



Near-optimal estimation of smooth transport maps with kernel sums-of-squares

Boris Muzellec, Adrien Vacher, Francis Bach, François-Xavier Vialard,
Alessandro Rudi

► To cite this version:

Boris Muzellec, Adrien Vacher, Francis Bach, François-Xavier Vialard, Alessandro Rudi. Near-optimal estimation of smooth transport maps with kernel sums-of-squares. SIAM Journal on Mathematics of Data Science, In press. hal-03466696

HAL Id: hal-03466696

<https://hal.science/hal-03466696>

Submitted on 6 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Near-optimal estimation of smooth transport maps with kernel sums-of-squares

Boris Muzellec ^{**} Adrien Vacher ^{°*} Francis Bach ^{**}
François-Xavier Vialard [°] Alessandro Rudi ^{**}

[°] LIGM, Université Gustave Eiffel, CNRS

^{*} INRIA Paris, 2 rue Simone Iff, 75012, Paris, France

^{*} ENS - Département d'Informatique de l'École Normale Supérieure,

^{*} PSL Research University, 2 rue Simone Iff, 75012, Paris, France

`boris.muzellec@inria.fr`, `adrien.vacher@u-pem.fr`, `francis.bach@inria.fr`,
`francois-xavier.vialard@u-pem.fr`, `alessandro.rudi@inria.fr`

Abstract

It was recently shown that under smoothness conditions, the squared Wasserstein distance between two distributions could be efficiently computed with appealing statistical error upper bounds. However, rather than the distance itself, the object of interest for applications such as generative modeling is the underlying optimal transport map. Hence, computational and statistical guarantees need to be obtained for the estimated maps themselves. In this paper, we propose the first tractable algorithm for which the statistical L^2 error on the maps nearly matches the existing minimax lower-bounds for smooth map estimation. Our method is based on solving the semi-dual formulation of optimal transport with an infinite-dimensional sum-of-squares reformulation, and leads to an algorithm which has dimension-free polynomial rates in the number of samples, with potentially exponentially dimension-dependent constants.

1 Introduction

Optimal transport (OT) provides a principled method to compare probability distributions, by finding the optimal way of coupling one to another based on a cost function on their supports. This optimization problem yields two useful quantities: the transport cost itself, which corresponds to the Wasserstein distance when the ground cost is a distance, and the minimizer, which is a map that pushes forward the first distribution onto the second, known as the transport map. While OT has gained attention lately in the statistics and machine learning communities due to the properties of the Wasserstein distance, transportation maps are playing an increasingly important role in data sciences. Indeed, many applications such as generative modeling (Arjovsky et al., 2017; Salimans et al., 2018; Bernton et al., 2017; Makkuva et al., 2020; Onken et al., 2021), domain adaptation (Courty et al., 2016, 2017), shape matching (Su et al., 2015; Feydy et al., 2017) or predicting cell trajectories (Schiebinger et al., 2019; Yang et al., 2020) among others can be formulated as the problem of finding a map from a reference distribution to a target distribution.

Over the past decades, efforts have particularly concentrated on the problem of computing OT distances. Two cases must be distinguished: the case of discrete measures (i.e., measures supported on a finite number of points), and the general case, including e.g. measures that admit a density with respect to the Lebesgue measure. For discrete measures supported on N points, OT distances and plans may be computed by solving a linear program (LP) in $O(N^3 \log(N))$ time using the network simplex algorithm (see, e.g., [Ahuja et al., 1993](#)). Using entropic optimal transport ([Cuturi, 2013](#)) as a proxy, this computational cost can be further reduced to $O(N^2)$. However, efficient methods for general measures remain to be found. The naive approach, which consists in using the OT distance between samples of the distributions (the so-called plugin estimator), fails as the dimension grows. Indeed, the sharpest statistical bounds known for the plugin estimator require $\varepsilon^{-\frac{2}{d}}$ samples to achieve precision ε ([Chizat et al., 2020](#)). Yet, theoretical estimators were recently derived in the case where the problem is smooth ([Weed and Berthet, 2019](#); [Hütter and Rigollet, 2021](#)), yielding rates of estimation of the OT distance that improve as the smoothness grows. However, the estimators presented in those works are not algorithmically tractable. Very recently, [Vacher et al. \(2021\)](#) closed this statistical-computational gap by designing an estimator of the squared Wasserstein distance relying on kernel sums-of-squares ([Marteau-Ferey et al. \(2020\)](#), and in particular [Rudi et al. \(2020\)](#)) that may be computed with polynomial dimension-free rates in the number of samples, but with a constant term that is potentially exponential in the dimension.

On the other hand, fewer results on computationally efficient statistical estimators of optimal maps are available in the literature. Compared to the problem of estimating OT distances, the estimation of an optimal transport map based on samples is particularly difficult since the estimated maps need to be evaluated on out-of-sample points. When smoothness assumptions are made, one can design theoretical statistical estimators of the transport maps that are statistically minimax optimal with respect to the smoothness for the L^2 error ([Hütter and Rigollet, 2021](#)), but such estimators cannot be computed in practice as they require projecting on the space of smooth and strongly convex functions, which is NP-hard. In a few empirical studies, the potentials were explicitly parameterized either by a neural network ([Seguy et al., 2018](#)) or by a Gaussian kernel ([Genevay et al., 2016](#)). However, neither computational nor statistical guarantees are provided in those works. Likewise, [Paty et al. \(2020\)](#) recently managed to derive an explicit algorithm when the potentials are assumed to have smooth gradients and to be strongly convex, but the corresponding rate of approximation is not known. A recent line of research studies estimators of optimal maps that rely on barycentric projection, which can be developed in regularized and non-regularized settings ([Gunsilius, 2018](#); [Deb et al., 2021](#); [Pooladian and Niles-Weed, 2021](#)). These estimators can be used without any further assumption on the source and target measures. Here again, either the rates do match minimax rates of estimation of [Hütter and Rigollet \(2021\)](#) without computationally feasible estimators, or the rates are not optimal. In particular, the estimator of [Pooladian and Niles-Weed \(2021\)](#) is based on entropic regularization and is computationally friendly, but only works for low regularity of the maps, and is not minimax optimal.

Contributions. The contributions of this paper are twofold. First, relying on the sum-of-squares (SoS) tight reformulation of OT proposed in [Vacher et al. \(2021\)](#), we bridge the statistical-computational gap on the estimation of the optimal potentials. Second, we propose

improvements on the algorithms proposed by [Vacher et al. \(2021\)](#) to enhance the practicality of SoS for optimal transport.

1. **An estimator for smooth OT potentials.** Relying on an SoS representation of Brenier dual constraints, we design an algorithm to compute an estimator of the OT potentials based on samples. Remarkably, the complexity of this algorithm has polynomial rates in the number of samples (with potentially exponential constants w.r.t. the underlying dimension). For this estimator, we derive a statistical upper bound on the L^2 error on the resulting maps that nearly matches the lower bound of [Hütter and Rigollet \(2021\)](#) when the smoothness is high. To achieve these nearly optimal rates, we refine the regularization that was proposed by [Vacher et al. \(2021\)](#) and carry out a careful analysis to bound the solutions of the empirical problem. The central argument consists in using the local strong convexity of the so-called semi-dual formulation of optimal transport. It enables us to relate the L^2 error both to the quality of the (scalar) approximation of the original problem and to the norm of the empirical minimizers. In turn, this allows us to obtain a statistical rate for the L^2 error with well-chosen regularizers.
2. **Algorithmic improvements.** We propose several improvements on the algorithms introduced by [Vacher et al. \(2021\)](#). First, as our objective is to compute an estimator of the OT potentials themselves, our algorithm should provide convergence guarantees for the minimizing solution, and not only the minimum. To do so, we reformulate the objective as a strongly convex optimization problem by using a square norm regularization instead of the trace, and by replacing hard constraints with quadratic penalties. A second advantage of this strongly convex relaxation is that it allows using first order optimization methods such as accelerated gradient descent, which are more scalable than the Newton methods employed by [Vacher et al. \(2021\)](#) whose cost quickly blows up with the number of samples. Next, we further reduce computational costs by introducing a Nyström approximation of the features in the dual constraints. Finally, we propose a criterion to choose the hyperparameters on which our estimator relies, which was lacking in the work of [Vacher et al.](#)

2 Notations and background

Let (μ, ν) be two probability distributions on two bounded domains $X, Y \subset \mathbb{R}^d$. We study the squared Wasserstein distance, whose dual Kantorovitch formulation is given by

$$\begin{aligned} W_2^2(\mu, \nu) = & \sup_{u, v \in C(\mathbb{R}^d)} \langle u, \mu \rangle + \langle v, \nu \rangle \\ \text{subject to } & \frac{\|x - y\|^2}{2} \geq u(x) + v(y), \quad \forall (x, y) \in X \times Y, \end{aligned} \tag{1}$$

where $C(\mathbb{R}^d)$ is the space of continuous functions on \mathbb{R}^d . For convenience, we shall use the equivalent Brenier formulation, where the Kantorovitch potentials u, v are replaced by the

Brenier potentials $f(\cdot) := \frac{\|\cdot\|^2}{2} - u(\cdot)$, $g(\cdot) := \frac{\|\cdot\|^2}{2} - v(\cdot)$

$$\begin{aligned} \text{OT}(\mu, \nu) = & \inf_{f, g \in C(\mathbb{R}^d)} \langle f, \mu \rangle + \langle g, \nu \rangle \\ \text{subject to } & f(x) + g(y) \geq \langle x, y \rangle, \quad \forall (x, y) \in X \times Y. \end{aligned} \quad (2)$$

These two quantities are related by $W_2^2(\mu, \nu) = \langle \frac{\|\cdot\|^2}{2}, \mu + \nu \rangle - \text{OT}(\mu, \nu)$.

Problem (2) (as well as problem (1)) is delicate to solve numerically due to the non-negativity constraint which has to be satisfied on a continuous set. For instance, a feasible idea consists in sampling the inequality constraint on a finite set and trying to extend it to the continuous set. Unfortunately, this strategy may only leverage Lipschitzness, even if the functions involved are smoother than Lipschitz, yielding an approximation rate of order $n^{-\frac{1}{2d}}$, $2d$ being the dimension of $X \times Y$. The method proposed by [Vacher et al. \(2021\)](#) overcomes this difficulty and is able to leverage the smoothness of the potentials. Following their work, we require assumptions on the support and the smoothness of the densities themselves.

Assumption 1 (*m*-times differentiable densities). *Let $m, d \in \mathbb{N}$. Let $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$.*

- a) μ, ν have densities. Their supports, resp. $X, Y \subset \mathbb{R}^d$ are convex, bounded and open with a Lipschitz boundary;*
- b) the densities are finite, bounded away from zero, with Lipschitz derivatives up to order m .*

Using Caffarelli's regularity theory ([De Philippis and Figalli, 2014](#)), Assumption 1 ensures that the Brenier potentials have a similar order of differentiability. In particular, defining the Sobolev space of order m over $Z \subset \mathbb{R}^d$

$$H^m(Z) := \left\{ f \in L^2(Z) \mid \|f\|_{H^m} := \sum_{|\alpha| \leq m} \|D^\alpha f\|_{L^2(Z)} < \infty \right\}, \quad (3)$$

the optimal Brenier potentials (f_*, g_*) belong to the Sobolev space $H^{m+2}(X)$ and $H^{m+2}(Y)$ respectively ([De Philippis and Figalli, 2014](#)). When $m > d/2 - 2$, the Sobolev embedding theorem gives that these spaces are continuously embedded in the space of continuous functions and thus are reproducing kernel Hilbert spaces (RKHS, see e.g. [Paulsen and Raghupathi, 2016](#); [Steinwart and Christmann, 2008](#)).

The approach proposed in [Vacher et al.](#) to estimate the Wasserstein distance is based on the kernel Sum-Of-Squares (SoS) approach and the tools introduced by [Rudi et al. \(2020\)](#) to deal with optimization problems subject to a dense set of inequalities. The procedure consists in two steps: (1) show that the optimization problem is equivalent to a problem where the inequality constraint is substituted by an equality constraint with respect to an SoS term; (2) consider an empirical version of the resulting problem with only a finite number of equality constraints. [Rudi et al. \(2020\)](#) shows, for the case of non-convex optimization, that this procedure is adaptive to the degree of differentiability of the objective function leading to rates that overcome the curse of dimensionality for very smooth objectives. [Vacher et al. \(2021\)](#) shows that for the case of the problem in (2), the two steps correspond to the following Eq. (4) and Eq. (5) as reported in the following theorem and below.

Theorem 1 (Vacher et al. (2021)). *Let $\mu \in \mathcal{P}(X), \nu \in \mathcal{P}(Y)$ satisfy Assumption 1 and assume that $m > d + 1$. Then, problem (2) can be reformulated as*

$$\begin{aligned} OT(\mu, \nu) = & \inf_{\substack{f, g \in C(\mathbb{R}^d) \\ A \in \mathbb{S}^+(H^m(X \times Y))}} \langle f, \mu \rangle + \langle g, \nu \rangle \\ & \text{subject to } f(x) + g(y) = \langle x, y \rangle + \langle \phi(x, y), A\phi(x, y) \rangle_{H^m(X \times Y)}, \end{aligned} \quad (4)$$

where $\mathbb{S}^+(H^m(X \times Y))$ is the set of linear, positive, self-adjoint operators on $H^m(X \times Y)$.

The key contribution of this representation theorem is to replace the inequality constraint by an equality constraint which is easier to deal with, as proposed and analyzed in Rudi et al. (2020). Given access to samples $x_1, \dots, x_n \sim \mu$ and $y_1, \dots, y_n \sim \nu$ with associated empirical distributions $\hat{\mu}$ and $\hat{\nu}$, we solve for $\lambda_1, \lambda_2 \geq 0$ the empirical problem

$$\begin{aligned} \min_{\substack{f \in H^{m+2}(X), g \in H^{m+2}(Y), \\ A \in \mathbb{S}_+(H^m(X \times Y))}} \langle f, \hat{\mu} \rangle + \langle g, \hat{\nu} \rangle + \lambda_1 \text{Tr}(A) + \lambda_2 (\|f\|_{H^{m+2}(X)}^2 + \|g\|_{H^{m+2}(Y)}^2) \\ \text{subject to } \forall j \in [n], \quad f(x_j) + g(y_j) - \langle x_j, y_j \rangle = \langle \phi(x_j, y_j), A\phi(x_j, y_j) \rangle_{H^m(X \times Y)}, \end{aligned} \quad (5)$$

where $\phi : X \times Y \mapsto H^m(X \times Y)$ is the feature map of $H^m(X \times Y)$. Using the techniques of Marteau-Ferey et al. (2020) and Rudi et al. (2020), the authors show that problem (5) that in the case where $m > 2d$, the energy of the empirical potentials can be controlled with high probability by setting $\lambda_1 = \lambda_2 \sim \frac{1}{\sqrt{n}}$, leading to a dimension-free approximation of $OT(\mu, \nu)$ at a $\frac{1}{\sqrt{n}}$ rate with high probability, with a computational complexity of $O(n^{3.5} \log(\frac{1}{\varepsilon}))$, for a precision of ε .

As problem (5) is strongly convex with respect to the potentials f, g , the uniqueness of the empirical potentials \hat{f}_n, \hat{g}_n is ensured. Further, their existence may also be proven. However, the question remains of recovering a rate of convergence of the empirical potentials toward the optimal Brenier potentials (f_*, g_*) with respect to some norm. Hütter and Rigollet (2021) prove that when $d > 3$, any estimator of the transport map \hat{T}_n can achieve at best an L^2 error that scales as

$$\mathbb{E}(\|\nabla f_* - \hat{T}_n\|_{L^2(\mu)}^2) \sim n^{-\frac{m+1}{m+d/2}}. \quad (6)$$

We prove in the next section that the solutions (\hat{f}_n, \hat{g}_n) of the problem (5) can nearly match this rate up to an exponent ε , where ε can be chosen arbitrary close to 0, for well-chosen regularization parameters λ_1, λ_2 .

3 Nearly optimal minimax rates

The key to derive the nearly minimax optimal rates is to have estimates on λ_1, λ_2 that are sharper than $\frac{1}{\sqrt{n}}$. Indeed, in the first step of the proof we show an upper-bound of the form

$$(\lambda_1 + \lambda_2)(\|\hat{f}_n\|_{H^{m+2}(X)} + \|\hat{g}_n\|_{H^{m+2}(Y)} + \hat{A}_n) \lesssim \frac{1}{\sqrt{n}}(\|\nabla \hat{f}_n - \nabla f_*\|_{L^2(\mu)}^2 + \|\nabla \hat{g}_n - \nabla g_*\|_{L^2(\nu)}^2)^\alpha,$$

where α is positive, decreases with the dimension d to zero and increases with the smoothness m and $(\hat{f}_n, \hat{g}_n, \hat{A}_n)$ are minimizers of problem (5). In order to control in turn the errors

$\|\nabla f - \nabla f_*\|_{L^2(\mu)}^2, \|\nabla g - \nabla g_*\|_{L^2(\nu)}^2$ with the regularizers λ_1, λ_2 , we need extra convexity. To this end, we introduce the semi-dual formulation of optimal transport (Brenier, 1987) which replaces the potential pair (f, g) by the couple (f, f^*) where f^* is the Fenchel-Legendre transform $f^*(y) = \sup_{x \in X} x^\top y - f(x)$. It reads

$$\text{OT}(\mu, \nu) = \inf_{f \in C(\mathbb{R}^d)} J_{\mu, \nu}(f), \quad (7)$$

where $J_{\mu, \nu}(f) = \langle f, \mu \rangle + \langle f^*, \nu \rangle$. We simply denote it by $J(f)$ when no confusion is possible. As shown in the next lemma, the functional $J(f)$ gains stronger convexity in comparison with the linear objective of Brenier's formulation.

Lemma 2. *The semi-dual objective $J(f)$ is convex in $f \in C(X)$. Assuming that there exists an optimal potential f_* such that ∇f_* pushes μ onto ν and that f is a convex C^1 function with M -Lipschitz gradient, we have*

$$\|\nabla f - \nabla f_*\|_{L^2(\mu)}^2 \leq 2M(J(f) - J(f_*)). \quad (8)$$

Note that there is no assumption on the measures μ and ν themselves. However, we assume the Lipschitz smoothness of the gradient of f and the existence of an optimal map, which is ensured by Brenier's theorem if μ has density w.r.t. the Lebesgue measure. A similar result is proven by Hütter and Rigollet (2021) and we include for completeness a simple proof in Appendix A. Thanks to this gain of convexity, we obtain an upper bound of the errors by the regularizers λ_1, λ_2 of the form

$$\|\nabla \hat{f}_n - \nabla f_*\|_{L^2(\mu)}^2 + \|\nabla \hat{g}_n - \nabla g_*\|_{L^2(\nu)}^2 \lesssim (\lambda_1 + \lambda_2) \frac{1}{\sqrt{n}}. \quad (9)$$

Thanks to these two connected bounds, we can calibrate the regularizers λ_1, λ_2 and eventually obtain rates sharper than $\frac{1}{\sqrt{n}}$.

Theorem 3. *Let $\delta, \varepsilon \in]0, 1]^2$ and assume that the regularizers are given by*

$$\lambda_n^1 = \lambda_n^2 = \lambda_n = \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}} + C_1 \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}}, \quad (10)$$

where C_1 is a constant that does not depend on n and δ . Denoting \hat{f}_n (resp. \hat{g}_n) the transport potential associated to μ (resp. ν) in problem (5), we have with probability at least $1 - \delta$ that for $n \geq n_0(X, Y, d, m)$,

$$\|\nabla \hat{f}_n - \nabla f_*\|_{L^2(\mu)}^2 + \|\nabla \hat{g}_n - \nabla g_*\|_{L^2(\nu)}^2 \leq C_2 \lambda_n, \quad (11)$$

where C_2 is a positive constant that is independent from n and δ but grows to infinity when ε goes to 0. In particular, when m is sufficiently large, the minimax rate is nearly attained:

$$\|\nabla \hat{f}_n - \nabla f_*\|_{L^2(\mu)}^2 + \|\nabla \hat{g}_n - \nabla g_*\|_{L^2(\nu)}^2 \leq C_2 \left(\frac{\log(\frac{2}{\delta})}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}}. \quad (12)$$

Proof Let us denote by $(\hat{f}_n, \hat{g}_n, \hat{A}_n)$ the solutions of the empirical problem (5) and by (f_*, g_*, A_*) the solutions of the deterministic problem (4). We define the energy of the potentials $R^2 = \|f_*\|_{H^{m+2}(X)}^2 + \|g_*\|_{H^{m+2}(Y)}^2$ and its empirical counterpart $\hat{R}_n^2 = \|\hat{f}_n\|_{H^{m+2}(X)}^2 + \|\hat{g}_n\|_{H^{m+2}(Y)}^2$. For probability measures (α, β) , we shall denote throughout the proof the linear objective of the dual as $\mathcal{E}_{\alpha, \beta}(f, g) = \langle f, \alpha \rangle + \langle g, \beta \rangle$. The strategy of the proof is as follows

1. We use the optimality conditions of the empirical and deterministic OT problems to upper bound λ_n and the norms of the empirical objects $(\hat{f}_n, \hat{g}_n, \hat{A}_n)$ by the errors $\|\nabla \hat{f}_n - f_*\|_{L^2(\mu)}, \|\nabla \hat{g}_n - g_*\|_{L^2(\nu)}$.
2. We use the strong convexity of the semi-dual to upper bound in turn the errors by the regularizer λ_n and the norms of $(\hat{f}_n, \hat{g}_n, \hat{A}_n)$.
3. We obtain two connected bounds and we optimize over λ_n to obtain sharp rates while maintaining the empirical objects $(\hat{f}_n, \hat{g}_n, \hat{A}_n)$ bounded.

Upper bound on the regularizer. The optimality conditions in the empirical problem (5) yield

$$\lambda_n(\text{Tr}(\hat{A}_n) - \text{Tr}(A_*) + \hat{R}_n^2 - R^2) \leq \mathcal{E}_{\hat{\mu}, \hat{\nu}}(f_*, g_*) - \mathcal{E}_{\hat{\mu}, \hat{\nu}}(\hat{f}_n, \hat{g}_n). \quad (13)$$

Conversely, we wish to test the empirical potentials against the measures μ, ν and use the optimality condition of the deterministic problem.

Proposition 4. *There exist constants C, G such that for all $0 < \delta < 1$, defining $\hat{\kappa}_{n, \delta} = C(\hat{R}_n + \text{Tr}(\hat{A}_n) + G)n^{-\frac{m-d}{2d}} \log(\frac{n}{\delta})^{\frac{m-d}{2d}}$ we have with probability at least $1 - \delta$ that for $n \geq n_0(X, Y, d, m)$, the empirical potentials $(\hat{f}_n, \hat{g}_n + \hat{\kappa}_{n, \delta})$ are admissible candidates:*

$$\forall (x, y) \in X \times Y, \hat{f}_n(x) + \hat{g}_n(y) + \hat{\kappa}_{n, \delta} \geq x^\top y. \quad (14)$$

The proof is left in Appendix A. We deduce that with probability at least $1 - \delta$

$$0 \leq \mathcal{E}_{\mu, \nu}(\hat{f}_n, \hat{g}_n) - \mathcal{E}_{\mu, \nu}(f_*, g_*) + \hat{\kappa}_{n, \delta}. \quad (15)$$

Adding the equation above with the empirical optimality conditions (13) yields

$$\begin{aligned} \lambda_n(\text{Tr}(\hat{A}_n) - \text{Tr}(A_*) + \hat{R}_n^2 - R^2) &\leq \mathcal{E}_{\hat{\mu}, \hat{\nu}}(f_*, g_*) - \mathcal{E}_{\hat{\mu}, \hat{\nu}}(\hat{f}_n, \hat{g}_n) + \mathcal{E}_{\mu, \nu}(\hat{f}_n, \hat{g}_n) - \mathcal{E}_{\mu, \nu}(f_*, g_*) \\ &\quad + \hat{\kappa}_{n, \delta} \\ &= \langle f_* - \hat{f}_n, \hat{\mu} \rangle + \langle g_* - \hat{g}_n, \hat{\nu} \rangle + \langle \hat{f}_n - f_*, \mu \rangle + \langle \hat{g}_n - g_*, \nu \rangle \\ &\quad + \hat{\kappa}_{n, \delta} \\ &= \langle \hat{f}_n - f_*, \mu - \hat{\mu} \rangle + \langle \hat{g}_n - g_*, \nu - \hat{\nu} \rangle + \hat{\kappa}_{n, \delta} \end{aligned}$$

Now we use the following lemma upper-bounding the linear forms of the r.h.s. with respect to the errors $\|\nabla \hat{f}_n - f_*\|_{L^2(\mu)}, \|\nabla \hat{g}_n - g_*\|_{L^2(\nu)}$.

Lemma 5. *Let $(u, v) \in H^{m+2}(X) \times H^{m+2}(X)$, $(\delta, \varepsilon) \in]0, 1[\times]0, 1[$ and μ be a probability measure over X . If μ has a density on X bounded away from zero, then there exists a constant C independent of n, δ, u, v such that with probability at least $1 - \delta$, we have*

$$\langle u - v, \mu - \hat{\mu} \rangle \leq C_\mu \frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(\Delta_{u,v,\mu}^{\frac{m+2-d/2-\varepsilon}{m+1}} R(u, v)^{\frac{d/2+\varepsilon-1}{m+1}} + \Delta_{u,v,\mu} \right), \quad (16)$$

where $\Delta_{u,v,\mu} = \|\nabla u - \nabla v\|_{L^2(\mu)}$ and $R(u, v) = \|u\|_{H^{m+2}(X)} + \|v\|_{H^{m+2}(X)}$.

The proof of this lemma is left in the Appendix A and is mainly an application of the Gagliardo-Nirenberg inequality and the Poincaré inequality. Applying the lemma to (f_*, \hat{f}_n, μ) and (g_*, \hat{g}_n, ν) , and using the fact that for $(\alpha, x, y) \in (\mathbb{R}^+)^3$, we have $x^\alpha + y^\alpha \leq 2(x + y)^\alpha$, we recover that for $n \geq n_0(X, Y, d, m)$ we have with probability at least $1 - 3\delta$

$$\begin{aligned} \lambda_n(\text{Tr}(\hat{A}_n) - \text{Tr}(A_*) + \hat{R}_n^2 - R^2) &\leq 2C' \frac{\log(\frac{2}{\delta})}{\sqrt{n}} (\hat{R}_n + R)^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} \\ &\quad + 2C' \frac{\log(\frac{2}{\delta})}{\sqrt{n}} a_n + \hat{\kappa}_{n,\delta}, \end{aligned} \quad (17)$$

where $C' = C_\mu + C_\nu$ and we denoted $a_n = \|\nabla \hat{f}_n - \nabla f_*\|_{L^2(\mu)} + \|\nabla \hat{g}_n - \nabla g_*\|_{L^2(\nu)}$.

Bounding the error a_n . First we bound the errors from above via the strong convexity of the semi-dual. Then, using the optimality conditions, we recover the same linear upper bound as in the previous paragraph.

a) Semi-dual upper bound. A straightforward way to upper-bound the errors would be to apply Lemma 2 to the empirical potentials \hat{f}_n, \hat{g}_n . Unfortunately, there is no guarantee that the empirical potentials are convex. Instead, we define for $t \in [0, 1]$ the following interpolating potential

$$\tilde{f}_n(t)(\cdot) = f_*(\cdot) + t(\hat{f}_n(\cdot) - f_*(\cdot)). \quad (18)$$

Since f_* and g_* are mutual Legendre transforms and that they both have Lipschitz gradients, (f_*, g_*) are strongly convex and we denote by γ a strong convexity constant of the optimal Brenier potentials. Hence for $t \leq \frac{\gamma}{2\|\hat{f}_n - f\|_{W^{2,\infty}(X)}}$, the interpolate $\tilde{f}_n(t)(\cdot)$ is $\frac{\gamma}{2}$ strongly convex. Hence we choose

$$\hat{t}_f = \min\left(1, \frac{\gamma}{2\|\hat{f}_n - f\|_{W^{2,\infty}(X)}}\right), \quad (19)$$

and we apply Lemma 2 to $\tilde{f}_n(\hat{t}_f)(\cdot)$, which yields

$$\|\nabla \tilde{f}_n(\hat{t}_f) - \nabla f_*\|_{L^2(\mu)}^2 \leq 2\hat{L}_f(J(\tilde{f}_n(\hat{t}_f)) - J(f_*)), \quad (20)$$

where $\hat{L}_f = \|\tilde{f}_n(\hat{t}_f)\|_{W^{2,\infty}(X)}$. By convexity of the semi-dual J , the r.h.s. is upper bounded by $2\hat{L}_f\hat{t}_f(J(\hat{f}_n) - J(f_*))$ and the l.h.s. is equal to $\hat{t}_f^2\|\nabla \hat{f}_n - \nabla f\|_{L^2(\mu)}^2$. Hence, we have

$$\|\nabla \hat{f}_n - \nabla f\|_{L^2(\mu)}^2 \leq \frac{2\hat{L}_f}{\hat{t}_f} (J(\hat{f}_n) - J(f_*)). \quad (21)$$

Finally, we upper-bound the quantity $\frac{\hat{L}_f}{\hat{t}_f}$.

Proposition 6. Denoting K_X the embedding constant such that $\forall u \in W^{2,\infty}(X), \|u\|_{W^{2,\infty}(X)} \leq K_X \|u\|_{H^{m+2}(X)}$, the quantity $\frac{\hat{L}_f}{\hat{t}_f}$ can be upper bounded as

$$\frac{\hat{L}_f}{\hat{t}_f} \leq 2K_X(\|\hat{f}_n\|_{H^{m+2}(X)} + \|f_*\|_{H^{m+2}(X)}) \left(1 + \frac{\|f_*\|_{H^{m+2}(X)}}{\gamma}\right). \quad (22)$$

The proof is left in Appendix A. Conversely, applying this result to \hat{g}_n , we obtain

$$\frac{\hat{L}_g}{\hat{t}_g} \leq 2K_Y(\|\hat{g}_n\|_{H^{m+2}(Y)} + \|g_*\|_{H^{m+2}(Y)}) \left(1 + \frac{\|g_*\|_{H^{m+2}(Y)}}{\gamma}\right), \quad (23)$$

where K_Y is the embedding constant from $H^{m+2}(Y)$ to $W^{2,\infty}(Y)$. In particular

$$\begin{aligned} \frac{\hat{L}_f}{\hat{t}_f} + \frac{\hat{L}_g}{\hat{t}_g} &\leq 2(K_X + K_Y)(\|\hat{f}_n\|_{H^{m+2}(X)} + \|\hat{g}_n\|_{H^{m+2}(Y)} + \|f_*\|_{H^{m+2}(X)} + \|g_*\|_{H^{m+2}(Y)}) \\ &\quad \times \left(1 + \frac{\|f_*\|_{H^{m+2}(X)} + \|g_*\|_{H^{m+2}(Y)}}{\gamma}\right) \end{aligned} \quad (24)$$

$$\leq 2\sqrt{2}(K_X + K_Y)(\hat{R}_n + R) \left(1 + \frac{\sqrt{2}R}{\gamma}\right). \quad (25)$$

b) Linear upper bound. Recall that the couple $(\hat{f}_n, \hat{g}_n + \hat{\kappa}_{n,\delta})$ are admissible candidates and in particular $J(\hat{f}_n) \leq \mathcal{E}_{\mu,\nu}(\hat{f}_n, \hat{g}_n + \hat{\kappa}_{n,\delta})$, which implies

$$J(\hat{f}_n) - J(f_*) \leq \mathcal{E}_{\mu,\nu}(\hat{f}_n, \hat{g}_n) - \mathcal{E}_{\mu,\nu}(f_*, g_*) + \hat{\kappa}_{n,\delta}. \quad (26)$$

Combining the equation above with the empirical optimality condition $0 \leq \mathcal{E}_{\hat{\mu},\hat{\nu}}(f_*, g_*) - \mathcal{E}_{\hat{\mu},\hat{\nu}}(\hat{f}_n, \hat{g}_n) + \lambda_n(\text{Tr}(A_*) + R^2)$ (13), we obtain

$$J(\hat{f}_n) - J(f_*) \leq \langle f_* - \hat{f}_n, \mu - \hat{\mu} \rangle + \langle g_* - \hat{g}_n, \nu - \hat{\nu} \rangle + \lambda_n(\text{Tr}(A_*) + R^2) + \hat{\kappa}_{n,\delta}. \quad (27)$$

As in the previous paragraph, we use Lemma 5 to upper-bound the linear forms of the r.h.s., and denoting $a_n = \|\nabla \hat{f}_n - \nabla f_*\|_{L^2(\mu)} + \|\nabla \hat{g}_n - \nabla g_*\|_{L^2(\nu)}$, we obtain

$$\begin{aligned} J(\hat{f}_n) - J(f_*) &\leq 2C' \frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left[(\hat{R}_n + R)^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n \right] \\ &\quad + \lambda_n(\text{Tr}(A_*) + R^2) + \hat{\kappa}_{n,\delta}, \end{aligned} \quad (28)$$

where C' is a constant independent of n, δ .

Combining the bounds on the error a_n . From the previous paragraphs, we obtain the following bounds

$$\begin{aligned} \lambda_n(\text{Tr}(\hat{A}_n) + \hat{R}_n^2) &\leq 2C' \frac{\log(\frac{2}{\delta})}{\sqrt{n}} [(\hat{R}_n + R)^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n] \\ &\quad + \lambda_n(\text{Tr}(A_*) + R^2) + \hat{\kappa}_{n,\delta}, \end{aligned} \quad (29)$$

$$a_n^2 \leq 2(\hat{R}_n + R)KC' \left(\frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left[(\hat{R}_n + R)^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n \right] + \lambda_n(\text{Tr}(A_*) + R^2) + \hat{\kappa}_{n,\delta} \right), \quad (30)$$

where we defined $K = \sqrt{2}(K_X + K_Y) \left(1 + \frac{2\sqrt{2}R}{\gamma} \right)$. The inequalities above mutually relate the empirical norms $\hat{R}_n, \text{Tr}(\hat{A}_n)$ and the error a_n . The next result shows that for a well-chosen regularizer λ_n , the empirical norms are bounded independently on δ and that for this choice of regularizer, the error a_n^2 is at most of the order λ_n .

Proposition 7. *Assume that Equations (29) and (30) hold for $n \geq n_0$. If we set the regularizer as*

$$\lambda_n = \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}} + C_1 \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}}, \quad (31)$$

where C_1 is the constant independent of n, δ , then $\hat{R}_n, \text{Tr}(\hat{A}_n)$ are bounded independently of δ and there exists a constant C_2 independent of n, δ , such that

$$a_n^2 \leq C_2 \lambda_n. \quad (32)$$

The proof is presented in Appendix A. Hence setting λ_n to the indicated value yields that with probability at least $1 - 3\delta$, for $n \geq n_0(X, Y, d, m)$ we have

$$a_n^2 \leq C \lambda_n, \quad (33)$$

with C a constant independent of n, δ . ■

Discussion and remarks. Let us make a few remarks regarding the comparison of our results to those of [Hütter and Rigollet \(2021\)](#) and [Vacher et al. \(2021\)](#).

Remark 8. *Our bound does not exactly match the lower-bound $n^{-\frac{m+1}{m+d/2}}$ of [Hütter and Rigollet \(2021\)](#) because of the extra exponent ε . Indeed, we cannot take $\varepsilon = 0$ since the constant C_2 would then go to infinity. However, it would be interesting to recover the growth rate of C_2 when ε goes to 0 ; if it does not increase too fast, we could hope for a trade-off in n on ε and maybe recover the exact lower-bound.*

Remark 9. *Note that the minimax rate is nearly attained only for the highly smooth case $m > 2d$, and that when m goes toward d we cannot even guarantee that empirical potentials converges to the original potentials. This gap compared with [Hütter and Rigollet \(2021\)](#) is explained by the fact that in their case, the authors directly solve the semi-dual problem while in our case, we approximate the cost constraint through the SoS reformulation which can only be achieved when the smoothness is sufficiently high.*

Remark 10. Comparing our results to [Vacher et al. \(2021\)](#), one may notice that we require lower values for the regularizers. In particular, one may wonder if this could lead to gains in terms of the estimation of the squared Wasserstein distance itself. This is not the case, as this effect is due to the fact that we are here considering the Brenier version of optimal transport. Indeed, it holds $W_2^2(\mu, \nu) = \langle \frac{\|\cdot\|^2}{2}, \mu + \nu \rangle - OT(\mu, \nu)$. Hence, to estimate $W_2^2(\mu, \nu)$ one must estimate the extra moment term, that fluctuates as $\frac{1}{\sqrt{n}}$.

4 Algorithms and Numerical Experiments

In this section, we provide and test algorithms to solve Eq. (5) and compute the estimators \hat{f} and \hat{g} . Following [Rudi et al. \(2020\)](#) and [Vacher et al. \(2021\)](#), we may rewrite Eq. (5) using a finite-dimensional representation of A that relies on finite-dimensional features $\Phi_i \in \mathbb{R}^\ell, i \in [n]$, and a PSD matrix $\mathbf{B} \in \mathbb{S}_+(\mathbb{R}^\ell)$:

$$\begin{aligned} \min_{\substack{f \in H^{m+2}(X), g \in H^{m+2}(Y), \\ \mathbf{B} \in \mathbb{S}_+(\mathbb{R}^\ell)}} \quad & \langle f, \hat{\mu} \rangle + \langle g, \hat{\nu} \rangle + \lambda_1 \text{Tr} \mathbf{B} + \lambda_2 (\|f\|_{H^{m+2}(X)}^2 + \|g\|_{H^{m+2}(Y)}^2) \\ \text{subject to} \quad & \forall j \in [\ell], \quad f(\tilde{x}_j) + g(\tilde{y}_j) - \langle \tilde{x}_j, \tilde{y}_j \rangle = \Phi_j^T \mathbf{B} \Phi_j. \end{aligned} \quad (34)$$

We refer to [Vacher et al. \(2021\)](#) for the definition of $\Phi_j, j \in [\ell]$ and the derivation of Eq. (34). Note that compared to Eq. (5) we have two sets of samples that do not necessarily coincide: $x_1 \dots x_n \in X$ and $y_1, \dots, y_n \in Y$ which are the support points of the empirical distributions $\hat{\mu}$ and $\hat{\nu}$, and $(\tilde{x}_1, \tilde{y}_1) \dots (\tilde{x}_\ell, \tilde{y}_\ell) \in X \times Y$ which are used to subsample the dual constraints. One can then apply the method proposed by [Vacher et al. \(2021\)](#) and solve (5) by using barrier methods on its dual formulation:

$$\begin{aligned} \min_{\gamma \in \mathbb{R}^\ell} \quad & \frac{1}{4\lambda_2} \gamma^T \mathbf{Q} \gamma - \frac{1}{2\lambda_2} \sum_{j=1}^{\ell} \gamma_j z_j + \frac{q^2}{2\lambda_2} - \frac{\delta}{\ell} \log \det \left(\sum_{j=1}^{\ell} \gamma_j \Phi_j \Phi_j^T + \lambda_1 \mathbf{I}_\ell \right) \\ \text{subject to} \quad & \sum_{j=1}^{\ell} \gamma_j \Phi_j \Phi_j^T + \lambda_1 \mathbf{I}_\ell \succeq 0, \end{aligned} \quad (35)$$

where $\mathbf{Q} \in \mathbb{R}^{\ell \times \ell}$ is defined as $\mathbf{Q}_{i,j} = k(\tilde{x}_i, \tilde{y}_i) + k(\tilde{x}_i, \tilde{y}_j), i, j \in [\ell]$, $z_j = \tilde{x}_j \cdot \tilde{y}_j + w_{\hat{\mu}}(\tilde{x}_j) + w_{\hat{\nu}}(\tilde{y}_j)$, $q^2 = \|w_{\hat{\mu}}(\tilde{x}_j)\|_{\mathcal{H}}^2 + \|w_{\hat{\nu}}(\tilde{y}_j)\|_{\mathcal{H}}^2$, and $\delta > 0$ is the barrier parameter.¹ The authors showed that solving this problem to precision ε has a $O(n^{3.5} \log \frac{1}{\varepsilon})$ computational cost. This method has two drawbacks: first, its cost is prohibitive when the number of samples becomes large, which is necessary to ensure better statistical approximation, and second, there are no guarantees on the convergence to the minimizers \hat{f} and \hat{g} as only the objective value is guaranteed to converge to the optimal value. Both issues can be alleviated by introducing a strongly convex and unconstrained relaxation of problem (5).

Let us now consider a variant of problem (5) with a squared norm penalty (instead of a

¹Compared to [Vacher et al. \(2021\)](#), the definition of z above is slightly different. This is due to the fact that we consider the ‘‘Brenier’’ formulation of dual OT (an infimum with an inner product cost), instead of the Kantorovich formulation (a supremum with a quadratic cost).

trace penalty) on \mathbf{B} :

$$\begin{aligned} \inf_{\substack{f \in \mathcal{H}_X, g \in \mathcal{H}_Y, \\ \mathbf{B} \in \mathbb{S}_+(\mathbb{R}^\ell)}} \quad & \langle f, \hat{w}_\mu \rangle_{\mathcal{H}_X} + \langle g, \hat{w}_\nu \rangle_{\mathcal{H}_Y} + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 + \lambda_2 (\|f\|_{\mathcal{H}_X}^2 + \|g\|_{\mathcal{H}_Y}^2) \\ \text{subject to} \quad & \forall j \in [\ell], \quad f(\tilde{x}_j) + g(\tilde{y}_j) - \langle \tilde{x}_j, \tilde{y}_j \rangle = \Phi_j^T \mathbf{B} \Phi_j. \end{aligned} \quad (36)$$

Lemma 11. Equation (36) admits the following dual problem:

$$\min_{\gamma \in \mathbb{R}^\ell} \quad \frac{1}{4\lambda_2} \gamma^T \mathbf{Q} \gamma - \frac{1}{2\lambda_2} \sum_{j=1}^{\ell} \gamma_j z_j + \frac{1}{2\lambda_1} \|(-\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T)_+\|_F^2 + \frac{q^2}{4\lambda_2}, \quad (37)$$

where $(\mathbf{M})_+$ denotes the positive part of $\mathbf{M} \in \mathbb{R}^\ell$, i.e. its projection on $\mathbb{S}_+(\mathbb{R}^\ell)$.

Proof in Appendix B, Page 26. Although the bounds described in Section 3 and by Vacher et al. (2021) rely on a trace control of \mathbf{B} , we conjecture that similar results could be derived using a square Hilbert-Schmidt norm control.

Faster algorithms with strong convexity. Replacing the trace penalty with a square norm penalty led to the *unconstrained* convex dual problem (37), without adding a barrier term. However, this problem is not necessarily strongly convex. Indeed, \mathbf{Q} may have arbitrarily small eigenvalues, and the positive part term vanishes when its argument is a negative matrix. We now propose a strongly convex relaxation of (37), obtained by replacing the constraints in (36) with a quadratic penalization. That is, for $\delta > 0$, we consider:

$$\begin{aligned} \inf_{\substack{f \in \mathcal{H}_X, g \in \mathcal{H}_Y, \\ \mathbf{B} \in \mathbb{S}_+(\mathbb{R}^d)}} \quad & \left\{ \langle f, \hat{w}_\mu \rangle_{\mathcal{H}_X} + \langle g, \hat{w}_\nu \rangle_{\mathcal{H}_Y} + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 + \lambda_2 (\|f\|_{\mathcal{H}_X}^2 + \|g\|_{\mathcal{H}_Y}^2) \right. \\ & \left. + \frac{\delta}{2\ell} \sum_{j=1}^{\ell} (f(\tilde{x}_j) + g(\tilde{y}_j) - \langle \tilde{x}_j, \tilde{y}_j \rangle - \Phi_j^T \mathbf{B} \Phi_j)^2 \right\}. \end{aligned} \quad (38)$$

Let us derive the corresponding dual problem.

Proposition 12. Strong duality holds, and problem (38) admits the following dual formulation:

$$\inf_{\gamma \in \mathbb{R}^\ell} \quad \frac{1}{4\lambda_2} \gamma^T \mathbf{Q} \gamma - \frac{1}{2\lambda_2} \sum_{j=1}^{\ell} \gamma_j z_j + \frac{1}{2\lambda_1} \|(-\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T)_+\|_F^2 + \frac{\ell}{2\delta} \|\gamma\|^2 + \frac{q^2}{4\lambda_2}, \quad (39)$$

with the following primal-dual relations

$$\begin{aligned} f &= \frac{1}{2\lambda_2} \left(\sum_{i=1}^{\ell} \gamma_i \phi_X(\tilde{x}_i) - \hat{w}_\mu \right) \\ g &= \frac{1}{2\lambda_2} \left(\sum_{i=1}^{\ell} \gamma_i \phi_Y(\tilde{y}_i) - \hat{w}_\nu \right). \end{aligned} \quad (40)$$

This problem is strongly convex, with constant $\alpha = \frac{\ell}{\delta} + \frac{1}{2\lambda_2} \lambda_{\min}(\mathbf{Q})$, and smooth, with constant $L = \frac{\ell}{\delta} + \frac{1}{2\lambda_2} \lambda_{\max}(\mathbf{Q}) + \frac{1}{\lambda_1} \lambda_{\max}(\mathbf{K} \circ \mathbf{K})$ (where $\mathbf{K} = \Phi^T \Phi$).

(39) is a tight relaxation of (37), as the dual of the constrained problem may be recovered by taking $\delta \rightarrow \infty$. As mentioned above, for most kernels the spectrum of \mathbf{Q} decays rapidly to zero. Therefore, problem (37) may not be considered strongly convex for practical purposes. On the other hand, problem (39) has a strong convexity constant that is bounded from under by $\frac{\delta}{\ell}$, thanks to the quadratic regularization term on γ that appears in Eq. (39). Hence, we may solve (39) using accelerated gradient descent algorithms. The bottleneck of the computation is the positive part, which is obtained by computing the SVD of a $\ell \times \ell$ matrix in $O(\ell^3)$ time. The size of this matrix can be reduced by considering a low-rank approximation of \mathbf{K} .

Note however that relaxing the constraints in (36) with a quadratic penalization introduces a gap with the theory developed in Section 3. However, we expect that bounds could be obtained in the case of quadratic penalties, much like soft constraints compared with hard constraints in kernel ridge regression. We leave this for future work.

Nyström approximation. The computational bottleneck of solving (39) is forming the matrix $\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T$ and computing the SVD at each iteration, both in $O(\ell^3)$ time. This computational cost can be reduced using a Nyström approximation (see e.g. Williams and Seeger, 2001) of the kernel matrix \mathbf{K} . Indeed, such an approximation is of the form

$$\mathbf{K}^{nys} = \mathbf{V} \mathbf{W}^{-1} \mathbf{V}^T, \quad (41)$$

where $\mathbf{W} \in \mathbb{R}^{r \times r}$ is the submatrix corresponding to $r < \ell$ randomly sampled columns (and corresponding rows) of \mathbf{K} , and $\mathbf{V} \in \mathbb{R}^{\ell \times r}$. Hence, writing $\mathbf{W} = \mathbf{L}^T \mathbf{L}$ with $\mathbf{L} \in \mathbb{R}^{r \times r}$, we may consider features $\Phi_i^{nys} \in \mathbb{R}^r, i \in [\ell]$ corresponding to the columns of

$$\mathbf{R}^{nys} \stackrel{\text{def}}{=} \mathbf{L}^{-T} \mathbf{V}^T, \quad (42)$$

verifying $(\mathbf{R}^{nys})^T \mathbf{R}^{nys} = \mathbf{K}^{nys}$. Replacing \mathbf{K} with \mathbf{K}^{nys} in the derivation of the dual leads to

$$\inf_{\gamma \in \mathbb{R}^{\ell}} \quad \frac{1}{4\lambda_2} \gamma^T \mathbf{Q} \gamma - \frac{1}{2\lambda_2} \sum_{j=1}^{\ell} \gamma_j z_j + \frac{1}{2\lambda_1} \left\| -\sum_{j=1}^{\ell} \gamma_j \Phi_j^{(nys)} (\Phi_j^{nys})^T \right\|_F^2 + \frac{\ell}{2\delta} \|\gamma\|^2 + \frac{q^2}{4\lambda_2}. \quad (43)$$

Forming the matrix $\sum_{i=1}^{\ell} \gamma_i \Phi_i^{(nys)} (\Phi_i^{nys})^T$ costs $O(r^2 \ell)$ (compared to $O(\ell^3)$ in (39)), and computing its SVD costs $O(r^3)$ (compared to $O(\ell^3)$). This brings down the cost of a gradient step to $O(r^3 + r^2 \ell)$, compared to $O(\ell^3)$ without using approximations.

The empirical results in Figures 1 to 3 tend to show that one can pick quite small ranks while retaining good performance: indeed, in those experiments the maximum employed rank is 100, whereas we increase the number of samples up to 1000. This is coherent with results on the Nyström approximation applied to statistical learning (Rudi et al., 2015). However, the effect of Nyström approximation is currently not measured in our statistical analysis.

Selecting the hyperparameters. Contrary to the optimization setting (Rudi et al., 2020), there is no direct way of assessing the quality of the estimator \hat{OT} (e.g. of the “smaller is better” type) for a given choice of hyperparameters, as the algorithm may output values

that are larger or smaller than the true value of OT. Alternatively, we may use the output potentials to estimate transportation maps \hat{T}_1, \hat{T}_2 from μ to ν and ν to μ , and compare the mapped points to the target distributions. That is, we compute

$$\begin{aligned}\hat{T}_1 : x \rightarrow \nabla f(x) &= \frac{1}{2\lambda_2} \left(\sum_{j=1}^{\ell} \gamma_j \nabla_x k(\tilde{x}_j, x) - \nabla_x \hat{w}_\mu(x) \right) \\ \hat{T}_2 : x \rightarrow \nabla f(x) &= \frac{1}{2\lambda_2} \left(\sum_{j=1}^{\ell} \gamma_j \nabla_x k(\tilde{y}_j, x) - \nabla_x \hat{w}_\nu(x) \right),\end{aligned}\tag{44}$$

and select the hyperparameters for which

$$\widehat{\text{MMD}} \stackrel{\text{def}}{=} \text{MMD} \left(\frac{1}{n_\mu} \sum_{i=1} \delta_{\hat{T}_1(x_i)}, \hat{\nu} \right) + \text{MMD} \left(\frac{1}{n_\nu} \sum_{i=1} \delta_{\hat{T}_2(y_i)}, \hat{\mu} \right)\tag{45}$$

is the smallest, where MMD is the maximum mean discrepancy (Gretton et al., 2012) for the kernel k , defined as

$$\text{MMD}(\mu, \nu) \stackrel{\text{def}}{=} \iint (k(x, x) - 2k(x, y) + k(y, y)) d(\mu \otimes \nu)(x, y).\tag{46}$$

The choice of MMD as a criterion has two main motivations: first, it is a divergence that can be efficiently computed and whose plugin estimator converges at rate $O(n^{-1/2})$, and second, as observed by Vacher et al. (2021) in Appendix F, when the filing samples $\tilde{x}_j, \tilde{y}_j, j \in [\ell]$ correspond to all pairs of samples in $\hat{\mu}$ and $\hat{\nu}$, the dual problem (35) may be reformulated as a regularized optimal transport problem with MMD marginal penalties.

Numerical results. We now report empirical evaluation of our estimators. In Figures 1 to 3, we sample data from two Gaussian distributions (whose covariances follow a Wishart distribution), and solve problem (39) for hyperparameters pairs (λ_1, λ_2) on the $\{10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\}^2$ grid. As explained above, we then compute the forward and backward transportation maps (44) corresponding to each (λ_1, λ_2) pair, and report the performance of the pair minimizing (45), varying the number of samples from 25 to 1000, averaging over 20 random draws. For each number of samples, we report the error of the OT distance estimator (like Vacher et al. (2021)) compared to the plugin estimator, the mean squared error of \hat{T}_1 and \hat{T}_2 compared to the ground truths T_1 and T_2 (which admit an analytical expression (see Takatsu, 2011)) evaluated on the samples from $\mu \otimes \nu$:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n |\hat{T}_1(x_i) - T_1(x_i)|^2 + \frac{1}{n} \sum_{i=1}^n |\hat{T}_2(y_i) - T_2(y_i)|^2,\tag{47}$$

the best hyperparameter pair and the MMD metric (45). We may obtain several takeaways from Figures 1 to 3. First, we see that gains of using the kernel SoS approach to OT increase with the dimension: while SoS-OT matches the performance of the sampled estimator in 2D (Fig. 1), it becomes increasingly more performant than the plugin estimator in 4D (Fig. 2) and 8D (Fig. 3). Second, we observe that the map estimator converges to the true map, and that the decrease of the L_2 error (47) is well correlated with the MMD criterion (45). Finally, as expected, the optimal regularization parameters (selected from a gridsearch) decrease as the number of sample increases.




Figure 1: 2D Gaussian data. (*Top left*) OT error $|\widehat{OT} - OT|$. (*Top right*) Best hyperparameters (selected via gridsearch), log scale. (*Bottom left*) Transportation map mean square error, log scale. (*Bottom right*) MMD between transported samples, log scale. Shaded areas correspond to \pm std. Algorithm: accelerated gradient descent ($\delta = 10^3$), Sobolev kernel ($s = 20$, bandwidth = 1), with Nystrom approximation (rank = 100). Filling pairs $(\tilde{x}_i, \tilde{y}_i)$ are drawn from $\mu \otimes \nu$. The number of filling pairs is equal to the number of μ and ν samples, reported on the x-axis.

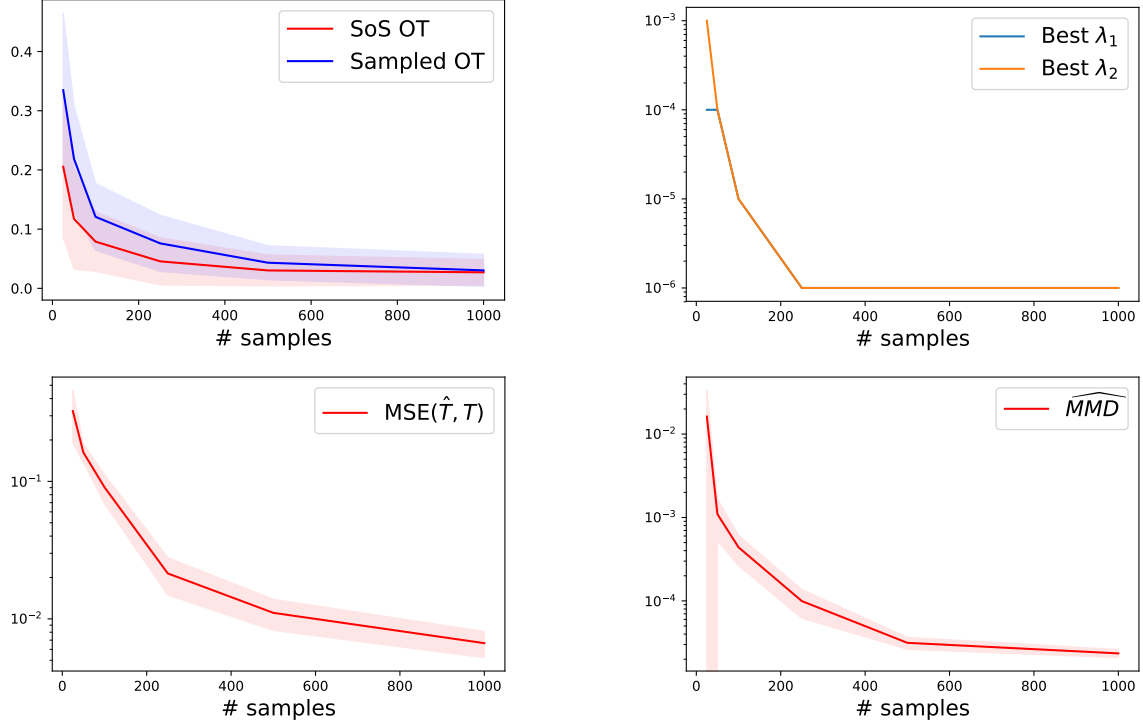


Figure 2: 4D Gaussian data. (*Top left*) OT error $|\widehat{OT} - OT|$. (*Top right*) Best hyperparameters (selected via gridsearch), log scale. (*Bottom left*) Transportation map mean square error, log scale. (*Bottom right*) MMD between transported samples, log scale. Shaded areas correspond to \pm std. Algorithm: accelerated gradient descent ($\delta = 10^3$), Sobolev kernel ($s = 20$, bandwidth = 1), with Nystrom approximation (rank = 100). Filling pairs $(\tilde{x}_i, \tilde{y}_i)$ are drawn from $\mu \otimes \nu$. The number of filling pairs is equal to the number of μ and ν samples, reported on the x-axis.

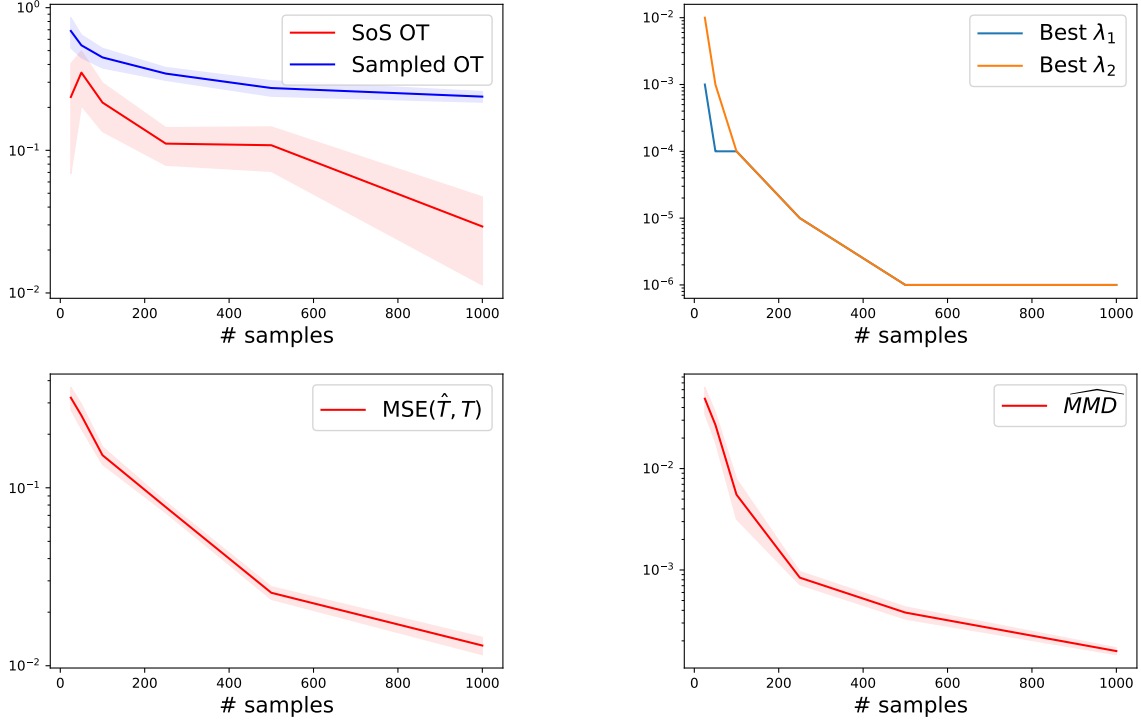


Figure 3: 8D Gaussian data. (Top left) OT error $|\widehat{OT} - OT|$, log scale. (Top right) Best hyperparameters (selected via gridsearch), log scale. (Bottom left) Transportation map mean square error, log scale. (Bottom right) MMD between transported samples, log scale. Shaded areas correspond to \pm std. Algorithm: accelerated gradient descent ($\delta = 10^3$), Sobolev kernel ($s = 20$, bandwidth = 1), with Nystrom approximation (rank = 100). Filling pairs $(\tilde{x}_i, \tilde{y}_i)$ are drawn from $\mu \otimes \nu$. The number of filling pairs is equal to the number of μ and ν samples, reported on the x-axis. Note that the std in the top left is actually decreasing, but that the size of the shaded area visually increases due to logarithmic scaling.

Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). We also acknowledge support from the European Research Council (grants SEQUOIA 724063 and REAL 947908), and Région Ile-de-France.

References

- Robert A. Adams and John J. F. Fournier. *Sobolev Spaces*. Elsevier, 2003.
- Ravindra K Ahuja, James B Orlin, and Thomas L Magnanti. *Network flows: theory, algorithms, and applications*. Prentice-Hall, 1993.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- Espen Bernton, Pierre E Jacob, Mathieu Gerber, and Christian P Robert. Inference in generative models using the wasserstein distance. *arXiv preprint arXiv:1701.05146*, 1(8):9, 2017.
- Yann Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *CR Acad. Sci. Paris Sér. I Math.*, 305:805–808, 1987.
- Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- Lenaïc Chizat, Pierre Roussillon, Flavien Léger, François-Xavier Vialard, and Gabriel Peyré. Faster Wasserstein distance estimation with the Sinkhorn divergence. *Advances in Neural Information Processing Systems*, 33, 2020.
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 39(9):1853–1865, 2016.
- Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. *Advances in Neural Information Processing Systems*, pages 3733–3742, 2017.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26:2292–2300, 2013.
- Guido De Philippis and Alessio Figalli. The Monge–Ampère equation and its link to optimal transportation. *Bulletin of the American Mathematical Society*, 51(4):527–580, 2014.
- Nabarun Deb, Promit Ghosal, and Bodhisattva Sen. Rates of estimation of optimal transport maps using plug-in estimators via barycentric projections. *Advances in Neural Information Processing Systems*, 34, 2021.

- Jean Feydy, Benjamin Charlier, François-Xavier Vialard, and Gabriel Peyré. Optimal transport for diffeomorphic registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 291–299. Springer, 2017.
- Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*, volume 30, 2016.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1): 723–773, 2012.
- Florian F Gunsilius. On the convergence rate of potentials of Brenier maps. *Econometric Theory*, pages 1–37, 2018.
- Jan-Christian Hütter and Philippe Rigollet. Minimax estimation of smooth optimal transport maps. *The Annals of Statistics*, 49(2):1166–1194, 2021.
- Ashok Makkuva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pages 6672–6681. PMLR, 2020.
- Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Non-parametric models for non-negative functions. *Advances in Neural Information Processing Systems*, 2020.
- Norman G. Meyers and William P. Ziemer. Integral inequalities of Poincaré and Wirtinger type for BV functions. *American Journal of Mathematics*, 99, 1977.
- Francis Narcowich, Joseph Ward, and Holger Wendland. Sobolev bounds on functions with scattered zeros, with applications to radial basis function surface fitting. *Mathematics of Computation*, 74(250):743–763, 2005.
- D Onken, S Wu Fung, Xingjian Li, and L Ruthotto. Ot-flow: Fast and accurate continuous normalizing flows via optimal transport. In *AAAI Conference on Artificial Intelligence*, volume 35, 2021.
- François-Pierre Paty, Alexandre d’Aspremont, and Marco Cuturi. Regularity as regularization: Smooth and strongly convex brenier potentials in optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pages 1222–1232. PMLR, 2020.
- Vern I. Paulsen and Mrinal Raghupathi. *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*, volume 152. Cambridge University Press, 2016.
- Aram-Alexandre Pooladian and Jonathan Niles-Weed. Entropic estimation of optimal transport maps. *arXiv preprint arXiv:2109.12004*, 2021.
- Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. *Advances in Neural Information Processing Systems*, 28: 1657–1665, 2015.

- Alessandro Rudi, Ulysse Marteau-Ferey, and Francis Bach. Finding global minima via kernel approximations. In *Arxiv preprint arXiv:2012.11978*, 2020.
- Tim Salimans, Dimitris Metaxas, Han Zhang, and Alec Radford. Improving GANs using optimal transport. In *International Conference on Learning Representations*, 2018.
- Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- Vivien Seguy, Bharath Bhushan Damodaran, Remi Flamary, Nicolas Courty, Antoine Rolet, and Mathieu Blondel. Large-scale optimal transport and mapping estimation. In *International Conference on Learning Representations*, 2018.
- Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer Science & Business Media, 2008.
- Zhengyu Su, Yalin Wang, Rui Shi, Wei Zeng, Jian Sun, Feng Luo, and Xianfeng Gu. Optimal mass transport for shape matching and comparison. *IEEE transactions on pattern analysis and machine intelligence*, 37(11):2246–2259, 2015.
- Asuka Takatsu. Wasserstein geometry of Gaussian measures. *Osaka Journal of Mathematics*, 48(4):1005 – 1026, 2011.
- Adrien Vacher, Boris Muzellec, Alessandro Rudi, Francis Bach, and Francois-Xavier Vialard. A dimension-free computational upper-bound for smooth optimal transport estimation. *Conference on Learning Theory*, 2021.
- Jonathan Weed and Quentin Berthet. Estimation of smooth densities in Wasserstein distance. In *Conference on Learning Theory*, pages 3118–3119, 2019.
- Christopher Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. *Advances in Neural Information Processing Systems*, 13:682–688, 2001.
- Karren Dai Yang, Karthik Damodaran, Saradha Venkatachalapathy, Ali C Soylemezoglu, GV Shivashankar, and Caroline Uhler. Predicting cell lineages using autoencoders and optimal transport. *PLoS computational biology*, 16(4):e1007828, 2020.

A Additional Proofs

A.1 Proof of lemma 2

Proof First note that the Legendre transform is pointwise convex. For (f, g) and $\lambda \in [0, 1]$, we have for all y

$$(\lambda f + (1 - \lambda)g)^*(y) = \sup_x x^\top y - (\lambda f(x) + (1 - \lambda)g(x)) \quad (48)$$

$$\leq \lambda f^*(y) + (1 - \lambda)g^*(y). \quad (49)$$

It follows that the semi-dual functional $J(f) = \langle f, \mu \rangle + \langle f^*, \nu \rangle$ is convex.

Denoting T_0 the optimal transport map from μ to ν , the semi-dual functional can be rewritten as $J(f) = \langle f, \mu \rangle + \langle f^* \circ T_0, \mu \rangle$. Now, the Fenchel-Young inequality on f gives for every couple $(x, y) \in X \times Y$ $y^\top x \leq f(x) + f^*(y)$, Equality holds for $y = \nabla f(x)$, and in this case, one also have $x \in \partial f^*(\nabla f(x))$. To simplify notations, let us denote $T(x) = \nabla f(x)$. We get

$$f(x) + f^*(T(x)) = T(x)^\top x. \quad (50)$$

Denoting f_0 an optimal potential, the optimality condition applied to f_0 gives $f_0(x) + f_0^*(T_0(x)) - T_0(x)^\top x = 0, \forall x \in \text{Supp}(\mu)$ and by integration $J(f_0) = \int T_0(x)^\top x \, d\mu$. Therefore, we have

$$J(f) - J(f_0) = \int f(x) + f^*(T_0(x)) - T_0(x)^\top x \, d\mu. \quad (51)$$

Subtracting Eq. (50) applied to f pointwise from the integrand in Eq. (51), we get

$$J(f) - J(f_0) = \int f^*(T_0(x)) - f^*(T(x)) - (T_0(x) - T(x))^\top x \, d\mu. \quad (52)$$

Applying the case of equality in Fenchel-Young (see above), we have $x \in \partial f^*(T(x))$ where ∂f^* is the sub-gradient of f^* . Hence, denoting $z_1 = T_0(x)$ and $z_2 = T(x)$, the integrand has the form

$$f^*(z_1) - f^*(z_2) - (z_1 - z_2)^\top \partial f^*(z_2). \quad (53)$$

Since f is a C^1 convex function with a M -Lipschitz gradient, then f^* is $\frac{1}{M}$ -strongly convex and (53) is lower bounded as

$$f^*(z_1) - f^*(z_2) - (z_1 - z_2)^\top \partial f^*(z_2) \geq \frac{1}{2M} \|z_1 - z_2\|^2, \quad (54)$$

and in particular, we recover

$$J(f) - J(f_0) \geq \frac{1}{2M} \int \|T_0(x) - T(x)\|^2 d\mu(x). \quad (55)$$

■

A.2 Proof of Proposition 4

Proof For $E \subset \mathbb{R}^d$ with lipschitz boundary and points $(e_i)_{1 \leq i \leq n}$, define the *fill distance* $h(E, e)$ as

$$h(E, e) = \sup_{z \in E} \min_{i \in [n]} \|z - e_i\|. \quad (56)$$

If f has smoothness m , the sampling inequalities (Narcowich et al., 2005) state that if $\forall i \in [n], f(e_i) = 0$ then there exists a positive constant B depending on the smoothness m , the dimension d and the geometry of E such that if $h \leq h_0(m, d, E)$

$$\sup_{e \in E} |f(e)| \leq B h^{m-d/2} \|f\|_{H^m(E)}. \quad (57)$$

Now, recall that at the optimum, the function

$$\hat{\alpha}_n(x, y) := \hat{f}_n(x) + \hat{g}_n(y) - x^\top y - \langle \phi(x, y), \hat{A}_n \phi(x, y) \rangle_{H^m(X \times Y)}$$

is such that $\forall i \in [n]$, $\hat{\alpha}_n(x_i, y_i) = 0$. Applying the previous result yields

$$\sup_{(x, y) \in X \times Y} |\hat{\alpha}_n(x, y)| \leq Bh(X \times Y, (\hat{x}, \hat{y}))^{m-d} \|\hat{\alpha}_n\|_{H^m(X \times Y)}. \quad (58)$$

Using Lemma 12 from [Vacher et al. \(2021\)](#), if $n \geq n_0(X, Y)$ there exists a universal constant C_0 such that we have with probability at least $1 - \delta$

$$h(X \times Y, (\hat{x}, \hat{y})) \leq C_0 \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{1}{2d}}. \quad (59)$$

There remains to upper bound $\|\hat{\alpha}_n\|_{H^m(X \times Y)}$

$$\begin{aligned} \|\hat{\alpha}_n\|_{H^m(X \times Y)} &\leq \|\hat{f}_n(\cdot)\|_{H^m(X \times Y)} + \|\hat{g}_n(\cdot)\|_{H^m(X \times Y)} + \|c\|_{H^m(X \times Y)} \\ &\quad + \|\langle \phi(\cdot, \cdot), \hat{A}_n \phi(\cdot, \cdot) \rangle_{H^m(X \times Y)}\|_{H^m(X \times Y)}, \end{aligned} \quad (60)$$

where $c(x, y) = x^\top y$. The term $\|\hat{f}_n(\cdot)\|_{H^m(X \times Y)}$ is upper bounded by $C_1 \|\hat{f}_n(\cdot)\|_{H^{m+2}(X)}$ where C_1 is the embedding constant of $H^{m+2}(X)$ in $H^m(X)$. Using Lemma 9 in [Rudi et al. \(2020\)](#), the term $\|\langle \phi(\cdot, \cdot), \hat{A}_n \phi(\cdot, \cdot) \rangle_{H^m(X \times Y)}\|_{H^m(X \times Y)}$ can be upper bounded by $C_2 \text{Tr}(\hat{A}_n)$ where C_2 is a constant depending on X, Y, m, d . Hence we obtain

$$\|\hat{\alpha}_n\|_{H^m(X \times Y)} \leq C_1 \|\hat{f}_n\|_{H^{m+2}(X)} + C_3 \|\hat{g}_n\|_{H^{m+2}(Y)} + C_2 \text{Tr}(\hat{A}_n) + G, \quad (61)$$

where C_3 is the embedding constant of $H^{m+2}(Y)$ in $H^m(Y)$ and $G = \|c\|_{H^m(X \times Y)} < \infty$. In particular, since $\forall (x, y) \in X \times Y$, $\langle \phi(x, y), \hat{A}_n \phi(x, y) \rangle_{H^m(X \times Y)} \geq 0$, we obtain that for $n \geq n_0(X, Y, m, d)$, we have with probability at least $1 - \delta$

$$\forall (x, y) \in X \times Y, \hat{f}_n(x) + \hat{g}_n(y) + C_0 C_4 \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}} (\sqrt{2} \hat{R}_n + \text{Tr}(\hat{A}_n) + \frac{G}{C_4}) \geq x^\top y, \quad (62)$$

where $C_4 = \max(C_1, C_2, C_3)$ and $\hat{R}_n^2 = \|\hat{f}_n\|_{H^{m+2}(X)}^2 + \|\hat{g}_n\|_{H^{m+2}(Y)}^2$. ■

A.3 Proof of Lemma 5

Proof In order to apply the Poincaré inequality, we need to re-normalize the function $u - v$ such that it integrates to 0. Denoting the residual $r = \int_X u(x) - v(x) dx$, we have that $\langle u - v, \mu - \hat{\mu} \rangle = \langle u - v - r, \mu - \hat{\mu} \rangle$ since μ and $\hat{\mu}$ have the same total mass. For a probability measure α and a RKHS H with kernel k , we define kernel mean embedding $w_\alpha = \mathbb{E}_\alpha(k(X, \cdot)) \in H$. With this notation, using the reproducing property, we can re write $\langle u - v - r, \mu - \hat{\mu} \rangle$ as $\langle u - v - r, w_\mu - w_{\hat{\mu}} \rangle_{H^{d/2+\varepsilon}(X)}$. Using Cauchy-Schwarz, we have the upper-bound

$$\langle u - v, \mu - \hat{\mu} \rangle \leq \|u - v - r\|_{H^{d/2+\varepsilon}(X)} \|w_\mu - w_{\hat{\mu}}\|_{H^{d/2+\varepsilon}(X)}. \quad (63)$$

We can apply the Pinelis inequality (see Proposition 2 of [Caponnetto and De Vito, 2007](#)) to control the second term. We have with probability at least $1 - \delta$

$$\|w_\mu - w_{\hat{\mu}}\|_{H^{d/2+\varepsilon}(X)} \leq k_{d,\varepsilon} \frac{\log(\frac{2}{\delta})}{\sqrt{n}}, \quad (64)$$

where $k_{d,\varepsilon}$ is a constant that goes to infinity as $1/\varepsilon$ when $\varepsilon \rightarrow 0$. Now let us deal with the first term. There exists a constant C such that

$$\|u - v - r\|_{H^{d/2+\varepsilon}(X)} \leq C(\|\nabla u - \nabla v\|_{H^{d/2+\varepsilon-1}(X)} + \|u - v - r\|_{L^2(X)}). \quad (65)$$

For the second term, Poincaré-Wirtinger inequality ([Meyers and Ziemer, 1977](#)) states that there exists a constant C' such that $\|u - v - r\|_{L^2(X)} \leq C'\|\nabla u - \nabla v\|_{L^2(X)}$. For the first term, we apply the Gagliardo-Nirenberg inequality which yields

$$\|\nabla u - \nabla v\|_{H^{d/2+\varepsilon-1}(X)} \leq C_1'' \|\nabla u - \nabla v\|_{H^{m+1}(X)}^{\frac{d/2+\varepsilon-1}{m+1}} \|\nabla u - \nabla v\|_{L^2(X)}^{1-\frac{d/2+\varepsilon-1}{m+1}} + C_2'' \|\nabla u - \nabla v\|_{L^2(X)}, \quad (66)$$

for C_1'', C_2'' constants independent of u, v and δ . Conversely, there exists C''' such that

$$\|\nabla u - \nabla v\|_{H^{m+1}(X)} \leq C''' \|u - v\|_{H^{m+2}(X)}. \quad (67)$$

Finally, denoting $\mu_0 = \inf_{x \in X} \frac{d\mu(x)}{d\lambda(x)}$, we have $\|\cdot\|_{L^2(X)} \leq \frac{\|\cdot\|_{L^2(\mu)}}{\mu_0}$ and we obtain with probability $1 - \delta$

$$\langle u - v, \mu - \hat{\mu} \rangle \leq K_0 \frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(K_1 \|u - v\|_{H^{m+2}(X)}^{\frac{d/2+\varepsilon-1}{m+1}} \|\nabla u - \nabla v\|_{L^2(\mu)}^{\frac{m+2-d/2-\varepsilon}{m+1}} + K_2 \|\nabla u - \nabla v\|_{L^2(\mu)} \right),$$

where $K_0 = k_{d,\varepsilon} C$, $K_1 = \frac{(C''')^{\frac{d/2+\varepsilon-1}{m+1}}}{\mu_0^{\frac{m+2-d/2-\varepsilon}{m+1}}}$ and $K_2 = \frac{C(C'+C_2'')}{\mu_0}$. Denoting $C_\mu = K_0(K_1 + K_2)$, we get

$$\langle u - v, \mu - \hat{\mu} \rangle \leq C_\mu \frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(\|u - v\|_{H^{m+2}(X)}^{\frac{d/2+\varepsilon-1}{m+1}} \|\nabla u - \nabla v\|_{L^2(\mu)}^{\frac{m+2-d/2-\varepsilon}{m+1}} + \|\nabla u - \nabla v\|_{L^2(\mu)} \right). \quad (68)$$

■

A.4 Proof of Proposition 6

Proof Recall that $\hat{L}_f = \|\tilde{f}_n(\hat{t}_f)\|_{W^{2,\infty}(X)}$ and that $\tilde{f}_n(\hat{t}_f) = f_* + \hat{t}_f(\hat{f}_n - f_*)$. Hence \hat{L}_f can be upper-bounded as

$$\hat{L}_f \leq \|f_*\|_{W^{2,\infty}(X)} + \hat{t}_f \|\hat{f}_n - f_*\|_{W^{2,\infty}(X)}, \quad (69)$$

and as a consequence $\frac{\hat{L}_f}{\hat{t}_f} \leq \frac{\|f_*\|_{W^{2,\infty}(X)}}{\hat{t}_f} + \|\hat{f}_n\|_{W^{2,\infty}(X)} + \|f_*\|_{W^{2,\infty}(X)}$. Now recall that \hat{t}_f is given by

$$\hat{t}_f = \min\left(1, \frac{\gamma}{2\|\hat{f}_n - f_*\|_{W^{2,\infty}(X)}}\right), \quad (70)$$

and in particular

$$\frac{1}{\hat{t}_f} \leq \frac{2\|\hat{f}_n - f_*\|_{W^{2,\infty}(X)}}{\gamma} + 1 \leq \frac{2\left(\|\hat{f}_n\|_{W^{2,\infty}(X)} + \|f_*\|_{W^{2,\infty}(X)}\right)}{\gamma} + 1, \quad (71)$$

which yields

$$\frac{\hat{L}_f}{\hat{t}_f} \leq \|f_*\|_{W^{2,\infty}(X)} \left(\frac{2\|f_*\|_{W^{2,\infty}(X)}}{\gamma} + 2 \right) + \|\hat{f}_n\|_{W^{2,\infty}(X)} \left(\frac{2\|f_*\|_{W^{2,\infty}(X)}}{\gamma} + 1 \right) \quad (72)$$

$$\leq 2(\|f_*\|_{W^{2,\infty}(X)} + \|\hat{f}_n\|_{W^{2,\infty}(X)}) \left(\frac{\|f_*\|_{W^{2,\infty}(X)}}{\gamma} + 1 \right) \quad (73)$$

$$\leq 2K_X(\|f_*\|_{H^{m+2}(X)} + \|\hat{f}_n\|_{H^{m+2}(X)}) \left(\frac{\|f_*\|_{W^{2,\infty}(X)}}{\gamma} + 1 \right), \quad (74)$$

where K_X is the embedding constant of $W^{2,\infty}(X)$ into $H^{m+2}(X)$ (Adams and Fournier, 2003). ■

A.5 Proof of Proposition 7

Proof The objective of the proof is to find a parameter λ_n as small as possible that keeps the quantities \hat{R}_n and $\text{Tr}(\hat{A}_n)$ bounded ; the rate in λ_n immediately follows the boundedness.

Denoting $b_n = \hat{R}_n + R$ and $c_n = \lambda_n(\text{Tr}(A_*) + R^2) + \hat{\kappa}_{n,\delta}$, we have the following system of inequalities

$$\begin{cases} a_n^2 \leq KC' b_n \left[\frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n \right) + c_n \right] \\ \lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n)) \leq (2C' + 1) \left[\frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n \right) + c_n \right]. \end{cases} \quad (75)$$

We want to split the analysis into two parts: the indexes for which

$$\begin{cases} \frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n \right) > c_n \text{ (Case 1)} \\ \frac{\log(\frac{2}{\delta})}{\sqrt{n}} \left(b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} + a_n \right) \leq c_n \text{ (Case 2)} \end{cases}. \quad (76)$$

Again, in the first case, we sub-split the analysis in two parts: the indexes for which

$$\begin{cases} a_n > b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} \text{ (Case 1a)} \\ a_n \leq b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} \text{ (Case 1b)} \end{cases}. \quad (77)$$

Case 1a. On these indexes, we can re-write Equation (75) as

$$\begin{cases} a_n^2 \leq 4KC' b_n \frac{\log(\frac{2}{\delta})}{\sqrt{n}} a_n \\ \lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n)) \leq 4(2C' + 1) \frac{\log(\frac{2}{\delta})}{\sqrt{n}} a_n, \end{cases} \quad (78)$$

Combining these two equations yields

$$\lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n)) \leq 16K(2C' + 1)C' \frac{\log(\frac{2}{\delta})^2}{n} b_n. \quad (79)$$

If we set $\lambda_n \geq \frac{\log(\frac{2}{\delta})}{n}$, we recover

$$\hat{R}_n^2 + \text{Tr}(\hat{A}_n) \leq 16K(2C' + 1)C'(\hat{R}_n + R), \quad (80)$$

which implies that \hat{R}_n and $\text{Tr}(\hat{A}_n)$ are bounded independently on δ and that

$$a_n^2 \leq (4KC'b_*)^2 \frac{\log(\frac{2}{\delta})^2}{n}, \quad (81)$$

where $b_* = \sup_n b_n$.

Case 1b. On these indexes, the inequalities we obtain are

$$\begin{cases} a_n^2 \leq 4KC' \frac{\log(\frac{2}{\delta})}{\sqrt{n}} b_n^{\frac{m+d/2+\varepsilon}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} \\ \lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n)) \leq 4(2C' + 1) \frac{\log(\frac{2}{\delta})}{\sqrt{n}} b_n^{\frac{d/2+\varepsilon-1}{m+1}} a_n^{\frac{m+2-d/2-\varepsilon}{m+1}}. \end{cases} \quad (82)$$

The first equation implies that $a_n^{\frac{m+d/2+\varepsilon}{m+1}} \leq 4KC' \frac{\log(\frac{2}{\delta})}{\sqrt{n}} b_n^{\frac{m+d/2+\varepsilon}{m+1}}$ which also gives

$$a_n^{\frac{m+2-d/2-\varepsilon}{m+1}} \leq (4KC')^{\frac{m+2-d/2-\varepsilon}{m+d/2+\varepsilon}} b_n^{\frac{m+2-d/2-\varepsilon}{m+1}} \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+2-d/2-\varepsilon}{2m+d+2\varepsilon}}. \quad (83)$$

Hence the upper bound we obtain on $\lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n))$ is

$$\lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n)) \leq 4(2C' + 1)(4KC')^{\frac{m+2-d/2-\varepsilon}{m+d/2+\varepsilon}} b_n \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{1}{2}(1 + \frac{m+2-d/2-\varepsilon}{m+d/2+\varepsilon})} \quad (84)$$

$$= 4(2C' + 1)(4KC')^{\frac{m+2-d/2-\varepsilon}{m+d/2+\varepsilon}} b_n \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}}. \quad (85)$$

Choosing $\lambda_n \geq \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}}$ gives that \hat{R}_n and $\text{Tr}(\hat{A}_n)$ are bounded (independently on δ) and yields the rate

$$a_n^2 \leq (4KC')^{\frac{m+1}{2m+d+2\varepsilon}} b_* \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}}. \quad (86)$$

Case 2. On these indexes, the second equation of (75) becomes

$$\lambda_n(\hat{R}_n^2 + \text{Tr}(\hat{A}_n)) \leq 2(2C' + 1)c_n. \quad (87)$$

Now recall that c_n is given by

$$c_n = \lambda_n(R^2 + \text{Tr}(A_*)) + C_1(\text{Tr}(\hat{A}_n) + \hat{R}_n + G) \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}}, \quad (88)$$

where C_1, G are constants that do not depend of n, δ . Hence, if we choose λ_n such that $\lambda_n \geq 4C_1(2C' + 1) \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}}$, the quantities \hat{R}_n and $\text{Tr}(\hat{A}_n)$ are bounded independently on δ and we recover a_n of the form

$$a_n^2 \leq K' \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}}, \quad (89)$$

where K' is a constant independent on n and δ .

Conclusion. If we set $\lambda_n = \left(\frac{\log(\frac{2}{\delta})^2}{n} \right)^{\frac{m+1}{m+d/2+\varepsilon}} + 4C_1(2C' + 1) \left(\frac{\log(\frac{n}{\delta})}{n} \right)^{\frac{m-d}{2d}}$, we have that in any case, the energies are bounded and we get

$$a_n^2 \leq C\lambda_n, \quad (90)$$

where C is a constant independent on n and δ . ■

B Proofs from Section 4

B.1 Proof of Lemma 11

Proof Not including the PSD constraints on \mathbf{B} , problem (34) admits the following Lagrangian:

$$\mathcal{L}(f, g, \mathbf{B}, \gamma) = \langle f, \hat{w}_\mu \rangle_{\mathcal{H}_X} + \langle g, \hat{w}_\nu \rangle_{\mathcal{H}_Y} + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 + \lambda_2 (\|f\|_{\mathcal{H}_X}^2 + \|g\|_{\mathcal{H}_Y}^2) \quad (91)$$

$$+ \sum_{j=1}^{\ell} \gamma_j (\Phi_j^T \mathbf{B} \Phi_j - f(\tilde{x}_j) - g(\tilde{y}_j) + \langle \tilde{x}_j, \tilde{y}_j \rangle). \quad (92)$$

Canceling the gradients in f and g , we get

$$\begin{aligned} f &= \frac{1}{2\lambda_2} \left(\sum_{i=1}^{\ell} \gamma_i \phi_X(\tilde{x}_i) - \hat{w}_\mu \right) \\ g &= \frac{1}{2\lambda_2} \left(\sum_{i=1}^{\ell} \gamma_i \phi_Y(\tilde{y}_i) - \hat{w}_\nu \right). \end{aligned} \quad (93)$$

Let us now derive the optimality condition on \mathbf{B} : completing the square, we have

$$\begin{aligned} \inf_{\mathbf{B} \in \mathbb{S}_+(\mathbb{R}^\ell)} \sum_{i=1}^{\ell} \gamma_i \Phi_i^T \mathbf{B} \Phi_i + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 &= \inf_{\mathbf{B} \in \mathbb{S}_+(\mathbb{R}^\ell)} \frac{\lambda_1}{2} \langle \mathbf{B}, \mathbf{B} + \frac{2}{\lambda_1} \sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T \rangle \\ &= -\frac{1}{2\lambda_1} \left\| \left(-\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T \right)_+ \right\|_F^2. \end{aligned} \quad (94)$$

Plugging Eq. (93) and Eq. (94) into Eq. (91), we obtain Eq. (37). ■

B.2 Proof of Proposition 12

Proof Dual formulation. It holds

$$\begin{aligned}
& \inf_{\substack{f \in \mathcal{H}_X, g \in \mathcal{H}_Y, \\ \mathbf{B} \in \mathbb{S}_+(\mathbb{R}^d)}} \langle f, \hat{w}_\mu \rangle_{\mathcal{H}_X} + \langle g, \hat{w}_\nu \rangle_{\mathcal{H}_Y} + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 + \lambda_2 (\|f\|_{\mathcal{H}_X}^2 + \|g\|_{\mathcal{H}_Y}^2) \\
& + \frac{\delta}{2\ell} \sum_{j=1}^{\ell} (f(\tilde{x}_j) + g(\tilde{y}_j) - \tilde{x}_j \cdot \tilde{y}_j - \Phi_j^T \mathbf{B} \Phi_j)^2, \\
& = \sup_{\gamma \in \mathbb{R}^\ell} \inf_{\substack{f \in \mathcal{H}_X, g \in \mathcal{H}_Y, \\ \mathbf{B} \in \mathbb{S}_+(\mathbb{R}^d)}} \langle f, \hat{w}_\mu \rangle_{\mathcal{H}_X} + \langle g, \hat{w}_\nu \rangle_{\mathcal{H}_Y} + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 + \lambda_2 (\|f\|_{\mathcal{H}_X}^2 + \|g\|_{\mathcal{H}_Y}^2) \\
& + \sum_{j=1}^{\ell} \gamma_j (\Phi_j^T \mathbf{B} \Phi_j - f(\tilde{x}_j) - g(\tilde{y}_j) + \tilde{x}_j \cdot \tilde{y}_j) - \frac{\ell}{2\delta} \|\gamma\|^2.
\end{aligned} \tag{95}$$

Indeed, the second problem above is a convex-concave min-max problem, and inverting min and max and solving for γ directly yields the original problem. Let us rewrite the inner inf as a function of γ . Canceling the gradients in f and g , we get

$$\begin{aligned}
f &= \frac{1}{2\lambda_2} \left(\sum_{i=1}^{\ell} \gamma_i \phi_X(\tilde{x}_i) - \hat{w}_\mu \right) \\
g &= \frac{1}{2\lambda_2} \left(\sum_{i=1}^{\ell} \gamma_i \phi_Y(\tilde{y}_i) - \hat{w}_\nu \right),
\end{aligned} \tag{96}$$

i.e. the primal-dual relations Eq. (40). Let us now derive the optimality condition on \mathbf{B} : as in the proof of Lemma 11, we have

$$\inf_{\mathbf{B} \in \mathbb{S}_+(\mathbb{R}^\ell)} \sum_{i=1}^{\ell} \gamma_i \Phi_i^T \mathbf{B} \Phi_i + \frac{\lambda_1}{2} \|\mathbf{B}\|_F^2 = -\frac{1}{2\lambda_1} \|(-\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T)_+\|_F^2. \tag{97}$$

Plugging Eq. (40) and Eq. (97) into Eq. (95), we obtain Eq. (39). Let us now derive the smoothness and strong convexity constants of

$$H(\gamma) \stackrel{\text{def}}{=} \frac{1}{4\lambda_2} \gamma^T \mathbf{Q} \gamma - \frac{1}{2\lambda_2} \sum_{j=1}^{\ell} \gamma_j z_j + \frac{1}{2\lambda_1} \|(-\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T)_+\|_F^2 + \frac{\ell}{2\delta} \|\gamma\|^2 + \frac{q^2}{4\lambda_2}. \tag{98}$$

Strong convexity. Let us recall that H is α -strongly convex i.f.f. $\forall \gamma, \gamma', \alpha \|\gamma - \gamma'\|^2 \leq \nabla \langle H(\gamma) - \nabla H(\gamma'), \gamma - \gamma' \rangle$. We have

$$\begin{aligned}
\langle H(\gamma) - \nabla H(\gamma'), \gamma - \gamma' \rangle &= \frac{\ell}{\delta} \|\gamma - \gamma'\|^2 + \frac{1}{2\lambda_2} (\gamma - \gamma')^T \mathbf{Q} (\gamma - \gamma') \\
&+ \frac{1}{\lambda_2} \text{diag}(\Phi^T ([\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T]_- - [\sum_{i=1}^{\ell} \gamma'_i \Phi_i \Phi_i^T]_-) \Phi)^T (\gamma - \gamma').
\end{aligned}$$

Given that the term with negative parts is non-negative but vanishes when the corresponding matrices are positive, we get the lower bound $\alpha \geq \frac{\ell}{\delta} + \frac{1}{2\lambda_2} \lambda_{\min}(\mathbf{Q})$.

Smoothness. Finally, we have

$$\begin{aligned} \|H(\gamma) - \nabla H(\gamma')\| &\leq \left\| \frac{\ell}{\delta}(\gamma - \gamma') \right\| + \frac{1}{2\lambda_2} \|\mathbf{Q}(\gamma - \gamma')\| \\ &\quad + \frac{1}{\lambda_2} \left\| \text{diag}(\Phi^T([\sum_{i=1}^{\ell} \gamma_i \Phi_i \Phi_i^T]_- - [\sum_{i=1}^{\ell} \gamma'_i \Phi_i \Phi_i^T]_-)\Phi) \right\|. \end{aligned}$$

Bounded those three terms independently yields

$$\|H(\gamma) - \nabla H(\gamma')\| \leq \left(\frac{\ell}{\delta} + \frac{1}{2\lambda_2} \lambda_{\max}(\mathbf{Q}) + \frac{1}{\lambda_1} \lambda_{\max}(\mathbf{K} \circ \mathbf{K}) \right) \|\gamma - \gamma'\|.$$

■