta, \cdots, 1-1 / m+\Delta\}$ with increasing distance, and the weights $w_{0}, \cdots, w_{t}$ be the solution to

$$
\sum_{j=0}^{t} w_{j}=1, \quad \sum_{j=0}^{t} w_{j}\left(x-x_{j}\right)^{\ell}=0, \quad \ell=1, \cdots, t
$$

Then there exists a constant $C>0$ depending only on $t$ such that

$$
\left|w_{0}\right| \leq C, \quad\left|w_{j}\right| \leq C \cdot m \Delta, \quad j=1,2, \cdots, t .
$$

Lemma 4 (Chernoff bound, Theorem 5.4 of [MU05]). For $X \sim B(n, \lambda / n)$ and any $\delta \in(0,1]$, we have

$$
\begin{aligned}
& \mathbb{P}(X \geq(1+\delta) \lambda) \leq\left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\lambda} \leq \exp \left(-\frac{\delta^{2} \lambda}{3}\right) \\
& \mathbb{P}(X \leq(1-\delta) \lambda) \leq\left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\lambda} \leq \exp \left(-\frac{\delta^{2} \lambda}{2}\right)
\end{aligned}
$$

Lemma 5 (Generalized Hardy-Littlewood maximal inequality, a slight variant of Lemma 3 of [JGH18]). Let $\mu_{1}, \mu_{2}$ be two Borel measures that are finite on the bounded Borel sets of $\mathbb{R}^{d}$. Then, for all $t>0$ and any bounded Borel set $A \subseteq \mathbb{R}^{d}$,

$$
\mu_{1}\left(\left\{x \in A: \sup _{\rho>0: B(x ; \rho) \subseteq A}\left(\frac{\mu_{2}(B(x ; \rho))}{\mu_{1}(B(x ; \rho))}\right) \geq t\right\}\right) \leq \frac{C_{d}}{t} \mu_{2}(A)
$$

Here $C_{d}>0$ is a constant that depends only on the dimension $d$.

## B Proof of Main Theorems

Throughout the proofs, we introduce the following notations: in addition to the asymptotic notation $a_{n}=\Theta\left(b_{n}\right)$, we also write $a_{n}=O\left(b_{n}\right)$ if $\limsup _{n \rightarrow \infty} a_{n} / b_{n}<\infty$, and $a_{n}=\Omega\left(b_{n}\right)$ if $b_{n}=$ $O\left(a_{n}\right)$. Similarly, the notations $\widetilde{O}(\cdot), \widetilde{\Omega}(\cdot)$ denote the respective meanings within logarithmic factors in $n$.

## B. 1 Proof of Theorem 1

The minimax risk of HTE estimation under the fixed design is a direct consequence of Theorems 3 and 4

## B.1. 1 Proof of Theorem 3

By the definition of the pseudo-observation $\hat{Y}^{1}(x)$ and the potential outcome model (1), we have

$$
\begin{equation*}
\hat{Y}^{1}\left(X_{i}^{0}\right)-Y_{0}^{i}=\tau\left(X_{i}^{0}\right)+b_{\mu}\left(X_{i}^{0}\right)+\tilde{\varepsilon}_{i}, \tag{8}
\end{equation*}
$$

with the new errors $\tilde{\varepsilon}_{i}$ defined as (with $t=\left\lfloor\beta_{\mu}\right\rfloor+1$ )

$$
\tilde{\varepsilon}_{i}=\sum_{j \in[t]^{d}} w_{j} \cdot \varepsilon_{j}^{1}-\varepsilon_{i}^{0}
$$

and the matching bias $b_{\mu}\left(X_{i}^{0}\right)$ is

$$
b_{\mu}\left(X_{i}^{0}\right)=\sum_{j \in[t]^{d}} w_{j} \cdot\left(\mu_{0}\left(x_{j}\right)+\tau\left(x_{j}\right)\right)-\left(\mu_{0}\left(X_{i}^{0}\right)+\tau\left(X_{i}^{0}\right)\right),
$$

where $w_{j}, x_{j}$ are the weights and the treatment covariates used to obtain pseudo-observations at $X_{i}^{0}$. By the smoothness property of $\mu_{1}=\mu_{0}+\tau$ in Assumption 1 , the Taylor expansion around $X_{i}^{0}$ gives

$$
\begin{aligned}
\mu_{1}\left(x_{j}\right)= & \sum_{\beta=0}^{\left\lfloor\beta_{\mu}\right\rfloor-1} \sum_{\sum_{i=1}^{d} \beta_{i}=\beta} \frac{\partial^{\beta} \mu_{1}\left(X_{i}^{0}\right)}{\partial x_{1}^{\beta_{1}} \cdots \partial x_{d}^{\beta_{d}}} \prod_{k=1}^{d} \frac{\left(x_{j, k}-X_{i, k}^{0}\right)^{\beta_{k}}}{\beta_{k}!} \\
& +\sum_{\sum_{i=1}^{d} \beta_{i}=\left\lfloor\beta_{\mu}\right\rfloor} \frac{\partial^{\beta} \mu_{1}\left(\xi_{j}\right)}{\partial x_{1}^{\beta_{1}} \cdots \partial x_{d}^{\beta_{d}}} \prod_{k=1}^{d} \frac{\left(x_{j, k}-X_{i, k}^{0}\right)^{\beta_{k}}}{\beta_{k}!}
\end{aligned}
$$

where $\xi_{j} \in \mathbb{R}^{d}$ is some point on the segment connecting $X_{i}^{0}$ and $x_{j}$. Hence, by the smoothness assumption of $\mu_{1}$, it holds that
$\left|\mu_{1}\left(x_{j}\right)-\sum_{\beta=0}^{\left\lfloor\beta_{\mu}\right\rfloor} \sum_{\sum_{i=1}^{d} \beta_{i}=\beta} \frac{\partial^{\beta} \mu_{1}\left(X_{i}^{0}\right)}{\partial x_{1}^{\beta_{1}} \cdots \partial x_{d}^{\beta_{d}}} \prod_{k=1}^{d} \frac{\left(x_{j, k}-X_{i, k}^{0}\right)^{\beta_{k}}}{\beta_{k}!}\right| \leq C \sum_{\left(\beta_{1}, \cdots, \beta_{d}\right) \in \Gamma} \prod_{k=1}^{d}\left|x_{j, k}-X_{i, k}^{0}\right|^{\beta_{k}}$,
where $C>0$ is an absolute constant, and

$$
\Gamma=\left\{\left(\beta_{1}, \cdots, \beta_{d}\right): \beta_{k} \geq 0, \sum_{k=1}^{d} \beta_{k}=\beta_{\mu}, \text { at most one of } \beta_{k} \text { is non-integer }\right\}
$$

is a finite set of indices. Now by the choice of each $w_{j}$ in (3), we have

$$
\begin{equation*}
\left|b_{\mu}\left(X_{i}^{0}\right)\right| \leq C \quad \sum_{\left(\beta_{1}, \cdots, \beta_{d}\right) \in \Gamma} \prod_{k=1}^{d}\left(\sum_{j=1}^{t}\left|w_{j, k}\right| \cdot\left|x_{j, k}-X_{i, k}^{0}\right|^{\beta_{k}}\right) \tag{9}
\end{equation*}
$$

Consider the term $\left|w_{j, k}\right| \cdot\left|x_{j, k}-X_{i, k}^{0}\right|^{\beta_{k}}$ for each index $j \in[t]$ and $k \in[d]$. If $x_{j, k}$ is the nearest neighbor of $X_{i, k}^{0}$ in the one-dimensional grid $\left\{\Delta_{k}, 1 / m+\Delta_{k}, \cdots, 1-1 / m+\Delta_{k}\right\}$, then $\left|x_{j, k}-X_{i, k}^{0}\right|=\Delta_{k}$ and $\left|w_{j, k}\right|=O(1)$ by Lemma 3. Otherwise, we have $\left|x_{j, k}-X_{i, k}^{0}\right| \leq t / m$ and $\left|w_{j, k}\right|=O\left(m \Delta_{k}\right)$ again by Lemma 3. Therefore, the following inequality always holds:

$$
\left|w_{j, k}\right| \cdot\left|x_{j, k}-X_{i, k}^{0}\right|^{\beta_{k}}=O\left(m^{-\beta_{k}}\left(m \Delta_{k}\right)^{\beta_{k} \wedge 1}\right)=O\left(m^{-\beta_{k}}\left(m\|\Delta\|_{\infty}\right)^{\beta_{k} \wedge 1}\right) .
$$

Consequently, for any $i \in[n]$ it holds that
$\left|b_{\mu}\left(X_{i}^{0}\right)\right|=O\left(C t^{d} \cdot|\Gamma| \max _{\left(\beta_{1}, \cdots, \beta_{d}\right) \in \Gamma}\left(\prod_{k=1}^{d} m^{-\beta_{k}}\left(m\|\Delta\|_{\infty}\right)^{\beta_{k} \wedge 1}\right)\right)=O\left(m^{-\beta_{\mu}}\left(m\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}\right)$,
which gives the claimed matching bias upper bound in the error decomposition of Theorem3
The rest of the proof follows from the classic analysis of the Nadaraya-Watson estimator applied to (8] (see, e.g. [Tsy09]), where the only difference is that the errors $\tilde{\varepsilon}_{i}$ are weakly dependent instead of being mutually independent, i.e. each $\tilde{\varepsilon}_{i}$ depends only on finitely many $\tilde{\varepsilon}_{j}$ 's.

## B.1.2 Proof of Theorem 4

To prove the lower bound $\Omega\left(\left(\sigma^{2} / n\right)^{\beta_{\tau} /\left(2 \beta_{\tau}+d\right)}\right)$, consider $\mu_{0} \equiv 0$. In this scenario, the outcomes of the control group are completely non-informative for estimating $\tau$, and the treatment outcome model (1) is precisely the nonparametric regression model for the $\beta_{\tau}$-Hölder smooth function $\tau$. Then by the standard minimax lower bound for nonparametric regression (see, e.g. [Tsy09]), the lower bound $\Omega\left(\left(\sigma^{2} / n\right)^{\beta_{\tau} /\left(2 \beta_{\tau}+d\right)}\right)$ is immediate.
Next we prove the lower bound $\Omega\left(n^{-\beta_{\mu} / d}\left(n^{1 / d}\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}\right)$, which corresponds to the matching bias. We first construct a feasible solution $\left(\mu_{0}, \tau\right)$ to the optimization problem (6). Without loss of generality we assume that $\Delta_{1}=\|\Delta\|_{\infty}$. Construct the following function $g$ on $[0,1]^{d}$ : for $x=\left(x_{1}, \cdots, x_{d}\right) \in[0,1]^{d}$,

$$
g(x)=c \cdot\left[x_{1}\left(1-x_{1}\right)\right]^{\beta_{\mu} \wedge 1}
$$

with some small constant $c>0$. We claim that $g$ is $\beta$-Hölder smooth on $[0,1]^{d}$ given a sufficiently small $c>0$. In fact, if $\beta_{\mu} \leq 1$, for all $x_{1}, x_{1}^{\prime} \geq 0$ simple algebra gives $\left|\left(x_{1}^{\prime}\right)^{\beta_{\mu}}-x_{1}^{\beta_{\mu}}\right| \leq\left|x_{1}^{\prime}-x_{1}\right|^{\beta_{\mu}}$, and therefore

$$
\begin{aligned}
\left|g(x)-g\left(x^{\prime}\right)\right| & =c \cdot\left|\left[x_{1}\left(1-x_{1}\right)\right]^{\beta_{\mu}}-\left[x_{1}^{\prime}\left(1-x_{1}^{\prime}\right)\right]^{\beta_{\mu}}\right| \\
& \leq c \cdot\left(\left(1-x_{1}\right)^{\beta_{\mu}}\left|x_{1}^{\beta_{\mu}}-\left(x_{1}^{\prime}\right)^{\beta_{\mu}}\right|+\left(x_{1}^{\prime}\right)^{\beta_{\mu}}\left|\left(1-x_{1}\right)^{\beta_{\mu}}-\left(1-x_{1}^{\prime}\right)^{\beta_{\mu}}\right|\right) \\
& \leq c \cdot\left(1 \cdot\left|x_{1}-x_{1}^{\prime}\right|^{\beta_{\mu}}+1 \cdot\left|x_{1}-x_{1}^{\prime}\right|^{\beta_{\mu}}\right) \\
& =2 c\left|x_{1}-x_{1}^{\prime}\right|^{\beta_{\mu}} \leq 2 c\left\|x-x^{\prime}\right\|_{2}^{\beta_{\mu}} .
\end{aligned}
$$

If $\beta_{\mu}>1$, the function $g(x)=c x_{1}\left(1-x_{1}\right)$ is a smooth function in $x$ and therefore $\beta_{\mu}$-Hölder smooth as well.
Recall that $(0,1]^{d}=\cup_{i=1}^{n}\left(X_{i}^{0}+\left(0, m^{-1}\right]^{d}\right)$ under the fixed design. Consider the following construction of $\mu_{0}$ :

$$
\mu_{0}(x)=h(x) \cdot \frac{1}{m^{\beta_{\mu}}} \sum_{i=1}^{n} g\left(m\left(x-X_{i}^{0}\right)\right) \mathbb{1}\left(m\left(x-X_{i}^{1}\right) \in(0,1]^{d}\right)
$$

where the function $g$ is defined above, and $h$ is an arbitrary smooth function on $\mathbb{R}^{d}$ supported on $[0,1]^{d}$ with $h(x) \equiv 1$ for $x \in[1 / 4,3 / 4]^{d}$. In other words, the baseline function $\mu_{0}$ is a dilation of the reference function $g$ with proper scaling to preserve $\beta_{\mu}$-smoothness, and the function $h$ preserves the value of $\mu_{0}$ in the center of $[0,1]^{d}$ while helps $\mu_{0}$ connect to zero smoothly at the boundary of $[0,1]^{d}$. Since multiplying a bounded smooth function does not decrease the smoothness parameter, it is straightforward to verify that $\mu_{0} \in \mathcal{H}_{d}\left(\beta_{\mu}\right)$ fulfills Assumption 1 .

We also construct the HTE $\tau$ as $\tau(x)=-m^{-\beta_{\mu}} g(m \Delta) \cdot h(x)$, where $h$ is the same smooth function used in the definition of $\mu_{0}$. By the smoothness of $h$, it is clear that $\tau \in \mathcal{H}_{d}\left(\beta_{\tau}\right)$ for any $\beta_{\tau}>0$. Moreover, the pair $\left(\mu_{0}, \tau\right)$ is a feasible solution to (6): the smoothness conditions have already been verified, and

$$
\begin{aligned}
\mu_{0}\left(X_{i}^{0}\right) & =h\left(X_{i}^{0}\right) \cdot m^{-\beta_{\mu}} g(0)=0 \\
\mu_{0}\left(X_{i}^{1}\right)+\tau\left(X_{i}^{1}\right) & =h\left(X_{i}^{1}\right) \cdot m^{-\beta_{\mu}} g(m \Delta)-m^{-\beta_{\mu}} g(m \Delta) \cdot h\left(X_{i}^{1}\right)=0
\end{aligned}
$$

Also, the construction of $g$ ensures that $|g(m \Delta)|=\Omega\left(\left(m\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}\right)$, and the fact $h(x) \equiv 1$ for $x \in[1 / 4,3 / 4]^{d}$ gives $\|h\|_{1}=\Omega(1)$. Consequently, we have

$$
\|\tau\|_{1}=m^{-\beta_{\mu}}|g(m \Delta)| \cdot\|h\|_{1}=\Omega\left(n^{-\frac{\beta_{\mu}}{d}}\left(n^{\frac{1}{d}}\|\Delta\|_{\infty}\right)^{\beta_{\mu} \wedge 1}\right) .
$$

Finally, note that the distributions of the outcomes (1) in two groups are the same under the above construction of $\left(\mu_{0}, \tau\right)$ and the naïve construction ( $\mu_{0}^{\prime} \equiv 0, \tau^{\prime} \equiv 0$ ), the standard two-point method yields to a minimax lower bound $\left\|\tau-\tau^{\prime}\right\|_{2}$, which is exactly the desired matching bias.

## B. 2 Proof of Theorem 2

The minimax risk of HTE estimation under the random design is a direct consequence of Theorems 5 and 6

## B.2.1 Proof of Theorem 5

For each $x_{0} \in[0,1]^{d}$, the construction of the two-stage nearest-neighbor based estimator $\hat{\tau}\left(x_{0}\right)$ in Algorithm 2 gives

$$
\begin{aligned}
\hat{\tau}\left(x_{0}\right)-\tau\left(x_{0}\right) & =\frac{1}{m_{2}} \sum_{i \in I_{2}}\left(\mu_{0}\left(X_{j(i)}^{1}\right)-\mu_{0}\left(X_{i}^{0}\right)\right)+\frac{1}{m_{2}} \sum_{i \in I_{2}}\left(\tau\left(X_{j(i)}^{1}\right)-\tau\left(x_{0}\right)\right)+\frac{1}{m_{2}} \sum_{i \in I_{2}}\left(\varepsilon_{j(i)}^{1}-\varepsilon_{i}^{0}\right) \\
& \triangleq b_{\mu}\left(x_{0}\right)+b_{\tau}\left(x_{0}\right)+s\left(x_{0}\right)
\end{aligned}
$$

where the above terms correspond to the matching bias incurred by $\mu$, estimation bias incurred by $\tau$, and the stochastic error incurred by the noises, respectively. As the final estimation error is measured
under $L_{1}\left(g_{0}\right)$, it suffices to upper bound the quantites $\mathbb{E}\left|b_{\mu}(X)\right|, \mathbb{E}\left|b_{\tau}(X)\right|$, and $\mathbb{E}|s(X)|$ for a fresh test covariate $X \sim g_{0}$, respectively.
We first upper bound the quantity $\mathbb{E}\left|b_{\mu}(X)\right|$. As in the description of the Algorithm 2 for $i \in I_{1}$, let $d_{i}=\min _{j \in[n]}\left\|X_{i}^{0}-X_{j}^{1}\right\|_{2}$ be the minimum Euclidean distance between the control covariate $X_{i}^{0}$ and all possible treatment covariates, and $D_{1}$ be the random variable representing the $m_{2}$-th smallest value of $\left(d_{i}\right)_{i \in I_{1}}$. Then by the smoothness condition of $\mu_{0}$, it is clear that

$$
\begin{equation*}
\mathbb{E}\left|b_{\mu}(X)\right| \lesssim \mathbb{E}\left[D_{1}^{\beta_{\mu}}\right] \leq\left(\mathbb{E}\left[D_{1}\right]\right)^{\beta_{\mu}} \tag{10}
\end{equation*}
$$

Hence, it remains to upper bound the expectation $\mathbb{E}\left[D_{1}\right]$. First, for any $x \in[0,1]^{d}, j \in[n]$ and $t \in(0, \sqrt{d})$, we have

$$
\mathbb{P}\left(\left\|X_{j}^{1}-x\right\|_{2} \geq t\right)=\int_{y \in[0,1]^{d}: y \notin B(x ; t)} g_{1}(y) d y \leq 1-\gamma_{d} t^{d} \cdot m\left[g_{1}\right](x),
$$

by the definition of the minimal function defined in (7), where $\gamma_{d}$ is the volume of the $d$-dimensional unit ball. Consequently, for each $i \in I_{1}$ and $0<t \lesssim n^{-1 / d}$, we have

$$
\begin{aligned}
\mathbb{P}\left(d_{i} \geq t\right) & =\mathbb{E}\left[\mathbb{P}\left(\left\|X_{1}^{1}-X_{i}^{0}\right\|_{2} \geq t \mid X_{i}^{0}\right)^{n}\right] \\
& \leq \mathbb{E}\left[\left(1-\gamma_{d} t^{d} \cdot m\left[g_{1}\right]\left(X_{i}^{0}\right)\right)^{n}\right] \\
& \stackrel{(\mathrm{a})}{\leq} \mathbb{E}\left[\exp \left(-n \gamma_{d} t^{d} \cdot m\left[g_{1}\right]\left(X_{i}^{0}\right)\right)\right] \\
& =\int_{[0,1]^{d}} g_{0}(x) \exp \left(-n \gamma_{d} t^{d} m\left[g_{1}\right](x)\right) d x \\
& \stackrel{\text { (b) }}{\leq} \int_{[0,1]^{d}} g_{0}(x) \exp \left(-\frac{n \gamma_{d} t^{d}}{\kappa} m\left[g_{0}\right](x)\right) d x \\
& \text { (c) } \exp \left(-\frac{c_{d} n t^{d}}{\kappa}\right)
\end{aligned}
$$

where (a) is due to the elementary inequality $1-x \leq \exp (-x)$, (b) follows from the upper bound on the density ratio, and (c) is due to Lemma 1 and the assumption $t \lesssim n^{-1 / d}$, with some absolute constant $c_{d}>0$. Now by the mutual independence of control and treatment covariates, the random variable $D_{1}$ is simply the $m_{2}$-th order statistic of $\left(d_{i}\right)_{i \in I_{1}}$, and therefore

$$
\begin{aligned}
\mathbb{P}\left(D_{1} \geq t\right) & =\mathbb{P}\left(\mathrm{B}\left(m_{1}, \mathbb{P}\left(d_{1} \geq t\right)\right) \geq m_{1}-m_{2}+1\right) \\
& \leq \mathbb{P}\left(\mathrm{B}\left(m_{1}, \exp \left(-c_{d} n t^{d} / \kappa\right)\right) \geq m_{1}-m_{2}+1\right) \\
& =\mathbb{P}\left(\mathrm{B}\left(m_{1}, 1-\exp \left(-c_{d} n t^{d} / \kappa\right)\right)<m_{2}\right) .
\end{aligned}
$$

Hence, for $t \geq C_{d}\left(\kappa m_{2} / n m_{1}\right)^{1 / d}$ with a large enough constant $C_{d}>0$, we have

$$
1-\exp \left(-c_{d} n t^{d} / \kappa\right) \leq \frac{m_{2}}{2 m_{1}}
$$

and therefore Chernoff bound (Lemma 4 gives the upper bound $\mathbb{P}\left(D_{1} \geq t\right)=\exp \left(-\Omega\left(m_{2}\right)\right)$, which is smaller than $O\left(n^{-1}\right)$ by the assumption that $m_{2}$ is at least logarithmic in $n$. Consequently,

$$
\mathbb{E}\left[D_{1}\right]=\int_{0}^{\sqrt{d}} \mathbb{P}\left(D_{1} \geq t\right) d t \leq C_{d}\left(\frac{\kappa m_{2}}{n m_{1}}\right)^{\frac{1}{d}}+\sqrt{d} \cdot \frac{1}{n} \lesssim\left(\frac{\kappa m_{2}}{n m_{1}}\right)^{\frac{1}{d}}
$$

and consequently (10) leads to

$$
\begin{equation*}
\mathbb{E}\left|b_{\mu}(X)\right| \lesssim\left(\frac{\kappa m_{2}}{n m_{1}}\right)^{\frac{\beta_{\mu}}{d}} \tag{11}
\end{equation*}
$$

We then upper bound the quantity $\mathbb{E}\left|b_{\tau}(X)\right|$. Let $D_{2}$ be the non-negative random variable representing the Euclidean distance between the query point $X$ and its $m_{1}$-nearest neighbor in the control group,
then standard nearest neighbor analysis (see, e.g. $\overline{\mathrm{BD} 15}$, Theorem 2.4]) gives $\mathbb{E}\left[D_{2}\right] \lesssim\left(m_{1} / n\right)^{1 / d}$. Consequently,

$$
\begin{align*}
\mathbb{E}\left|b_{\tau}(X)\right| & \leq \frac{1}{m_{2}} \sum_{i \in I_{2}} \mathbb{E}\left(\left|\tau\left(X_{j(i)}^{1}\right)-\tau\left(X_{i}^{0}\right)\right|+\left|\tau\left(X_{i}^{0}\right)-\tau(X)\right|\right) \\
& \lesssim \mathbb{E}\left[D_{1}^{\beta_{\tau}}\right]+\mathbb{E}\left[D_{2}^{\beta_{\tau}}\right] \\
& \leq\left(\mathbb{E}\left[D_{1}\right]\right)^{\beta_{\tau}}+\left(\mathbb{E}\left[D_{2}\right]\right)^{\beta_{\tau}} \\
& \lesssim\left(\frac{\kappa m_{2}}{n m_{1}}\right)^{\frac{\beta_{\tau}}{d}}+\left(\frac{m_{1}}{n}\right)^{\frac{\beta_{\tau}}{d}} \\
& \lesssim\left(\frac{m_{1}}{n}\right)^{\frac{\beta_{\tau}}{d}} \tag{12}
\end{align*}
$$

where the last inequality follows from the assumption that $\kappa m_{2} \leq m_{1}$.
Finally we deal with the stochastic error $\mathbb{E}|s(X)|$. Clearly, due to the mutual independence of the noises in Assumption 3, we have

$$
\mathbb{E}\left[s(X)^{2}\right] \lesssim \frac{\sigma^{2}}{m_{2}^{2}} \mathbb{E}\left[\sum_{j \in[n]} S_{j}^{2}\right]
$$

where $S_{j}=\sum_{i \in I_{2}} \mathbb{1}(i \rightarrow j)$ is the total number of times $X_{j}^{1}$ is chosen to be the nearest neighbor, and $i \rightarrow j$ denotes that $X_{j}^{1}$ is the nearest neighbor of $X_{i}^{0}$ among the treatment covariates. Consequently,

$$
\begin{align*}
\mathbb{E}\left[\sum_{j \in[n]} S_{j}^{2}\right] & =\mathbb{E}\left[\sum_{i \in I_{2}} \sum_{i^{\prime} \in I_{2}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right] \\
& =\mathbb{E}\left[\sum_{i \in I_{2}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j)\right]+\mathbb{E}\left[\sum_{i \neq i^{\prime} \in I_{2}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right] \\
& =m_{2}+\mathbb{E}\left[\sum_{i \neq i^{\prime} \in I_{2}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right], \tag{13}
\end{align*}
$$

where the last identity follows from $\sum_{j \in[n]} \mathbb{1}(i \rightarrow j)=1$. To deal with the second term, first assume that the control and treatment covariates have the same distribution. Then

$$
\mathbb{E}\left[\sum_{i \neq i^{\prime} \in I_{2}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right] \leq \mathbb{E}\left[\sum_{i \in I_{2}} \sum_{i^{\prime} \in[n], i^{\prime} \neq i} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right]
$$

and we only need to upper bound each quantity $\mathbb{E}\left[\mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right]$ for $i \in I_{2}, i^{\prime} \in[n] \backslash\{i\}$, and $j \in[n]$. By the i.i.d. assumption, the covariates $X_{1}^{1}, \cdots, X_{n}^{1}, X_{i^{\prime}}^{0}$ are i.i.d. conditioning on $X_{i}^{0}$. This observation motivates us to consider the following problem: let $Y_{0}=x \in[0,1]^{d}$ be any fixed point, and $Y_{1}, \cdots, Y_{n+1}$ be i.i.d. distributed as $X_{1}^{1}$. For $i \in[n+1] \cup\{0\}$ and $j \in[n+1]$ with $i \neq j$, let $d(i, j) \in[n]$ be the integer $d$ such that $Y_{j}$ is the $d$-th nearest neighbor of $Y_{i}$ among $\left\{Y_{1}, \cdots, Y_{i-1}, Y_{i+1}, \cdots, Y_{n+1}\right\}$. Consequently, identifying $X_{i}^{0}, X_{i^{\prime}}^{0}, X_{j}^{1}$ as $Y_{0}, Y_{n+1}$, and $Y_{1}$ respectively, the i.i.d. assumption gives

$$
\mathbb{E}\left[\mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right) \mid X_{i}^{0}=x\right] \leq \mathbb{E}[\mathbb{1}(d(0,1) \leq 2) \mathbb{1}(d(n+1,1)=1)]
$$

To upper bound the RHS, note that the exchangeability of $\left(X_{1}, \cdots, X_{n+1}\right)$ yields

$$
\begin{aligned}
\mathbb{E}[\mathbb{1}(d(0,1) \leq 2) \mathbb{1}(d(n+1,1)=1)] & =\frac{1}{n(n+1)} \sum_{i \neq j \in[n+1]} \mathbb{E}[\mathbb{1}(d(0, i) \leq 2) \mathbb{1}(d(j, i)=1)] \\
& =\frac{1}{n(n+1)} \mathbb{E}\left[\sum_{i=1}^{n+1} \mathbb{1}(d(0, i) \leq 2) \sum_{j \neq i} \mathbb{1}(d(j, i)=1)\right] \\
& \stackrel{(\text { a })}{\leq} \frac{1}{n(n+1)} \mathbb{E}\left[\sum_{i=1}^{n+1} \mathbb{1}(d(0, i) \leq 2) \cdot c_{d}\right] \\
& \stackrel{(\mathrm{b})}{=} \frac{2 c_{d}}{n(n+1)}
\end{aligned}
$$

where (a) follows from the deterministic inequality $\sum_{j \neq i} \mathbb{1}(d(j, i)=1) \leq c_{d}$ in GKOV17] Lemma C.1], i.e. each $X_{j}$ can only be the nearest neighbor of at most $c_{d}$ points, with constant $c_{d}$ depending only on $d$. The identity (b) trivially follows from the definition of nearest-neighbors and $d(0, i)$. Consequently, combining the above displays and (13), we arrive at

$$
\mathbb{E}\left[\sum_{j \in[n]} S_{j}^{2}\right] \leq\left(1+2 c_{d}\right) m_{2}
$$

and therefore

$$
\begin{equation*}
\mathbb{E}[|s(X)|] \leq \sqrt{\mathbb{E}\left[s(X)^{2}\right]} \lesssim \frac{\sigma}{\sqrt{m_{2}}} \tag{14}
\end{equation*}
$$

In the general case, as $I_{2}$ is essentially a uniformly random subset of $I_{1}$, we have

$$
\begin{aligned}
\mathbb{E}\left[\sum_{i \neq i^{\prime} \in I_{2}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right] & \lesssim\left(\frac{m_{2}}{m_{1}}\right)^{2} \mathbb{E}\left[\sum_{i \neq i^{\prime} \in I_{1}} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right] \\
& \leq\left(\frac{m_{2}}{m_{1}}\right)^{2} \mathbb{E}\left[\sum_{i \in I_{1}} \sum_{i^{\prime} \in[n], i^{\prime} \neq i} \sum_{j \in[n]} \mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right)\right]
\end{aligned}
$$

Moreover, by the likelihood ratio assumption, a simple change of measure also gives

$$
\mathbb{E}\left[\mathbb{1}(i \rightarrow j) \mathbb{1}\left(i^{\prime} \rightarrow j\right) \mid X_{i}^{0}=x\right] \leq \kappa \cdot \mathbb{E}[\mathbb{1}(d(0,1) \leq 2) \mathbb{1}(d(n+1,1)=1)]
$$

Consequently, as $m_{1} \geq \kappa m_{2}$, we again arrive at the same upper bound 14 , and the claimed Theorem 5 follows from (11), (12), and (14).

## B.2.2 Proof of Theorem 6

We first show that to prove the lower bound of the minimax rate, it suffices to consider the case where both $g_{0}$ and $g_{1}$ are uniform over $[0,1]^{d}$, and there are $n$ and $m=n / \kappa$ (assumed to be an integer as $\kappa \leq n$ ) observations for the control and treatment groups, respectively. To see this, note that we may choose $g_{0}$ to be the uniform distribution on $[0,1]^{d}$, and $g_{1}$ be $1 / \kappa$ in a smaller cube and $\kappa$ otherwise. Note that the smaller cube has the volume $\kappa /(\kappa+1)$, which is at least $\Omega(1)$. Besides, the conditional distribution of $g_{1}$ restricted to the smaller cube is again uniform, and with high probability there are $O(n)$ control observations and $O(n / \kappa)$ treatment observations in the smaller cube. Consequently, restricting everything including the minimax risk to the smaller cube and using a proper scaling of constant factor, the desired reduction holds for the purpose of proving the lower bound. Furthermore, in this case the $L_{1}\left(g_{0}\right)$ norm coincides with the usual $L_{1}$ norm.
As in the proof of Theorem 4 , after choosing $\mu \equiv 0$, the lower bound $\Omega\left(\left(\sigma^{2} / m\right)^{\beta_{\tau} /\left(2 \beta_{\tau}+d\right)}\right)$ follows directly from the known results in nonparametric estimation. Hence, using the relationship $m=n / \kappa$, the last error term of Theorem 6 is established. For the other two errors, we will need to construct appropriate functions in $\mathcal{H}_{d}\left(\beta_{\mu}\right)$ and $\mathcal{H}_{d}\left(\beta_{\tau}\right)$, and by multiplying a smooth function supported on
$[0,1]^{d}$ (as in the proof of Theorem 4 we will not consider the boundary effects in the construction below.
First we prove the lower bound $\widetilde{\Omega}\left(\left(\kappa / n^{2}\right)^{1 / d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}\right)$ via providing another feasible solution to the optimization problem (6). Specifically, let $\mu_{0}, \tau$ be the Hölder smooth functions with their respective smoothness parameters $\beta_{\mu}, \beta_{\tau}$ assured by Lemma 2 based on the following value specifications:

$$
\mu_{0}\left(X_{i}^{0}\right)=0, \quad \mu_{0}\left(X_{i}^{1}\right)=-\tau\left(X_{i}^{1}\right)=c \cdot \min _{j \in[m]}\left(\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}+\min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}}\right)
$$

for all $i \in[m]$, where $c>0$ is a small constant. However, before applying Lemma 2 , we need to show that the conditions of Lemma 2 hold for the above value specifications.
For the HTE function $\tau$, using the Lipschitz property of the minimum $\left|\min _{k} a_{k}-\min _{k} b_{k}\right| \leq$ $\max _{k}\left|a_{k}-b_{k}\right|$, we have

$$
\begin{aligned}
& \left|\tau\left(X_{i}^{1}\right)-\tau\left(X_{i^{\prime}}^{1}\right)\right| \\
& \leq c \cdot \max _{j \in[m]}\left|\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}+\min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}}-\left(\left\|X_{i^{\prime}}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}+\min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}}\right)\right| \\
& =c \cdot \max _{j \in[m]}\left|\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}-\left\|X_{i^{\prime}}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}\right| \\
& \leq c \cdot\left\|X_{i}^{1}-X_{i^{\prime}}^{1}\right\|_{2}^{\beta_{\tau}}
\end{aligned}
$$

where the last step follows from $\left||a|^{\beta_{\tau}}-|b|^{\beta_{\tau}}\right| \leq|a-b|^{\beta_{\tau}}$ for all $\beta_{\tau} \in[0,1]$. Hence, for constant $c>0$ sufficiently small, the condition of Lemma 2 holds for $\tau$.
For the baseline function $\mu_{0}$, using the same analysis we have $\left|\mu_{0}\left(X_{i}^{1}\right)-\mu_{0}\left(X_{i^{\prime}}^{1}\right)\right|=O\left(\| X_{i}^{1}-\right.$ $\left.X_{i^{\prime}}^{1} \|_{2}^{\beta_{\tau}}\right)=O\left(\left\|X_{i}^{1}-X_{i^{\prime}}^{1}\right\|_{2}^{\beta_{\mu}}\right)$ for the treatment-treatment pairs. However, for $\mu_{0}$ we also need to verify the condition of Lemma 2 for control-control and treatment-control pairs. The the controlcontrol pairs are easy to verify: $\left|\mu_{0}\left(X_{i}^{0}\right)-\mu_{0}\left(X_{i^{\prime}}^{0}\right)\right|=0$ always holds. As for the treatment-control pairs, we have

$$
\begin{aligned}
\left|\mu_{0}\left(X_{i}^{1}\right)-\mu_{1}\left(X_{i^{\prime}}^{0}\right)\right| & =\left|\mu_{0}\left(X_{i}^{1}\right)\right| \\
& =c \cdot \min _{j \in[m]}\left(\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}+\min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}}\right) \\
& \leq c \cdot \min _{k \in[n]}\left\|X_{i}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}} \\
& \leq c \cdot\left\|X_{i}^{1}-X_{i^{\prime}}^{0}\right\|_{2}^{\beta_{\mu}}
\end{aligned}
$$

as desired. Hence, both functions $\mu_{0}$ and $\tau$ can be extended to Hölder smooth functions on $[0,1]^{d}$ with desired smoothness parameters, and $\left(\mu_{0}, \tau\right)$ is a feasible solution to the optimization problem (6).

Next we provide a lower bound of $\|\tau\|_{1}$ for the above construction of $\tau$. Choose a bandwidth

$$
h_{1}=\widetilde{\Theta}\left((m n)^{-\frac{\beta_{\mu}}{d\left(\beta_{\tau}+\beta_{\mu}\right)}}\right)
$$

and define

$$
I=\left\{i \in[m]: \min _{j \in[n]}\left\|X_{i}^{1}-X_{j}^{0}\right\|_{2} \leq \frac{1}{h_{1}}\left(\frac{1}{m n \log n}\right)^{\frac{1}{d}}\right\}
$$

The next lemma shows that the cardinality of $I$ is upper bounded by $1 /\left(3 h^{d}\right)$ with high probability.
Lemma 6. If $h_{1} \leq c(\log n)^{-2 / d}$ with constant $c>0$ small enough, then with probability at least $1-n^{-3}$ we have $|I| \leq 1 /\left(3 h^{d}\right)$ for sufficiently large $n$.

Proof. As the control covariates follow the uniform distribution, for any $x \in[0,1]^{d}$ we have

$$
\mathbb{P}\left(\left\|x-X_{j}^{0}\right\|_{2} \geq \frac{1}{h_{1}(m n \log n)^{1 / d}}\right) \geq 1-v_{d}\left(\frac{1}{h_{1}(m n \log n)^{1 / d}}\right)^{d}=1-\frac{v_{d}}{m n h_{1}^{d} \log n}
$$

where $v_{d}$ is the volume of the unit ball in $d$ dimensions. Hence, for each individual $i \in[m]$ in the treatment group, we have
$\mathbb{P}\left(E_{i}\right) \triangleq \mathbb{P}\left(\min _{j \in[n]}\left\|X_{i}^{1}-X_{j}^{0}\right\|_{2} \geq \frac{1}{h_{1}(m n \log n)^{1 / d}}\right) \geq\left(1-\frac{v_{d}}{m n h_{1}^{d} \log n}\right)^{n} \geq 1-\frac{v_{d}}{m h_{1}^{d} \log n}$.
Hence, the target quantity $\sum_{i=1}^{m} \mathbb{1}\left(E_{i}^{c}\right)$ is a Binomial random variable with number of observations $m$ and $\mathbb{P}\left(E_{i}^{c}\right) \leq v_{d} /\left(m h_{1}^{d} \log n\right)$. Using the first inequality of Lemma 4 with $\delta=1, \lambda=m \mathbb{P}\left(E_{i}^{c}\right)$ gives

$$
\mathbb{P}\left(\sum_{i=1}^{m} \mathbb{1}\left(E_{i}^{c}\right) \geq \frac{2 v_{d}}{h_{1}^{d} \log n}\right) \leq \exp \left(-\frac{v_{d}}{3 h_{1}^{d} \log n}\right)
$$

Therefore, by using $h_{1} \leq c(\log n)^{-2 / d}$ with a small enough $c>0$, the claimed results hold for sufficiently large $n$.

By Lemma 6 we know that $|I| \leq 1 /\left(3 h_{1}^{d}\right)$ with probability at least $1-n^{-3}$. We also define the set of good indices as follows:

$$
I^{\star}=\left\{i \in[m]: \min _{j \in I}\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}>\frac{h_{1}}{2}\right\}
$$

We claim that for any $i \in I^{\star}$, we have $\left|\tau\left(X_{i}^{1}\right)\right|=\widetilde{\Omega}\left(h_{1}^{\beta_{\tau}}\right)$. In fact, for any $j \notin I$, the definition of $I$ gives

$$
\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}+\min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}} \geq \min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}}>\left(\frac{1}{h_{1}(m n \log n)^{1 / d}}\right)^{\beta_{\mu}}=\widetilde{\Omega}\left(h_{1}^{\beta_{\tau}}\right)
$$

where the last identity follows from the choice of $h_{1}$. As for the other case $j \in I$, the definition of $I^{\star}$ gives

$$
\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}+\min _{k \in[n]}\left\|X_{j}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}} \geq\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}} \geq \min _{j \in I}\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\tau}}=\Omega\left(h_{1}^{\beta_{\tau}}\right)
$$

as claimed. Hence, all $\left|\tau\left(X_{i}^{1}\right)\right|$ with index $i \in I^{\star}$ have a large magnitude, and it remains to show that $\left|I^{\star}\right|$ is large with high probability. In fact, conditioning on $|I| \leq 1 /\left(3 h_{1}^{d}\right)$, the following set

$$
A:=\left\{x \in[0,1]^{d}: \min _{j \in I}\left\|X_{i}^{1}-x\right\|_{2} \leq \frac{h_{1}}{2}\right\}
$$

has $d$-dimensional volume at most $|I| \cdot h_{1}^{d} \leq 1 / 3$. Note that $X_{i}^{1}$ is uniformly distributed, the region $A$ has probability at most $1 / 3$. Hence, by the standard Binomial tail bound (cf. Lemma 4), we conclude that with probability at least $1-2 n^{-3}$, the number of treatment covariates $X_{i}^{1}$ falling into $A$ is at most $m / 2$. Consequently, we have $\left|I^{\star}\right| \geq m / 2$. In other words, with high probability, at least a constant fraction of $X_{i}^{1}$ satisfies $\left|\tau\left(X_{i}^{1}\right)\right|=\widetilde{\Omega}\left(h_{1}^{\beta_{\tau}}\right)$, and therefore $\|\tau\|_{1}=\widetilde{\Omega}\left(h_{1}^{\beta_{\tau}}\right)$ with high probability as well. Now the two-point method applied to $\left(\mu_{0}, \tau\right)$ and $\left(\mu_{0}^{\prime} \equiv 0, \tau^{\prime} \equiv 0\right)$ together with the choice of $h_{1}$ gives the desired lower bound $\widetilde{\Omega}\left(\left(\kappa / n^{2}\right)^{1 / d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}\right)$.
Next we prove the final lower bound $\widetilde{\Omega}\left(\left(\kappa \sigma^{2} / n^{2}\right)^{1 /\left(2+d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)\right)}\right)$. To this end, we construct exponentially many hypotheses $\left(\mu_{0}^{v}, \tau^{v}\right)$ indexed by $v \in\{ \pm 1\}^{M}$ and apply Fano's inequality. Specifically, let $h_{1}, h_{2}$ be the bandwidths given as follows:

$$
\begin{aligned}
& h_{1}=\widetilde{\Theta}\left(\frac{1}{(m n)^{1 / d}}\left(\frac{\sigma}{\sqrt{m n}}\right)^{-\frac{2 \beta_{\tau}}{2 \beta_{\mu} \beta_{\tau}+d\left(\beta_{\mu}+\beta_{\tau}\right)}}\right) \\
& h_{2}=\widetilde{\Theta}\left(\left(\frac{\sigma}{\sqrt{m n}}\right)^{\frac{2 \beta_{\mu}}{2 \beta_{\mu} \beta_{\tau}+d\left(\beta_{\mu}+\beta_{\tau}\right)}}\right)
\end{aligned}
$$

and $M=h_{2}^{-d}$. Fix any smooth function $g$ supported on $[0,1]^{d}$, then the standard dilation analysis yields that

$$
\tau^{v}(x)=c \cdot h_{2}^{\beta_{\tau}} \sum_{i=1}^{M} v_{i} g\left(\frac{x-x_{i}}{h_{2}}\right)
$$

for sufficiently small constant $c>0$ belongs to the Hölder ball $\mathcal{H}_{d}\left(\beta_{\tau}\right)$ for all $v \in\{ \pm 1\}^{M}$, where $x_{1}, \cdots, x_{M}$ are vertices of the small cubes so that $\left\{x_{1}, \cdots, x_{M}\right\}+\left[0, h_{2}\right]^{d}=[0,1]^{d}$. Note that

$$
\left\|\tau^{v}-\tau^{v^{\prime}}\right\|_{1}=c h_{2}^{\beta_{\tau}}\|g\|_{1} \cdot d_{\mathrm{H}}\left(v, v^{\prime}\right)
$$

where $d_{\mathrm{H}}\left(v, v^{\prime}\right)=M^{-1} \sum_{i=1}^{M} \mathbb{1}\left(v_{i} \neq v_{i}^{\prime}\right)$ is the normalized Hamming distance between binary vectors $v$ and $v^{\prime}$. By the Gilbert-Varshamov bound, there exists $\mathcal{V}_{0} \subseteq\{ \pm 1\}^{M}$ with $\left|\mathcal{V}_{0}\right| \geq 2^{M / 8}$ such that $d_{\mathrm{H}}\left(v, v^{\prime}\right) \geq 1 / 5$ for all $v, v^{\prime} \in \mathcal{V}_{0}$. Let $V$ be a random variable uniformly distributed in $\mathcal{V}_{0}$, and $Y$ be the collection of all outcomes in (1) based on $\tau^{v}$ and $\mu_{0}^{v}$ defined below. The standard Fano-type arguments (see, e.g. [Tsy09]) give

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\mathrm{random}}(\kappa)=\Omega\left(h_{2}^{\beta_{\tau}}\left(1-\frac{I(V ; Y)+\log 2}{M / 8}\right)\right) .
$$

Next we give an explicit construction of $\mu_{0}^{v}$ for each $v \in\{ \pm 1\}^{M}$, and upper bound the mutual information $I(V ; Y)$.

The baseline functions $\mu_{0}^{v}$, also indexed by $v \in\{ \pm 1\}^{M}$, are chosen to offset the differences of $\tau^{v}$ on the treatment covariates. Motivated by Lemma 2 we choose the reference function

$$
\tilde{\mu}_{0}\left(X_{i}^{1}\right)=c \cdot \min _{j \in[m]}\left(\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2}^{\beta_{\mu}}+\min _{k \in[n]}\left\|X_{i}^{1}-X_{k}^{0}\right\|_{2}^{\beta_{\mu}}\right)
$$

and

$$
\mu_{0}^{v}\left(X_{i}^{1}\right)=-\operatorname{sign}\left(\tau^{v}\left(X_{i}^{1}\right)\right) \cdot\left(\left|\tau^{v}\left(X_{i}^{1}\right)\right| \wedge \tilde{\mu}_{0}\left(X_{i}^{1}\right)\right),
$$

and $\mu_{0}^{v}\left(X_{i}^{0}\right)=0$ for any $i \in[n]$. As before, using similar arguments we conclude that the above value specifications satisfy the conditions of Lemma 2, and thus each $\mu_{0}^{v}$ can be extended to a function in $\mathcal{H}_{d}\left(\beta_{\mu}\right)$ for each $v$. Since distributions of the control outcomes do not depend on $v$, $\mu_{0}^{v}\left(X_{i}^{1}\right)+\tau^{v}\left(X_{i}^{1}\right)=0$ whenever $\tilde{\mu}_{0}\left(X_{i}^{1}\right) \geq c h_{2}^{\beta_{\tau}}\|g\|_{\infty} \geq\left|\tau^{v}\left(X_{i}^{1}\right)\right|$, and all noises are normal distributed with variance $\sigma^{2}$, we have

$$
\begin{aligned}
I(V ; Y) & \leq \max _{v \neq v^{\prime} \in \mathcal{V}_{0}} D_{\mathrm{KL}}\left(P_{Y \mid V=v} \| P_{Y \mid V=v^{\prime}}\right) \\
& =\max _{v \neq v^{\prime} \in \mathcal{V}_{0}} \sum_{i=1}^{m} \frac{\left(\mu_{0}^{v}\left(X_{i}^{1}\right)+\tau_{0}^{v}\left(X_{i}^{1}\right)-\mu_{0}^{v^{\prime}}\left(X_{i}^{1}\right)-\tau_{0}^{v^{\prime}}\left(X_{i}^{1}\right)\right)^{2}}{2 \sigma^{2}} \\
& \leq 4 c^{2} h_{2}^{2 \beta_{\tau}}\|g\|_{\infty}^{2} \cdot \sum_{i=1}^{m} \mathbb{1}\left(\tilde{\mu}_{0}\left(X_{i}^{1}\right) \leq c h_{2}^{\beta_{\tau}}\|g\|_{\infty}\right) .
\end{aligned}
$$

By the choices of $h_{1}$ and $h_{2}$ above, we have $h_{2}^{\beta_{\tau}}=\widetilde{\Omega}\left(\left(h_{1}(m n)^{1 / d}\right)^{-\beta_{\mu}}\right)$. Hence, similar arguments of Lemma6lead to

$$
\left|\left\{i \in[m]: \min _{j \in[n]}\left\|X_{i}^{1}-X_{j}^{0}\right\|_{2} \wedge \min _{j \neq i}\left\|X_{i}^{1}-X_{j}^{1}\right\|_{2} \leq \frac{1}{h_{1}}\left(\frac{1}{m n \log n}\right)^{\frac{1}{d}}\right\}\right|=O\left(h_{1}^{-d}\right)
$$

with probability at least $1-n^{-3}$. Also, for any $i \in[m]$ not belonging to the above set of indices, it is straightforward to see that $\tilde{\mu}_{0}\left(X_{i}^{1}\right)=\widetilde{\Omega}\left(\left(h_{1} n^{2 / d}\right)^{-\beta_{\mu}}\right)$. Hence, with probability at least $1-n^{-3}$ we have $\sum_{i=1}^{n} \mathbb{1}\left(\tilde{\mu}_{0}\left(X_{i}^{1}\right) \leq c h_{2}^{\beta_{\tau}}\|g\|_{\infty}\right)=O\left(h_{1}^{-d}\right)$. Therefore,

$$
I(V ; Y)=O\left(\frac{h_{2}^{2 \beta_{\tau}}}{h_{1}^{d} \sigma^{2}}\right)=O\left(h_{2}^{-d}\right)=O(M)
$$

holds with high probability over the random covariates. Plugging the above upper bound of the mutual information into the Fano's inequality, we conclude that

$$
R_{n, d, \beta_{\mu}, \beta_{\tau}, \sigma}^{\mathrm{random}}(\kappa)=\widetilde{\Omega}\left(h_{2}^{\beta_{\tau}}\right)=\widetilde{\Omega}\left(\left(\frac{\kappa \sigma^{2}}{n^{2}}\right)^{\frac{1}{2+d\left(\beta_{\mu}^{-1}+\beta_{\tau}^{-1}\right)}}\right),
$$

as desired.

## C Proof of Main Lemmas

## C. 1 Proof of Lemma 1

We first establish the following inequality: for any $\varepsilon>0$, there exists a constant $C_{d}>0$ depending only on $d$ such that

$$
\begin{equation*}
\int_{[0,1]^{d}} f(x) \mathbb{1}(m[f](x) \leq \varepsilon) d x \leq C_{d} \varepsilon \tag{15}
\end{equation*}
$$

To establish (15), we apply the maximal inequality in Lemma 5 to the set $A=[-\sqrt{d}, 1+\sqrt{d}]^{d}$ and measures $\mu_{1}(d x)=f(x) \lambda(d x), \mu_{2}(d x)=\lambda(d x)$, with $\lambda$ being the $d$-dimensional Lebesgue measure. Then clearly for every $x \in[0,1]^{d}$, we have

$$
\sup _{\rho>0: B(x ; \rho) \subseteq A}\left(\frac{\mu_{2}(B(x ; \rho))}{\mu_{1}(B(x ; \rho))}\right) \geq \frac{1}{m[f](x)} .
$$

Consequently, choosing $t=1 / \varepsilon$ in Lemma 5 gives

$$
\begin{aligned}
\int_{[0,1]^{d}} f(x) \mathbb{1}(m[f](x) \leq \varepsilon) d x & \leq \mu_{1}\left(\left\{x \in A: \sup _{\rho>0: B(x ; \rho) \subseteq A}\left(\frac{\mu_{2}(B(x ; \rho))}{\mu_{1}(B(x ; \rho))}\right) \geq \frac{1}{\varepsilon}\right\}\right) \\
& \leq C_{d}^{\prime} \varepsilon \cdot \mu_{2}(A)=C_{d} \varepsilon
\end{aligned}
$$

establishing (15).
Next we make use of (15) to prove Lemma 1 For $\lambda>e C_{d}$, let $\varepsilon$ be a random variable following an exponential distribution with rate $\lambda$. Then clearly $\mathbb{P}(m[f](x) \leq \varepsilon)=\exp (-\lambda \cdot m[f](x))$, and taking expectation at both sides of (15) leads to

$$
\int_{[0,1]^{d}} f(x) \exp (-\lambda \cdot m[f](x)) d x \leq C_{d} \cdot \mathbb{E}[\varepsilon]=\frac{C_{d}}{\lambda}
$$

as claimed. For $\lambda \leq e C_{d}$, note that applying the above arguments gives that for any $t \geq 1$,

$$
\int_{[0,1]^{d}} f(x) \exp (-t \lambda \cdot m[f](x)) d x \leq \frac{C_{d}}{t \lambda}
$$

Hence, using the fact that $f$ is a density and the Hölder's inequality, we have

$$
\begin{aligned}
\int_{[0,1]^{d}} f(x) \exp (-\lambda \cdot m[f](x)) d x & \leq \inf _{t \geq 1}\left(\int_{[0,1]^{d}} f(x) \exp (-t \lambda \cdot m[f](x)) d x\right)^{\frac{1}{t}} \\
& \leq \inf _{t \geq 1}\left(\frac{C_{d}}{t \lambda}\right)^{\frac{1}{t}}=\exp \left(-\frac{\lambda}{e C_{d}}\right)
\end{aligned}
$$

where the minimizer in the last step is $t=e C_{d} / \lambda \geq 1$.

## C. 2 Proof of Lemma 2

The necessity result simply follows from the Hölder ball condition, and it remains to prove the sufficiency. Without loss of generality we assume that $L=1$. We first show that given $\left\{\left(x_{i}, y_{i}\right)\right\}_{i \in[n]}$ satisfying the condition of Lemma 2 , for any other $x \in \mathbb{R}^{d} \backslash\left\{x_{i}\right\}_{i \in[n]}$ there exists a value specification $y \in \mathbb{R}$ of $f(x)$ such that the same condition holds for the $(n+1)$ points $\left\{\left(x_{i}, y_{i}\right)\right\}_{i \in[n]} \cup\{(x, y)\}$, i.e.

$$
\left|y_{i}-y\right| \leq\left\|x_{i}-x\right\|_{2}^{\beta}, \quad \forall i \in[n] .
$$

To prove the existence of $y$, it suffices to show that

$$
\min _{i \in[n]}\left\{y_{i}+\left\|x_{i}-x\right\|_{2}^{\beta}\right\} \geq \max _{j \in[n]}\left\{y_{j}-\left\|x_{j}-x\right\|_{2}^{\beta}\right\} .
$$

In fact, for $i, j \in[n]$,

$$
\begin{aligned}
\left(y_{i}+\left\|x_{i}-x\right\|_{2}^{\beta}\right)-\left(y_{j}-\left\|x_{j}-x\right\|_{2}^{\beta}\right) & =\left(y_{i}-y_{j}\right)+\left\|x_{i}-x\right\|_{2}^{\beta}+\left\|x_{j}-x\right\|_{2}^{\beta} \\
& \geq-\left\|x_{i}-x_{j}\right\|_{2}^{\beta}+\left\|x_{i}-x\right\|_{2}^{\beta}+\left\|x_{j}-x\right\|_{2}^{\beta} \\
& \geq 0
\end{aligned}
$$

Hence, we may construct the $(n+1)$-th point based on $n$ points, and repeating this procedure leads to a construction of $f$ on any countable dense subset $\mathcal{B} \subseteq \mathbb{R}^{d}$. Now the proof is completed by taking

$$
f(x)=\lim _{x_{n} \in \mathcal{B}, x_{n} \rightarrow x} f\left(x_{n}\right)
$$

## C. 3 Proof of Lemma 3

For $t=0$, the lemma is obvious. In the following we consider $t \geq 1$ and write the linear equations in the matrix form

$$
\left[\begin{array}{cccc}
1 & 1 & \cdots & 1 \\
z_{0} & z_{1} & \cdots & z_{t} \\
\vdots & \vdots & \vdots & \vdots \\
z_{0}^{t} & z_{1}^{t} & \cdots & z_{t}^{t}
\end{array}\right]\left[\begin{array}{c}
w_{0} \\
w_{1} \\
\vdots \\
w_{t}
\end{array}\right]=\left[\begin{array}{c}
1 \\
0 \\
\vdots \\
0
\end{array}\right] .
$$

where $z_{j}=x-x_{j}, 0 \leq j \leq t$. Since the above square matrix is a Vandermonde matrix, the weight vector $\left(w_{0}, \cdots, w_{t}\right)$ exists and is unique. Now by the Cramer's rule, for $i \geq 1$ we have

$$
\begin{aligned}
\left|w_{i}\right| & \left.=\left|\frac{\operatorname{det}\left[\begin{array}{ccccc}
1 & \cdots & 1 & \cdots & 1 \\
z_{0} & \cdots & 0 & \cdots & z_{t} \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
z_{0}^{t} & \cdots & 0 & \cdots & z_{t}^{t}
\end{array}\right]}{\left.\operatorname{det}\left[\begin{array}{ccccc}
1 & \cdots & 1 & \cdots & 1 \\
z_{0} & \cdots & z_{i} & \cdots & z_{t} \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
z_{0}^{t} & \cdots & z_{i}^{t} & \cdots & z_{t}^{t}
\end{array}\right] \right\rvert\,} \prod_{0 \leq j \leq t, j \neq i}\right| \frac{z_{j}}{z_{j}-z_{i}} \right\rvert\, \\
& =\frac{\Delta}{\left|z_{0}-z_{i}\right|} \cdot \prod_{1 \leq j \leq t, j \neq i}\left|\frac{z_{j}}{z_{j}-z_{i}}\right| \leq m \Delta \cdot(t+1)^{t-1},
\end{aligned}
$$

where the last inequality follows from $\left|z_{j}\right| \leq(t+1) / m$ and $\left|z_{j}-z_{i}\right| \geq 1 / m$ for $i \neq j$. The proof is completed by noting that

$$
\left|w_{0}\right| \leq 1+\sum_{i=1}^{t}\left|w_{i}\right| \leq 1+m \Delta \cdot t(t+1)^{t-1} \leq 1+t(t+1)^{t-1}=: C
$$

