On the Private Estimation of Smooth Transport Maps

Clément Lalanne [1] **Franck Iutzeler** [1] **Jean-Michel Loubes** [2 1] **Julien Chhor** [3]

Abstract

Estimating optimal transport maps between two distributions from respective samples is an important element for many machine learning methods. To do so, rather than extending discrete transport maps, it has been shown that estimating the Brenier potential of the transport problem and obtaining a transport map through its gradient is near minimax optimal for smooth problems. In this paper, we investigate the private estimation of such potentials and transport maps with respect to the distribution samples. We propose a differentially private transport map estimator achieving an L^2 error of at most $n^{-1} \vee n^{-\frac{2\alpha}{2\alpha-2+d}} \vee (n\epsilon)^{-\frac{2\alpha}{2\alpha+d}}$ up to poly-logarithmic terms where n is the sample size, ϵ is the desired level of privacy, α is the smoothness of the true transport map, and d is the dimension of the feature space. We also provide a lower bound for the problem.

1. Introduction

Given two probability measures P, Q on $\mathbb{R}^d$, the question of "how to optimally move the mass" from P to Q, i.e. to find a map $T_0 : \mathbb{R}^d \to \mathbb{R}^d$ that solves the Monge (Monge, 1781) problem

$$T_0 \in \operatorname{argmin} \int_{\mathbb{R}^d} \|T(x) - x\|^2 dP(x) \quad \text{s.t.} \quad T_\# P = Q \tag{1}$$

where $T_\# P$ denotes the *push-forward* of P by T is of high interest in both theoretical and more numerical branches of mathematics. The problem is referred to as an *Optimal Transport Problem*, and the optimal mapping T_0 is referred to as the *Optimal Transport Map*; we refer the reader to the textbooks (Villani, 2003; Villani et al., 2009; Santambrogio, 2016; Peyré & Cuturi, 2019) for a general introduction to optimal transport and its computational aspects.

Because of its utility in measuring geometrical discrepancies between measures, and because of recent algorithmic developments (Cuturi, 2013; Altschuler et al., 2017; Dvurechensky et al., 2018), it has become a standard tool in Computer Science (Feydy et al., 2017; Lavenant et al., 2018; Solomon et al., 2015; 2016), Machine Learning (Alaux et al., 2018; Alvarez-Melis et al., 2018; Arjovsky et al., 2017; Cañas & Rosasco, 2012; Gordaliza et al., 2019; Flamary et al., 2018; Genevay et al., 2018; Grave et al., 2019; Janati et al., 2019; Montavon et al., 2016; Schmitz et al., 2018; Staib et al., 2017; Le et al., 2024) and Statistics (del Barrio et al., 2024; Del Barrio et al., 2024; Cazelles et al., 2018; del Barrio et al., 2019; Klatt et al., 2020; Kroshnin et al., 2019; Panaretos & Zemel, 2019; Ramdas et al., 2017; Rigollet & Weed, 2018; Seguy & Cuturi, 2015; Tameling & Munk, 2018; Weed & Berthet, 2019; Zemel & Panaretos, 2019).

When dealing with real-world data, the true distributions P and Q are often accessible only through samples (Courty et al., 2017a; 2014; 2017b; Damodaran et al., 2018; Forrow et al., 2019; Perrot et al., 2016; Seguy et al., 2018; Hütter & Rigollet, 2021). In this article, we suppose that we have access to

$$X_1, \dots, X_n \overset{\text{i.i.d.}}{\sim} P \quad \text{and} \quad Y_1, \dots, Y_n \overset{\text{i.i.d.}}{\sim} Q$$

such that the X_is and the Y_is are mutually independent.[1] In this scenario, the problem (1) cannot be solved directly to obtain an optimal transport map; instead, it must be estimated using the available samples. To do so, a crude approach would be to replace P and Q with their empirical counterparts. This approach has two main drawbacks: first, it does not specify how to move points across the entire support of P; second, it is affected by the curse of dimensionality (Niles-Weed & Rigollet, 2022). To resolve these issues, (Hütter & Rigollet, 2021) proposed incorporating smoothness and regularity assumptions into the optimal transport map T_0 and leveraging functional estimators.

[1]Institut de Mathématiques de Toulouse, UMR5219, Université de Toulouse, CNRS, UPS, F-31062 Toulouse Cedex 9, France [2]INRIA, France [3]Toulouse School of Economics, Université Toulouse Capitole, France. Correspondence to: Clément Lalanne <clement.lalanne@math.univ-toulouse.fr>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

[1]For simplicity, we take the same number of samples for both distributions. Our results naturally extend to the case where the two sample sizes are different, at the cost of using more involved notation.

At the same time, deriving estimators from real user data raises new challenges, especially regarding privacy. It is well-documented that sharing statistics based on such data without adequate protections can lead to serious privacy leakages (Narayanan & Shmatikov, 2006; Backstrom et al., 2007; Fredrikson et al., 2015; Dinur & Nissim, 2003; Homer et al., 2008; Loukides et al., 2010; Narayanan & Shmatikov, 2008; Sweeney, 2000; Wagner & Eckhoff, 2018; Sweeney, 2002).

In the context of transport maps estimation, ensuring privacy is highly beneficial for applications involving infringement of fundamental rights as in bias analysis of machine learning algorithms. For instance, in (De Lara et al., 2024) or (Black et al., 2020), optimal transport maps enable to build counterfactual individuals to assess whether an algorithm is prone to discrimination and to mitigate its effect.

To mitigate these concerns, differential privacy (DP) (Dwork et al., 2006) has become the benchmark to ensure privacy protection. DP introduces randomness into computations, ensuring that the released statistics are determined not solely by the dataset but also by the added randomness. This limits the influence of any single data point on the outcome, thereby preserving privacy. Major organizations, including the US Census Bureau (Abowd, 2018), Google (Erlingsson et al., 2014), Apple (Thakurta et al., 2017), and Microsoft (Ding et al., 2017), have embraced this methodology.

This work investigates the problem of estimating smooth optimal transport maps from samples under differentially privacy. In particular, we prove that the optimal estimation rate is expected to degrade under specific regimes.

1.1. Contributions

The main contributions of this article can be summarized as follows:

- **A Differentially Private Estimator for Transport Maps.** We introduce a private mechanism (under so-called pure differential privacy) which is aimed at solving the problem of smooth optimal transport maps estimation (see Section 3). This estimator is then adapted to be implementable in practice in Section 6.

- **Statistical Upper Bound.** We analyze this estimator in Section 4 and provide an upper bound on its convergence rate.

- **Minimax Lower Bound.** We provide in Section 5 a minimax lower bound on the private optimal transport map problem in order to characterize the difficulty of the problem.

1.2. Related Work

Differential Privacy and Statistics/Learning. Over the past decade, there has been growing interest in estimating various quantities while ensuring differential privacy. Examples of relevant works include, but are not limited to, the following references (Wasserman & Zhou, 2010; Barber & Duchi, 2014; Diakonikolas et al., 2015; Karwa & Vadhan, 2018; Bun et al., 2019; 2021; Kamath et al., 2019; Biswas et al., 2020; Kamath et al., 2020; Acharya et al., 2021; Lalanne, 2023; Aden-Ali et al., 2021; Cai et al., 2019; Brown et al., 2021; Cai et al., 2019; Kamath et al., 2022a; Lalanne et al., 2023c;d; Lalanne & Gadat, 2024; Singhal, 2023; Kamath et al., 2023; 2022b).

More specifically, this article falls into the scope of private nonparametric statistics (Tsybakov, 2009), where the quantities of interest live in *infinite* dimensional vector spaces, necessitating the use of approximation techniques alongside estimation. The problem of non-parametric density estimation under differential privacy has been studied in various works (Wasserman & Zhou, 2010; Barber & Duchi, 2014; Lalanne et al., 2023b; Lalanne & Gadat, 2024), both in the setting of this article and in the setup of local privacy (without a trusted server) (Evfimievski et al., 2003; Kasiviswanathan et al., 2008; Duchi et al., 2013; 2016; Butucea et al., 2019; Kroll, 2021; Schluttenhofer & Johannes, 2022; Györfi & Kroll, 2023). Additional works include studies on nonparametric regression (Berrett et al., 2021; Györfi & Kroll, 2022), nonparametric tests (Lam-Weil et al., 2022), interactions between robustness and privacy (Chhor & Sentenac, 2023) and recent advances in locally private Bayesian modeling (Beraha et al., 2023).

To the best of our knowledge, our work is the first to address private estimation of smooth transport maps with statistical guarantees. A noteworthy exception is the article (Xian et al., 2024) which uses elements of optimal transport in a private way in order to obtain fairness with statistical guarantees. Their approach is discussed in the next paragraph.

Differential Privacy and Optimal Transport A line of work (Rakotomamonjy & Ralaivola, 2021), building on the ideas from (Harder et al., 2021) provides privacy guarantees for the values of the sliced Wasserstein distance and MMD, respectively.

On another note, other research has explored task-specific private methods utilizing optimal transport. For instance, (Sebag et al., 2023) employs the sliced Wasserstein distance for data generation using an approach based on gradient flows. Additionally, (Tien et al., 2019) addresses differentially private domain adaptation via optimal transport by perturbing the optimal coupling between noisy datasets.

More recently, (Xian et al., 2024) investigated fair and private regression, using tools from optimal transport theory. More precisely, they build private histogram estimators of the distributions P and Q and then compute the maps to their Wasserstein barycenter (which has applications in fairness). Their work is not directly related to our study, since the authors focus on fairness guarantees rather than on the problem of learning transport maps.

Beyond these efforts, optimal transport has also been explored within novel privacy paradigms outside the scope of our work (Pierquin et al., 2024; Kawamoto & Murakami, 2019; Yang et al., 2024).

1.3. Notation

The symbols $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{R}$ and $\mathbb{C}$ are respectively used to refer to the sets of natural numbers (including 0), relative numbers, real numbers, and complex numbers. For any $k \in \mathbb{N}$, $\mathcal{C}^k(\mathcal{S})$ denotes the set of functions from a space $\mathcal{S}$ to $\mathbb{C}$ that are k times continuously differentiable, and $\mathcal{C}^\infty(\mathcal{S})$ is defined as $\cap_{k\in\mathbb{N}}\mathcal{C}^k(\mathcal{S})$. Whenever applicable, ∇ and ∇^2 are used to refer to gradient and Hessian operators, respectively. For a subset S of a normed vector space, $|S|$ refers to $\sup_{x\in S}\|x\|$ for the inherited norm. For any set S, $\#(S)$ denotes its cardinality. For a family F of vectors, $\mathrm{Span}(F)$ refers to the smallest vector space containing F. For a closed convex set of an Hilbert space, Proj_C refers to the convex projection onto C. The Hölder norm of order α on S (see Appendix B in (Hütter & Rigollet, 2021)) is denoted by $\|T\|_{C^\alpha(S)}$. For a measure μ, $\|\cdot\|_{L^2(\mu)}$ is the usual L^2 norm with reference measure μ. With a small overlap of notation, for a measurable $S \subset \mathbb{R}^d$, $\|\cdot\|_{L^2(S)}$ refers to the L^2 norm with reference measure $\mathbb{1}_S \cdot \lambda$ where λ is Lebesgue's measure. The asymptotic regimes are considered when $n \to +\infty$ and $n\epsilon \to +\infty$.

2. Smooth Transport Maps Estimation

In this section, we cover some basic theory on optimal transport and the minimax estimation of smooth transport maps. Following (Hütter & Rigollet, 2021), we set the feature space[2] $\Omega = [0,1]^d$ and define $\tilde{\Omega} = [-1,2]^d$. The notation M (resp. R) will refer to a positive constant strictly greater than 2 (resp. $\|\text{ Identity }\|_{C^\alpha(\tilde{\Omega})}$) throughout the paper. The constants will hide terms that depend on Ω, $\tilde{\Omega}$, M, α and d, which was also the case in (Hütter & Rigollet, 2021).

2.1. Semi-Dual Problem and Smoothness

In the case where P and Q are supported in Ω and are absolutely continuous with respect to Lebesgue's measure, Brenier's theorem (Brenier, 1991) ensures that the solution to Problem (1) is a transport map $T_0 = \nabla f_0$ where f_0 is a convex function. Furthermore, f_0 can be equivalently obtained by solving the semi-dual problem

$$
\begin{aligned}
f_0 \in \operatorname{argmin} & \underbrace{\int f(x)dP(x) + \int f^*(y)dQ(y)}_{=:S(f)} \\
\text{s.t.} \quad & f \in L^1(P)
\end{aligned} \tag{2}
$$

where for all y

$$
f^*(y) = \sup_{x\in\tilde{\Omega}} \langle x, y\rangle - f(x) \tag{3}
$$

is the Fenchel-Legendre transform of f;[3] see e.g. (Rüschendorf & Rachev, 1990). In this problem, f is called the *(Brenier) potential*.

In order to estimate transport maps through Brenier potentials and obtain convergence rates, we need to enforce some regularity conditions, which we detail below. First, we require the first probability distribution to be absolutely continuous with lower and upper bounded density.

Definition 2.1 (Admissible source distributions)**.** We denote by $\mathcal{M}$ the set of probability measures P on $\mathbb{R}^d$ that are supported on Ω, that are absolutely continuous w.r.t. Lebesgue's measure and whose density ρ_P verifies $\frac{1}{M} \le \rho_P(x) \le M$ for almost all x in Ω.

Furthermore, we control the bias by imposing smoothness assumptions on the optimal transport map. The following definition introduces the class of admissible smooth transport maps.

Definition 2.2 (Admissible smooth transport maps)**.** We denote by $\mathcal{T}$ the set of differentiable mappings $T : \tilde{\Omega} \to \mathbb{R}^d$ such that $T = \nabla f$ for some differentiable convex function $f : \tilde{\Omega} \to \mathbb{R}^d$ and

- $\forall x \in \tilde{\Omega}, \quad \|T(x)\| \le M,$
- $\forall x \in \tilde{\Omega}, \frac{1}{M} \preceq \nabla^2 f(x) \preceq M,$
- $P_{\#T}(\Omega) = 1.$

Furthermore, for $\alpha > 1$ and $R > 1$, we define

$$
\mathcal{T}_\alpha(R) = \left\{ T \in \mathcal{T} : \begin{cases} T \text{ is } \lfloor \alpha \rfloor \text{ times differentiable,} \\ \text{and } \|T\|_{C^\alpha(\tilde{\Omega})} \le R \end{cases} \right\}. \tag{4}
$$

In words, admissible transport maps have to be bounded (first item), come from a smooth and strongly convex potential (second item), and have a stable support (third item). In

[2] As in (Hütter & Rigollet, 2021), Ω can be replaced by any bounded and connected Lipschitz domain. In this case, $\tilde{\Omega}$ can be replaced by an hypercube containing Ω in its interior.

[3] We restrict the sup to $\tilde{\Omega}$ because we are considering extended functions as in Appendix A in (Hütter & Rigollet, 2021).

addition, we will assume that the optimal transport map belongs to the set $\mathcal{T}_\alpha(R)$ of admissible transport maps whose Hölder norm of order α is bounded by R.

2.2. Empirical Semi-Dual

Hütter & Rigollet (2021) observed that the objective in (2) consists of the sum of an expectation over P and an expectation over Q, which can thus be approximated using Monte-Carlo averages. This led them to propose an estimator of the form

$$\hat{T}_0 := \nabla \hat{f}_0 \ , \tag{5}$$

where

$$\begin{aligned} \hat{f}_0 \in \text{argmin} \quad & \underbrace{\frac{1}{n}\sum_{i=1}^n f(X_i) + \frac{1}{n}\sum_{i=1}^n f^*(Y_i)}_{=:\hat{S}(f|X_{1:n},Y_{1:n})} \\ \text{s.t.} \quad & f \in \hat{V}_0 \end{aligned} \tag{6}$$

and where $\hat{V}_0 \subset L^2(\mathbb{R}^d)$ is a functional space to be specified later. Note that S and $\hat{S}$ are linked by the relation

$$\forall f, \quad S(f) = \mathbb{E}_{X_{1:n},Y_{1:n}}\left(\hat{S}(f|X_{1:n},Y_{1:n})\right) \ . \tag{7}$$

Controlling the deviations of $\hat{S}$ from its expectation, as well as using $S(f)-S(f_0)$ effectively as a for the transport maps suboptimality, typically requires regularity of the admissible potentials. This is formalized in the following definition mirroring Definition 2.2.

Definition 2.3 (Admissible potentials)**.** We denote by $\mathcal{X}(M)$ the set of twice continuously differentiable functions $f : \tilde{\Omega} \to \mathbb{R}$ such that $\forall x \in \tilde{\Omega}$

- $|f(x)| \le 2M^2$,
- $\|\nabla f(x)\| \le M$,
- $\frac{1}{M} \preceq \nabla^2 f(x) \preceq M$.

Next, we specify how $\hat{V}_0$ can be chosen to obtain an implementable estimator.

2.3. Wavelet Decompositions and Subspace Approximations

In line with standard approaches in nonparametric estimation, and in order to exploit the smoothness of the optimal map, Hütter & Rigollet (2021) suggested using successive approximations over nested subspaces $V_1 \subset V_2 \subset \cdots \subset V_J \subset \cdots \subset L^2(\mathbb{R}^d)$ of $L^2(\mathbb{R}^d)$, where $J \ge 1$ denotes an integer controlling the resolution of the approximation space.

Given a pair ψ_M and ψ_F of smooth enough mother and father wavelets with compact supports, one can define a wavelet Hilbert basis of $L^2(\mathbb{R})$ of the form

$$(\Psi_k^{j,g})_{j\in\mathbb{N},k\in\mathbb{Z}^d,g\in G^j} \tag{8}$$

where $\forall j$, G^j is finite. The exact construction is presented in Appendix B of (Hütter & Rigollet, 2021), with complements given in the present Appendix A. We do not include it in the main body of this article as we believe that it is not essential to understanding the core message.

Importantly, since we are interested in guarantees in $L^2(\tilde{\Omega})$ rather than $L^2(\mathbb{R}^d)$, and because the supports of the father and mother wavelets are compact, the basis can be restricted to the indices such that $k \in K_j$ for some *finite* $K_j \subset \mathbb{Z}^d$.

Thus, defining

$$V_J := \text{Span}\left((\Psi_k^{j,g})_{j\le J,k\in K_j,g\in G^j}\right) \tag{9}$$

yields a sequence of *finite-dimensional* nested approximation spaces.

The estimator proposed by Hütter & Rigollet (2021) is the solution of (6) where $\hat{V}_0$ is replaced with $V_J \cap \mathcal{X}(2M)$. We denote the solution of the resulting problem and its associated transport map by $\hat{f}_J$ and $\hat{T}_J := \nabla \hat{f}_J$, respectively.

Finally, measuring the *utility* (or error) of a candidate transport map T by its squared $L^2(P)$ distance

$$d(T,T_0)^2 := \int \|T - T_0\|^2 dP \ , \tag{10}$$

we can guarantee that such a transport map estimator has a (near) minimax optimal rate.

Theorem 2.4 (Th. 2 of (Hütter & Rigollet, 2021))**.** *For any $\alpha > 1$, the minimax rate of smooth optimal transport map estimation is*

$$\inf_{\hat{T}} \sup_{P\in\mathcal{M},T_0\in\mathcal{T}_\alpha} \mathbb{E}_{X_{1:n},Y_{1:n}}\left(d(T,\hat{T}_0)^2\right) \gtrsim \frac{1}{n} \vee n^{-\frac{2\alpha}{2\alpha-2+d}} \ . \tag{11}$$

Moreover, if $P \in \mathcal{M}$ and $T_0 \in \mathcal{T}_\alpha$, the estimator $\hat{T}_J := \nabla \hat{f}_J$ defined above achieves this rate up to polylogarithmic factors. Here, the constants also hide a dependence on R.

2.4. Towards the Analysis of a Private Estimator

In this paper, we extend this approach to the context of differential privacy. We derive a private estimator $\hat{f}_{\text{priv}}$ for the optimal Brenier potential and deduce the corresponding private optimal transport map $\hat{T}_{\text{priv}} = \nabla \hat{f}_{\text{priv}}$. To understand how privacy will affect the utility of the estimation, we propose the following risk decomposition.

Lemma 2.5. *Suppose that there exists a random variable* Priv *independent of all the other sources of randomness such that* $\hat{f}_{\text{priv}}$ *is* $(X_{1:n}, Y_{1:n}, \text{Priv})$*-measurable and there exists* $U \geq 0$ *that is* Priv*-measurable such that*

$$\underbrace{\hat{S}(\hat{f}_{\text{priv}}|X_{1:n}, Y_{1:n}) - \hat{S}(\hat{f}_J|X_{1:n}, Y_{1:n})}_{\text{Suboptimality on the semi-dual}} \leq U \tag{12}$$

almost surely. If $\hat{f}_{\text{priv}} \in V_J \cap \mathcal{X}(2M)$ *almost surely, then*

$$\begin{aligned}\mathbb{E}_{X_{1:n},Y_{1:n}} &\left(d(\hat{T}_{\text{priv}}, T_0)^2\right) \\ &\lesssim U + \frac{J2^{J(d-2)}\ln(1+Cn)}{n} + \frac{1}{n} \\ &\quad + \inf_{T \in V_J \cap \mathcal{X}(2M)} d(T, T_0)^2\end{aligned} \tag{13}$$

Priv*-almost surely for some positive constant* C*. Here,* $\mathbb{E}_{X_{1:n},Y_{1:n}}$ *denotes the expectation w.r.t. the data, i.e. conditional on* Priv.

Proof. See Appendix C.1. □

In the next section, we propose a private estimator $\hat{f}_{\text{priv}}$ and explain how to control $\hat{S}(\hat{f}_{\text{priv}}|X_{1:n}, Y_{1:n}) - \hat{S}(\hat{f}_J|X_{1:n}, Y_{1:n})$. In particular, this estimator will satisfy the hypotheses of Lemma 2.5.

3. Private Estimator

This section presents some background on differential privacy and introduces our private estimator.

3.1. Differential Privacy

Given a (randomized) mechanism M (i.e. a conditional kernel of probabilities), $\text{dom}\,(M)$ denotes its domain (i.e. the set of admissible inputs) and $\text{codom}\,(M)$ refers to its codomain (i.e. the set of admissible outputs). The set $\text{dom}\,(M)$ represents the space of all possible data sets and is equipped with a binary symmetric relation $\cdot \sim \cdot$ called *neighboring* relations. Informally, two datasets are neighboring if they contain the same data except for at most one individual. Formally, we say that $(X_{1:n}, Y_{1:n}) \sim (X'_{1:n}, Y'_{1:n})$ if $d_{\text{ham}}\left((X_{1:n}, Y_{1:n}), (X'_{1:n}, Y'_{1:n})\right) \leq 1$, where $d_{\text{ham}}\,(\cdot,\cdot)$ to denote the Hamming distance,

Definition 3.1 (Differential Privacy (Dwork et al., 2006))**.** Given $\epsilon \geq 0$, a randomized mechanism $M : \text{dom}\,(M) \to \text{codom}\,(M)$ is ϵ-differentially private (or ϵ-DP) if for all $D \sim D' \in \text{dom}\,(M)$ and all measurable $S \subset \text{codom}\,(M)$ we have

$$\mathbb{P}_M\left(M(D) \in S\right) \leq e^{\epsilon}\mathbb{P}_M\left(M(D') \in S\right) \ . \tag{14}$$

Here, $\epsilon \geq 0$ is a parameter that controls the privacy level: Lower values of ϵ correspond to more private mechanisms M.

We will rely on the standard *report noisy max* mechanism (Dwork & Roth, 2014; McSherry & Talwar, 2007; McKenna & Sheldon, 2020; Ding et al., 2021). We give its version instantiated with Laplace noise below.

Lemma 3.2 (Report Noisy Argmin with Laplace Noise)**.** *Let* $f_1, \ldots, f_N$ *be queries with sensitivities that are uniformly bounded by* Δ*, that is, for any of neighboring datasets* $D \sim D'$ *and for any* $i \in \{1, \ldots, n\}$*,*

$$|f_i(D) - f_i(D')| \leq \Delta \ . \tag{15}$$

Then, if $L_1, \ldots, L_N$ *are independent and identically distributed random variables with standard Laplace distribution, the mechanism* $\hat{\imath}$ *defined as*

$$\hat{\imath}(D) \in \operatorname{argmin}\left\{f_i(D) + \frac{2\Delta}{\epsilon}L_i\right\} \tag{16}$$

is ϵ*-DP, with the convention that we return a random index in the* argmin *when it is not unique.*

In particular, we will use this mechanism as a building block to construct a private estimator of $\hat{f}_J$.

3.2. Noisy semi-dual estimator

In this section, we construct our private estimator, building on the empirical semi-dual one presented above.

The first step is to provide an upper bound on the sensitivity of $\hat{S}(f|X_{1:n}, Y_{1:n})$, which is established in the lemma below.

Lemma 3.3 (Sensitivity of the semi-dual objective)**.** *For any* $f \in \mathcal{X}(2M)$*, it holds that*

$$\begin{aligned}\sup_{(X_{1:n},Y_{1:n})\sim(X'_{1:n},Y'_{1:n})} &\left|\hat{S}(f|X_{1:n}, Y_{1:n}) - \hat{S}(f|X'_{1:n}, Y'_{1:n})\right| \\ &\leq \frac{2\|f\|_\infty \vee 2|\tilde{\Omega}|^2}{n} \ .\end{aligned} \tag{17}$$

Proof. See Appendix D.2. □

Now, recalling Section 2.3, let $C_{J,M}$ be a *finite* subset of $V_J \cap \mathcal{X}(2M)$. We can enumerate this set as

$$C_{J,M} = \{f_1, \ldots, f_{\#(C_{J,M})}\} \ . \tag{18}$$

Considering independent random variables $L_1, \ldots, L_{\#(C_{J,M})}$ with Laplace distribution that are also independent of all the other sources of randomness, we define

$$\begin{aligned}\hat{\imath}_{\text{priv}}(D) \in \operatorname{argmin}&\left\{\hat{S}(f_i|D) + \frac{32M^2 \vee 36d}{n\epsilon}L_i\right\} \ . \\ \text{for} \quad & i \in \{1, \ldots, \#\left(C_{J,M}\right)\}\end{aligned} \tag{19}$$

Theorem 3.4 (Privacy of the noisy semi-dual estimator)**.** *The mechanism returning the pair*

$$(\hat{f}_{\text{priv}} := f_{\hat{i}_{\text{priv}}}, \hat{T}_{\text{priv}} := \nabla \hat{f}_{\text{priv}}) \tag{20}$$

is ϵ-DP.

Proof. It follows as a direct consequence of Lemma 3.2 and of Lemma 3.3 that $\hat{i}_{\text{priv}}$ is ϵ-DP. Hence, $(\hat{f}_{\text{priv}}, \hat{T}_{\text{priv}})$ is ϵ-DP by the post-processing property of differential privacy (Dwork & Roth, 2014). □

We will use the estimator defined in Equation (20) as our estimator for both the optimal Brenier potential and the optimal transport map.

A question that remains unanswered in this section is how to choose J and $C_{J,M}$. So far, we have treated these quantities as hyperparameters. The next section explains how to choose them in order to optimize the bias-variance trade-off.

4. Statistical Upper Bound

This section details the construction of $C_{J,M}$ and gives an upper bound on the utility of the proposed private estimator.

4.1. Covering construction

We now construct the set $C_{J,M}$ such that it forms a δ-*covering* of the space $V_J \cap \mathcal{X}(2M)$ with respect to the empirical criterion $\hat{S}(\cdot|X_{1:n}, Y_{1:n})$. Specifically, we aim to ensure that $C_{J,M} \subset V_J \cap \mathcal{X}(2M)$ and that

$$\begin{aligned} \forall f \in & V_J \cap \mathcal{X}(2M), \exists f_c \in C_{J,M} \text{ s.t.} \\ & |\hat{S}(f|X_{1:n}, Y_{1:n}) - \hat{S}(f_c|X_{1:n}, Y_{1:n})| \leq \delta \ . \end{aligned} \tag{21}$$

A first step towards this goal is to observe that $\hat{S}(\cdot|X_{1:n}, Y_{1:n})$ is Lipschitz-continuous with respect to the L^∞ norm, as shown in the following lemma.

Lemma 4.1. *For any $f_1, f_2 \in \mathcal{X}(2M)$,*

$$\left|\hat{S}(f_1|X_{1:n}, Y_{1:n}) - \hat{S}(f_2|X_{1:n}, Y_{1:n})\right| \leq 2\|f_1 - f_2\|_\infty \tag{22}$$

for any dataset $(X_{1:n}, Y_{1:n})$.

Proof. See Appendix E.1 □

Thus, a $\delta > 0$-covering of $V_J \cap \mathcal{X}(2M)$ for $\hat{S}(\cdot|X_{1:n}, Y_{1:n})$ can be built from a $\frac{\delta}{2}$-covering of $V_J \cap \mathcal{X}(2M)$ for $\|\cdot\|_\infty$. The following lemma bounds the cardinality of a specific covering of this form.

Lemma 4.2. *There exists a δ-covering of $V_J \cap \mathcal{X}(2M)$ for $\|\cdot\|_\infty$ of cardinality at most N where*

$$\ln(N) \lesssim 2^{Jd} \ln\left(\frac{C 2^{Jd/2}}{\delta} + 1\right) \tag{23}$$

for some constant C that does not depend on J or δ

Proof. See Appendix E.2 □

4.2. Main result

We now have all the pieces to bound the error of $\hat{T}_{\text{priv}}$. Indeed, we can apply Lemma 2.5 with an lower bound on the empirical semi-dual error (U) that depends on (i) the resolution δ of the covering, and (ii) the uncertainty created by privacy $\max\left\{|L_1|, \ldots, |L_{\#(C_{J,M})}|\right\}$. The remaining approximation bias $\inf_{T \in V_J \cap \mathcal{X}(2M)} d(T, T_0)^2$ is controlled as in (Hütter & Rigollet, 2021).

Theorem 4.3. *Let $P \in \mathcal{M}, T_0 \in \mathcal{T}_\alpha$. By choosing δ appropriately, the estimator $\hat{T}_{\text{priv}}$ has a utility satisfying*

$$\begin{aligned} & \mathbb{E}\left(\int \|T_0 - \hat{T}_{\text{priv}}\|^2 dP\right) \\ & \lesssim R^2 2^{-2J\alpha} + \frac{J 2^{J(d-2)} \ln(n)}{n} + \frac{1}{n} + \frac{2^{Jd} \ln(n\epsilon)}{n\epsilon} \end{aligned} \tag{24}$$

for J larger than a sufficiently large constant that also depend on R.

Proof. See Appendix E.3 □

Ignoring poly-logarithmic terms and treating R as a constant for conciseness, we can optimize over J to obtain an asymptotic upper bound on the error as follows

$$\begin{aligned} & \mathbb{E}\left(\int \|T_0 - \hat{T}_{\text{priv}}\|^2 dP\right) \\ & \quad \lesssim n^{-1} \vee n^{-\frac{2\alpha}{2\alpha - 2 + d}} \vee (n\epsilon)^{-\frac{2\alpha}{2\alpha + d}} \ . \end{aligned} \tag{25}$$

Remark 4.4. It is possible to obtain similar bounds that hold with high probability with minimal adaptations to the proofs.

5. Lower Bound

Packing arguments from (Hütter & Rigollet, 2021) as well as coupling arguments (Acharya et al., 2018; 2021; Lalanne et al., 2023a) can be applied to handle the DP nature of the constraint, and yield the following lower bound.

Theorem 5.1 (Lower Bound)**.** *Asymptotically,*

$$\begin{aligned} & \inf_{\hat{T}_{\text{priv}}} \sup_{P \in \mathcal{M}, T_0 \in \mathcal{T}_\alpha} \mathbb{E}\left(\int \|T_0 - \hat{T}_{\text{priv}}\|^2 dP\right) \\ & \qquad \gtrsim \frac{1}{n} \vee n^{-\frac{2\alpha}{2\alpha - 2 + d}} \vee (n\epsilon)^{-\frac{2\alpha}{\alpha - 1 + d}} \end{aligned} \tag{26}$$

where the infimum is taken over all estimators $\hat{T}_{\text{priv}}$'s that are ϵ-DP.

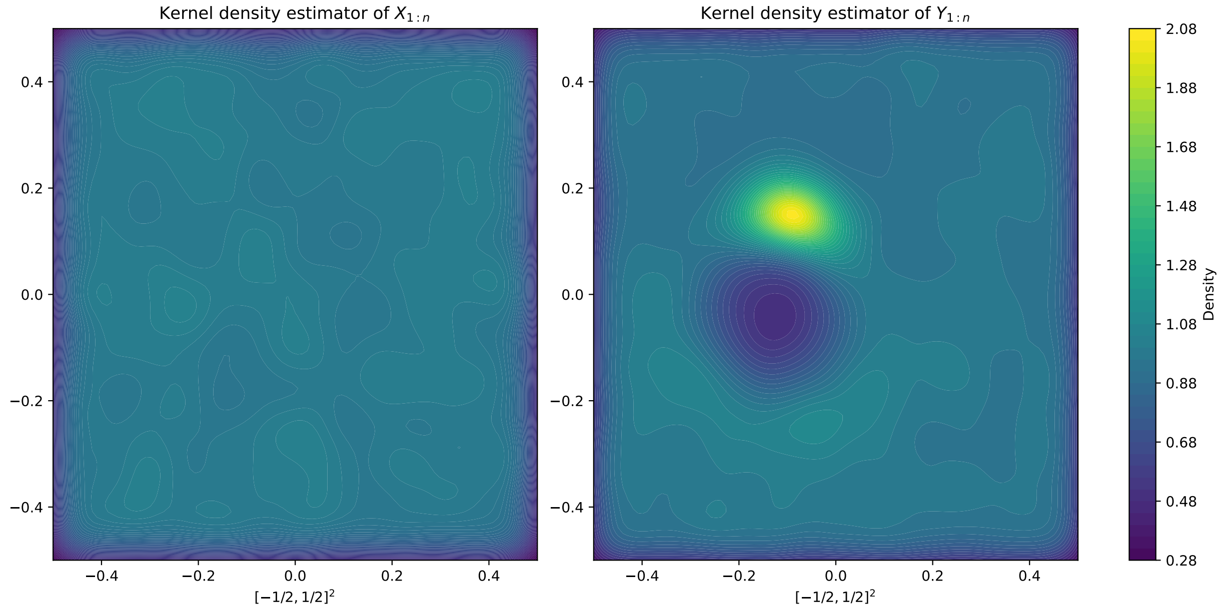

$n = 200000$, $\alpha_1 = \alpha_2 = 0.005$, $\sigma = \sigma_1 = \sigma_2 = 0.1$.

Figure 1. Kernel density estimators of the samples used in the experiments

Proof. See Appendix F.1 □

We observe that the term in $n\epsilon$ in the lower bound above does not match the corresponding term in the upper bound (25). It remains unclear which, if any, of these bounds is tight. However, it is worth noting that similar coupling arguments have given optimal lower bounds in other non-parametric problems (Lalanne et al., 2023b; Lalanne & Gadat, 2024). We thus believe that the lower bound is more likely to be optimal than the upper bound.

Following this conjecture, the suboptimality may arise from a suboptimal analysis of the estimator, a suboptimal choice of a covering, or limitations of the proposed estimator itself. Furthermore, our estimator is, to the best of our knowledge, the only one that has been proposed to solve the problem.

Another consequence of this lower bound is that the estimation is provably degraded by privacy (even with the best estimator) when $\epsilon \lesssim n^{-\frac{\alpha-1}{2\alpha-2+d}} \wedge n^{-\frac{\alpha+1-d}{2\alpha}}$.

6. Experiments

This section focuses on how to numerically approximate $\hat{f}_{\text{priv}}$ and $\hat{T}_{\text{priv}}$ defined in (20) and presents numerical results to illustrate the proposed method.

6.1. Discretization

The main difficulty in computing $\hat{f}_{\text{priv}}$ and $\hat{T}_{\text{priv}}$ is that the Fenchel-Legendre transform f^* of a candidate Brenier potential f may not always have a closed-form expression. To address this issue, (Hütter & Rigollet, 2021) approximate f^* by restricting the supremum in its definition to a grid. This subsection develops similar ideas to adapt this method to the context of differential privacy.

Let $\tilde{\Omega}^{(\text{grid})} \subset \tilde{\Omega}$ be a *finite* approximation set. In practice it is taken to be a uniform grid. We then define

$$f^*_{(\text{grid})}(y) = \sup_{x \in \tilde{\Omega}^{(\text{grid})}} \langle x, y \rangle - f(x), \quad \forall y \in \tilde{\Omega}. \tag{27}$$

Note that the sensitivity analysis in Lemma 3.3 still holds under this approximation of the Fenchel dual.

We represent potentials f by storing their discretization $f_{(\text{grid})}$ on the grid $\tilde{\Omega}^{(\text{grid})}$. Their gradients are approximated using finite differences, denoted by $\nabla_{(\text{grid})} f_{(\text{grid})}$.

Since we now only store the value of the potential and its (approximate) dual on a grid, we need to clip the points to this grid. We thus denote by $X^{(\text{grid})}_{1:n}$ and $Y^{(\text{grid})}_{1:n}$ the modified versions of the datasets $X_{1:n}$ and $Y_{1:n}$ where each point is replaced with its closest neighbor in $\tilde{\Omega}^{(\text{grid})}$.

Finally, if one wishes to apply Lemma 3.3 but cannot guarantee that $f \in \mathcal{X}(2M)$, we propose introducing a clipping constant C in order to safeguard privacy.

This procedure leads to the following practical algorithm: Start with a finite family of candidate discretized potentials

$$F = \left\{ f_{(\text{grid}),1}, \ldots, f_{(\text{grid}),N} \right\} , \tag{28}$$

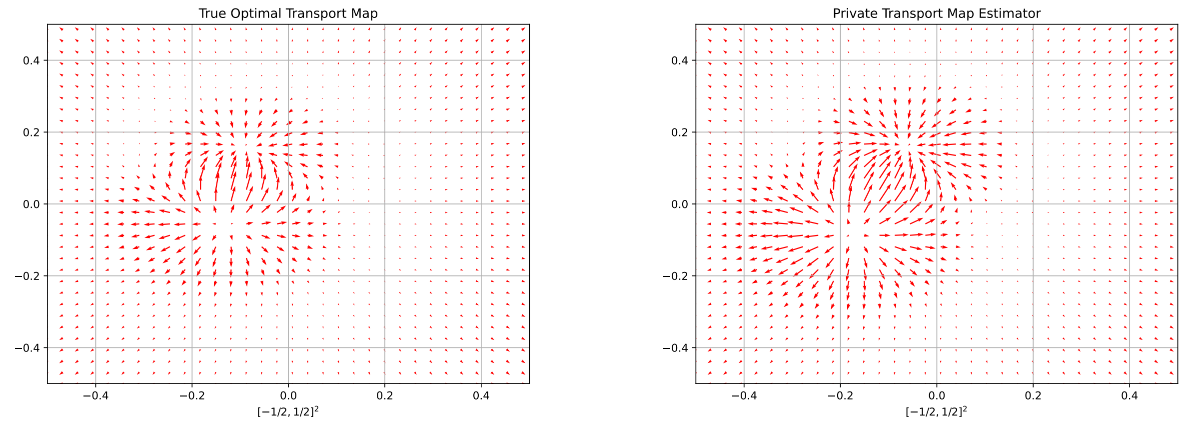

On the left: The true optimal transport map between the underlying distributions in the Gaussian attraction / repulsion model described in Section 6.2. On the right: The transport map that is privately estimated using the approximation algorithm described in Section 6.1. For the numerical values, $\tilde{\Omega}^{(\mathrm{grid})}$ is a uniform grid in $[-1/2, 1/2]^2$ with 64^2 points, $\epsilon = 1.$, $n = 200000$, $\alpha_1 = \alpha_2 = 0.005$, $\sigma = \sigma_1 = \sigma_2 = 0.1$, $T = 2000$, and $C = 0.25$.

Figure 2. Optimal Transport Map VS Estimated Private Map

then compute

$$\hat{i} \in \operatorname{argmin}\left\{ \hat{S}^C_{(\mathrm{grid})}(f_{(\mathrm{grid}),i} | X^{(\mathrm{grid})}_{1:n}, Y^{(\mathrm{grid})}_{1:n}) + \frac{4C}{n\epsilon} L_i \right\} \tag{29}$$

where for any f, $X_{1:n}$ and $Y_{1:n}$

$$\begin{aligned} \hat{S}^C_{(\mathrm{grid})}(f | X_{1:n}, Y_{1:n}) := \frac{1}{n} \sum_{i=1}^{n} \mathrm{Proj}_{[-C,C]}\, f(X_i) \\ + \frac{1}{n} \sum_{i=1}^{n} \mathrm{Proj}_{[-C,C]}\, f^*_{(\mathrm{grid})}(Y_i) \end{aligned} \tag{30}$$

and where $L_1, \ldots, L_N$ are independent Laplace random variables. By Lemma 3.2, $\hat{i}$ is ϵ-DP, and by post-processing, so is $\hat{T} := \nabla_{(\mathrm{grid})} f_{(\mathrm{grid}),\hat{i}}$, which we use as an estimator.

6.2. Gaussian Attraction / Repulsion Model

We illustrate the behavior of the estimator $\hat{T}$ defined in Section 6.1 on the following model, which we refer to as the Gaussian attraction/repulsion model.

We begin with a simple quadratic potential $f_{\mathrm{quad}}(x) = \frac{1}{2}\|x\|^2$ for which the mapping induced by its gradient is the identity. For any μ and $\sigma > 0$, denoting the Gaussian potential by $f_{\mathcal{N}}(x|\mu, \sigma) = e^{-\|x-\mu\|^2/(2\sigma^2)}$, we consider potentials of the form

$$\begin{aligned} f(x|\alpha_1, \alpha_2, \mu_1, \mu_2, \sigma_1, \sigma_2) = \\ f_{\mathrm{quad}}(x) + \alpha_1 f_{\mathcal{N}}(x|\mu_1, \sigma_1) - \alpha_2 f_{\mathcal{N}}(x|\mu_2, \sigma_2) \end{aligned} \tag{31}$$

where $\alpha_1, \alpha_2, \sigma_1, \sigma_2$ are positive numbers and $\mu_1, \mu_2 \in \mathbb{R}^d$ are location parameters.

Then, a random potential is generated by sampling two independent copies $\mu_1, \mu_2 \sim \mathcal{N}(0, \sigma I_2)$. This defines the "true" potential, which is the target of the estimation.

We then sample n i.i.d. uniform random variables $X_{1:n}$ over $[-1/2, 1/2]^2$, and generate the observations $Y_{1:n}$ by applying the gradient of the previously defined potential to n new i.i.d. uniform random variables over $[-1/2, 1/2]^2$. This is illustrated in Figure 1, where we represent two kernel density estimators constructed on the datasets $X_{1:n}$ and $Y_{1:n}$, respectively.

We then apply the algorithm defined in Section 6.1 where F is a collection of T independent potentials that are generated according to the same distribution as the true potential.

The result of the algorithm and the optimal transport map are represented in Figure 2. This simulation confirms that the estimated transport map is close to the optimal one. The code used to generate those results can be found at `/https://github.com/clemlal/PrivateSmoothTransportMapRNArgMin`.

7. Conclusion

In this paper, we have introduced a private estimator for the problem of privately estimating smooth optimal transport maps. We have analyzed the utility of this estimator and derived a lower bound on the difficulty of the statistical problem. Finally, we proposed an adaptation of this theoretical estimator that is implementable in practice.

This work opens up several research directions. One of them is the quest for an optimal private mechanism, and another one is the design of numerically efficient algorithms. In both cases, we believe that investigating *regularization* techniques such as *entropic regularization* would be an important step forward as it could help stabilize the problem, reduce the amount of privacy noise, and enable the use of numerically efficient algorithms such as Sinkhorn's algorithm (Cuturi, 2013).

Acknowledgements

This paper has been partially funded by the Agence Nationale de la Recherche under grants ANR-17-EURE-0010 (Investissements d'Avenir program), ANR-23-CE23-0029 Regul-IA, and ANR-24-CE23-1529 MAD. The authors also acknowledge the support of the AI Cluster ANITI (ANR-19-PI3A-0004). In addition, we thank the anonymous reviewers for the constructive discussions that led, in particular, to the addition of Appendix G and Appendix H.

Impact Statement

This paper presents work whose goal is to advance the field of Privacy-Preserving Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Abowd, J. M. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2867–2867, 2018.

Acharya, J., Sun, Z., and Zhang, H. Differentially private testing of identity and closeness of discrete distributions. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 6879–6891, 2018. URL /https://proceedings.neurips.cc/paper/2018/hash/7de32147a4f1055bed9e4faf3485a84d-Abstract.html.

Acharya, J., Sun, Z., and Zhang, H. Differentially private Assouad, Fano, and Le Cam. In Feldman, V., Ligett, K., and Sabato, S. (eds.), *Algorithmic Learning Theory, 16-19 March 2021, Virtual Conference, Worldwide*, volume 132 of *Proceedings of Machine Learning Research*, pp. 48–78. PMLR, 2021. URL /http://proceedings.mlr.press/v132/acharya21a.html.

Aden-Ali, I., Ashtiani, H., and Kamath, G. On the sample complexity of privately learning unbounded high-dimensional gaussians. In Feldman, V., Ligett, K., and Sabato, S. (eds.), *Algorithmic Learning Theory, 16-19 March 2021, Virtual Conference, Worldwide*, volume 132 of *Proceedings of Machine Learning Research*, pp. 185–216. PMLR, 2021. URL /http://proceedings.mlr.press/v132/aden-ali21a.html.

Alaux, J., Grave, E., Cuturi, M., and Joulin, A. Unsupervised hyperalignment for multilingual word embeddings. *CoRR*, abs/1811.01124, 2018. URL /http://arxiv.org/abs/1811.01124.

Altschuler, J. M., Weed, J., and Rigollet, P. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 1964–1974, 2017. URL /https://proceedings.neurips.cc/paper/2017/hash/491442df5f88c6aa018e86dac21d3606-Abstract.html.

Alvarez-Melis, D., Jaakkola, T. S., and Jegelka, S. Structured optimal transport. In Storkey, A. J. and Pérez-Cruz, F. (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1771–1780. PMLR, 2018. URL /http://proceedings.mlr.press/v84/alvarez-melis18a.html.

Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 214–223. PMLR, 2017. URL /http://proceedings.mlr.press/v70/arjovsky17a.html.

Backstrom, L., Dwork, C., and Kleinberg, J. M. Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography. In Williamson, C. L., Zurko, M. E., Patel-Schneider, P. F., and Shenoy, P. J. (eds.), *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pp. 181–190. ACM, 2007. doi: /10.1145/1242572.1242598. URL /https://doi.org/10.1145/1242572.1242598.

Barber, R. F. and Duchi, J. C. Privacy and statistical risk: Formalisms and minimax bounds. *CoRR*, abs/1412.4451, 2014. URL /http://arxiv.org/abs/1412.4451.

Beraha, M., Favaro, S., and Rao, V. Mcmc for bayesian nonparametric mixture modeling under differential privacy. *arXiv preprint arXiv:2310.09818*, 2023.

Berrett, T. B., Györfi, L., and Walk, H. Strongly universally consistent nonparametric regression and classification with privatised data. *Electronic Journal of Statistics*, 15 (1):2430 – 2453, 2021. doi: /10.1214/21-EJS1845. URL /https://doi.org/10.1214/21-EJS1845.

Biswas, S., Dong, Y., Kamath, G., and Ullman, J. R. Coinpress: Practical private mean and covariance estimation. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL /https://proceedings.neurips.cc/paper/2020/hash/a684eceee76fc522773286a895bc8436-Abstract.html.

Black, E., Yeom, S., and Fredrikson, M. Fliptest: fairness testing via optimal transport. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 111–121, 2020.

Brenier, Y. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.

Brown, G., Gaboardi, M., Smith, A. D., Ullman, J. R., and Zakynthinou, L. Covariance-aware private mean estimation without private covariance estimation. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 7950–7964, 2021. URL /https://proceedings.neurips.cc/paper/2021/hash/42778ef0b5805a96f9511e20b5611fce-Abstract.html.

Bun, M., Kamath, G., Steinke, T., and Wu, Z. S. Private hypothesis selection. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 156–167, 2019. URL /https://proceedings.neurips.cc/paper/2019/hash/9778d5d219c5080b9a6a17bef029331c-Abstract.html.

Bun, M., Kamath, G., Steinke, T., and Wu, Z. S. Private hypothesis selection. *IEEE Trans. Inf. Theory*, 67 (3):1981–2000, 2021. doi: /10.1109/TIT.2021.3049802. URL /https://doi.org/10.1109/TIT.2021.3049802.

Butucea, C., Dubois, A., Kroll, M., and Saumard, A. Local differential privacy: Elbow effect in optimal density estimation and adaptation over besov ellipsoids. *CoRR*, abs/1903.01927, 2019. URL /http://arxiv.org/abs/1903.01927.

Cai, T. T., Wang, Y., and Zhang, L. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *CoRR*, abs/1902.04495, 2019. URL /http://arxiv.org/abs/1902.04495.

Cañas, G. D. and Rosasco, L. Learning probability measures with respect to optimal transport metrics. In Bartlett, P. L., Pereira, F. C. N., Burges, C. J. C., Bottou, L., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, pp. 2501–2509, 2012. URL /https://proceedings.neurips.cc/paper/2012/hash/c54e7837e0cd0ced286cb5995327d1ab-Abstract.html.

Cazelles, E., Seguy, V., Bigot, J., Cuturi, M., and Papadakis, N. Geodesic PCA versus log-pca of histograms in the wasserstein space. *SIAM J. Sci. Comput.*, 40(2), 2018. doi: /10.1137/17M1143459. URL /https://doi.org/10.1137/17M1143459.

Chhor, J. and Sentenac, F. Robust estimation of discrete distributions under local differential privacy. In *International Conference on Algorithmic Learning Theory*, pp. 411–446. PMLR, 2023.

Courty, N., Flamary, R., and Tuia, D. Domain adaptation with regularized optimal transport. In Calders, T., Esposito, F., Hüllermeier, E., and Meo, R. (eds.), *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part I*, volume 8724 of *Lecture Notes in Computer Science*, pp. 274–289. Springer, 2014. doi: /10.1007/978-3-662-44848-9_18. URL /https://doi.org/10.1007/978-3-662-44848-9_18.

Courty, N., Flamary, R., Habrard, A., and Rakotomamonjy, A. Joint distribution optimal transportation for domain adaptation. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017a. URL /https://proceedings.neurips.cc/paper_files/paper/2017/file/0070d23b06b1486a538c0eaa45dd167a-Paper.pdf.

Courty, N., Flamary, R., Tuia, D., and Rakotomamonjy, A. Optimal transport for domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(9):1853–1865, 2017b. doi: /10.1109/TPAMI.2016.2615921. URL /https://doi.org/10.1109/TPAMI.2016.2615921.

Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. In Burges, C. J. C., Bottou, L., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pp. 2292–2300, 2013. URL /https://proceedings.neurips.cc/paper/2013/hash/af21d0c97db2e27e13572cbf59eb343d-Abstract.html.

Damodaran, B. B., Kellenberger, B., Flamary, R., Tuia, D., and Courty, N. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y. (eds.), *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, volume 11208 of *Lecture Notes in Computer Science*, pp. 467–483. Springer, 2018. doi: /10.1007/978-3-030-01225-0_28. URL /https://doi.org/10.1007/978-3-030-01225-0_28.

De Lara, L., González-Sanz, A., Asher, N., Risser, L., and Loubes, J.-M. Transport-based counterfactual models. *Journal of Machine Learning Research*, 25(136):1–59, 2024.

del Barrio, E., Gordaliza, P., Lescornel, H., and Loubes, J. Central limit theorem and bootstrap procedure for wasserstein's variations with an application to structural relationships between distributions. *J. Multivar. Anal.*, 169:341–362, 2019. doi: /10.1016/J.JMVA.2018.09.014. URL /https://doi.org/10.1016/j.jmva.2018.09.014.

del Barrio, E., González-Sanz, A., and Loubes, J.-M. Central limit theorems for general transportation costs. In *Annales de l'Institut Henri Poincare (B) Probabilites et statistiques*, volume 60, pp. 847–873. Institut Henri Poincaré, 2024.

Del Barrio, E., González Sanz, A., and Loubes, J.-M. Central limit theorems for semi-discrete wasserstein distances. *Bernoulli*, 30(1):554–580, 2024.

Diakonikolas, I., Hardt, M., and Schmidt, L. Differentially private learning of structured discrete distributions. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 2566–2574, 2015. URL /https://proceedings.neurips.cc/paper/2015/hash/2b3bf3eee2475e03885a110e9acaab61-Abstract.html.

Ding, B., Kulkarni, J., and Yekhanin, S. Collecting telemetry data privately. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 3571–3580, 2017. URL /https://proceedings.neurips.cc/paper/2017/hash/253614bbac999b38b5b60cae531c4969-Abstract.html.

Ding, Z., Kifer, D., E., S. M. S. N., Steinke, T., Wang, Y., Xiao, Y., and Zhang, D. The permute-and-flip mechanism is identical to report-noisy-max with exponential noise. *CoRR*, abs/2105.07260, 2021. URL /https://arxiv.org/abs/2105.07260.

Dinur, I. and Nissim, K. Revealing information while preserving privacy. In Neven, F., Beeri, C., and Milo, T. (eds.), *Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 9-12, 2003, San Diego, CA, USA*, pp. 202–210. ACM, 2003. doi: /10.1145/773153.773173. URL /https://doi.org/10.1145/773153.773173.

Duchi, J. C., Jordan, M. I., and Wainwright, M. J. Local privacy and statistical minimax rates. In *51st Annual Allerton Conference on Communication, Control, and Computing, Allerton 2013, Allerton Park & Retreat Center, Monticello, IL, USA, October 2-4, 2013*, pp. 1592. IEEE, 2013. doi: /10.1109/Allerton.2013.6736718. URL /https://doi.org/10.1109/Allerton.2013.6736718.

Duchi, J. C., Wainwright, M. J., and Jordan, M. I. Minimax optimal procedures for locally private estimation. *CoRR*, abs/1604.02390, 2016. URL /http://arxiv.org/abs/1604.02390.

Dvurechensky, P. E., Gasnikov, A. V., and Kroshnin, A. Computational optimal transport: Complexity by accelerated gradient descent is better than by sinkhorn's algorithm. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1366–1375. PMLR, 2018. URL /http://proceedings.mlr.press/v80/dvurechensky18a.html.

Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. doi: /10.1561/0400000042. URL /https://doi.org/10.1561/0400000042.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. D. Calibrating noise to sensitivity in private data analysis. In Halevi, S. and Rabin, T. (eds.), *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pp. 265–284. Springer, 2006. doi: /10.1007/11681878_14. URL /https://doi.org/10.1007/11681878_14.

Erlingsson, Ú., Pihur, V., and Korolova, A. RAPPOR: randomized aggregatable privacy-preserving ordinal response. In Ahn, G., Yung, M., and Li, N. (eds.), *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pp. 1054–1067. ACM, 2014. doi: /10.1145/2660267.2660348. URL /https://doi.org/10.1145/2660267.2660348.

Evfimievski, A., Gehrke, J., and Srikant, R. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '03, pp. 211–222, New York, NY, USA, 2003. Association for Computing Machinery. ISBN 1581136706. doi: /10.1145/773153.773174. URL /https://doi.org/10.1145/773153.773174.

Feydy, J., Charlier, B., Vialard, F., and Peyré, G. Optimal transport for diffeomorphic registration. In Descoteaux, M., Maier-Hein, L., Franz, A. M., Jannin, P., Collins, D. L., and Duchesne, S. (eds.), *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017 - 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I*, volume 10433 of *Lecture Notes in Computer Science*, pp. 291–299. Springer, 2017. doi: /10.1007/978-3-319-66182-7_34. URL /https://doi.org/10.1007/978-3-319-66182-7_34.

Flamary, R., Cuturi, M., Courty, N., and Rakotomamonjy, A. Wasserstein discriminant analysis. *Mach. Learn.*, 107 (12):1923–1945, 2018. doi: /10.1007/S10994-018-5717-1. URL /https://doi.org/10.1007/s10994-018-5717-1.

Forrow, A., Hütter, J., Nitzan, M., Rigollet, P., Schiebinger, G., and Weed, J. Statistical optimal transport via factored couplings. In Chaudhuri, K. and Sugiyama, M. (eds.), *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2454–2465. PMLR, 2019. URL /http://proceedings.mlr.press/v89/forrow19a.html.

Fredrikson, M., Jha, S., and Ristenpart, T. Model inversion attacks that exploit confidence information and basic countermeasures. In Ray, I., Li, N., and Kruegel, C. (eds.), *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pp. 1322–1333. ACM, 2015. doi: /10.1145/2810103.2813677. URL /https://doi.org/10.1145/2810103.2813677.

Genevay, A., Peyré, G., and Cuturi, M. Learning generative models with sinkhorn divergences. In Storkey, A. J. and Pérez-Cruz, F. (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1608–1617. PMLR, 2018. URL /http://proceedings.mlr.press/v84/genevay18a.html.

Gordaliza, P., del Barrio, E., Gamboa, F., and Loubes, J. Obtaining fairness using optimal transport theory. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2357–2365. PMLR, 2019. URL /http://proceedings.mlr.press/v97/gordaliza19a.html.

Grave, E., Joulin, A., and Berthet, Q. Unsupervised alignment of embeddings with wasserstein procrustes. In Chaudhuri, K. and Sugiyama, M. (eds.), *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1880–1890. PMLR, 2019. URL /http://proceedings.mlr.press/v89/grave19a.html.

Györfi, L. and Kroll, M. On rate optimal private regression under local differential privacy. *arXiv preprint arXiv:2206.00114*, 2022.

Györfi, L. and Kroll, M. Multivariate density estimation from privatised data: universal consistency and minimax rates. *Journal of Nonparametric Statistics*, 0(0): 1–23, 2023. doi: /10.1080/10485252.2022.2163634. URL /https://doi.org/10.1080/10485252.2022.2163634.

Harder, F., Adamczewski, K., and Park, M. Dp-merf: Differentially private mean embeddings with randomfeatures for practical privacy-preserving data generation. In *International conference on artificial intelligence and statistics*, pp. 1819–1827. PMLR, 2021.

Homer, N., Szelinger, S., Redman, M., Duggan, D., Tembe, W., Muehling, J., Pearson, J. V., Stephan, D. A., Nelson, S. F., and Craig, D. W. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS Genet*, 4 (8):e1000167, 2008.

Hütter, J.-C. and Rigollet, P. Minimax estimation of smooth optimal transport maps. *The Annals of Statistics*, 49(2): 1166 – 1194, 2021. doi: /10.1214/20-AOS1997. URL /https://doi.org/10.1214/20-AOS1997.

Janati, H., Cuturi, M., and Gramfort, A. Wasserstein regularization for sparse multi-task regression. In Chaudhuri, K. and Sugiyama, M. (eds.), *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1407–1416. PMLR, 2019. URL /http://proceedings.mlr.press/v89/janati19a.html.

Kallenberg, O. Lectures on the coupling method (torgny lindvall). *SIAM Review*, 35(3):525–527, 1993. doi: /10.1137/1035121. URL /https://doi.org/10.1137/1035121.

Kamath, G., Li, J., Singhal, V., and Ullman, J. R. Privately learning high-dimensional distributions. In Beygelzimer, A. and Hsu, D. (eds.), *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pp. 1853–1902. PMLR, 2019. URL /http://proceedings.mlr.press/v99/kamath19a.html.

Kamath, G., Singhal, V., and Ullman, J. R. Private mean estimation of heavy-tailed distributions. In Abernethy, J. D. and Agarwal, S. (eds.), *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pp. 2204–2235. PMLR, 2020. URL /http://proceedings.mlr.press/v125/kamath20a.html.

Kamath, G., Liu, X., and Zhang, H. Improved rates for differentially private stochastic convex optimization with heavy-tailed data. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 10633–10660. PMLR, 2022a. URL /https://proceedings.mlr.press/v162/kamath22a.html.

Kamath, G., Mouzakis, A., and Singhal, V. New lower bounds for private estimation and a generalized fingerprinting lemma. In *NeurIPS*, 2022b. URL /http://papers.nips.cc/paper_files/paper/2022/hash/9a6b278218966499194491f55ccf8b75-Abstract-Conference.html.

Kamath, G., Mouzakis, A., Regehr, M., Singhal, V., Steinke, T., and Ullman, J. R. A bias-variance-privacy trilemma for statistical estimation. *CoRR*, abs/2301.13334, 2023. doi: /10.48550/ARXIV.2301.13334. URL /https://doi.org/10.48550/arXiv.2301.13334.

Karwa, V. and Vadhan, S. P. Finite sample differentially private confidence intervals. In Karlin, A. R. (ed.), *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pp. 44:1–44:9. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi: /10.4230/LIPIcs.ITCS.2018.44. URL /https://doi.org/10.4230/LIPIcs.ITCS.2018.44.

Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., and Smith, A. What can we learn privately? In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pp. 531–540, 2008. doi: /10.1109/FOCS.2008.27.

Kawamoto, Y. and Murakami, T. Local obfuscation mechanisms for hiding probability distributions. In Sako, K., Schneider, S. A., and Ryan, P. Y. A. (eds.), *Computer Security - ESORICS 2019 - 24th European Symposium on Research in Computer Security, Luxembourg, September 23-27, 2019, Proceedings, Part I*, volume 11735 of *Lecture Notes in Computer Science*, pp. 128–148. Springer, 2019. doi: /10.1007/978-3-030-29959-0_7. URL /https://doi.org/10.1007/978-3-030-29959-0_7.

Klatt, M., Tameling, C., and Munk, A. Empirical regularized optimal transport: Statistical theory and applications. *SIAM J. Math. Data Sci.*, 2(2):419–443, 2020. doi: /10.1137/19M1278788. URL /https://doi.org/10.1137/19M1278788.

Kroll, M. On density estimation at a fixed point under local differential privacy. *Electronic Journal of Statistics*, 15 (1):1783 – 1813, 2021. doi: /10.1214/21-EJS1830. URL /https://doi.org/10.1214/21-EJS1830.

Kroshnin, A., Spokoiny, V. G., and Suvorikova, A. L. Statistical inference for bures–wasserstein barycenters. *The Annals of Applied Probability*, 2019. URL /https://api.semanticscholar.org/CorpusID:88524508.

Lalanne, C. *On the tradeoffs of statistical learning with privacy*. Theses, Ecole normale supérieure de lyon - ENS LYON, October 2023. URL /https://theses.hal.science/tel-04379624.

Lalanne, C. and Gadat, S. Privately Learning Smooth Distributions on the Hypercube by Projections. In *ICML 2024 - 41st International Conference on Machine Learning*, pp. 39 p., Vienna, Austria, July 2024. URL /https://hal.science/hal-04549279.

Lalanne, C., Garivier, A., and Gribonval, R. On the Statistical Complexity of Estimation and Testing under Privacy Constraints. *Transactions on Machine Learning Research Journal*, April 2023a. URL /https://hal.science/hal-03794374.

Lalanne, C., Garivier, A., and Gribonval, R. About the cost of central privacy in density estimation. *Transactions on Machine Learning Research*, 2023b. ISSN 2835-8856. URL /https://openreview.net/forum?id=uq29MIWvIV.

Lalanne, C., Garivier, A., and Gribonval, R. Private Statistical Estimation of Many Quantiles. In *ICML 2023 - 40th International Conference on Machine Learning*, Honolulu, United States, July 2023c. URL /https://hal.science/hal-03986170.

Lalanne, C., Gastaud, C., Grislain, N., Garivier, A., and Gribonval, R. Private Quantiles Estimation in the Presence of Atoms. *Information and Inference*, August 2023d. doi: /10.1093/imaiai/iaad030. URL /https://hal.science/hal-03572701.

Lam-Weil, J., Laurent, B., and Loubes, J.-M. Minimax optimal goodness-of-fit testing for densities and multinomials under a local differential privacy constraint. *Bernoulli*, 28(1):579–600, 2022.

Lavenant, H., Claici, S., Chien, E., and Solomon, J. Dynamical optimal transport on discrete surfaces. *ACM Trans. Graph.*, 37(6):250, 2018. doi: /10.1145/3272127.3275064. URL /https://doi.org/10.1145/3272127.3275064.

Le, T., Nguyen, K., Sun, S., Han, K., Ho, N., and Xie, X. Diffeomorphic mesh deformation via efficient optimal transport for cortical surface reconstruction. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL /https://openreview.net/forum?id=gxhRR8vUQb.

Loukides, G., Denny, J. C., and Malin, B. A. The disclosure of diagnosis codes can breach research participants' privacy. *J. Am. Medical Informatics Assoc.*, 17(3):322–327, 2010. doi: /10.1136/jamia.2009.002725. URL /https://doi.org/10.1136/jamia.2009.002725.

Manole, T., Balakrishnan, S., Niles-Weed, J., and Wasserman, L. Plugin estimation of smooth optimal transport maps, 2024. URL /https://arxiv.org/abs/2107.12364.

McKenna, R. and Sheldon, D. Permute-and-flip: A new mechanism for differentially private selection. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL /https://proceedings.neurips.cc/paper/2020/hash/01e00f2f4bfcbb7505cb641066f2859b-Abstract.html.

McSherry, F. and Talwar, K. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pp. 94–103. IEEE Computer Society, 2007. doi: /10.1109/FOCS.2007.41. URL /https://doi.org/10.1109/FOCS.2007.41.

Monge, G. *Mémoire sur la théorie des déblais et des remblais*. De l'Imprimerie Royale, 1781.

Montavon, G., Müller, K., and Cuturi, M. Wasserstein training of restricted boltzmann machines. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3711–3719, 2016. URL /https://proceedings.neurips.cc/paper/2016/hash/728f206c2a01bf572b5940d7d9a8fa4c-Abstract.html.

Narayanan, A. and Shmatikov, V. How to break anonymity of the netflix prize dataset. *CoRR*, abs/cs/0610105, 2006. URL /http://arxiv.org/abs/cs/0610105.

Narayanan, A. and Shmatikov, V. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (S&P 2008), 18-21 May 2008, Oakland, California, USA*, pp. 111–125. IEEE Computer Society, 2008. doi: /10.1109/SP.2008.33. URL /https://doi.org/10.1109/SP.2008.33.

Niles-Weed, J. and Rigollet, P. Estimation of Wasserstein distances in the Spiked Transport Model. *Bernoulli*, 28 (4):2663 – 2688, 2022. doi: /10.3150/21-BEJ1433. URL /https://doi.org/10.3150/21-BEJ1433.

Panaretos, V. M. and Zemel, Y. Statistical aspects of wasserstein distances. *Annual Review of Statistics and Its Application*, 6(Volume 6, 2019): 405–431, 2019. ISSN 2326-831X. doi: /https://doi.org/10.1146/annurev-statistics-030718-104938. URL /https://www.annualreviews.org/content/journals/10.1146/annurev-statistics-030718-104938.

Perrot, M., Courty, N., Flamary, R., and Habrard, A. Mapping estimation for discrete optimal transport. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 4197–4205, 2016. URL /https://proceedings.neurips.cc/paper/2016/hash/26f5bd4aa64fdadf96152ca6e6408068-Abstract.html.

Peyré, G. and Cuturi, M. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. ISSN 1935-8237. doi: /10.1561/2200000073. URL /http://dx.doi.org/10.1561/2200000073.

Pierquin, C., Bellet, A., Tommasi, M., and Boussard, M. Rényi pufferfish privacy: General additive noise mechanisms and privacy amplification by iteration via shift reduction lemmas. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL /https://openreview.net/forum?id=VZsxhPpu9T.

Rakotomamonjy, A. and Ralaivola, L. Differentially private sliced wasserstein distance. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8810–8820. PMLR, 2021. URL /http://proceedings.mlr.press/v139/rakotomamonjy21a.html.

Ramdas, A., Trillos, N. G., and Cuturi, M. On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017. doi: /10.3390/E19020047. URL /https://doi.org/10.3390/e19020047.

Randal-Williams, O. Algebraic topology. *Lecture Notes*, 2024. URL /https://www.dpmms.cam.ac.uk/~or257/teaching/notes/at.pdf.

Rigollet, P. and Weed, J. Entropic optimal transport is maximum-likelihood deconvolution. *Comptes Rendus Mathematique*, 356(11):1228–1235, 2018. ISSN 1631-073X. doi: /https://doi.org/10.1016/j.crma.2018.10.010. URL /https://www.sciencedirect.com/science/article/pii/S1631073X18302802.

Rüschendorf, L. and Rachev, S. T. A characterization of random variables with minimum l2-distance. *Journal of multivariate analysis*, 32(1):48–54, 1990.

Santambrogio, F. *Optimal Transport for Applied Mathematicians*. Birkhäuser Cham/Springer, 2016. doi: /10.1007/978-3-319-20828-2.

Schluttenhofer, S. and Johannes, J. Adaptive pointwise density estimation under local differential privacy, 2022.

Schmitz, M. A., Heitz, M., Bonneel, N., Mboula, F. M. N., Coeurjolly, D., Cuturi, M., Peyré, G., and Starck, J. Wasserstein dictionary learning: Optimal transport-based unsupervised nonlinear dictionary learning. *SIAM J. Imaging Sci.*, 11(1):643–678, 2018. doi: /10.1137/17M1140431. URL /https://doi.org/10.1137/17M1140431.

Sebag, I., Pydi, M. S., Franceschi, J., Rakotomamonjy, A., Gartrell, M., Atif, J., and Allauzen, A. Differentially private gradient flow based on the sliced wasserstein distance for non-parametric generative modeling. *CoRR*, abs/2312.08227, 2023. doi: /10.48550/ARXIV.2312.08227. URL /https://doi.org/10.48550/arXiv.2312.08227.

Seguy, V. and Cuturi, M. Principal geodesic analysis for probability measures under the optimal transport metric. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 3312–3320, 2015. URL /https://proceedings.neurips.cc/paper/2015/hash/f26dab9bf6a137c3b6782e562794c2f2-Abstract.html.

Seguy, V., Damodaran, B. B., Flamary, R., Courty, N., Rolet, A., and Blondel, M. Large scale optimal transport and mapping estimation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL /https://openreview.net/forum?id=B1zlp1bRW.

Singhal, V. A polynomial time, pure differentially private estimator for binary product distributions. *CoRR*, abs/2304.06787, 2023. doi: /10.48550/ARXIV.2304.06787. URL /https://doi.org/10.48550/arXiv.2304.06787.

Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., and Guibas, L. J. Convolutional wasserstein distances: efficient optimal transportation on geometric domains. *ACM Trans. Graph.*, 34(4):66:1–66:11, 2015. doi: /10.1145/2766963. URL /https://doi.org/10.1145/2766963.

Solomon, J., Peyré, G., Kim, V. G., and Sra, S. Entropic metric alignment for correspondence problems. *ACM Trans. Graph.*, 35(4):72:1–72:13, 2016. doi: /10.1145/2897824.2925903. URL /https://doi.org/10.1145/2897824.2925903.

Staib, M., Claici, S., Solomon, J., and Jegelka, S. Parallel streaming wasserstein barycenters. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 2647–2658, 2017. URL /https://proceedings.neurips.cc/paper/2017/hash/253f7b5d921338af34da817c00f42753-Abstract.html.

Sweeney, L. Simple demographics often identify people uniquely. *Health (San Francisco)*, 671(2000):1–34, 2000.

Sweeney, L. k-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.*, 10(5):557–570, 2002. doi: /10.1142/S0218488502001648. URL /https://doi.org/10.1142/S0218488502001648.

Tameling, C. and Munk, A. Computational strategies for statistical inference based on empirical optimal transport. In *2018 IEEE Data Science Workshop, DSW 2018, Lausanne, Switzerland, June 4-6, 2018*, pp. 175–179. IEEE, 2018. doi: /10.1109/DSW.2018.8439912. URL /https://doi.org/10.1109/DSW.2018.8439912.

Thakurta, A. G., Vyrros, A. H., Vaishampayan, U. S., Kapoor, G., Freudiger, J., Sridhar, V. R., and Davidson, D. Learning new words. *Granted US Patents*, 9594741, 2017.

Tien, N. L., Habrard, A., and Sebban, M. Differentially private optimal transport: Application to domain adaptation. In Kraus, S. (ed.), *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pp. 2852–2858. ijcai.org, 2019. doi: /10.24963/IJCAI.2019/395. URL /https://doi.org/10.24963/ijcai.2019/395.

Tsybakov, A. B. *Introduction to Nonparametric Estimation*. Springer series in statistics. Springer, 2009. ISBN 978-0-387-79051-0. doi: /10.1007/b13794. URL /https://doi.org/10.1007/b13794.

Villani, C. *Topics in Optimal Transportation*. Graduate studies in mathematics. American Mathematical Society, 2003. ISBN 9781470418045. URL /https://books.google.fr/books?id=MyPjjgEACAAJ.

Villani, C. et al. *Optimal transport: old and new*, volume 338. Springer, 2009. ISBN 978-3-540-71049-3. doi: /10.1007/978-3-540-71050-9.

Wagner, I. and Eckhoff, D. Technical privacy metrics: A systematic survey. *ACM Comput. Surv.*, 51(3):57:1–57:38, 2018. doi: /10.1145/3168389. URL /https://doi.org/10.1145/3168389.

Wasserman, L. A. and Zhou, S. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010. doi: /10.1198/jasa.2009.tm08651. URL /https://doi.org/10.1198/jasa.2009.tm08651.

Weed, J. and Berthet, Q. Estimation of smooth densities in wasserstein distance. In Beygelzimer, A. and Hsu, D. (eds.), *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pp. 3118–3119. PMLR, 2019. URL /http://proceedings.mlr.press/v99/weed19a.html.

Xian, R., Li, Q., Kamath, G., and Zhao, H. Differentially private post-processing for fair regression. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL /https://openreview.net/forum?id=JNeeRjKbuH.

Yang, C., Qi, J., and Zhou, A. Wasserstein differential privacy. In Wooldridge, M. J., Dy, J. G., and Natarajan, S. (eds.), *Thirty-Eighth AAAI Conference on Artificial*

Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada, pp. 16299–16307. AAAI Press, 2024. doi: /10.1609/AAAI.V38I15.29565. URL /https://doi.org/10.1609/aaai.v38i15.29565.

Zemel, Y. and Panaretos, V. M. Fréchet means and Procrustes analysis in Wasserstein space. *Bernoulli*, 25(2): 932 – 976, 2019. doi: /10.3150/17-BEJ1009. URL /https://doi.org/10.3150/17-BEJ1009.

A. Details on Wavelet Decomposition

In this article, we adopt the same conventions for functional spaces as in (Hütter & Rigollet, 2021). In particular, we encourage the reader to consult Appendix B in (Hütter & Rigollet, 2021) for formal definitions of the functional spaces considered in this paper.

A notable difference from (Hütter & Rigollet, 2021) is that all the functions considered in this work are supported on $\tilde{\Omega}$, which is a hypercube rather than a more general Lipschitz domain. Therefore, we can restrict the wavelet basis to functions whose support is included in $\tilde{\Omega}$ only, thus avoiding overlaps with the complement: for any function in this basis, its support is either included in $\tilde{\Omega}$ or in $\bar{\tilde{\Omega}} \cup \partial\tilde{\Omega}$.

As a consequence, if $f \in V_J$, and if $(\gamma_k^{j,g})_{k,j,g}$ denotes its wavelet coefficients in the corresponding wavelet basis, then

$$\|f\|_{L^2(\tilde{\Omega})}^2 = \|(\gamma_k^{j,g})_{k,j,g}\|_2^2 \ . \tag{32}$$

B. Technical Tools

Lemma B.1. *There exists an open neighborhood O of 0 in $\mathbb{R}^{q\times q}$ such that*

$$M \in O \quad \implies \quad |\det(I_q + M) - 1 - \mathrm{tr}(M)| \leq \|M\| \ , \tag{33}$$

where O is only affected by the choice of $\|\cdot\|$.

Proof. By computing the partial derivatives of $M \mapsto \det(I_q + M)$ in the canonical basis of $\mathbb{R}^{q\times q}$, we immediately see that the gradient of $M \mapsto \det(I_q + M)$ at 0 for the Euclidean structure induced by the Frobenius inner product is I_d. Consequently, since $M \mapsto \det(I_q + M)$ is $\mathcal{C}^1$ (it has a polynomial expression in the coefficients of M),

$$\det(I_q + M) = 1 + \mathrm{tr}(M) + o(\|M\|) \ . \tag{34}$$

The equivalence of norms in finite dimension allows us to conclude. □

Lemma B.2. *Let $L_1, \ldots, L_M$ be M independent random variables with Laplace distribution (that is, with pdf $t \mapsto \frac{1}{2}e^{-|t|}$ w.r.t. Lebesgue's measure). Then*

$$\mathbb{E}\left(\max_{i=1,\ldots,N} |L_i|\right) \leq 1 + \ln(N) \ . \tag{35}$$

Proof. Let $\delta \geq 0$, by independence,

$$\begin{aligned}\mathbb{P}\left(\max_{i=1,\ldots,N} |L_i| \leq \delta\right) &= \mathbb{P}\left(\bigcap_{i=1,\ldots,N} (|L_i| \leq \delta)\right) \\ &= \prod_{i=1}^{N} \mathbb{P}\left((|L_i| \leq \delta)\right) \\ &= \left(1 - e^{-\delta}\right)^N \ .\end{aligned} \tag{36}$$

Thus,

$$\begin{aligned}\mathbb{E}\left(\max_{i=1,\dots,N}|L_i|\right) &= \int_{[0,+\infty)} \mathbb{P}\left(\max_{i=1,\dots,N}|L_i|>\delta\right)d\delta \\ &= \int_{(0,+\infty)}\left(1-\left(1-e^{-\delta}\right)^N\right)d\delta \\ &\overset{u=e^{-\delta}}{=} \int_{(0,1)}\left(1-(1-u)^N\right)\frac{du}{u} \\ &\overset{(1-u)^N\geq \max(0,1-Nu)}{=} \int_{(0,1)}\left(1-\max(0,1-Nu)\right)\frac{du}{u} \\ &= \int_0^{\frac{1}{N}}\left(1-\max(0,1-Nu)\right)\frac{du}{u} + \int_{\frac{1}{N}}^1\left(1-\max(0,1-Nu)\right)\frac{du}{u} \\ &\leq 1+\ln(N)\ .\end{aligned} \tag{37}$$

□

Lemma B.3. *Let $f \in V_J$, and denote by $(\gamma_k^{j,g})_{k,j,g}$ its wavelet coefficients for compactly supported mother and father wavelets, and with support adequation to $\tilde{\Omega}$ as described in Appendix A. Then*

$$\|(\gamma_k^{j,g})_{k,j,g}\|_\infty \leq \|(\gamma_k^{j,g})_{k,j,g}\|_2 \lesssim \|f\|_\infty \lesssim 2^{\frac{Jd}{2}}\|(\gamma_k^{j,g})_{k,j,g}\|_\infty\ . \tag{38}$$

Proof.

$$\begin{aligned}\|f\|_\infty &= \sqrt{\|f^2\|_\infty} \\ &\geq \sqrt{\frac{1}{\mathrm{Vol}\left(\tilde{\Omega}\right)}\int_{\tilde{\Omega}} f^2} \\ &\overset{\text{Parseval}}{=} \sqrt{\frac{1}{\mathrm{Vol}\left(\tilde{\Omega}\right)}\|(\gamma_k^{j,g})_{k,j,g}\|_2^2} \\ &\gtrsim \|(\gamma_k^{j,g})_{k,j,g}\|_2 \\ &\geq \|(\gamma_k^{j,g})_{k,j,g}\|_\infty\ .\end{aligned} \tag{39}$$

The inequality

$$\|f\|_\infty \lesssim 2^{\frac{Jd}{2}}\|(\gamma_k^{j,g})_{k,j,g}\|_\infty \tag{40}$$

comes from Lemma 24 in (Hütter & Rigollet, 2021).

□

C. Proofs of Section 2

C.1. Proof of Lemma 2.5

Following subsection 5.4 in (Hütter & Rigollet, 2021), we use $\|\cdot\| := \|\cdot\|_{L^2(P)}$ and we define

$$\forall f,\quad S_0(f) := S(f) - S(f_0),\quad \hat{S}_0(f) := \hat{S}(f|X_{1:n},Y_{1:n}) - \hat{S}(f_0|X_{1:n},Y_{1:n})\ , \tag{41}$$

$$\forall \tau \geq 0,\quad \mathcal{F}_J(\tau^2) := \{f \in V_J \cap \mathcal{X}(2M) : S_0(f) \leq \tau^2\}\ , \tag{42}$$

and

$$\bar{f} \in \mathrm{argmin}_{f\in V_J\cap\mathcal{X}(2M)}\, S_0(f)\ , \tag{43}$$

Furthermore, we define $\hat{f} = \hat{f}_J \in V_J \cap \mathcal{X}(2M)$, and, for $\sigma > 0$, we also let

$$\hat{f}_s = s\hat{f}_{\text{priv}} + (1-s)\bar{f},\quad \text{where}\quad s = \frac{\sigma}{\sigma + \|\nabla\hat{f}_{\text{priv}} - \nabla\bar{f}\|}\ . \tag{44}$$

Following the same steps as in Proposition 11 from (Hütter & Rigollet, 2021), we obtain

$$\|\nabla \hat{f}_s - \nabla \bar{f}\| = s\|\nabla \hat{f}_{\text{priv}} - \nabla \bar{f}\| = \frac{\sigma\|\nabla \hat{f}_{\text{priv}} - \nabla \bar{f}\|}{\sigma + \|\nabla \hat{f}_{\text{priv}} - \nabla \bar{f}\|} \leq \sigma \tag{45}$$

and

$$S_0(\hat{f}_s) \leq 4M(\sigma^2 + \|\nabla \bar{f} - \nabla f_0\|^2) =: \tau^2 \ . \tag{46}$$

It follows that, $\hat{f}_s \in \mathcal{F}_J(\tau^2)$ and, therefore, $\bar{f} \in \mathcal{F}_J(\tau^2)$ since we have $S_0(\bar{f}) \leq S_0(\bar{f}_s) \leq \tau^2$ by the definition of $\bar{f}$. We now turn to controlling the error of the semi-dual. We have

$$\begin{aligned} \hat{S}(\hat{f}_s) &= s\hat{S}(\hat{f}_{\text{priv}}) + (1-s)\hat{S}(\bar{f}) \\ &\leq s\left(\hat{S}(\hat{f}_{\text{priv}}|X_{1:n}, Y_{1:n}) - \hat{S}(\hat{f}|X_{1:n}, Y_{1:n})\right) + s\hat{S}(\hat{f}) + (1-s)\hat{S}(\bar{f}) \\ &\leq U + \hat{S}(\bar{f}) \end{aligned} \tag{47}$$

where the last inequality comes from the optimality of $\hat{f}$ for the empirical semi-dual problem.

Hence,

$$S_0(\hat{f}_s) \leq S_0(\bar{f}) + U + 2\sup_{f \in \mathcal{F}_J(\tau^2)} |S_0(f) - \hat{S}_0(f)| \ . \tag{48}$$

The rest of the proof follows the same lines as the end of Section 5.4 in (Hütter & Rigollet, 2021), replacing $S_0(\bar{f})$ with $S_0(\bar{f}) + U$ and $\hat{f}$ with $\hat{f}_{\text{priv}}$.

D. Proofs of Section 3

D.1. Proof of Lemma 3.2

This result is folklore in the literature of differential privacy, we include its proof for the sake of completeness. To simplify the notation, let us look at the case where the scaling factor in front of the Laplace random variables in (15) is one. The general case will be obtained by scaling. Let $D \sim D'$ be a pair of neighboring datasets, let $k \in \{1, \ldots, N\}$. We will condition on $(L_{k'})_{k' \neq k}$. All that follow holds almost surely.

$$\mathbb{P}\left(\hat{i}(D) = k|(L_{k'})_{k'\neq k}\right) = \mathbb{P}\left(\cap_{k'\neq k}\left(f_k(D) + L_k < f_{k'}(D) + L_{k'}\right)|(L_{k'})_{k'\neq k}\right) \ . \tag{49}$$

Notice that we took strict inequalities ($<$) because L_k is absolutely continuous w.r.t. Lebesgue's measure conditionally on $(L_{k'})_{k'\neq k}$, and hence, ties occur only with null probability.

Thus,

$$\begin{aligned} \mathbb{P}\left(\hat{i}(D) = k|(L_{k'})_{k'\neq k}\right) &= \mathbb{P}\left(L_k < \min_{k'\neq k}\left(f_{k'}(D) - f_k(D) + L_{k'}\right)\middle|(L_{k'})_{k'\neq k}\right) \\ &= \frac{1}{2}\int_{-\infty}^{\min_{k'\neq k}(f_{k'}(D) - f_k(D) + L_{k'})} e^{-|t|}dt \ . \end{aligned} \tag{50}$$

Hence, two cases occur

Case A: $\min_{k'\neq k}\left(f_{k'}(D) - f_k(D) + L_{k'}\right) \leq 0$

$$\mathbb{P}\left(\hat{i}(D) = k|(L_{k'})_{k'\neq k}\right) = \frac{1}{2}e^{\min_{k'\neq k}(f_{k'}(D) - f_k(D) + L_{k'})} \ . \tag{51}$$

Case B: $\min_{k'\neq k}\left(f_{k'}(D) - f_k(D) + L_{k'}\right) > 0$

$$\mathbb{P}\left(\hat{i}(D) = k|(L_{k'})_{k'\neq k}\right) = 1 - \frac{1}{2}e^{-\min_{k'\neq k}(f_{k'}(D) - f_k(D) + L_{k'})} \ . \tag{52}$$

Now, we can look at the ratio $\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)}$. First, we can notice that

$$\left|\left(\min_{k'\neq k}\left(f_{k'}(D)-f_k(D)+L_{k'}\right)\right)-\left(\min_{k'\neq k}\left(f_{k'}(D')-f_k(D')+L_{k'}\right)\right)\right|\leq 2\Delta\ . \tag{53}$$

Indeed, since the sensitivities of $(f_k)_k$ are uniformly bounded, for any $k'\neq k$,

$$f_{k'}(D)-f_k(D)\leq\left(f_{k'}(D')+\Delta\right)-\left(f_k(D')-\Delta\right)=f_{k'}(D')-f_k(D')+2\Delta\ , \tag{54}$$

and taking the minimum yields half of the result. The other half is given by swapping D and D'.

Then, by using $t(D)$ as a short for $\min_{k'\neq k}\left(f_{k'}(D)-f_k(D)+L_{k'}\right)$, we can look at the following four cases:

Case 1: $t(D),t(D')\leq 0$

$$\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)}\overset{(51)}{=}\frac{\frac{1}{2}e^{t(D)}}{\frac{1}{2}e^{t(D')}}\leq e^{|t(D)-t(D')|}\overset{(53)}{\leq}e^{2\Delta}\ . \tag{55}$$

Case 2: $t(D)\leq 0<t(D')$

$$\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)}\overset{(51)\&(52)}{=}\frac{\frac{1}{2}e^{t(D)}}{1-\frac{1}{2}e^{-t(D')}}=e^{t(D')-t(D)}\frac{\frac{1}{2}e^{t(D)-t(D')}}{e^{-t(D)}-\frac{1}{2}e^{-t(D')-t(D)}}\ . \tag{56}$$

Since $t(D)\leq 0<t(D')$,

$$e^{t(D)-t(D')}\leq 1\ . \tag{57}$$

Furthermore, a simple study of the function $x\mapsto e^{-x}-\frac{1}{2}e^{-t(D')-x}$ on $(-\infty,0]$ shows that it is non-increasing (because $t(D')>0$), and thus that

$$e^{-t(D)}-\frac{1}{2}e^{-t(D')-t(D)}\geq 1-\frac{1}{2}e^{-t(D')}\geq\frac{1}{2}\ . \tag{58}$$

In the end,

$$\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)}\leq e^{t(D')-t(D)}\overset{(53)}{\leq}e^{2\Delta}\ . \tag{59}$$

Case 3: $t(D')\leq 0<t(D)$

$$\begin{aligned}\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)}&\overset{(51)\&(52)}{=}\frac{1-\frac{1}{2}e^{-t(D)}}{\frac{1}{2}e^{t(D')}}=e^{t(D)-t(D')}\frac{e^{-t(D)}-\frac{1}{2}e^{-t(D')-t(D)}}{\frac{1}{2}}\\&\leq e^{t(D)-t(D')}\frac{e^{-t(D)}-\frac{1}{2}e^{-t(D)}}{\frac{1}{2}}=e^{t(D)-t(D')}\frac{\frac{1}{2}e^{-t(D)}}{\frac{1}{2}}\\&\leq e^{t(D)-t(D')}\overset{(53)}{\leq}e^{2\Delta}\ .\end{aligned} \tag{60}$$

Case 4: $t(D), t(D') > 0$

$$\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)} \overset{(52)}{=} \frac{1-\frac{1}{2}e^{-t(D)}}{1-\frac{1}{2}e^{-t(D')}}$$
$$\leq \min\left\{ e^{t(D')-t(D)} \underbrace{\frac{e^{-t(D')}-\frac{1}{2}e^{-t(D)-t(D')}}{e^{-t(D)}-\frac{1}{2}e^{-t(D')-t(D)}}}_{\leq 1 \text{ if } t(D')\geq t(D)}, e^{t(D)-t(D')} \underbrace{\frac{e^{-t(D)}-\frac{1}{2}e^{-2t(D)}}{e^{-t(D')}-\frac{1}{2}e^{-2t(D')}}}_{\leq 1 \text{ if } t(D)\geq t(D')} \right\} \tag{61}$$
$$\leq e^{|t(D)-t(D')|}$$
$$\overset{(53)}{\leq} e^{2\Delta} \ .$$

Thus, we have shown that $\frac{\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)}{\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)} \leq e^{2\Delta}$.

Integrating over $(L_{k'})_{k'\neq k}$ yields

$$\mathbb{E}_{(L_{k'})_{k'\neq k}}\left(\mathbb{P}\left(\hat{i}(D)=k|(L_{k'})_{k'\neq k}\right)\right) \leq e^{2\Delta}\mathbb{E}_{(L_{k'})_{k'\neq k}}\left(\mathbb{P}\left(\hat{i}(D')=k|(L_{k'})_{k'\neq k}\right)\right) \ , \tag{62}$$

which translates to

$$\mathbb{P}\left(\hat{i}(D)=k\right) \leq e^{2\Delta}\mathbb{P}\left(\hat{i}(D')=k\right) \ . \tag{63}$$

Finally, since this is true for any k, $\hat{i}$ is ϵ-DP when $\Delta \leq \frac{\epsilon}{2}$. This concludes the proof by scaling appropriatetly.

D.2. Proof of Lemma 3.3

Let $(X_{1:n}, Y_{1:n}) \sim (X'_{1:n}, Y'_{1:n})$ be two neighboring datasets. By definition, they contain the same data points except for one individual. Suppose first that $Y'_{1:n} = Y'_{1:n}$ and that, for some $i \in [n]$, the datasets $X_{1:n}$ and $X'_{1:n}$ differ in the i-th coordinate only. Then the term $\frac{2\|f\|_\infty}{n}$ is obtained by noting that, by the triangular inequality

$$\left|\hat{S}(f|(X_{1:n}, Y_{1:n})) - \hat{S}(f|(X_{1:n}, Y_{1:n}))\right| = \frac{1}{n}\left|f(X_i) - f(X'_i)\right| \leq \frac{2\|f\|_\infty}{n}.$$

In the opposite case, it holds that $X_{1:n} = X'_{1:n}$ and that the datasets $Y_{1:n}$ and $Y'_{1:n}$ differ in one of the coordinates $i \in [n]$. We first prove that f^* is $|\tilde{\Omega}|$-Lipschitz on $\tilde{\Omega}$.

Let $y_1, y_2 \in \tilde{\Omega}$, let $x \in \tilde{\Omega}$,

$$\begin{aligned} \langle x, y_1\rangle - f(x) &= \langle x, y_1 - y_2\rangle + \langle x, y_2\rangle - f(x) \\ &\overset{\text{Cauchy-Schwarz}}{\leq} \|x\|\|y_1 - y_2\| + \langle x, y_2\rangle - f(x) \\ &\leq |\tilde{\Omega}|\|y_1 - y_2\| + \langle x, y_2\rangle - f(x) \ . \end{aligned} \tag{64}$$

Thus, taking the $\sup$ with respect to x on the left-hand side and on the right-hand side yields

$$f^*(y_1) \leq |\tilde{\Omega}|\|y_1 - y_2\| + f^*(y_2) \ , \tag{65}$$

and since y_1 and y_2 play symmetric roles,

$$f^*(y_2) \leq |\tilde{\Omega}|\|y_1 - y_2\| + f^*(y_1) \ . \tag{66}$$

Hence, $\hat{S}(f|X_{1:n}, Y_{1:n})$ is $\frac{|\tilde{\Omega}|}{n}$-Lipschitz w.r.t. any of the Ys, and therefore, changing the value of any of the Ys can only affect the value of $\hat{S}(f|X_{1:n}, Y_{1:n})$ by at most $\frac{2|\tilde{\Omega}|^2}{n}$ by a triangular inequality.

E. Proofs of Section 4

E.1. Proof of Lemma 4.1

Let $f_1, f_2 \in \mathcal{X}$. We observe that $\|f_1^* - f_2^*\|_\infty \leq \|f_1 - f_2\|_\infty$, since for any $y \in \tilde{\Omega}$,

$$\begin{aligned} f_1^*(y) - f_2^*(y) &=: \sup_x \Big\{ \langle x, y \rangle - f_1(x) \Big\} + \inf_{x'} \Big\{ -\langle x', y \rangle + f_2(x') \Big\} \\ &= \sup_x \inf_{x'} \langle x - x', y \rangle + f_2(x') - f_1(x) \\ &\leq \sup_x \inf_{x'=x} \langle x - x', y \rangle + f_2(x') - f_1(x) \\ &= \sup_x \ f_2(x) - f_1(x). \end{aligned}$$

Thus,

$$\begin{aligned} \left| \hat{S}(f_1 | X_{1:n}, Y_{1:n}) - \hat{S}(f_2 | X_{1:n}, Y_{1:n}) \right| &= \left| \frac{1}{n} \sum_{i=1}^n (f_1(X_i) - f_2(X_i)) + \frac{1}{n} \sum_{i=1}^n (f_1^*(Y_i) - f_2^*(Y_i)) \right| \\ &\leq \|f_1 - f_2\|_\infty + \|f_1^* - f_2^*\|_\infty \\ &\leq 2\|f_1 - f_2\|_\infty \ . \end{aligned} \tag{67}$$

E.2. Proof of Lemma 4.2

By Lemma B.3, we have

$$V_J \cap \mathcal{X}(2M) \subset B_{V_J, \|\cdot\|_{\infty,J}}(0, CM^2) \tag{68}$$

for a positive constant C, where $B_{V_J, \|\cdot\|_{\infty,J}}(0, r)$ refers to the closed ball of V_J with center 0 and radius $r \geq 0$ for the norm $\|\cdot\|_{\infty,J}$ (the infinite norm of the vector of coefficients in the wavelet basis).

Since this is a ball in a finite-dimensional space for a distance in infinite norm, it is possible to build a δ-covering of this ball with a δ-grid on the coefficients of the wavelet decomposition, which is of cardinality

$$C_J := \left(\left\lceil \frac{2CM^2}{\delta} \right\rceil \right)^{\dim(V_J)} . \tag{69}$$

Furthermore, because of the bounded support of the father and mother wavelets, $\dim(V_J) \lesssim 2^{Jd}$ (see page 30 in (Hütter & Rigollet, 2021)). Thus,

$$C_J \leq \left(\left\lceil \frac{2CM^2}{\delta} \right\rceil \right)^{C' 2^{Jd}} \tag{70}$$

where C' is a constant that does not depend on J and δ.

Let us consider the following construction: for each element of this covering, if there exists a point in $V_J \cap \mathcal{X}(2M)$ at distance not more than δ, we replace the original element by this new point (we choose any point that satisfies this condition in case of non uniqueness). Else, we remove this element if no point in $V_J \cap \mathcal{X}(2M)$ satisfies this condition. By the triangular inequality, this construction yields a 2δ covering of $V_J \cap \mathcal{X}(2M)$ for $\|\cdot\|_{\infty,J}$ of cardinality at most C_J.

Furthermore, by Lemma B.3, a $\frac{\delta}{2^{Jd/2}}$-covering of $V_J \cap \mathcal{X}(2M)$ for $\|\cdot\|_{\infty,J}$ is a δ-covering of $V_J \cap \mathcal{X}(2M)$ for the usual infinite norm $\|\cdot\|_\infty$.

Scaling δ appropriately concludes the proof.

E.3. Proof of Theorem 4.3

We will control

$$\hat{S}(\hat{f}_{\text{priv}} | X_{1:n}, Y_{1:n}) - \hat{S}(\hat{f}_J | X_{1:n}, Y_{1:n}) \ . \tag{71}$$

We work conditionally on $X_{1:n}, Y_{1:n}$. All the relations below hold almost surely.

Since $C_{J,M}$ is a δ-covering of $V_J \cap \mathcal{X}(2M)$ with respect to $\|\cdot\|_\infty$, it is also a 2δ-covering of $V_J \cap \mathcal{X}(2M)$ with respect to $\hat{S}(\cdot|X_{1:n}, Y_{1:n})$ by Lemma 4.1. Thus,

$$\begin{aligned}
&\hat{S}(\hat{f}_{\text{priv}}|X_{1:n}, Y_{1:n}) - \hat{S}(\hat{f}_J|X_{1:n}, Y_{1:n}) \\
&\quad= \min_{i\in\{1,\dots,\#(C_{J,M})\}} \left\{ \hat{S}(f_i|D) + \frac{32M^2 \vee 36d}{n\epsilon} L_i \right\} - \hat{S}(\hat{f}_J|X_{1:n}, Y_{1:n}) \\
&\quad\leq \min_{i\in\{1,\dots,\#(C_{J,M})\}} \left\{ \hat{S}(f_i|X_{1:n}, Y_{1:n}) - \hat{S}(\hat{f}_J|X_{1:n}, Y_{1:n}) \right\} + \max_i \left| \frac{32M^2 \vee 36d}{n\epsilon} L_i \right| \\
&\quad\leq 2 \min_{i\in\{1,\dots,\#(C_{J,M})\}} \|f_i - \hat{f}_J\|_\infty + \max_i \left| \frac{32M^2 \vee 36d}{n\epsilon} L_i \right| \\
&\quad\leq \underbrace{2\delta + \max_i \left| \frac{32M^2 \vee 36d}{n\epsilon} L_i \right|}_{=:D} .
\end{aligned} \tag{72}$$

By integrating over the source of privacy and by Lemma B.2,

$$\begin{aligned}
\mathbb{E}_{L_{1:\#\left(C_{J,M}\right)}}(D) &\lesssim \delta + \frac{\ln(\#\left(C_{J,M}\right))}{n\epsilon} \\
&\lesssim \delta + \frac{2^{Jd} \ln\left(\frac{C2^{Jd/2}}{\delta} + 1\right)}{n\epsilon} ,
\end{aligned} \tag{73}$$

where $C \geq 0$ is a constant independent of J and δ. Fixing $\delta = \frac{C2^{Jd/2}}{n\epsilon}$ yields

$$\mathbb{E}_{L_{1:\#\left(C_{J,M}\right)}}(D) \lesssim \frac{2^{Jd} \ln(n\epsilon)}{n\epsilon} . \tag{74}$$

By Lemma 2.5 and subsection 5.5 in (Hütter & Rigollet, 2021) in order to control the bias,

$$\mathbb{E}\left(d(\hat{T}_{\text{priv}}, T_0)^2\right) \lesssim \frac{2^{Jd}\ln(n\epsilon)}{n\epsilon} + \frac{J2^{J(d-2)}\ln(n)}{n} + \frac{1}{n} + R^2 2^{-2J\alpha} \tag{75}$$

for J larger than a suitable constant.

F. Proofs of Section 5

F.1. Proof of Theorem 5.1

The proof of the lower bounding terms $\frac{1}{n}$ and $n^{-\frac{2\alpha}{2\alpha-2+d}}$ uses standard packing arguments (Tsybakov, 2009) with a specific packing for the problem at hand. It may be found in (Hütter & Rigollet, 2021). We only provide the proof for the additional term $(n\epsilon)^{-\frac{2\alpha}{2\alpha+d}}$. It is based on packing arguments for differential privacy standardized via couplings in (Acharya et al., 2018; 2021; Lalanne et al., 2023a). We start by building a packing similar to the one from (Hütter & Rigollet, 2021), which is later analyzed with TV distances instead of the traditional KL divergences in order to highlight the effect of differential privacy.

Packing construction. Let m be an integer that will be specified later in the proof and let $N = m^d$. We consider the grid $\left(\left(\frac{k_1}{m+1}, \dots, \frac{k_d}{m+1}\right)\right)_{1 \leq k_1, \dots, k_d \leq m}$. This grid has cardinality N, therefore we can enumerate its elements as $(p_1, \dots, p_N) \in (\mathbb{R}^d)^N$. By construction, we have that

$$\forall i, j \in \{1, \dots, N\}, \quad i \neq j \implies \|p_i - p_j\|_\infty \geq \frac{1}{m+1} \, . \tag{76}$$

Now, we consider a C^∞ "bump" function B defined over $\mathbb{R}$ and supported on $[-1, 1]$, taking positive values over $(-1, 1)$. For such a function, there exists $x_0 \in (-1, 1)$ such that $B(x_0) \neq 0$ and $B'(x_0) \neq 0$.

For some sufficiently small constant $a > 0$ whose value will be fixed later, we define the function

$$\psi(x_1, \dots, x_d) = a \prod_{i=1}^{d} B\left(\frac{x_i}{2}\right) . \tag{77}$$

and for some $h \in (0, 1]$ and any $\theta \in \{0, 1\}^N$, we define the potential $\phi_\theta : \tilde{\Omega} \to \mathbb{R}$ as

$$\phi_\theta(\cdot) := \frac{1}{2}\| \cdot \|^2 + h^{\alpha+1} \sum_{i=1}^{N} \theta_i \psi\left(\frac{\cdot - p_i}{h}\right) \, . \tag{78}$$

To ensure that the functions $\psi\left(\frac{\cdot - p_i}{h}\right)$ for $i \in \{1, \dots, N\}$ have disjoint supports, we take $h < \frac{1}{2(m+1)} \lesssim \frac{1}{m}$.

First, we show that for any $\theta \in \{0, 1\}^N$, $\nabla\phi_\theta \in \mathcal{T}_\alpha$:

- (Bounded transport map) $\forall \theta \in \{0, 1\}^N, \forall x \in \tilde{\Omega}$,

$$\nabla\phi_\theta(x) = x + h^\alpha \sum_{i=1}^{N} \theta_i \nabla\psi\left(\frac{x - p_i}{h}\right) \, , \tag{79}$$

  and since the supports are disjoint, only one term in the sum is non-zero, which yields

$$\forall \theta \in \{0, 1\}^N, \forall x \in \tilde{\Omega}, \quad \|\nabla\phi_\theta(x) - x\| \leq h^\alpha \|\nabla\psi\|_\infty . \tag{80}$$

  Choosing a small enough ensures that $\forall \theta \in \{0, 1\}^N, \forall x \in \tilde{\Omega}, \|T(x)\| \leq M$.

- (Strong convexity of the potential) $\forall \theta \in \{0, 1\}^N, \forall x \in \tilde{\Omega}$,

$$\nabla^2\phi_\theta(x) = I_d + h^{\alpha-1} \sum_{i=1}^{N} \theta_i \nabla^2\psi\left(\frac{x - p_i}{h}\right) \tag{81}$$

  and thus, since the supports are disjoint,

$$\forall \theta \in \{0, 1\}^N, \forall x \in \tilde{\Omega}, \quad \|\nabla^2\phi_\theta(x) - I_d\|_{\text{op}} \leq h^{\alpha-1} \sup \|\nabla^2\psi\|_{\text{op}} \, , \tag{82}$$

  where the right-hand side can be made arbitrarily close to 0 by taking a small enough (since $h^{\alpha-1} \leq 1$), which ensures that $\forall \theta \in \{0, 1\}^N, \forall x \in \tilde{\Omega}, \frac{1}{M} \preceq \nabla^2\phi_\theta(x) \preceq M$.

- (Stability of the support) Let $\theta \in \{0,1\}^N$. It follows from (80) that, when $h \leq \frac{1}{m+1}$, one can take a small enough to ensure that

$$\forall i, \nabla\phi_\theta \left(B_\infty\left(p_i, \frac{h}{2}\right)\right) \subset B_\infty\left(p_i, h\right) \ . \tag{83}$$

Furthermore, since outside the balls of the form $\left(B_\infty\left(p_i, \frac{h}{2}\right)\right)$, $\nabla\phi_\theta$ is the identity, we have that $\nabla\phi_\theta(\Omega) \subset \Omega$ and that

$$\forall i, \nabla\phi_\theta\left(B_\infty\left(p_i, h\right)\right) \subset B_\infty\left(p_i, h\right) \ . \tag{84}$$

We now give additional details on the properties of $\nabla\phi_\theta$. We have proved above that ϕ_θ is strongly convex for any $\theta \in \{0,1\}^N$. It follows that $\nabla\phi_\theta$ is injective and that, for any $x \in \tilde{\Omega}$, the Hessian $\nabla^2\phi_\theta(x) \succ 0$ is invertible. Since $\nabla\phi_\theta$ is $\mathcal{C}^\infty$ over $\tilde{\Omega}$, the global inversion theorem states that $\nabla\phi_\theta$ is a $\mathcal{C}^\infty$ diffeomorphism from $\tilde{\Omega}$ to $\nabla\phi_\theta(\tilde{\Omega})$.

We may also see that the inclusion $\nabla\phi_\theta(\Omega) \subset \Omega$ is in fact an equality using arguments borrowed from basic algebraic topology. First, $\nabla\phi_\theta$ is the identity over $\partial\Omega$. Suppose for the sake of contradiction that $\Omega \setminus \nabla\phi_\theta(\Omega) \neq \emptyset$ and let $x \in \Omega \setminus \nabla\phi_\theta(\Omega)$ ($x \notin \partial\Omega$ since $\nabla\phi_\theta$ is the identity over $\partial\Omega$). Then, letting r_x be a retraction of $\Omega \setminus \{x\}$ on $\partial\Omega$ (which exists), we obtain that $r_x \circ \nabla\phi_\theta$ is a retraction of Ω (which is homeomorphic to $B(0,1)$) to $\partial\Omega$ (which is homeomorphic to $\partial B(0,1)$), which is impossible by Brouwer's theorem (see Theorem 7.6.2 in (Randal-Williams, 2024)).

As a consequence,

$$\forall\theta \in \{0,1\}^N, \quad \nabla\phi_\theta(\Omega) = \Omega \ . \tag{85}$$

By injectivity of the functions ϕ_θ's, this immediately yields

$$\forall\theta \in \{0,1\}^N, \quad \left(\nabla\phi_\theta\right)^{-1}(\Omega) = \Omega \ . \tag{86}$$

Finally, the same arguments applied to the supports of the functions $\psi\left(\frac{\cdot - p_i}{h}\right)$ ensure that, for a small enough $a > 0$, we have

$$\forall\theta \in \{0,1\}^N, \forall i, \quad \nabla\phi_\theta\left(B_\infty\left(p_i, h\right)\right) = B_\infty\left(p_i, h\right) \ . \tag{87}$$

By injectivity of the functions ϕ_θ's, we similarly obtain

$$\forall\theta \in \{0,1\}^N, \forall i, \quad \left(\nabla\phi_\theta\right)^{-1}\left(B_\infty\left(p_i, h\right)\right) = B_\infty\left(p_i, h\right) \ . \tag{88}$$

- (Smoothness) Let $\theta \in \{0,1\}^N$. Using multi-index notation (see (Hütter & Rigollet, 2021)),

$$\|\nabla\phi_\theta\|_{C^\alpha(\tilde{\Omega})} = \sum_{i=1}^{d}\left(\sum_{|b|\leq\alpha}\underbrace{\|\partial^b\left(\nabla\phi_\theta\right)_i\|_\infty}_{=:A_{i,b}} + \sum_{|b|=\lfloor\alpha\rfloor}\underbrace{\sup_{x\neq y, x,y\in\tilde{\Omega}}\frac{\left|\partial^b\left(\nabla\phi_\theta\right)_i(x) - \partial^b\left(\nabla\phi_\theta\right)_i(y)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}}}_{=:B_{i,b}}\right) \ . \tag{89}$$

Furthermore,

$$A_{i,b} = \left\|\partial^b \,\text{Coord.}_i + h^{\alpha-|b|}\sum_{j=1}^{N}\theta_j\partial^b(\nabla\psi)_i\left(\frac{\cdot - p_j}{h}\right)\right\|_\infty \tag{90}$$

where Coord.$_i$ is the projection along the i^{th} canonical vector. Thus, $A_{i,b}$ can be made as close as we want to $\left\|\partial^b \,\text{Coord.}_i\right\|_\infty$ granted that the constant a is taken small enough.

Finally,

$$
\begin{aligned}
B_{i,b} = &\sup_{x\neq y, x,y\in\tilde{\Omega}} \\
&\frac{\left|\partial^b \text{ Coord.}_i(x) + h^{\alpha-|b|}\sum_{j=1}^N \theta_j \partial^b(\nabla\psi)_i\left(\frac{x-p_j}{h}\right) - \partial^b \text{ Coord.}_i(y) + h^{\alpha-|b|}\sum_{j=1}^N \theta_j \partial^b(\nabla\psi)_i\left(\frac{y-p_j}{h}\right)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} \\
&= \sup_{x\neq y, x,y\in\tilde{\Omega}} \frac{\left|h^{\alpha-|b|}\sum_{j=1}^N \theta_j \partial^b(\nabla\psi)_i\left(\frac{x-p_j}{h}\right) - h^{\alpha-|b|}\sum_{j=1}^N \theta_j \partial^b(\nabla\psi)_i\left(\frac{y-p_j}{h}\right)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} \\
&\leq h^{\alpha-|b|} \sup_{x\neq y, x,y\in\tilde{\Omega}} \frac{\sum_{j=1}^N \left|\partial^b(\nabla\psi)_i\left(\frac{x-p_j}{h}\right) - \partial^b(\nabla\psi)_i\left(\frac{y-p_j}{h}\right)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} \\
&\overset{\text{Disjoint Supports \& Continuity}}{\leq} h^{\alpha-|b|} \max_i \sup_{x\neq y, x,y\in\tilde{\Omega}} \frac{\left|\partial^b(\nabla\psi)_i\left(\frac{x-p_j}{h}\right) - \partial^b(\nabla\psi)_i\left(\frac{y-p_j}{h}\right)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} \\
&\leq h^{\alpha-|b|} \max_i \sup_{x\neq y, x,y\in\mathbb{R}^d} \frac{\left|\partial^b(\nabla\psi)_i\left(\frac{x-p_j}{h}\right) - \partial^b(\nabla\psi)_i\left(\frac{y-p_j}{h}\right)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} \\
&= h^{\alpha-|b|} \max_i h^{-(\alpha-\lfloor\alpha\rfloor)} \sup_{x\neq y, x,y\in\mathbb{R}^d} \frac{\left|\partial^b(\nabla\psi)_i(x) - \partial^b(\nabla\psi)_i(y)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} \\
&= \max_i \sup_{x\neq y, x,y\in\mathbb{R}^d} \frac{\left|\partial^b(\nabla\psi)_i(x) - \partial^b(\nabla\psi)_i(y)\right|}{|x-y|^{\alpha-\lfloor\alpha\rfloor}} .
\end{aligned} \tag{91}
$$

Hence, $B_{i,b}$ can be made as small as we want to 0 granted that the constant a is taken small enough.

All in all, if the constant a is taken small enough, $\sup_\theta \|\nabla\phi_\theta\|_{C^\alpha(\tilde{\Omega})}$ can be made as close as we want to $\|$ Identity $\|_{C^\alpha(\tilde{\Omega})}$, which guarantees that $\sup_\theta \|\nabla\phi_\theta\|_{C^\alpha(\tilde{\Omega})} \leq R$ if a is small enough.

Similarly as in (Hütter & Rigollet, 2021), we consider the family of statistics $(S_\theta := P^{\otimes n} \otimes Q_\theta^{\otimes n})_{\theta\in\{0,1\}^N}$ on the problem where

$$
P = \text{Unif}([0,1]^d) , \tag{92}
$$

and

$$
\forall \theta \in \{0,1\}^N, \quad Q_\theta = P_{\#\nabla\phi_\theta} . \tag{93}
$$

It is thus possible to simplify the lower-bounding problem as

$$
\inf_{\hat{T}} \sup_{P\in\mathcal{M}, T_0\in\mathcal{T}_\alpha} \mathbb{E}\left(\int \|T_0 - \hat{T}\|^2 dP\right) \geq \inf_{\hat{T}} \sup_{\theta\in\{0,1\}^N} \mathbb{E}_{S_\theta}\left(\int \|\nabla\phi_\theta - \hat{T}\|^2 dP\right) \tag{94}
$$

where the supremum is now taken over a finite family only.

We verify that this family of statistics forms a packing for the pseudo metric defined below (with some overlap in the notation)

$$
d(S_{\theta_1}, S_{\theta_2})^2 := d(\theta_1, \theta_2)^2 := \int \|\nabla\phi_{\theta_1} - \nabla\phi_{\theta_2}\|^2 dP \tag{95}
$$

Lemma F.1. *For any* $\theta_1, \theta_2 \in \{0,1\}^N$,

$$
d(S_{\theta_1}, S_{\theta_2})^2 \gtrsim d_{\text{ham}}(\theta_1, \theta_2)\, h^{2\alpha+d} . \tag{96}
$$

Proof. Let $\theta_1, \theta_2 \in \{0,1\}^N$, then

$$\begin{aligned} \int_{[0,1]^d} \|\nabla\phi_{\theta_1} - \nabla\phi_{\theta_2}\|^2 &\overset{\text{Disjoint supports}}{=} \sum_{i=1}^N \int_{[0,1]^d} \left\| (\theta_1^i - \theta_2^i) h^\alpha \nabla\psi\left(\frac{x-p_i}{h}\right)\right\|^2 dx \\ &\overset{\text{Change of variables}}{=} d_{\text{ham}}(\theta_1,\theta_2)\, h^{2\alpha+d} \underbrace{\int \|\nabla\psi\|^2}_{>0} \\ &\gtrsim d_{\text{ham}}(\theta_1,\theta_2)\, h^{2\alpha+d} \ . \end{aligned} \tag{97}$$

□

From line 1 to line 2, we used that $|\theta_1^i - \theta_2^i|$ is equal to 1 only $d_{\text{ham}}(\theta_1,\theta_2)$ times, and is it equal to 0 in the other cases. The change of variable $u = \frac{x-p_i}{h}$ in each of the remaining terms yields the term h^d.

Upper bounding the TV distances. To proceed with the lower-bound argument, it remains to lower-bound the testing difficulty between the hypotheses of this packing. Techniques for bounding the testing difficulties usually require bounding the KL-divergences between the statistical models of the packing and a reference measure. This is the approach proposed in (Hütter & Rigollet, 2021), which leads to a lower bound of the order of $\frac{1}{n} \vee n^{-\frac{2\alpha}{2\alpha-2+d}}$. However, under differential privacy, the test difficulty is characterized by the TV distances between the *untensorized* marginals of the statistical models. This part of the proof is dedicated to controlling these terms.

Lemma F.2. *For any $\theta_1, \theta_2 \in \{0,1\}^N$, it holds that*

$$\text{TV}(Q_{\theta_1}, Q_{\theta_2}) \lesssim d_{\text{ham}}(\theta_1,\theta_2)\, h^{\alpha-1+d} \ . \tag{98}$$

Proof. By the change of variables formula, for any $\theta \in \{0,1\}^N$, the density of Q_θ with respect to P is

$$\frac{dQ_\theta}{dP}(y) = \frac{\mathbb{1}_{[0,1]^d}((\nabla\phi_\theta)^{-1}(y))}{\det(\nabla^2\phi_\theta((\nabla\phi_\theta)^{-1}(y)))} \tag{99}$$

for P-almost all y. Let $\theta_1, \theta_2 \in \{0,1\}^N$,

$$\begin{aligned} \text{TV}(Q_{\theta_1}, Q_{\theta_2}) &= \frac{1}{2}\int_{(-1,2)^d} \left| \frac{dQ_{\theta_1}}{dP}(y) - \frac{dQ_{\theta_2}}{dP}(y)\right| dP(y) \\ &= \frac{1}{2}\int_{[0,1]^d} \left| \frac{\mathbb{1}_{[0,1]^d}((\nabla\phi_{\theta_1})^{-1}(y))}{\det(\nabla^2\phi_{\theta_1}((\nabla\phi_{\theta_1})^{-1}(y)))} - \frac{\mathbb{1}_{[0,1]^d}((\nabla\phi_{\theta_2})^{-1}(y))}{\det(\nabla^2\phi_{\theta_2}((\nabla\phi_{\theta_2})^{-1}(y)))}\right| dy \ , \end{aligned} \tag{100}$$

and by stability of $[0,1]^d$ and of its complement by either $\nabla\phi_{\theta_1}$ or $\nabla\phi_{\theta_2}$, this expression can be further simplified as

$$\text{TV}(Q_{\theta_1}, Q_{\theta_2}) = \frac{1}{2}\int\limits_{[0,1]^d} \left| \frac{1}{\det(\nabla^2\phi_{\theta_1}((\nabla\phi_{\theta_1})^{-1}(y)))} - \frac{1}{\det(\nabla^2\phi_{\theta_2}((\nabla\phi_{\theta_2})^{-1}(y)))}\right| dy \ . \tag{101}$$

Then, since both $\nabla\phi_{\theta_1}$ and $\nabla\phi_{\theta_2}$ are the identity outside balls of the form $B_\infty(p_i,h)$, and since these balls and their complements are stable by $\nabla\phi_{\theta_1}$ and by $\nabla\phi_{\theta_2}$, this can be rewritten as

$$\text{TV}(Q_{\theta_1}, Q_{\theta_2}) = \frac{1}{2}\sum_{i=1}^N \int\limits_{B_\infty(p_i,h)} \left| \frac{1}{\det(\nabla^2\phi_{\theta_1}((\nabla\phi_{\theta_1})^{-1}(y)))} - \frac{1}{\det(\nabla^2\phi_{\theta_2}((\nabla\phi_{\theta_2})^{-1}(y)))}\right| dy \ . \tag{102}$$

This expression finally further simplifies as

$$\begin{aligned} &\text{TV}(Q_{\theta_1}, Q_{\theta_2}) = \\ &\quad \frac{1}{2}\sum_{i=1}^N \mathbb{1}_{\theta_1^{(i)} \neq \theta_2^{(i)}} \underbrace{\int\limits_{B_\infty(p_i,h)} \left| \frac{1}{\det(\nabla^2\phi_{\theta_1}((\nabla\phi_{\theta_1})^{-1}(y)))} - \frac{1}{\det(\nabla^2\phi_{\theta_2}((\nabla\phi_{\theta_2})^{-1}(y)))}\right| dy}_{:=\Delta_i} \ . \end{aligned} \tag{103}$$

We treat the case $\theta_1^{(i)} = 0$, $\theta_2^{(i)} = 1$, the other case being treated similarly. In this case,

$$\begin{aligned}\Delta_i &= \int_{B_\infty(p_i,h)} \left| 1 - \frac{1}{\det(\nabla^2 \phi_{\theta_2}((\nabla\phi_{\theta_2})^{-1}(y)))} \right| dy \\ &= \int_{B_\infty(p_i,h)} \left| 1 - \frac{1}{\det(\nabla^2 \phi_{\theta_2}(x))} \right| |\det(\nabla^2\phi_{\theta_2}(x))| dx \ ,\end{aligned} \tag{104}$$

where the last line comes from a change of variables and from the stability of $B_\infty(p_i, h)$ and of its complement.

For any $x \in B_\infty(p_i, h)$,

$$\nabla^2\phi_{\theta_2}(x) = I_d + h^{\alpha-1}\nabla^2\psi\left(\frac{x-p_i}{h}\right) \ . \tag{105}$$

By Lemma B.1, we see that if a is chosen small enough, then we have $|\det(\nabla^2\phi_{\theta_2}(\cdot))| \leq 2$ uniformly over $B_\infty(p_i, h)$. Finally, and again by Lemma B.1, if a is chosen small enough, $\forall x \in B_\infty(p_i, h)$,

$$\left|1 - \frac{1}{\det(\nabla^2\phi_{\theta_2}(x))}\right| \leq h^{\alpha-1}\left(\left|\mathrm{tr}\left(\nabla^2\psi\left(\frac{x-p_i}{h}\right)\right)\right| + \left\|\nabla^2\psi\left(\frac{x-p_i}{h}\right)\right\|\right) \ , \tag{106}$$

which entails, by a change of variables,

$$\Delta_i \leq \frac{h^{\alpha-1+d}}{2}\int \left(\left|\mathrm{tr}\left(\nabla^2\psi\right)\right| + \left\|\nabla^2\psi\right\|\right) \lesssim h^{\alpha-1+d} \ . \tag{107}$$

Plugging this result back into (103) yields

$$\begin{aligned}\mathrm{TV}\left(Q_{\theta_1}, Q_{\theta_2}\right) &\lesssim h^{\alpha-1+d}\sum_{i=1}^{N}\mathbb{1}_{\theta_1^{(i)}\neq\theta_2^{(i)}} \\ &= d_{\mathrm{ham}}\left(\theta_1, \theta_2\right) h^{\alpha-1+d} \ .\end{aligned} \tag{108}$$

□

Private distributional tests and private Assouad method.

Lemma F.3 (Assouad's Lemma)**.** *Assume that there exists $\tau > 0$ such that for any $\theta_1, \theta_2 \in \{0,1\}^N$,*

$$d(S_{\theta_1}, S_{\theta_2})^2 \gtrsim d_{\mathrm{ham}}\left(\theta_1, \theta_2\right)\tau. \tag{109}$$

Let $\hat{T}$ be any estimator and define $\hat{\theta} = \mathrm{argmin}_{\theta\in\{0,1\}^N}\, d(\hat{T}, \nabla\phi_\theta)$. Then it holds that

$$\sup_{\theta\in\{0,1\}^N}\mathbb{E}_{S_\theta}\left(\int \|\nabla\phi_\theta - \hat{T}\|^2 dP\right) \gtrsim \tau\sum_{i=1}^{N}\left(\mathbb{P}_{\theta_{-i}}(\hat{\theta}^i \neq 0) + \mathbb{P}_{\theta_{+i}}(\hat{\theta}^i \neq 1)\right), \tag{110}$$

where $\mathbb{P}_{\theta_{+i}}$ and $\mathbb{P}_{\theta_{-i}}$ are the mixture distributions

$$\mathbb{P}_{\theta_{+i}} := \frac{1}{2^{N-1}}\sum_{\theta:\theta^{(i)}=1} S_\theta, \quad \textit{and} \quad \mathbb{P}_{\theta_{-i}} := \frac{1}{2^{N-1}}\sum_{\theta:\theta^{(i)}=0} S_\theta \ . \tag{111}$$

Note that in (110) there is a second layer or randomness that is implicit, and that is w.r.t. the estimator itself (for privacy for instance).

Proof. The proof can be directly adapted from (Acharya et al., 2021) by substituting the notation with that used in the present paper. □

Lemma F.4. *If $\hat{T}$ satisfies ϵ-DP, then for any i,*

$$\begin{aligned}&\mathbb{P}_{\theta_{-i}}(\hat{\theta}^i \neq 0) + \mathbb{P}_{\theta_{+i}}(\hat{\theta}^i \neq 1) \geq \\ &\frac{1}{2^{N-1}} \sum_{\theta^1,\dots,\theta^{i-1},\theta^{i+1}\dots,\theta^N \in\{0,1\}} e^{-n\epsilon \mathrm{TV}\left(Q_{(\theta^1,\dots,\theta^{i-1},0,\theta^{i+1}\dots,\theta^N)},Q_{(\theta^1,\dots,\theta^{i-1},1,\theta^{i+1}\dots,\theta^N)}\right)} ,\end{aligned} \tag{112}$$

where $\hat{\theta}$ is built from $\hat{T}$ as in Lemma F.3.

Proof. This proof builds on coupling arguments from (Acharya et al., 2018; 2021; Lalanne et al., 2023a). Let us consider the coupling $\mathcal{C}$ that selects $\theta^1, \dots, \theta^{i-1}, \theta^{i+1} \dots, \theta^N \in \{0, 1\}$ uniformly at random, and then samples accordingly

$$X_1, \dots, X_n \sim P \tag{113}$$

and

$$(Y_1, Y_1'), \dots, (Y_n, Y_n') \sim Q \tag{114}$$

where Q is the optimal coupling between $Q_{(\theta^1,\dots,\theta^{i-1},0,\theta^{i+1}\dots,\theta^N)}$ and $Q_{(\theta^1,\dots,\theta^{i-1},1,\theta^{i+1}\dots,\theta^N)}$. Here, the optimality means that if $(Y, Y') \sim Q$, then

$$\mathbb{P}(Y = Y') = 1 - \mathrm{TV}\left(Q_{(\theta^1,\dots,\theta^{i-1},0,\theta^{i+1}\dots,\theta^N)}, Q_{(\theta^1,\dots,\theta^{i-1},1,\theta^{i+1}\dots,\theta^N)}\right) . \tag{115}$$

The existence of such a coupling is well known (see, *e.g.* (Kallenberg, 1993)). Furthermore, we let $X_1', \dots, X_n'$ be copies of $X_1, \dots, X_n$.

Then $((X_1, \dots, X_n, Y_1, \dots, Y_n), (X_1', \dots, X_n', Y_1', \dots, Y_n'))$ (whose distribution is $\mathcal{C}$) is a coupling between $\mathbb{P}_{\theta_{+i}}$ and $\mathbb{P}_{\theta_{-i}}$.

Thus, we may write

$$\begin{aligned}&\mathbb{P}_{\theta_{-i}}(\hat{\theta}^i \neq 0) + \mathbb{P}_{\theta_{+i}}(\hat{\theta}^i \neq 1) = \\ &\quad\mathbb{E}_{\mathcal{C}}\left(\mathbb{P}(\hat{\theta}^i \neq 0 | X_1, \dots, X_n, Y_1, \dots, Y_n) + \mathbb{P}(\hat{\theta}^i \neq 1 | X_1', \dots, X_n', Y_1', \dots, Y_n')\right)\end{aligned} \tag{116}$$

where the inner randomness (i.e. inside the expectation) is only w.r.t. $\hat{\theta}^i$. Since $\hat{T}$ is ϵ-DP, then so is $\hat{\theta}^i$ by post-processing.

Thus, by group privacy and the characteristic property of differential privacy,

$$
\begin{aligned}
&\mathbb{P}_{\theta_{-i}}(\hat{\theta}^i \neq 0) + \mathbb{P}_{\theta_{+i}}(\hat{\theta}^i \neq 1) = \\
&\quad \mathbb{E}_{\mathcal{C}}\Big(\mathbb{P}(\hat{\theta}^i((X_1,\dots,X_n,Y_1,\dots,Y_n)) \neq 0) + \mathbb{P}(\hat{\theta}^i((X_1',\dots,X_n',Y_1',\dots,Y_n')) \neq 1)\Big) \\
&\quad \geq \mathbb{E}_{\mathcal{C}}\Bigg(e^{-\epsilon d_{\mathrm{ham}}\left((X_1,\dots,X_n,Y_1,\dots,Y_n),(X_1',\dots,X_n',Y_1',\dots,Y_n')\right)}\mathbb{P}(\hat{\theta}^i((X_1',\dots,X_n',Y_1',\dots,Y_n')) \neq 0) \\
&\qquad\qquad + \mathbb{P}(\hat{\theta}^i((X_1',\dots,X_n',Y_1',\dots,Y_n')) \neq 1)\Bigg) \\
&\quad \geq \mathbb{E}_{\mathcal{C}}\Bigg(e^{-\epsilon d_{\mathrm{ham}}\left((X_1,\dots,X_n,Y_1,\dots,Y_n),(X_1',\dots,X_n',Y_1',\dots,Y_n')\right)} \\
&\qquad\qquad \underbrace{\Big(\mathbb{P}(\hat{\theta}^i((X_1',\dots,X_n',Y_1',\dots,Y_n')) \neq 0) + \mathbb{P}(\hat{\theta}^i((X_1',\dots,X_n',Y_1',\dots,Y_n')) \neq 1)\Big)}_{=1}\Bigg) \\
&\quad = \mathbb{E}_{\mathcal{C}}\left(e^{-\epsilon d_{\mathrm{ham}}\left((X_1,\dots,X_n,Y_1,\dots,Y_n),(X_1',\dots,X_n',Y_1',\dots,Y_n')\right)}\right) \\
&\quad = \mathbb{E}_{\theta^1,\dots,\theta^{i-1},\theta^{i+1}\dots,\theta^N}\Bigg(\\
&\qquad\qquad \mathbb{E}_{(X_1,\dots,X_n,Y_1,\dots,Y_n)(X_1',\dots,X_n',Y_1',\dots,Y_n')}\left(e^{-\epsilon d_{\mathrm{ham}}\left((X_1,\dots,X_n,Y_1,\dots,Y_n),(X_1',\dots,X_n',Y_1',\dots,Y_n')\right)}\right)\Bigg) \\
&\quad \overset{\text{Jensen}}{\geq} \mathbb{E}_{\theta^1,\dots,\theta^{i-1},\theta^{i+1}\dots,\theta^N}\Bigg(\\
&\qquad\qquad e^{-\epsilon\mathbb{E}_{(X_1,\dots,X_n,Y_1,\dots,Y_n)(X_1',\dots,X_n',Y_1',\dots,Y_n')}\left(d_{\mathrm{ham}}\left((X_1,\dots,X_n,Y_1,\dots,Y_n),(X_1',\dots,X_n',Y_1',\dots,Y_n')\right)\right)}\Bigg) \\
&\quad = \frac{1}{2^{N-1}} \sum_{\theta^1,\dots,\theta^{i-1},\theta^{i+1}\dots,\theta^N\in\{0,1\}} e^{-n\epsilon\mathrm{TV}\left(Q_{(\theta^1,\dots,\theta^{i-1},0,\theta^{i+1}\dots,\theta^N)},Q_{(\theta^1,\dots,\theta^{i-1},1,\theta^{i+1}\dots,\theta^N)}\right)},
\end{aligned}
\tag{117}
$$

where the last line comes from a simple computation. □

Conclusion of the proof. As a consequence of Lemma F.1, Lemma F.2, Lemma F.3 and Lemma F.4,

$$
\inf_{\hat{T}} \sup_{\theta\in\{0,1\}^N} \mathbb{E}_{S_\theta}\left(\int \|\nabla\phi_\theta - \hat{T}\|^2 dP\right) \gtrsim N h^{2\alpha+d} e^{-Cn\epsilon h^{\alpha-1+d}}, \tag{118}
$$

where $C > 0$ is a positive constant. Finally, since $N \asymp h^{-d}$, taking $h \asymp (n\epsilon)^{-\frac{1}{\alpha-1+d}}$ yields

$$
\inf_{\hat{T}} \sup_{\theta\in\{0,1\}^N} \mathbb{E}_{S_\theta}\left(\int \|\nabla\phi_\theta - \hat{T}\|^2 dP\right) \gtrsim (n\epsilon)^{-\frac{2\alpha}{\alpha-1+d}}, \tag{119}
$$

and (94) then concludes the proof.

G. Practical algorithm

This section provides further details on how to obtain a numerically tractable algorithm derived from the ideas presented in Section 3 and Section 4.

G.1. Choosing J

First, we determine J to achieve the best bias-variance tradeoff. Using Equation (24), we set

$$
J^* = \min\left(\left\lceil \frac{\log_2(n)}{d-2+2\alpha}\right\rceil, \left\lceil \frac{\log_2(n\epsilon)}{d+2\alpha}\right\rceil\right).
$$

This choice asymptotically yields an error

$$\lesssim J^* R^2(\log_2(n) + \log_2(n\epsilon)) \times \text{Upper-Bound Equation (25)},$$

where we explicitly track the scaling in R and retain logarithmic terms.

G.2. Constructing the Covering & Enforcing Conditions

With J^* fixed, we construct the covering. From Equation (39), controlling the infinite functional norm by the vector's infinite norm requires computing $\sqrt{\mathrm{Vol}(\tilde{\Omega})}$. In the paper, we assume $\mathrm{Vol}(\tilde{\Omega}) = 3^d$, but since $\tilde{\Omega}$ only needs to be a hypercube containing Ω in its interior, we can instead set

$$\mathrm{Vol}(\tilde{\Omega}) = (1+\gamma)^d$$

for any fixed $\gamma > 0$.

In Lemma 4.2 and Theorem 4.3, we first construct a $\delta = C/(n\epsilon)$ covering in vector infinite norm for the space of wavelet coefficients up to J^*, within a radius CM^2, where

$$C = 2\sqrt{\mathrm{Vol}(\tilde{\Omega})}.$$

This covering is straightforward since it requires only axis-wise discretization in the space of wavelets coefficients.

The challenging part follows: the conditions on Hessian eigenvalues and the potential's norm/gradient may not hold. We enforce them via the reasoning in the end of Appendix E.2. While this step is theoretically sound for obtaining the upper bound, we believe it is intractable in practice.

To address this, we can approximate using a grid, as in Section 6. This does not affect the privacy analysis but may introduce numerical errors that vanish as the discretization step decreases. Candidate potentials are now represented by their grid discretization and live in the span of wavelet discretizations up to level J^*.

For each candidate in the initial covering, we check whether a nearby potential (within distance δ) satisfies the conditions on the Hessian, gradient, and potential itself, following Definition 2.3 and the end of Appendix E.2 in a discretized manner. At this stage, existence testing reduces to a convex optimization problem—minimizing the infinite norm of a vector under linear constraints, since numerical differentiation schemes are linear operators. A numerical solver can efficiently handle this step. The set $C_{J,M}$ is thus constructed using this quasi-projection approach (see Appendix E.2) applied to each candidate potential from the initial covering.

G.3. Conclusion

This concludes the practical implementation of our estimator. While computationally expensive, it was designed for theoretical purposes. However, we emphasize that the selection rule in Section 3 remains valid for any set of convex potentials with bounded infinite norm—ensuring meaningfulness in optimal transport theory and privacy analysis.

The complexity mainly lies in generating candidate potentials (which is not due to privacy but rather to the difficulty of the non-parametric nature of the potential estimation as in (Hütter & Rigollet, 2021)), whereas the selection procedure itself is more practical. Thus, with a realistic potential generator (e.g., leveraging prior knowledge), our private selector should yield strong utility in practice.

H. Potential alternative approaches

H.1. Plugin estimators

An alternative approach could be to estimate privately both the input and output distributions, and to consider the optimal transport between those estimates. Such a technique is usually referred to as a *plugin* estimator.

For our problem, it typically requires additional assumptions on the regularity of the source or the target distributions (Manole et al., 2024). We are confident that such a technique could lead to good results under this more restrictive setup, yet it would require extra work as it requires stability of the estimates in W_2 distance, whereas most results on density estimation under differential privacy are stated either in terms of infinite or L_2 norms (Wasserman & Zhou, 2010; Barber & Duchi, 2014; Lalanne et al., 2023b; Lalanne & Gadat, 2024).

H.2. Gradient-based approaches

The semi-dual problem is a finite-dimensional convex optimization problem once discretized in the space of wavelet coefficients. Differentially private algorithms for convex optimization (such as DP-SGD) could, in principle, be applied directly.

However, the problem is that f^* (the Fenchel conjugate which appears in the optimization problem) does not admit a closed-form expression in terms of the coefficients, akin to its gradient. In addition, it is not guaranteed that the problem benefits from gradient smoothness or strong convexity, also limiting approaches such as objective perturbation.