

Optimal Locally Private Nonparametric Classification with Public Data

Yuheng Ma

YMA@RUC.EDU.CN

School of Statistics

Renmin University of China

100872 Beijing, China

Hanfang Yang

HYANG@RUC.EDU.CN

Center for Applied Statistics, School of Statistics

Renmin University of China

100872 Beijing, China

Editor: Po-Ling Loh

Abstract

In this work, we investigate the problem of public data assisted non-interactive Local Differentially Private (LDP) learning with a focus on non-parametric classification. Under the posterior drift assumption, we for the first time derive the mini-max optimal convergence rate with LDP constraint. Then, we present a novel approach, the locally differentially private classification tree, which attains the mini-max optimal convergence rate. Furthermore, we design a data-driven pruning procedure that avoids parameter tuning and provides a fast converging estimator. Comprehensive experiments conducted on synthetic and real data sets show the superior performance of our proposed methods. Both our theoretical and experimental findings demonstrate the effectiveness of public data compared to private data, which leads to practical suggestions for prioritizing non-private data collection.

Keywords: Decision Tree, Local Differential Privacy, Non-parametric Classification, Posterior Drift, Public Data

1 Introduction

Local differential privacy (LDP) (Kairouz et al., 2014; Duchi et al., 2018), which is a variant of differential privacy (DP) (Dwork et al., 2006), has gained considerable attention in recent years, particularly among industrial developers (Erlingsson et al., 2014; Tang et al., 2017). Unlike central DP, which relies on a trusted curator who has access to the raw data, LDP assumes that each sample is possessed by a data holder, and each holder privatizes their data before it is collected by the curator. Although this setting provides stronger protection, learning with perturbed data requires more samples (Duchi et al., 2018) compared to the central setting. Moreover, basic techniques such as principal component analysis (Wang and Xu, 2020; Bie et al., 2022), data standardization (Bie et al., 2022), and tree partitioning (Wu et al., 2022) are troublesome or even prohibited. Consequently, LDP introduces challenges for various machine learning tasks that are otherwise considered straightforward, including

density estimation (Duchi et al., 2018), mean estimation (Duchi et al., 2018), Gaussian estimation (Joseph et al., 2019), and change-point detection (Berrett and Yu, 2021).

Fortunately, in certain scenarios, private estimation performance can be enhanced with an additional public data set. From an empirical perspective, numerous studies demonstrated the effectiveness of using public data (Papernot et al., 2017, 2018; Yu et al., 2021b, 2022; Nasr et al., 2023; Gu et al., 2023). In most cases, the public data is out-of-distribution and collected from other related sources. Yet, though public data from an identical distribution has been extensively studied (Bassily et al., 2018; Alon et al., 2019; Kairouz et al., 2021; Amid et al., 2022; Ben-David et al., 2023; Lowy et al., 2023; Wang et al., 2023a), only a few works systematically described the out-of-distribution relationship. Focusing on unsupervised problems, i.e. Gaussian estimation, Bie et al. (2022) discussed the gain of using public data with respect to the total variation distance between private and public data distributions. As for supervised learning, Gu et al. (2023) discussed the similarities between data sets to guide the selection of appropriate public data. Ma et al. (2023) leveraged public data to create informative partitions for LDP regression trees. Neither of them discussed the relationship between regression functions of private and public data distributions. Aiming at the gap between theory and practice in supervised private learning with public data, we pose the following intriguing questions:

1. *Theoretically, when is labeled public data beneficial for LDP learning?*
2. *Under such conditions, how to effectively leverage public data?*

Hyper-parameter tuning is an essential yet challenging task in private learning (Papernot and Steinke, 2021; Mohapatra et al., 2022) and remains an unsolved problem in local differential privacy. Common strategies, such as cross-validation and information criteria, remain unavailable due to their requirement for multiple queries to the training data. Tuning by omitting a validation set is restrictive in terms of the size of the potential hyperparameter space (Ma et al., 2023). Moreover, leveraging public data necessitates more complex models, which often results in an increased number of hyperparameters to tune. Consequently, models with public data may face difficulties in parameter selection, which leads to the third consideration:

3. *Can data-driven hyperparameter tuning approaches be derived?*

In this work, we answer the above questions from the perspective of non-parametric classification with the non-interactive LDP constraint. Though parametric methods enjoy better privacy-utility trade-off (Duchi et al., 2018) under certain assumptions, they are vulnerable to model misspecification. Under LDP setting where we have limited prior knowledge of the data, non-parametric methods ensure the worst-case performance. As for private classification, the local setting remains rarely explored compared to the central setting. A notable reason is that most gradient-based methods (Song et al., 2013; Abadi et al., 2016) are prohibited due to the high demand for memory and communication capacity on the terminal machine (Tramèr et al., 2022). We consider non-interactive LDP, where the communication between the curator and the data holders is limited to a single round. This type of method is favored by practitioners (Smith et al., 2017; Zheng et al., 2017; Daniely

and Feldman, 2019) since they avoid multi-round protocols that are prohibitively slow in practice due to network latency.

Under such background, we summarize our contributions by answering the questions:

- For the first question, we propose to adopt the framework of *posterior drift*, a setting in transfer learning (Cai and Wei, 2021), for our analysis. Specifically, given the target distribution P and the source distribution Q , we require $P_X = Q_X$, i.e. they have the same marginal distributions. Moreover, their Bayes decision rule are identical, i.e. $(P(Y = 1|X) - 1/2) \cdot (Q(Y = 1|X) - 1/2) \geq 0$. The assumption covers a wide range of distributions and includes $P = Q$ as a special case. Theoretically, we for the first time establish a mini-max lower bound for nonparametric LDP learning with public data.
- To answer the second question, we propose the *Locally differentially Private Classification Tree* (LPCT), a novel algorithm that leverages both public and private data. Specifically, we first create a tree partition on public data and then estimate the regression function through a weighted average of public and private data. Besides inheriting the merits of tree-based models such as efficiency, interpretability, and extensivity to multiple data types, LPCT is superior from both theoretical and empirical perspectives. Theoretically, we show that LPCT attains the mini-max optimal convergence rate. Empirically, we conduct experiments on both synthetic and real data sets to show the effectiveness of LPCT over state-of-the-art competitors.
- To answer the last question, we propose the *Pruned Locally differentially Private Classification Tree* (LPCT-prune), a data-driven classifier that is free of parameter tuning. Specifically, we query the privatized information with a sufficient tree depth and conduct an efficient pruning procedure. Theoretically, we show that LPCT-prune achieves the optimal convergence rate in most cases and maintains a reasonable convergence rate otherwise. Empirically, we illustrate that LPCT-prune with a default parameter performs comparable to LPCT with the best parameters.

This article is organized as follows. In Section 2, we provide a comprehensive literature review. Then we formally define our problem setting and introduce our methodology in Section 3. The obtained theoretical results are presented in Section 4. Then we develop the data-driven estimator in Section 5. Extensive experiments are conducted on synthetic data and real data in Section 6 and 7 respectively. The conclusion and discussions are in Section 8.

2 Related Work

2.1 Private Learning with Public Data

Leveraging public data in private learning is a research topic that is both theoretically and practically meaningful to closing the gap between private and non-private methods. There is a long list of works addressing the issue of private learning with public data from a central differential privacy perspective. Through unlabeled public data, Papernot et al. (2017, 2018) fed knowledge privately into student models, whose theoretical benefits are established by Liu et al. (2021). Empirical investigations have demonstrated the effectiveness of

pretraining on public data and fine-tuning privately on sensitive data (Li et al., 2022; Yu et al., 2022; Ganesh et al., 2023; Yu et al., 2023). Using the information obtained in public data, Wang and Zhou (2020); Kairouz et al. (2021); Yu et al. (2021a); Zhou et al. (2021); Amid et al. (2022); Nasr et al. (2023) investigated preconditioning or adaptively clipping the private gradients, which reduce the required amount of noise in DP-SGD and accelerate its convergence. Theoretical works such as Bassily et al. (2018); Alon et al. (2019); Bassily et al. (2020) studied sample complexity bounds for PAC learning and query release based on the VC dimension of the function space. Bie et al. (2022) used public data to standardize private data and showed that the sample complexity of Gaussian mean estimation can be augmented in the sense of the range parameter. More recently, by relating to sample compression schemes, Ben-David et al. (2023) presented sample complexities for several problems including Gaussian mixture estimation. Ganesh et al. (2023) explained the necessity of public data by investigating the loss landscape.

In contrast, there is less attention focused on the local setting. Su et al. (2023) considered half-space estimation with an auxiliary unlabeled public data. They construct several weak LDP classifiers and label the public data by majority vote. They established sample complexity that is linear in the dimension and polynomial in other terms for both private and public data. Wang et al. (2023a) employed public data to estimate the leading eigenvalue of the covariance matrix, which will serve as a reference for data holders to clip their statistics. Given the learned clipping threshold, the sample complexity for GLM with LDP for a general class of functions is improved to polynomial concerning error, dimension, and privacy budget, which is shown impossible without public data (Smith et al., 2017; Dagan and Feldman, 2020). Both Wang et al. (2023a) and Su et al. (2023) considered the unlabeled data, while our work further leverages the information contained in the labels if labeled public data is available. More recently, Ma et al. (2023) enhanced LDP regression by using public data to create an informative partition. However, the relationship between the source and target distributions is vaguely defined. Also, their theoretical results are no better than using just private data. Lowy et al. (2023) studied mean estimation, DP-SCO, and DP-ERM with public data. They assume that the public data and the private data are identically distributed. Also, the theoretical results are established under sequentially-interactive LDP. Despite the distinctions in the research problem, their conclusions are analogous to ours: the asymptotic optimal error can be achieved by using either the private or public estimation solely, while weighted average estimation achieves better constant as well as empirical performance.

2.2 Locally Private Classification

LDP classification with parametric assumptions can be viewed as a special case of LDP-ERM (Smith et al., 2017; Zheng et al., 2017; Wang et al., 2018; Daniely and Feldman, 2019; Dagan and Feldman, 2020; Wang et al., 2023a). Most works provide evidence of hardness on non-interactive LDP-ERM. Non-parametric methods possess a worse privacy-utility trade-off (Duchi et al., 2018) and are considered to be harder than parametric problems. Zheng et al. (2017) investigated the kernel ridge regression and established a bound of estimation error of the order $n^{-1/4}$. Their results only apply to losses that are strongly convex and smooth. The works most related to ours is that of Berrett and Butucea (2019); Berrett

et al. (2021). Berrett and Butucea (2019) studied non-parametric classification with the Hölder smoothness assumption. They proposed a histogram-based method using the Laplace mechanism and showed that this approach is mini-max optimal under their assumption. The strong consistency of this approach is established in Berrett et al. (2021) as a by-product. Our results include their conclusion as a special case and hold in a stronger sense, i.e. “uniformly with high probability” instead of “in expectation”.

2.3 Private Hyperparameter Tuning

Researchers highlighted the importance and difficulty of parameter tuning under DP constraint (Papernot and Steinke, 2021; Wang et al., 2023b) focusing on central DP. Chaudhuri and Vinterbo (2013) explored this problem under strong stability assumptions. Liu and Talwar (2019) presented DP algorithms for selecting the best parameters in several candidates. Compared to the original algorithm, their method suffers a factor of 3 in the privacy parameter, meaning that 2/3 of the privacy budget is spent on the selection process. Papernot and Steinke (2021) improved the result with a much tighter privacy bound in Renyi-DP by running the algorithm random times and returning the best result. Recently, Wang et al. (2023b) proposed a framework that reduces DP parameter selection to a non-DP counterpart. This approach leverages existing hyperparameter optimization methods and outperforms the uniform sampling approaches (Liu and Talwar, 2019; Papernot and Steinke, 2021). Ramaswamy et al. (2020) proposed to tune hyperparameters on public data sets. Mohapatra et al. (2022) argued that, under a limited privacy budget, DP optimizers that require less tuning are preferable. As far as we know, there is a lack of research focusing on the LDP parameter selection. Butucea et al. (2020) investigated adaptive parameter selection in non-parametric density estimation, and provided a data-driven wavelet estimator.

3 Locally Differentially Private Classification Tree

In this section, we introduce the *Locally differentially Private Classification Tree* (LPCT). After explaining necessary notations and problem setting in Section 3.1, we propose our decision tree partition rule in Section 3.2. Finally, Section 3.3 introduces our privacy mechanism and proposes the final algorithm.

3.1 Problem Setting

We introduce necessary notations. For any vector x , let x^i denote the i -th element of x . Recall that for $1 \leq p < \infty$, the L_p -norm of $x = (x^1, \dots, x^d)$ is defined by $\|x\|_p := (|x^1|^p + \dots + |x^d|^p)^{1/p}$. Let $x_{i:j} = (x_i, \dots, x_j)$ be the slicing view of x in the i, \dots, j position. Throughout this paper, we use the notation $a_n \lesssim b_n$ and $a_n \gtrsim b_n$ to denote that there exist positive constant c and c' such that $a_n \leq cb_n$ and $a_n \geq c'b_n$, for all $n \in \mathbb{N}$. In addition, we denote $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $b_n \lesssim a_n$. Let $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. Besides, for any set $A \subset \mathbb{R}^d$, the diameter of A is defined by $\text{diam}(A) := \sup_{x, x' \in A} \|x - x'\|_2$. Let $f_1 \circ f_2$ represent the composition of functions f_1 and f_2 . Let $A \times B$ be the Cartesian product of sets where $A \in \mathcal{X}_1$ and $B \in \mathcal{X}_2$. For measure P on \mathcal{X}_1 and Q on \mathcal{X}_2 , define the product measure $P \otimes Q$ on $\mathcal{X}_1 \times \mathcal{X}_2$ as $P \otimes Q(A \times B) = P(A) \cdot Q(B)$. For an integer k ,

denote the k -fold product measure on \mathcal{X}_1^k as P^k . Let the standard Laplace random variable have probability density function $e^{-|x|}/2$ for $x \in \mathbb{R}$.

It is legitimate to consider binary classification while an extension of our results to multi-class classification is straightforward. Suppose there are two unknown probability measures, the private measure P and the public measure Q on $\mathcal{X} \times \mathcal{Y} = [0, 1]^d \times \{0, 1\}$. We observe n_P i.i.d. private samples $D^P = \{(X_1^P, Y_1^P), \dots, (X_{n_P}^P, Y_{n_P}^P)\}$ drawn from the target distribution P , and n_Q i.i.d. public samples $D^Q = \{(X_1^Q, Y_1^Q), \dots, (X_{n_Q}^Q, Y_{n_Q}^Q)\}$ drawn from the source distribution Q . The data points from the distributions P and Q are also mutually independent. Our goal is to classify under the target distribution P . Given the observed data, we construct a classifier $\hat{f} : \mathcal{X} \rightarrow \mathcal{Y}$ which minimizes the classification risk under the target distribution P :

$$\mathcal{R}_P(\hat{f}) \triangleq \mathbb{P}_{(X,Y) \sim P}(Y \neq \hat{f}(X)).$$

Here $\mathbb{P}_{(X,Y) \sim P}(\cdot)$ means the probability when (X, Y) are drawn from distribution P . The Bayes risk, which is the smallest possible risk with respect to P , is given by $\mathcal{R}_P^* := \inf\{\mathcal{R}_P(f) | f : \mathcal{X} \rightarrow \mathcal{Y} \text{ is measurable}\}$. In such binary classification problems, the regression functions are defined as

$$\eta_P^*(x) \triangleq P(Y = 1 | X = x) \quad \text{and} \quad \eta_Q^*(x) \triangleq Q(Y = 1 | X = x), \quad (1)$$

which represent the conditional distributions $P_{Y|X}$ and $Q_{Y|X}$. The function that achieves Bayes risk with respect to P is called Bayes function, namely, $f_P^*(x) := \mathbf{1}(\eta_P^*(x) > 1/2)$.

We consider the following setting. The estimator \hat{f} is considered as a random function with respect to both D^P and D^Q , while its construction process with respect to D^P is *locally differentially private* (LDP). The rigorous definition of LDP is as follows.

Definition 1 (Local Differential Privacy) *Given data $\{(X_i^P, Y_i^P)\}_{i=1}^{n_P}$, a privatized information Z_i , which is a random variable on \mathcal{S} , is released based on (X_i^P, Y_i^P) and Z_1, \dots, Z_{i-1} . Let $\sigma(\mathcal{S})$ be the σ -field on \mathcal{S} . Z_i is drawn conditional on (X_i^P, Y_i^P) and $Z_{1:i-1}$ via the distribution $R_i(S | X_i^P = x, Y_i^P = y, Z_{1:i-1} = z_{1:i-1})$ for $S \in \sigma(\mathcal{S})$. Then the mechanism $R = \{R_i\}_{i=1}^{n_P}$ is sequentially-interactive ε -locally differentially private (ε -LDP) if*

$$\frac{R_i(S | X_i^P = x, Y_i^P = y, Z_{1:i-1} = z_{1:i-1})}{R_i(S | X_i^P = x', Y_i^P = y', Z_{1:i-1} = z_{1:i-1})} \leq e^\varepsilon$$

for all $1 \leq i \leq n_P, S \in \sigma(\mathcal{S}), x, x' \in \mathcal{X}, y, y' \in \mathcal{Y}$, and $z_{1:i-1} \in \mathcal{S}^{i-1}$. Moreover, if

$$\frac{R_i(S | X_i^P = x, Y_i^P = y)}{R_i(S | X_i^P = x', Y_i^P = y')} \leq e^\varepsilon$$

for all $1 \leq i \leq n_P, S \in \sigma(\mathcal{S}), x, x' \in \mathcal{X}, y, y' \in \mathcal{Y}$, then R is non-interactive ε -LDP.

This formulation is widely adopted (Duchi et al., 2018; Berrett and Butucea, 2019). In contrast to central DP where the likelihood ratio is taken concerning some statistics of all data, LDP requires individuals to guarantee their own privacy by considering the likelihood ratio of each (X_i^P, Y_i^P) . Once the view z is provided, no further processing can reduce

the deniability about taking a value (x, y) since any outcome z is nearly as likely to have come from some other initial value (x', y') . Sequentially-interactive mechanisms fulfill the requirements for a wide range of methods including gradient-based approaches. However, they can prove to be prohibitively slow in practice. Conversely, non-interactive mechanisms are quite efficient but suffer from poorer performance (Smith et al., 2017). Our analysis does not encompass the more general fully-interactive mechanisms due to their complex nature (Duchi et al., 2018).

Besides the privacy guarantee of P data, we consider the existence of additional public data from distribution Q. To depict the relationship between P and Q, we consider the *posterior drift* setting in transfer learning literature (Cai and Wei, 2021). Under the posterior drift model, let $P_X = Q_X$ be the identical marginal distribution. The main difference between P and Q lies in the regression functions $\eta_P^*(x)$ and $\eta_Q^*(x)$. Specifically, let $\eta_Q^*(x) = \phi(\eta_P^*(x))$ for some strictly increasing link function $\phi(\cdot)$ with $\phi(1/2) = 1/2$. The condition that ϕ is strictly increasing leads to the situation where those X that are more likely to be labeled $Y = 1$ under P are more likely to be labeled $Y = 1$ under Q. The assumption $\phi(1/2) = 1/2$ guarantees that those X that are non-informative under P are the same under Q, and more importantly, the sign of $\eta_P^* - 1/2$ and $\eta_Q^* - 1/2$ are identical.

3.2 Decision Tree Partition

In this section, we formalize the decision tree partition process. A binary tree partition π is a disjoint cover of \mathcal{X} obtained by recursively split grids into subgrids. While our methodology applies to any tree partition, it can be challenging to use general partitions such as the original CART (Breiman, 1984) for theoretical analysis. Following Cai et al. (2023); Ma et al. (2023), we propose a new splitting rule called the *max-edge partition rule*. This rule is amenable to theoretical analysis and can also achieve satisfactory practical performance. In the non-interactive setting, the private data set is unavailable during the partition process and the partition is created solely on the public data. Given public data set $\{(X_i^Q, Y_i^Q)\}_{i=1}^{n_Q}$, the partition rule is stated as follows:

- Let $A_{(0)}^1 := [0, 1]^d$ be the initial rectangular cell and $\pi_0 := \{A_{(0)}^j\}_{j \in \mathcal{I}_0}$ be the initialized cell partition. $\mathcal{I}_0 = \{1\}$ stands for the initialized index set. In addition, let $p \in \mathbb{N}$ represent the maximum depth of the tree. The parameter is fixed beforehand by the user and possibly depends on n .
- Suppose we have obtained a partition π_{i-1} of \mathcal{X} after $i - 1$ steps of the recursion. Let $\pi_i = \emptyset$. In the i -th step, let $A_{(i-1)}^j \in \pi_{i-1}$ be $\times_{\ell=1}^d [a_\ell, b_\ell]$ for $j \in \mathcal{I}_{i-1}$. We choose the edge to be split among the longest edges. The index set of longest edges is defined as

$$\mathcal{M}_{(i-1)}^j = \left\{ k \mid |b_k - a_k| = \max_{\ell=1, \dots, d} |b_\ell - a_\ell|, k = 1, \dots, d \right\}.$$

- Assume we split along the ℓ -th dimension for $\ell \in \mathcal{M}_{(i-1)}^j$, $A_{(i-1)}^j$ is then partitioned into a left sub-cell $A_{(i-1)}^{j,0}(\ell)$ and a right sub-cell $A_{(i-1)}^{j,1}(\ell)$ along the midpoint of the chosen dimension, where $A_{(i-1)}^{j,0}(\ell) = \{x \mid x \in A_{(i-1)}^j, x^\ell < (a_\ell + b_\ell)/2\}$ and

Algorithm 1: Max-edge Partition

Input: Public data D^Q , depth p .

Initialization: $\pi_0 = [0, 1]^d$.

for $i = 1$ **to** p **do**

$\pi_i = \emptyset$

for j **in** \mathcal{I}_{i-1} **do**

 Select ℓ as in (2).

$\pi_i = \pi_i \cup \{A_{(i-1)}^{j,0}(\ell), A_{(i-1)}^{j,1}(\ell)\}$.

end

end

Output: Partition π_p

$A_{(i-1)}^{j,1}(\ell) = A_{(i-1)}^j / A_{(i-1)}^{j,0}(\ell)$. Then the dimension to be split is chosen by

$$\min_{\ell \in \mathcal{M}_{(i-1)}^j} \arg \min \text{score} \left(A_{(i-1)}^{j,0}(\ell), A_{(i-1)}^{j,1}(\ell), D^Q \right), \quad (2)$$

where the $\text{score}(\cdot)$ function can be criteria for decision tree classifiers such as the Gini index or information gain. The \min operator implies that if the $\arg \min$ function returns multiple indices, we select the smallest index. Thus, the partition process is feasible even if there is no public data, where all axes have the same score.

- Once ℓ is selected, let $\pi_i = \pi_i \cup \{A_{(i-1)}^{j,0}(\ell), A_{(i-1)}^{j,1}(\ell)\}$.

The complete process is presented in Algorithm 1 and illustrated in Figure 1. For each grid, the partition rule selects the midpoint of the longest edges that achieves the largest score reduction. This procedure continues until either node has no samples or the depth of the tree reaches its limit. We also define relationships between nodes which are necessary in Section 5. For node $A_{(i)}^j$, we define its parent node as

$$\text{parent}(A_{(i)}^j) := A_{(i-1)}^{j'} \quad \text{s.t.} \quad A_{(i-1)}^{j'} \in \pi_{i-1}, \quad A_{(i)}^j \subset A_{(i-1)}^{j'}.$$

For $i' < i$, the ancestor of a node $A_{(i)}^j$ with depth i' is then defined as

$$\text{ancestor}(A_{(i)}^j, i') = \text{parent}^{i-i'}(A_{(i)}^j) = \underbrace{\text{parent} \circ \cdots \circ \text{parent}}_{i-i'}(A_{(i)}^j),$$

which is the grid in $\pi_{i'}$ that contains $A_{(i)}^j$.

3.3 Privacy Mechanism for Any Partition

This section focuses on the privacy mechanism based on general tree partitions. We first introduce the necessary definitions and then present our private estimation based on the Laplacian mechanism in (8). To avoid confusion, we refer the readers to a clear summarization of the defined estimators in Appendix A.1.

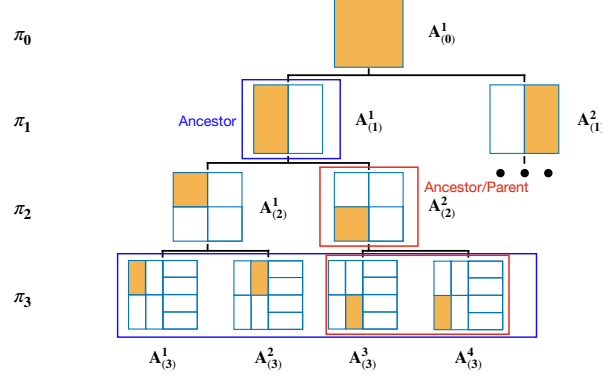


Figure 1: Partition created by the max-edge rule. The areas filled with orange represent the corresponding $A_{(i)}^j$ and the blue lines represent the partition boundaries after each level of partitioning. Red boxes contain an ancestor with depth 2, which is also a parent. Blue boxes contain an ancestor with depth 1.

For index set \mathcal{I} , let $\pi = \pi_p = \{A_{(p)}^j\}_{j \in \mathcal{I}_p}$ be the tree partition created by Algorithm 1. A *population classification tree estimation* is defined as

$$\bar{\eta}_P^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{\int_{A_{(p)}^j} \eta_P^*(x') dP_X(x')}{\int_{A_{(p)}^j} dP_X(x')} \quad (3)$$

The label is inferred using the *population tree classifier* which is defined as $\bar{f}_P^\pi(x) = \mathbf{1}(\bar{\eta}_P^\pi(x) > 1/2)$. Here, we let $0/0 = 0$ by definition.

To get an empirical estimator given the data set $D = \{(X_1^P, Y_1^P), \dots, (X_{n_P}^P, Y_{n_P}^P)\}$, we estimate the numerator and the denominator of (3) separately. To estimate the denominator, each sample (X_i^P, Y_i^P) contributes a one-hot vector $U_i^P \in \{0, 1\}^{|\mathcal{I}_p|}$ where the j -th element of U_i^P is $\mathbf{1}\{X_i^P \in A_{(p)}^j\}$. Then an estimation of $\int_{A_{(p)}^j} dP_X(x')$ is $\frac{1}{n_P} \sum_{i=1}^{n_P} U_i^{Pj}$, which is the number of samples in $A_{(p)}^j$ divided by n_P . Analogously, let $V_i^P = Y_i \cdot U_i^P \in \{0, 1\}^{|\mathcal{I}_p|}$. An estimation of $\int_{A_{(p)}^j} \eta_P^*(x') dP_X(x')$ is $\frac{1}{n_P} \sum_{i=1}^{n_P} V_i^{Pj}$, which is the sum of the labels in $A_{(p)}^j$ divided by n_P . Combining the pieces, a *classification tree estimation* is defined as

$$\hat{\eta}_P^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{\sum_{i=1}^{n_P} V_i^{Pj}}{\sum_{i=1}^{n_P} U_i^{Pj}}. \quad (4)$$

The corresponding classifier is defined as $\hat{f}_P^\pi(x) = \mathbf{1}(\hat{\eta}_P^\pi(x) > 1/2)$. In other words, $\hat{\eta}_P^\pi(x)$ estimates $\eta_P^*(x)$ by the average of the responses in the cell. In the non-private setting, each data holder prepares U_i^P and V_i^P according to the partition π and sends it to the curator. Then the curator aggregates the transmission following (4).

To protect the privacy of each data, we propose to estimate the numerator and denominator of the population regression tree using a privatized method. Specifically, U_i^P and V_i^P are combined with the Laplace mechanism (Dwork et al., 2006) before being sent to the

curator. For i.i.d. standard Laplace random variables ζ_i^j and ξ_i^j , let

$$\tilde{U}_i^P := U_i^P + \frac{4}{\varepsilon} \zeta_i = U_i^P + (\frac{4}{\varepsilon} \zeta_i^1, \dots, \frac{4}{\varepsilon} \zeta_i^{|\mathcal{I}_p|}). \quad (5)$$

Then a privatized estimation of $\int_{A_{(p)}^j} dP_X(x)$ is $\frac{1}{n_P} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj}$. Analogously, let

$$\tilde{V}_i^P := V_i^P + \frac{4}{\varepsilon} \xi_i = V_i^P + (\frac{4}{\varepsilon} \xi_i^1, \dots, \frac{4}{\varepsilon} \xi_i^{|\mathcal{I}_p|}). \quad (6)$$

An estimation of $\int_{A_{(p)}^j} \eta_P^*(x) dP_X(x)$ is $\frac{1}{n_P} \sum_{i=1}^{n_P} \tilde{V}_i^{Pj}$. Then using the privatized information $(\tilde{U}_i^P, \tilde{V}_i^P), i = 1, \dots, n_P$, we can estimate the regression function as

$$\tilde{\eta}_P^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj}}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj}}. \quad (7)$$

The procedure is also derived by Berrett and Butucea (2019). As an alternative, one can use random response (Warner, 1965) since both U and V are binary vectors. The following proposition shows that the estimation procedures satisfy LDP.

Proposition 2 *Let $\pi = \{A^j\}_{j \in \mathcal{I}}$ be any partition of \mathcal{X} with $\cup_{j \in \mathcal{I}} A^j = \mathcal{X}$ and $A^i \cap A^j = \emptyset, i \neq j$. Then the privacy mechanism defined in (5) and (6) is non-interactive ε -LDP.*

In the case where one also has access to the Q-data in addition to the P-data, the Q data can be used to help the classification task under the target distribution P and should be taken into consideration. The utilization of P data must satisfy the local DP constraint while the Q data can be arbitrarily adopted. We accommodate the existing private estimation (7) and non-private estimation (4) by taking a weighted average based on both estimators. Denote the encoded information from D^P and D^Q as (U_i^P, V_i^P) and (U_i^Q, V_i^Q) , respectively. We use the same partition for P estimator and Q estimator. Namely, (4) and (7) are constructed via an identical π . When taking the average, data from the different distributions should have different weights since the signal strengths are different between P and Q. To make the classification at $x \in \mathcal{X}$, the *Locally differentially Private Classification Tree estimation* (LPCT) is defined as follows:

$$\tilde{\eta}_{P,Q}^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}}. \quad (8)$$

Here, we let $\lambda \geq 0$. The *locally private tree classifier* is $\tilde{f}_{P,Q}^\pi(x) = \mathbf{1}(\tilde{\eta}_{P,Q}^\pi(x) > 1/2)$. In each leaf node $A_{(p)}^j$, the classifier is assigned a label based on the weighted average of estimations from P and Q data. The parameter λ serves to balance the relative contributions of both P and Q. A higher value of λ indicates that the final predictions are predominantly influenced by the public data, and vice versa.

The estimator defined in (8) easily adapts to solve the non-private transfer learning problem in Cai and Wei (2021) if we define

$$\hat{\eta}_{P,Q}^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}}. \quad (9)$$

Similar to the nearest neighbor estimator in Cai and Wei (2021), (9) estimates by taking weighted averages of labels from nearby samples. Our theoretical results in Section 4.2 imply that (9) is also rate optimal in terms of rate for non-private transfer learning. A key advantage for incorporating (9) is that, given the partition, it stores only the prediction values at each node, which makes it convenient to inject noise. Conversely, the nearest neighbor estimator (similarly, the kernel estimator) requires a query to the entire training data set for each single testing sample, which prohibits any modification towards a private estimator. See Kroll (2021) for an example of such limitations.

4 Theoretical Results

In this section, we present the obtained theoretical results. We present the matching lower and upper bounds of excess risk in Section 4.1 and 4.2. Finally, we demonstrate that LPCT is free from the unknown range parameter in Section 4.3. All technical proofs can be found in Appendix B.

4.1 Mini-max Convergence Rate

We first give the necessary assumptions on the distribution P and Q . Then we provide the mini-max rate under the assumptions.

Assumption 1 *Let $\alpha \in (0, 1]$. Assume the regression function $\eta_P^* : \mathcal{X} \rightarrow \mathbb{R}$ is α -Hölder continuous, i.e. there exists a constant $c_L > 0$ such that for all $x_1, x_2 \in \mathcal{X}$, $|\eta_P^*(x_1) - \eta_P^*(x_2)| \leq c_L \|x_1 - x_2\|^\alpha$. Also, assume that the marginal density functions of P and Q are upper and lower bounded, i.e. $\underline{c} \leq dP_X(x) \leq \bar{c}$ for some $\underline{c}, \bar{c} > 0$.*

Assumption 2 (Margin Assumption) *Let $\beta \geq 0, C_\beta > 0$. Assume*

$$P_X \left(\left| \eta_P^*(X) - \frac{1}{2} \right| < t \right) \leq C_\beta t^\beta.$$

Assumption 3 (Relative Signal Exponent) *Let the relative signal exponent $\gamma \in (0, \infty)$ and $C_\gamma \in (0, \infty)$. Assume that*

$$\left(\eta_Q^*(x) - \frac{1}{2} \right) \left(\eta_P^*(x) - \frac{1}{2} \right) \geq 0, \text{ and } \left| \eta_Q^*(x) - \frac{1}{2} \right| \geq C_\gamma \left| \eta_P^*(x) - \frac{1}{2} \right|^\gamma.$$

Assumption 1 and 2 are standard conditions that have been widely used for non-parametric classification problems (Audibert and Tsybakov, 2007; Samworth, 2012; Chaudhuri and Dasgupta, 2014). Assumption 3 depicts the similarity between the conditional distribution of public data and private data. When γ is small, the signal strength of Q is strong compared to P . In the extreme instance of $\gamma = 0$, $\eta_Q^*(x)$ consistently remains distant from $1/2$ by a constant. In contrast, when γ is large, there is not much information contained in Q . The special case of $\gamma = \infty$ allows $\eta_Q^*(x)$ to always be $1/2$ and is thus non-informative. We present the following theorem that specifies the mini-max lower bound with LDP constraints on P data under the above assumptions. The proof is based on first constructing two function classes and then applying the information inequalities for local privacy (Duchi et al., 2018) and Assouad's Lemma (Tsybakov, 2009).

Theorem 3 Denote the function class of (P, Q) satisfying Assumption 1, 2, and 3 by \mathcal{F} . Then for any estimator \hat{f} that is sequentially-interactive ε -LDP, there holds

$$\inf_{\hat{f}} \sup_{\mathcal{F}} \mathbb{E}_{X,Y} [\mathcal{R}_P(\hat{f}) - \mathcal{R}_P^*] \gtrsim \left(n_P (e^\varepsilon - 1)^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{\alpha(\beta+1)}{2\alpha+2d}}. \quad (10)$$

If $n_Q = 0$, our results reduce to the special case of LDP non-parametric classification with only private data. In this case, the mini-max lower bound is of the order $(n_P(e^\varepsilon - 1)^2)^{-\frac{\alpha(\beta+1)}{2\alpha+2d}}$. Previously, Berrett and Butucea (2019) provided an estimator that reaches the mini-max optimal convergence rate for small ε . If $n_P = 0$, the result recovers the learning rate $n_Q^{-\frac{\alpha(\beta+1)}{2\gamma\alpha+d}}$ established for transfer learning under posterior drift (Cai and Wei, 2021). In this case, we have only source data from Q and want to generalize the classifier on the target distribution P .

The mini-max convergence rate (10) is the same as if one fits an estimator using only P data with sample size $n_P + (e^\varepsilon - 1)^{-2} \cdot n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$. The contribution of Q data is substantial when $n_P \ll (e^\varepsilon - 1)^{-2} \cdot n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$. Otherwise, the convergence rate is not significantly better than using P data only. The quantity $(e^\varepsilon - 1)^{-2} \cdot n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$, which is referred to as effective sample size, represent the theoretical gain of using Q data with sample size n_Q under given ε, α, d , and γ . Notably, the smaller γ is, the more effective the Q data becomes, and the greater gain is acquired by using Q data. Regarding different values of γ , our work covers the setting in previous work when a small amount of effective public samples are required (Bie et al., 2022; Ben-David et al., 2023), and when a large amount of noisy/irrelevant public samples are required (Ganesh et al., 2023; Yu et al., 2023). Also, the public data is more effective for small ε , i.e. under stronger privacy constraints. Intuitively, in the presence of stringent privacy constraints, the utilization of Q data is as effective as incorporating a substantial amount of P data. Furthermore, even for a constant ε , our effective sample size significantly exceeds $n_Q^{\frac{2\alpha+d}{2\gamma\alpha+d}}$ as in the non-private scenario (Cai and Wei, 2021). This suggests that, in comparison to transfer learning in the non-private setting, public data assumes greater value in the private setting.

4.2 Convergence Rate of LPCT

In the following, we show the convergence rate of LPCT reaches the lower bound with proper parameters. The proof is based on a carefully designed decomposition in Appendix B.3. Define the quantity

$$\delta := \left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} + \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{d}{2\alpha+2d}}. \quad (11)$$

Theorem 4 Let π be constructed through the max-edge partition in Section 3.2 with depth p . Let $\tilde{f}_{P,Q}^\pi$ be the locally private two-sample tree classifier. Suppose Assumption 1, 2, and

3 hold. Then, for $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{d}{4\alpha+2d}}$, if $\lambda \asymp \delta^{\frac{(\gamma\alpha-\alpha)\wedge(2\gamma\alpha-d)}{d}}$ and $2^p \asymp \delta^{-1}$, there holds

$$\mathcal{R}_P(\tilde{f}_{P,Q}^\pi) - \mathcal{R}_P^* \lesssim \delta^{\frac{\alpha(\beta+1)}{d}} = \left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} + \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{\alpha(\beta+1)}{2\alpha+2d}} \quad (12)$$

with probability $1 - 3/(n_P + n_Q)^2$ with respect to $P^{n_P} \otimes Q^{n_Q} \otimes R$ where R is the joint distribution of privacy mechanisms in (5) and (6).

Compared to Theorem 3, the LPCT with the best parameter choice reaches the minimax lower bound up to a logarithm factor. Both parameters λ and p should be assigned a value of the correct order to achieve the optimal trade-off. The depth p increases when either n_P or n_Q grows. This is intuitively correct since the depth of a decision tree should increase with a larger number of samples. The strength ratio λ , on the other hand, is closely related to the value of γ . As explained earlier, a small γ guarantees a stronger signal in η_Q^* . Thus, the value of λ is decreasing for γ . In this case, when the signal strength of η_Q^* is large, we rely more on the estimation by Q data by assigning a larger ratio to it.

We note that when ε is large, there is a gap between $(e^\varepsilon - 1)^2$ in the lower bound and ε^2 in the upper bound. The problem can be mitigated by employing the random response mechanism (Warner, 1965) instead of the Laplace mechanism, as mentioned in Section 3.3. The theoretical analysis generalizes analogously. Here, we do not address this issue as we want to focus on public data. The theorem restricts $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{d}{4\alpha+2d}}$. When $\varepsilon \lesssim n_P^{-\frac{\alpha}{2\alpha+2d}}$, the estimator is drastically perturbed by the random noise and the convergence rate is no longer optimal. As ε gets larger, the upper bound of excess risk decreases until $\varepsilon \asymp n_P^{\frac{d}{4\alpha+2d}}$. When $\varepsilon \gtrsim n_P^{\frac{d}{4\alpha+2d}}$, the random noise is negligible and the convergence rate remains to be

$$\left(\frac{n_P}{\log(n_P + n_Q)} + \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+d}{2\gamma\alpha+d}} \right)^{-\frac{\alpha(\beta+1)}{2\alpha+d}},$$

which recovers the rate of non-private learning as established in Cai and Wei (2021).

We discuss the order of ratio λ . When $\gamma \geq \frac{d}{\alpha} - 1$, we have $\lambda \asymp \delta^{\frac{(\gamma-1)\alpha}{d}}$. This rate is exactly the same as in the conventional transfer learning case (Cai and Wei, 2021) and is enough to achieve the optimal trade-off. Yet, under privacy constraints, this choice of λ will fail when $\gamma < \frac{d}{\alpha} - 1$. Roughly speaking, if we let $\lambda \asymp \delta^{\frac{(\gamma-1)\alpha}{d}}$ in this case, there is a probability such that the noise $\frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j$ in (8) will dominate the non-private density estimation $\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}$. When the domination happens, the point-wise estimation is decided by random noises, and thus the overall excess risk fails to converge. To alleviate this issue, we assign a larger value $\lambda \asymp \delta^{\frac{2\gamma\alpha-d}{d}} > \delta^{\frac{(\gamma-1)\alpha}{d}}$ to ensure $\frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j \lesssim \sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}$. As an interpretation, in the private setting, we place greater reliance on the estimation by public data compared to the non-private setting.

Last, we discuss the computational complexity of LPCT. We first consider the average computation complexity of LPCT. The training stage consists of two parts. Similar to the

standard decision tree algorithm (Breiman, 1984), the partition procedure takes $\mathcal{O}(pn_Q d)$ time. The computation of (8) takes $\mathcal{O}(p(n_P + n_Q))$ time. From the proof of Theorem 4, we know that $2^p \asymp \delta^{-1}$, where $\delta^{-1} \lesssim n_P + n_Q$. Thus the training stage complexity is around $\mathcal{O}((n_P + n_Q d) \log(n_P + n_Q))$ which is linear. Since each prediction of the decision tree takes $\mathcal{O}(p)$ time, the test time for each test instance is no more than $\mathcal{O}(\log(n_P + n_Q))$. As for storage complexity, since LPCT only requires the storage of the tree structure and the prediction value at each node, the space complexity of LPCT is $\mathcal{O}(\delta^{-1})$ which is also sub-linear. In short, LPCT is an efficient method in time and memory.

4.3 Eliminating the Range Parameter

We discuss how public data can eliminate the range parameter contained in the convergence rate. When \mathcal{X} is unknown, one must select a predefined range $[-r, r]^d$ for creating partitions. Only then do the data holders have a reference for encoding their information. In this case, the convergence rate of Berrett and Butucea (2019) becomes $(r^{2d}/(n_P \varepsilon^2))^{\frac{(1+\beta)\alpha}{2\alpha+2d}} \vee (1 - \int_{[-r,r]^d} dP(\mathbf{x}))$, which is slower than the known $\mathcal{X} = [0, 1]^d$ for any r . A small r means that a large part of the informative \mathcal{X} will be ignored (the second term is large), whereas a large r may create too many redundant cells (the first term is large). See Györfi and Kroll (2022) for detailed derivation. Without much prior knowledge, this parameter can be hard to tune. The removal of the range parameter is first discussed in Bie et al. (2022), where $d+1$ public data is required. Our result shows that we do not need to tune this range parameter with the help of public data.

Theorem 5 *Suppose the domain $\mathcal{X} = \times_{j=1}^d [a^j, b^j]$ is unknown. Let $\hat{a}^j = \min_i X_i^j$ and $\hat{b}^j = \max_i X_i^j$. Define the min-max scaling map $\mathcal{M} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ where $\mathcal{M}(x)^j = (x^j - \hat{a}^j)/(\hat{b}^j - \hat{a}^j)$. The min-max scaled D^P and D^Q is written as $\mathcal{M}(D^P)$ and $\mathcal{M}(D^Q)$. Then for $\tilde{f}_{P,Q}^\pi$ trained within $[0, 1]^d$ on $\mathcal{M}(D^P)$ and $\mathcal{M}(D^Q)$ as in Theorem 4, we have*

$$\mathcal{R}_P \left((\tilde{f}_{P,Q}^\pi \cdot \mathbf{1}_{[0,1]^d}) \circ \mathcal{M} \right) - \mathcal{R}_P^* \lesssim \delta^{\frac{\alpha(\beta+1)}{d}} + \frac{\log n_Q}{n_Q}$$

with probability $1 - 3/(n_P + n_Q)^2 - d/n_Q^2$ with respect to $P^{n_P} \otimes Q^{n_Q} \otimes R$.

Compared to Theorem 4, the upper bound of Theorem 5 increases by $\log n_Q/n_Q$, which is a minor term in most cases. Moreover, Theorem 5 fails with additional probability d/n_Q^2 since we are estimating $[a^j, b^j]$ with $[\hat{a}^j, \hat{b}^j]$. Both changes indicate that sufficient public data is required to resolve the unknown range issue.

5 Data-driven Tree Structure Pruning

In this section, we develop a data-driven approach for parameter tuning and investigate its theoretical properties. In Section 5.1, we derive a selection scheme in the spirit of Lepski's method (Lepskii, 1992), a classical method for adaptively choosing the hyperparameters. In Section 5.2, we make slight modifications to the scheme and illustrate the effectiveness of the amended approach.




Figure 2: Illustration of pruning process. The yellow area filled at node $A_{(i)}^j$ means that for $x \in A_{(i)}^j$, we use the average of labels in the yellow area instead of $A_{(i)}^j$ itself. The prediction in $A_{(3)}^1$ is pruned to its depth - 1 ancestor; the prediction in $A_{(3)}^2$ is pruned to its depth - 2 ancestor; the prediction in $A_{(3)}^3$ is pruned to its depth - 3 ancestor, i.e. not pruned.

5.1 A Naive Strategy

We first develop a private analog of Lepski's method (Lepskii, 1992) for LPCT. There are two parameters: the depth p and the relative ratio λ , both of which have an optimal value depending on the unknown α and γ . We do not set the parameters to specific values. Instead, we first query the privatized information with a sufficient tree depth p_0 . Then, we introduce a greedy procedure to locally prune the tree to an appropriate depth and select λ based on local information.

We provide a preliminary result that will guide the derivation of the pruning procedure. For max-edge partition π_{p_0} with depth p_0 and node $A_{(p_0)}^j$, the *depth k ancestor private tree estimation* is defined as

$$\tilde{\eta}_{P,Q}^{\pi_k}(x) = \sum_{j \in \mathcal{I}_{p_0}} \mathbf{1}\{x \in A_{(p_0)}^j\} \frac{\sum_{\text{ancestor}(A_{(p_0)}^{j'},k)=\text{ancestor}(A_{(p_0)}^j,k)} \left(\sum_{i=1}^{n_P} \tilde{V}_i^{Pj'} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj'} \right)}{\sum_{\text{ancestor}(A_{(p_0)}^{j'},k)=\text{ancestor}(A_{(p_0)}^j,k)} \left(\sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj'} \right)} \quad (13)$$

and consequently the classifier is defined as $\tilde{f}_{P,Q}^{\pi_k}(x) = \mathbf{1}(\tilde{\eta}_{P,Q}^{\pi_k}(x) > 1/2)$. See Figure 2(a) for an illustration of estimation at ancestor nodes.

Proposition 6 *Let π be constructed through the max-edge partition in Section 3.2 with depth p_0 . Let $\tilde{\eta}_{P,Q}^{\pi_k}$ be the tree estimator with depth k ancestor node in (13). Let $\hat{\eta}_{P,Q}^{\pi_k}$ be the non-private tree estimator with partition π_k in (9). Define the quantity*

$$r_k^j := \sqrt{\frac{4C_1 (2^{p_0-k+3}\varepsilon^{-2}n_P) \vee \left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} \right) + C_1 \lambda^2 \sum_{j'} \sum_{i=1}^{n_P} U_i^{Qj'}}{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj} \right)^2}} \quad (14)$$

where the summation over j' is with respect to all descendent nodes of the depth- k ancestor, i.e. $\text{ancestor}(A_{(p_0)}^{j'},k) = \text{ancestor}(A_{(p_0)}^j,k)$. The constant C_1 is specified in the proof. Then,

for all λ and p_0 , if $x \in A_{(p_0)}^j$, there holds

$$\left| \tilde{\eta}_{P,Q}^{\pi_k}(x) - \mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_k}(x) \right] \right| \leq r_k^j$$

with probability $1 - 3/(n_P + n_Q)^2$ with respect to $P_{Y^P|X^P}^{n_P} \otimes Q_{Y^Q|X^Q}^{n_Q} \otimes R$ where R is the joint distribution of privacy mechanisms in (5) and (6). $\mathbb{E}_{Y|X}$ is taken with respect to $P_{Y|X}^{n_P} \otimes Q_{Y|X}^{n_Q}$.

The above proposition indicates that if $\tilde{\eta}_{P,Q}^{\pi_k}(x) - 1/2 \geq r_k^j$, then there holds

$$\mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_k}(x) \right] - \frac{1}{2} = \mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_k}(x) \right] - \tilde{\eta}_{P,Q}^{\pi_k}(x) + \tilde{\eta}_{P,Q}^{\pi_k}(x) - \frac{1}{2} \geq 0$$

with probability $1 - 3/(n_P + n_Q)^2$. Similar conclusion holds when $\tilde{\eta}_{P,Q}^{\pi_k}(x) - 1/2 < r_k^j$. In other words, when we have

$$\frac{|\tilde{\eta}_{P,Q}^{\pi_k}(x) - \frac{1}{2}|}{r_k^j} \geq 1, \quad (15)$$

the population version of ancestor estimation is of the same sign as the sample version estimation. As long as (15) holds, it is enough to use the sample version estimation, and our goal is to find the best k for $\mathbb{E}_{Y|X}[\hat{\eta}_{P,Q}^{\pi_k}(x)]$ to approximate $\eta_P^*(x)$. Under continuity assumption 1, the approximation error is bounded by the radius of the largest cell (see Lemma 15). Consequently, one can show that $\mathcal{R}_P(\mathbf{1}(\mathbb{E}_{Y|X}[\hat{\eta}_{P,Q}^{\pi_k}(x)] > 1/2))$ is monotonically decreasing with respect to k . Thus, our goal is to find the largest k such that (15) holds. On the other hand, (15) is dependent on λ . We select the best possible λ in order to let (15) hold for each j and k , i.e.

$$\lambda_k^j = \arg \max_{\lambda} \frac{|\tilde{\eta}_{P,Q}^{\pi_k}(x) - 1/2|}{r_k^j} \quad (16)$$

for $x \in A_{(p_0)}^j$. The optimization problem in (16) has a closed-form solution that can be computed efficiently and explicitly. The derivation of the closed-form solution is postponed to Section A.2.

Based on the above analysis, we can perform the following pruning procedure. We first query with a sufficient depth, i.e. we select p_0 large enough, create a partition on D^Q , and receive the information from data holders. Then we prune back by the following procedure:

- For $j \in \mathcal{I}_{p_0}$, we do the following operation: for $k = p_0, \dots, 2$ and r_k^j defined in (14), we assign λ_k^j as in (16). Then if (15) holds, let $k_j = k$; else, $k \leftarrow k - 1$.
- If k_j is not assigned, we select $k_j = \arg \max_k |\tilde{\eta}_{P,Q}^{\pi_k}(x) - 1/2|/r_k^j$.
- Assign $\tilde{\eta}_{P,Q}^{prune}(x) = \tilde{\eta}_{P,Q}^{\pi_{k_j}}(x)$ for all $x \in A_{(p_0)}^j$.

This process can be done efficiently. The exact process is illustrated in Figure 2. We first calculate estimations as well as the relative signal strength at all nodes. Then we trace back from each leaf node to the root and find the node that maximizes the statistic (15). The prediction value at the leaf node is assigned as the prediction at the ancestor node with the maximum statistic value. In total, the pruning causes additional time complexity at most $\mathcal{O}(2^{p_0})$, which is ignorable compared to the original complexity as long as $2^{p_0} \lesssim n_P + dn_Q$.

5.2 Pruned LDP Classification Tree

In the non-private case, i.e. when $\varepsilon = \infty$, $\tilde{\eta}_{P,Q}^{prune}$ is essentially the Lepski's method (Lepskii, 1992; Cai and Wei, 2021). However, in the presence of privacy concerns, a tricky term $2^{p_0-k+3}\varepsilon^{-2}n_P$ arises in (14) due to the Laplace noises. Specifically, if $2^{p_0-k+3}\varepsilon^{-2}n_P \geq \sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}$, the first term in the numerator of (14) contains no information other than the level of noise. As a result, the theoretical analysis of Lepski's method fails to adapt.

To deal with the issue, we simultaneously perform two pruning procedures using information of solely P or Q data. We first query with a large depth $p_0 = \lfloor \frac{d}{2+2d} \log_2(n_P\varepsilon^2 + n_Q^{\frac{2+2d}{d}}) \rfloor$ and half of the budget $\varepsilon/2$. When $2^{p_0-k+3}\varepsilon^{-2}n_P \geq \sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}$, we compare the values of $|\tilde{\eta}_{P,Q}^{\pi_k}(x) - 1/2|/r_k^j$ with $\lambda = 0$ and $\lambda = \infty$, associating to P and Q data, respectively. If either of the statistics is larger than 1, we terminate the pruning. Else, we compare the values of the statistics. If the statistic associated with Q is larger, we return the estimation $\hat{f}_Q^{\pi_k}$. In this case, we guarantee that the pruned classifier is at least as good as the estimation using only Q data. If the statistic associated with P is larger, we proceed the pruning to $k-1$. If we reach $k \leq \lfloor \frac{d}{2+2d} \log_2(n_P\varepsilon^2) \rfloor$ for any j , we terminate the whole algorithm and return (7) with depth $\lfloor \frac{d}{2+2d} \log_2(n_P\varepsilon^2) \rfloor$, which costs another $\varepsilon/2$. In this case, we ensure Q data is unimportant and adopt a terminating condition that is close to the optimal value $\lfloor \frac{d}{2\alpha+2d} \log_2(n_P\varepsilon^2) \rfloor$. The detailed algorithm is presented in Algorithm 2. Moreover, we show that the adjusted pruning procedure will result in a near-optimal excess risk.

Theorem 7 *Let $\tilde{f}_{P,Q}^{prune} = \mathbf{1}(\tilde{\eta}_{P,Q}^{prune} \geq 1/2)$ be the pruned LPCT defined in Algorithm 2. Suppose Assumption 1, 2, and 3 hold. Let δ be defined in (11). Then, for $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{d}{4+2d}}$, with probability $1 - 4/(n_P + n_Q)^2$ with respect to $P^{n_P} \otimes Q^{n_Q} \otimes R$, we have*

$$\begin{aligned} (i) \quad & \text{if } \frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \lesssim \left(\frac{n_Q}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}, \quad \mathcal{R}_P(\tilde{f}_{P,Q}^{prune}) - \mathcal{R}_P^* \lesssim \delta^{\frac{\alpha(\beta+1)}{d}}; \\ (ii) \quad & \text{if } \frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \gtrsim \left(\frac{n_Q}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}, \quad \mathcal{R}_P(\tilde{f}_{P,Q}^{prune}) - \mathcal{R}_P^* \lesssim \delta^{\frac{\alpha(\alpha+d)(\beta+1)}{d(1+d)}}. \end{aligned}$$

The conclusion of Theorem 7 is twofold. On one hand, when Q data dominates P data, i.e. $\frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \lesssim \left(\frac{n_Q}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$, the pruned estimator remains rate-optimal. In this case, the estimator is appropriately pruned to the correct depth. On the other hand, when P data dominates Q data, the convergence rate suffers degradation. The degradation factor $(\alpha+d)/(1+d)$ is mild, and even diminishes when $\alpha = 1$, i.e. η_P^* is Lipschitz continuous. In this case, there are two possible regions in which our conclusion is derived. If $\left(\frac{n_Q}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \lesssim \left(\frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2+d}}$, the estimator triggers termination and is computed with depth $\lfloor \frac{d}{2+2d} \log_2(n_P\varepsilon^2) \rfloor$. If the opposite, i.e. $\frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \gtrsim \left(\frac{n_Q}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \gtrsim \left(\frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2+d}}$, pruning fails to identify the optimal depth $\lfloor \frac{d}{2\alpha+2d} \log_2(n_P\varepsilon^2) \rfloor$. Yet, it also does not trigger termination at $\lfloor \frac{d}{2+2d} \log_2(n_P\varepsilon^2) \rfloor$. In this region, we take a shortcut to avoid intricate analysis, while the true convergence rate is more complex yet faster than that stated in Theorem 7. Notably, one drawback of this shortcut is the discontinuity of the upper bound at the boundary $\frac{n_P\varepsilon^2}{\log(n_P+n_Q)} \asymp \left(\frac{n_Q}{\log(n_P+n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$.

Algorithm 2: Pruned LPCT

Input: Private data $D^P = \{(X_i^P, Y_i^P)\}_{i=1}^{n_P}$, public data $D^Q = \{(X_i^Q, Y_i^Q)\}_{i=1}^{n_Q}$.

Initialization: $p_0 = \lfloor \frac{d}{2+2d} \log_2(n_P \varepsilon^2 + n_Q^{\frac{2+2d}{d}}) \rfloor$, **flag** = 0.

Create π with depth p_0 on D^Q . Query (5) and (6) with $\varepsilon/2$.

for $A_{(p_0)}^j$ **in** π **do**

for $k = p_0, \dots, 1$ **do**

if $2^{p_0-k+3}\varepsilon^{-2}n_P \geq \sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}$ **then**

$$r_k^{Qj} = \sqrt{\frac{4 \log(n_P + n_Q)}{\sum_{j'} \sum_{i=1}^{n_P} U_i^{Qj'}}}, v_k^{Qj} = |\hat{\eta}_Q^{\pi_k} - 1/2|/r_k^{Qj}$$

$$r_k^{Pj} = \sqrt{\frac{2^{p_0-k+5}\varepsilon^{-2}n_P \log(n_P + n_Q)}{(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'})^2}}, v_k^{Pj} = |\hat{\eta}_P^{\pi_k} - 1/2|/r_k^{Pj}$$

if $v_k^{Qj} \leq v_k^{Pj}$ **then**

$$s_k^j = \tilde{\eta}_P^{\pi_k}, r_k^j = r_k^{Pj}, v_k^j = v_k^{Pj}$$

if $k \leq \lfloor \frac{d}{2+2d} \log_2(n_P \varepsilon^2) \rfloor$ **then**

$k_j = k$, **flag** = 1, **break**

end

else

$$s_k^j = \hat{\eta}_Q^{\pi_k}, r_k^j = r_k^{Qj}, v_k^j = v_k^{Qj}$$

end

else

$$r_k^j = \sqrt{\frac{(32 \sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} + 4\lambda^2 \sum_{j'} \sum_{i=1}^{n_P} U_i^{Qj'}) \log(n_P + n_Q)}{(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'})^2}}$$

$$v_k^j = \max_{\lambda} |\hat{\eta}_{P,Q}^{\pi_k} - 1/2|/r_k^j, s_k^j = \tilde{\eta}_{P,Q}^{\pi_k} \text{ with } \lambda = \lambda_k^j \text{ in (16)}$$

end

if $v_k^j \geq 1$ **then**

$k_j = k$, **break**

end

end

end

if **flag** **then**

$p = \lfloor \frac{d}{2+2d} \log_2(n_P \varepsilon^2) \rfloor$. Create π_p on D^Q . Query (5) and (6) with $\varepsilon/2$.

Output: $\tilde{\eta}_{P,Q}^{prune} = \tilde{\eta}_P^{\pi_p}$.

else

Output: $\tilde{\eta}_{P,Q}^{prune} = \sum_j s_k^j \cdot \mathbf{1}\{A_{(p_0)}^j\}$.

end

Algorithm 2 lacks true adaptiveness, given that its upper bound suffers a significant loss factor of $(\alpha + d)/(1 + d)$. Conventional adaptive methods typically incur a loss factor of logarithmic order (Butucea et al., 2020; Cai and Wei, 2021). Following the approach outlined in Butucea et al. (2020), one potential avenue for improvement involves fitting multiple trees for $p = \lfloor \frac{d}{2+2d} \log_2(n_P \varepsilon^2) \rfloor, \dots, \lfloor \frac{d}{2+2d} \log_2(n_P \varepsilon^2 + n_Q^{\frac{2+2d}{d}}) \rfloor$ and designing a test statistic to zero out the suboptimal candidates.

In summary, this proposed scheme offers benefits from two perspectives: (1) It avoids multiple queries to the data. All data holders send their privatized messages at most twice. (2) It selects sensitive parameters p and λ in a data-driven manner. In sacrifice, the bound of excess risk is reasonably loosened compared to the model with the best parameters, and one more query to the private data is required.

6 Experiments on Synthetic Data

We validate our theoretical findings by comparing LPCT with several of its variants on synthetic data in this section. All experiments are conducted on a machine with 72-core Intel Xeon 2.60GHz and 128GB memory. Reproducible codes are available on GitHub¹.

6.1 Simulation Design

We motivate the design of simulation settings. While conducting experiments across various combinations of γ , n_P , and n_Q , our primary focus is on two categories of data: (1) both n_Q and γ are small, i.e. we have a small amount of high-quality public data (2) both n_Q and γ are large, i.e. we have a large amount of low-quality public data. When γ is small but n_Q is large, using solely public data yields sufficient performance, rendering private estimation less meaningful. Conversely, when γ is large but n_Q is small, incorporating public data into the estimation offers limited assistance. Moreover, we require n_P to be much larger than n_Q since practical data collection becomes much easier with a privacy guarantee.

We consider the following pair of distributions P and Q . The marginal distributions $P_X = Q_X$ are both uniform distributions on $\mathcal{X} = [0, 1]^2$. The regression function of P is

$$\eta_P(x) = \frac{1}{2} + \text{sign}(x^1 - \frac{1}{3}) \cdot \text{sign}(x^2 - \frac{1}{3}) \cdot \left(\left| \frac{(3x^1 - 1)(3x^2 - 1)}{10} \right| \cdot \max(0, 1 - |2x^2 - 1|) \right)^{\frac{1}{10}}$$

while the regression function of Q with parameter γ is

$$\eta_Q(x) = \frac{1}{2} + \frac{2}{5} \cdot \left(\frac{5}{2} \cdot \left(\eta_P(x) - \frac{1}{2} \right) \right)^{\gamma}.$$

It can be easily verified that the constructed distributions above satisfy Assumption 1, 2, and 3. Throughout the following experiments, we take γ in 0.5 and 5. For better illustration, their regression functions are plotted in Figure 3.

We conduct experiments with privacy budgets $\varepsilon \in \{0.5, 1, 2, 4, 8, 1000\}$. Here, $\varepsilon = 1000$ can be regarded as the non-private case which marks the performance limitation of our methods. The other budget values cover all commonly seen magnitudes of privacy budgets.

1. <https://github.com/Karlmeyh/LPCT>.




Figure 3: Contour plots of the simulation distributions.

We take classification accuracy as the evaluation metric. In the simulation studies, the comparison methods and their abbreviation are as follows:

- **LPCT** is the proposed algorithm with the max-edge partition rule. We select the parameters among $p \in \{1, \dots, 8\}$ and $\lambda \in \{0.1, 0.5, 1, 2, 5, 10, 50, 100, 200, 300, 400, 500, 750, 1000, 1250, 1500, 2000\}$. We choose the Gini index as the reduction criterion.
- **LPCT-P** and **LPCT-Q** are the estimation associated to P and Q data, respectively. Their parameter grids are the same as LPCT.
- **LPCT-R** is the proposed algorithm using the random max-edge partition (Cai et al., 2023). Specifically, we randomly select an edge among the longest edges to cut at the midpoint. Its parameter grids are the same as LPCT.
- **LPCT-prune** is the pruned locally private classification tree in Algorithm 2. We use constants that yield a tighter bound than (14).

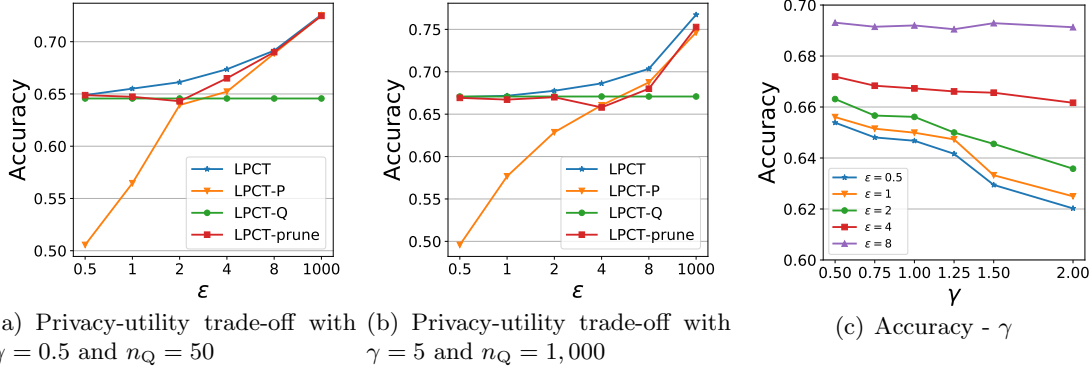
For each model, we report the best result over its parameter grids, with the best result determined based on the average result of at least 100 replications.

6.2 Simulation Results

In Section 6.2.1 and 6.2.2, we analyze the influence of underlying parameters and sample sizes to validate the theoretical findings. In Section 6.2.3, we analyze the parameters of LPCT to understand their behaviors. We illustrate various ways that public data benefits the estimation in Section 6.2.4.

6.2.1 INFLUENCE OF UNDERLYING PARAMETERS

Privacy-utility trade-off We analyze how the privacy budget ε influences the quality of prediction in terms of accuracy. We generate 10,000 private samples, with 50 and 1,000 public samples for $\gamma = 0.5$ and $\gamma = 5$, respectively. The results are presented in Figure 4(a) and 4(b). In both cases, the performances of all models increase as ε increases. Furthermore, the performance of LPCT is consistently better than both LPCT-P and LPCT-Q, demonstrating the effectiveness of our proposed method in combining both data. LPCT-prune performs reasonably worse than LPCT. In the high privacy regime ($\varepsilon \leq 1$), LPCT-P is almost non-informative and our estimator performs similarly to LPCT-Q. In the low privacy

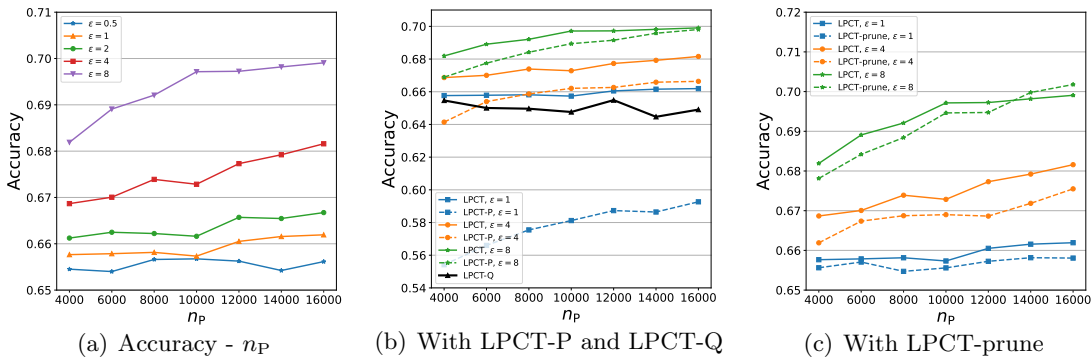

 Figure 4: Illustration of relationship between accuracy and underlying parameters, i.e. ϵ and γ .

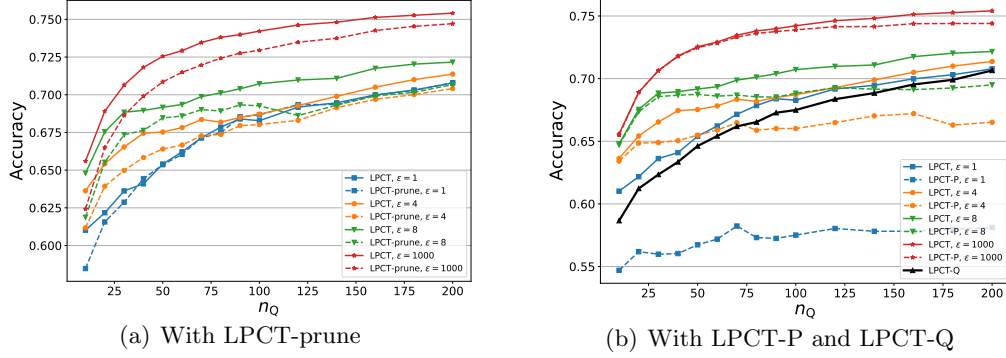
regime ($\epsilon \geq 8$), the performance of our estimator improves rapidly along with LPCT-P. In the medium regime, neither LPCT-P nor LPCT-Q obviously outperforms each other, and there is a clear performance gain from using LPCT.

Influence of γ We analyze how the relative signal exponent γ influences the performance. We set $\gamma = 0.5$, $n_P = 10,000$ and $n_Q = 50$, while varying γ in $\{0.5, 0.75, 1, 1.25, 1.5, 2\}$. In Figure 4(c), we can see that the accuracy of LPCT is decreasing with respect to γ for all ϵ , which aligns with Theorem 4. Since γ only influences the performance through Q data, its effect is less apparent when the P data takes the dominance, for instance $\epsilon = 8$.

6.2.2 INFLUENCE OF SAMPLE SIZES

Accuracy - n_P We conduct experiments to investigate the influence of the private sample size n_P . We fix $n_Q = 50$ and $\gamma = 0.5$ while varying n_P from 4,000 to 16,000. The results are presented in Figure 5(a). Additionally, we compare the performance of LPCT-P and LPCT-Q to analyze the contribution of both data sets, as shown in Figure 5(b). Moreover, we compare LPCT and LPCT-prune in Figure 5(c). We refrain from invoking $\tilde{\eta}_P^{\pi_P}$ even if the `flag` = 1 condition is triggered, as the termination induces sudden changes in the continuous curves, thereby rendering the findings less evident. In Figure 5(a), except for


 Figure 5: Analysis of accuracy with respect to n_P .


 Figure 6: Analysis of accuracy with respect to n_Q .

$\varepsilon = 0.5$, the accuracy improves as n_P increases for all ε values. Moreover, the rate of improvement is more pronounced for higher values of ε . This observation aligns with the convergence rate established in Theorem 4. In 5(b), LPCT always outperforms methods using single data set, i.e. LPCT-P and LPCT-Q. The improvement of LPCT over LPCT-Q is more significant for large ε . In Figure 5(c), we justify that LPCT-prune performs analogously to LPCT and its performance improves as ε and n_P increase.

Accuracy - n_Q We conduct experiments to investigate the influence of the private sample size n_Q . We keep $n_P = 10,000$ and $\gamma = 0.5$ while varying n_Q from 10 to 200. The classification accuracy monotonically increases with respect to n_Q for all methods. In Figure 6(a), the performance of LPCT-prune is closer to LPCT when n_Q gets larger. As explained in Section 5.2, the pruning is appropriately conducted when Q data takes the dominance, and is disturbed when P data takes the dominance. In Figure 6(b), it is worth noting that the performance improvement is more significant for small values of n_Q , specifically when $n_Q \leq 50$. Moreover, the performance of LPCT-P also improves in this region and remains approximately unchanged for larger n_Q . This can be attributed to the information gained during partitioning, which is further validated in Figure 9.

6.2.3 ANALYSIS OF MODEL PARAMETERS

Parameter analysis of depth We conduct experiments to investigate the choice of depth. Under $(n_P, n_Q) = (10000, 50)$ and $\gamma = 0.5$, we plot the best performance of LPCT and LPCT-prune when fixing p and p_0 to certain values, respectively in Figure 7(a) and 7(b). To avoid obscuring empirical findings, we refrain from invoking $\tilde{\eta}_P^{\pi_P}$ even if `flag` = 1 is triggered, since the experiment is conducted with respect to p_0 . In Figure 7(a), the accuracy first increases as p increases until it reaches a certain value. Then the accuracy begins to decrease as p increases. This confirms the trade-off observed in Theorem 4. Moreover, the p that maximizes the accuracy is increasing with respect to ε . This is compatible with the theory since the optimal choice $p \asymp \log(n_P \varepsilon^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}})$ is monotonically increasing with respect to ε . As in Figure 7(b), the accuracy increases with p_0 at first, whereas the performance remains unchanged or decreases slightly for p_0 larger than the optimum. Note

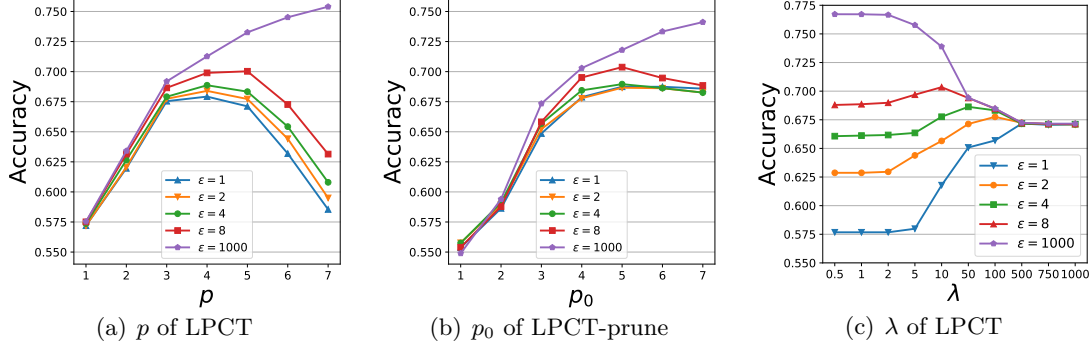


Figure 7: Parameter analysis.

that the axes of 7(a) and 7(b) are aligned. LPCT-prune behaves analogously to LPCT in terms of accuracy and optimal depth choice.

Parameter analysis of λ We conduct experiments to investigate the choice of λ in terms of accuracy. Under $(n_P, n_Q) = (10000, 1000)$ and $\gamma = 5$, we plot in Figure 7(c) the best performance of LPCT when fixing λ to certain values. In Figure 7(c), the relation between accuracy and λ is inverted U-shaped for $\epsilon = 2, 4, 8$, while is monotonically increasing and decreasing for $\epsilon = 1$ and 1000, respectively. The results indicate that a properly chosen λ is necessary as stated in Theorem 4. Moreover, the λ that maximizes the accuracy is decreasing with respect to ϵ , which justify the fact that we should rely more on public data with higher privacy constraint.

6.2.4 BENEFITS OF PUBLIC DATA

Range parameter We demonstrate the effectiveness of range estimation using public data when the domain is unknown. Consider the case where $\mathcal{X} = [0, a]^d$ and a is a random variable uniformly distributed in $[1, 5]$. We compare the following scenarios: (I) a is known. (II) a is estimated using public data as shown in Theorem 5. (III) a is unknown. Case (I) is equivalent to $\mathcal{X} = [0, 1]^d$. In case (II), we min-max scale the samples from the estimated




Figure 8: Analysis of range parameter. In 8(a), we have $a = 2.9$. The black square is the range estimated by public data in case (II). The grey squares represent the grid of \bar{a} in case (III).

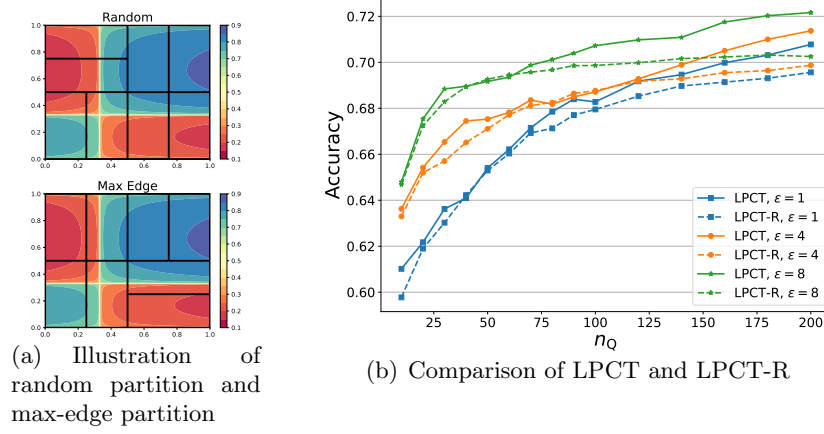


Figure 9: Illustration of partition with public data.

range to $[0, 1]^d$. In case (III), we scale the samples from $[0, \bar{a}]^d$ to $[0, 1]^d$ and select the best \bar{a} among $\bar{a} \in \{1, 1.5, \dots, 5\}$. In Figure 8(a), we showcase the scaling schemes in (II) and (III). The estimated range from public data is much more concise than the best \bar{a} . Given no prior knowledge, the \bar{a} chosen in practice can perform significantly worse. In Figure 8(b) and 8(c), we observe that the performance of the estimated range matches that of (I) as n_Q increases, whereas the unknown range significantly degrades performance. The performance gap between the known and estimated range is more substantial when n_Q is small. This observation is compatible with the conclusion drawn in Theorem 5 emphasizing the need for sufficient public data to resolve the unknown range issue.

Partition We demonstrate the effectiveness of partitioning with information of public data by comparing LPCT with LPCT-R. In Figure 9(a), the showcase of $p = 3$ is presented. The random partition is possibly created as the top figure, where the upper left and the lower right cells are hard to distinguish as 0 or 1. Using sufficient public data, the max-edge partition is bound to create an informative partition like the lower figure. Moreover, Figure 9(b) shows that LPCT outperforms LPCT-R.

7 Experiments on Real Data

7.1 Data Description

We conduct experiments on real-world data sets to show the superiority of LPCT. We utilize 11 real data sets that are available online. Each of these data sets contains sensitive information and necessitates privacy protection. Furthermore, these data sets cover a wide range of sample sizes and feature dimensions commonly encountered in real-world scenarios. When creating D^P and D^Q from the original data set D , some data sets have explicit criteria for splitting, reflecting real-world considerations. In these cases, we have $P \neq Q$. For others, the split is performed randomly and we have $P = Q$. A summary of key information for these data sets after pre-processing can be found in Table 1. Following the analysis in

Table 1: Information of real data sets.

Dataset	P = Q	n_P	n_Q	d	Area
Anonymity	✓	453	50	14	Social
Diabetes	✓	2,054	200	8	Medical
Email	✓	3,977	200	3,000	Text
Employ	✓	2,225	200	8	Business
Rice	✓	240	3,510	7	Image
Census	×	41,292	3,144	46	Social
Election	×	1,752	222	9	Social
Employee	×	3,319	504	9	Business
Jobs	×	46,218	14,318	11	Social
Landcover	×	7,097	131	28	Image
Taxi	×	3,297,639	117,367	93	Social

Section 4.3, we min-max scale each data to $[0, 1]^d$ according to public data. Detailed data information and pre-processing procedures are provided in Appendix C.1.

7.2 Competitors

We compare LPCT with several state-of-the-art competitors. For each model, we report the best result over its parameter grids based on an average of 20 replications. The tested methods are

- **LPCT** is described in Section 6. In real data experiments, we enlarge the grids to $p \in \{1, 2, 3, 4, 5, 6, 7, 8, 10, 12, 14, 16\}$ and $\lambda \in \{0.1, 0.5, 1, 2, 5, 10, 50, 100, 200, 300, 400, 500, 750, 1000, 1250, 1500, 2000\}$ due to the large sample size. Since the max-edge rule is easily affected by useless features or categorical features (Ma et al., 2023), we boost the performance of LPCT by incorporating the criterion reduction scheme from the original CART (Breiman, 1984) to the tree construction. Since Proposition 2 holds for any partition, the method is also ε -LDP. We choose the Gini index as the reduction criterion.
- **LPCT-prune** is the pruned locally private classification tree in Algorithm 2. We also adopt the criterion reduction scheme for tree construction.
- **CT** is the original classification tree (Breiman, 1984) with the same grid as LPCT. We define two approaches, **CT-W** and **CT-Q**. For CT-W, we fit CT on the whole data, namely P and Q data together, with no privacy protection. Its accuracy will serve as an upper bound. For CT-Q, we fit CT on Q data only. Its performance is taken into comparison. We use implementation by `sklearn` (Pedregosa et al., 2011) with default parameters while varying `max_depth` in $\{1, 2, \dots, 16\}$.
- **PHIST** is the locally differentially private histogram estimator proposed by Berrett and Butucea (2019). We fit PHIST on private data solely as it does not leverage public data. We choose the number of splits h^{-1} along each axis in $\{1, 2, 3, 4, 5, 6\}$.
- **LPDT** is the locally private decision tree originally proposed for regression (Ma et al., 2023). The approach creates a partition on Q data and utilizes P data for estimation. We adapt the strategy to classification. We select p in the same range as LPCT.

- **PGLM** is the locally private generalized linear model proposed by Wang et al. (2023a) with logistic link function. Note that their privacy result is under (ε, δ) -LDP. We let $\delta = 1/n_P^2$ for a fair comparison.
- **PSGD** is the locally private stochastic gradient descent developed in Lowy et al. (2023). The presented algorithm is in a sequentially interactive manner, meaning that the privatized information from X_i depends on the information from X_{i-1} . We adopt the warm start paradigm where we first find a minimizer on the public data and then initialize the training process with the minimizer. Following Lowy et al. (2023), we choose the parameters of the privacy mechanism, i.e. PrivUnitG, by the procedure they provided. We let $C = 1$ and learning rate $\eta \in \{10^{-5}, 10^{-4}, \dots, 1\}$. We consider two models: **PSGD-L** is the linear classifier using the logistic loss; **PSGD-N** is the non-linear neural network classifier with one hidden layer and ReLU activation function. The number of neurons is selected in $d \in \{d/2, d, 2d\}$.

7.3 Experiment Results

We first illustrate the substantial utility gain brought by public data. We choose **Diabetes** and **Employee**, representing the case of $P = Q$ and $P \neq Q$, respectively. On each data, we randomly select a portion of private data and 80 public samples to train three models: LPCT, LPDT, and PHIST. PHIST does not leverage any information of public data, while LPDT benefits from public data only through the partition. LPCT utilizes public data in both partition and estimation. In Figure 10, the accuracy increases with respect to the extent of public data utilization for both identically and non-identically distributed public data. Moreover, the accuracy gain is significant. For instance, in Figure 10(b), 80 public samples and $0.2 \times 3319 \approx 664$ private samples achieve the same accuracy as using $0.6 \times 3319 \approx 1991$ private samples alone. The observation yields potential practical clues. For instance, when data-collecting resources are limited, organizations can put up rewards to get a small amount of public data rather than collecting a large amount of private data.

Next, we report the averaged accuracy for $\varepsilon \in \{0.5, 2, 8\}$ in Table 2 on each real data set. Here, CT-W is not private and serves as a reference. To ensure significance, we adopt the Wilcoxon signed-rank test (Wilcoxon, 1992) with a significance level of 0.05 to check if

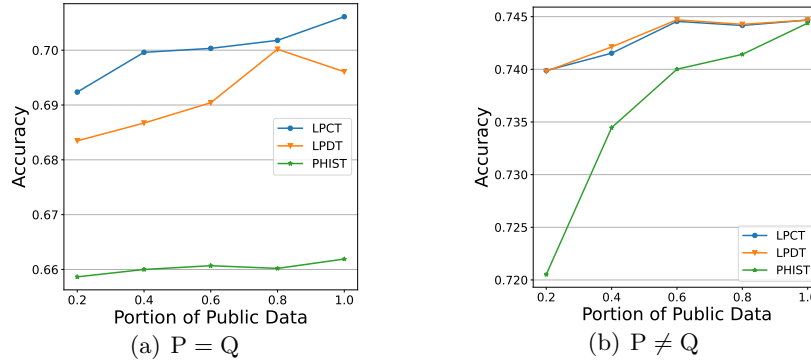


Figure 10: Illustration of utility gain brought by public data.

the result is significantly better. For each data set, trials that exceed a memory of 100GB or a running time of one hour are terminated.

From Table 2, we can see that LPCT-based methods generally outperform their competitors, both in the sense of average performance (rank sum) and best performance (number of No.1). Compared to LPCT, LPCT-prune performs reasonably worse due to its data-driven nature. In practice, however, all other methods require parameter tuning which limits the privacy budget for the final estimation. Thus, their real performance is expected to drop down while that of LPCT-prune remains. Methods with the CART rule substantially outperform those with the max-edge rule, as the CART rule is robust to useless features or categorical features and can fully employ data information. We note that on some data sets, the results of LPCT remain the same for all ϵ . This is the case where public data takes dominance and the estimator is barely affected by private data. Moreover, we observe that linear classifiers such as PSDG-L and PGLM perform poorly on some data sets. This is attributed to the non-linear underlying nature of these data sets.

We compare the efficiency of the methods by reporting the total running time in Table 3. PSGD-L and PSGD-N are generally slower and even take more than one hour for a single round on `taxi` data. Compared to these sequentially-interactive LDP methods, non-interactive methods achieve acceptable running time. We argue that the performance gap between LPCT and CT is caused by `cython` acceleration in `sklearn`. Moreover, PHIST causes memory leakage on three data sets. Though histogram-based methods and tree-based methods possess the same storage complexity theoretically (Ma et al., 2023), LPCT avoids memory leakage when d is large. The results show that LPCT is an efficient method.

8 Conclusion and Discussions

In this work, we address the under-explored problem of public data assisted LDP learning by investigating non-parametric classification under the posterior drift assumption. We propose the locally private classification tree that achieves the mini-max optimal convergence rate. Moreover, we introduce a data-driven pruning procedure that avoids parameter tuning and produces an estimator with a fast convergence rate. Extensive experiments on synthetic and real data sets validate the superior performance of these methods. Both theoretical and experimental results suggest that public data is significantly more effective compared to data under privacy protection. This finding leads to practical suggestions. For instance, prioritizing non-private data collection should be considered. When data-collecting resources are limited, organizations can offer rewards to encourage users to share non-private data instead of collecting a large amount of privatized data.

We summarize the ways that public data contributes to LPCT. The benefits lie in four-fold: (1) Elimination of range parameter. Since the domain is unknown in the absence of the raw data, the convergence rate and performance of partition-based estimators are degraded by a range parameter. For LPCT, public data helps to eliminate the range parameter both theoretically and empirically, as demonstrated in Section 4.3 and 6.2.4, respectively. (2) Implicit information in partition. Through the partition created on public data, LPCT implicitly feeds information from public data to the estimation. The performance gain is explained empirically in Section 6.2.4. (3) Augmented estimation. In the weighted average estimation, the public data reduces the instability caused by perturbations and

Table 2: Accuracy (10^{-1}) over real data sets. The best results are marked with *. The results that are not significantly worse than the best result are **bolded**. Due to memory limitation, PHIST is corrupted on three data sets. PSGD-L and PSGD-N exceed the running time of 3,600 seconds on one data set and are terminated. These results are marked with -.

(a) $\varepsilon = 0.5$

Datasets	CT-W	LPCT		LPCT-prune		CT-Q	PHIST	LPDT	PGLM	PSGD-L	PSGD-N
		Maxedge	CART	Maxedge	CART						
Anonymity	8.346	8.346*	8.346*	8.309	8.309	8.346*	8.346*	7.546	5.899	7.790	8.123
Diabetes	9.964	7.067	7.590*	6.661	7.529	7.301	6.676	6.305	5.499	6.582	6.448
Email	9.261	7.332	7.487*	7.205	7.185	7.283	-	7.112	5.136	6.435	6.292
Employ	9.048	7.403	8.646*	6.684	6.362	8.577	5.704	5.644	4.874	5.742	5.595
Rice	8.992	9.039	9.183*	9.053	9.173	9.117	5.467	6.583	4.886	8.517	8.142
Census	8.685	8.539	8.743*	8.471	8.528	8.285	-	7.195	5.104	8.310	8.405
Election	6.408	5.610	5.846	5.259	5.773	5.958*	5.060	5.089	4.948	5.207	5.074
Employee	7.648	7.397*	7.131	7.381	7.040	7.151	6.660	6.098	5.280	6.493	6.387
Jobs	7.863	5.628	7.814*	5.561	7.793	7.599	5.327	5.441	4.942	5.530	5.732
Landcover	9.404	8.615	8.960*	8.299	8.719	6.670	8.362	8.366	4.955	8.368	8.349
Taxi	9.601	5.647	9.589*	5.637	9.589*	9.490	-	5.438	4.912	-	-
Rank Sum	36	18*	54	36	38	83	84	106	70	80	
No. 1	2	10*	2	5	5	1	0	0	0	0	

(b) $\varepsilon = 2$

Datasets	CT-W	LPCT		LPCT-prune		CT-Q	PHIST	LPDT	PGLM	PSGD-L	PSGD-N
		Maxedge	CART	Maxedge	CART						
Anonymity	8.346	8.346*	8.346*	8.086	8.168	8.346*	8.346*	8.346*	5.899	7.965	8.039
Diabetes	9.964	7.274	7.590*	6.661	7.412	7.301	6.676	6.676	5.499	6.653	6.674
Email	9.261	7.350	7.485*	7.224	7.459	7.283	-	7.181	5.136	7.136	7.203
Employ	9.048	7.702	9.121*	7.005	6.578	8.577	6.174	7.181	4.874	5.786	5.792
Rice	8.992	9.027	9.183*	9.053	9.183*	9.117	5.467	8.458	4.886	9.050	8.750
Census	8.685	8.537	8.743*	8.460	8.633	8.285	-	8.017	5.104	8.432	8.471
Election	6.408	5.648	5.850	5.327	5.595	5.958*	5.056	5.114	4.948	5.323	5.111
Employee	7.648	7.437*	7.175	7.362	7.175	7.151	6.472	6.951	5.280	6.400	6.383
Jobs	7.863	5.628	7.821*	5.597	7.778	7.599	5.333	5.591	4.942	6.788	5.934
Landcover	9.404	8.715	8.960*	8.623	8.734	6.670	8.367	8.493	4.955	8.553	8.369
Taxi	9.601	5.659	9.590*	5.643	9.580	9.490	-	5.536	4.912	-	-
Rank Sum	36	18*	54	36	38	83	84	106	70	80	
No. 1	3	10*	1	5	4	1	1	0	1	0	

(c) $\varepsilon = 8$

Datasets	CT-W	LPCT		LPCT-prune		CT-Q	PHIST	LPDT	PGLM	PSGD-L	PSGD-N
		Maxedge	CART	Maxedge	CART						
Anonymity	8.346	8.346*	8.346*	8.221	8.309	8.346*	8.346*	8.346*	5.395	8.338	8.154
Diabetes	9.964	7.503	7.930*	6.648	7.433	7.301	6.676	7.496	5.588	7.244	7.101
Email	9.261	7.355	7.520	7.251	7.491	7.283	-	7.326	5.271	8.446	7.060
Employ	9.048	7.619	9.068*	7.393	7.492	8.577	9.043	7.572	5.237	5.795	5.798
Rice	8.992	9.011	9.183*	9.053	9.137	9.117	8.858	9.017	6.283	9.017	8.875
Census	8.685	8.557	8.743*	8.404	8.671	8.285	-	8.527	5.389	8.530	8.154
Election	6.408	5.894	6.082*	5.321	5.563	5.958	5.015	5.892	5.462	5.765	5.705
Employee	7.648	7.441*	7.173	7.384	7.105	7.151	7.437	7.439	4.930	6.775	6.883
Jobs	7.863	5.650	7.861*	5.584	7.794	7.599	5.992	5.641	6.802	7.712	7.414
Landcover	9.404	8.737	8.963	8.620	8.745	6.670	8.367	8.695	5.551	9.137*	8.948
Taxi	9.601	5.666	9.610*	5.637	9.579	9.490	-	5.620	5.087	-	-
Rank Sum	42	18*	66	36	49	90	68	104	54	78	
No. 1	2	8*	1	4	4	2	2	0	3	0	

Table 3: Running time (s) on real data sets.

Datasets	CT-W	LPCT		LPCT-prune		CT-Q	PHIST	LPDT	PGLM	PSGD-L	PSGD-N
		Maxedge	CART	Maxedge	CART						
Anonymity	3e-3	1e-2	6e-3	9e-3	1e-2	2e-3	1e-2	2e-3	2e-3	8e-1	1e0
Diabetes	1e-2	1e-1	3e-2	3e-2	3e-2	3e-3	3e-2	6e-3	3e-3	3e0	6e0
Email	3e0	1e0	3e-2	2e-1	1e0	6e-2	-	1e0	4e0	8e0	3e3
Employ	5e-3	9e-2	2e-1	1e0	9e-1	3e-3	3e-2	4e-2	4e-3	5e0	7e0
Rice	1e-2	2e-2	1e-2	1e-2	2e-2	2e-2	5e-3	6e-3	8e-3	2e0	2e0
Census	6e-1	1e0	5e0	6e-1	9e-1	3e-1	-	2e-1	5e-2	2e1	3e1
Election	1e-2	6e-2	3e-2	3e-2	3e-2	3e-3	1e0	3e-2	3e-3	3e0	5e0
Employee	6e-3	8e-2	1e-1	4e-2	1e-1	3e-3	5e-2	5e-2	4e-3	7e0	1e2
Jobs	1e-1	1e0	1e0	6e-1	6e-1	3e-2	8e-1	5e-1	3e-2	1e2	1e2
Landcover	2e-1	1e-1	8e-2	9e-2	1e-1	5e-3	2e-1	9e-2	5e-3	1e1	2e1
Taxi	1e2	2e2	7e1	4e1	4e1	1e0	-	1e2	1e0	-	-

produces more accurate estimations. Theoretically, the convergence rate can be significantly improved. Empirically, even a small amount of public data improves the performance significantly (Section 6 and 7). (4) Data-driven pruning. As shown in Section 5.2, we are capable of deriving the data-driven approach only with the existence of public data.

We discuss the potential future works and improvements. The assumption that P_X and Q_X are identical can be mitigated. Similar to Ma et al. (2023), we only need the density ratio to be bounded. Moreover, the marginal densities are not necessarily lower bounded, if we add a parameter to guarantee that sufficient samples are contained in each cell (Ma et al., 2023). In practice, \mathcal{X}_P and \mathcal{X}_Q can be distinctive (Yu et al., 2021b; Gu et al., 2023). This problem is related to domain adaption (Redko et al., 2020) and can serve as future work direction. Determining whether a data set satisfies Assumption 3 is of independent interest, as demonstrated in Gu et al. (2023); Yu et al. (2023).

Acknowledgments and Disclosure of Funding

We would like to thank the reviewers and action editor for their help and advice, which led to a significant improvement of the article. The research is supported by the Special Funds of the National Natural Science Foundation of China (Grant No. 72342010). Yuheng Ma is supported by the Outstanding Innovative Talents Cultivation Funded Programs 2024 of Renmin University of China. This research is also supported by Public Computing Cloud, Renmin University of China.

Appendix A. Methodology

A.1 Estimator Summarization

Function	Meaning	Defined
η_P^*, η_Q^*	true regression function of P and Q	(1)
$\bar{\eta}_P^\pi$	population tree estimation	(3)
$\hat{\eta}_P^\pi, \hat{\eta}_Q^\pi$	non-private tree estimation of P and Q	(4)
$\tilde{\eta}_P^\pi$	private tree estimation based on P	(7)
$\tilde{\eta}_{P,Q}^\pi$	private tree estimation	(8)
$\hat{\eta}_{P,Q}^\pi$	non-private two-sample tree estimation	(9)
$\hat{\eta}_{P,Q}^{\pi_k}$	ancestor private tree estimation	(13)
$\tilde{\eta}_{P,Q}^{prune}$	pruned private tree estimation	Algorithm 2

A.2 Derivation of the Closed-form Solution for Pruning Procedure

Remind that

$$r_k^j := \sqrt{\frac{4C_1 (2^{p_0-k+3}\varepsilon^{-2}n_P) \vee \left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} \right) + C_1 \lambda^2 \sum_{j'} \sum_{i=1}^{n_P} U_i^{Qj'}}{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj} \right)^2}}.$$

For each $\text{ancestor}(A_{(p_0)}^j, k)$, define the quantity

$$w_k^j = \frac{\sum_{\text{ancestor}(A_{(p_0)}^{j'}, k) = \text{ancestor}(A_{(p_0)}^j, k)} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}}{\sum_{\text{ancestor}(A_{(p_0)}^{j'}, k) = \text{ancestor}(A_{(p_0)}^j, k)} \left(\sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj'} \right)}.$$

Also, let $u_k = 4 (2^{p_0-k+3}\varepsilon^{-2}n_P) \vee \left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} \right)$. When the signs of $\tilde{\eta}_P^{\pi_k}(x) - \frac{1}{2}$ and $\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}$ are identical, the object in (16) becomes

$$\begin{aligned} & \sqrt{C_1} \cdot \frac{|\tilde{\eta}_{P,Q}^{\pi_k}(x) - \frac{1}{2}|}{r_k^j} \\ &= \frac{\left| w_k^j \left(\tilde{\eta}_P^{\pi_k}(x) - \frac{1}{2} \right) + (1 - w_k^j) \cdot \left(\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right) \right|}{\sqrt{\frac{w_k^{j2} u_k}{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} \right)^2} + \frac{(1-w_k^j)^2}{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'}}}} \\ &\leq \sqrt{\frac{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} \right)^2}{u_k} \left(\tilde{\eta}_P^{\pi_k}(x) - \frac{1}{2} \right)^2 + \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \cdot \left(\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right)^2} \end{aligned}$$

where the inequality follows from Cauchy's inequality. The maximum is attained when

$$\begin{aligned} & \frac{w_k^{j^2} u_k}{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}\right)^2} \cdot \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \cdot \left(\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}\right)^2 \\ &= \frac{(1 - w_k^j)^2}{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'}} \cdot \frac{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}\right)^2}{u_k} \left(\hat{\eta}_P^{\pi_k}(x) - \frac{1}{2}\right)^2. \end{aligned}$$

By bringing in w_k^j , this leads to the solution

$$\lambda = \frac{u_k}{\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}} \cdot \frac{\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}}{\hat{\eta}_P^{\pi_k}(x) - \frac{1}{2}}.$$

When $\hat{\eta}_P^{\pi_k}(x) - \frac{1}{2}$ and $\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}$ have different signs, the object in (16) is maximized at the extremum, i.e. when $\lambda = 0$ or $\lambda = \infty$. The maximum value is

$$\sqrt{\frac{\max\left(\frac{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}\right)^2}{u_k} \left(\hat{\eta}_P^{\pi_k}(x) - \frac{1}{2}\right)^2, \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \cdot \left(\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}\right)^2\right)}{C_1}}.$$

Appendix B. Proofs

B.1 Proof of Proposition 2

Proof of Proposition 2 Since all mechanisms R_i are completely the same, we use R for notation simplicity. For each conditional distribution $R\left(\tilde{U}_i, \tilde{V}_i | X_i = x, Y_i = y\right)$, we can compute the density ratio as

$$\begin{aligned} & \sup_{x, x' \in \mathbb{R}^d, y, y' \in \{0,1\}, S \in \sigma(S)} \frac{R\left((\tilde{U}, \tilde{V}) \in S | X_i = x, Y_i = y\right)}{R\left((\tilde{U}, \tilde{V}) \in S | X_i = x', Y_i = y'\right)} \\ & \leq \sup_{x, x' \in \mathbb{R}^d} \frac{dR(\tilde{U} | X_i = x)}{dR(\tilde{U} | X_i = x')} \cdot \sup_{x, x' \in \mathbb{R}^d, y, y' \in \{0,1\}} \frac{dR(\tilde{V} | X_i = x, Y_i = y)}{dR(\tilde{V} | X_i = x', Y_i = y')}. \end{aligned} \quad (17)$$

The first part is characterized by the Laplace mechanism. Since the conditional density is identical if x and x' belongs to a same $A_{(p)}^j$, we have

$$\sup_{x, x' \in \mathbb{R}^d} \frac{dR(\tilde{U} | X_i = x)}{dR(\tilde{U} | X_i = x')} = \sup_{u, u' \in \{0,1\}^{|I|}} \frac{dR(\tilde{U} | U_i = u)}{dR(\tilde{U} | U_i = u')}.$$

The right-hand side can be computed as

$$\frac{dR(\tilde{U} | U_i = u)}{dR(\tilde{U} | U_i = u')} = \frac{\prod_j \exp\left(\frac{\varepsilon}{4} |\tilde{u}_i^j - u^j|\right)}{\prod_j \exp\left(\frac{\varepsilon}{4} |\tilde{u}_i^j - u'^j|\right)} \leq \prod_j \exp\left(\frac{\varepsilon}{4} |u^j - u'^j|\right).$$

By definition, u and u' are one-hot vectors and differ at most on two entries. As a result,

$$\sup_{x, x' \in \mathbb{R}^d} \frac{dR(\tilde{U}|X_i = x)}{dR(\tilde{U}|X_i = x')} \leq \exp\left(\frac{\varepsilon}{2}\right) \quad (18)$$

For the second part, note that v and v' also differ at most on two entries. Then with analogous steps, we have

$$\sup_{x, x' \in \mathbb{R}^d, y, y' \in \{0, 1\}} \frac{dR(\tilde{V}|X_i = x, Y_i = y)}{dR(\tilde{V}|X_i = x', Y_i = y')} \leq \exp\left(\frac{\varepsilon}{2}\right). \quad (19)$$

Bringing (18) and (19) into (17) yields the desired conclusion. \blacksquare

B.2 Proof of Theorem 3

Proof of Theorem 3 The proof is analogous to Cai and Wei (2021, Theorem 3.2). We first construct two classes of distributions P_σ and Q_σ for $\sigma \in \{-1, 1\}^m$. We then combine the information inequality under privacy constraint (Duchi et al., 2018) and Assouad's Lemma (Tsybakov, 2009). We first define two classes of distributions by construction. Define the useful quantity

$$r = c_r \left(n_P (e^\varepsilon - 1)^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{1}{2\alpha+2d}}, \quad w = c_w r^d, \quad m = \lfloor c_m r^{\alpha\beta-d} \rfloor \quad (20)$$

where c_r , c_w , and c_m are some universal constants. Let

$$G = \{(6k_1r, 6k_2r, \dots, 6k_dr), k_i = 1, 2, \dots, \lfloor (6r)^{-1} \rfloor, i = 1, 2, \dots, d\}.$$

be a grid of $\lfloor (6r)^{-1} \rfloor^d$ points in the unit cube. Denote $x_1, x_2, \dots, x_{\lfloor (6r)^{-1} \rfloor^d}$ as points in G . Since $m \ll \lfloor (6r)^{-1} \rfloor^d$, we are interested in $B(x_k, 2r)$, $k = 1, 2, \dots, m$ which are balls with radius $2r$ centered at x_k . Let $B_c = [0, 1]^d \setminus \bigcup_{k=1}^m B(x_k, 2r)$ denote the points that aren't in any of the m balls. Define function $g(\cdot)$ on $[0, \infty)$:

$$g(z) = \begin{cases} 1 & 0 \leq z < 1 \\ 2 - z & 1 \leq z < 2 \\ 0 & z \geq 2 \end{cases}$$

And define

$$h_P(z) = C_\alpha r^\alpha g^\alpha(z/r), \quad h_Q(z) = C_\gamma C_\alpha^\gamma r^{\alpha\gamma} g^{\alpha\gamma}(z/r)$$

where C_α, C_γ are some universal constants. Define the hypercube \mathcal{H} of pairs (P_σ, Q_σ) by

$$\mathcal{H} = \{(P_\sigma, Q_\sigma), \sigma = (\sigma_1, \sigma_2, \dots, \sigma_m) \in \{-1, 1\}^m\}$$

with P_σ and Q_σ defined as follows. For conditional density, we let

$$\begin{aligned}\eta_{P,\sigma}(x) &= \begin{cases} \frac{1}{2}(1 + \sigma_k h_P(\|x - x_k\|)) & \text{if } x \in B(x_k, 2r) \text{ for some } k = 1, 2, \dots, m \\ \frac{1}{2} & \text{if } x \in B_c. \end{cases} \\ \eta_{Q,\sigma}(x) &= \begin{cases} \frac{1}{2}(1 + \sigma_k h_Q(\|x - x_k\|)) & \text{if } x \in B(x_k, 2r) \text{ for some } k = 1, 2, \dots, m \\ \frac{1}{2} & \text{if } x \in B_c. \end{cases}\end{aligned}$$

Let $P_{\sigma,X}, Q_{\sigma,X}, \sigma \in \{-1, 1\}^m$ all have the same marginal distribution on \mathcal{X} , with density $dP_X(x)$ defined as follows:

$$dP_X(x) = \begin{cases} w \cdot r^{-d} & \text{if } x \in B(x_k, r) \text{ for some } k = 1, 2, \dots, m \\ c_{B_c}(1 - mw) & \text{if } x \in B_c \\ 0 & \text{otherwise.} \end{cases}$$

where c_{B_c} is a normalizing constant that makes dP well-defined. Similar to Cai and Wei (2021), with a proper choice of all constants, \mathcal{H} is a subset of \mathcal{F} that satisfies 3, 1, and 2. Then it suffices to show that

$$\inf_{\hat{f}} \sup_{\mathcal{H}} \mathbb{E}_{X,Y} [\mathcal{R}_{P_\sigma}(\hat{f}) - \mathcal{R}_{P_\sigma}^*] \gtrsim \left(n_P (e^\varepsilon - 1)^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{(\beta+1)\alpha}{2\alpha+2d}}.$$

Let $TV(\cdot, \cdot)$ and $H(\cdot, \cdot)$ denote the total variation distance and the Hellinger distance between measures, respectively. Let $Ham(\cdot, \cdot)$ denote the hamming distance between ordered indices. Let σ and σ' be two indices in $\{-1, 1\}^m$ such that $Ham(\sigma, \sigma') = 1$, i.e. they differ at only one element σ_k and σ'_k . We compute the total variation between the induced P_σ and $P_{\sigma'}$. By Scheffe's theorem (Tsybakov, 2009), there holds

$$TV(P_\sigma, P_{\sigma'}) = \frac{1}{2} \int_{\mathcal{X} \times \mathcal{Y}} |dP_\sigma(x, y) - dP_{\sigma'}(x, y)|.$$

Since the marginal distributions for both measures are identical, we have

$$\begin{aligned}TV(P_\sigma, P_{\sigma'}) &= \int_{\mathcal{X}} |\eta_{P,\sigma}(x) - \eta_{P,\sigma'}(x)| dP_{\sigma,X}(x) \\ &= \int_{B(x_k, r)} C_\alpha r^\alpha \cdot w \cdot r^{-d} dx = C_\alpha w r^\alpha.\end{aligned}$$

According to Cai and Wei (2021), we have

$$H^2(Q_\sigma, Q_{\sigma'}) = C_\gamma^2 C_\alpha^{2\gamma} w r^{2\alpha\gamma}. \quad (21)$$

Let conditional distribution $R = \{R_i\}_{i=1}^{n_P}$ be any transformation on $\mathcal{X}^{n_P} \times \mathcal{Y}^{n_P}$ that is sequentially-interactive ε -local differentially private as defined in Definition 1. Denote $RP_\sigma^{n_P}$ by the random variable mapped compositionally through R and $P_\sigma^{n_P}$. Therefore, for $Ham(\sigma, \sigma') = 1$, there holds

$$H^2(RP_\sigma^{n_P} \times Q_\sigma^{n_Q}, RP_{\sigma'}^{n_P} \times Q_{\sigma'}^{n_Q}) \leq H^2(RP_\sigma^{n_P}, RP_{\sigma'}^{n_P}) + H^2(Q_\sigma^{n_Q} \times Q_{\sigma'}^{n_Q}). \quad (22)$$

For the first term in (22), combining Tsybakov (2009, Lemma 2.4) and Duchi et al. (2018, Corollary 3) yields

$$H^2(\text{RP}_{\sigma'}^{n_P}, \text{RP}_{\sigma}^{n_P}) \leq 4(e^\alpha - 1)^2 n_P TV^2(P_\sigma, P_{\sigma'}) = 4C_\alpha^2 (e^\alpha - 1)^2 n_P w^2 r^{2\alpha}.$$

This together with (20) lead to

$$\begin{aligned} H^2(\text{RP}_{\sigma'}^{n_P}, \text{RP}_{\sigma}^{n_P}) &\leq 4C_\alpha^2 c_w^2 (e^\varepsilon - 1)^2 n_P r^{2\alpha+2d} \\ &= 4C_\alpha^2 c_w^2 c_r^{2\alpha+2d} (e^\varepsilon - 1)^2 n_P \left(n_P (\varepsilon - 1)^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-1} \leq 4C_\alpha^2 c_w^2 c_r^{2\alpha+2d}. \end{aligned} \quad (23)$$

For the second term in (22), (21) implies

$$H^2(Q_\sigma^{n_Q} \times Q_{\sigma'}^{n_Q}) \leq n_Q H^2(Q_\sigma \times Q_{\sigma'}) = C_\gamma^2 C_\alpha^{2\gamma} n_Q w r^{2\alpha\gamma}.$$

Again, bringing in (20) leads to

$$\begin{aligned} H^2(Q_\sigma^{n_Q} \times Q_{\sigma'}^{n_Q}) &\leq c_w C_\gamma^2 C_\alpha^{2\gamma} n_Q r^{2\alpha\gamma+d} = C_\gamma^2 C_\alpha^{2\gamma} c_w n_Q \left(n_P (e^\varepsilon - 1)^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{2\gamma\alpha+d}{2\alpha+2d}} \\ &\leq C_\gamma^2 C_\alpha^{2\gamma} c_w. \end{aligned} \quad (24)$$

Bringing (23), (24) into (22), we have

$$H^2(\text{RP}_{\sigma'}^{n_P} \times Q_\sigma^{n_Q}, \text{RP}_{\sigma}^{n_P} \times Q_{\sigma'}^{n_Q}) \leq \frac{1}{4}$$

as long as the constants are small enough such that $\max(4C_\alpha^2 c_w^2 c_r^{2\alpha+2d}, C_\gamma^2 C_\alpha^{2\gamma} c_w) \leq 1/8$. Let

$$\hat{\sigma} = \arg \min_{\sigma \in \{-1, 1\}^m} \mathcal{R}_{P_\sigma}(\hat{f}) - \mathcal{R}_{P_\sigma}^*.$$

Apply Tsybakov (2009, Theorem 2.12), there holds

$$\inf_{\hat{\sigma}} \max_{\sigma} E_{X,Y \sim P_\sigma} \text{Ham}(\hat{\sigma}, \sigma) \geq \frac{m}{2} \left(1 - \frac{\sqrt{2}}{2} \right). \quad (25)$$

Next, we reduce the mini-max risk into this quantity. Since

$$\mathcal{R}_{P_\sigma}(f_{\hat{\sigma}}^*) - \mathcal{R}^* \leq \mathcal{R}_{P_\sigma}(\hat{f}) - \mathcal{R}^* + \mathcal{R}_{P_{\hat{\sigma}}}(\hat{f}) - \mathcal{R}^* \leq 2 \left(\mathcal{R}_{P_\sigma}(\hat{f}) - \mathcal{R}^* \right),$$

for any estimator \hat{f} , we can reduce the excess risk as

$$\begin{aligned} \mathbb{E}_{X,Y} \left[\mathcal{R}_{P_\sigma}(\hat{f}) - \mathcal{R}_{P_\sigma}^* \right] &\geq \frac{1}{2} \mathbb{E}_{X,Y} \left[\mathcal{R}_{P_\sigma}(f_{\hat{\sigma}}^*) - \mathcal{R}_{P_\sigma}^* \right] \\ &= \frac{1}{2} \mathbb{E}_{X,Y} \left[\int_{\mathcal{X}} \left| \eta_\sigma(x) - \frac{1}{2} \right| \mathbf{1} \{ f_\sigma^*(x) \neq f_{\hat{\sigma}}^*(x) \} dP_{\sigma,X}(x) \right] \\ &= \frac{1}{4} \mathbb{E}_{X,Y} \left[\sum_{\sigma_k \neq \hat{\sigma}_k} \int_{B(x_k, r)} c_\alpha r^\alpha dP_{\sigma,X}(x) \right] = \frac{c_\alpha w r^\alpha}{4} \mathbb{E}_{X,Y} [\text{Ham}(\sigma, \hat{\sigma})]. \end{aligned}$$

This together with (25) yield

$$\begin{aligned} \inf_{\hat{f}} \max_{\mathcal{H}} \mathbb{E}_{X,Y} \left[\mathcal{R}_{P_\sigma}(\hat{f}) - \mathcal{R}_{P_\sigma}^* \right] &\geq \frac{1}{8} \left(1 - \frac{\sqrt{2}}{2}\right) c_\alpha w r^\alpha m \\ &\gtrsim r^{(\beta+1)\alpha} = \left(n_P (e^\varepsilon - 1)^2 + n_Q^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{(\beta+1)\alpha}{2\alpha+2d}}. \end{aligned}$$

■

B.3 Proof of Theorem 4

To prove Theorem 4, we rely on the following decomposition.

$$\tilde{\eta}_{P,Q}^\pi(x) - \frac{1}{2} = \underbrace{\tilde{\eta}_{P,Q}^\pi(x) - \hat{\eta}_{P,Q}^\pi(x)}_{\text{Privatized Error}} + \underbrace{\hat{\eta}_{P,Q}^\pi(x) - \bar{\eta}_{P,Q}^\pi(x)}_{\text{Sample Error}} + \underbrace{\bar{\eta}_{P,Q}^\pi(x) - \frac{1}{2}}_{\text{Signal Strength}} \quad (26)$$

where we define

$$\hat{\eta}_{P,Q}^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \quad (27)$$

$$\bar{\eta}_{P,Q}^\pi(x) = \sum_{j \in \mathcal{I}_p} \mathbf{1}\{x \in A_{(p)}^j\} \frac{n_P \int_{A_{(p)}^j} \eta_P^*(x') dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} \eta_Q^*(x') dQ_X(x')}{n_P \int_{A_{(p)}^j} dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} dQ_X(x')}. \quad (28)$$

Correspondingly, we let $\hat{f}_{P,Q}^\pi(x) = \mathbf{1}\left(\hat{\eta}_{P,Q}^\pi(x) > \frac{1}{2}\right)$ and $\bar{f}_{P,Q}^\pi(x) = \mathbf{1}\left(\bar{\eta}_{P,Q}^\pi(x) > \frac{1}{2}\right)$.

Loosely speaking, the first error term quantifies the degradation brought by adding noises to the estimator, which we call the privatized error. The second term corresponds to the expected estimation error brought by the randomness of the data, which we call the sample error. The last term is called signal strength, which quantifies the uncertainty of the population regression function. In general, we want the absolute value of the third term to be large, (i.e. the population regression function is sure about the label) while the first and the second terms to be small (i.e. the noise is small). In the proof of Theorem 4, we specify a region Ω_n containing points with strong signal strength. We show that the first two terms are dominated by the signal strength on Ω_n . The following three lemmas provide bounds for each of the three terms.

Lemma 8 (Bounding of Privatized Error) *Let π be constructed through the max-edge partition in Section 3.2 with depth p . Let $\tilde{\eta}_{P,Q}^\pi$ be the locally private two-sample tree classifier in (8). Let $\hat{\eta}_{P,Q}^\pi$ be defined in (27). Then for $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{d}{4\alpha+2d}}$, under the condition of Theorem 4, there holds*

$$\|\tilde{\eta}_{P,Q}^\pi - \hat{\eta}_{P,Q}^\pi\|_\infty \lesssim \frac{2^p \cdot \sqrt{n_P \log(n_P + n_Q)}}{\varepsilon(n_P + \lambda n_Q)}$$

with probability $P^{n_P} \otimes Q^{n_Q} \otimes R$ at least $1 - 3/(n_P + n_Q)^2$.

Lemma 9 (Bounding of Sample Error) *Let π be constructed through the max-edge partition in Section 3.2 with depth p . Let $\widehat{\eta}_{P,Q}^\pi$ and $\overline{\eta}_{P,Q}^\pi$ be defined in (27) and (28). Then for $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{\alpha}{4\alpha+2d}}$, under the condition of Theorem 4, there holds*

$$\|\widehat{\eta}_{P,Q}^\pi - \overline{\eta}_{P,Q}^\pi\|_\infty \lesssim \sqrt{\frac{2^p(n_P + \lambda^2 n_Q) \cdot \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2}}$$

with probability $P^{n_P} \otimes Q^{n_Q} \otimes R$ at least $1 - 3/(n_P + n_Q)^2$.

Lemma 10 (Bounding of Signal Strength) *Let π be constructed through the max-edge partition in Section 3.2 with depth p . Define the region*

$$\Delta_n := \left\{ x \in \mathcal{X} \mid |\eta_P^*(x) - 1/2| \geq c_\Delta \cdot 2^{-p\alpha/d} \right\} \quad (29)$$

for some constant c_Δ large enough. Let $\overline{\eta}_{P,Q}^\pi$ be defined in (28). Then under the condition of Theorem 4, there holds

$$\left(\overline{\eta}_{P,Q}^\pi(x) - \frac{1}{2} \right) \cdot \left(\eta_P^*(x) - \frac{1}{2} \right) > 0, \quad \text{and} \quad \left| \overline{\eta}_{P,Q}^\pi(x) - \frac{1}{2} \right| \gtrsim \frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q}$$

for all $x \in \Delta_n$.

We also need the following technical lemma.

Lemma 11 *Suppose $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{\alpha}{4\alpha+2d}}$. If we let*

$$\lambda \asymp \delta^{\frac{(\gamma-1)\alpha}{d} + (\frac{(\gamma+1)\alpha}{d} - 1) \wedge 0} \asymp \left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} + \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{(\gamma\alpha-\alpha) \wedge (2\gamma\alpha-d)}{2\alpha+2d}},$$

then the following statements hold:

$$\frac{\varepsilon^2(n_P + \lambda^2 n_Q)}{\log(n_P + n_Q)} \gtrsim \delta^{-\frac{2\alpha+2d}{d}}, \quad \frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2} \lesssim \delta, \quad \frac{\sqrt{n_P \varepsilon^{-2} \log(n_P + n_Q)}}{(n_P + \lambda n_Q)} \lesssim \delta.$$

All proofs of the above lemmas are provided in Section B.4.

Proof of Theorem 4 Remind the definition in (29). Theorem 4 takes $2^p \asymp \delta^{-1}$ and $\lambda \asymp \delta^{(\gamma-1)\alpha/d}$. We prove the theorem by separately considering the region

$$\Omega_n := \left\{ x \in \mathcal{X} \mid |\eta_P^*(x) - 1/2| \geq c_\Omega \cdot \delta^{\alpha/d} \right\}$$

for constant c_Ω such that $\Omega_n \subset \Delta_n$. In this case, suppose we have $2^{-p} = C_p \delta$. Then it suffices to take $c_\Omega = 2C_p^{\alpha/d} \cdot c_\Delta$.

For $x \in \Omega_n$, we show that $\tilde{f}_{P,Q}^\pi(x) = f_P^*(x)$. Without loss of generality, let $f_P^*(x) = 1$. By Lemma 8, 9, and 10 we have

$$\begin{aligned} \tilde{f}_{P,Q}^\pi(x) - \frac{1}{2} &\geq \bar{\eta}_{P,Q}^\pi(x) - \frac{1}{2} - \|\bar{\eta}_{P,Q}^\pi - \hat{\eta}_{P,Q}^\pi\|_\infty - \|\hat{\eta}_{P,Q}^\pi - \tilde{\eta}_{P,Q}^\pi\|_\infty \\ &\gtrsim \frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q} - \frac{2^p \cdot \sqrt{n_P}}{\varepsilon(n_P + \lambda n_Q)} - \sqrt{\frac{2^p(n_P + \lambda^2 n_Q) \cdot \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2}} \end{aligned} \quad (30)$$

with probability $1 - 3/(n_P + n_Q)^2$ with respect to $P^{n_P} \otimes Q^{n_Q} \otimes R$. We first compare the signal strength with the privatized error. From the conditions in Theorem 4, we have the fact that $\lambda \geq 2^{-(\gamma-1)p\alpha/d}$.

$$\frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q} \cdot \left(\frac{2^p \cdot \sqrt{n_P}}{\varepsilon(n_P + \lambda n_Q)} \right)^{-1} \geq 2^{-p\alpha/d} \cdot 2^{-p} \cdot \varepsilon \cdot \frac{n_P + \lambda^2 n_Q}{\sqrt{n_P \log(n_P + n_Q)}}.$$

By bringing in the value of δ and applying Lemma 11 (the first conclusion), there holds

$$\begin{aligned} 2^{-p\alpha/d} \cdot 2^{-p} \cdot \varepsilon \cdot \frac{n_P + \lambda^2 n_Q}{\sqrt{n_P \log(n_P + n_Q)}} &\gtrsim C_p^{\frac{\alpha+d}{d}} \cdot \delta^{-\frac{\alpha+d}{d}} \cdot n_P^{-\frac{1}{2}} \cdot \varepsilon^{-1} \cdot \sqrt{\log(n_P + n_Q)} \\ &\gtrsim C_p^{\frac{\alpha+d}{d}} \cdot \left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} \right)^{\frac{1}{2}} \cdot n_P^{-\frac{1}{2}} \cdot \varepsilon^{-1} \cdot \sqrt{\log(n_P + n_Q)} \\ &= C_p^{\frac{\alpha+d}{d}}. \end{aligned}$$

Thus, for sufficiently large C_p , we have

$$\frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q} \geq 2 \frac{2^p \cdot \sqrt{n_P}}{\varepsilon(n_P + \lambda n_Q)}. \quad (31)$$

Analogously, we compare the signal strength with the sample error, which is

$$\begin{aligned} &\frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q} \cdot \left(\frac{2^p(n_P + \lambda^2 n_Q) \cdot \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2} \right)^{-\frac{1}{2}} \\ &= 2^{-p\alpha/d} \cdot \delta^{\frac{1}{2}} \cdot \sqrt{n_P + \lambda^2 n_Q} \cdot \log^{-\frac{1}{2}}(n_P + n_Q). \end{aligned}$$

By bringing in the value of δ and applying Lemma 11 (the first conclusion), there holds

$$2^{-p\alpha/d} \cdot \delta^{\frac{1}{2}} \cdot \sqrt{n_P + \lambda^2 n_Q} \cdot \log^{-\frac{1}{2}}(n_P + n_Q) \gtrsim \delta^{-\frac{1}{2}} \cdot \varepsilon^{-1} \gtrsim 1.$$

Thus, for sufficiently large n_P and n_Q , we have

$$\frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q} \geq 2 \sqrt{\frac{2^p(n_P + \lambda^2 n_Q) \cdot \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2}}. \quad (32)$$

Combined with (31) and (32), (30) yields $\tilde{\eta}_{P,Q}^\pi(x) - \frac{1}{2} \geq 0$, which indicates that $\tilde{f}_{P,Q}^\pi(x) = f_P^*(x)$ for all $x \in \Omega_n$.

For Ω_n^c , by Assumption 2, we have

$$\int_{\Omega_n^c} \left| \eta_P^*(x) - \frac{1}{2} \right| \mathbf{1} \left\{ \tilde{f}_{P,Q}^\pi(x) \neq f_P^*(x) \right\} dP_X(x) \leq c_\Omega \delta^{\alpha/d} \int_{\Omega_n^c} dP_X(x) \leq c_\Omega^{1+\beta} C_\beta \delta^{(\beta+1)\alpha/d}.$$

Together, there holds

$$\begin{aligned} \mathcal{R}_P(\tilde{f}_{P,Q}^\pi) - \mathcal{R}_P^* &= \int_{\mathcal{X}} \left| \eta_P^*(x) - \frac{1}{2} \right| \mathbf{1} \left\{ \tilde{f}_{P,Q}^\pi(x) \neq f_P^*(x) \right\} dP_X(x) \\ &= \int_{\Omega_n^c} \left| \eta_P^*(x) - \frac{1}{2} \right| dP_X(x) \lesssim \delta^{(\beta+1)\alpha/d}. \end{aligned}$$

■

B.4 Proofs of Lemmas in Section B.3

To prove lemmas in Section B.3, we first present several technical results. Their proof will be found in Section B.4

Lemma 12 *Suppose $\xi_i, i = 1, \dots, n$ are independent sub-exponential random variables with parameters (a, b) (Wainwright, 2019, Definition 2.9). Then there holds*

$$\mathbb{P} \left[\left| \frac{1}{n} \sum_{i=1}^n \xi_i - \mathbb{E} \frac{1}{n} \sum_{i=1}^n \xi_i \right| \geq t \right] \leq 2e^{-\frac{nt^2}{2a^2}}.$$

for any $t > 0$. Moreover, a standard Laplace random variable is sub-exponential with parameters $(\sqrt{2}, 1)$. A union bound argument yields that

$$\mathbb{P} \left[\left| \frac{1}{n} \sum_{i=1}^n \xi_i^j - \mathbb{E} \frac{1}{n} \sum_{i=1}^n \xi_i^j \right| \geq t, \text{ for } j = 1, \dots, 2^p \right] \leq 2e^{p \log 2 - \frac{nt^2}{2a^2}}.$$

Lemma 13 *Let π be a partition generated from Algorithm 1 with depth p . Suppose Assumption 1 holds. Suppose $D^P = \{(X_1^P, Y_1^P), \dots, (X_{n_P}^P, Y_{n_P}^P)\}$ and $D^Q = \{(X_1^Q, Y_1^Q), \dots, (X_{n_Q}^Q, Y_{n_Q}^Q)\}$ are drawn i.i.d. from P and Q , respectively. Let $0 < \lambda \lesssim (n_P + n_Q) / \log(n_P + n_Q)$. Then for all $A_{(p)}^j$ in π_p , there holds*

$$\begin{aligned} & \left| \frac{1}{n_P + \lambda n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{X_i^P \in A_{(p)}^j\} + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{X_i^Q \in A_{(p)}^j\} \right) - \int_{A_{(p)}^j} dP_X(x') \right| \\ & \lesssim \sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + \lambda n_Q)^2}} \end{aligned}$$

with probability $P^{n_P} \otimes Q^{n_Q}$ at least $1 - 1/(n_P + n_Q)^2$.

Lemma 14 *Let π be a partition generated from Algorithm 1 with depth p . Suppose Assumption 1 holds. Suppose $D^P = \{(X_1^P, Y_1^P), \dots, (X_{n_P}^P, Y_{n_P}^P)\}$ and $D^Q = \{(X_1^Q, Y_1^Q), \dots, (X_{n_Q}^Q, Y_{n_Q}^Q)\}$ are drawn i.i.d. from P and Q , respectively. Let $0 < \lambda \lesssim (n_P + n_Q)/\log(n_P + n_Q)$. Then for all $A_{(p)}^j$ in π_p , there holds*

$$\frac{1}{n_P + \lambda n_Q} \left| \left(\sum_{i=1}^{n_P} Y_i^P \mathbf{1}\{X_i^P \in A_{(p)}^j\} + \lambda \sum_{i=1}^{n_Q} Y_i^Q \mathbf{1}\{X_i^Q \in A_{(p)}^j\} \right) - \left(n_P \int_{A_{(p)}^j} \eta_P^*(x') dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} \eta_Q^*(x') dQ_X(x') \right) \right| \lesssim \sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + \lambda n_Q)^2}}$$

with probability $P^{n_P} \otimes Q^{n_Q}$ at least $1 - 1/(n_P + n_Q)^2$.

Lemma 15 *Let $\pi_i := \{A_{(i)}^j, j \in \{1, \dots, 2^p\}\}$ be the partition generated by Algorithm 1 with depth i . Then for any $A_{(i)}^j$, there holds*

$$2^{-1} \sqrt{d} \cdot 2^{-i/d} \leq \text{diam}(A_{(i)}^j) \leq 2 \sqrt{d} \cdot 2^{-i/d}.$$

Proof of Lemma 8 We intend to bound

$$\|\tilde{\eta}_{P,Q}^\pi - \hat{\eta}_{P,Q}^\pi\|_\infty = \max_{j \in \mathcal{I}} \left| \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} - \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \right|$$

For each $j \in \mathcal{I}$ and any $x \in A_{(p)}^j$, the error can be decomposed as

$$\begin{aligned} & \left| \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} - \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \right| \\ & + \left| \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} - \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \right| \\ & \leq \underbrace{\left| \frac{\sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} - \frac{\frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j}{\left(\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right) \left(\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right)} \right|}_{(I)} \\ & + \underbrace{\left| \frac{\frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \right|}_{(II)}. \end{aligned}$$

For the first term in (I), we can apply Lemma 12 and 14 to get

$$\begin{aligned}
 & \left| \sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj} \right| \\
 & \leq \left| \sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj} \right| + \left| \frac{4}{\varepsilon} \sum_{i=1}^{n_P} \xi_i^j \right| \\
 & \lesssim (n_P + \lambda n_Q) \left(\sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + \lambda n_Q)^2}} + 2^{-p} + \frac{\sqrt{n_P \varepsilon^{-2} \log(n_P + n_Q)}}{n_P + \lambda n_Q} \right)
 \end{aligned}$$

with probability at least $1 - 2/(n_P + n_Q)^2$, where we used the fact that

$$\int_{A_{(p)}^j} dP_X(x) \asymp 2^{-p}, \quad \int_{A_{(p)}^j} \eta_P^*(x) dP_X(x) \asymp 2^{-p}, \quad \int_{A_{(p)}^j} \eta_Q^*(x) dQ_X(x) \asymp 2^{-p}$$

Using Lemma 11 (the second and the third conclusions), we know that the second term is the dominating term, indicating that

$$\left| \sum_{i=1}^{n_P} \tilde{V}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj} \right| \lesssim (n_P + \lambda n_Q) \cdot 2^{-p}.$$

Analogously, by Lemma 13, there holds

$$\begin{aligned}
 \left| \sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right| & \gtrsim (n_P + \lambda n_Q) \left(2^{-p} - \sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + \lambda n_Q)^2}} \right) \\
 & \asymp (n_P + \lambda n_Q) \cdot 2^{-p}
 \end{aligned} \tag{33}$$

as well as

$$\begin{aligned}
 \left| \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right| & \geq \left| \sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right| - \left| \frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j \right| \\
 & \gtrsim (n_P + \lambda n_Q) \left(2^{-p} - \frac{\sqrt{n_P \varepsilon^{-2} \log(n_P + n_Q)}}{n_P + \lambda n_Q} \right) \\
 & \asymp (n_P + \lambda n_Q) \cdot 2^{-p}.
 \end{aligned}$$

The last step is by Lemma 11 (the second and the third conclusions), knowing that 2^{-p} is the dominating term. Combining these three pieces, there holds

$$(I) \lesssim \sqrt{n_P \varepsilon^{-2} \log(n_P + n_Q)} \cdot ((n_P + \lambda n_Q) \cdot 2^{-p})^{-1} \asymp \frac{2^p \cdot \sqrt{n_P \log(n_P + n_Q)}}{\varepsilon(n_P + \lambda n_Q)}. \tag{34}$$

For (II), (33) and Lemma 12 together yields

$$(II) \lesssim \frac{2^p \cdot \sqrt{n_P \log(n_P + n_Q)}}{\varepsilon(n_P + \lambda n_Q)}. \tag{35}$$

(34) and (35) together lead to the desired conclusion. \blacksquare

Proof of Lemma 9 Let $\eta_{P,Q}^*(x) = (n_P \eta_P^*(x) + \lambda n_Q \eta_Q^*(x)) / (n_P + \lambda n_Q)$. We intend to bound

$$\|\widehat{\eta}_{P,Q}^\pi - \bar{\eta}_{P,Q}^\pi\|_\infty = \max_{j \in \mathcal{I}} \left| \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} - \frac{(n_P + \lambda n_Q) \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x')}{(n_P + \lambda n_Q) \int_{A_{(p)}^j} dP_X(x')} \right|$$

For each $j \in \mathcal{I}$ and any $x \in A_{(p)}^j$, the error can be decomposed as

$$\begin{aligned} & \left| \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj}}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} - \frac{(n_P + \lambda n_Q) \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x')}{(n_P + \lambda n_Q) \int_{A_{(p)}^j} dP_X(x')} \right| \\ & \leq \underbrace{\left| \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x') \right| \cdot \left| \frac{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} - (n_P + \lambda n_Q) \int_{A_{(p)}^j} dP_X(x')}{\left(\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right) \int_{A_{(p)}^j} dP_X(x')} \right|}_{(I)} \\ & + \underbrace{\left| \frac{\sum_{i=1}^{n_P} V_i^{Pj} + \lambda \sum_{i=1}^{n_Q} V_i^{Qj} - (n_P + \lambda n_Q) \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x')}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \right|}_{(II)} \end{aligned}$$

We bound two terms separately. For the first term in (I), Assumption 1 yield

$$\int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x') \lesssim 2^{-p}$$

For the second term, (33) and Lemma 13 yields

$$\begin{aligned} & \left| \frac{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} - (n_P + \lambda n_Q) \int_{A_{(p)}^j} dP_X(x')}{\left(\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj} \right) \int_{A_{(p)}^j} dP_X(x')} \right| \\ & \lesssim (n_P + \lambda n_Q) \cdot \sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + \lambda n_Q)^2}} \cdot ((n_P + \lambda n_Q) \cdot 2^{-2p})^{-1} \\ & \asymp 2^p \cdot \sqrt{\frac{2^p (n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2}} \end{aligned}$$

with probability $1 - 1/(n_P + n_Q)^2$. Together, these two pieces yield

$$(I) \lesssim \sqrt{\frac{2^p (n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2}}. \quad (36)$$

Similarly for (II), we have

$$(II) \lesssim \sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + \lambda n_Q)^2}} \cdot 2^p \asymp \sqrt{\frac{2^p (n_P + \lambda^2 n_Q) \cdot \log(n_P + n_Q)}{(n_P + \lambda n_Q)^2}}. \quad (37)$$

(36) and (37) together yield the desired conclusion. \blacksquare

Proof of Lemma 10 Without loss of generality, consider $f_P^*(x) = 1$ for any $x \in \Delta_n$. By Assumption 3, we have

$$\bar{\eta}_{P,Q}^\pi(x) \geq \frac{n_P \int_{A_{(p)}^j} \eta_P^*(x') dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} \left(\frac{1}{2} + C_\gamma (\eta_P^*(x') - \frac{1}{2})^\gamma\right) dQ_X(x')}{n_P \int_{A_{(p)}^j} dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} dQ_X(x')} \quad (38)$$

If Assumption 1 holds, for any $x' \in A_{(p)}^j$, we have

$$\begin{aligned} \eta_P^*(x') - \frac{1}{2} &\geq \eta_P^*(x) - \frac{1}{2} - |\eta_P^*(x) - \eta_P^*(x')| \\ &\geq \eta_P^*(x) - \frac{1}{2} - c_L \|x - x'\|^\alpha \geq \eta_P^*(x) - \frac{1}{2} - 2c_L \sqrt{d} 2^{-p\alpha/d} \geq (c_\Delta - 2c_L \sqrt{d}) 2^{-p\alpha/d} \end{aligned}$$

where the second last and the last inequalities come from Lemma 15 and the definition of Δ_n , respectively. For sufficiently large c_Δ , we have $\eta_P^*(x') \geq 1/2 + 2^{-p\alpha/d}$ for all $x' \in A_{(p)}^j$, and thus (38) yields

$$\begin{aligned} \bar{\eta}_{P,Q}^\pi(x) &\geq \frac{1}{2} + \frac{n_P \int_{A_{(p)}^j} 2^{-p\alpha/d} dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} C_\gamma 2^{-p\gamma\alpha/d} dQ_X(x')}{n_P \int_{A_{(p)}^j} dP_X(x') + \lambda n_Q \int_{A_{(p)}^j} dQ_X(x')} \\ &\geq \frac{1}{2} + \frac{\underline{c}^2 \cdot (1 \vee C_\gamma)}{\bar{c}^2} \frac{n_P 2^{-p\alpha/d} + \lambda n_Q 2^{-p\gamma\alpha/d}}{n_P + \lambda n_Q} \end{aligned}$$

which yields the desired result. Note that the factor \underline{c}/\bar{c} comes from the lower bound of ratio $\int_{A_{(p)}^j} dQ_X(x') / \int_{A_{(p)}^j} dP_X(x')$. \blacksquare

Proof of Lemma 11 We prove the statement by discussing four cases depending on the value of n_P , n_Q , and γ , α , d .

Case (i): $\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} \geq \left(\frac{n_Q}{\log(n_P + n_Q)}\right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$ and $\frac{(\gamma+1)\alpha}{d} \geq 1$. We prove the three conclusions by turn. The first conclusion follows from

$$\begin{aligned} \frac{n_P + \lambda^2 n_Q}{\delta^{-(2\alpha+2d)/d}} &\geq n_P \cdot \delta^{\frac{2\alpha+2d}{d}} \\ &\gtrsim n_P \cdot \left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)}\right)^{-\frac{2\alpha+2d}{2\alpha+2d}} \gtrsim \frac{\log(n_P + n_Q)}{\varepsilon^2}. \end{aligned}$$

For the second conclusion, let $n_Q = n_P^a$ for some constant a . Note that we have $\frac{2\gamma\alpha+d}{2\alpha+2d} \geq a > 0$. By replacing all terms with equivalent order of n_P , it suffices to prove

$$\begin{aligned} & n_P + n_P^{a-\frac{2(\gamma-1)\alpha}{2\alpha+2d}} \cdot \varepsilon^{-\frac{2(\gamma-1)\alpha}{\alpha+d}} \\ & \lesssim n_P^{\frac{4\alpha+3d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{d}{\alpha+d}} + n_P^{1+a-\frac{(\gamma-1)\alpha+d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{(\gamma-1)\alpha+d}{\alpha+d}} + n_P^{2a-\frac{2(\gamma-1)\alpha+d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{2(\gamma-1)\alpha+d}{\alpha+d}} \end{aligned}$$

The left-hand side is dominated respectively by $n_P \lesssim n_P^{\frac{4\alpha+3d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{d}{\alpha+d}}$ and

$$n_P^{a-\frac{2(\gamma-1)\alpha}{2\alpha+2d}} \cdot \varepsilon^{-\frac{2(\gamma-1)\alpha}{\alpha+d}} \lesssim n_P^{a-\frac{2(\gamma-1)\alpha}{2\alpha+2d}+\frac{1}{2}+\frac{\gamma\alpha}{2\alpha+2d}} \cdot \varepsilon^{-\frac{(\gamma-1)\alpha+d}{\alpha+d}} = n_P^{1+a-\frac{(\gamma-1)\alpha+d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{(\gamma-1)\alpha+d}{\alpha+d}}.$$

Here, the inequality is drawn by

$$n_P^{\frac{1}{2}} \cdot \varepsilon^{-\frac{d-\alpha}{\alpha+d}} \geq n_P^{\frac{1}{2}(1-\frac{d-\alpha}{\alpha+d})}, \quad n_P^{\frac{\gamma\alpha}{2\alpha+2d}} \cdot \varepsilon^{\frac{\gamma\alpha}{\alpha+d}} \gtrsim 1.$$

Note that we omit the log factor since the polynomial part of the left-hand side is strictly dominated by the right-hand side, yielding that the conclusion holds with sufficiently large n_P and n_Q . For the last statement, the same procedure yields that it suffices to show

$$n_P^{\frac{1}{2}} \lesssim \left(n_P^{\frac{2\alpha+d}{2\alpha+2d}} + n_P^{a-\frac{(\gamma-1)\alpha+d}{2\alpha+2d}} \right) \cdot \varepsilon^{\frac{\alpha}{\alpha+d}},$$

which holds since $n_P^{\frac{\alpha}{2\alpha+2d}} \cdot \varepsilon^{\frac{\alpha}{\alpha+d}} \gtrsim n_P^{\frac{d}{\alpha+d} \cdot \frac{\alpha}{2\alpha+2d}}$.

Case (ii): $\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} \geq \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$ and $\frac{(\gamma+1)\alpha}{d} < 1$. Compared to case (i), the value is λ is multiplied by a factor $\delta^{(\frac{(\gamma+1)\alpha}{d}-1) \wedge 0} > 1$. Thus, the first conclusion holds since the value of the left-hand side increases. The third conclusion holds since the value of the left-hand side decreases. For the second conclusion, by replacing all terms with equivalent order of n_P , it suffices to prove

$$\begin{aligned} & n_P + n_P^{a-\frac{4\gamma\alpha-2d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{4\gamma\alpha-2d}{\alpha+d}} \\ & \lesssim n_P^{\frac{4\alpha+3d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{d}{\alpha+d}} + n_P^{1+a-\frac{2\gamma\alpha}{2\alpha+2d}} \cdot \varepsilon^{-\frac{2\gamma\alpha}{\alpha+d}} + n_P^{2a-\frac{4\gamma\alpha-d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{4\gamma\alpha-d}{\alpha+d}}. \end{aligned}$$

For the left-hand side, the first term is dominated as $n_P \lesssim n_P^{\frac{4\alpha+3d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{d}{\alpha+d}}$. Since $\frac{-2\gamma\alpha+2d}{2\alpha+2d} < 1$,

we have $n_P \cdot (n_P \cdot \varepsilon^2)^{\frac{2\gamma\alpha-2d}{2\alpha+2d}} \gtrsim n_P$. Thus, the second term of left-hand side is controlled by $n_P^{a-\frac{4\gamma\alpha-2d}{2\alpha+2d}} \cdot \varepsilon^{-\frac{4\gamma\alpha-2d}{\alpha+d}} \lesssim n_P^{1+a-\frac{2\gamma\alpha}{2\alpha+2d}} \cdot \varepsilon^{-\frac{2\gamma\alpha}{\alpha+d}}$.

Case (iii): $\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} < \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$ and $\frac{(\gamma+1)\alpha}{d} \geq 1$. We now have $a > \frac{2\gamma\alpha+d}{2\alpha+2d}$. There holds

$$\begin{aligned} & \frac{n_P + \lambda^2 n_Q}{\delta^{-(2\alpha+2d)/d}} \geq n_Q \cdot \delta^{\frac{4\gamma\alpha-2d+2\alpha+2d}{d}} \\ & \asymp n_Q \cdot \left(\left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}} \right)^{-\frac{2\gamma\alpha}{2\alpha+2d}} \\ & = \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{d}{2\gamma\alpha+d}} \cdot \varepsilon^2 \cdot \frac{\log(n_P + n_Q)}{\varepsilon^2} \gtrsim \frac{\log(n_P + n_Q)}{\varepsilon^2} \end{aligned}$$

for sufficiently large $n_P + n_Q$, which implies the first conclusion. For the second conclusion, we want

$$n_P + n_P^{a \frac{2\alpha+d}{2\gamma\alpha+d}} \lesssim n_P^{a \frac{2\gamma\alpha}{2\gamma\alpha+d}} + n_P^{1+a \frac{(\gamma+1)\alpha+d}{2\gamma\alpha+d}} + n_P^{a \frac{2(\gamma+1)\alpha+d}{2\gamma\alpha+d}}.$$

The left-hand side is dominated by the term $n_P^{a \frac{2\alpha+d}{2\gamma\alpha+d}}$. Thus, the conclusion holds since there holds $a \frac{2\alpha+d}{2\gamma\alpha+d} < a \frac{2(\gamma+1)\alpha+d}{2\gamma\alpha+d}$ for all a, γ, α , and d . For the last statement, the same procedure yields that it suffices to show

$$n_P^{\frac{1}{2}} \lesssim \left(n_P^{1-a \frac{d}{2\gamma\alpha+d}} + n_P^{a-a \frac{2\gamma\alpha}{2\gamma\alpha+d}} \right) \cdot \varepsilon.$$

Thus, the conclusion holds since the following conditions hold:

$$\frac{1}{2} \leq 1 - \frac{d}{2\alpha+2d} < 1 - a \frac{d}{2\gamma\alpha+d}, \quad \text{and} \quad n_P^{\frac{\alpha}{2\alpha+2d}} \cdot \varepsilon \gtrsim 1.$$

Case (iv): $\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} < \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$ and $\frac{(\gamma+1)\alpha}{d} < 1$. Compared to case (iii), the value is λ is multiplied by a factor $\delta^{(\frac{(\gamma+1)\alpha}{d}-1)\wedge 0} > 1$. Thus, the first conclusion holds since the value of the left-hand side increases. The third conclusion holds since the value of the left-hand side decreases. For the second conclusion, by replacing all terms with equivalent order of n_P , it suffices to prove

$$n_P + n_P^{a \frac{-2\gamma\alpha+3d}{2\gamma\alpha+d}} \lesssim n_P^{a \frac{2\gamma\alpha}{2\gamma\alpha+d}} + n_P^{1+a-a \frac{2\gamma\alpha}{2\gamma\alpha+d}} + n_P^{2a-a \frac{4\gamma\alpha-d}{2\gamma\alpha+d}}.$$

We have $n_P \lesssim n_P^{1+\frac{ad}{2\gamma\alpha+d}} = n_P^{1+a-a \frac{2\gamma\alpha}{2\gamma\alpha+d}}$ and $n_P^{a \frac{-2\gamma\alpha+3d}{2\gamma\alpha+d}} \lesssim n_P^{\frac{3ad}{2\gamma\alpha+d}} = n_P^{2a-a \frac{4\gamma\alpha-d}{2\gamma\alpha+d}}$. ■

B.5 Proofs of Results in Section B.4

Proof of Lemma 12 The conclusion follows from (2.18) in Wainwright (2019). ■

In the subsequent proof, we define the empirical measure $D := \frac{1}{n} \sum_{i=1}^n \delta_{(X_i, Y_i)}$ given samples $D = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$, where δ is the Dirac function. To conduct our analysis, we need to introduce the following fundamental descriptions of *covering number* which enables an approximation of an infinite set by finite subsets.

Definition 16 (Covering Number) Let (\mathcal{X}, d) be a metric space and $A \subset \mathcal{X}$. For $\varepsilon > 0$, the ε -covering number of A is denoted as

$$\mathcal{N}(A, d, \varepsilon) := \min \left\{ n \geq 1 : \exists x_1, \dots, x_n \in \mathcal{X} \text{ such that } A \subset \bigcup_{i=1}^n B(x_i, \varepsilon) \right\},$$

where $B(x, \varepsilon) := \{x' \in \mathcal{X} : d(x, x') \leq \varepsilon\}$.

To prove Lemma 17, we need the following fundamental lemma concerning the VC dimension of random partitions in Section 3.2. To this end, let $p \in \mathbb{N}$ be fixed and π_p be a partition of \mathcal{X} with the number of splits p and $\pi_{(p)}$ denote the collection of all partitions π_p .

Lemma 17 *Let $\tilde{\mathcal{A}}$ be the collection of all cells $\times_{i=1}^d [a_i, b_i]$ in \mathbb{R}^d . For all $0 < \varepsilon < 1$, there exists a universal constant C such that*

$$\mathcal{N}(\mathbf{1}_{\tilde{\mathcal{A}}}, \|\cdot\|_{L_1(Q)}, \varepsilon) \leq C(2d+1)(4e)^{2d+1}(1/\varepsilon)^{2d}.$$

Proof of Lemma 17 The result follows directly from Theorem 9.2 in Kosorok (2008). ■

Before we proceed, we list the well-known Bernstein's inequality that will be used frequently in the proofs.

Lemma 18 (Bernstein's inequality) *Let $B > 0$ and $\sigma > 0$ be real numbers, and $n \geq 1$ be an integer. Furthermore, let ξ_1, \dots, ξ_n be independent random variables satisfying $\mathbb{E}_P \xi_i = 0$, $\|\xi_i\|_\infty \leq B$, and $\mathbb{E}_P \sum_{i=1}^n \xi_i^2 \leq \sigma^2$ for all $i = 1, \dots, n$. Then for all $\tau > 0$, we have*

$$P\left(\frac{1}{n} \sum_{i=1}^n \xi_i \geq \sqrt{\frac{2\sigma^2\tau}{n^2}} + \frac{2B\tau}{3n}\right) \leq e^{-\tau}.$$

Proof of Lemma 13 Since P and Q share the same marginal distribution, it is equivalent to prove

$$\begin{aligned} & \left| \frac{1}{n_P + n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{X_i \in A_{(p)}^j\} + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{X_i \in A_{(p)}^j\} \right) - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{A_{(p)}^j} dP_X(x') \right| \\ & \lesssim \sqrt{\frac{(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{2^p (n_P + n_Q)^2}} \end{aligned}$$

holds with probability $P^{n_P+n_Q}$ at least $1 - 1/(n_P + n_Q)^2$ for $X_i \sim P_X$. Let $\tilde{\mathcal{A}}$ be the collection of all cells $\times_{i=1}^d [a_i, b_i]$ in \mathbb{R}^d . Applying Lemma 17 with $(D_X + P_X)/2$, there exists an ε -net $\{\tilde{A}^k\}_{k=1}^K \subset \tilde{\mathcal{A}}$ with

$$K \leq C(2d+1)(4e)^{2d+1}(1/\varepsilon)^{2d} \quad (39)$$

such that for any $j \in \mathcal{I}_p$, there exist some $k \in \{1, \dots, K\}$ such that

$$\|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1((D_X+P_X)/2)} \leq \varepsilon,$$

Since

$$\begin{aligned} & \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1((D_X+P_X)/2)} \\ & = 1/2 \cdot \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(D_X)} + 1/2 \cdot \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(P_X)}, \end{aligned}$$

we get

$$\|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(D_X)} \leq 2\varepsilon, \quad \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(P_X)} \leq 2\varepsilon. \quad (40)$$

Consequently, by the definition of the covering number and the triangle inequality, for any $j \in \mathcal{I}_p$, there holds

$$\begin{aligned} & \left| \frac{1}{n_P + n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{x \in A_{(p)}^j\}(X_i) + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{x \in A_{(p)}^j\}(X_i) \right) - \int_{\tilde{A}_{(p)}^j} dP_X(x') \right| \\ & \leq \left| \frac{1}{n_P + n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{x \in \tilde{A}^k\}(X_i) + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{x \in \tilde{A}^k\}(X_i) \right) - \int_{\tilde{A}^k} dP_X(x') \right| \\ & \quad + \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(D_X)} + \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(P_X)} \\ & \leq \left| \frac{1}{n_P + n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{x \in \tilde{A}^k\}(X_i) + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{x \in \tilde{A}^k\}(X_i) \right) - \int_{\tilde{A}^k} dP_X(x') \right| + 4\varepsilon. \end{aligned}$$

Therefore, we get

$$\begin{aligned} & \sup_{j \in \mathcal{I}_p} \left| \frac{1}{n_P + n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{x \in A_{(p)}^j\}(X_i) + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{x \in A_{(p)}^j\}(X_i) \right) - \int_{\tilde{A}_{(p)}^j} dP_X(x') \right| \quad (41) \\ & \leq \sup_{1 \leq k \leq K} \left| \frac{1}{n_P + n_Q} \left(\sum_{i=1}^{n_P} \mathbf{1}\{x \in \tilde{A}^k\}(X_i) + \lambda \sum_{i=1}^{n_Q} \mathbf{1}\{x \in \tilde{A}^k\}(X_i) \right) - \int_{\tilde{A}^k} dP_X(x') \right| + 4\varepsilon. \quad (42) \end{aligned}$$

For any fixed $1 \leq k \leq K$, let the random variable ξ_i be defined by $\xi_i := \mathbf{1}\{X_i \in \tilde{A}^k\} - \int_{\tilde{A}^k} dP_X(x')$ for $i = 1, \dots, n_P$ and $\xi_i := \lambda(\mathbf{1}\{X_i \in \tilde{A}^k\} - \int_{\tilde{A}^k} dP_X(x'))$ for $i = n_P + 1, \dots, n_P + n_Q$. Then we have $\mathbb{E}_{P_X} \xi_i = 0$, $\|\xi\|_\infty \leq 1 \vee \lambda$, and $\mathbb{E}_{P_X} \sum_{i=1}^{n_P+n_Q} \xi_i^2 \leq (n_P + \lambda^2 n_Q) \int_{\tilde{A}^k} dP_X(x')$. According to Assumption 1, there holds $\mathbb{E}_{P_X} \sum_{i=1}^{n_P+n_Q} \xi_i^2 \leq (n_P + \lambda^2 n_Q) \cdot \bar{c} \cdot 2^{-p}$. Applying Bernstein's inequality in Lemma 18, we obtain

$$\begin{aligned} \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in \tilde{A}^k\} - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{\tilde{A}^k} dP_X(x') \right| & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p} \tau (n_P + \lambda^2 n_Q)}{(n_P + n_Q)^2}} \\ & \quad + \frac{2\tau \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)} \end{aligned}$$

with probability $P^{n_P+n_Q}$ at least $1 - 2e^{-\tau}$. Then the union bound together with the covering number estimate (39) implies that

$$\begin{aligned} & \sup_{1 \leq k \leq K} \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in \tilde{A}^k\} - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{\tilde{A}^k} dP_X(x') \right| \\ & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p} (n_P + \lambda^2 n_Q) (\tau + \log(2K))}{(n_P + n_Q)^2}} + \frac{2(\tau + \log(2K)) \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)} \end{aligned}$$

with probability $P^{n_P+n_Q}$ at least $1 - e^{-\tau}$. Let $\tau = 2 \log(n_P + n_Q)$ and $\varepsilon = 1/(n_P + n_Q)$. Then for any $n_P + n_Q > N_1 := (2C) \wedge (2d + 1) \wedge (4e)$, we have $\tau + \log(2K) = 2 \log n + \log(2C) + \log(2d + 1) + (2d + 1) \log(4e) + 2d \log n \leq (4d + 5) \log(n_P + n_Q)$. Therefore, we have

$$\begin{aligned} & \sup_{1 \leq k \leq K} \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in \tilde{A}^k\} - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{\tilde{A}^k} dP_X(x') \right| \\ & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p}(n_P + \lambda^2 n_Q)(4d + 5) \log(n_P + n_Q)}{(n_P + n_Q)^2}} + \frac{2(4d + 5) \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)} \end{aligned}$$

with probability $P^{n_P+n_Q}$ at least $1 - 1/(n_P + n_Q)^2$. This together with (41) yields that

$$\begin{aligned} & \sup_{j \in \mathcal{I}_p} \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{x \in A_{(p)}^j\} - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{\tilde{A}_{(p)}^j} dP_X(x') \right| \\ & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p}(4d + 5)(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{(n_P + n_Q)^2}} + \frac{2(4d + 5) \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)} + \frac{4}{n_P + n_Q}. \end{aligned}$$

We now compare to see which term is dominating. Apply Lemma 11 (the first conclusion), we have

$$\frac{2^{-p}(n_P + \lambda^2 n_Q) \cdot \log(n_P + n_Q)}{(1 \vee \lambda^2) \cdot \log^2(n_P + n_Q)} \gtrsim \frac{\delta^{-\frac{2\alpha+d}{d}}}{\varepsilon^2 \cdot (1 \vee \lambda^2)}.$$

When $\gamma \leq \frac{d}{\alpha} - 1$ and $\gamma \leq \frac{d}{2\alpha}$,

$$\begin{aligned} \frac{\delta^{-\frac{2\alpha+d}{d}}}{\varepsilon^2 \cdot (1 \vee \lambda^2)} &= \varepsilon^{-2} \cdot \delta^{-\frac{2\alpha+d}{d} - \frac{(\gamma-1)\alpha}{d}} > \varepsilon^{-2} \cdot \delta^{-\frac{2\alpha+3d}{2d}} \geq \varepsilon^{-2} \cdot (n_P \cdot \varepsilon^2)^{\frac{2\alpha+3d}{4\alpha+4d}} \\ &= (n_P \cdot \varepsilon^{-\frac{4\alpha+2d}{2\alpha+3d}})^{\frac{2\alpha+3d}{4\alpha+4d}} \gtrsim 1. \end{aligned}$$

When $\gamma \leq \frac{d}{\alpha} - 1$ and $\gamma > \frac{d}{2\alpha}$,

$$\frac{\delta^{-\frac{2\alpha+d}{d}}}{\varepsilon^2 \cdot (1 \vee \lambda^2)} = \varepsilon^{-2} \cdot \delta^{-\frac{2\alpha+d}{d}} \geq \varepsilon^{-2} \cdot (n_P \cdot \varepsilon^2)^{\frac{2\alpha+d}{2\alpha+2d}} = (n_P \cdot \varepsilon^{-\frac{2d}{2\alpha+d}})^{\frac{2\alpha+d}{2\alpha+2d}} \gtrsim 1.$$

When $\gamma > \frac{d}{\alpha} - 1$ and $\gamma \geq 1$, the derivation is the same as the last situation. When $\gamma > \frac{d}{\alpha} - 1$ and $\gamma < 1$,

$$\frac{\delta^{-\frac{2\alpha+d}{d}}}{\varepsilon^2 \cdot (1 \vee \lambda^2)} = \varepsilon^{-2} \cdot \delta^{-\frac{2\alpha+d}{d} - \frac{2\gamma\alpha-d}{d}} \geq \varepsilon^{-2} \cdot (n_P \cdot \varepsilon^2)^{\frac{2\alpha}{d}} = (n_P \cdot \varepsilon^{-\frac{4\alpha-2d}{2\alpha}})^{\frac{2\alpha}{d}} \gtrsim 1.$$

These derivations imply that the first term is dominating the upper bound, which implies the desired result. \blacksquare

Proof of Lemma 14 For notation simplicity, let $(X_i, Y_i) \sim P$ for $i = 1, \dots, n_P$ and $(X_i, Y_i) \sim Q$ for $i = n_P + 1, \dots, n_P + n_Q$. Let $\tilde{Y}_i = Y_i$ for $i = 1, \dots, n_P$ and $\tilde{Y}_i = \lambda \cdot Y_i$ for $i = n_P + 1, \dots, n_P + n_Q$. Let $\eta_{P,Q}^*(x) = (n_P \eta_P^*(x) + \lambda n_Q \eta_Q^*(x)) / (n_P + \lambda n_Q)$. Let $\tilde{\mathcal{A}}$ be the collection of all cells $\times_{i=1}^d [a_i, b_i]$ in \mathbb{R}^d . Then there exists an ε -net $\{\tilde{A}^k\}_{k=1}^K \subset \tilde{\mathcal{A}}$ with K bounded by (39) such that for any $j \in \mathcal{I}_p$, (40) holds for some $k \in \{1, \dots, K\}$. Consequently, by the definition of the covering number and the triangle inequality, for any $j \in \mathcal{I}_p$, there holds

$$\begin{aligned}
 & \left| \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in A_{(p)}^j\} \tilde{Y}_i - \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x') \right| \\
 & \leq \left| \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in \tilde{A}^k\} \tilde{Y}_i - \int_{\tilde{A}^k} \eta_{P,Q}^*(x') dP_X(x') \right| \\
 & \quad + \int_{\mathbb{R}^d} |\mathbf{1}\{x' \in A_{(p)}^j\} - \mathbf{1}\{x' \in \tilde{A}^k\}| |\eta_{P,Q}^*(x')| dP_X(x') + \sum_{i=1}^{n_P+n_Q} |\mathbf{1}\{X_i \in \tilde{A}^k\} - \mathbf{1}\{X_i \in A_{(p)}^j\}| \\
 & \leq \left| \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in \tilde{A}^k\} \tilde{Y}_i - \int_{\tilde{A}^k} \eta_{P,Q}^*(x') dP_X(x') \right| \\
 & \quad + \max_{1 \leq i \leq n} \cdot \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(D_X)} + \|\mathbf{1}\{x \in A_{(p)}^j\} - \mathbf{1}\{x \in \tilde{A}^k\}\|_{L_1(P_X)} \\
 & \leq \left| \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in \tilde{A}^k\} \tilde{Y}_i - \int_{\tilde{A}^k} \eta_{P,Q}^*(x') dP_X(x') \right| + 4\varepsilon. \tag{43}
 \end{aligned}$$

where the last inequality follow from the condition $\mathcal{Y} = \{0, 1\}$.

For any fixed $1 \leq k \leq K$, let the random variable ξ_i be defined by $\xi_i := \mathbf{1}\{X_i \in \tilde{A}^k\} \tilde{Y}_i - \int_{\tilde{A}^k} \eta_{P,Q}^*(x') dP_X(x')$ for $i = 1, \dots, n_P$ and $\tilde{\xi}_i := \mathbf{1}\{X_i \in \tilde{A}^k\} \tilde{Y}_i - \lambda \int_{\tilde{A}^k} \eta_Q^*(x') dP_X(x')$ for $i = n_P + 1, \dots, n_P + n_Q$. Then we have $\mathbb{E} \xi_i = 0$, $\|\xi\|_\infty \leq 1 \vee \lambda$, and $\mathbb{E} \sum_{i=1}^{n_P+n_Q} \tilde{\xi}_i^2 \leq (n_P + \lambda^2 n_Q) \int_{\tilde{A}^k} dP_X(x')$. According to Assumption 1, there holds $\mathbb{E} \sum_{i=1}^{n_P+n_Q} \tilde{\xi}_i^2 \leq \bar{c} \cdot (n_P + \lambda^2 n_Q) \cdot 2^{-p}$. Applying Bernstein's inequality in Lemma 18, we obtain

$$\begin{aligned}
 & \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in A_{(p)}^j\} \tilde{Y}_i - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x') \right| \\
 & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p} \tau (n_P + \lambda^2 n_Q)}{(n_P + n_Q)^2}} + \frac{2\tau \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)}
 \end{aligned}$$

with probability $P^{n_P+n_Q}$ at least $1 - 2e^{-\tau}$. Similar to the proof of Lemma 13, one can show that for any $n \geq N_1$, there holds

$$\begin{aligned}
 & \sup_{1 \leq k \leq K} \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in A_{(p)}^j\} \tilde{Y}_i - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x') \right| \\
 & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p} \tau (n_P + \lambda^2 n_Q)}{(n_P + n_Q)^2}} + \frac{2\tau \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)}
 \end{aligned}$$

with probability $\mathbb{P}^{n_P+n_Q}$ at least $1 - 1/(n_P + n_Q)^2$. This together with (43) yields that

$$\begin{aligned} & \left| \frac{1}{n_P + n_Q} \sum_{i=1}^{n_P+n_Q} \mathbf{1}\{X_i \in A_{(p)}^j\} \tilde{Y}_i - \frac{n_P + \lambda n_Q}{n_P + n_Q} \int_{A_{(p)}^j} \eta_{P,Q}^*(x') dP_X(x') \right| \\ & \leq \sqrt{\frac{\bar{c} \cdot 2^{1-p}(4d+5)(n_P + \lambda^2 n_Q) \log(n_P + n_Q)}{(n_P + n_Q)^2}} + \frac{2(4d+5) \cdot 1 \vee \lambda \cdot \log(n_P + n_Q)}{3(n_P + n_Q)} + \frac{4}{n_P + n_Q}. \end{aligned}$$

This implies the desired result by analogous derivations as in the last part of the proof of Lemma 13. \blacksquare

Proof of Lemma 15 According to the max-edge partition rule, when the depth of the tree i is a multiple of dimension d , each cell of the tree partition is a high-dimensional cube with a side length $2^{-i/d}$. On the other hand, when the depth of the tree p is not a multiple of dimension d , we consider the max-edge tree partition with depth $\lfloor i/d \rfloor$ and $\lceil i/d \rceil$, whose corresponding side length of the higher dimensional cube is $2^{-\lfloor i/d \rfloor}$ and $2^{-\lceil i/d \rceil}$. Note that in the splitting procedure of max-edge partition, the side length of each sub-rectangle decreases monotonically with the increase of p , so the side length of a random tree partition cell is between $2^{-\lceil i/d \rceil}$ and $2^{-\lfloor i/d \rfloor}$. This implies that

$$\sqrt{d} \cdot 2^{-\lceil i/d \rceil} \leq \text{diam}(A_{(i)}^j) \leq \sqrt{d} \cdot 2^{-\lfloor i/d \rfloor}$$

Since $i/d - 1 \leq \lfloor i/d \rfloor \leq \lceil i/d \rceil \leq i/d + 1$, we immediately get $2^{-1}\sqrt{d} \cdot 2^{-i/d} \leq \text{diam}(A_{(i)}^j) \leq 2\sqrt{d} \cdot 2^{-i/d}$. \blacksquare

B.6 Proof of Theorem 5

Proof of Theorem 5 We let $\tilde{f}_{P,Q}^{\pi,\mathcal{M}} := (\tilde{f}_{P,Q}^{\pi} \cdot \mathbf{1}_{[0,1]^d}) \circ \mathcal{M}$ and $X \sim P_X$. The excess risk can be decomposed by

$$\begin{aligned} \mathcal{R}_P(\tilde{f}_{P,Q}^{\pi,\mathcal{M}}) - \mathcal{R}_P^* &= \int_{[0,1]^d} L(x, y, \tilde{f}_{P,Q}^{\pi,\mathcal{M}}(x)) dP(x, y) - \int_{[0,1]^d} L(x, y, f^*(x)) dP(x, y) \\ &+ \int_{[0,1]^{dc}} L(x, y, \tilde{f}_{P,Q}^{\pi,\mathcal{M}}(x)) dP(x, y) - \int_{[0,1]^{dc}} L(x, y, f^*(x)) dP(x, y) \\ &\lesssim \delta^{\frac{\alpha(\beta+1)}{d}} + \mathbb{P}(\mathcal{M}(X) \notin \times_{j=1}^d [0, 1]) \end{aligned}$$

where δ is defined in Theorem 4. Note that

$$\mathbb{P}(\mathcal{M}(X) \notin \times_{j=1}^d [0, 1]) = \mathbb{P}(X \notin \times_{j=1}^d [\hat{a}^j, \hat{b}^j]) \leq \sum_{j=1}^d \mathbb{P}(X^j \notin [\hat{a}^j, \hat{b}^j]).$$

Denote the CDF of X^j as F^j . Since

$$\begin{aligned} \mathbb{P}\left(\left|F^j(\hat{a}^j) - F^j(a^j)\right| > \frac{2 \log n_Q}{n_Q}\right) &= \mathbb{P}\left(F^j(\hat{a}^j) > \frac{2 \log n_Q}{n_Q}\right) \\ &= \left(1 - \frac{2 \log n_Q}{n_Q}\right)^{n_Q} = \frac{1}{n_Q^2} \end{aligned}$$

and similar result holds for b^j , we have $\mathbb{P}(X^j \notin [\hat{a}^j, \hat{b}^j]) \lesssim \frac{\log n_Q}{n_Q}$ with probability $1 - 1/n_Q^2$. Thus, there holds $\mathbb{P}(X \notin \times_{j=1}^d [\hat{a}^j, \hat{b}^j]) \lesssim \frac{\log n_Q}{n_Q}$ with probability $1 - d/n_Q^2$. \blacksquare

B.7 Proof of Proposition 6

Proof of Proposition 6 We prove the conclusion for $k = p_0$ first. When $k = p_0$ and $x \in A_{(k)}^j$, the term $\tilde{\eta}_{P,Q}^{\pi_k}(x) - \mathbb{E}_{Y|X} [\hat{\eta}_{P,Q}^{\pi_k}(x)]$ can be written as

$$\begin{aligned} & \frac{1}{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \cdot \left(\overbrace{\sum_{i=1}^{n_P} (V_i^{Pj} - \mathbb{E}_{Y^P|X^P}[V_i^{Pj}])}^{(I)} + \overbrace{\lambda \sum_{i=1}^{n_Q} (V_i^{Qj} - \mathbb{E}_{Y^Q|X^Q}[V_i^{Qj}])}^{(II)} \right) \\ & + \underbrace{\frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j - \frac{\sum_{i=1}^{n_P} \mathbb{E}_{Y^P|X^P}[V_i^{Pj}] + \lambda \sum_{i=1}^{n_Q} \mathbb{E}_{Y^Q|X^Q}[V_i^{Qj}]}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \cdot \frac{4}{\varepsilon} \sum_{i=1}^{n_P} \zeta_i^j}_{(III)}. \end{aligned}$$

We bound the three parts separately. Since $\mathbb{E}_{Y^P|X^P}[V_i^{Pj}] \leq U_i^{Pj}$, there holds

$$\left| \frac{\sum_{i=1}^{n_P} \mathbb{E}_{Y^P|X^P}[V_i^{Pj}] + \lambda \sum_{i=1}^{n_Q} \mathbb{E}_{Y^Q|X^Q}[V_i^{Qj}]}{\sum_{i=1}^{n_P} U_i^{Pj} + \lambda \sum_{i=1}^{n_Q} U_i^{Qj}} \right| \leq 1.$$

Applying Lemma 12, (III) is bounded by

$$|(III)| \leq 2\sqrt{6\varepsilon^{-2}n_P \log(n_P + n_Q)} \quad (44)$$

with probability R at least $1 - 1/(n_P + n_Q)^2$ for all j . The proofs for (I) and (II) are analogous to the proof of Lemma 14 which utilizes Bernstein's inequality with the VC dimension of possible squares on \mathbb{R}^d . (I) is an independent sum of Bernoulli random variables with at most $\sum_{i=1}^{n_P} U_i^{Pj}$ non-zero elements. Thus by Bernstein's inequality (Lemma 18), we have

$$|(I)| \leq \sqrt{\frac{\sum_{i=1}^{n_P} U_i^{Pj} \cdot \tau}{2}}$$

with probability $\mathbb{P}_{Y^P|X^P}^{n_P}$ at least $1 - e^{-\tau}$. The order of p_0 yields that $p_0 \leq \log(n_P \varepsilon^2 + n_Q)$. Since there are at most 2^{p_0+1} cells in the union of all π_k , there holds

$$\begin{aligned} |(I)| & \leq \sqrt{\frac{\sum_{i=1}^{n_P} U_i^{Pj} \cdot (\tau + (p_0 + 1) \log 2)}{2}} \leq \sqrt{\sum_{i=1}^{n_P} U_i^{Pj} (2 \log(n_P + n_Q) + \log \varepsilon)} \quad (45) \\ & \leq \sqrt{3 \sum_{i=1}^{n_P} U_i^{Pj} \log(n_P + n_Q)} \end{aligned}$$

with probability $P_{Y^P|X^P}^{n_P}$ at least $1 - 1/(n_P + n_Q)^2$. By Lemma 12, for all j , with probability R at least $1 - 1/(n_P + n_Q)^2$, we have

$$\left| \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} - \sum_{i=1}^{n_P} U_i^{Pj} \right| \leq 2\sqrt{3\varepsilon^{-2}n_P \log(n_P + n_Q)}.$$

If $\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \geq 8\varepsilon^{-2}n_P$, there holds

$$\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \geq 8\varepsilon^{-2}n_P - 2\sqrt{3\varepsilon^{-2}n_P \log(n_P + n_Q)} \geq 6\varepsilon^{-2}n_P,$$

as well as

$$\begin{aligned} \frac{4}{3} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} &\geq \sum_{i=1}^{n_P} U_i^{Pj} - 2\sqrt{3\varepsilon^{-2}n_P \log(n_P + n_Q)} + \frac{1}{3} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \\ &\geq \sum_{i=1}^{n_P} U_i^{Pj} - \sqrt{3 \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \log(n_P + n_Q)} + \frac{1}{3} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \geq \sum_{i=1}^{n_P} U_i^{Pj} \end{aligned}$$

for sufficiently large n_P . In this case, we have

$$\begin{aligned} |(I)| + |(III)| &\leq \sqrt{3 \sum_{i=1}^{n_P} U_i^{Pj} \log(n_P + n_Q)} + 2\sqrt{6\varepsilon^{-2}n_P \log(n_P + n_Q)} \\ &\leq 2\sqrt{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \log(n_P + n_Q)} + 2\sqrt{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \log(n_P + n_Q)} \\ &= 4\sqrt{\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \log(n_P + n_Q)}. \end{aligned} \tag{46}$$

If $\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} < 8\varepsilon^{-2}n_P$, there holds

$$\sum_{i=1}^{n_P} U_i^{Pj} \leq \sum_{i=1}^{n_P} \tilde{U}_i^{Pj} + 2\sqrt{3\varepsilon^{-2}n_P \log(n_P + n_Q)} \leq 8\varepsilon^{-2}n_P + \frac{8}{3}\varepsilon^{-2}n_P$$

for sufficiently large n_P . In this case, we get

$$\begin{aligned} |(I)| + |(III)| &\leq \sqrt{3 \sum_{i=1}^{n_P} U_i^{Pj} \log(n_P + n_Q)} + 2\sqrt{6\varepsilon^{-2}n_P \log(n_P + n_Q)} \\ &\leq 2\sqrt{8\varepsilon^{-2}n_P \log(n_P + n_Q)} + 2\sqrt{8\varepsilon^{-2}n_P \log(n_P + n_Q)} \\ &= 4\sqrt{8\varepsilon^{-2}n_P \log(n_P + n_Q)}. \end{aligned} \tag{47}$$

Combing (46) and (47), we get

$$|(I)| + |(III)| \leq 4 \sqrt{(8\varepsilon^{-2}n_P) \vee \left(\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \right) \log(n_P + n_Q)}. \quad (48)$$

Similarly to (45), we have

$$\begin{aligned} |(II)| &\leq \lambda \sqrt{\frac{\sum_{i=1}^{n_P} U_i^{Qj} \cdot (\tau + (p_0 + 1) \log 2)}{2}} \leq \lambda \sqrt{\sum_{i=1}^{n_P} U_i^{Qj} (2 \log(n_P + n_Q) + \log \varepsilon)} \\ &\leq \lambda \sqrt{3 \sum_{i=1}^{n_P} U_i^{Qj} \log(n_P + n_Q)} \end{aligned} \quad (49)$$

with probability $Q_{Y^Q|X^Q}^{n_Q}$ at least $1 - 1/(n_P + n_Q)^2$. Combining (48) and (49), we get

$$\begin{aligned} |(I) + (II) + (III)| &\leq \left(4 \sqrt{(8\varepsilon^{-2}n_P) \vee \left(\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \right)} + \lambda \sqrt{3 \sum_{i=1}^{n_P} U_i^{Qj}} \right) \cdot \sqrt{\log(n_P + n_Q)} \\ &\leq \sqrt{4C_1 (8\varepsilon^{-2}n_P) \vee \left(\sum_{i=1}^{n_P} \tilde{U}_i^{Pj} \right) + C_1 \lambda^2 \sum_{i=1}^{n_P} U_i^{Qj}}. \end{aligned}$$

with probability $P_{Y^P|X^P}^{n_P} \otimes Q_{Y^Q|X^Q}^{n_Q} \otimes R$ at least $1 - 3/(n_P + n_Q)^2$. Here, we define constant C_1 for notation simplicity

$$C_1 = 8 \log(n_P + n_Q).$$

We now prove the conclusion for $k < p_0$ with a slight modification on the above statement. For the term (II) , summing among all nodes with the same ancestor yields the sum of $2^{p_0-k} \cdot n_P$ Laplace random variables, for which we have

$$|(I) + (II) + (III)| \leq \sqrt{4C_1 (2^{p_0-k+3}\varepsilon^{-2}n_P) \vee \left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'} \right) + C_1 \lambda^2 \sum_{j'} \sum_{i=1}^{n_P} U_i^{Qj'}}.$$

The proof is completed by bringing the bound into the decomposition in the beginning. ■

B.8 Proof of Theorem 7

In the proof, depth p_0, p, k are all integers. Without loss of generality, we only write p_0, p, k equal some order which will only cause difference in constants.

Proof of Theorem 7 In the following proof, we give several version of upper bounds of $\left| \tilde{\eta}_{P,Q}^{prune}(x) - \mathbb{E}_{Y|X} [\hat{\eta}_{P,Q}^{prune}(x)] \right| \lesssim b$ for each fixed $x \in \Omega_n$. Then, we consider the region Ω_n which is

$$\Omega_n := \left\{ x \in \mathcal{X} \left| \left| \eta_P^*(x) - \frac{1}{2} \right| \geq c_\Omega \cdot b \right. \right\}.$$

For some c_Ω , by techniques analogous to the proof of Theorem 4, we can show that all x in Ω is assigned to the same label as the Bayes classifier. For each fixed x , let A^{j_k} denote the partition cell in π_k that contains x . We will discuss the stopping time of the pruning process and the resulting point-wise error. We prove the conclusion under two cases respectively.

Case (i): $\left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)}\right) \lesssim \left(\frac{n_Q}{\log(n_P + n_Q)}\right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$. In this case, $\delta^{-1} \lesssim n_Q^{\frac{d}{2\gamma\alpha+d}} \lesssim n_Q \leq \lfloor \frac{d}{2+2d} \log_2 \left(n_P \varepsilon^2 + n_Q^{\frac{2+2d}{d}} \right) \rfloor$. Without loss of generality, we let $\eta_P^*(x) - \frac{1}{2} \geq 0$. Following the analysis in Section A.2, the final statistics will satisfy

$$\begin{aligned} \frac{|\hat{\eta}_{P,Q}^{\pi_k}(x) - \frac{1}{2}|}{r_k^j} &= \sqrt{\frac{\left(\sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}\right)^2}{u_k} \left(\hat{\eta}_P^{\pi_k}(x) - \frac{1}{2}\right)^2 + \sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left(\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}\right)^2} \\ &\geq \sqrt{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left|\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2}\right|}. \end{aligned}$$

By Proposition 6, there exists a constant \bar{C}_p such that

$$\left|\hat{\eta}_Q^{\pi_k}(x) - \mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x)\right| \leq \bar{C}_p \sqrt{\frac{\log(n_P + n_Q)}{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'}}} \quad (50)$$

with probability $P_{Y^P|X^P}^{n_P} \otimes Q_{Y^Q|X^Q}^{n_Q} \otimes R$ at least $1 - 3/(n_P + n_Q)^2$. Consider the region

$$\Omega_n := \left\{x \in \mathcal{X} \left| \left| \eta_P^*(x) - \frac{1}{2} \right| \geq C_\Omega \cdot \delta^{\alpha/d} \right\}.$$

Let $2^{-k} = \delta$. For each $x \in \Delta_n$, by Lemma 10, there exist $\underline{C}_{p1}, \underline{C}_{p2}$ such that,

$$\min_{x \in A^{j_k}} \mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \geq \underline{C}_p^1 \left(\mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right) \geq \underline{C}_{p1} (C_\Omega - \underline{C}_{p2}) \delta^{\gamma\alpha/d}. \quad (51)$$

Also, by Lemma 13, there exists constant \underline{C}_{p3} such that

$$\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \geq \underline{C}_{p3} \delta \cdot n_Q \quad (52)$$

with probability $P_{X^P}^{n_P} \otimes Q_{X^Q}^{n_Q}$ at least $1 - 1/(n_P + n_Q)^2$. (51) and (52) together lead to

$$\begin{aligned} \sqrt{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left(\mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right)} &\geq \underline{C}_{p1} \underline{C}_{p3}^{1/2} (C_\Omega - \underline{C}_{p2}) \sqrt{n_Q} \cdot \delta^{\frac{2\gamma\alpha+d}{2d}} \\ &\geq 2^{\frac{2\gamma\alpha+d}{2d}} \underline{C}_{p1} \underline{C}_{p3}^{1/2} (C_\Omega - \underline{C}_{p2}) \sqrt{\log(n_P + n_Q)}. \end{aligned} \quad (53)$$

For sufficiently large constant C_Ω , (53) yields that the choice of $2^{-k} = \delta$ leads to

$$\sqrt{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left(\mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right)} \geq 3\bar{C}_p \sqrt{\log(n_P + n_Q)}. \quad (54)$$

Combining (50) and (54), there holds the following

$$\begin{aligned} & \sqrt{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left(\hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right)} \\ & \geq \sqrt{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left(\mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x) - \frac{1}{2} \right)} - \sqrt{\sum_{j'} \sum_{i=1}^{n_Q} U_i^{Qj'} \left| \hat{\eta}_Q^{\pi_k}(x) - \mathbb{E}_{Y|X} \hat{\eta}_Q^{\pi_k}(x) \right|} \\ & > \bar{C}_p \sqrt{\log(n_P + n_Q)}. \end{aligned}$$

for the choice of $2^k = \delta^{-1}$. Thus, for $x \in \Omega$ and $x \in A_{(p_0)}^j$, the terminating condition will be triggered at least at $2^k = \delta^{-1}$, indicating that $2^{k_j} \geq \delta^{-1}$ with probability $P^{n_P} \otimes Q^{n_Q} \otimes R$ at least $1 - 4/(n_P + n_Q)^2$. Thus, in this case, we have

$$\mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_{k_j}}(x) \right] - \frac{1}{2} = \mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_{k_j}}(x) \right] - \tilde{\eta}_{P,Q}^{\pi_{k_j}}(x) + \tilde{\eta}_{P,Q}^{\pi_{k_j}}(x) - \frac{1}{2} \geq r_k^j - r_k^j \geq 0, \quad (55)$$

which implies that the signs of $\mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_{k_j}}(x) \right]$ and $\tilde{\eta}_{P,Q}^{\pi_{k_j}}(x)$ are identical. By steps analogous to proof of Proposition 10, there holds $\mathbb{E}_{Y|X} \left[\hat{\eta}_{P,Q}^{\pi_{k_j}}(x) \right] - \frac{1}{2} > 0$ for $x \in \Omega_n$. Thus we show that, for each $x \in \Omega_n$ and $x \in A_{(p_0)}^j$, the label is assigned correctly as the Bayes classifier. Therefore, by Assumption 2, we have

$$\begin{aligned} \mathcal{R}_P(\tilde{f}_{P,Q}^{prune}) - \mathcal{R}_P^* &= \int_{\mathcal{X}} \left| \eta_P^*(x) - \frac{1}{2} \right| \mathbf{1} \left\{ \tilde{f}_{P,Q}^{\pi}(x) \neq f_P^*(x) \right\} dP_X(x) \\ &= \int_{\Omega_n^c} \left| \eta_P^*(x) - \frac{1}{2} \right| dP_X(x) \lesssim \delta^{(\beta+1)\alpha/d}. \end{aligned}$$

Case (ii) $\left(\frac{n_P \varepsilon^2}{\log(n_P + n_Q)} \right) \gtrsim \left(\frac{n_Q}{\log(n_P + n_Q)} \right)^{\frac{2\alpha+2d}{2\gamma\alpha+d}}$. We first note that, with $n_P^{-\frac{\alpha}{2\alpha+2d}} \lesssim \varepsilon \lesssim n_P^{\frac{d}{4+2d}}$, there holds $2^{p_0-k+3} \varepsilon^{-2} n_P \geq \sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}$ with high probability, i.e. Algorithm 2 will determine true in first if block. Since $\varepsilon \lesssim n_P^{\frac{d}{4+2d}}$, there holds $2^{p_0} \geq (n_P \varepsilon^2)^{\frac{d}{2+2d}} \gtrsim \varepsilon^2$. By Lemma 13, when $\lambda = 0$, we have $\frac{1}{n_P} \sum_{j'} \sum_{i=1}^{n_P} U_i^{Pj'} \asymp 2^{-k}$. Together, we have

$$2^{p_0-k+3} \varepsilon^{-2} n_P \gtrsim \sum_{j'} \sum_{i=1}^{n_P} U_i^{Pj'}. \quad (56)$$

For the Laplace noise, there holds

$$2^{p_0-k} \varepsilon^{-2} n_P \gtrsim \sqrt{2^{p_0-k} \varepsilon^{-2} n_P \log(n_P + n_Q)} \gtrsim \frac{4}{\varepsilon} \sum_{j'} \sum_{i=1}^{n_P} \zeta_i^{Pj'} \quad (57)$$

due to Lemma 12. Adding (56) and (57) shows $2^{p_0-k+3}\varepsilon^{-2}n_P \geq \sum_{j'} \sum_{i=1}^{n_P} \tilde{U}_i^{Pj'}$ holds with probability $1 - 2/(n_P + n_Q)^2$. Under this case, if $k_j > \lfloor \frac{d}{2+2d} \log_2(n_P \varepsilon^2) \rfloor$ for all $j \in \mathcal{I}_{p_0}$, we end up with $\tilde{\eta}_{P,Q}^{prune} = \sum_j s_k^j \cdot \mathbf{1}\{A_{(p_0)}^j\}$. We define

$$\Omega_n := \left\{ x \in \mathcal{X} \mid \left| \eta_P^*(x) - \frac{1}{2} \right| \geq C_\Omega \cdot \delta^{\alpha(\alpha+d)/d(1+d)} \right\}$$

for some C_Ω large enough. Assume $\eta_P^*(x) > 1/2$. If x' belongs to the same A^{k_j} as x , then

$$\eta_P^*(x') - \frac{1}{2} \geq \eta_P^*(x) - \frac{1}{2} - C' 2^{-k_j \alpha/d} \geq \eta_P^*(x) - \frac{1}{2} - C' \delta^{\alpha(\alpha+d)/d(1+d)} \geq 0$$

for some C' . Analogous to (55), all $x \in \Omega_n$ are correctly classified and thus

$$\begin{aligned} \mathcal{R}_P(\tilde{f}_{P,Q}^{prune}) - \mathcal{R}_P^* &= \int_{\mathcal{X}} \left| \eta_P^*(x) - \frac{1}{2} \right| \mathbf{1}\left\{ \tilde{f}_{P,Q}^\pi(x) \neq f_P^*(x) \right\} dP_X(x) \\ &= \int_{\Omega_n^c} \left| \eta_P^*(x) - \frac{1}{2} \right| dP_X(x) \lesssim \delta^{(\beta+1)\alpha(\alpha+d)/d(1+d)}. \end{aligned}$$

If there exists some $j \in \mathcal{I}_{p_0}$ such that the termination is triggered, i.e. we will end up with $\tilde{f}_P^{\pi_P}$, the proof becomes analogous to that of Theorem 4. Specifically, the trade off (30) will have a larger approximation error $\delta^{\alpha(\alpha+d)/d(1+d)}$ and smaller sample error as well as privacy error. Thus, the final rate is dominated by the approximation error, which becomes $\delta^{(\beta+1)\alpha(\alpha+d)/d(1+d)}$. In Case (ii), the overall incidents happen with probability $P^{n_P} \otimes Q^{n_Q} \otimes R$ at least $1 - 4/(n_P + n_Q)^2$. ■

Appendix C. Experiment Details

C.1 Real Datasets Pre-processing

We present additional information on the data sets including data source and pre-processing details. The data sets are sourced from **UCI Machine Learning Repository** (Dua and Graff, 2017) and **Kaggle**.

Anonymity: The Anonymity data set (piAI, 2019) contains poll results about Internet privacy from a large number of people, including answers about whether the interviewee believes United States law provides reasonable privacy protection for Internet users or their conservativeness about Internet privacy. We use other variables to predict whether the interviewee has ever tried to mask his/her identity when using the Internet. We randomly select 50 samples as the public data.

Diabetes: This data set from Pore (2023) contains a diverse range of health-related attributes, meticulously collected to aid in the development of predictive models for identifying individuals at risk of diabetes. We randomly select 200 samples as the public data.

Email: The data set is used for spam or not-spam classification (Biswas, 2019). For each row, the count of each word (column) in that email (row) is stored in the respective cells.

Thus, information regarding all emails is stored in a compact data frame. We randomly select 200 samples as the public data.

Employ: The data set (Hamoutn, 2022) was collected from different agencies in the Philippines consisting of mock interview performances of students. The final task is to decide whether the student got employed. We randomly select 200 samples as the public data.

Rice: The Rice data set (Cammeo and Osmancik, 2019) contains morphological features obtained for two grains of rice to classify them. We randomly select 300 samples as the public data.

Census: This data set is the well-known census data (Becker and Kohavi, 1996) and is used to predict whether income exceeds 50,000 per year. We select the samples whose native country is the U.S. to be public data and others to be private data.

Election: This data set from (Winner and Thacker, 2022) contains county-level demographics and whether or not Bill Clinton won each county during the 1992 U.S. presidential election. The goal of this data set is to predict if Clinton won a county. We use samples whose state is TX as public data and the rest as private data.

Employee: The Employee data (Elmetwally, 2023) contains information about employees in a company, including their educational backgrounds, work history, and employment-related factors. We use these variables to predict whether the employee leaves or not. We select the sample to be public if the employee joined the company before 2012.

Jobs: The data (Tankha, 2023) contains a comprehensive collection of information regarding job applicants and their respective employability scores to predict whether the applicant has been hired. We select the samples whose native country is the U.S. to be public data and others to be private data.

Landcover: The data is from Johnson and Iizuka (2016), which contains Landsat time-series satellite imagery information on given pixels and their corresponding land cover class labels (farm, forest, water, etc.) obtained from multiple sources. We classify whether the image comes from a farm or a forest. Within this data set, there are two kinds of label sources: labels obtained from Open-StreetMap and labels that are accurately labeled by experts. We take the latter as public data and the former as private data.

Taxi: The Taxi data set was obtained from the Differential Privacy Temporal Map Challenge (DrivenData, 2021), which aims to develop algorithms that preserve data utility while guaranteeing individual privacy protection. The data set contains quantitative and categorical information about taxi trips in Chicago, including time, distance, location, payment, and service provider. These features include the identification number of taxis (*taxi_id*), time of each trip (*seconds*), the distance of each trip (*miles*), index of the zone where the trip starts (*pca*), index of the zone where the trip ends (*dca*), service provider (*company*), the method used to pay for the trip (*payment_type*), amount of tips (*tips*) and fares (*fare*). We use the other variable to predict whether the passenger is paying by credit card or by cash. We select data from one company as public data and others as private data.

References

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016.
- Noga Alon, Raef Bassily, and Shay Moran. Limits of private learning with access to public data. *Advances in neural information processing systems*, 32, 2019.
- Ehsan Amid, Arun Ganesh, Rajiv Mathews, Swaroop Ramaswamy, Shuang Song, Thomas Steinke, Vinith M Suriyakumar, Om Thakkar, and Abhradeep Thakurta. Public data-assisted mirror descent for private model training. In *International Conference on Machine Learning*. PMLR, 2022.
- Jean-Yves Audibert and Alexandre B Tsybakov. Fast learning rates for plug-in classifiers. *Annals of statistics*, 35(2):608–633, 2007.
- Raef Bassily, Om Thakkar, and Abhradeep Guha Thakurta. Model-agnostic private learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- Raef Bassily, Albert Cheu, Shay Moran, Aleksandar Nikolov, Jonathan Ullman, and Steven Wu. Private query release assisted by public data. In *International Conference on Machine Learning*, pages 695–703. PMLR, 2020.
- Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996. URL <https://doi.org/10.24432/C5XW20>.
- Shai Ben-David, Alex Bie, Clement Louis Canonne, Gautam Kamath, and Vikrant Singhal. Private distribution learning with public data: The view from sample compression. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Thomas Berrett and Cristina Butucea. Classification under local differential privacy. *arXiv preprint arXiv:1912.04629*, 2019.
- Thomas B Berrett, László Györfi, and Harro Walk. Strongly universally consistent nonparametric regression and classification with privatised data. *Electronic Journal of Statistics*, 15:2430–2453, 2021.
- Tom Berrett and Yi Yu. Locally private online change point detection. *Advances in Neural Information Processing Systems*, 34:3425–3437, 2021.
- Alex Bie, Gautam Kamath, and Vikrant Singhal. Private estimation with public data. In *Advances in Neural Information Processing Systems*, 2022.
- Balaka Biswas. Email spam classification dataset csv, 2019. URL <https://www.kaggle.com/datasets/balaka18/email-spam-classification-dataset-csv/data>.
- L Breiman. Classification and regression trees. *The Wadsworth & Brooks/Cole*, 1984.

- Cristina Butucea, Amandine Dubois, Martin Kroll, and Adrien Saumard. Local differential privacy: Elbow effect in optimal density estimation and adaptation over besov ellipsoids. *Bernoulli*, 26(3):1727–1764, 2020.
- T Tony Cai and Hongji Wei. Transfer learning for nonparametric classification: Minimax rate and adaptive classifier. *The Annals of Statistics*, 49(1), 2021.
- Yuchao Cai, Yuheng Ma, Yiwei Dong, and Hanfang Yang. Extrapolated random tree for regression. In *Proceedings of the 40th International Conference on Machine Learning*, pages 3442–3468. PMLR, 2023.
- Cammeo and Osmancik. Rice. UCI Machine Learning Repository, 2019. URL <https://doi.org/10.24432/C5MW4Z>.
- Kamalika Chaudhuri and Sanjoy Dasgupta. Rates of convergence for nearest neighbor classification. *Advances in Neural Information Processing Systems*, 27, 2014.
- Kamalika Chaudhuri and Staal A Vinterbo. A stability-based validation procedure for differentially private machine learning. *Advances in Neural Information Processing Systems*, 26, 2013.
- Yuval Dagan and Vitaly Feldman. Interaction is necessary for distributed learning with privacy or communication constraints. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 450–462, 2020.
- Amit Daniely and Vitaly Feldman. Locally private learning without interaction requires separation. *Advances in neural information processing systems*, 32, 2019.
- DrivenData. Differential privacy temporal map challenge: Sprint 3 (pre-screened arena), 2021. URL <https://www.drivendata.org/competitions/77/deid2-sprint-3-prescreened/page/332/>.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- John Duchi, Michael Jordan, and Martin Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521): 182–201, 2018.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- Tawfik Elmetwally. Employee dataset, 2023. URL <https://www.kaggle.com/datasets/tawfikelmetwally/employee-dataset/data>.
- Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067, 2014.

- Arun Ganesh, Mahdi Haghifam, Milad Nasr, Sewoong Oh, Thomas Steinke, Om Thakkar, Abhradeep Guha Thakurta, and Lun Wang. Why is public pretraining necessary for private model training? In *International Conference on Machine Learning*, pages 10611–10627. PMLR, 2023.
- Xin Gu, Gautam Kamath, and Zhiwei Steven Wu. Choosing public datasets for private machine learning via gradient subspace distance. *arXiv preprint arXiv:2303.01256*, 2023.
- László Györfi and Martin Kroll. On rate optimal private regression under local differential privacy. *arXiv preprint arXiv:2206.00114*, 2022.
- Anas Hamoutn. Students’ employability dataset - philippines, 2022. URL <https://www.kaggle.com/datasets/anashamoutni/students-employability-dataset/data>.
- Brian A Johnson and Kotaro Iizuka. Integrating openstreetmap crowdsourced data and landsat time-series imagery for rapid land use/land cover (lulc) mapping: Case study of the laguna de bay area of the philippines. *Applied Geography*, 67:140–149, 2016.
- Matthew Joseph, Janardhan Kulkarni, Jieming Mao, and Steven Z Wu. Locally private gaussian estimation. *Advances in Neural Information Processing Systems*, 32, 2019.
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. *Advances in neural information processing systems*, 27, 2014.
- Peter Kairouz, Monica Ribero Diaz, Keith Rush, and Abhradeep Thakurta. (nearly) dimension independent private erm with adagrad rates via publicly estimated subspaces. In *Conference on Learning Theory*, pages 2717–2746. PMLR, 2021.
- Michael R. Kosorok. *Introduction to Empirical Processes and Semiparametric Inference*. Springer Series in Statistics. Springer, New York, 2008.
- Martin Kroll. On density estimation at a fixed point under local differential privacy. *Electronic Journal of Statistics*, 15:1783–1813, 2021.
- OV Lepskii. Asymptotically minimax adaptive estimation. i: Upper bounds. optimally adaptive estimates. *Theory of Probability & Its Applications*, 36(4):682–697, 1992.
- Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. In *The Tenth International Conference on Learning Representations*, 2022.
- Chong Liu, Yuqing Zhu, Kamalika Chaudhuri, and Yu-Xiang Wang. Revisiting model-agnostic private learning: Faster rates and active learning. *The Journal of Machine Learning Research*, 22(1):11936–11979, 2021.
- Jingcheng Liu and Kunal Talwar. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019.
- Andrew Lowy, Zeman Li, Tianjian Huang, and Meisam Razaviyayn. Optimal differentially private learning with public data. *arXiv preprint arXiv:2306.15056*, 2023.

- Yuheng Ma, Han Zhang, Yuchao Cai, and Hanfang Yang. Decision tree for locally private estimation with public data. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Shubhankar Mohapatra, Sajin Sasy, Xi He, Gautam Kamath, and Om Thakkar. The role of adaptive optimizers for honest private hyperparameter selection. In *Proceedings of the aaai conference on artificial intelligence*, volume 36, pages 7806–7813, 2022.
- Milad Nasr, Saeed Mahloujifar, Xinyu Tang, Prateek Mittal, and Amir Houmansadr. Effectively using public data in privacy preserving machine learning. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202. PMLR, 2023.
- Nicolas Papernot and Thomas Steinke. Hyperparameter tuning with renyi differential privacy. In *International Conference on Learning Representations*, 2021.
- Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. In *International Conference on Learning Representations*, 2017.
- Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. Scalable private learning with pate. In *International Conference on Learning Representations*, 2018.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- piAI. Internet privacy poll, 2019. URL <https://www.kaggle.com/datasets/econdata/internet-privacy-poll>.
- Nandita Pore. Healthcare diabetes dataset, 2023. URL <https://www.kaggle.com/datasets/nanditapore/healthcare-diabetes/data>.
- Swaroop Ramaswamy, Om Thakkar, Rajiv Mathews, Galen Andrew, H Brendan McMahan, and Françoise Beaufays. Training production language models without memorizing user data. *arXiv preprint arXiv:2009.10031*, 2020.
- Ievgen Redko, Emilie Morvant, Amaury Habrard, Marc Sebban, and Younès Bennani. A survey on domain adaptation theory: learning bounds and theoretical guarantees. *arXiv preprint arXiv:2004.11829*, 2020.
- Richard J Samworth. Optimal weighted nearest neighbour classifiers. *The Annals of Statistics*, pages 2733–2763, 2012.
- Adam Smith, Abhradeep Thakurta, and Jalaj Upadhyay. Is interaction necessary for distributed private learning? In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 58–77. IEEE, 2017.

- Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. Stochastic gradient descent with differentially private updates. In *2013 IEEE global conference on signal and information processing*, pages 245–248. IEEE, 2013.
- Jinyan Su, Jinhui Xu, and Di Wang. On pac learning halfspaces in non-interactive local privacy model with public unlabeled data. In *Asian Conference on Machine Learning*, pages 927–941. PMLR, 2023.
- Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. Privacy loss in apple’s implementation of differential privacy on macos 10.12. *arXiv preprint arXiv:1709.02753*, 2017.
- Ayush Tankha. Employability classification of over 70,000 job applicants, 2023. URL <https://www.kaggle.com/datasets/ayushtankha/70k-job-applicants-data-human-resource/data>.
- Florian Tramèr, Gautam Kamath, and Nicholas Carlini. Considerations for differentially private learning with large-scale public pretraining. *arXiv preprint arXiv:2212.06470*, 2022.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York, 2009. ISBN 978-0-387-79051-0. doi: 10.1007/b13794. URL <http://hfhhc391f4815d8064db7sqbf60x0xf60c6kpo.faxb.libproxy.ruc.edu.cn/10.1007/b13794>. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Di Wang and Jinhui Xu. Principal component analysis in the local differential privacy model. *Theoretical computer science*, 809:296–312, 2020.
- Di Wang, Marco Gaboardi, and Jinhui Xu. Empirical risk minimization in non-interactive local differential privacy revisited. *Advances in Neural Information Processing Systems*, 31, 2018.
- Di Wang, Lijie Hu, Huanyu Zhang, Marco Gaboardi, and Jinhui Xu. Generalized linear models in non-interactive local differential privacy with public data. *Journal of Machine Learning Research*, 24(132):1–57, 2023a.
- Hua Wang, Sheng Gao, Huanyu Zhang, Weijie J Su, and Milan Shen. Dp-hypo: An adaptive private framework for hyperparameter optimization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b.
- Jun Wang and Zhi-Hua Zhou. Differentially private learning with small public data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2020.
- Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

- Frank Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics*, pages 196–202. Springer, 1992.
- Larry Winner and Jared Thacker. 1992 u.s presidential election, 2022. URL <https://www.kaggle.com/datasets/jwt0024/1992-us-presidential-election>.
- Xiaotong Wu, Lianying Qi, Jiaquan Gao, Genlin Ji, and Xiaolong Xu. An ensemble of random decision trees with local differential privacy in edge computing. *Neurocomputing*, 485:181–195, 2022.
- Da Yu, Huishuai Zhang, Wei Chen, and Tie-Yan Liu. Do not let privacy overbill utility: Gradient embedding perturbation for private learning. In *International Conference on Learning Representations*, 2021a.
- Da Yu, Huishuai Zhang, Wei Chen, Jian Yin, and Tie-Yan Liu. Large scale private learning via low-rank reparametrization. In *International Conference on Machine Learning*, pages 12208–12218. PMLR, 2021b.
- Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A. Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, Sergey Yekhanin, and Huishuai Zhang. Differentially private fine-tuning of language models. In *The Tenth International Conference on Learning Representations*, 2022.
- Da Yu, Sivakanth Gopi, Janardhan Kulkarni, Zinan Lin, Saurabh Naik, Tomasz Lukasz Religa, Jian Yin, and Huishuai Zhang. Selective pre-training for private fine-tuning. *arXiv preprint arXiv:2305.13865*, 2023.
- Kai Zheng, Wenlong Mou, and Liwei Wang. Collect at once, use effectively: Making non-interactive locally private learning possible. In *International Conference on Machine Learning*, pages 4130–4139. PMLR, 2017.
- Yingxue Zhou, Steven Wu, and Arindam Banerjee. Bypassing the ambient dimension: Private sgd with gradient subspace identification. In *International Conference on Learning Representations*, 2021.