

Statistica Sinica Preprint No: SS-2022-0186

Title	On Rate Optimal Private Regression Under Local Differential Privacy
Manuscript ID	SS-2022-0186
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202022.0186
Complete List of Authors	László Györfi and Martin Kroll
Corresponding Authors	Martin Kroll
E-mails	martin.kroll@uni-bayreuth.de

ON RATE OPTIMAL PRIVATE REGRESSION UNDER LOCAL DIFFERENTIAL PRIVACY

László Györfi and Martin Kroll

Budapest University of Technology and Economics

Ruhr-Universität Bochum

Universität Bayreuth

Abstract: We consider the problem of estimating a regression function from anonymized data in the framework of local differential privacy. We propose a novel partitioning estimate of the regression function, derive a rate of convergence for the excess prediction risk over Hölder classes, and prove a matching lower bound. In contrast to the existing literature on the problem, the so-called strong density assumption on the design distribution is obsolete.

Key words and phrases: Nonparametric regression, local differential privacy, partitioning estimate, rate of convergence, minimax lower bound.

1. Introduction

Let $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$ be a pair of random variables with explanatory variable $X \in \mathbb{R}^d$ and real-valued response Y satisfying $\mathbf{E}[Y^2] < \infty$. We denote by

μ the distribution of X , that is, $\mu(A) = \mathbf{P}(X \in A)$ for all measurable sets $A \subseteq \mathbb{R}^d$. Then, the *regression function*

$$m(x) = \mathbf{E}[Y|X = x]$$

is well defined for μ -almost all $x \in \mathbb{R}^d$. For any measurable function $g: \mathbb{R}^d \rightarrow \mathbb{R}$, we have

$$\mathbf{E}[(g(X) - Y)^2] = \mathbf{E}[(m(X) - Y)^2] + \mathbf{E}[(m(X) - g(X))^2],$$

and therefore, setting

$$L^* = \mathbf{E}[(m(X) - Y)^2],$$

it follows that

$$\mathbf{E}[(g(X) - Y)^2] = L^* + \int (m(x) - g(x))^2 \mu(dx).$$

Thus, measuring the performance of an estimator \hat{m} of m using the loss function

$$L(m, \hat{m}) := \int (m(x) - \hat{m}(x))^2 \mu(dx)$$

may be interpreted as the excess prediction risk at a new design point X , distributed according to the design measure μ .

In this study, we consider piecewise constant estimators of the regression function m based on cubic partitions. Let $\mathcal{P}_h = \{A_{h,1}, A_{h,2}, \dots\}$ be such a cubic partition of \mathbb{R}^d , with cubic cells $A_{h,j}$ of volume h^d . The raw data \mathcal{D}_n

are assumed to be independent and identically distributed (i.i.d.) copies of the random vector (X, Y) ,

$$\mathcal{D}_n := \{(X_1, Y_1), \dots, (X_n, Y_n)\}. \quad (1.1)$$

Put

$$\nu_n(A_{h,j}) = \frac{1}{n} \sum_{i=1}^n Y_i \mathbf{1}_{\{X_i \in A_{h,j}\}}$$

and

$$\mu_n(A_{h,j}) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{X_i \in A_{h,j}\}}.$$

Then, a standard regression estimate is defined by

$$m_n(x) = \frac{\nu_n(A_{h,j})}{\mu_n(A_{h,j})},$$

for any $x \in A_{h,j}$, with the usual convention that $0/0 := 0$. Theorem 4.3 in the monograph (Györfi et al., 2002) states an upper bound on the rate of convergence for this partitioning estimate for Lipschitz continuous regression functions. Extending it to the more general case of Hölder continuous functions is straightforward, and yields the following result.

Theorem 1.1. *If the function*

$$\sigma^2(x) := \mathbf{Var}(Y|X = x)$$

is bounded, m is (β, C) -Hölder smooth with index $0 < \beta \leq 1$, that is,

$$|m(x) - m(x')| \leq C\|x - x'\|^\beta, \quad x, x' \in \mathbb{R}^d, \quad (1.2)$$

and X is bounded, then

$$\mathbf{E} \left[\int (m(x) - m_n(x))^2 \mu(\mathrm{d}x) \right] \lesssim \frac{1}{nh^d} + h^{2\beta}. \quad (1.3)$$

In particular, the choice $h = h_n \asymp n^{-1/(2\beta+d)}$ realizes the best compromise between the two antagonistic terms on the right-hand side of (1.3), and the resulting rate is $n^{-2\beta/(2\beta+d)}$. Standard arguments for nonparametric lower bounds, for instance based on Assouad's lemma, show that this rate is indeed optimal.

The main purpose of this study is to provide an analogue of Theorem 1.1 for the case when the raw data \mathcal{D}_n are not directly accessible but only a suitably anonymized surrogate. More precisely, the anonymized data must satisfy a *local differential privacy* (LDP) condition. Our work is motivated by the recent work (Berrett et al., 2021), which provides a first step in this direction. (Berrett et al., 2021) considers a private partitioning estimate, and derives the upper bound $n^{-1/(d+1)}$ on the rate of convergence for Lipschitz continuous functions ($\beta = 1$). However, this rate is established under a quite restrictive assumption on the design distribution μ (called the *strong density assumption* (SDA) in (Berrett et al., 2021)). Moreover, it was con-

jectured that the rate of convergence may be arbitrarily slow when the SDA is not fulfilled. In this paper, we show that this conjecture does not hold, and propose an estimator that attains the rate $n^{-\beta/(\beta+d)}$ without the SDA. Note that we do not even need the existence of a Lebesgue density for μ . We complement our upper bound by proving a minimax lower bound, showing that the rate $n^{-\beta/(\beta+d)}$ is indeed optimal. This agrees with what can be expected from similar problems, such as nonparametric density estimation, in which a similar deterioration of the rate of convergence has been found (Duchi et al., 2018; Butucea et al., 2020). The same phenomenon has also been observed for classification problems (Berrett and Butucea, 2019).

The rest of the paper is organized as follows. In Section 2, we recap the notion of LDP and introduce a suitable anonymization of the raw data that generates locally differentially private data. In Section 3, we introduce a modification of the classical partitioning estimate of the regression function that is based only on the availability of the anonymized data, and derive a convergence rate for this estimator. In Section 4 we prove a matching lower bound that coincides with the upper one. All proofs are gathered in Section 5.

2. Anonymization of the raw data

In this section, we briefly recall the definition of LDP before we describe our privacy mechanism. In the language of probability theory, *non-interactive* privacy mechanisms are given by conditional distributions Q_i , for $i = 1, \dots, n$, that draw privatized data Z_i from potentially different measurable spaces $(\mathcal{Z}_i, \mathcal{L}_i)$. More precisely, given the raw data $(X_i, Y_i) = (x_i, y_i)$, one draws Z_i according to a probability measure defined as $Q_i(A|(X_i, Y_i) = (x_i, y_i))$, for any $A \in \mathcal{L}_i$. Such a non-interactive mechanism is *local*, because any data holder can independently generate privatized data.

For a privacy parameter $\alpha \in [0, \infty]$, any non-interactive privacy mechanism is said to be an α -*locally differentially private* mechanism if the condition

$$\frac{Q_i(A|(X_i, Y_i) = (x, y))}{Q_i(A|(X_i, Y_i) = (x', y'))} \leq \exp(\alpha) \quad (2.1)$$

is satisfied for any $A \in \mathcal{L}_i$ and all potential values $(x, y), (x', y')$ of the raw data. The set of all α -locally differentially private mechanisms is denoted by \mathcal{Q}_α .

We now state the specific privacy mechanism we consider for the anonymization of the raw data \mathcal{D}_n in (1.1). Our approach follows the Laplace perturbation technique already considered in (Duchi et al., 2018; Berrett and

Butucea, 2019; Butucea et al., 2020; Berrett et al., 2021; Györfi and Kroll, 2023). In order to define the privatized data, we first choose a closed Euclidean ball $B = \{x \in \mathbb{R}^d : \|x\| \leq r\}$ of radius $r > 0$ centred at the origin. Furthermore, let $A_{h,1}, A_{h,2}, \dots$ be a partition of \mathbb{R}^d consisting of cubic cells with volume h^d , for some $h = h_n > 0$. Without loss of generalization, we assume that the cells are numbered such that $A_{h,j} \cap B \neq \emptyset$ when $j \leq N_n$, for some nonnegative integer N_n , and $A_{h,j} \cap B = \emptyset$ otherwise. For a threshold $T > 0$ and any $x \in \mathbb{R}$, we define $[x]_T = \max\{-T, \min\{x, T\}\}$. In our privacy setup, the data holder of the i th datum (X_i, Y_i) generates and transmits to the statistician the data

$$Z_{ij} := [Y_i]_T \mathbf{1}_{\{X_i \in A_{h,j}\}} + \sigma_Z \varepsilon_{ij}, \quad j = 1, \dots, N_n, \quad (2.2)$$

and

$$W_{ij} := \mathbf{1}_{\{X_i \in A_{h,j}\}} + \sigma_W \zeta_{ij}, \quad j = 1, \dots, N_n, \quad (2.3)$$

with noise levels $\sigma_Z, \sigma_W > 0$, and $\varepsilon_{ij}, \zeta_{ij}$ ($i = 1, \dots, n, j = 1, \dots, N_n$) are independent centered Laplace random variables with unit variance. This means that the individual with index i generates noisy data for any cell $A_{h,j}$ that has a nontrivial intersection with the ball B . The noise levels σ_Z and σ_W should be chosen sufficiently large in dependence on the desired privacy level α to make the overall mechanism satisfy α -LDP. It is shown

in (Berrett et al., 2021, p. 2438) that the choices

$$\sigma_W^2 = 32/\alpha^2 \quad \text{and} \quad \sigma_Z^2 = 32T^2/\alpha^2 \quad (2.4)$$

ensure α -LDP.

Remark 2.1. The privacy mechanism defined by (2.2) and (2.3) is not the only possibility. For instance, one could build on the randomized response technique discussed in (Duchi et al., 2018), and we conjecture that data privatized in this way will attain the same rates of convergence as derived below. However, because our work is motivated by (Berrett et al., 2021), we use Laplace perturbation.

3. Rate of convergence

For a threshold $t > 0$, the work (Berrett et al., 2021) considers the estimator

$$\tilde{m}_n(x) = \frac{\tilde{\nu}_n(A_{h,j})}{\tilde{\mu}_n(A_{h,j})} \mathbf{1}_{\{\tilde{\mu}_n(A_{h,j}) \geq t\}} \mathbf{1}_{\{j \leq N_n\}} \quad \text{when } x \in A_{h,j}, \quad (3.1)$$

where

$$\tilde{\nu}_n(A_{h,j}) = \frac{1}{n} \sum_{i=1}^n Z_{ij} \quad \text{and} \quad \tilde{\mu}_n(A_{h,j}) = \frac{1}{n} \sum_{i=1}^n W_{ij}.$$

In (Berrett et al., 2021), the convergence rate $n^{-1/(d+1)}$ is derived (up to a logarithmic term) for Lipschitz continuous functions by specializing (3.1) with $h = h_n \asymp n^{-1/(2d+2)}$ and $t = t_n \asymp h_n^d / \sqrt{\log n}$. However, their proof is

essentially based on the validity of the strong density assumption (SDA), which means that if $\mu(A_{h,j}) > 0$, then

$$\mu(A_{h,j}) \geq ch^d, \quad j = 1, \dots, N_n, \quad (3.2)$$

for some constant $c > 0$. Moreover, instead of taking (3.2) as an assumption, it was deduced mistakenly from the existence of a density that is lower bounded from below on its support. Apart from this minor flaw in the proof, imposing (3.2) is rather artificial, because it is a condition on the design distribution μ and on the relationship between the distribution and the sets $A_{h,j}$ of the chosen partition. Because this condition is rarely justified in practice, it is desirable to eliminate it from the prerequisites. In general, without the SDA, the convergence rate $n^{-1/(d+1)}$ is not attainable using the estimator of (Berrett et al., 2021) (see Remark 3.2 below).

In the following, we introduce a novel estimator and bound the private rate of convergence without assuming the SDA. This disproves the conjecture of (Berrett et al., 2021) that the rate of convergence of any estimate can be arbitrarily slow when the SDA does not hold.

The idea for the general estimator is to include a further modification that, in some sense, enforces condition (3.2) to hold (see Remark 3.3 below for more details). We again depart from the privatized data (2.2) and (2.3).

This already guarantees LDP, because no further data that depend on the raw data are used in what follows.

In order to define our novel estimator, let λ_n denote the uniform distribution on $A_n := \bigcup_{j=1}^{N_n} A_{h,j}$, that is, for any Borel set A ,

$$\lambda_n(A) = \frac{\lambda(A \cap A_n)}{\lambda(A_n)},$$

where λ denotes the Lebesgue measure. We now define our final estimator $[\widehat{m}_n]_T$ by

$$[\widehat{m}_n]_T(x) := (-T) \vee (\widehat{m}_n(x) \wedge T),$$

where

$$\widehat{m}_n(x) = \frac{\widetilde{\nu}_n(A_{h,j})}{\widehat{\mu}_n(A_{h,j})} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t\}} \mathbf{1}_{\{j \leq N_n\}} \quad \text{when } x \in A_{h,j},$$

and

$$\begin{aligned} \widehat{\mu}_n(A_{h,j}) &= \frac{3}{4} \widetilde{\mu}_n(A_{h,j}) + \frac{1}{4} \lambda_n(A_{h,j}) \\ &= \left[\frac{3}{4} \widetilde{\mu}_n(A_{h,j}) + \frac{1}{4N_n} \right] \mathbf{1}_{\{j \leq N_n\}}. \end{aligned}$$

The following result states upper risk bounds for the estimators \widehat{m}_n and $[\widehat{m}_n]_T$. Its proof is deferred to Section 5.1.

Theorem 3.1. *Assume that m satisfies (1.2) with $0 < \beta \leq 1$, and that both X and Y are bounded, such that $|X| \leq r$ and $|Y| \leq T$. Consider the*

estimator \widehat{m}_n with $t = t_n = \lambda_n(A_{h,1})/8 = 1/(8N_n)$. Then,

$$\mathbf{E} \left[\int (m(x) - \widehat{m}_n(x))^2 \mu(dx) \right] \lesssim \frac{1}{nh_n^d} + h_n^{2\beta} + \frac{\sigma_Z^2}{nt_n^2} + \frac{\sigma_W^2}{nt_n^2} + \exp\left(-\frac{8nt_n^2}{9}\right). \quad (3.3)$$

As a consequence, taking $h_n \asymp \max\{(1/(n\alpha^2))^{1/(2\beta+2d)}, (1/n)^{1/(2\beta+d)}\}$ and σ_W and σ_Z as in (2.4), yields

$$\mathbf{E} \left[\int (m(x) - [\widehat{m}_n]_T(x))^2 \mu(dx) \right] \lesssim (n^{-\frac{2\beta}{2\beta+d}} \vee (n\alpha^2)^{-\frac{\beta}{\beta+d}}) \wedge 1, \quad (3.4)$$

where the numerical constant hidden in the \lesssim notation depends on r, T , and d .

Remark 3.2. The estimate \widetilde{m}_n defined by (3.1) cannot achieve the rate of convergence in (3.4) without assuming the SDA. In order to see this, consider the case of no privatization and constant, noiseless observations; that is, $Y = C$ a.s. for a constant $C \neq 0$. Then, the estimator \widetilde{m}_n satisfies

$$\begin{aligned} & \int \mathbf{E}[(m(x) - \widetilde{m}_n(x))^2] \mu(dx) \\ &= C^2 \sum_{j=1}^{N_n} \mathbf{E}[\mathbf{1}_{\{\mu_n(A_{h,j}) < t_n\}}] \mu(A_{h,j}) \\ &\geq C^2 \sum_{j=1}^{N_n} \mathbf{P}(\mu_n(A_{h,j}) < t_n) \mathbf{1}_{\{\mu(A_{h,j}) < 2t_n\}} \mu(A_{h,j}) \\ &\geq C^2 \sum_{j=1}^{N_n} \mathbf{P}(\mu_n(A) - \mu(A) < -t_n) \mathbf{1}_{\{\mu(A_{h,j}) < 2t_n\}} \mu(A_{h,j}). \end{aligned}$$

Applying of Hoeffding's inequality or a normal approximation easily shows that the probability in the last line is bounded from below by $1/2$, for n sufficiently large. Consider now the case $d = 1$, and assume that the distribution μ has a density f on $[0, 1]$ satisfying $f(x) = x$, for $0 \leq x \leq \delta$, with some $\delta > 0$. Let the regular partition be given by $[0, h_n), [h_n, 2h_n), \dots$. For n sufficiently large, we have $\mathbf{1}_{\{\mu(A_{h_n,j}) < 2t_n\}} = 1$ for those cells contained in $[0, \delta]$ where $xh_n < 2t_n$.

Consequently, using the choice $t_n \asymp h_n/\sqrt{\log n}$ suggested in (Berrett et al., 2021), we obtain the lower bound

$$\begin{aligned} \int \mathbf{E}[(m(x) - \tilde{m}_n(x))^2] \mu(dx) &\geq \frac{C^2}{2} \int_0^{\lfloor \delta/h_n \rfloor \cdot h_n} \mathbf{1}_{\{xh_n < 2t_n\}} x dx \\ &= \frac{C^2}{2} \int_0^{\lfloor \delta/h_n \rfloor \cdot h_n} \mathbf{1}_{\{x < 2/\sqrt{\log n}\}} x dx \\ &= \frac{C^2}{2} \int_0^{2/\sqrt{\log n}} x dx \\ &\geq \frac{C^2}{\log n}, \end{aligned}$$

which is a much slower rate than the one in (3.4).

Remark 3.3. The rationale behind the construction of the estimator \hat{m}_n comes from the auxiliary model obtained by replacing the raw data \mathcal{D}_n with

$$\mathcal{D}'_n = \{(X'_1, Y'_1), \dots, (X'_n, Y'_n)\},$$

where with probability $3/4$, one has $(X'_i, Y'_i) = (X_i, Y_i)$, and with probabil-

ity $1/4$, one has $X'_i \sim \lambda_n$ and $Y'_i = 0$ independently for each $i = 1, \dots, n$. In this mixture model, condition (3.2) holds by construction. Recently, in the context of density estimation under LDP the derivation of the optimal convergence rates has been reduced to such a mixture model (see the proof of Proposition 7 in (Sart, 2022)). The definition of our estimator \hat{m}_n is motivated by this approach, but our definition does not rely on any additional randomization, and replacing $\tilde{\mu}$ with $\hat{\mu}$ in the definition of \hat{m}_n may be interpreted as some kind of regularization.

4. Lower bound

In order to prove a lower bound, we restrict ourselves to a specific instance of the general regression model (1). This submodel is chosen to be sufficiently complex to rule out inference with an essentially faster rate than that obtained in Theorem 3.1. More precisely, we consider the regression model with a generic observation $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$ obeying the model

$$Y = m(X) + \eta, \quad (4.1)$$

where X is distributed uniformly on $[0, 1]^d$, the noise η is distributed uniformly on $[-1/2, 1/2]$, and the regression function m belongs to the Hölder class $\mathcal{F}(\beta, C, M)$, defined as the set of functions satisfying (1.2), with support contained in $[0, 1]^d$, and bounded from above by some constant $M > 0$.

For the lower bound, we allow the potential privacy mechanism to belong to the class \mathcal{Q}_α of *sequentially interactive* privacy mechanisms that generalizes the class of non-interactive mechanisms introduced in Section 2. More precisely, given any ordering of the raw data X_i , privatized data Z_i are generated according to conditional probability measures $Q_i(\cdot | (X_i, Y_i) = (x_i, y_i), Z_1 = z_1, \dots, Z_{i-1} = z_{i-1})$. That is, the value Z_i depends on X_i , and on privatized data created previously by other data holders. In this more general case, condition (2.1) is replaced with

$$\frac{Q_i(A | (X_i, Y_i) = (x, y), Z_{i-1} = z_{i-1}, \dots, Z_1 = z_1)}{Q_i(A | (X_i, Y_i) = (x', y'), Z_{i-1} = z_{i-1}, \dots, Z_1 = z_1)} \leq \exp(\alpha),$$

which must be satisfied for any $A \in \mathcal{Z}_i$, $z_j \in \mathcal{Z}_j$, for $j = 1, \dots, i - 1$, and all potential values $(x, y), (x', y')$ of the raw data.

For this general setup, we can prove the following lower bound result.

Its proof, based on Assouad's lemma, is given in Section 5.2.

Theorem 4.1. *The following holds:*

$$\inf_{\tilde{m} \in \mathcal{Q}_\alpha} \sup_{m \in \mathcal{F}(\beta, C, M)} \mathbf{E} \left[\int (\tilde{m}(x) - m(x))^2 dx \right] \gtrsim (n(e^\alpha - 1)^2)^{-\beta/(\beta+d)} \wedge 1,$$

where the supremum is taken over all admissible regression functions from the Hölder class $\mathcal{F}(\beta, C, M)$, and the infimum is taken over all estimators \tilde{m} based on a private sample of size n of raw data from model (4.1), as well as all potentially sequentially interactive privacy mechanisms $Q \in \mathcal{Q}_\alpha$.

Combining Theorem 4.1 with the fact that the convergence rate under LDP cannot be faster than the nonprivate rate $n^{-2\beta/(2\beta+d)}$, the rate of convergence derived in Theorem 3.1 is essentially optimal.

Remark 4.2. The lower bound is established for the specific design measure μ given by the uniform distribution on $[0, 1]^d$. The proof can be adapted easily to design measures μ with support $[0, 1]^d$ and Lebesgue density bounded away from zero.

Remark 4.3. Similarly, the theorem is established for the special error distribution given by the uniform distribution on $[-1/2, 1/2]$. This choice permits explicit calculations of certain total variation distances needed in the proof of the lower bound. It is an interesting open question whether the same lower bound for private estimation holds already when the response variable in the raw data is noise-free, that is, $Y_i = m(X_i)$, for $i = 1, \dots, n$, because this would rule out the possibility of error distributions with a faster convergence rate. Note that our proof of Theorem 4.1 does not apply in this case.

5. Proofs

5.1 Proof of Theorem 3.1 (Upper bound)

Loosely speaking, the proof of the upper bound decomposes the overall risk into three terms that are bounded separately. The first term (estimate (5.1)) captures the privatization of the response variable in (2.2), and its upper bound contains the noise level σ_Z . The second term (estimate (5.3)) yields the classical bound that holds already for the raw data without privatization, and the third and last term (estimate (5.4)) is the contribution of the privatization of the covariate values in (2.3), which contains the noise level σ_W .

We start the proof with the decomposition

$$\widehat{m}_n = \widehat{m}'_1 + \widehat{m}'_2,$$

where, for $x \in A_{h,j}$, we set

$$\widehat{m}'_1(x) = \frac{\frac{\sigma_Z}{n} \sum_{i=1}^n \varepsilon_{ij}}{\widehat{\mu}_n(A_{h,j})} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \mathbf{1}_{\{j \leq N_n\}}$$

and

$$\widehat{m}'_2(x) = \frac{\nu_n(A_{h,j})}{\widehat{\mu}_n(A_{h,j})} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \mathbf{1}_{\{j \leq N_n\}}$$

(recall that we assume that $|Y| \leq T$, which implies that $[Y_i]_T = Y_i$, for all

5.1 Proof of Theorem 3.1 (Upper bound)

$i = 1, \dots, n$). Using this notation, proving (3.3) reduces to showing that

$$\int \mathbf{E}[(\widehat{m}'_1(x))^2] \mu(dx) \leq \frac{\sigma_Z^2}{nt_n^2} \quad (5.1)$$

and

$$\begin{aligned} \int \mathbf{E}[(m(x) - \widehat{m}'_2(x))^2] \mu(dx) &\lesssim \frac{1}{nh_n^d} + h_n^{2\beta} \\ &+ \frac{\sigma_W^2 \vee \sigma_Z^2}{nt_n^2} + \exp\left(-\frac{8nt_n^2}{9}\right). \end{aligned} \quad (5.2)$$

Proof of (5.1): We have

$$\begin{aligned} \int \mathbf{E}[(\widehat{m}'_1(x))^2] \mu(dx) &= \sum_{j=1}^{N_n} \mathbf{E} \left[\frac{(\frac{\sigma_Z}{n} \sum_{i=1}^n \varepsilon_{ij})^2}{(\widehat{\mu}_n(A_{h,j}))^2} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \right] \mu(A_{h,j}) \\ &\leq \frac{\sigma_Z^2}{nt_n^2} \sum_{j=1}^{N_n} \mu(A_{h,j}) \\ &\leq \frac{\sigma_Z^2}{nt_n^2}. \end{aligned}$$

Proof of (5.2): Let m'_n be the modification of m_n , where μ_n is replaced with $\mu'_n = \frac{3}{4}\mu_n + \frac{1}{4}\lambda_n$. Then, similarly to the proof of Theorem 4.3 in (Györfi et al., 2002) we can show that

$$\mathbf{E} \left[\int (m(x) - m'_n(x))^2 \mu(dx) \right] \leq \frac{C_1 T^2}{nh_n^d} + C_2 h_n^{2\beta}, \quad (5.3)$$

where $C_1 = C_1(d, \mu)$ (more precisely, this constant depends on the measure μ only through its support) and $C_2 = C_2(d)$. In order to show (5.2), it is sufficient to show that

$$\mathbf{E} \left[\int (\widehat{m}'_2(x) - m'_n(x))^2 \mu(dx) \right] \lesssim \exp\left(-\frac{8nt_n^2}{9}\right) + \frac{\sigma_W^2}{nt_n^2}. \quad (5.4)$$

5.1 Proof of Theorem 3.1 (Upper bound)

In order to prove this bound, note that

$$\int (\widehat{m}'_2(x) - m'_n(x))^2 \mu(dx) \leq J_n + \frac{16T^2}{9} \mu(\mathbb{R}^d \setminus A_n), \quad (5.5)$$

where

$$J_n = \sum_{j=1}^{N_n} (\nu_n(A_{h,j}))^2 \left(\frac{1}{\mu'_n(A_{h,j})} - \frac{1}{\widehat{\mu}_n(A_{h,j})} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \right)^2 \mu(A_{h,j}).$$

Because $|X| \leq r$, the support of μ is contained in A_n , and consequently the second term on the right-hand side of (5.5) vanishes. Therefore, it is sufficient to find a bound for $\mathbf{E}[J_n]$. Using that $3\mu_n(A_{h,j})/4 \leq \mu'_n(A_{h,j})$,

$$\begin{aligned} J_n &\leq T^2 \sum_{j=1}^{N_n} (\mu_n(A_{h,j}))^2 \left(\frac{1}{\mu'_n(A_{h,j})} - \frac{1}{\widehat{\mu}_n(A_{h,j})} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \right)^2 \mu(A_{h,j}) \\ &\leq \frac{16T^2}{9} \sum_{j=1}^{N_n} \left(1 - \frac{\mu'_n(A_{h,j})}{\widehat{\mu}_n(A_{h,j})} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \right)^2 \mu(A_{h,j}) \\ &= \frac{16T^2}{9} \sum_{j=1}^{N_n} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) < t_n\}} \mu(A_{h,j}) \\ &\quad + \frac{16T^2}{9} \sum_{j=1}^{N_n} \left(1 - \frac{\mu'_n(A_{h,j})}{\widehat{\mu}_n(A_{h,j})} \right)^2 \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \mu(A_{h,j}). \end{aligned}$$

Therefore,

$$\mathbf{E}[J_n] \leq \mathbf{E}[J_{n,1}] + \mathbf{E}[J_{n,2}], \quad (5.6)$$

where, setting $\mu' = \frac{3}{4}\mu + \frac{1}{4}\lambda_n$,

$$\begin{aligned} J_{n,1} &= \frac{64T^2}{27} \sum_{j=1}^{N_n} \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) < t_n\}} \mu'(A_{h,j}), \quad \text{and} \\ J_{n,2} &= \frac{16T^2}{9} \sum_{j=1}^{N_n} \left(1 - \frac{\mu'_n(A_{h,j})}{\widehat{\mu}_n(A_{h,j})} \right)^2 \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \mu(A_{h,j}). \end{aligned}$$

5.1 Proof of Theorem 3.1 (Upper bound)

In order to deal with the expectation of $J_{n,1}$, note that $\mu'(A_{h,j}) \geq \lambda_n(A_{h,j})/4 = 1/(4N_n) \geq 2t_n$ holds for our choice of t_n , by definition. Thus,

$$\begin{aligned} \mathbf{P}(\widehat{\mu}_n(A_{h,j}) < t_n) &= \mathbf{P}(\widehat{\mu}_n(A_{h,j}) < t_n, \mu'(A_{h,j}) \geq 2t_n) \\ &\leq \mathbf{P}(\mu'(A_{h,j}) - \widehat{\mu}_n(A_{h,j}) \geq t_n) \\ &\leq P_{1,j} + P_{2,j}, \end{aligned}$$

where

$$\begin{aligned} P_{1,j} &= \mathbf{P}\left(\mu(A_{h,j}) - \mu_n(A_{h,j}) \geq \frac{2t_n}{3}\right), \\ P_{2,j} &= \mathbf{P}\left(\frac{\sigma_W}{n} \sum_{i=1}^n \zeta_{ij} \geq \frac{2t_n}{3}\right). \end{aligned}$$

Applying Hoeffding's inequality in the formulation taken from (Boucheron et al., 2013), Theorem 2.8, yields

$$P_{1,j} \leq \exp\left(-\frac{8nt_n^2}{9}\right),$$

whereas Chebyshev's inequality implies that

$$P_{2,j} \leq \frac{9\sigma_W^2}{4nt_n^2}.$$

Hence,

$$\mathbf{E}[J_{n,1}] \leq \frac{64T^2}{27} \exp\left(-\frac{8nt_n^2}{9}\right) + \frac{16T^2\sigma_W^2}{3nt_n^2}.$$

5.2 Proof of Theorem 4.1 (Lower bound)

Furthermore, we have

$$\begin{aligned} \mathbf{E}[J_{n,2}] &= \frac{16T^2}{9} \sum_{j=1}^{N_n} \mathbf{E} \left\{ \left(\frac{\widehat{\mu}_n(A_{h,j}) - \mu'_n(A_{h,j})}{\widehat{\mu}_n(A_{h,j})} \right)^2 \mathbf{1}_{\{\widehat{\mu}_n(A_{h,j}) \geq t_n\}} \right\} \mu(A_{h,j}) \\ &\leq \frac{16T^2}{9t_n^2} \sum_{j=1}^{N_n} \mathbf{E}[(\widehat{\mu}_n(A_{h,j}) - \mu'_n(A_{h,j}))^2] \mu(A_{h,j}) \\ &\leq \frac{16T^2 \sigma_W^2}{9nt_n^2}. \end{aligned}$$

Putting the bounds obtained for $\mathbf{E}[J_{n,1}]$ and $\mathbf{E}[J_{n,2}]$ into (5.6) yields (5.4), which proves (3.3). In addition, (3.4) follows from (3.3) by considering that

$$t_n \asymp \lambda_n(A_{h,1}) = h_n^d / \lambda(A_n) \asymp h_n^d / r_n^d,$$

and that the term of order $\exp(-8nt_n^2/9)$ in (3.3) is negligible. The truncation in the definition of $[\widehat{m}_n]_T$ guarantees that the risk is at least bounded by a constant depending on T , r , and the dimension d .

5.2 Proof of Theorem 4.1 (Lower bound)

The overall strategy to establish the stated private lower bound is similar to that for the classical lower bound that holds for estimators defined in terms of the raw data. Indeed, we borrow the reduction to the pairwise comparison of certain hypotheses parameterized using the corners of a high-dimensional hypercube, and the construction of these hypotheses, from Chapter 2.6.1 in (Tsybakov, 2009). In order to apply Assouad's lemma under privacy

5.2 Proof of Theorem 4.1 (Lower bound)

constraints, one has to use a suitable bound for the Kullback–Leibler divergence of the privatized data under different hypothetical regression functions. Such a bound is derived by combining the information theoretical inequality (5.7) with a bound for the (squared) total variation distance of the raw data. Inequality (5.7) is essential for proving lower bounds in the density estimation problem (Butucea et al., 2020; Györfi and Kroll, 2023).

We start by introducing some notation. Let $K_0: \mathbb{R} \rightarrow [0, \infty)$ be a C^∞ -function, such that (1.2) is satisfied with constant equal to 1, $\|K_0\|_\infty \leq 1$, and $\text{supp}(K_0) \subseteq [0, 1]$. For $x = (x_1, \dots, x_d) \in [0, 1]^d$, define the function $K: [0, 1]^d \rightarrow \mathbb{R}$ as $K(x) = \min_{i=1, \dots, d} K_0(x_i)$. We restrict the complexity of the overall problem by restricting ourselves to a finite set of hypotheses parameterized by $\theta \in \Theta := \{0, 1\}^{k^d}$, for some positive integer k , that will be specified below. Set $c = C \wedge M$. For any $j = (j_1, \dots, j_d) \in \{0, \dots, k-1\}^d$, define the function K_j by

$$K_j(x) = ck^{-\beta} K(kx_1 - j_1, \dots, kx_d - j_d).$$

It is readily checked that $\text{supp}(K_j) \subseteq B_j := \times_{i=1}^d [j_i/k, (j_i + 1)/k]$. For any $\theta \in \Theta$, we consider the candidate regression function

$$m_\theta = \sum_j \theta_j K_j,$$

5.2 Proof of Theorem 4.1 (Lower bound)

where the sum is taken over all multi-indices $j \in \{0, \dots, k-1\}^d$. By construction, m_θ belongs to $\mathcal{F}(\beta, C, M)$, for any $\theta \in \Theta$. Let us now assume that the raw data have been privatized by means of an arbitrary privacy mechanism $Q \in \mathcal{Q}_\alpha$, and let \tilde{m} be any estimator defined in terms of the outcome Z of Q . We denote \mathbf{P}_θ as the distribution of the tuple (X_1, Y_1) , and $Q\mathbf{P}_\theta^n$ as the distribution of $Z = (Z_1, \dots, Z_n)$ when the true regression function is m_θ . We also write \mathbf{E}_θ for the expectation operator in this case.

After these preliminaries, we start the proof with the observation that for any $\theta \in \Theta$,

$$\begin{aligned} \mathbf{E}_\theta \left[\int_{[0,1]^d} (\tilde{m}(x) - m_\theta(x))^2 dx \right] &= \sum_j \mathbf{E}_\theta \left[\int_{B_j} (\tilde{m}(x) - m_\theta(x))^2 dx \right] \\ &= \sum_j \mathbf{E}_\theta [\rho_j^2(\tilde{m}, \theta_j)], \end{aligned}$$

where

$$\rho_j(\tilde{m}, \theta_j) = \left(\int_{B_j} (\tilde{m}(x) - \theta_j K_j(x))^2 dx \right)^{1/2}.$$

Putting $\hat{\theta}_j = \arg \min_{t \in \{0,1\}} \rho_j(\tilde{m}, t)$, we have

$$\rho_j(\tilde{m}, \theta_j) \geq \frac{\|K_j\|_2}{2} \cdot |\hat{\theta}_j - \theta_j|.$$

Hence, using that $\|K_j\|_2^2 = c^2 k^{-2\beta-d} \|K\|_2^2$, we obtain

$$\mathbf{E}_\theta \left[\int_{[0,1]^d} (\tilde{m}(x) - m_\theta(x))^2 dx \right] \geq \frac{c^2 \|K\|_2^2}{4} k^{-2\beta-d} \mathbf{E}_\theta [\rho(\hat{\theta}, \theta)],$$

5.2 Proof of Theorem 4.1 (Lower bound)

where $\widehat{\theta} = (\widehat{\theta}_j)$ and $\rho(\theta, \theta') = \sum_j \mathbf{1}_{\{\theta_j \neq \theta'_j\}}$ denotes the Hamming distance between θ and θ' . Consequently,

$$\begin{aligned} & \sup_{m \in \mathcal{F}(\beta, C, M)} \mathbf{E} \left[\int_{[0,1]^d} (\widetilde{m}(x) - m(x))^2 dx \right] \\ & \geq \sup_{\theta \in \Theta} \mathbf{E}_{\theta} \left[\int_{[0,1]^d} (\widetilde{m}(x) - m_{\theta}(x))^2 dx \right] \\ & \geq \frac{c^2 \|K\|_2^2}{4} k^{-2\beta-d} \inf_{\widehat{\theta}} \sup_{\theta \in \Theta} \mathbf{E}_{\theta} [\rho(\widehat{\theta}, \theta)]. \end{aligned}$$

In order to bound the quantity $\inf_{\widehat{\theta}} \sup_{\theta \in \Theta} \mathbf{E}_{\theta} [\rho(\widehat{\theta}, \theta)]$, we use Statement (iv) of Theorem 2.12 in (Tsybakov, 2009), which relies on a finite bound on the Kullback–Leibler distance $K(Q\mathbf{P}_{\theta}^n, Q\mathbf{P}_{\theta'}^n)$ for θ, θ' such that $\rho(\theta, \theta') = 1$. In order to obtain such a bound, first note that Equation (14) in (Duchi et al., 2018) yields

$$K(Q\mathbf{P}_{\theta}^n, Q\mathbf{P}_{\theta'}^n) \leq 4n(e^{\alpha} - 1)^2 V^2(\mathbf{P}_{\theta}, \mathbf{P}_{\theta'}), \quad (5.7)$$

where $V(\mathbf{P}, \mathbf{Q})$ denotes the total variation distance between two probability measures. Thus, it remains to find a bound for $V(\mathbf{P}_{\theta}, \mathbf{P}_{\theta'})$. In order to bound this quantity, note that under model (4.1), the vector (X, Y) has a Lebesgue density φ_m equal to one on the set

$$\{(x, y) \in \mathbb{R}^{d+1} : x \in [0, 1]^d, y \in [m(x) - 1/2, m(x) + 1/2]\},$$

and equal to zero otherwise. Now, let θ, θ' be such that $\rho(\theta, \theta') = 1$. Then,

5.2 Proof of Theorem 4.1 (Lower bound)

a direct calculation using Scheffé's Theorem yields

$$\begin{aligned} V(\mathbf{P}_\theta, \mathbf{P}_{\theta'}) &= \frac{1}{2} \int_{\mathbb{R}^{d+1}} |\varphi_{m_\theta}(x, y) - \varphi_{m_{\theta'}}(x, y)| dx dy \\ &\leq k^{-d} \|K_j\|_\infty \\ &= ck^{-d-\beta} \|K\|_\infty. \end{aligned}$$

Combining this bound with (5.7) for $k \asymp (n(e^\alpha - 1))^{1/(2\beta+2d)} \vee 1$ yields

$$K(Q\mathbf{P}_\theta^n, Q\mathbf{P}_{\theta'}^n) \lesssim 1.$$

By applying Theorem 2.12, Statement (iv), of (Tsybakov, 2009), we obtain

$$\sup_{m \in \mathcal{F}(\beta, C, M)} \mathbf{E} \left[\int_{[0,1]^d} (\tilde{m}(x) - m(x))^2 dx \right] \gtrsim (n(e^\alpha - 1)^2)^{-\beta/(\beta+d)} \wedge 1,$$

which implies the claim, because $Q \in \mathcal{Q}_\alpha$ and \tilde{m} are arbitrary.

Acknowledgements

The research of Martin Kroll was supported by the German Research Foundation (DFG) under the grant DFG DE 502/27-1. The research of László Györfi was supported by the National Research, Development and Innovation Fund of Hungary under the 2019-1.1.1-PIACI-KFI-2019-00018 funding scheme.

REFERENCES

References

- Berrett, T. B. and C. Butucea (2019). Classification under local differential privacy. *Pub. Inst. Stat. Univ. Paris, Special issue in honour of Denis Bosq's 80th birthday 63*, 191–205.
- Berrett, T. B., L. Györfi, and H. Walk (2021). Strongly universally consistent nonparametric regression and classification with privatised data. *Electron. J. Stat.* 15(1), 2430–2453.
- Boucheron, S., G. Lugosi, and P. Massart (2013). *Concentration inequalities*. Oxford University Press, Oxford. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.
- Butucea, C., A. Dubois, M. Kroll, and A. Saumard (2020). Local differential privacy: elbow effect in optimal density estimation and adaptation over Besov ellipsoids. *Bernoulli* 26(3), 1727–1764.
- Duchi, J. C., M. I. Jordan, and M. J. Wainwright (2018). Minimax optimal procedures for locally private estimation. *J. Amer. Statist. Assoc.* 113(521), 182–201.
- Györfi, L., M. Kohler, A. Krzyżak, and H. Walk (2002). *A distribution-free theory of nonparametric regression*. Springer Series in Statistics. Springer-Verlag, New York.
- Györfi, L. and M. Kroll (2023, jan). Multivariate density estimation from privatised data: universal consistency and minimax rates. *Journal of Nonparametric Statistics*, 1–23.
- Sart, M. (2022). Density estimation under local differential privacy and hellinger loss. Preprint. Available at <https://hal.archives-ouvertes.fr/hal-03471018/>.

REFERENCES

Tsybakov, A. B. (2009). *Introduction to nonparametric estimation*. Springer Series in Statistics.

Springer, New York.

Department of Computer Science and Information Theory, Budapest University of Technology
and Economics, Magyar Tudósok krt. 2., Budapest, H-1117, Hungary.

E-mail: gyorfi@cs.bme.hu

Ruhr-Universität Bochum, Fakultät für Mathematik, Universitätsstraße 150. D-44801 Bochum.

Universität Bayreuth, Fakultät für Mathematik, Physik und Informatik, D-95440 Bayreuth.

E-mail: martin.kroll@uni-bayreuth.de