



HAL
open science

On the Concentration of the Minimizers of Empirical Risks

Paul Escande

► **To cite this version:**

| Paul Escande. On the Concentration of the Minimizers of Empirical Risks. 2023. hal-04054021v1

HAL Id: hal-04054021

<https://hal.science/hal-04054021v1>

Preprint submitted on 31 Mar 2023 (v1), last revised 10 Sep 2024 (v4)

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.

On the Concentration of the Minimizers of Empirical Risks

Paul Escande *

March 31, 2023

Abstract

Obtaining guarantees on the convergence of the minimizers of empirical risks to the ones of the true risk is a fundamental matter in statistical learning. Instead of deriving guarantees on the usual estimation error, the goal of this paper is to provide concentration inequalities on the distance between the sets of minimizers of the risks for a broad spectrum of estimation problems. In particular, the risks are defined on metric spaces through probability measures that are also supported on metric spaces. A particular attention will therefore be given to include unbounded spaces and non-convex cost functions that might also be unbounded. This work identifies a set of assumptions allowing to describe a regime that seem to govern the concentration in many estimation problems, where the empirical minimizers are stable. This stability can then be leveraged to prove parametric concentration rates in probability and in expectation. The assumptions are verified, and the bounds showcased, on a selection of estimation problems such as barycenters on metric space with positive or negative curvature, subspaces of covariance matrices, regression problems and entropic-Wasserstein barycenters.

1 Introduction

Let (\mathbb{X}, ρ, μ) be a metric probability space¹ and let (\mathcal{M}, ϑ) . Many estimation problems in statistics or in machine learning can be regarded as solving the following generic optimization problem

$$\phi_* \in \arg \min_{\phi \in \mathcal{M}} \mathcal{J}(\phi) = \mathbb{E}_{X \sim \mu}[l(\phi, X)] + \mathcal{R}(\phi) \quad (1)$$

where $l : \mathcal{M} \times \mathbb{X} \rightarrow \mathbb{R}$ is a cost function or loss and $\mathcal{R} : \mathcal{M} \rightarrow \mathbb{R} \cup \{+\infty\}$ is an optional regularization function promoting the a priori structure of the solution. The function $\mathcal{J} : \mathcal{M} \rightarrow \mathbb{R}$ is oftentimes referred to as the risk or the objective function. For convenience, the data fidelity term $\ell : \mathcal{M} \rightarrow \mathbb{R}$ will be defined by $\ell(\phi) = \mathbb{E}_{X \sim \mu}[l(\phi, X)]$.

In practice, the measure μ is rarely known thus hindering the knowledge of \mathcal{J} . However, samples drawn from μ are usually accessible. Let $\mathbf{X} = (X_i)_{i=1}^n \in \mathbb{X}^n$ be such samples that will further be assumed to be independent. The main idea of the Empirical Risk Minimization

*Institut de Mathématiques de Toulouse; UMR 5219, Université de Toulouse, CNRS ; UPS, F-31062 Toulouse Cedex 9, France. paul.escande@math.univ-toulouse.fr

¹a measurable space \mathbb{X} whose Borel σ -algebra is induced by ρ and is endowed with a probability measure μ .

(ERM) [72] is to use these samples to build $\widehat{\ell}$, an estimator of ℓ , as $\widehat{\ell}(\phi) = \frac{1}{n} \sum_{i=1}^n l(\phi, X_i)$ and to solve

$$\widehat{\phi} \in \arg \min_{\phi \in \mathcal{M}} \widehat{\mathcal{F}}(\phi) = \widehat{\ell}(\phi) + \mathcal{R}(\phi) \quad (2)$$

in place of (1).

Understanding when the estimator $\widehat{\phi}$ is legit is a fundamental matter in statistical learning. In this work, we will be particularly interested in proving high probability bounds of kind

$$\mathbb{P} \left(\vartheta(\arg \min \widehat{\mathcal{F}}; \arg \min \mathcal{F}) < g(\delta, n) \right) \geq 1 - \delta \quad \text{for } \delta \in [0, 1], \quad (3)$$

where $\vartheta(A; B)$ will be a suitable distance between sets of \mathcal{M} and $g : [0, 1] \times \mathbb{N} \rightarrow \mathbb{R}_+$ going to zero as n goes to infinity for all fixed δ . This question is very similar to the Probably Approximately Correct (PAC) framework [71] with the important distinction that it proves bound on the estimation (or generalization) error $\mathcal{F}(\widehat{\phi}) - \inf \mathcal{F}$, also known as the generalization error, instead of a metric on the minimizers. More specifically, this work attempts to identify a set of assumptions from which optimal concentration rates of kind (3) can be derived for a broad class of estimation problems. In essence, it is proposed here to understand the underlying mechanisms responsible for the concentration guarantees of many estimation problems.

These fundamental questions have a long history and have led to a remarkably rich theory. In order to precisely contextualize this work, it seemed essential to dedicate a section to present its history and some important works (see Section 1.2). Prior to this literature review, some occurrences of the estimation problems (2) solved in place of (1) are stated with the aim of describing the taxonomy of problems encountered in applications.

1.1 Applications

Some occurrences of (2) are briefly discussed in this section to describe the taxonomy of problems encountered in applications. They will be studied in detail in Section 6.

Barycenters Let $(\mathcal{M}, \vartheta) = (\mathbb{X}, \rho)$ be a metric space, $l(\phi, x) = \rho(x, y)^2$ and \mathcal{R} be identically zero. The problems (1) and (2) become

$$\phi_* \in \arg \min_{\phi \in \mathbb{X}} \int_{\mathbb{X}} d^2(\phi, x) d\mu(x), \quad \widehat{\phi} \in \arg \min_{\phi \in \mathbb{X}} \frac{1}{n} \sum_{i=1}^n d^2(\phi, x_i), \quad (4)$$

which are the Fréchet means of μ and $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ respectively [30]. The properties of the risks \mathcal{F} and $\widehat{\mathcal{F}}$ depend on the convexity and the curvature of the space \mathcal{M} . The risks are known to be strongly geodesically convex when the space \mathcal{M} are of non-positive curvature in the sense of Alexandrov [67]. Otherwise, the risks might not be convex since the distance function might exhibit some concavity properties [4].

Subspace estimation Let $(\mathbb{X}, \rho) = (\mathbb{R}^d, \|\cdot\|_2)$, $\mathcal{M} = \mathbb{S}^{d-1}$ endowed with the great circle distance and \mathcal{R} be identically zero. The problem of estimating the leading direction of variability of $\text{Cov}(\mu) = \mathbb{E}[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^T]$, the covariance of μ , amounts to solving

$$\widehat{\phi} \in \arg \max_{\phi \in \mathbb{S}^{d-1}} \langle \text{Cov}(\mu_n) \phi, \phi \rangle, \quad \text{in place of} \quad \phi_* \in \arg \max_{\phi \in \mathbb{S}^{d-1}} \langle \text{Cov}(\mu) \phi, \phi \rangle$$

which can be seen as instance an of (1) and (2) with the proper choice of l displayed in Section 6.2, and \mathcal{R} being identically zero. These problems are non-convex.

Supervised learning – LASSO Let $\mathbb{Y} = \mathbb{R}^d$ and $\mathbb{V} = \mathbb{R}$ be respectively a space of inputs and outputs. Let $\mathbb{X} = \mathbb{Y} \times \mathbb{V}$ and $\mathcal{M} = \mathbb{R}^m$ endowed with the Euclidean distance. Let $\theta : \mathbb{R}^d \rightarrow \mathbb{R}^m$ be feature vectors, $l : \mathbb{X} \rightarrow \mathbb{R}_+$ defined for all $(y, v) \in \mathbb{X}$ as $l(\phi, (y, v)) = \frac{1}{2}(\langle \phi, \theta(y) \rangle - v)^2$ and $\mathcal{R} = \lambda \|\cdot\|_1$ for $\lambda > 0$. The problems (1) and (2) instantiate to

$$\phi_* \in \arg \min_{\phi \in \mathbb{R}^m} \frac{1}{2} \int_{\mathbb{X}} (\langle \phi, \theta(y) \rangle - v)^2 d\mu(y, v) + \lambda \|\phi\|_1,$$

and

$$\widehat{\phi} \in \arg \min_{\phi \in \mathbb{R}^m} \frac{1}{2} \|V - \Theta\phi\|_2^2 + \lambda \|\phi\|_1,$$

with $V \in \mathbb{R}^n$ being the vectors of outputs $V = (v_i)_{i=1}^n$ and $\Theta \in \mathbb{R}^{n \times m}$ defined by $\Theta_{i,j} = \theta(y_i)_j$. The risks \mathcal{J} and $\widehat{\mathcal{J}}$ are convex, but they may fail to be strongly convex if the matrix $\text{Cov}(\theta_{\#}\mu)$ (respectively with μ_n) has a non-empty kernel. In this case, the risks possess a quadratic growth around the minimizers in the orthogonal of the kernel and a linear growth in the kernel driven by the regularization.

These three examples are simple yet important to understand the assumptions needed to describe the concentration phenomenon of most of optimization problems. In particular, the loss l or the objectives \mathcal{J} and $\widehat{\mathcal{J}}$ cannot be assumed to be strongly convex or even convex. Similarly, the spaces \mathcal{M} and \mathbb{X} at stake might be unbounded.

1.2 Problem history and related works

Solving (2) in place (1) is related to the principle of ERM or of Structural Risk Minimization (SRM). The study of theoretical guarantees for the estimator $\widehat{\phi}$, in terms of bounds of kind (3) for the estimation error, has a long history which can possibly be tracked back to the 1970's with the works [74, 73]. Two main classes of analyses can probably be identified: the techniques using uniform convergence and the one from algorithm stability. They are described in the following two paragraphs and will be linked to the question at stake in the third paragraph i.e. deriving bounds of type (3) for the metric distance between the minimizers. While being insufficient to fully explain the optimal error bounds that are observed in (3), their presentation was found important to articulate the contributions of this work. Moreover, the tools used to achieve the error bounds for (3) share some similarities with these previous works. Finally, it should be mentioned that the connection between uniform convergence and stability have been investigated in various works such as [65] and reference therein.

It is assumed in this section that the minimizers of \mathcal{J} and $\widehat{\mathcal{J}}$ exist and are unique. The minimal value of \mathcal{J} will be referred to as \mathcal{J}_* with similar notation for other functions.

Uniform convergence The techniques in the uniform convergence class can be tracked back to the 1970's with the works of Vapnik and Chervonenkis [74, 73, 72] and the PAC framework proposed in the 1980's [71]. Since then, their seminal works have generated a large body of literature and a mature theory. A complete introduction of these notions can be found in good textbooks such as [26, 63, 77, 75, 6] to name a few.

The starting point of these works is the decomposition of the estimation error

$$\begin{aligned}
\mathcal{J}(\hat{\phi}) - \mathcal{J}_* &\leq \mathcal{J}(\hat{\phi}) - \mathcal{J}(\phi_*) + \widehat{\mathcal{J}}(\phi_*) - \widehat{\mathcal{J}}(\hat{\phi}) \quad \text{since } \hat{\phi} \in \arg \min \widehat{\mathcal{J}} \\
&= \ell(\hat{\phi}) - \ell(\phi_*) + \widehat{\ell}(\phi_*) - \widehat{\ell}(\hat{\phi}) \\
&\leq 2 \sup_{\phi \in \mathcal{M}} \Delta(\phi) \qquad \text{where } \Delta(\phi) = |\ell(\phi) - \widehat{\ell}(\phi)|
\end{aligned} \tag{5}$$

where the uniform bound is taken to remove the dependency between $\hat{\phi}$ and $\widehat{\ell}$ that was preventing the $(\ell(\hat{\phi}, X_i))_{i=1}^n$ to be independent. This allows to use standard concentration tools such as the McDiarmid's inequality [51] (assuming for simplicity that the cost function $0 \leq l \leq M$ for some $M > 0$) leading to the following generic bound on the estimation error

$$\mathcal{J}(\hat{\phi}) - \mathcal{J}_* \leq 2\mathbb{E} \left[\sup_{\phi \in \mathcal{M}} \Delta(\phi) \right] + M \sqrt{\frac{2}{n} \log \left(\frac{2}{\delta} \right)} \tag{6}$$

with high probability at least $1 - \delta$. The downside of this uniform bound is to control $\mathbb{E} [\sup_{\phi \in \mathcal{M}} \Delta(\phi)]$ (i.e. the expectation of the supremum of a random process) which can be quite large and hindering any exploitation of a potential structure of the minimizers. Yet, this control can be achieved in mainly two ways.

First, via symmetrization, the later can further be bounded by the Rademacher complexity or the Gaussian complexity of the function class $\{(\ell(\phi, x_i))_{i=1}^n \mid \phi \in \mathcal{M}\}$ [8]. Note that when the values of l are discrete, the Rademacher complexity can be linked to the VC dimension of the class of functions [20]. The Rademacher complexities of the typical function classes encountered in estimation behave as c/\sqrt{n} , where $c > 0$ is a constant depending on the problem [77, 63].

Second, when the cost l exhibits some smoothness properties (e.g. Lipschitz) with respect to the variable ϕ , the quantity $\mathbb{E} [\sup_{\phi \in \mathcal{M}} \Delta(\phi)]$ can alternatively be bounded using chaining techniques such as Dudley's chaining or the Talagrand's generic chaining [68, 75]. This leads to an upper-bound dependent on a measure of complexity of \mathcal{M} , e.g the entropy of \mathcal{M} , linked to the covering numbers of \mathcal{M} , or the Talagrand's γ_2 functional. These complexities typically behave as $n^{-1/(2s)}$ with $s = 1$ for simple spaces such as Euclidean spaces, or doubling spaces, and with $s > 1$ for more complex spaces [2, 62].

To summarize, the bounds on the estimation error obtained with these techniques of uniform convergence decrease as $n^{-1/2}$. Without further assumption on the cost function l , this rate is optimal [52].

Algorithm stability Another body of literature studies the generalization of learning algorithms through stability. These ideas can be tracked back to the 1970s as well with [60] and later popularized in the founding article [21]. The notions involved in these works will be introduced here for the case of the empirical minimizer, although they can be applied to the broader setting of learning algorithms².

For convenient notation, let $\mathbf{X} = (X_1, \dots, X_n) \in \mathbb{X}^n$ be the random vector containing the n independent sample points, and let $Y \in \mathbb{X}$ be an i.i.d. copy of X_1 . The random

²More precisely, a learning algorithm is a function $\mathcal{A} : \mathbb{X}^n \rightarrow \mathcal{M}$ that maps the given data $\mathbf{X} = (x_i)_{i=1}^n$ to an element of \mathcal{M} which can be the minimizer $\hat{\phi}$ of the empirical risk $\widehat{\mathcal{J}}$ or an approximation of it obtained using any algorithm (e.g gradient descent).

vector \mathbf{X}' is obtained by swapping one element of \mathbf{X} with Y so that $\mathbf{X}' = (Y, X_2, \dots, X_n)$. Furthermore, the empirical risk associated to \mathbf{X}' will be denote by $\tilde{\mathcal{J}}$ and its minimizer by $\tilde{\phi}$. It is important to remark at this point that the risk $\hat{\mathcal{J}}$ and therefore its minimizers are invariant under reordering of the components of the vector \mathbf{X} . The perturbation of \mathbf{X} by swapping any element can therefore be though as exchanging the first component.

The starting point of these works is a different decomposition of the estimation error

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} \left[\mathcal{J}(\hat{\phi}) - \mathcal{J}_* \right] &\leq \mathbb{E}_{\mathbf{X}} \left[\mathcal{J}(\hat{\phi}) - \hat{\mathcal{J}}(\hat{\phi}) \right] && \text{using (5)} \\ &= \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y[l(\hat{\phi}, Y)] - l(\hat{\phi}, X_1) \right] && \text{since } \mathbb{E}[l(\hat{\phi}, X_i)] = \mathbb{E}[l(\hat{\phi}, X_j)] \text{ for all } i, j \\ &= \mathbb{E}_{\mathbf{X}, Y} \left[l(\tilde{\phi}, X_1) - l(\hat{\phi}, X_1) \right] && \text{with } \tilde{\phi} = \arg \min \tilde{\mathcal{J}} \end{aligned} \tag{7}$$

because the roles of Y and X_1 can be switched as being i.i.d. This decomposition emphasizes the link between the estimation error and the *stability* of the empirical minimizer i.e. if it does not vary much when a sample of the data \mathbf{X} is perturbed or removed. Various notions of stability have been developed e.g. the hypothesis stability [60, 37] or the uniform one [21]. The latter is probably the most popular one since it can be used to obtain Gaussian concentration bounds. The empirical minimizer is said to be ϵ -uniform stable if there exists $\epsilon > 0$ such that $|l(\hat{\phi}, x) - l(\tilde{\phi}, x)| \leq \epsilon$ for all $x \in \mathbb{X}$. This stability can be leverage to directly bound the expected estimation error as $\mathbb{E}_{\mathbf{X}} \left[\mathcal{J}(\hat{\phi}) - \mathcal{J}_* \right] \leq \epsilon$, and obtain [21, Theorem 12] that

$$\mathcal{J}(\hat{\phi}) - \mathcal{J}_* \leq \epsilon + (4n\epsilon + M) \sqrt{\frac{1}{n} \log \left(\frac{1}{\delta} \right)} \tag{8}$$

with high probability $1 - \delta$ whenever $0 \leq l \leq M$. This bound is informative as long as the stability $\epsilon \leq cn^{-1}$ for some constant $c > 0$. In contrast of (6), the stability of the empirical minimizers avoids the uniform bound on \mathcal{M} , which is desirable since the empirical solutions may exhibit a very strong structure. Importantly, the bound (8) is *independent* of the complexity of \mathcal{M} .

Note that (8) is also obtained using the McDiarmid's inequality and requires the boundedness of the cost l . This might narrow the range of application of these results. This assumption can however be dropped at the expense of a stronger notion of stability and a bounded sub-Gaussian or sub-exponential diameter of \mathbb{X} [41, 50].

Recently, a series of papers [27, 28, 22] improved (8) to

$$\mathcal{J}(\hat{\phi}) - \mathcal{J}_* \leq \epsilon \log(n) \log \left(\frac{1}{\delta} \right) + M \sqrt{\frac{1}{n} \log \left(\frac{1}{\delta} \right)} \tag{9}$$

with high probability $1 - \delta$. This was considered as a breakthrough in the algorithm stability community as the bound becomes informative as long as $\epsilon \leq cn^{-1/2}$ which extends the validity of the result to a broader variety of applications.

In addition of theses refinements, it is relevant for this work to mention that the estimation error can still be bounded when the stability of the empirical minimizers happens only on a subset of large measure [44, 59].

Obtaining concentration bounds on ϑ Despite a wide interest in bounding the estimation error, little is however known regarding the concentration of the minimizers in the sense of (3), involving a distance on the minimizers.

The first observation is that when a error bound of kind

$$\tau \vartheta(\widehat{\phi}, \phi_*)^\beta \leq \mathcal{F}(\widehat{\phi}) - \mathcal{F}_* \quad (10)$$

is available for some $\beta > 0$ and $\tau > 0$, the concentration of $\vartheta(\widehat{\phi}, \phi_*)$ can be derived using the results of the two previous paragraphs. A direct combination of (10) and (9) in the example of barycenters in a compact and convex $\Omega \subset \mathbb{R}^d$ would give

$$\|\widehat{\phi} - \phi_*\|_2^2 = \mathcal{F}(\widehat{\phi}) - \mathcal{F}_* \leq \epsilon \log(n) \log\left(\frac{1}{\delta}\right) + M \sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} \quad (11)$$

with high probability $1 - \delta$. Assuming for example that the random variables $X \sim \mu$ are sub-Gaussian with variance proxy $\|X\|_{\psi_2}$ (see Preliminaries section or [75]), their concentration property would instead lead to

$$\|\widehat{\phi} - \phi_*\|_2^2 \leq cd \|X\|_{\psi_2}^2 \frac{1}{n} \log\left(\frac{2}{\delta}\right) \quad (12)$$

with high probability $1 - \delta$ and some universal constant $c > 0$ [36]. The rate in n^{-1} in (12) is called *parametric*. Even if the stability of the empirical minimizer can be shown to be $\epsilon = O(n^{-1})$, the sampling error in $\sqrt{n^{-1} \log(2\delta^{-1})}$ would slow down the rate nevertheless. The same slow rate is obtained with the uniform convergence (6).

This gap is known in the literature [64, 65, 39] and is due to the fact that uniform convergence and stability techniques do not incorporate any structure of the problem. In particular, the behavior of \mathcal{F} around its minimizers can be leveraged to obtain faster concentration rates on the estimation error. When the cost function l is strongly convex and Lipschitz, the uniform stability is known to behave as $\epsilon = O(n^{-1})$ [65]. Yet necessary, this improvement is still vain since the rate in (9) would now be controlled by the sampling error in $\sqrt{n^{-1} \log(2\delta^{-1})}$. In [64] the authors asked if a bound on the estimation error for strongly convex and Lipschitz l with rate in $O(n^{-1})$ is possible. Ten years later, this question was answered positively in [39] using an intricate decomposition of the estimation error and an extension of the McDiarmid's inequality so that (9) becomes in this case

$$\mathcal{F}(\widehat{\phi}) - \mathcal{F}_* \leq c \frac{\log(n)}{n} \log\left(\frac{1}{\delta}\right) \quad (13)$$

with high probability $1 - \delta$ where c is a constant depending on the Lipschitz constant and the constant of strong convexity of l . Up to the $\log(n)$ factor, this last result is in par with the optimal rate for the barycenter problem on Euclidean spaces (12). However, the strong convexity and the Lipschitz assumptions on l are rather restrictive and are not satisfied by the examples presented above in Section 1.1.

Recently, two works [62, 2] have explored the concentration of $\vartheta(\widehat{\phi}, \phi_*)$ using uniform convergence techniques. Both assume a growth of \mathcal{F} around its minimizers of type (10) without assuming the strong convexity of l . In [2], concentration bounds are derived in the case where l and $\mathcal{M} = \mathbb{X}$ are bounded, exploiting a Hölder condition for $\phi \mapsto l(\phi, x)$ holding uniformly

on x . This hypothesis is usually accessible in the bounded setting, but might fail when the spaces are unbounded. In [62], the analysis is more general where \mathcal{M} can be different from \mathbb{X} and with \mathcal{M} , \mathbb{X} and l possibly unbounded. To deal with this additional difficulty, the author introduces an interesting weaker notion of smoothness for l through a quadruple inequality, that can be interpreted as a Hölder condition on $\phi \mapsto l(\phi, x)$ with a constant that varies with x . A similar condition will be used in the present work. While providing first insights on the question, the bounds obtained depend on the complexity of \mathcal{M} which might lead to rates suffering from the curse of dimensionality. This can be illustrated through the barycenter problem in non-positively curved spaces. It is known, without using uniform convergence techniques, that parametric rate holds regardless of the complexity of \mathcal{M} [46]. However, the rates obtained in [62] depends on the complexity of \mathcal{M} which translates into parametric rates for doubling metric measure spaces \mathcal{M} and rates that are cursed by the dimension when the metric space is more complex [2, 62].

1.3 Contributions

The review of the literature indicates that available results are insufficient to provide optimal concentration rates for $\vartheta(\widehat{\mathcal{S}}; \mathcal{S})$ for a wide variety of estimation problems, such as those described in Section 1.1. This work attempts to fill that gap by providing an alternative view point of this problem that allows to settle a set of assumptions from which optimal concentration rates can be proven. These assumptions fit most of estimation problems encountered in application where the objective function can be non-convex and defined on unbounded sets. Importantly, the rates obtained are parametric, and do not depend on the complexity of the underlying space.

The paper is organized as follows. Notation and preliminary facts are gathered in the Section 2. The Section 3 introduces the different assumptions needed to prove the concentration rates that will be presented in Section 4. These results will be proved in the Section 5 while being applied to some practical estimation problems in Section 6. To improve the readability of the core of the document, many technical results are postponed to the appendices.

2 Preliminaries

2.1 General facts on random variables

For $p > 0$, the p -th norm of a random variable X on the probability measure space (\mathbb{X}, ρ, μ) is defined by $\|X\|_p = \mathbb{E}[|X|^p]^{1/p}$. This definition can be extended to the case $p = +\infty$ with the essential supremum $\|X\|_\infty = \text{ess sup } |X|$.

For $q \geq 1$, the q -exponential Orlicz norm of a random variable X on the probability space (\mathbb{X}, ρ, μ) is defined by $\|X\|_{\psi_q}$ as $\|X\|_{\psi_q} = \inf_{c>0} \{\mathbb{E}[\exp(|X|/c^q)] \leq 2\}$. The case $q = 1$ corresponds to sub-exponential random variables while $q = 2$ corresponds to sub-Gaussian ones. A random variable having a finite $\|X\|_{\psi_q}$ admits a tail satisfying $\mathbb{P}(|X| \geq t) \leq 2 \exp\left(-\frac{t^q}{\|X\|_{\psi_q}^q}\right)$ and its moments grow as $\|X\|_p \leq \|X\|_{\psi_q} p^{1/q}$. These properties can be shown to be equivalent, see [75] for more details. Importantly, $\|X^2\|_{\psi_1} = \|X\|_{\psi_2}^2$ and $\|XY\|_{\psi_1} \leq \|X\|_{\psi_2} \|Y\|_{\psi_2}$ [75, Lemma 2.7.6 and 2.7.7]. When $X \in \mathbb{R}^d$ is a random vector, its ψ_q norm is defined

by $\|X\|_{\psi_q} = \sup_{v \in \mathbb{S}^{d-1}} \|\langle X, v \rangle\|_{\psi_q}$. In particular, $\| \|X\| \|_{\psi_2} \leq \sqrt{d} \|X\|_{\psi_2}$ [75, Section 3.4] and $\| \|X\|_2 \|_{\psi_1} \leq \| \|X\|_1 \|_{\psi_1} \leq d \|X\|_{\psi_1}$ see [79].

Following [19, Section 2.4], a random variable X is said to be sub-gamma on the right tail with variance factor σ^2 and scale parameter M if its logarithmic moment generating function satisfies

$$\log \mathbb{E} [\exp (t(X - \mathbb{E}X))] \leq \frac{t^2 \sigma^2}{2(1 - Mt)}$$

for every $t \in (0, 1/M)$. The set of all such random variables is denoted by $\text{sub}\Gamma_+(\sigma^2, M)$. The random variable X is furthermore said sub-gamma on the left tail with variance factor σ^2 and scale parameter M if $-X \in \text{sub}\Gamma_+(\sigma^2, M)$. Similarly, the set of all such random variables is denoted by $\text{sub}\Gamma_-(\sigma^2, M)$. Lastly, a random variable X is said to be sub-gamma with variance factor σ^2 and scale parameter M if $X \in \text{sub}\Gamma_+(\sigma^2, M) \cap \text{sub}\Gamma_-(\sigma^2, M)$. The set of all such random variables is denoted by $\text{sub}\Gamma(\sigma^2, M)$. These sub-gamma properties can be characterized in terms of the tail. A random variable $X \in \text{sub}\Gamma_+(\sigma^2, M)$ satisfies

$$\mathbb{P} \left(X > \sqrt{2\sigma^2 t} + Mt \right) \leq \exp(-t)$$

with a similar tail on $-X$ for $X \in \text{sub}\Gamma_-(\sigma^2, M)$. In particular, this tail condition implies that

$$\mathbb{P}(X > t) \leq \exp \left(-\frac{-t^2}{2(\sigma^2 + Mt)} \right).$$

2.2 Notation

In this work, the function $l : \mathcal{M} \times \mathbb{X} \rightarrow \mathbb{R}$ is assumed to be such that $l(\cdot, \phi)$ are measurable and $\mathbb{E}[|l(X, \phi)|] < +\infty$ for every $\phi \in \mathcal{M}$.

Let $(X_i)_{i=1}^n$ be random variables sampled independently from μ . The vector of \mathbb{X}^n will be denoted with bold font i.e., $\mathbf{X} = (X_1, \dots, X_n)$ for random vectors and $\mathbf{x} = (x_1, \dots, x_n)$ for vectors. Let Y be an i.i.d. copy of X_1 . The random vector \mathbf{X}' is defined by $\mathbf{X}' = (Y, X_2, \dots, X_n)$ i.e. it is obtained by swapping the first element of \mathbf{X} .

For convenient short hand notation, we will use $\mathcal{S} = \arg \min \mathcal{F}$ while \mathcal{F}_* will refer to its minimal value. The empirical risk depends on the samples \mathbf{X} . This dependency will not be made explicit but will be clear from the context.

Let \mathcal{E}_1 and \mathcal{E}_2 two subsets of \mathcal{M} , the notion of similarity between two sets used in this work is

$$\vartheta(\mathcal{E}_1; \mathcal{E}_2) = \sup_{\phi \in \mathcal{E}_1} \vartheta(\phi_1, \mathcal{E}_2) = \sup_{\phi_1 \in \mathcal{E}_1} \inf_{\phi_2 \in \mathcal{E}_2} \vartheta(\phi_1, \phi_2).$$

Note that $\vartheta(\cdot; \cdot)$ is not a metric since it is not symmetric and $\vartheta(\mathcal{E}_1; \mathcal{E}_2) = 0$ only implies that $\mathcal{E}_1 \subseteq \mathcal{E}_2$. However, $\vartheta(\cdot; \cdot)$ satisfies the triangle inequality.

To further simplify the notation, we write $[\mathcal{F} \leq M] = \{\phi \in \mathcal{M} \mid \mathcal{F}(\phi) \leq M\}$ and similar notation can be guessed from the context, e.g., $[\mathcal{F} \geq m] = \{\phi \in \mathcal{M} \mid \mathcal{F}(\phi) \geq m\}$ and $[m \leq \mathcal{F} \leq M] = \{\phi \in \mathcal{M} \mid m \leq \mathcal{F}(\phi) \leq M\}$.

A metric space (\mathcal{M}, ϑ) is said geodesic if for any two points $\phi, \psi \in \mathcal{M}$, there exists a path linking them whose length equals $\vartheta(\phi, \psi)$. For a proper introduction on these concepts see for example [67].

Let (\mathcal{M}, ϑ) be a geodesic space. It has positive curvature in the sense of Alexandrov, if for every constant speed geodesic $\gamma : [0, 1] \rightarrow \mathcal{M}$ and every $\phi \in \mathcal{M}$ the concavity inequality holds:

$$\vartheta^2(\gamma(t), \phi) \geq (1-t)\vartheta^2(\gamma(0), \phi) + t\vartheta^2(\gamma(1), \phi) - t(1-t)\vartheta^2(\gamma(0), \gamma(1)). \quad (14)$$

Similarly, \mathcal{M} is said to be a non positively curved space, in the sense of Alexandrov, if the converse inequality holds. The equality holds when \mathcal{M} is a Hilbert space.

3 Assumptions

In this section, the assumptions needed to derive the main concentration result are described.

3.1 Existence of minimizers

Studying the concentration of the minimizers of \mathcal{F} and $\widehat{\mathcal{F}}$ is possible only if those exist. Their existence is formalized by the following assumption.

Assumption 1. *The function $\mathcal{F} : \mathcal{M} \rightarrow \mathbb{R} \cup \{+\infty\}$ is proper i.e. such that $\text{dom}\mathcal{F} = \{\phi \in \mathcal{M} \mid \mathcal{F}(\phi) \neq +\infty\} \neq \emptyset$ and have at least one minimizer. Furthermore, for all $n \geq 1$ and almost every $\mathbf{x} \in \mathbb{X}^n$, the empirical risks $\widehat{\mathcal{F}}$, associated to \mathbf{x} , are also assumed to be proper with at least one minimizer. Its minimizers are all assumed to be measurable.*

When the function \mathcal{F} (respectively $\widehat{\mathcal{F}}$) is lower semi-continuous and either \mathcal{M} is bounded or \mathcal{F} (respectively $\widehat{\mathcal{F}}$) is coercive, at least one minimizer exists.

3.2 Local Hölder-type error bound

The following assumption describes the growth of \mathcal{F} locally around its minimizers and is a generalization of (10).

Assumption 2. *There exist $\mathcal{F}_0 \in (\mathcal{F}_*, +\infty]$, $\beta \geq 1$ and $\tau > 0$ such that for all $\phi \in \mathcal{M}$,*

$$\phi \in (\mathcal{F}_* < \mathcal{F} \leq \mathcal{F}_0] \implies \tau\vartheta(\phi, \mathcal{S})^\beta \leq \mathcal{F}(\phi) - \mathcal{F}_* \quad (15)$$

The bound (15) is a central ingredient in this work as it avoids pathological oscillations of \mathcal{F} around its minimizers and it provides leverage to use concentration techniques on the function values $\mathcal{F}(\widehat{\phi}) - \mathcal{F}_*$ to bound ϑ .

It is important to mention that many functions used in applications satisfy Assumption 2. For example, any strongly convex function \mathcal{F} satisfies the assumption with $\beta = 2$ and $\mathcal{F}_0 = +\infty$. The assumption can still hold even when the function is not convex. An argument in this direction is that when $\mathcal{M} = \mathbb{R}^d$ and \mathcal{F} is a semi-algebraic function [17], the bound (15) holds on any compact subset [49, Theorem II]. Semi-algebraic functions contains for examples piecewise polynomials, square root, quotients, norms, relu, rank norm to name a few. As a consequence of Tarski Seidenberg theorem [9, Chapter 2], these functions are furthermore stable under many operations such as differentiation or composition. For these reasons, the Assumption 2 covers most of the functions encountered in applications where $\mathcal{M} = \mathbb{R}^d$. Extension of this result for semi-analytic functions to Riemannian analytic manifolds can be found in [14, Theorem 6.4].

This assumption appears in many communities under different names such as local Hölder-type bound in functional analysis [48], Łojasiewicz inequality in real algebraic geometry [49] or Łojasiewicz error bound inequality in optimization [18]. In the particular case of Fréchet means where $\mathbb{X} = \mathcal{M}$ and $l = \rho^2$, the Assumption 2 is known as variance inequality [67]. It is also called a margin condition in [7] or a low noise assumption in [2].

Remark 1. *The Assumption 2 is equivalent to assuming that there exist $\mathcal{J}_0 \in (\mathcal{J}_*, +\infty]$ and $\beta > 0$ with the quantity*

$$\tau = \inf \left\{ \frac{\mathcal{F}(\phi) - \mathcal{J}_*}{\vartheta(\phi, \mathcal{S})^\beta}, \quad \phi \in [\mathcal{F} \leq \mathcal{J}_0] \right\}$$

being positive. Such a quantity might however be delicate to compute and only a positive lower-bound of τ can be accessible. The constant β^{-1} is often referred to as the Łojasiewicz exponent in the optimization literature [18]. To the best of the author's knowledge, there is no particular name for τ , so it will be named Łojasiewicz constant of \mathcal{F} .

3.3 Convergence to the basin

To control $\vartheta(\widehat{\mathcal{S}}; \mathcal{S})$ leveraging Assumption 2, the minimizers of $\widehat{\mathcal{F}}$ have to fall in the level-set $[\mathcal{F} \leq \mathcal{J}_0]$. This event will be assumed to hold with high probability.

Assumption 3. *There exists $\eta : \mathbb{N} \rightarrow [0, 1]$ with $\lim_{n \rightarrow +\infty} \eta(n) = 0$ such that*

$$\mathbb{P} \left(\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} \mathcal{F}(\widehat{\phi}) > \mathcal{J}_0 \right) \leq \eta(n).$$

This assumption is rather natural since as the number of points increases, the empirical risk $\widehat{\mathcal{F}}$ is expected to better approach the true risk. The value of $\mathcal{F}(\widehat{\phi})$ is therefore expected to converge to the minimum \mathcal{J}_* so that $\widehat{\phi}$ would fall into the level-set $[\mathcal{F} \leq \mathcal{J}_0]$.

The technique used to derive an Assumption 3, that is optimal, depends on the context and is therefore left to be checked for each application. Let us however mention some techniques that can be used. Following (5), if $\sup_{\phi \in \mathcal{M}} \Delta(\phi) \leq t$ holds with high probability, e.g. when the Rademacher complexity of the function class or the entropy of \mathcal{M} are bounded, then setting $t = (\mathcal{J}_0 - \mathcal{J}_*)/2$ would imply the Assumption 3. When concentration-based argument are not available, for example when the diameter of \mathcal{M} is unbounded, the small-ball method [53, 54] might be used to prove that Assumption 3 holds. Finally, the convergence to the basin might also be proved considering approaches like [62].

3.4 Łojasiewicz constant of $\widehat{\mathcal{F}}$

Let $\mathbf{X} \in \mathbb{X}^n$ and define with a slight abuse of notation $\mathbf{X}' = (y, X_2, \dots, X_n)$ for a fixed $y \in \text{supp}(\mu)$. Let $\widetilde{\mathcal{F}}$ respectively $\widetilde{\mathcal{S}}$ refer to the empirical risk and empirical set of minimizers associated to \mathbf{X}' , so that for a realization of \mathbf{X} , $\widetilde{\mathcal{S}}$ implicitly depends on y . Let $\widehat{\tau}$ be the random variable defined by

$$\widehat{\tau} = \inf \left\{ \frac{\widehat{\mathcal{F}}(\phi) - \widehat{\mathcal{J}}_*}{\vartheta(\phi, \widehat{\mathcal{S}})^\beta}, \quad \phi \in \bigcup_y \left([\mathcal{F} \leq \mathcal{J}_0] \cap \widetilde{\mathcal{S}} \right) \right\}.$$

The constant $\hat{\tau}$ quantifies the behavior of $\hat{\mathcal{F}}$, around its minimizers, regarding only the minimizers of the empirical risk obtained by perturbing exactly one sample point of \mathbf{X} .

The next assumption will be crucial in proving a notion of stability of the empirical minimizers.

Assumption 4. *There exists $\kappa : \mathbb{N} \rightarrow [0, 1]$ with $\lim_{n \rightarrow +\infty} \kappa(n) = 0$ such that*

$$\mathbb{P}\left(\hat{\tau} < \frac{\tau}{2}\right) \leq \kappa(n).$$

Remark 2. *Similarly to Remark 1, it might only be possible to compute a lower bound to $\hat{\tau}$. In this assumption, the lower bound however has to be sufficiently large to be larger than $\frac{1}{2}\tau$ with high probability. Actually, the only hypothesis needed is that the lower-bound is larger than a constant with high probability. The constant here is chosen to be $\tau/2$ in order not to overload the bounds of Theorem 2.*

Remark 3. *This Assumption 4 is weaker than supposing that*

$$\inf \left\{ \frac{\hat{\mathcal{F}}(\phi) - \hat{\mathcal{F}}_*}{\vartheta(\phi, \hat{\mathcal{S}})^\beta}, \quad \phi \in [\mathcal{F} \leq \mathcal{F}_0] \right\} \geq \frac{1}{2}\tau$$

with high probability, which involves the infimum on the whole level-set. The proof of the main result only requires the local error-bound (15) to hold for $\hat{\mathcal{F}}$ regarding only the minimizers of $\hat{\mathcal{F}}$.

A general treatment of the phenomenon assumed in the Assumption 4 and providing the weakest conditions under which it holds is, to the knowledge of the author, an open question and is postponed to future works. However, a systematic analysis shows, in the examples, that this assumption is verified. Note that in some applications, ℓ might be convex and a regularization \mathcal{R} is added in order to ensure the uniqueness of the minimizer and to improve the resolution of (1) with numerical methods. This regularization is usually strongly convex and thus drives the Łojasiewicz constants of \mathcal{F} and $\hat{\mathcal{F}}$.

3.5 Smoothness of l

The concentration phenomenon can occur when the integrand l possesses some smoothness. The smoothness needed is described in the following assumption.

Assumption 5. *There exist $\alpha > 0$, a pseudometric³ $a : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_+$ satisfying $\|a\|_{\psi_1} < +\infty$, and a subset $\Upsilon \subseteq \mathcal{M}$ such that*

$$\phi, \psi \in \Upsilon \implies l(\phi, x) - l(\psi, x) - l(\phi, y) + l(\psi, y) \leq a(x, y)\vartheta(\phi, \psi)^\alpha \quad (16)$$

for μ -almost every $x, y \in \mathbb{X}$. The set Υ should contain \mathcal{S} and it is further assumed that there exists $\iota : \mathbb{N} \rightarrow [0, 1]$ with $\lim_{n \rightarrow +\infty} \iota(n) = 0$ such that $\mathbb{P}(\hat{\mathcal{S}} \not\subseteq \Upsilon) \leq \iota(n)$.

³A pseudometric on \mathbb{X} is a function from $\mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_+$ satisfying the same axioms of a metric except the implication $a(x, y) \implies x = y$.

The assumption made here can be interpreted as a Hölder condition on $l(\cdot, x)$ with a constant that may vary with x . Since \mathbb{X} is a metric space, this constant has to be measured with respect to a reference point. This translates to assuming a Hölder condition on $l(\cdot, x) - l(\cdot, y)$ with constant $a(x, y)$. This function a is required to be sub-exponential.

Remark that when $l(x, \cdot)$ satisfy a Hölder condition i.e. when there exist $K > 0$ and $\alpha > 0$ such that for all $x \in \mathbb{X}$, $|l(\phi, x) - l(\psi, x)| \leq K\vartheta(\phi, \psi)^\alpha$ for all $\psi, \phi \in \mathcal{M}$ then Assumption 5 holds with $a(x, y) = 2K$ and $\iota = 0$. Such a bound is oftentimes available when $\text{diam}(\mathbb{X}) < +\infty$.

The inequality in (16), when holding on the whole space $\Upsilon = \mathcal{M}$, is also called a *quadruple inequality* in [62]. In the Fréchet mean setting i.e. when $\mathcal{M} = \mathbb{X}$ and $l = \rho^2$ and if this condition holds for all $\phi, \psi \in \mathcal{M}$ with $a(x, y) = 2\rho(x, y)$, then Assumption 5 characterizes metric spaces with non-positive curvature [67, Definition 2.1].

Remark 4. *The class of pseudometric functions includes distances and additively separable functions of the form $a(x, y) = (h(x) + h(y)) \mathbb{1}_{x \neq y}$ with $h : \mathbb{X} \rightarrow \mathbb{R}_+$.*

4 Main results

It is now time to present the main results of this work. The concentration in probability of $\vartheta(\widehat{\mathcal{S}}; \mathcal{S})$ will be derived leveraging the following result.

Theorem 1. *Let $f : \mathbb{X}^n \rightarrow \mathbb{R}$ and $\mathcal{B} \subseteq \mathbb{X}^n$ such that $p = \mathbb{P}(\mathbf{X} \notin \mathcal{B}) \leq 3/4$. Assume there exist a pseudometric $b : \mathbb{X} \rightarrow \mathbb{X} \rightarrow \mathbb{R}^+$ with $\|b\|_{\psi_1} < +\infty$ such that*

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq \sum_{i=1}^n b(x_i, y_i), \quad \text{for all } \mathbf{x}, \mathbf{y} \in \mathcal{B}. \quad (17)$$

Then

$$f(\mathbf{X}) - \mathbb{E}[f | \mathbf{X} \in \mathcal{B}] \leq 4n\|b\|_{\psi_1}\sqrt{p} + e\|b\|_{\psi_1} \left(2\sqrt{n \log\left(\frac{1}{\delta}\right)} + \log\left(\frac{1}{\delta}\right) \right)$$

with probability at least $1 - p - \delta$.

This result is a generalization of McDiarmid's inequality when the differences are sub-exponential with high probability. In the literature, the derivation of such inequalities when differences are assumed bounded with high-probability has been studied in a few works [51, 43, 44, 25]. This result is a combination of [50] and [25] for which a proof is given in Section 5.

In the present case of the concentration of $\vartheta(\widehat{\mathcal{S}}; \mathcal{S})$, the subset $\mathcal{A} \subseteq \mathbb{X}^n$ of interest, that will be used as \mathcal{B} in the Theorem, will be defined by

$$\mathcal{A} = \left\{ \mathcal{F}(\widehat{\phi}) \leq \mathcal{F}_0 \right\} \cap \left\{ \widehat{\tau} \geq \frac{1}{2}\tau \right\} \cap \left\{ \widehat{\mathcal{S}} \subseteq \Upsilon \right\}. \quad (18)$$

The remainder of this section consists in computing the conditional expectation and checking the condition (17). Both will be achieved by taking advantage of the *stability* of the empirical minimizers in \mathcal{A} , in the sense presented in the next paragraph. In the rest of this section, the probability of being outside of \mathcal{A} will be referred to as $p_n = \mathbb{P}(\mathbf{X} \notin \mathcal{A})$ and is controlled using Assumptions 3, 4 and 5 by $p_n \leq \eta(n) + \kappa(n) + \iota(n)$.

4.1 The empirical minimizers are stable

The cornerstone of the results obtained in this work is the stability of the empirical minimizers of $\widehat{\mathcal{F}}$ when a data point of \mathbf{X} is perturbed. Conditionally on the event $\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}$, the minimizers in $\widehat{\mathcal{S}}$ and $\widetilde{\mathcal{S}}$ fall into the level set $(\mathcal{F}_* < \mathcal{F} \leq \mathcal{F}_0]$. Furthermore, conditionally on this event, the Łojasiewicz constant $\widehat{\tau} \geq \frac{1}{2}\tau$ so that, for a minimizer $\widetilde{\phi} \in \widetilde{\mathcal{S}}$,

$$\begin{aligned}
\frac{\tau}{2}\vartheta(\widetilde{\phi}, \widehat{\mathcal{S}})^\beta &\leq \widehat{\mathcal{F}}(\widetilde{\phi}) - \widehat{\mathcal{F}}_* \\
&\leq \inf_{\phi \in \widehat{\mathcal{S}}} \widehat{\mathcal{F}}(\widetilde{\phi}) - \widehat{\mathcal{F}}(\phi) + \widetilde{\mathcal{F}}(\phi) - \widetilde{\mathcal{F}}(\widetilde{\phi}) && \text{since } \widetilde{\phi} \text{ is a minimizer of } \widetilde{\mathcal{F}} \\
&= \inf_{\phi \in \widehat{\mathcal{S}}} \widehat{\ell}(\widetilde{\phi}) - \widehat{\ell}(\phi) + \widetilde{\ell}(\phi) - \widetilde{\ell}(\widetilde{\phi}) && \text{since the terms in } \mathcal{R} \text{ cancel out} \\
&= \inf_{\phi \in \widehat{\mathcal{S}}} \frac{1}{n} \left(l(\widetilde{\phi}, X_1) - l(\phi, X_1) + l(\phi, Y) - l(\widetilde{\phi}, Y) \right) \\
&\leq \frac{1}{n} a(X_1, Y) \vartheta(\widehat{\phi}, \widetilde{\mathcal{S}})^\alpha && \text{from Assumption 5.}
\end{aligned}$$

Since this bound holds for any $\widetilde{\phi} \in \widetilde{\mathcal{S}}$, the following lemma is derived.

Lemma 1. *Conditionally on the event $\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}$, it holds that*

$$\vartheta(\widetilde{\mathcal{S}}; \widehat{\mathcal{S}})^{\beta-\alpha} \leq \frac{2}{\tau} \frac{1}{n} a(X_1, Y).$$

This notion of stability of the empirical minimizers is, in essence, similar to the notions of stability of [21]. Since the probability p_n is controlled by Assumptions 3, 4 and 5, the stability happens with high-probability, and echoes the works [44].

This notion of stability can then be exploited to bound the expectation of $\vartheta^\beta(\widehat{\mathcal{S}}; \mathcal{S})$ conditionally on $\mathbf{X} \in \mathcal{A}$.

Lemma 2. *Under Assumptions 1, 2, 3, 4 and 5 with the further condition that $p_n \leq 3/4$ then*

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A} \right] \leq L^{\frac{\beta}{\beta-\alpha}} \left(c_1 n^{-\frac{\alpha}{\beta-\alpha}} + K \sqrt{p_n} \right),$$

with

$$\begin{aligned}
L &= \tau^{-1} \|a\|_{\psi_1} & K &= 2^{\max(0, \alpha-1)+2} \left(L^{\frac{-\alpha}{\beta-\alpha}} \text{diam}(\mathcal{S})^\alpha + c_2 \right) \\
c_1 &= c_1(\alpha, \beta) = \left(\frac{4\beta}{\beta-\alpha} \right)^{\frac{\beta}{\beta-\alpha}} & c_2 &= c_2(\alpha, \beta) = 2 \max \left(\frac{4\alpha}{\beta-\alpha}, 1 \right)^{\frac{\alpha}{\beta-\alpha}}
\end{aligned}$$

Similarly to (7), the proof is essentially based on switching the variables X and X_1 in the expectation. This step allows to leverage the stability of the empirical minimizers of Lemma 1. However, due to the conditioning on $\mathbf{X} \in \mathcal{A}$, the random variables do not share the same distribution anymore so that extra care must be taken. Note that the function l or the spaces \mathcal{M} or \mathbb{X} can be unbounded which require additional precautions compared to the usual argument in [21, 27, 28, 22]. For these reasons, the proof of this lemma is postponed to Section 5.3.

The next Corollary is an instance of Lemma 2 for particular values of α and β .

Corollary 1. When $\beta = 2$ and $\alpha = 1$, the bound of Lemma 2 reads

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A} \right] \leq 2^6 L^2 n^{-1} + 4 (L \operatorname{diam}(\mathcal{S}) + 8L^2) \sqrt{p_n}.$$

When \mathcal{S} contains a unique minimizer, the bound further simplifies to

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A} \right] \leq 2^6 L^2 n^{-1} + 2^5 L^2 \sqrt{p_n}.$$

4.2 Concentration in probability

We now state the main results of this paper.

Theorem 2. Suppose Assumptions 1, 2, 3, 4, 5 hold, and assume furthermore that $p_n \leq 3/4$. For $0 < \beta - \alpha < 2$, let $q = \min(\beta - \alpha, 1)$, $Q = \max(\beta - \alpha, 1)$ and $s = \min\left(1, \frac{2}{\beta - \alpha} - 1\right)$, then

$$\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^q \leq L^{\frac{1}{Q}} \left(c_1^{\frac{q}{\beta}} n^{-\frac{\alpha}{\beta Q}} + 2^{\frac{1}{Q}} e \left(2\sqrt{\frac{1}{n^s} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n^s} \log\left(\frac{1}{\delta}\right) \right) \right) + C p_n^{\frac{q}{2\beta}} \quad (19)$$

with probability at least $1 - p_n - \delta$ where $C = L^{\frac{1}{Q}} \left(K^{\frac{q}{\beta}} + 2^{2+\frac{1}{Q}} n^{1-\frac{1}{Q}} \right)$ where the constants c_1 , K and L are defined in Lemma 2.

Proof. Let $f(\mathbf{X}) = \vartheta(\widehat{\mathcal{S}}; \mathcal{S})^q$. The proof relies on the fact that

$$\mathbb{P}(f(\mathbf{X}) > t) \leq \mathbb{P}(f(\mathbf{X}) - \mathbb{E}[f \mid \mathbf{X} \in \mathcal{A}] > (t - M)_+)$$

for all $t > 0$ and where M is such that $\mathbb{E}[f \mid \mathbf{X} \in \mathcal{A}] \leq M$. From this observation, the McDiarmid's inequality of Theorem 1 is applied to bound $f(\mathbf{X}) - \mathbb{E}[f \mid \mathbf{X} \in \mathcal{A}]$ with high probability.

The upper-bound M is obtained using Lemma 2 as

$$\begin{aligned} \mathbb{E}[f(\mathbf{X}) \mid \mathbf{X} \in \mathcal{A}] &\leq \mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A} \right]^{\frac{q}{\beta}} && \text{using Jensen's inequality } (q \leq \beta) \\ &\leq L^{\frac{1}{Q}} \left(c_1^{\frac{q}{\beta}} n^{-\frac{\alpha}{\beta Q}} + K^{\frac{q}{\beta}} p_n^{\frac{q}{2\beta}} \right) && \text{since } t > 0 \mapsto t^{\frac{q}{\beta}} \text{ is subadditive,} \end{aligned}$$

where we also have used the fact that $q/(\beta - \alpha) = 1/Q$.

The remaining of the proof consists in constructing the pseudometric b satisfying the conditions of Theorem 1. With a slight abuse of notation let $\widehat{\mathcal{S}}$ and $\widetilde{\mathcal{S}}$ be the sets of minimizers of $\widehat{\mathcal{F}}$ defined from \mathbf{x} and $\widetilde{\mathcal{F}}$ from \mathbf{y} respectively. When $\mathbf{x}, \mathbf{y} \in \mathcal{A}$,

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq \left| \vartheta(\widehat{\mathcal{S}}; \mathcal{S})^q - \vartheta(\widetilde{\mathcal{S}}; \mathcal{S})^q \right| \leq \vartheta(\widehat{\mathcal{S}}; \widetilde{\mathcal{S}})^q \leq \sum_{i=1}^n \left(\frac{2}{n\tau} a(x_i, y_i) \right)^{\frac{1}{Q}}$$

using the triangle inequality combined with Lemma 1. Therefore, the choice $b = \left(\frac{2}{n\tau} a\right)^{\frac{1}{Q}}$ satisfies the assumptions of Theorem 1 so that

$$\begin{aligned} f(\mathbf{X}) - \mathbb{E}[f \mid \mathbf{X} \in \mathcal{A}] &\leq 2^{2+\frac{1}{Q}} L^{\frac{1}{Q}} n^{1-\frac{1}{Q}} \sqrt{p_n} + 2^{\frac{1}{Q}} e L^{\frac{1}{Q}} n^{-\frac{1}{Q}} \left(2\sqrt{n \log\left(\frac{1}{\delta}\right)} + \log\left(\frac{1}{\delta}\right) \right) \\ &\leq 2^{2+\frac{1}{Q}} L^{\frac{1}{Q}} n^{1-\frac{1}{Q}} \sqrt{p_n} + 2^{\frac{1}{Q}} e L^{\frac{1}{Q}} \left(2\sqrt{\frac{1}{n^s} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n^s} \log\left(\frac{1}{\delta}\right) \right) \end{aligned}$$

with probability at least $1 - p_n - \delta$ with $s = \min\left(1, \frac{2}{\beta - \alpha} - 1\right)$. This bound combined with the bound on M together with the fact that $\frac{q}{2\beta} < \frac{1}{2}$ proves the claimed result. \square

The notable case $\beta = 2$ and $\alpha = 1$ is a direct consequence of this result.

Corollary 2. *When $\beta = 2$ and $\alpha = 1$, the bound of Theorem 2 reads*

$$\vartheta(\widehat{\mathcal{S}}; \mathcal{S}) \leq 8L \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n} \log\left(\frac{1}{\delta}\right) \right) + Cp_n^{\frac{1}{4}}$$

with probability at least $1 - p_n - \delta$, and $C = \left(14L + 2\sqrt{L \text{diam}(\mathcal{S})}\right)$.

The following comments can be drawn from these result. As long as β and α are independent of the complexity of \mathcal{M} , the rate of convergence will also be *dimension-free*. This comes at the expense of additional terms driven by the probability of not falling into the basin \mathcal{A} where the empirical minimizers are stable. When p_n decays sufficiently fast, the quantity $n^{1-\frac{1}{Q}} \sqrt{p_n}$ goes to zero faster than n^{-s} and becomes negligible for large enough n . Actually when $\beta - \alpha \leq 1$, the quantity $n^{1-\frac{1}{Q}} = 1$ so that $\sqrt{p_n}$ should decay faster than $n^{-s} = n^{-\frac{1}{2}}$ or $n^{-\frac{\alpha}{\beta Q}} = n^{-\frac{\alpha}{\beta}}$ which is a rather mild condition since in the considered applications, p_n will decay as $\exp(-cn)$ for some $c > 0$.

When $\beta - \alpha$ approaches 2, the rate deteriorates. It is not clear to the author if this phenomena is to be expected or is an artifact introduced by the tools used in the proofs. On the one hand, this can be due to a parameter β getting larger meaning that the functions \mathcal{J} and $\widehat{\mathcal{J}}$ are becoming flatter, which naturally slows down the rate. On the other hand, this could also be interpreted as a lack of regularity of the problem since α would be too small compared to β . In case the rate becomes too slow, techniques like [62] might be able to provide a usable rate.

The optimality of the rates Theorem 2 is difficult to be assessed because of its level of generality. However, its application to various scenarii provides the optimal parametric rates independently of the complexity of the underlying space \mathcal{M} .

4.3 Concentration in expectation

When $p_n = 0$, the Lemma 2 used to prove Theorem 2 also provides a bound on the expectation of $\vartheta^\beta(\widehat{\mathcal{S}}; \mathcal{S})$ of kind

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \right] \leq c_1 (\tau^{-1} \|a\|_{\psi_1})^{\frac{\beta}{\beta-\alpha}} n^{-\frac{\alpha}{\beta-\alpha}}.$$

with $c_1 = c_1(\alpha, \beta)$ defined in Lemma 2. Additional developments are however needed to obtain a bound on the expectation $\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \right]$ for $r > \beta$ or whenever $p_n > 0$. The total law of expectation gives

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \right] \leq \mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \mid \mathbf{X} \in \mathcal{A} \right] + \mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \mathbf{1}_{\mathbf{X} \notin \mathcal{A}} \right].$$

so that a bound can be obtained when the expectation of $\vartheta^r \mathbf{1}_{\mathbf{X} \notin \mathcal{A}}$ can be bounded. On the other hand, the expectation of ϑ^r conditionally on $\mathbf{X} \in \mathcal{A}$ for $r > \beta$ can be obtained by integrating its tail (19). The complete computations are omitted in this paper as the applications detailed in Section 6 will not heavily rely on this result.

5 Proof of the main results

It is now time to prove the supporting results needed for the derivation of Theorem 2.

5.1 General concentration results

The first result is adapted from [50], a generalization of the McDiarmid's inequality.

Definition 1 ([50, Definition 1]). *For a function $f : \mathbb{X}^n \rightarrow \mathbb{R}$, a vector $\mathbf{x} \in \mathbb{X}^n$ and a random vector $\mathbf{X} \in \mathbb{X}^n$ with its components sampled independently, the k -th centered conditional version of f is the random variable*

$$\begin{aligned} f_k(\mathbf{X})(\mathbf{x}) &= f(x_1, \dots, x_{k-1}, X_k, x_{k+1}, \dots, x_n) - \mathbb{E}[f(x_1, \dots, x_{k-1}, X'_k, x_{k+1}, \dots, x_n)] \\ &= \mathbb{E}[f(x_1, \dots, x_{k-1}, X_k, x_{k+1}, \dots, x_n) - f(x_1, \dots, x_{k-1}, X'_k, x_{k+1}, \dots, x_n) | X_k]. \end{aligned}$$

The conditional version f_k describes the fluctuations of f in its k -th component, given the other variables $(x_i)_{i \neq k}$. Some properties of the conditional version of f are extracted from [50] to improve its introduction. Note that $f_k(\mathbf{X}) : \mathbf{x} \in \mathbb{X}^n \mapsto f_k(\mathbf{X})(\mathbf{x})$ is a random function that does not depend on the k -th component of \mathbf{x} . Considering a norm $\|\cdot\|_\psi$ on random variables, the function $\|f_k(\mathbf{X})\|_\psi : \mathbf{x} \in \mathbb{X}^n \mapsto \mathbb{R}_+$ is defined by $\|f_k(\mathbf{X})\|_\psi(\mathbf{x}) = \|f_k(\mathbf{X})(\mathbf{x})\|_\psi$. Remark that $\|f_k(\mathbf{X})\|_\psi(\mathbf{X})$ is a random variable and that $f_k(\mathbf{X})(\mathbf{X}) = f(\mathbf{X}) - \mathbb{E}[f(\mathbf{X}) | X_1, \dots, X_{k-1}, X_{k+1}, \dots, X_n]$. The quantity $\| \|f_k(\mathbf{X})\|_\psi \|_\infty$ refers to its essential supremum. If two \mathbf{X}, \mathbf{Y} are i.i.d. then $\|f_k(\mathbf{X})\|_\psi = \|f_k(\mathbf{Y})\|_\psi$ and $\|f_k(\mathbf{X})\|_\psi(\mathbf{X})$ is i.i.d. to $\|f_k(\mathbf{X})\|_\psi(\mathbf{Y})$.

The next result provides a McDiarmid-type inequality for function with sub-exponential conditional versions.

Theorem 3 (McDiarmid's inequality [50, Theorem 4]). *Let $f : \mathbb{X}^n \rightarrow \mathbb{R}$ and $\mathbf{X} \in \mathbb{X}^n$ a random vector with independent components. Then, for any $t > 0$,*

$$f(\mathbf{X}) - \mathbb{E}f \leq e \left(2\sigma \sqrt{\log \left(\frac{1}{\delta} \right)} + M \log \left(\frac{1}{\delta} \right) \right) \quad (20)$$

with probability at least $1 - \delta$, with $\sigma^2 = \left\| \sum_k \|f_k(\mathbf{X})\|_{\psi_1}^2 \right\|_\infty$ and $M = \max_k \| \|f_k(\mathbf{X})\|_{\psi_1} \|_\infty$.

Proof. Since this result is not exactly stated in this form, we provide a short proof in the Appendix A. \square

The quantities f_k , σ^2 and M are the substitutes to the differences, variance and scale respectively for the standard McDiarmid's inequality.

5.2 Proof of Theorem 1

The Theorem 1 can now be proved using Theorem 3.

Proof. Let $m = \mathbb{E}[f(\mathbf{X}) | \mathbf{X} \in \mathcal{B}]$. By the total law of probability,

$$\mathbb{P}(f(\mathbf{X}) - m > t) \leq \mathbb{P}(f(\mathbf{X}) - m > t, \mathbf{X} \in \mathcal{B}) + \mathbb{P}(\mathbf{X} \notin \mathcal{B}).$$

The inequality of Theorem 3 cannot be directly applied to $\mathbb{P}(f(\mathbf{X}) - m, \mathbf{X} \in \mathfrak{B})$ since the components of \mathbf{X} are no longer independent conditionally on $\mathbf{X} \in \mathfrak{B}$.

Similarly to [25], the main idea of the proof is to define an extension of f to the whole domain that preserves its conditional versions. Let $B : \mathbb{X}^n \times \mathbb{X}^n \rightarrow \mathbb{R}_+$ be the pseudometric defined by

$$B(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n b(x_i, y_i)$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{X}^n$. The extension $\bar{f} : \mathbb{X}^n \rightarrow \mathbb{R}$ of f will be defined by

$$\bar{f}(\mathbf{x}) = \inf_{\mathbf{y} \in \mathfrak{B}} f(\mathbf{y}) + B(\mathbf{x}, \mathbf{y})$$

for all $\mathbf{x} \in \mathbb{X}^n$. This extension was also used in the Kirszbraun's theorem [38, 33] to extend the range of Lipschitz functions.

Lemma 3. *The extension function \bar{f} satisfies*

- $\bar{f}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in \mathfrak{B}$.
- $|\bar{f}(\mathbf{x}) - \bar{f}(\mathbf{y})| \leq B(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{X}^n$.

The proof of the Theorem continues using this extension. Let $M = \mathbb{E}[\bar{f}(\mathbf{X})]$. By definition of \bar{f} ,

$$\mathbb{P}(f(\mathbf{X}) - m > t, \mathbf{X} \in \mathfrak{B}) \leq \mathbb{P}(\bar{f}(\mathbf{X}) - m > t) = \mathbb{P}(\bar{f}(\mathbf{X}) - M > t + m - M).$$

It remains to find a bound on $M - m$. The expectation M can be decomposed as follows

$$M = \mathbb{E}[\bar{f}(\mathbf{X})\mathbf{1}_{\mathbf{X} \in \mathfrak{B}}] + \mathbb{E}[\bar{f}(\mathbf{X})\mathbf{1}_{\mathbf{X} \notin \mathfrak{B}}] \leq \mathbb{P}(\mathbf{X} \in \mathfrak{B})m + \mathbb{E}[\bar{f}(\mathbf{X})\mathbf{1}_{\mathbf{X} \notin \mathfrak{B}}].$$

By definition of $\bar{f}(\mathbf{x}) = \inf_{\mathbf{y} \in \mathfrak{B}} f(\mathbf{y}) + B(\mathbf{x}, \mathbf{y}) \leq \mathbb{E}[f(\mathbf{Y}) + B(\mathbf{x}, \mathbf{Y}) \mid \mathbf{Y} \in \mathfrak{B}]$ so that

$$\mathbb{E}[\bar{f}(\mathbf{X})\mathbf{1}_{\mathbf{X} \notin \mathfrak{B}}] \leq m\mathbb{P}(\mathbf{X} \notin \mathfrak{B}) + \mathbb{E}[\mathbb{E}[B(\mathbf{X}, \mathbf{Y}) \mid \mathbf{Y} \in \mathfrak{B}]\mathbf{1}_{\mathbf{X} \notin \mathfrak{B}}]$$

Therefore,

$$M - m \leq \frac{1}{\mathbb{P}(\mathbf{Y} \in \mathfrak{B})} \mathbb{E}[B(\mathbf{X}, \mathbf{Y})\mathbf{1}_{\mathbf{X} \notin \mathfrak{B}, \mathbf{Y} \in \mathfrak{B}}]$$

where the order of integration on B has been interchanged by Tonelli's theorem. An application of Cauchy-Schwartz inequality gives

$$\begin{aligned} \mathbb{E}[B(\mathbf{X}, \mathbf{Y})\mathbf{1}_{\mathbf{X} \notin \mathfrak{B}, \mathbf{Y} \in \mathfrak{B}}] &\leq \|B\|_2 \mathbb{P}(\mathbf{X} \notin \mathfrak{B}, \mathbf{Y} \in \mathfrak{B})^{1/2} \\ &\leq 2n\|b\|_{\psi_1} \mathbb{P}(\mathbf{X} \notin \mathfrak{B}, \mathbf{Y} \in \mathfrak{B})^{1/2} \end{aligned}$$

The independence of \mathbf{X} and \mathbf{Y} leads to

$$M - m \leq 2n\|b\|_{\psi_1} \sqrt{\frac{p}{1-p}}.$$

In the regime where $p \leq 3/4$, $p^{1/2}(1-p)^{-1/2} \leq 2\sqrt{p}$, leading to

$$M - m \leq 4n\|b\|_{\psi_1} \sqrt{p}$$

which implies

$$\mathbb{P}(\bar{f}(\mathbf{X}) - M > t + m - M) \leq \mathbb{P}(\bar{f}(\mathbf{X}) - M > t - 4n\|b\|_{\psi_1}\sqrt{p}).$$

To conclude the Theorem 3 is used on \bar{f} since

$$\begin{aligned} \bar{f}_k(\mathbf{X})(\mathbf{x}) &= \mathbb{E}[\bar{f}(x_1, \dots, x_{k-1}, X_k, x_{k+1}, \dots, x_n) - \bar{f}(x_1, \dots, x_{k-1}, X'_k, x_{k+1}, \dots, x_n) | X_k] \\ &\leq \mathbb{E}[b(X_k, X'_k) | X_k] \end{aligned}$$

so that $\|\bar{f}_k(\mathbf{X})(\mathbf{x})\|_{\psi_1} \leq \|\mathbb{E}[b(X_k, X'_k) | X_k]\|_{\psi_1} \leq \|b\|_{\psi_1}$ by [50, Lemma 6]. Therefore,

$$\bar{f}(\mathbf{X}) - M \leq 4n\|b\|_{\psi_1}\sqrt{p} + e \left(2\sigma \sqrt{\log\left(\frac{1}{\delta}\right)} + M \log\left(\frac{1}{\delta}\right) \right)$$

with probability at least $1 - p - \delta$ where $\sigma^2 = n\|b\|_{\psi_1}^2$ and $M = \|b\|_{\psi_1}$. \square

Remark 5. In Theorem 1, the function b has to be a pseudometric while Theorem 3 does not assume any further structure on b . It is not clear whether Theorem 1 holds dropping the pseudometric assumption.

5.3 Proof of Lemma 2

The goal of this section is to prove the Lemma 2 providing the bound on the conditional expectation $\mathbb{E}[\vartheta(\hat{\mathcal{S}}; \mathcal{S})^\beta | \mathbf{X} \in \mathcal{A}]$.

To start this section, it is important to describe some properties of the subset \mathcal{A} . First, the subset \mathcal{A} , as the empirical risk $\hat{\mathcal{J}}$, is invariant under reordering of the components of the vector \mathbf{X} . This property will be used several times in the proof. The section of \mathcal{A} is defined for $\mathbf{X} \in \mathcal{A}$ by $I(\mathbf{X}) = \{x \in \mathbb{X} | (x, X_2, \dots, X_n) \in \mathcal{A}\}$. The subset \mathcal{A} and its section satisfy the following properties

$$\begin{aligned} \{\mathbf{X} \in \mathcal{A}\} \cap \{\mathbf{X}' \in \mathcal{A}\} &= \{\mathbf{X} \in \mathcal{A}\} \cap \{Y \in I(\mathbf{X})\} \\ \{\mathbf{X} \in \mathcal{A}\} \cap \{\mathbf{X}' \notin \mathcal{A}\} &= \{\mathbf{X} \in \mathcal{A}\} \cap \{Y \notin I(\mathbf{X})\} \end{aligned} \tag{21}$$

recalling that $\mathbf{X}' = (Y, X_2, \dots, X_n)$ where Y is an i.i.d. copy of X_1 . The events $\mathbf{X} \in \mathcal{A}$ and $\mathbf{X} \notin \mathcal{A}$ are negatively correlated, meaning that $\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}) \leq \mathbb{P}(\mathbf{X} \in \mathcal{A})\mathbb{P}(\mathbf{X}' \notin \mathcal{A})$, see Lemma 12. This property quantifies the intuitive idea that, once \mathbf{X} fall in \mathcal{A} with large probability, the probability that the set of points \mathbf{X}' (i.e. \mathbf{X} with one sample point swapped) fall outside \mathcal{A} is small.

Proof of Lemma 2. The proof starts with the usual decomposition for any $\hat{\phi} \in \hat{\mathcal{S}}$

$$\begin{aligned} \tau\vartheta(\hat{\phi}, \mathcal{S})^\beta &\leq \mathcal{J}(\hat{\phi}) - \mathcal{J}_* && \text{using Assumption 2 and } \mathbf{X} \in \mathcal{A} \\ &= \mathbb{E}_Y[l(\hat{\phi}, Y)] + \mathcal{R}(\hat{\phi}) - \mathcal{J}_*. \end{aligned}$$

The standard argument of (7) cannot be directly applied here since Y and the X_i do not share the same distribution anymore due to the conditioning on $\mathbf{X} \in \mathcal{A}$. It is however possible to split the expectation on Y into the domains (i) $Y \in I(\mathbf{X})$ where Y will be swappable with

any X_i and (ii) $Y \notin I(\mathbf{X})$ where Y and the X_i do not share the same distribution. Fortunately, the event $Y \notin I(\mathbf{X})$ will have a small measure. This leads to the following decomposition

$$\begin{aligned} \tau \vartheta(\widehat{\phi}, \mathcal{S})^\beta &\leq e_1 + e_2 \quad \text{with,} \\ e_1 &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) + \mathcal{R}(\widehat{\phi}) - \widehat{\mathcal{F}}_* \right) \mathbf{1}_{Y \in I(\mathbf{X})} \right] = \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - \widehat{\ell}(\widehat{\phi}) \right) \mathbf{1}_{Y \in I(\mathbf{X})} \right] \\ e_2 &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) + \mathcal{R}(\widehat{\phi}) - \widehat{\mathcal{F}}_* \right) \mathbf{1}_{Y \notin I(\mathbf{X})} \right] + \widehat{\mathcal{F}}_* - \mathcal{F}_*. \end{aligned}$$

so that

$$\tau \mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A} \right] \leq E_1 + E_2, \quad \text{with, } E_i = \mathbb{E} \left[\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} e_i \mid \mathbf{X} \in \mathcal{A} \right]$$

Bounding E_1 Since $\sup_\phi \mathbb{E}[g(\phi, X)] \leq \mathbb{E}[\sup_\phi g(\phi, X)]$ for any measurable $g : \mathcal{M} \times \mathbb{X} \rightarrow \mathbb{R}$, the term E_1 can further be bounded by

$$\begin{aligned} E_1 &\leq \mathbb{E}_{\mathbf{X}} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{E}_Y \left[\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} (l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_i)) \mathbf{1}_{Y \in I(\mathbf{X})} \right] \mid \mathbf{X} \in \mathcal{A} \right] \\ &= \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} (l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1)) \mathbf{1}_{Y \in I(\mathbf{X})} \right] \mid \mathbf{X} \in \mathcal{A} \right] \end{aligned}$$

since \mathcal{A} is invariant under permutations of the components. The next step of the proof consists in bringing out the stability of the minimizers of the empirical risk in E_1 . Let $\widehat{\ell}$, $\widehat{\mathcal{F}}$ and $\widehat{\mathcal{S}}$ be the empirical risks and the minimizers associated to \mathbf{X}' . Remarking that,

$$l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1) = n(\widetilde{\ell}(\widehat{\phi}) - \widehat{\ell}(\widehat{\phi})) = n(\widetilde{\mathcal{F}}(\widehat{\phi}) - \widehat{\mathcal{F}}(\widehat{\phi})) = n(\widetilde{\mathcal{F}}(\widehat{\phi}) - \widehat{\mathcal{F}}_*)$$

since the terms in \mathcal{R} cancel out, and using the fact that \mathbf{X} and \mathbf{X}' have the same distribution gives

$$\mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\widetilde{\mathcal{F}}_* \mathbf{1}_{Y \in I(\mathbf{X})} \right] \mathbf{1}_{\mathbf{X} \in \mathcal{A}} \right] = \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\widehat{\mathcal{F}}_* \mathbf{1}_{Y \in I(\mathbf{X})} \right] \mathbf{1}_{\mathbf{X} \in \mathcal{A}} \right],$$

(see Lemma 9). This leads to the bound

$$E_1 \leq \frac{n}{1-p} \mathbb{E}_{\mathbf{X}, Y} \left[\frac{1}{n} a(Y, X_1) \vartheta(\widehat{\mathcal{S}}; \widetilde{\mathcal{S}})^\alpha \mathbf{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}} \right]$$

with an argument similar to Lemma 1. Note that the order of integration has been interchanged using Tonelli's theorem. Therefore, using Lemma 1 shows that

$$\begin{aligned} E_1 &\leq \frac{1}{1-p} \mathbb{E} \left[a(Y, X_1) (2\tau^{-1}n^{-1}a(Y, X_1))^{\frac{\alpha}{\beta-\alpha}} \mathbf{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}} \right] \\ &\leq \frac{(2\tau^{-1}n^{-1})^{\frac{\alpha}{\beta-\alpha}}}{1-p} \left\| a^{\frac{\beta}{\beta-\alpha}} \right\|_2 \sqrt{\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A})} \end{aligned} \tag{22}$$

using Cauchy-Schwartz inequality. By definition of p_n

$$\frac{\sqrt{\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A})}}{1-p_n} \leq \frac{1}{\sqrt{1-p_n}} \leq 2 \tag{23}$$

in the regime where $p_n \leq 3/4$. Furthermore, since $\frac{\beta}{\beta-\alpha} \geq 1$

$$\left\| a^{\frac{\beta}{\beta-\alpha}} \right\|_2 = \|a\|_{\frac{2\beta}{\beta-\alpha}} \leq \left(\frac{2\beta}{\beta-\alpha} \|a\|_{\psi_1} \right)^{\frac{\beta}{\beta-\alpha}}. \quad (24)$$

Combining (22), (23) and (24) leads to

$$E_1 \leq 2(2\tau^{-1}n^{-1})^{\frac{\alpha}{\beta-\alpha}} \left(\frac{2\beta}{\beta-\alpha} \|a\|_{\psi_1} \right)^{\frac{\beta}{\beta-\alpha}}$$

Bounding E_2 The second term emerges from the domain where Y and X_i do not have the same distribution. It can be bounded by

$$E_2 \leq 2^{\max(0, \alpha-1)+2} \tau^{-\frac{\alpha}{\beta-\alpha}} \|a\|_{\psi_1}^{\frac{\beta}{\beta-\alpha}} \sqrt{p_n} \left(\left(\tau \|a\|_{\psi_1}^{-1} \right)^{\frac{\alpha}{\beta-\alpha}} \text{diam}(\mathcal{S})^\alpha + c_2(\alpha, \beta) \right)$$

using Lemma 10, with $c_2(\alpha, \beta) = 2 \max\left(\frac{4\alpha}{\beta-\alpha}, 1\right)^{\frac{\alpha}{\beta-\alpha}}$.

Combining the bounds on E_1 and E_2 gives the desired result. \square

The bound of E_2 makes the diameter of \mathcal{S} appear, requiring the set of minimizers to be bounded. This is not too restrictive in practice, but it is not clear if it is an artifact of the proof or a price to be paid when splitting the domain of integration of Y

6 Applications

In this section, the main results of the Section 4 are applied on the examples presented in the introduction.

6.1 Barycenters in Hadamard spaces

First introduced in [30], Fréchet means as defined in (4) have found numerous applications in geometry [67, 76], statistics [57, 56], data science [11, 58, 15] and more. Note that the behavior of the empirical counterpart has been studied in specific scenarios for example in Riemannian manifold [13], in Wasserstein spaces [2, 46] or recently in non-positively curved metric spaces [23]. This section is dedicated to Fréchet means defined on Hadamard spaces.

Let $(\mathcal{M}, \vartheta) = (\mathbb{X}, \rho)$ a Hadamard space defined by

Definition 2. A Hadamard space is a complete geodesic metric space (\mathbb{X}, ρ) fulfilling

$$\rho(\phi, x)^2 - \rho(\psi, x)^2 - \rho(\phi, y)^2 + \rho(\psi, y)^2 \leq 2\rho(x, y)\rho(\phi, \psi) \quad (25)$$

for all $x, y, \phi, \psi \in \mathbb{X}$.

The property (25) characterizes CAT(0)-spaces (see e.g. [12, Corollary 3] or [67, Proposition 2.4]) i.e. geodesic metric spaces with non-positive curvature in the sense of Alexandrov. The class of Hadamard spaces describes a large variety of spaces such as Euclidean spaces, Hilbert spaces, complete metric trees, complete simply-connected Riemannian manifolds with

non-positive sectional curvature and more. An introduction to Hadamard spaces can be found in [5].

The computation of the barycenter (or Fréchet mean) of μ consists in minimizing $\mathcal{F}(\phi) = \int_{\mathbb{X}} \rho^2(\phi, x) d\mu(x)$ which corresponds to $l(\phi, x) = \rho(\phi, x)^2$ in (1).

It is known that the barycenter of a measure on a Hadamard space exists and is unique provided that there exists a $\phi \in \mathbb{X}$ such that $\mathbb{E}_{X \sim \mu}[\rho(\phi, X)^2] < +\infty$ [67, Proposition 4.3]. Furthermore, [67, Proposition 4.4] shows that Assumption 2 holds with $\mathcal{J}_0 = +\infty$, $\tau = 1$ and $\beta = 2$.

Assuming that $\|\rho\|_{\psi_1}$ is finite implies in particular that both moments $\mathbb{E}_{X \sim \mu}[\rho(\phi, X)^2]$ and $\mathbb{E}_{X \sim \mu_n}[\rho(\phi, X)^2]$ are finite for some $\phi \in \mathbb{X}$ so that all the assumptions are satisfied as detailed by the next Lemma.

Lemma 4. *Assuming that $\|\rho\|_{\psi_1} < +\infty$, then this example satisfies Assumptions 1, 2, 3, 4 and 5 with*

\mathcal{J}_0	β	τ	α	a	$\ a\ _{\psi_1}$	η	κ	ι
$+\infty$	2	1	1	2ρ	$2\ \rho\ _{\psi_1}$	0	0	0

Corollary 3. *Assuming that $\|\rho\|_{\psi_1} < +\infty$ then*

$$\rho(\phi_*, \hat{\phi}) \leq 16\|\rho\|_{\psi_1} \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n} \log\left(\frac{1}{\delta}\right) \right)$$

with probability at least $1 - \delta$ and

$$\mathbb{E} \left[\rho(\phi_*, \hat{\phi})^2 \right] \leq 2^8 \|\rho\|_{\psi_1}^2 n^{-1}.$$

Similar results are obtained in [46] and [23].

Remark 6. *In the Euclidean case, the quantity $\|\rho\|_{\psi_1}$ can be bounded by*

$$\|\rho\|_{\psi_1} \leq 2\|X\|_{\psi_1} \leq 2d\|X\|_{\psi_1}$$

(see Preliminaries or [79]) so that the condition $\|\rho\|_{\psi_1} < +\infty$ is met for sub-exponential random variables. In Hilbert spaces, this condition is however more subtle to grasp. Some examples of random functions fulfilling this condition are given in [47, Proposition 5.3]

Remark 7. *When $\mathcal{M} = \mathbb{X}$ is further restricted to be an Euclidean space $\mathcal{M} = \mathbb{X} = \mathbb{R}^d$ and endowed with the usual ℓ_2 norm, the bound of Corollary 3 can be compared with straight calculations. In this context, the unique barycenter is $\phi_* = \mathbb{E}_{\mu}[X]$ (in the Pettis integral sense) and the empirical one $\hat{\phi} = (1/n) \sum_{i=1}^n X_i$. The concentration property of sub-exponential variable gives*

$$\|\phi_* - \hat{\phi}\|_2 \leq 2d\|X\|_{\psi_1} \left(\frac{1}{n} \log\left(\frac{2}{\delta}\right) + \sqrt{\frac{2}{n} \log\left(\frac{2}{\delta}\right)} \right)$$

with high probability $1 - \delta$. Moreover, further computations shows that

$$\mathbb{E} \left[\|\phi_* - \hat{\phi}\|_2^2 \right] = \frac{1}{n} \mathbb{E} \left[\|X - \mathbb{E}[X]\|_2^2 \right] \leq \frac{16}{n} \|X\|_{\psi_1}^2 = \frac{16d^2}{n} \|X\|_{\psi_1}^2 \quad (26)$$

The concentration in probability and expectation obtained via Corollary 3 are therefore optimal up to universal constants.

6.2 Largest eigenvector of covariance matrices

Principal components analysis (PCA) [55, 35] is a dimension reduction technique. It consists in computing the eigensystems of the covariance matrix of the data to identify the subspaces that contain the largest variability. In the majority of high-dimensional analyses, the data is represented as points $(y_i)_{i=1}^n$ sampled from the unknown probability measure ν in \mathbb{R}^d . The covariance matrix $\text{Cov}(\nu)$ can thus be estimated only from the samples represented by the discrete measure $\nu_n = \frac{1}{n} \sum_{i=1}^n \delta_{y_i}$. One important question concerns the stability of the subspaces of $\text{Cov}(\nu_n)$ with respect to those of $\text{Cov}(\nu)$. The case of the leading direction of variability i.e. the largest eigenvector will be analyzed in this section.

Let \mathbb{R}^d endowed with the Euclidean distance and equipped with a probability measure ν . The covariance matrix of ν is defined as

$$\text{Cov}(\nu) = \int_{\mathbb{R}^d} (y - \mathbb{E}_\nu[Y])(y - \mathbb{E}_\nu[Y])^T d\nu(x).$$

Let $\lambda_1 \geq \dots \geq \lambda_m$ be the eigenvalues of $\text{Cov}(\nu)$ and ϕ_* one of the two eigenvectors associated to the largest eigenvalue. By the Courant–Fischer min-max principle, ϕ_* is defined by

$$\phi_* \in \arg \max_{\phi \in \mathbb{S}^{d-1}} \langle \text{Cov}(\nu)\phi, \phi \rangle. \quad (27)$$

This optimization problem falls in the setting of (1) observing that

Lemma 5. *The cost function \mathcal{F} of (27) can be written as*

$$\mathcal{F}(\phi) = - \int_{\mathbb{X}^2} \langle y, \phi \rangle^2 - \langle y, \phi \rangle \langle z, \phi \rangle d\nu(y) d\nu(z).$$

This problem fits in the formalism of (1) with $\mathcal{M} = \mathbb{S}^{d-1}$ endowed with the great circle distance $\vartheta(\phi, \psi) = \arccos \langle \phi, \psi \rangle_{\mathbb{R}^d}$, $\mathbb{X} = \mathbb{R}^d \times \mathbb{R}^d$, $\mu = \nu \otimes \nu$ and $l(\phi, x) = l(\phi, (y, z)) = \langle y, \phi \rangle^2 - \langle y, \phi \rangle \langle z, \phi \rangle$. In this context, the set of minimizers of \mathcal{F} is $\mathcal{S} = \{\pm \phi_*\}$ and $\vartheta(\phi, \mathcal{S}) = \arccos |\langle \phi, \phi_* \rangle|$.

Lemma 6. *Assuming $\|Y\|_{\psi_2} < +\infty$ then this example satisfies Assumptions 1, 2, 3, 4 and 5 with*

\mathcal{J}_0	β	τ	α	a	$\ a\ _{\psi_1}$	η	κ	ι
$+\infty$	2	$\frac{4}{\pi^2}(\lambda_1 - \lambda_2)$	1	a	$8d\ Y\ _{\psi_2}^2$	0	$d \exp(-cn)$	0

with

$$a(x, x') = 2 (\|y\|_2^2 + \|y'\|_2^2 + \|y\|_2 \|z\|_2 + \|y'\|_2 \|z'\|_2)$$

$$c = \frac{(\lambda_1 - \lambda_2)^2}{32e^2 d^2 \|Y\|_{\psi_2}^4 + 8ed \|Y\|_{\psi_2}^2 (\lambda_1 - \lambda_2)}$$

Proof. See Section B.2. □

Furthermore since $\text{diam}(\mathcal{M}) = \pi$, the following bounds hold.

Corollary 4. *Assuming $\|Y\|_{\psi_2} < +\infty$ then*

$$\arccos |\langle \widehat{\phi}, \phi_* \rangle| \leq 8L \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n} \log\left(\frac{1}{\delta}\right) \right) + C_1 d^{\frac{1}{4}} e^{-\frac{c}{4}n}$$

with probability at least $1 - e^{-cn} - e^{-t}$ for all $t \geq 0$ and

$$\mathbb{E} \left[\arccos^2 |\langle \widehat{\phi}, \phi_* \rangle| \right] \leq 2^6 L^2 n^{-1} + C_2 d^{\frac{1}{2}} e^{-\frac{c}{2}n}$$

with the constant c defined in Lemma 6 and where

$$L = \frac{2d\pi^2 \|Y\|_{\psi_2}^2}{(\lambda_1 - \lambda_2)}, \quad C_1 = 14L + 2\sqrt{L\pi}, \quad C_2 = 4(L\pi + 8L^2) + \pi^2.$$

Remark 8. *The usual way to obtain bounds between ϕ_* and $\widehat{\phi}$ is to use matrix perturbation tools such as the Davis-Kahan Theorem [78, Theorem 2] which gives*

$$\inf_{\phi_* \in \mathcal{S}} \|\widehat{\phi} - \phi_*\|_2 \leq \sqrt{2} \sin(\arccos |\langle \phi_*, \widehat{\phi} \rangle|) \leq 2\sqrt{2} \frac{\|\text{Cov}(\nu_n) - \text{Cov}(\nu)\|_{2 \rightarrow 2}}{\lambda_1 - \lambda_2}$$

and then bound the deviation of $\text{Cov}(\nu_n)$ from $\text{Cov}(\nu)$ using for example the matrix Bernstein inequality of Lemma 13 giving

$$\begin{aligned} \mathbb{P} \left(\inf_{\phi_* \in \mathcal{S}} \|\widehat{\phi} - \phi_*\|_2 > t \right) &\leq d \exp \left(\frac{-nt^2}{4\pi^{-4}e^2 L^2 + 2\sqrt{2}e\pi^{-2}Lt} \right) \\ &\leq d \exp \left(\frac{-nt^2}{L^2 + Lt} \right) \end{aligned}$$

On the other hand, the bound of Corollary 4 implies that

$$\mathbb{P} \left(\inf_{\phi_* \in \mathcal{S}} \|\widehat{\phi} - \phi_*\|_2 > t - 8Ln^{-\frac{1}{2}} - C_1 e^{-\frac{c}{4}n} \right) \leq \exp(-cn) + \exp \left(\frac{-nt^2}{256L^2 + 16Lt} \right).$$

which are, up to universal constants, similar.

Remark 9. *This analysis can also be applied to the eigenvector associated to the smallest eigenvalue and replacing $\lambda_1 - \lambda_2$ by $\lambda_{d-1} - \lambda_d$. The analysis can be extended to a whole subspace of dimension $r \leq d$ by setting $\mathcal{M} = V_r(\mathbb{R}^d)$ the Stiefel manifold of dimension r in \mathbb{R}^d .*

6.3 LASSO

The goal of supervised learning is to predict outputs $V \in \mathbb{V} \subseteq \mathbb{R}$ from inputs $Y \in \mathbb{Y} \subseteq \mathbb{R}^d$ given the knowledge of n pairs $x_i = (y_i, v_i) \in \mathbb{Y} \times \mathbb{V}$. In the usual machine learning framework, these pairs are assumed to be i.i.d. sampled from a probability distribution μ on $\mathbb{X} = \mathbb{Y} \times \mathbb{V}$.

The marginal $\mu_{\mathbb{Y}}$, of μ along \mathbb{Y} , models the randomness in the inputs. A common model for the randomness in the outputs is to assume that there exists a measurable function $f_* : \mathbb{Y} \rightarrow \mathbb{V}$ such that $V = f_*(Y) + \epsilon$ where ϵ is the noise of the model represented by the marginal $\mu_{\mathbb{V}}$. For instance, the noiseless case is obtained when $\mu_{\mathbb{V}}$ is the Dirac in $f_*(X)$.

As minimizing on the set of measurable functions is intractable, the solution is sought in an a priori “good” class of functions parameterized by a $\phi \in \mathbb{R}^m$. A simple but powerful idea used in machine learning consists in seeking the solution in the space of linear functions with sparse parametrization $\mathcal{H} = \{f \mid f(y) = \langle \phi, \theta(y) \rangle, \text{ for all } k\text{-sparse } \phi \in \mathbb{R}^m\}$ where $\theta : \mathbb{R}^d \rightarrow \mathbb{R}^m$ are the feature vectors.

Finding the best predictor amounts to solving the problem

$$\min_{\phi \in \mathbb{R}^m} \mathcal{F}(\phi) = \frac{1}{2} \int_{\mathbf{Y} \times \mathbf{V}} (\langle \phi, \theta(y) \rangle - v)^2 d\mu(y, v) + \lambda \|\phi\|_1. \quad (28)$$

with \mathcal{F} defined on $\mathcal{M} = \mathbb{R}^m$ endowed with the Euclidean distance $\vartheta = \|\cdot - \cdot\|_2$. In this case the regularization $\mathcal{R} = \lambda \|\cdot\|_1$ and the corresponding $l : \mathbb{R}^m \times \mathbb{R}^{d+1} \rightarrow \mathbb{R}_+$ is defined by $l(\phi, (y, v)) = (\langle \phi, \theta(y) \rangle - v)^2$.

These linear regression problems appears in statistics and machine learning and called LASSO (Least Absolute Shrinkage and Selection Operator) [69]. It is a very mature problem that has been thoroughly studied, extended and applied in various contexts. Understanding the convergence of the estimator $\hat{\phi}$ towards the exact solution ϕ_* is an important question addressed in the literature see e.g. [40].

Lemma 7. *Assuming that $\|\theta(Y)\|_{\psi_2} < +\infty$ and $\|V\|_{\psi_2} < +\infty$ then this example satisfies Assumptions 1, 2, 3, 3 and 5 with*

\mathcal{J}_0	β	τ	α	a	$\ a\ _{\psi_1}$	η	κ	ι
$\mathbb{E}_{\mu_{\mathbf{V}}}[V^2]\lambda^{-1}$	2	(32)	1	(41)	(42)	$2e^{-c_1 n}$	$me^{-c_2 n}$	η

with c_1 and c_2 displayed in the proof.

Proof. See Section B.3 for the proof and the various expressions. □

Corollary 5. *Assuming that $\|\theta(Y)\|_{\psi_2} < +\infty$ and $\|V\|_{\psi_2} < +\infty$ then*

$$\|\hat{\phi} - \phi_*\|_2 \leq 8L \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n} \log\left(\frac{1}{\delta}\right) \right) + 14Lp_n^{\frac{1}{4}}$$

with probability at least $1 - p_n - e^{-t}$ and

$$\mathbb{E} \left[\|\hat{\phi} - \phi_*\|_2^2 \right] \leq L^2 2^6 n^{-1} + C_2 p_n^{-\frac{1}{2}}$$

where $C_2 = 2^5 + (2\|V\|_{\psi_2}^2 + 2^5\|V\|_{\psi_2}^4)\lambda^{-2}$ and $L = \tau^{-1}\|a\|_{\psi_1}$.

Proof. See Section B.3. □

6.4 Wasserstein barycenter

Studying barycenters in Hadamard spaces was quite straightforward due to their non-positive curvature that implied Assumptions 1, 2 (with $\mathcal{J}_0 = +\infty$), 4 and 5 (see Section 6.1). The situation is more delicate in spaces of non-negative curvature, and it represents a limitation of our analysis that cannot be applied verbatim. Deeper deeper investigations are required to obtain concentration bounds in this case. A perfect example of such spaces, which has drawn attention in the last decade, is the Wasserstein space. A brief introduction is presented here, but complete treatments can be found [76, 4, 3] for example.

The Wasserstein space Given (H, h) a Polish space, we let $\mathcal{P}_2(H)$ be the set of all Borel probability measures ϕ on H such that $\mathbb{E}_{X \sim \phi}[h(X, y)^2]$ for all $y \in H$. Define $\Pi_{\phi, \psi}$, the set of couplings between two measures ϕ and ψ in $\mathcal{P}_2(H)$, as the set of probability measures π on $H \times H$ having ϕ and ψ as marginals. The Wasserstein distance between ϕ and ψ is defined as

$$W_2^2(\phi, \psi) = \inf_{\pi \in \Pi_{\phi, \psi}} \mathbb{E}_{(X, Y) \sim \pi} [h(X, Y)^2]. \quad (29)$$

This section focuses on the case where H is a Hilbert space endowed with the distance h associated to the inner product i.e. the spaces $\mathcal{M} = \mathbb{X} = (\mathcal{P}_2(H), W_2)$. The Wasserstein space \mathcal{M} is known to be geodesic and with non-negative curvature [4, Section 7.3]. Let μ be a probability measure on $\mathcal{P}_2(\mathcal{M})$, its barycenters are defined as the Fréchet means

$$\phi_* \in \arg \min_{\phi \in \mathcal{P}_2(\mathcal{M})} \mathcal{F}(\phi) = \int W_2^2(\phi, \psi) d\mu(\psi) \quad (30)$$

i.e. setting $l(\phi, \psi) = W_2^2(\phi, \psi)$ in the formalism of (1). The existence of the barycenter and its empirical counterpart is guaranteed in most reasonable settings [1, 45], e.g. when (H, h) is a separable locally compact geodesic space.

The difficulties of non-negative curvature As observed in [4, Theorem 7.3.2], the barycenter function (30) and its empirical counterpart are not strongly convex. This is due to the positive curvature of $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, and actually \mathcal{F} and $\widehat{\mathcal{F}}$ exhibit some concavity property, in the sense of (14). Therefore, the functions \mathcal{F} and $\widehat{\mathcal{F}}$ cannot be strongly convex and the global error bound of Assumption 2 with $\mathcal{F} = +\infty$ cannot be obtained via these means.

The lack of strong convexity of the barycenter functions does not imply that a global error bound of Assumption 2 with $\mathcal{F} = +\infty$ cannot hold as it represents a weaker notion of strong convexity (see Section 6.2 for an example of non-convex function with a global error bound). The situation becomes more complex when observing that such a bound cannot however hold for all measure $\mu \in \mathcal{P}_2(\mathcal{M})$ since this would imply that \mathcal{M} has non-positive curvature [67, Theorem 4.9]. Furthermore, such a bound cannot hold uniformly at the vicinity of the barycenters, since the curvature of the balls of \mathcal{M} of any radius is unbounded [66, Section 2.5].

Similarly, Assumption 5 cannot hold for all $x, y, \phi, \psi \in \mathcal{M}$ since this would imply that \mathcal{M} has a non-positive curvature (see Section 6.1). It would however hold for any ϕ, ψ on any bounded subset of \mathcal{M} .

When convergence can be proved After having drawn a catastrophic picture of the situation, proving that the empirical barycenters converge to the exact ones with parametric rate is surprisingly still possible.

Additional work is however required to characterize measures $\mu \in \mathcal{P}_2(\mathcal{M})$ for which convergence can be proved. The following derivation are from [2, 46] from which we state the results relevant for the purposes of this section while referring to them for a complete treatment.

The first ingredient is the fact that the tangent space $\mathcal{T}_{\phi_*} \mathcal{M}$ of \mathcal{M} at the barycenter ϕ_* can be identified to a Hilbert space [46, Theorem 7]. The distance induced by the metric will be denoted $\|\cdot\|_{\phi_*}$. Second they show (in Corollary 16) that when the measure $\mu \in \mathcal{P}_2(\mathcal{M})$ has one barycenter ϕ_* such that every $\phi \in \text{supp}(\mu)$ is the pushforward of ϕ_* by a the gradient of an α -strongly convex and β -smooth function $\zeta_{\phi_* \rightarrow \phi}$ (the Kantorovich potential), i.e. $\phi =$

$\nabla\zeta_{\phi_* \rightarrow \phi} \# \phi_*$, the function \mathcal{F} satisfies the global error bound of Assumption 2 with $\beta = 2$, $\tau = \frac{1}{\beta - \alpha - 1}$ and $\mathcal{J}_0 = +\infty$. Importantly, under this assumption $\tau \vartheta(\widehat{\phi}, \mathcal{S})^2 \leq \|\frac{1}{n} \log_{\phi_*} X_i\|_{\phi_*}^2$. In other words, the problem of finding a barycenter of μ can be lifted to finding a barycenter of $\log_{\phi_*} \# \mu$ on the tangent space $\mathcal{T}_{\phi_*} \mathcal{M}$, which has a Hilbert structure, from which parametric rates from Section 6.1 can be used.

6.5 Entropic-Wassertein barycenters

Establishing convergence guarantees on the empirical Wasserstein barycenters towards their exact counterpart appeared to be a delicate task (see Section 6.4). This obstacle was due to the non-negative curvature of $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ that lead to an Assumption 2 difficult to check for general measures.

To circumvent these difficulties, it was proposed in [16, 42] to seek a regularized barycenter by adding an entropic regularization to the Wasserstein variance function. In [16], the authors proved the existence and uniqueness of the regularized barycenters as well as convergence of their empirical regularized barycenters. Recently, [24] derived regularity properties of the entropic-barycenters and established central limit theorems.

Let Ω be a bounded subset of \mathbb{R}^d endowed with the Euclidean distance and define $\mathcal{M} = \mathcal{P}_2^{ac}(\Omega)$ where $\mathcal{P}_2^{ac}(\Omega)$ refers to the subset of measures of $\mathcal{P}_2(\Omega)$ that are absolute continuous with respect to the Lebesgue measure. The space \mathcal{M} will be endowed with the total variation distance, which coincides with the ℓ^1 norm on $\mathcal{P}_2^{ac}(\Omega)$ i.e. $\vartheta = \|\cdot - \cdot\|_1$. This choice might be surprising at first but will be substantiated later in this section. Furthermore, let $(\mathbb{X}, \rho) = (\mathcal{P}_2(\Omega), W_2)$. The entropic-barycenter of a probability measure $\mu \in \mathcal{P}_2(\mathcal{M})$ is defined as proposed in [16] as the minimizer of

$$\mathcal{F}(\phi) = \frac{1}{2} \int_{\mathcal{P}_2(\Omega)} W_2^2(\phi, \psi) d\mu(\psi) + \lambda \mathcal{R}(\phi)$$

where $\lambda > 0$ is the regularization parameter and \mathcal{R} is the negative entropy defined for every $\phi \in \mathcal{P}_2(\Omega)$ by

$$\mathcal{R}(\phi) = \begin{cases} \int_{\Omega} (f(x)(\log(f(x)) - 1) + 1) dx & \text{if } \phi = f dx \text{ and } \int_{\Omega} f(x) dx = 1 \\ +\infty & \text{otherwise.} \end{cases}$$

This choice of regularization enforces the Wasserstein barycenter to be absolutely continuous with respect to the Lebesgue measure. The measures ϕ_* and $\widehat{\phi}$, as the minimizer of \mathcal{F} and $\widehat{\mathcal{F}}$ respectively, therefore admit a density. We will slightly abuse notation and use the same one for the measures and their densities. Any probability measure μ on \mathcal{M} satisfying $\int_{\mathcal{P}_2(\Omega)} \mathbb{E}_{X \sim \psi} [X^2] d\mu(\psi) < +\infty$, admits a unique entropic-barycenter (see Proposition 2.1 in [24]).

While the Wasserstein distance squared is not geodesically convex (see Section 6.4), it is convex with respect to the linear structure in $\mathcal{P}_2(\Omega)$ (i.e. considering paths like $\phi_t = t\phi_1 + (1-t)\phi_0$ for $\phi_0, \phi_1 \in \mathcal{P}_2(\Omega)$ and $t \in [0, 1]$) (see e.g. Proposition 7.17 in [61]). This leads to a convex functions \mathcal{F} and $\widehat{\mathcal{F}}$. Furthermore, the negative entropy is strongly convex, for the linear structure, with respect to the total variation distance on $\mathcal{P}_2^{ac}(\Omega)$. These two properties imply the local error bounds for \mathcal{F} and $\widehat{\mathcal{F}}$ i.e. Assumptions 2 and 4.

Checking Assumption 5 is more delicate since one has to bound the left hand side of (16) with $l = \frac{1}{2}W_2^2$ by a quantity involving the total variation which endows $\mathcal{P}_2^{ac}(\Omega)$ with a different topology than with W_2 .

Lemma 8. *Assuming that Ω is bounded, this example satisfies Assumptions 1, 2, 3, 4 and 5 with*

\mathcal{J}_0	β	τ	α	a	$\ a\ _{\psi_1}$	η	κ	ι
$+\infty$	2	2	1	$4 \text{diam}(\Omega)^2$	$4 \text{diam}(\Omega)^2$	0	0	0

Proof. See Section B.4. □

Corollary 6. *Assuming that Ω is bounded then*

$$\|\widehat{\phi} - \phi_*\|_1 < 16 \text{diam}(\Omega)^2 \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n} \log\left(\frac{1}{\delta}\right) \right)$$

with probability at least $1 - \delta$ and

$$\mathbb{E} \left[\|\widehat{\phi} - \phi_*\|_1^2 \right] \leq 2^8 \text{diam}^4(\Omega) n^{-1}.$$

As a comparison, in [16] a bound of kind $\mathbb{E}[\text{KL}(\widehat{\phi}|\phi_*)^2] \leq cn^{-1}$ is obtained, where KL refers to the Kullback-Leibler divergence, for $d = 1$ and some constant $c > 0$. The general case $d \geq 2$ requires additional smoothness assumption. Expressing this bound in terms of total variation via the Pinsker's inequality would give $\mathbb{E}[\|\widehat{\phi} - \phi_*\|_1^4] \leq cn^{-1}$. The bound obtained here in Corollary 6 appears to be faster and does not require additional smoothness assumptions to deal with $d \geq 2$.

7 Discussion

The concentration of the empirical minimizers of the risk (2) to the ones of (1) is a fundamental question in statistical learning. Instead of the usual guarantees on the estimation error, concentration inequalities on the distance between the sets of minimizers are investigated.

In particular, this work identifies a set of assumptions, that is rather weak, and should fit a wide variety of estimations problems e.g. with spaces that are unbounded and infinite dimensional and cost functions that are unbounded and non-convex. Therefore, these assumptions allow to describe a mechanism that seems to govern the concentration of the empirical minimizers in many estimation problems. More precisely, there exists a regime $\mathbf{X} \in \mathcal{A}$, happening with high-probability, where the empirical minimizers are stable. This stability was leveraged in Theorem 2 to obtain parametric concentration rates. The result was obtained using a new McDiarmid-type inequality for functions with sub-exponential differences on subsets of large measure. The assumptions were verified on a selection of estimation problems showing their relevance and showcasing the optimality of the bounds of the Theorem 2.

Even tough this work establishes the regime of stability of the empirical minimizers and the machinery to obtain parametric concentration rates from it, the set of assumptions has to be checked for each problem, which is quite unsatisfactory in the quest of user-friendly

theorems. While some clues were provided for the general treatments of Assumptions 3 and 5, the most interesting and probably the most open ones seems to be the concentration of the Łojasiewicz constant in Assumption 4. Additionally to the technical questions raised in the core of the paper, Theorem 2 leaves aside the case $\beta = \alpha$ which might be of great interest for applications where the cost function l has a Lipschitz continuous gradient.

A Additional proofs

A.1 General concentration results

Proof of Theorem 3. The Theorem 4 in [50] is not exactly stated in the form of (20). More precisely, it states that

$$\mathbb{P}(f(\mathbf{X}) - \mathbb{E}f \geq t) \leq \exp\left(-\frac{t^2}{4e^2\sigma^2 + 2eMt}\right).$$

for all $t \geq 0$ which is slightly weaker than (20). However, a closer inspection of the proofs shows that

$$\log \mathbb{E} \left[e^{t(f - \mathbb{E}f)} \right] \leq \frac{t^2 e^2 \sigma^2}{1 - teM}$$

for $t \leq (eM^{-1})$. This shows that $f(\mathbf{X}) \in \text{sub}\Gamma_+(2e^2\sigma^2, eM)$ which implies the tail (20), see Section 2 or [19, Section 2.4] for more details. \square

A.2 Supporting results for the proof of Lemma 2

Lemma 9.

$$\mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\left(\tilde{\mathcal{F}}_* - \hat{\mathcal{F}}_* \right) \mathbf{1}_{Y \in I(\mathbf{X})} \right] \mathbf{1}_{\mathbf{X} \in \mathcal{A}} \right] = 0$$

Proof. This result will be clear, as long as we are able to exchange the order of integration, because $\mathbf{X} \in \mathcal{A}$ and $\mathbf{X}' \in \mathcal{A}$ share the same distribution. This will be achieved checking that

$$\mathbb{E}_{\mathbf{X}, Y} \left[\left| \tilde{\mathcal{F}}_* - \hat{\mathcal{F}}_* \right| \mathbf{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}} \right] < +\infty.$$

The quantity is bounded by

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}, Y} \left[\left| \tilde{\mathcal{F}}_* - \inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) + \inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) - \inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) + \inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) - \hat{\mathcal{F}}_* \right| \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ & \leq \mathbb{E}_{\mathbf{X}, Y} \left[\left| \tilde{\mathcal{F}}_* - \inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) + \inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) - \inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) + \inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) - \hat{\mathcal{F}}_* \right| \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ & \leq \mathbb{E}_{\mathbf{X}, Y} \left[\left(\inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) - \tilde{\mathcal{F}}_* \right) \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] + \mathbb{E}_{\mathbf{X}, Y} \left[\left(\inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) - \hat{\mathcal{F}}_* \right) \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ & \quad + \mathbb{E}_{\mathbf{X}, Y} \left[\left| \inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) - \inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) \right| \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \end{aligned}$$

The third term is finite since

$$\begin{aligned} \mathbb{E}_{\mathbf{X}, Y} \left[\left| \inf_{\phi_* \in \mathcal{S}} \tilde{\mathcal{F}}(\phi_*) - \inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{F}}(\phi_*) \right| \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] & \leq 2 \mathbb{E}_{\mathbf{X}, Y} \left[\left| \hat{\mathcal{F}}(\phi_*) \right| \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \quad \text{for some } \phi_* \in \mathcal{S} \\ & \leq \frac{2}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}, Y} \left[|l(\phi_*, X_i)| \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \end{aligned}$$

together with the assumption that $\mathbb{E}_X [|l(\phi_*, X)|] < +\infty$. The first two terms can further be bounded in the same fashion so that only one is detailed. Due to the optimality of ϕ_* and leveraging Assumption 5

$$\begin{aligned} \mathbb{E}_{\mathbf{X}, Y} \left[\left(\inf_{\phi_* \in \mathcal{S}} \widehat{\mathcal{F}}(\phi_*) - \widehat{\mathcal{F}}_* \right) \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] &\leq \mathbb{E}_{\mathbf{X}, Y} \left[\left(\inf_{\phi_* \in \mathcal{S}} \widehat{\mathcal{F}}(\phi_*) - \widehat{\mathcal{F}}(\widehat{\phi}) + \mathcal{F}(\widehat{\phi}) - \mathcal{F}(\phi_*) \right) \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ &\leq \mathbb{E}_{\mathbf{X}, Y} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{E}_Z [a(Z, X_i)] \vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ &= \mathbb{E}_{\mathbf{X}, Y} \left[\mathbb{E}_Z [a(Z, X_i)] \vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ &\leq \|a\|_2 \|\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha \mathbf{1}_{\mathbf{X} \in \mathcal{A}}\|_2 \end{aligned}$$

by an application of Cauchy-Schwartz inequality. The latter quantity is finite since $\|a\|_2 \leq 2\|a\|_{\psi_1}$ and using the bounds on the moments of $\vartheta(\widehat{\mathcal{S}}; \mathcal{S})$ of Lemma 11.

Therefore, in virtue of Fubini's theorem,

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\left(\widetilde{\mathcal{F}}_* - \widehat{\mathcal{F}}_* \right) \mathbf{1}_{Y \in I(\mathbf{X})} \right] \mathbf{1}_{\mathbf{X} \in \mathcal{A}} \right] &= \mathbb{E}_{\mathbf{X}, Y} \left[\left(\widetilde{\mathcal{F}}_* - \widehat{\mathcal{F}}_* \right) \mathbf{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{A}} \right] \\ &= 0 \end{aligned}$$

since \mathbf{X} and \mathbf{X}' share the same distribution. □

Lemma 10. *The term E_2 from the proof of Lemma 2 can be bounded by*

$$E_2 \leq 2^{\max(0, \alpha-1)+2} \left(\|a\|_{\psi_1} \text{diam}(\mathcal{S})^\alpha + 2\|a\|_{\psi_1}^{\frac{\beta}{\beta-\alpha}} \left(\tau^{-1} \max \left(\frac{4\alpha}{\beta-\alpha}, 1 \right) \right)^{\frac{\alpha}{\beta-\alpha}} \right) \sqrt{pn}$$

Proof. The term e_2 can further be bounded, for any $\phi_* \in \mathcal{S}$ by

$$\begin{aligned} e_2 &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) + \mathcal{R}(\widehat{\phi}) - l(\phi_*, Y) - \mathcal{R}(\phi_*) + \widehat{\mathcal{F}}(\phi_*) - \widehat{\mathcal{F}}_* \right) \mathbf{1}_{Y \notin I(\mathbf{X})} \right] \\ &\quad + \mathbb{E}_Y \left[\left(l(\phi_*, Y) + \mathcal{R}(\phi_*) - \widehat{\mathcal{F}}(\phi_*) \right) \mathbf{1}_{Y \notin I(\mathbf{X})} \right] + \widehat{\mathcal{F}}_* - \mathcal{F}_* \\ &= e_{2,1} + e_{2,2} \end{aligned}$$

with

$$\begin{aligned} e_{2,1} &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - l(\phi_*, Y) + \widehat{\mathcal{F}}(\phi_*) - \widehat{\mathcal{F}}(\widehat{\phi}) \right) \mathbf{1}_{Y \notin I(\mathbf{X})} \right] \\ e_{2,2} &= \mathbb{E}_Y \left[\left(l(\phi_*, Y) + \mathcal{R}(\phi_*) - \widehat{\mathcal{F}}(\phi_*) \right) \mathbf{1}_{Y \notin I(\mathbf{X})} \right] + \widehat{\mathcal{F}}_* - \mathcal{F}_* \end{aligned}$$

Therefore

$$E_2 \leq E_{2,1} + E_{2,2} \quad \text{with, } E_{2,i} = \mathbb{E} \left[\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} e_{2,i} \mid \mathbf{X} \in \mathcal{A} \right]$$

The remaining of the proof consists in bounding each term.

Bounding $E_{2,1}$ The term $e_{2,1}$ can further be bounded by

$$\begin{aligned}
e_{2,1} &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_Y \left[(l(\widehat{\phi}, Y) - l(\phi_*, Y) + l(\phi_*, X_i) - l(\widehat{\phi}, X_i)) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_Y \left[a(Y, X_i) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] \sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} \vartheta(\widehat{\phi}, \phi_*)^\alpha \quad \text{using Assumption 5} \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_Y \left[a(Y, X_i) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] 2^{\max(0, \alpha-1)} \left(\vartheta(\widehat{\phi}, \mathcal{S})^\alpha + \text{diam}(\mathcal{S})^\alpha \right)
\end{aligned}$$

This bounds allows to upper bound $E_{2,1}$ by

$$\begin{aligned}
E_{2,1} &\leq 2^{\max(0, \alpha-1)} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[a(Y, X_i) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] \left(\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha + \text{diam}(\mathcal{S})^\alpha \right) \mid \mathbf{X} \in \mathcal{A} \right] \\
&= 2^{\max(0, \alpha-1)} \frac{1}{1-p} \mathbb{E}_{\mathbf{X}, Y} \left[\left(\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha + \text{diam}(\mathcal{S})^\alpha \right) a(Y, X_1) \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}} \right]
\end{aligned}$$

using Tonelli's theorem to swap the order of integration and (21). The term

$$\begin{aligned}
\mathbb{E}_{\mathbf{X}, Y} \left[a(Y, X_1) \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}} \right] &\leq \|a\|_2 \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \\
&\leq 2\|a\|_{\psi_1} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2}
\end{aligned}$$

using Cauchy-Schwartz inequality. Similarly, two successive applications of the Cauchy-Schwartz inequality give

$$\begin{aligned}
&\mathbb{E}_{\mathbf{X}, Y} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha a(Y, X_1) \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}} \right] \\
&\leq \|a\|_4 \|\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}}\|_4 \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \\
&\leq \|a\|_4 \|\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^{\beta-\alpha} \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}}\|_{\max\left(\frac{4\alpha}{\beta-\alpha}, 1\right)^{\frac{\alpha}{\beta-\alpha}}} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \\
&\leq \|a\|_4 \left(\tau^{-1} \|a\|_{\max\left(\frac{4\alpha}{\beta-\alpha}, 1\right)} \right)^{\frac{\alpha}{\beta-\alpha}} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \quad \text{using Lemma 11} \\
&\leq 4\|a\|_{\psi_1} \left(\tau^{-1} \max\left(\frac{4\alpha}{\beta-\alpha}, 1\right) \|a\|_{\psi_1} \right)^{\frac{\alpha}{\beta-\alpha}} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2}
\end{aligned}$$

On the other hand, the Lemma 12 allows to show that

$$\frac{\sqrt{\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})}}{1-p_n} \leq \frac{\sqrt{p_n(1-p_n)}}{p_n} \leq 2\sqrt{p_n}$$

since $\sqrt{p_n/(1-p_n)} \leq 2\sqrt{p_n}$ for $p_n \leq 3/4$. Combining all these bounds leads to

$$E_{2,1} \leq 2^{\max(0, \alpha-1)+1} \sqrt{p_n} \left(2\|a\|_{\psi_1} \text{diam}(\mathcal{S})^\alpha + 4\|a\|_{\psi_1}^{\frac{\beta}{\beta-\alpha}} \left(\tau^{-1} \max\left(\frac{4\alpha}{\beta-\alpha}, 1\right) \right)^{\frac{\alpha}{\beta-\alpha}} \right)$$

Bounding $E_{2,2}$ Since $\widehat{\mathcal{F}}_* \leq \widehat{\mathcal{F}}(\phi_*)$ for all $\phi_* \in \mathcal{S}$ the term $e_{2,2}$ can further be bounded by

$$e_{2,2} \leq \mathbb{E}_Y \left[\left(\widehat{\mathcal{F}}(\phi_*) - l(\phi_*, Y) \right) \mathbb{1}_{Y \in I(\mathbf{X})} \right] \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_Y \left[(l(\phi_*, X_i) - l(\phi_*, Y)) \mathbb{1}_{Y \in I(\mathbf{X})} \right]$$

Taking the conditional expectation gives

$$\mathbb{E} [e_{2,2} | \mathbf{X} \in \mathcal{A}] \leq \frac{1}{1-p} \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[(l(\phi_*, X_1) - l(\phi_*, Y)) \mathbb{1}_{Y \in I(\mathbf{X})} \right] \mathbb{1}_{\mathbf{X} \in \mathcal{A}} \right]$$

because \mathcal{A} is invariant under permutations of the components of \mathbf{X} .

Moreover, since $\mathbb{E}_Y [|l(\phi_*, Y)|] < +\infty$, the order of integration can be switched in virtue of Fubini's theorem, leading to

$$E_{2,2} \leq \frac{1}{1-p} \mathbb{E}_{\mathbf{X}, Y} \left[(l(\phi_*, X_1) - l(\phi_*, Y)) \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}} \right] = 0$$

since \mathbf{X} and \mathbf{X}' have the same distribution.

Combining the bounds on $E_{2,1}$ and $E_{2,2}$ gives the desired result. \square

Lemma 11. *Let \mathcal{A} be the subset defined by (18) and assume Assumptions 1, 2, 3, 4 and 5 hold, then conditionally on \mathcal{A}*

$$\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^{\beta-\alpha} \leq \frac{1}{\tau n} \sum_{i=1}^n \mathbb{E}_X [a(X, X_i)]$$

Furthermore, for all $p \geq 1$

$$\left\| \vartheta(\widehat{\mathcal{S}}; \mathcal{S})^{\beta-\alpha} \mathbb{1}_{\mathbf{X} \in \mathcal{A}} \right\|_p \leq \frac{\|a\|_p}{\tau}.$$

Proof. The conditioning on $\mathbf{X} \in \mathcal{A}$ implies that all $\widehat{\phi} \in \widehat{\mathcal{S}}$ belong to the level-set $[\mathcal{F} \leq \mathcal{F}_0]$. Let $\widehat{\phi} \in \widehat{\mathcal{S}}$, the Assumption 2 gives

$$\begin{aligned} \tau \vartheta(\widehat{\phi}, \mathcal{S})^\beta &\leq \mathcal{F}(\widehat{\phi}) - \mathcal{F}_* \\ &\leq \inf_{\phi_* \in \mathcal{S}} \mathcal{F}(\widehat{\phi}) - \mathcal{F}_* + \widehat{\mathcal{F}}(\phi_*) - \widehat{\mathcal{F}}(\widehat{\phi}) && \text{since } \widehat{\phi} \in \widehat{\mathcal{S}} \\ &= \inf_{\phi_* \in \mathcal{S}} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_X \left[l(\widehat{\phi}, X) - l(\phi_*, X) + l(\phi_*, X_i) - l(\widehat{\phi}, X_i) \right] && \text{by Assumption 5} \\ &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_X [a(X, X_i)] \inf_{\phi_* \in \mathcal{S}} \vartheta(\widehat{\phi}, \phi_*)^\alpha \end{aligned}$$

Therefore,

$$\tau \vartheta(\widehat{\phi}, \mathcal{S})^{\beta-\alpha} \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_X [a(X, X_i)]$$

holds for any $\widehat{\phi} \in \widehat{\mathcal{S}}$ leading to the expected result. The bound on the moments is obtained remarking that

$$\begin{aligned} \left\| \frac{1}{\tau n} \sum_{i=1}^n \mathbb{E}_X[a(X, X_i)] \mathbf{1}_{\mathbf{X} \in \mathcal{A}} \right\|_p &\leq \left\| \frac{1}{\tau n} \sum_{i=1}^n \mathbb{E}_X[a(X, X_i)] \right\|_p \\ &\leq \frac{1}{\tau n} \sum_{i=1}^n \|\mathbb{E}_X[a(X, X_i)]\|_p \\ &\leq \frac{1}{\tau n} \sum_{i=1}^n (\mathbb{E}_{\mathbf{X}, X}[a(X, X_i)^p])^{1/p} \end{aligned}$$

using Jensen's inequality. The final result is obtained remarking that the samples $(X_i)_{i=1}^n$ are i.i.d. \square

A.3 Negative correlation of $\mathbf{X} \in \mathcal{A}$ and $\mathbf{X}' \in \mathcal{A}'$

Lemma 12. *Let $\mathbf{X} = (X_1, \dots, X_n)$ and $\mathbf{X}' = (Y, X_2, \dots, X_n)$ where $(X_i)_{i=1}^n$ and Y are i.i.d. samples of μ . Let \mathcal{A} be the subset of \mathbb{X}^n defined in (18) then*

$$\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}) \leq \mathbb{P}(\mathbf{X} \in \mathcal{A})\mathbb{P}(\mathbf{X}' \notin \mathcal{A})$$

Proof. The probability of the intersection can be decomposed as follows

$$\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}) = \mathbb{E}[\mathbf{1}_{\mathbf{X} \in \mathcal{A}} \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}}] = \mathbb{E}[\mathbf{1}_{\mathbf{X} \in \mathcal{A}}] \mathbb{E}[\mathbf{1}_{\mathbf{X}' \notin \mathcal{A}}] + \text{Cov}(\mathbf{1}_{\mathbf{X} \in \mathcal{A}}, \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}})$$

Let $\mathbf{Z} = (X_2, \dots, X_n)$. In virtue of the law of total covariance

$$\text{Cov}(\mathbf{1}_{\mathbf{X} \in \mathcal{A}}, \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}}) = \mathbb{E}[\text{Cov}(\mathbf{1}_{\mathbf{X} \in \mathcal{A}}, \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}} | \mathbf{Z})] + \text{Cov}(\mathbb{E}[\mathbf{1}_{\mathbf{X} \in \mathcal{A}} | \mathbf{Z}], \mathbb{E}[\mathbf{1}_{\mathbf{X}' \notin \mathcal{A}} | \mathbf{Z}]).$$

Since conditionally on \mathbf{Z} , the random variables X_1 and Y are independent, the covariance $\text{Cov}(\mathbf{1}_{\mathbf{X} \in \mathcal{A}}, \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}} | \mathbf{Z}) = 0$. Let $g : \mathbb{X}^{n-1} \rightarrow \mathbb{R}$ be the function defined by $g(\mathbf{Z}) = \mathbb{E}[\mathbf{1}_{\mathbf{X}' \notin \mathcal{A}} | \mathbf{Z}]$.

$$\text{Cov}(\mathbb{E}[\mathbf{1}_{\mathbf{X} \in \mathcal{A}} | \mathbf{Z}], \mathbb{E}[\mathbf{1}_{\mathbf{X}' \notin \mathcal{A}} | \mathbf{Z}]) = \text{Cov}(g(\mathbf{Z}), 1 - g(\mathbf{Z})) \leq 0$$

by the FKG inequality [29]. Therefore $\text{Cov}(\mathbf{1}_{\mathbf{X} \in \mathcal{A}}, \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}}) \leq 0$ which gives the desired result. \square

B Additional proofs for Applications

B.1 Concentration of sample covariance matrices

Lemma 13. *Let $Z \in \mathbb{R}^d$ be a sub-Gaussian random vector and let $(Z_i)_{i=1}^n$ be n i.i.d. copy of Z . Let $A = \mathbb{E}[ZZ^T]$ and $\widehat{A} = \frac{1}{n} \sum_{i=1}^n Z_i Z_i^T$ then*

$$\mathbb{P}(\|A - \widehat{A}\| > t) \leq d \exp\left(-\frac{nt^2}{2(e^2 d^2 \|Z_i\|_{\psi_2}^4 + ed \|Z_i\|_{\psi_2}^2 t)}\right)$$

Proof. This result is a consequence of [70, Theorem 6.2]

Theorem 4. Let $(S_i)_{i=1}^n$ be n independent, self-adjoint symmetric random matrices of dimension d . Assume that $\mathbb{E}[S_i] = 0$ and there exist R and n matrices $(B_i)_{i=1}^n$ such that $\mathbb{E}[S_i^p] \preceq \frac{p!}{2} R^{p-2} B_i^2$ for all $p \geq 2$ then

$$\mathbb{P} \left(\lambda_{\max} \left(\sum_{i=1}^n S_i \right) > t \right) \leq d \exp \left(-\frac{t^2}{2(\sigma^2 + Rt)} \right)$$

with $\sigma^2 = \left\| \sum_{i=1}^n B_i^2 \right\|$.

This result will be used by setting $S_i = \frac{1}{n} (Z_i Z_i^T - \mathbb{E}[Z Z^T])$. This choice implies that $\mathbb{E}[S_i] = 0$ and $\sum_{i=1}^n S_i = \hat{A} - A$. Moreover, since S_i are self-adjoint matrices, the spectral norm of the matrix $A - \hat{A}$ is equal to its largest eigenvalue. Since each $S_i \preceq \|S_i\| \text{Id}_d$, the Bernstein condition can be obtained by bounding the moments of $\|S_i\|$ because $\mathbb{E}[S_i^p] \leq \mathbb{E}[\|S_i\|^p] \text{Id}_d$. This is achieved remarking that

$$\begin{aligned} \|S_i\| &= \lambda_{\max}(S_i) = \sup_{v \in \mathbb{S}^{d-1}} \langle S_i v, v \rangle = \frac{1}{n} \sup_{v \in \mathbb{S}^{d-1}} (\langle Z_i, v \rangle^2 - \mathbb{E}[\langle Z, v \rangle^2]) \\ &\leq \frac{1}{n} \sup_{v \in \mathbb{S}^{d-1}} \langle Z_i, v \rangle^2 \leq \frac{1}{n} \|Z_i\|_2^2 \end{aligned}$$

so that $\| \|S_i\| \|_{\psi_1} \leq \frac{1}{n} \| \|Z_i\| \|_{\psi_2}^2 \leq \frac{d}{n} \|Z_i\|_{\psi_2}^2$ leading to $\mathbb{E}[\|S_i\|^p] \leq \frac{p!}{2} \left(e \frac{d}{n} \|Z_i\|_{\psi_2}^2 \right)^p$. The Theorem 4 can now be applied with $R = e \frac{d}{n} \|Z_i\|_{\psi_2}^2$ and $B_i = e \frac{d}{n} \|Z_i\|_{\psi_2}^2 \text{Id}_d$ with the corresponding $\sigma^2 = \frac{1}{n} e^2 d^2 \|Z_i\|_{\psi_2}^4$. \square

B.2 Largest eigenvector of covariance matrices

Proof of Lemma 5. Remark that

$$\langle \text{Cov}(\nu) \phi, \phi \rangle = \int_{\mathbb{R}^d} |\langle y - \mathbb{E}_\nu[Y], \phi \rangle|^2 d\nu(y) = \int_{\mathbb{R}^d} \langle y, u \rangle^2 d\nu(y) - \langle \mathbb{E}_\nu[Y], \phi \rangle^2.$$

Writing $\mathbb{E}_\nu[Y] = \int_{\mathbb{R}^d} y d\nu(y)$ gives

$$\begin{aligned} \langle \mathbb{E}_\nu[Y], \phi \rangle^2 &= \left\langle \int_{\mathbb{X}} y d\nu(y), \phi \right\rangle \left\langle \int_{\mathbb{X}} z d\nu(z), \phi \right\rangle \\ &= \int_{\mathbb{R}^d \times \mathbb{R}^d} \langle y, \phi \rangle \langle z, \phi \rangle d\nu(y) d\nu(z) \end{aligned}$$

which ends the proof. \square

Proof of Lemma 6. The tangent space $\mathcal{T}_\phi \mathbb{S}^{d-1} = \{v \in \mathbb{R}^d \mid \langle \phi, v \rangle_{\mathbb{R}^d} = 0\}$ and the metric is defined by $\mathfrak{g}_\phi(u, v) = \langle u, v \rangle_{\mathbb{R}^d}$, the Euclidean scalar product, for all $u, v \in \mathcal{T}_\phi \mathbb{S}^{d-1}$. With this metric, the logarithm map reads $\log_\phi(\psi) = (\psi - \langle \psi, \phi \rangle \phi) \frac{\theta}{\sin(\theta)}$ with $\theta = \arccos \langle \phi, \psi \rangle$. The exponential map is given by $\exp_\phi(\xi) = \cos(\|\xi\|_2) \phi + \sin(\|\xi\|_2) \frac{\xi}{\|\xi\|_2}$ for $\xi \in \mathcal{T}_\phi \mathbb{S}^{d-1}$. Furthermore, the geodesic from ϕ to ψ reads $\gamma(t) = \cos(\|\log_\phi(\psi)t\|) \phi + \sin(\|\log_\phi(\psi)t\|) \frac{\log_\phi(\psi)}{\|\log_\phi(\psi)\|}$. The geodesic distance between ϕ and ψ is therefore $\vartheta(\phi, \psi) = \arccos \langle \phi, \psi \rangle_{\mathbb{R}^d}$. One can check that $\|\phi - \psi\|_2 \leq \vartheta(\phi, \psi) \leq \frac{\pi}{2} \|\phi - \psi\|_2$. The proof consists in checking the assumptions one by one. To simplify notation the matrix A will refer to $\text{Cov}(\nu)$.

Assumption 1 The functions \mathcal{F} and $\widehat{\mathcal{F}}$ are continuous and defined on a compact domain $\mathcal{M} = \mathbb{S}^{d-1}$. Therefore, each of them admits at least one minimizer.

Assumption 2 Let $\phi \in \mathcal{S} = \{\pm u_1\}$ where u_1 refer to the eigenvector associated to the largest eigenvalue λ_1 . Remark that any point on the sphere is at distance at most $\pi/2$ of one of the minimizers i.e. $\vartheta(\psi, \mathcal{S}) \leq \pi/2$ for all $\psi \in \mathbb{S}^{d-1}$.

For a given $\psi \in \mathbb{S}^{d-1}$ we fix ϕ as the closest element of \mathcal{S} to ψ (since $\mathcal{S} = \{\pm u_1\}$). Let $\xi = \log_\phi(\psi)$ and one has that $\|\xi\|_\phi = \|\xi\|_2 \leq \frac{\pi}{2}$. The bound is obtained using

$$\begin{aligned} & \mathcal{F}(\exp_\phi(\xi)) - \mathcal{F}(\phi) \\ &= - \left(\cos^2(\|\xi\|) \langle A\phi, \phi \rangle + 2 \cos(\|\xi\|) \sin(\|\xi\|) \left\langle A\phi, \frac{\xi}{\|\xi\|} \right\rangle + \sin^2(\|\xi\|) \left\langle A \frac{\xi}{\|\xi\|}, \frac{\xi}{\|\xi\|} \right\rangle \right) \end{aligned}$$

Observing that $\langle A\phi, \xi \rangle = \langle \Lambda U^T \phi, U^T \xi \rangle = \lambda_d \langle \pm u_1, \xi \rangle = 0$ since $\xi \in \mathcal{T}_\phi \mathbb{S}^{d-1}$ we obtain

$$\begin{aligned} \mathcal{F}(\exp_\phi(\xi)) - \mathcal{F}(\phi) &= \frac{\sin^2(\|\xi\|)}{\|\xi\|_2^2} (\|\xi\|_2^2 \langle A\phi, \phi \rangle - \langle A\xi, \xi \rangle) \\ &\geq \frac{4}{\pi^2} (\lambda_1 - \lambda_2) \|\xi\|_2^2. \end{aligned}$$

The last inequality is obtained using Lemma 14 and from the fact that the function $t \mapsto \sin^2(t)/t^2$ is lower bounded by $\frac{4}{\pi^2}$ on $[0, \frac{\pi}{2}]$. Since the norm of the tangent vector $\|\xi\|_2$ is equal to the geodesic distance $\vartheta(\psi, \mathcal{S})$, this lead to

$$\frac{4}{\pi^2} (\lambda_1 - \lambda_2) \vartheta(\psi, \mathcal{S})^2 \leq \mathcal{F}(\psi) - \mathcal{F}_*$$

for all $\psi \in \mathcal{M}$. So that Assumption 2 holds for $\mathcal{F}_0 = +\infty$, $\beta = 2$ and $\tau = \frac{4}{\pi^2} (\lambda_1 - \lambda_2)$.

Assumption 3 Since Assumption 2 holds for $\mathcal{F}_0 = +\infty$, Assumption 3 holds with $\eta = 0$ since $\mathbb{P}(\mathcal{F}(\phi) > \mathcal{F}_0) = 0$.

Assumption 4 Let $\widehat{\lambda}_1 \geq \dots \geq \widehat{\lambda}_m$ be the eigenvalues of $\widehat{A} = \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}_{\mu_n}[X])(X_i - \mathbb{E}_{\mu_n}[X])^T$. The derivation of Assumption 2 can also be carried on $\widehat{\mathcal{F}}$ so that

$$\frac{4}{\pi^2} (\widehat{\lambda}_1 - \widehat{\lambda}_2) \vartheta(\psi, \widehat{\mathcal{S}})^2 \leq \widehat{\mathcal{F}}(\psi) - \widehat{\mathcal{F}}_*$$

holds for $\phi \in \mathcal{M}$. This bound can be used as long as $(\widehat{\lambda}_1 - \widehat{\lambda}_2) > 0$.

In virtue of the Weyl's inequality the distance between the eigenvalues can be bounded by $|\lambda_i - \widehat{\lambda}_i| \leq \|A - \widehat{A}\|_{2 \rightarrow 2}$ for all $1 \leq i \leq d$ so that the eigengap $\widehat{\lambda}_1 - \widehat{\lambda}_2$ can be lower bounded by

$$\begin{aligned} \widehat{\lambda}_1 - \widehat{\lambda}_2 &\geq \lambda_1 - \lambda_2 - |\widehat{\lambda}_1 - \lambda_1| - |\widehat{\lambda}_2 - \lambda_2| \\ &\geq \lambda_1 - \lambda_2 - 2\|A - \widehat{A}\|_{2 \rightarrow 2}. \end{aligned}$$

When $2\|A - \widehat{A}\|_{2 \rightarrow 2} \leq \frac{1}{2}(\lambda_1 - \lambda_2)$ the eigengap $\widehat{\lambda}_1 - \widehat{\lambda}_2 \geq \frac{1}{2}(\lambda_1 - \lambda_2)$. This event happens with high probability since by Lemma 13:

$$\mathbb{P} \left(\|A - \widehat{A}\|_{2 \rightarrow 2} > \frac{1}{4}(\lambda_1 - \lambda_2) \right) \leq d \exp \left(- \frac{n(\lambda_1 - \lambda_2)^2}{32e^2 d^2 \|X\|_{\psi_2}^4 + 8ed \|X\|_{\psi_2}^2 (\lambda_1 - \lambda_2)} \right).$$

Assumption 5 Let $\phi, \psi \in \mathcal{M}$ and $x, x' \in \mathbb{X}$ such that $x \neq x'$, we have

$$\begin{aligned}
& l(\phi, x) - l(\psi, x) - l(\phi, x') + l(\psi, x') \\
&= \langle y, \phi - \psi \rangle \langle y, \psi + \phi \rangle + \langle y', \psi - \phi \rangle \langle y', \psi + \phi \rangle \\
&\quad - \langle y, \phi - \psi \rangle \langle z, \phi \rangle - \langle y, \psi \rangle \langle z, \phi - \psi \rangle \\
&\quad + \langle y', \phi - \psi \rangle \langle z', \phi \rangle + \langle y', \psi \rangle \langle z', \phi - \psi \rangle \\
&\leq 2 (\|y\|_2^2 + \|y'\|_2^2 + \|y\|_2 \|z\|_2 + \|y'\|_2 \|z'\|_2) \|\phi - \psi\|_2 \\
&\leq 2 (\|y\|_2^2 + \|y'\|_2^2 + \|y\|_2 \|z\|_2 + \|y'\|_2 \|z'\|_2) \vartheta(\phi, \psi)
\end{aligned}$$

using $(a^2 - b^2) = (a - b)(a + b)$ for the first equality and where the the first inequality is obtained with the Cauchy-Schwartz inequality together with the fact that ψ and ϕ have unit norm. The function defined by

$$a(x, x') = 2 (\|y\|_2^2 + \|y'\|_2^2 + \|y\|_2 \|z\|_2 + \|y'\|_2 \|z'\|_2) \mathbf{1}_{x \neq x'}$$

for $x, x' \in \mathbb{X}$ is a pseudometric and therefore is a candidate for Assumption 5 to hold. Since X and X' share the same distribution, the sub-exponential norm of a can be bounded by

$$\begin{aligned}
\|a\|_{\psi_1} &\leq 4 \| \|Y\|^2 \|_{\psi_1} + 4 \| \|Y\| \|Z\| \|_{\psi_1} \\
&\leq 4 \| \|Y\| \|_{\psi_2}^2 + 4 \| \|Y\| \| \|Z\| \|_{\psi_2} \\
&\leq 8 \| \|Y\| \|_{\psi_2}^2
\end{aligned}$$

since Z has the same distribution as Y . Finally, using the fact that the sub-Gaussian norm $\| \|Y\| \|_{\psi_2} \leq \sqrt{d} \|Y\|_{\psi_2}$, the Assumption 5 is therefore satisfied with $\Upsilon = \mathcal{M}$, $\alpha = 1$ and the function a that is sub-exponential when X is sub-Gaussian. □

Lemma 14. Let $A \in \mathbb{R}^{d \times d}$ be a real symmetric matrix. Let $U = (u_1, \dots, u_d) \in \mathbb{R}^{d \times d}$ and $\lambda_1 \geq \dots \geq \lambda_d$ such that $A = U \Lambda U^T$ with $\Lambda = \text{diag}(\lambda_i)_{i=1}^d$. Let $v \in \mathbb{R}^d$ orthogonal to u_1 then

$$\|v\|_2^2 \langle Au_1, u_1 \rangle - \langle Av, v \rangle \geq (\lambda_1 - \lambda_2) \|v\|_2^2.$$

Proof. We have

$$\begin{aligned}
\|v\|_2^2 \langle Au_1, u_1 \rangle - \langle Av, v \rangle &= \|v\|_2^2 \langle \Lambda U^T u_1, U^T u_1 \rangle - \langle \Lambda U^T v, U^T v \rangle \\
&= \|v\|_2^2 \lambda_1 - \sum_{i=2}^d \lambda_i \langle u_i, v \rangle^2
\end{aligned}$$

since v is orthogonal to u_1 . The vector v can be decomposed in the orthogonal basis as $v = \sum_{i=2}^d \langle u_i, v \rangle u_i$ and therefore has a norm $\|v\|_2^2 = \sum_{i=2}^d \langle u_i, v \rangle^2$. We therefore conclude by remarking that

$$\|v\|_2^2 \lambda_1 - \sum_{i=2}^d \lambda_i \langle u_i, v \rangle^2 \geq \lambda_1 \|v\|_2^2 - \lambda_2 \|v\|_2^2.$$

□

B.3 LASSO

Proof of Lemma 7. Prior to the proof, some general observations on the problem are stated.

The functions \mathcal{F} and $\widehat{\mathcal{F}}$ are strictly convex and are piecewise polynomials of degree 2. Such functions satisfy a Łojasiewicz error bound of Assumption 2 with $\beta = 2$ [18, Corollary 9] for any $\mathcal{F}_0 > \mathcal{F}_*$, with an adequate $\tau(\mathcal{F}_0)$. This general result however does not specify the values of the constants \mathcal{F}_0 and τ which are important to derive proper bounds, and to check Assumption 4. Those values can be obtained from an intricate geometric analysis of \mathcal{F} that is summarized in Lemma 15.

This lemma involves Hoffman's constants [34, 32] which were first introduced in the study of error bounds for the solutions of systems of linear inequalities. Hoffman's constants can be defined in several equivalent ways, the following will be used [10]. Let $C \in \mathbb{R}^{k,l}$ be a matrix, its Hoffman's constant is defined by

$$H = \max_{\substack{I \subseteq \{1, \dots, k\} \\ C_I \text{ full row rank}}} \frac{1}{\sigma_{|I|}(C_I)} \quad (31)$$

where $C_I \in \mathbb{R}^{|I| \times l}$ refers to the submatrix built from the rows of indexes in I . This constant plays an important role in the derivation of error bounds for the LASSO problem as shown in Lemma 15.

Assumption 1 Since the functions \mathcal{F} and $\widehat{\mathcal{F}}$ are strictly convex and coercive, each of them has a unique minimizer.

Assumption 2 The following lemma will be used to prove this assumption.

Lemma 15 ([18, Lemma 10]). *Let $\mathcal{L} : \mathbb{R}^q \rightarrow \mathbb{R}$ defined by $\mathcal{L}(\phi) = \frac{1}{2} \|A\phi - b\|_2^2 + \lambda \|\phi\|_1$ where $\lambda > 0$, $b \in \mathbb{R}^p$ and A is a matrix of size $p \times q$. Let $R_* > 0$ be chosen such that the unique minimizer of \mathcal{L} , satisfies $\|\phi_*\|_1 \leq R_*$. Fix $R > R_*$, then*

$$\tau \vartheta(\phi, \phi_*)^2 \leq \mathcal{L}(\phi) - \mathcal{L}_*$$

for all $\phi \in \mathbb{R}^q$ such that $\|\phi\|_1 \leq R$ where

$$\tau = (4H^2 [1 + \lambda R + (R\|A\| + \|b\|)(4R\|A\| + \|b\|)])^{-1} \quad (32)$$

and $H = H(A, \lambda) > 0$ is the Hoffman's constant (31) of the matrix

$$C(A, \lambda) = \begin{pmatrix} E & -1_{\mathbb{R}^{2^q \times 1}} \\ 0_{\mathbb{R}^{1 \times q}} & 1 \\ A & 0_{\mathbb{R}^{p \times 1}} \\ 0_{\mathbb{R}^{1 \times q}} & \lambda \end{pmatrix} \in \mathbb{R}^{(2^q+1+p+1) \times (q+1)}. \quad (33)$$

where E is the matrix of size $2^q \times q$ with rows being all the possible distinct vectors of the form $(\pm 1, \dots, \pm 1)$.

Proof. This results is adapted from [18, Lemma 10]. Its statement only differs from the choice of R_* . In [18, Lemma 10] it is chosen to be $R_* = \frac{1}{2} \|b\|_2^2$ by observing that $\lambda \|\phi_*\|_1 \leq \mathcal{L}(\phi_*) \leq \mathcal{L}(0) = \frac{1}{2} \|b\|_2^2$. This choice is arbitrary, and any other R_* (satisfying $\|\phi_*\|_1 \leq R_*$) can be chosen. A more appropriate choice will ease the proof of the other assumptions. \square

In order to apply this lemma, the function \mathcal{F} has to be written in the form $\mathcal{F}(\phi) = \mathcal{L}(\phi) + c$ where \mathcal{L} corresponds to Lemma 15 and c is a constant term. Let

$$\text{Cov}_\theta(\mu_Y) = \int_Y \theta(y)\theta(y)^T d\mu_Y(y) \in \mathbb{R}^{m \times m}$$

It can be shown that

$$\begin{aligned} \mathcal{f}(\phi) &= \frac{1}{2} \langle \text{Cov}_\theta(\mu_Y)\phi, \phi \rangle - \langle \phi, \mathbb{E}_\mu[\theta(Y)V] \rangle + \frac{1}{2} \mathbb{E}_{\mu_V}[V^2] \\ &= \frac{1}{2} \langle \text{Cov}_\theta(\mu_Y)(\phi - \phi_0), (\phi - \phi_0) \rangle + \frac{1}{2} \mathbb{E}_{\mu_V}[V^2] - \frac{1}{2} \|\text{Cov}_\theta(\mu_Y)^{1/2} \phi_0\|_2^2. \end{aligned}$$

with ϕ_0 being chosen as $\text{Cov}_\theta(\mu_Y)\phi_0 = \mathbb{E}_\mu[\theta(Y)V]$. Such a vector always exists since it is a solution of the linear least square problem

$$\min_{\phi \in \mathbb{R}^m} \mathcal{f}(\phi). \quad (34)$$

Minimizing \mathcal{F} is therefore equivalent to minimizing \mathcal{L} defined by

$$\mathcal{L}(\phi) = \frac{1}{2} \left\| \text{Cov}_\theta(\mu_Y)^{1/2} \phi - b \right\|_2^2 + \lambda \|\phi\|_1.$$

setting $b = \text{Cov}_\theta(\mu_Y)^{1/2} \phi_0$. Remark at this point that since ϕ_0 is the solution of (34)

$$0 \leq \mathcal{f}(\phi_0) \leq \frac{1}{2} \mathbb{E}_{\mu_V}[V^2] - \frac{1}{2} \|\text{Cov}_\theta(\mu_Y)^{1/2} \phi_0\|_2^2$$

so that $\|b\|_2^2 \leq \mathbb{E}_{\mu_V}[V^2]$.

Finally, observe that $\lambda \|\phi_*\|_1 \leq \mathcal{F}(\phi_*) \leq \mathcal{F}(0) = \mathcal{f}(0) = \frac{1}{2} \mathbb{E}_{\mu_V}[V^2]$. The Lemma 15 can now be applied with $A = \text{Cov}_\theta(\mu_Y)^{1/2}$ and $R = \mathbb{E}_{\mu_V}[V^2]/\lambda$ leading to a bound

$$\tau \vartheta(\phi, \phi_*)^2 \leq \mathcal{L}(\phi) - \mathcal{L}_* = \mathcal{F}(\phi) - \mathcal{F}_*$$

as long as $\|\phi\|_1 \leq R$ for τ with expression given in (32). Using the fact that $\|b\|_2^2 \leq \mathbb{E}_{\mu_V}[V^2]$,

$$\tau \geq \left(4H^2 \left[1 + \lambda R + (R\|A\| + \sqrt{\mathbb{E}_{\mu_V}[V^2]})(4R\|A\| + \sqrt{\mathbb{E}_{\mu_V}[V^2]}) \right] \right)^{-1} \quad (35)$$

This lower-bound will be used in place of τ . This choice is not optimal but it will ease the proof of Assumption 4.

The choice $\mathcal{F}_0 = \mathbb{E}_{\mu_V}[V^2]/\lambda$ ensures that $\mathcal{F}(\phi) \leq \mathcal{F}_0$ implies that $\phi_1 \leq R$ so that the assumption is verified.

Assumption 3 Similarly $\widehat{\mathcal{F}}(\widehat{\phi}) \leq \widehat{\mathcal{F}}(0) = \frac{1}{2} \frac{1}{n} \sum_{i=1}^n V_i^2$ so that $\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} \widehat{\mathcal{F}}(\widehat{\phi}) \leq \mathcal{F}_0$ as soon as $\frac{1}{n} \sum_{i=1}^n V_i^2 \leq 2\mathbb{E}_{\mu_V}[V^2]$. This event happens with high probability since using a Bernstein inequality

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n V_i^2 - \mathbb{E}_{\mu_V}[V^2] \right| > t \right) \leq 2 \exp \left(- \frac{nt^2}{2(e^2 \|V^2\|_{\psi_1}^2 + e \|V^2\|_{\psi_1} t)} \right) \quad (36)$$

for all $t > 0$, because the random vector V^2 is sub-exponential with $\|V^2\|_{\psi_1} = \|V\|_{\psi_2}^2$. Choosing $t = \mathbb{E}_{\mu_{\mathbf{V}}}[V^2]$ therefore shows that the assumption is fulfilled with $\eta(n) = 2e^{-c_1 n}$ with

$$c_1 = \frac{\|V\|_2^4}{2(e^2\|V\|_{\psi_2}^4 + e\|V\|_{\psi_2}^2\|V\|_2^2)}.$$

Assumption 4 The Assumption will be proved checking that

$$\tilde{\tau} = \inf \left\{ \frac{\widehat{\mathcal{F}}(\phi) - \widehat{\mathcal{F}}_*}{\vartheta(\phi, \widehat{\mathcal{S}})^2} \mid \phi \in [\mathcal{F} \leq \mathcal{F}_0] \right\}$$

is greater than $\frac{1}{2}\tau$ with high probability. This will imply that Assumption 4 holds since $\widehat{\tau} \geq \tilde{\tau}$.

In this proof, the matrix A will refer to $\text{Cov}_{\theta}(\mu_{\mathbf{Y}})^{\frac{1}{2}}$ and \widehat{A} to its empirical counterpart. Similarly $Z = \mathbb{E}_{\mu_{\mathbf{V}}}[V^2]$ and $\widehat{Z} = \frac{1}{n} \sum_{i=1}^n V_i^2$.

Note that $\tilde{\tau}$ is defined in a similar fashion as (35) by

$$\tilde{\tau} = \left(4\widehat{H}^2 \left[1 + \lambda R + \left(R\|\widehat{A}\| + \sqrt{\widehat{Z}} \right) \left(4R\|\widehat{A}\| + \sqrt{\widehat{Z}} \right) \right] \right)^{-1}, \quad (37)$$

so that the high-probability bound will be obtained by:

- Proving that $\sqrt{\widehat{Z}} \leq 2^{1/4}\sqrt{Z}$. This is equivalent to proving that $\widehat{Z} \leq 2^{1/2}Z$. This happens under the event $|\widehat{Z} - Z| \leq (2^{1/2} - 1)Z$.
- Proving that $\widehat{H}^2 \leq 2^{1/2}H^2$. Using definition of H , the Hoffman's constant of the matrix $C(A, \lambda)$ defined in (33), and under the assumption that it is positive, it can be proved that

$$\lambda_{|I|}(C(A, \lambda)_I^T C(A, \lambda)_I) = \sigma_{|I|}^2(C(A, \lambda)_I) \geq \frac{1}{H^2} \quad (38)$$

for all subset of rows $I \subseteq \{1, \dots, 2^m + m + 2\}$. In virtue of Weyl's inequality, for all $1 \leq i \leq |I|$ and all subsets I ,

$$\begin{aligned} r_{I,i} &= |\lambda_i(C(A, \lambda)_I^T C(A, \lambda)_I) - \lambda_i(C(\widehat{A}, \lambda)_I^T C(\widehat{A}, \lambda)_I)| \\ &\leq \|C(A, \lambda)_I^T C(A, \lambda)_I - C(\widehat{A}, \lambda)_I^T C(\widehat{A}, \lambda)_I\| \\ &\leq \|C(A, \lambda)^T C(A, \lambda) - C(\widehat{A}, \lambda)^T C(\widehat{A}, \lambda)\| \end{aligned}$$

by the interlacing property of eigenvalues. Further computations show that

$$r_{I,i} \leq \left\| \begin{pmatrix} A^T A - \widehat{A}^T \widehat{A} & 0 \\ 0 & 0 \end{pmatrix} \right\| = \|A^T A - \widehat{A}^T \widehat{A}\| = \|\text{Cov}_{\theta}(\mu_{\mathbf{Y}}) - \text{Cov}_{\theta}(\mu_{\mathbf{Y},n})\|$$

In particular, under the event $\|\text{Cov}_{\theta}(\mu_{\mathbf{Y}}) - \text{Cov}_{\theta}(\mu_{\mathbf{Y},n})\| \leq \frac{1}{\epsilon H^2}$

$$\begin{aligned} \lambda_{|I|}(C(\widehat{A}, \lambda)_I^T C(\widehat{A}, \lambda)_I) &\geq \lambda_{|I|}(C(A, \lambda)_I^T C(A, \lambda)_I) - \frac{1}{\epsilon H} \\ &\geq \left(1 - \frac{1}{\epsilon}\right) \lambda_{|I|}(C(A, \lambda)_I^T C(A, \lambda)_I), \end{aligned} \quad (39)$$

using (38). This implies that under this event, the submatrices $C(\widehat{A}, \lambda)_I$ share the row rank of their counterpart $C(A, \lambda)_I$. Therefore, the set of constraints in the maximization problem (31) are equal. The equation (39) furthermore implies that $\widehat{H}^2 \leq \frac{\epsilon}{\epsilon-1} H^2$. Choosing $\epsilon = \frac{2^{1/2}}{2^{1/2}-1}$ gives the desired bound.

- Proving that $\|\widehat{A}\| \leq 2^{1/4}\|A\|$. This is equivalent to proving that $\|\text{Cov}_\theta(\mu_{\mathbb{Y}}, n)\| \leq 2^{1/2}\|\text{Cov}_\theta(\mu_{\mathbb{Y}})\|$ which will also be proved using the Weyl's inequality as

$$\left| \|\text{Cov}_\theta(\mu_{\mathbb{Y}})\| - \|\text{Cov}_\theta(\mu_{\mathbb{Y}, n})\| \right| \leq \|\text{Cov}_\theta(\mu_{\mathbb{Y}}) - \text{Cov}_\theta(\mu_{\mathbb{Y}, n})\|.$$

The expected bound happens under the event $\|\text{Cov}_\theta(\mu_{\mathbb{Y}}) - \text{Cov}_\theta(\mu_{\mathbb{Y}, n})\| \leq (2^{1/2} - 1)\|\text{Cov}_\theta(\mu_{\mathbb{Y}})\|$.

If all these inequalities hold, then

$$\begin{aligned} \tilde{\tau} &\geq \left(4\sqrt{2}H^2 \left[1 + \lambda R + \left(2^{1/4}R\|A\| + 2^{1/4}\sqrt{Z} \right) \left(2^{1/4}4R\|A\| + 2^{1/4}\sqrt{Z} \right) \right] \right)^{-1} \\ &\geq \left(4\sqrt{2}H^2 \left[\sqrt{2}(1 + \lambda R) + \left(2^{1/4}R\|A\| + 2^{1/4}\sqrt{Z} \right) \left(2^{1/4}4R\|A\| + 2^{1/4}\sqrt{Z} \right) \right] \right)^{-1} \\ &\geq \frac{1}{2}\tau. \end{aligned}$$

Using a union bound, together with (36) and Lemma 13 shows that

$$\begin{aligned} \mathbb{P} \left(\tilde{\tau} < \frac{1}{2}\tau \right) &\leq \mathbb{P} \left(|\widehat{Z} - Z| \geq (\sqrt{2} - 1)Z \right) + \mathbb{P} \left(\|\text{Cov}_\theta(\mu_{\mathbb{Y}}) - \text{Cov}_\theta(\mu_{\mathbb{Y}, n})\| \geq c \right) \\ &\leq 2e^{-n \frac{(\sqrt{2}-1)^2 \|V\|_2^4}{2(e^2 \|V\|_{\psi_2}^4 + e \|V\|_{\psi_2}^2 (\sqrt{2}-1) \|V\|_2^2)}} + me^{-n \frac{c^2}{2(e^2 m^2 \|\theta(Y)\|_{\psi_2}^4 + em \|\theta(Y)\|_{\psi_2}^2) c}} \end{aligned} \quad (40)$$

with $c = \min \left(\frac{\sqrt{2}-1}{\sqrt{2}} \frac{1}{H^2}, (\sqrt{2}-1)\|\text{Cov}_\theta(\mu_{\mathbb{Y}})\| \right)$. The assumption is therefore satisfied with $\kappa(n) = me^{-c_2 n}$ with c_2 being the minimum of the two exponents in (40).

Assumption 5 First, let us identify a candidate set Υ for which Assumption 5 holds. The inclusion of \mathcal{S} and of $\widehat{\mathcal{S}}$, with high probability, will be checked in the second part of the proof.

Let $R = \mathbb{E}_{\mu_{\mathbb{V}}} [V^2] / \lambda$ and $\Upsilon = \{\phi \in \mathcal{M} \mid \|\phi\|_2 \leq R\}$. For $\phi, \psi \in \Upsilon$ and $x = (y, v), x' = (y', v') \in \mathbb{Y} \times \mathbb{V}$ such that $x \neq x'$, one has

$$\begin{aligned} &l(\phi, x) - l(\psi, x') - l(\phi, x) + l(\psi, x') \\ &= \langle \phi - \psi, \theta(y) \rangle (\langle \phi + \psi, \theta(y) \rangle - 2v) - \langle \psi - \phi, \theta(y') \rangle (\langle \phi + \psi, \theta(y') \rangle - 2v') \\ &\leq \|\phi - \psi\|_2 (\|\theta(y)\|_2 |\langle \phi + \psi, \theta(y) \rangle - 2v| + \|\theta(y')\|_2 |\langle \phi + \psi, \theta(y') \rangle - 2v'|) \\ &\leq \|\phi - \psi\|_2 (\|\theta(y)\|_2 (2R\|\theta(y)\|_2 + 2|v|) + \|\theta(y')\|_2 (2R\|\theta(y')\|_2 + 2|v'|)). \end{aligned}$$

where we successively used the identity $(a^2 - b^2) = (a - b)(a + b)$, the Cauchy-Schwartz inequality. The function a defined by

$$a(x, x') = (\|\theta(y)\|_2 (2R\|\theta(y)\|_2 + 2|v|) + \|\theta(y')\|_2 (2R\|\theta(y')\|_2 + 2|v'|)) \mathbf{1}_{x \neq x'} \quad (41)$$

for $x, x' \in \mathbb{X}$ is a pseudometric. Moreover, since X and X' share the same distribution,

$$\begin{aligned}
\|a\|_{\psi_1} &\leq 4 \|\theta(Y)\|_2 (2R\|\theta(Y)\|_2 + 2|V|) \|_{\psi_1} \\
&\leq 4R \|\theta(Y)\|_2^2 \|_{\psi_1} + 4 \|\theta(Y)\|_2 \|V\|_{\psi_1} \\
&\leq 4R \|\theta(Y)\|_2^2 \|_{\psi_2} + 4 \|\theta(Y)\|_2 \|V\|_{\psi_2} \\
&\leq 4Rm \|\theta(Y)\|_{\psi_2}^2 + 4\sqrt{m} \|\theta(Y)\|_{\psi_2} \|V\|_{\psi_2}.
\end{aligned} \tag{42}$$

Therefore, the set Υ , $\alpha = 1$ and the function a are candidates for Assumption 5 to hold.

It is known from the proof of Assumption 2 that the norm of $\phi \in \mathcal{S}$ is $\|\phi\|_1 \leq R$. Using the fact that $\|\cdot\|_2 \leq \|\cdot\|_1$, we obtain $\mathcal{S} \subset \mathcal{M}$. Similarly, for $\phi \in \widehat{\mathcal{S}}$, $\|\phi\|_2 \leq \frac{1}{2n} \sum_{i=1}^n V_i^2$ so that under the event $\frac{1}{n} \sum_{i=1}^n V_i \leq 2\mathbb{E}_{\mu_V}[V^2]$, the norm $\|\phi\|_2 \leq R$. This event holds with probability at least $1 - \iota(n)$ with $\iota(n) = \eta(n)$, see (36) in the proof of Assumption 3. □

Proof. Proof of Corollary 5 The proof of the first part is a direct application of Theorem 2 with the constants of Lemma 7. The second part required a little more work.

The total expectation can be split accordingly to (4.3) where the first term can be bounded using Lemma 2. The second term is bounded as follows. Due to the the optimality of ϕ_* and $\widehat{\phi}$

$$\begin{aligned}
\|\phi_*\|_2 &\leq \|\phi_*\|_1 \leq \frac{1}{2\lambda} \mathbb{E}[V^2] \\
\|\widehat{\phi}\|_2 &\leq \|\widehat{\phi}\|_1 \leq \frac{1}{2\lambda} \frac{1}{n} \sum_{i=1}^n V_i^2,
\end{aligned}$$

see the proof of Lemma 7. Therefore,

$$\begin{aligned}
\mathbb{E} \left[\|\phi_* - \widehat{\phi}\|_2^2 \mathbf{1}_{\mathbf{X} \notin \mathcal{A}} \right] &\leq \frac{1}{2} \mathbb{E}[V^2]^2 \lambda^{-2} p_n + \frac{1}{2} \mathbb{E}[V^4 \mathbf{1}_{\mathbf{X} \notin \mathcal{A}}] \\
&\leq \frac{1}{2} \mathbb{E}[V^2]^2 \lambda^{-2} p_n + \frac{1}{2} \mathbb{E}[V^8]^{\frac{1}{2}} \sqrt{p_n} \\
&\leq (2\|V\|_{\psi_2}^2 + 2^5 \|V\|_{\psi_2}^4) \lambda^{-2} \sqrt{p_n}.
\end{aligned}$$

□

B.4 Entropic-Wasserstein barycenters

Proof of Lemma 8. The proof consist in checking the assumptions one by one.

Assumption 1 Since Ω is bounded there exists $M > 0$ such that $\int_{\mathcal{P}_2(\Omega)} \mathbb{E}_{X \sim \psi}[X^2] d\mu(\psi) \leq M^2$ for all measure $\mu \in \mathcal{P}_2(\mathcal{M})$. Using proposition Proposition 2.1 in [24] allows to conclude that ϕ_* and $\widehat{\phi}$ exist and are unique.

Assumption 2 The function $W_2^2(\cdot, \psi)$ is convex from Proposition 7.17 of [61] for all $\psi \in \mathcal{P}_2(\Omega)$ for the linear structure on $\mathcal{P}_2(\Omega)$. This means that

$$W_2^2(\phi_t, \psi) \leq (1-t)W_2^2(\phi_0, \psi) + tW_2^2(\phi_1, \psi)$$

for any pair of measures $\phi_0, \phi_1 \in \mathcal{P}_2(\Omega)$ and for any $t \in [0, 1]$ where $\phi_t = (1-t)\phi_0 + t\phi_1$. Integrating with respect to μ or μ_n shows that the functions \mathcal{F} and $\widehat{\mathcal{F}}$ are convex.

On the other hand, some computations show that

$$\mathcal{R}(\phi) = \mathcal{R}(\psi) + \langle \nabla \mathcal{R}(\psi), \phi - \psi \rangle + \text{KL}(\phi|\psi)$$

with $\nabla \mathcal{R}(\psi) = \log(\psi)$ and KL being the Kullback-Liebler divergence between ϕ and ψ defined by

$$\text{KL}(\phi|\psi) = \int_{\Omega} \log \left(\frac{\phi(z)}{\psi(z)} \right) \phi(z) dz.$$

The Pinsker's inequality states that $\|\phi - \psi\|_1 \leq \sqrt{\frac{1}{2} \text{KL}(\phi|\psi)}$. This means that \mathcal{R} is 2-strongly convex with respect to the total variation distance for the linear structure.

Hence \mathcal{F} and $\widehat{\mathcal{F}}$ are both 2-strongly convex with respect to the total variation distance so that

$$2\|\phi - \phi_*\|_1^2 \leq \mathcal{F}(\phi) - \mathcal{F}_*$$

for all $\phi \in \mathcal{P}_2^{ac}(\Omega)$. The similar inequality holds for $\widehat{\mathcal{F}}$.

Assumption 3 Since (2) is verified with $\mathcal{J}_0 = +\infty$, this assumption is verified with $\eta = 0$

Assumption 4 Since $\widehat{\mathcal{F}}$ is 2-strongly convex with respect to the total variation distance, it holds that $\widehat{\tau} = \tau$ with probability 1. This implies that the assumption is verified with $\kappa = 0$.

Assumption 5 Prior to proving this assumption, the Kantorovich duality is briefly introduced, see [76] for more details. Given two measures ϕ and ψ and two maps $g \in L^1(\phi)$ and $h \in L^1(\psi)$ satisfying $g(y) + h(z) \geq \langle y, z \rangle$ for ϕ -a.e. y and ψ -a.e. z , then by definition of W_2

$$\frac{1}{2}W_2^2(\phi, \psi) \geq \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - g \right) d\phi + \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - h \right) d\psi.$$

Due the Kantorovich duality, it is known that the equality holds for some pair of $(g, h) = (\zeta_{\phi \rightarrow \psi}, \zeta_{\phi \rightarrow \psi}^*)$ with $\zeta_{\phi \rightarrow \psi}$ being a lower semi-continuous proper convex function and where $\zeta_{\phi \rightarrow \psi}^* = \zeta_{\psi \rightarrow \phi}$ is the convex conjugate of $\zeta_{\phi \rightarrow \psi}$. The map $\zeta_{\phi \rightarrow \psi}$ is called the Kantorovich potential of the pair of measures (ϕ, ψ) and therefore satisfies

$$\frac{1}{2}W_2^2(\phi, \psi) = \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\phi \rightarrow \psi} \right) d\phi + \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\psi \rightarrow \phi} \right) d\psi.$$

Back to the proof of the assumption, let $\phi_0, \phi_1 \in \mathcal{M}$ and $\psi_0, \psi_1 \in \mathbb{X}$. By Kantorovich duality [76],

$$\begin{aligned} \frac{1}{2}W_2^2(\phi_0, \psi_0) &= \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\psi_0 \rightarrow \phi_0} \right) d\psi_0 + \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\phi_0 \rightarrow \psi_0} \right) d\phi_0 \\ \frac{1}{2}W_2^2(\phi_1, \psi_0) &\geq \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\psi_0 \rightarrow \phi_0} \right) d\psi_0 + \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\phi_0 \rightarrow \psi_0} \right) d\phi_1 \end{aligned}$$

so that

$$\frac{1}{2}W_2^2(\phi_0, \psi_0) - \frac{1}{2}W_2^2(\phi_1, \psi_0) \leq \int_{\Omega} \left(\frac{\|\cdot\|^2}{2} - \zeta_{\phi_0 \rightarrow \psi_0} \right) d(\phi_0 - \phi_1).$$

Therefore,

$$\begin{aligned} & \frac{1}{2}W_2^2(\phi_0, \psi_0) - \frac{1}{2}W_2^2(\phi_1, \psi_0) + \frac{1}{2}W_2^2(\phi_1, \psi_1) - \frac{1}{2}W_2^2(\phi_0, \psi_1) \\ & \leq \int_{\Omega} (\zeta_{\phi_1 \rightarrow \psi_1} - \zeta_{\phi_0 \rightarrow \psi_0}) d(\phi_0 - \phi_1) \end{aligned}$$

By definition, the Kantorovich ζ are $2 \operatorname{diam}(\Omega)$ -Lipschitz so that

$$\begin{aligned} \frac{1}{2}W_2^2(\phi_0, \psi_0) - \frac{1}{2}W_2^2(\phi_1, \psi_0) + \frac{1}{2}W_2^2(\phi_1, \psi_1) - \frac{1}{2}W_2^2(\phi_0, \psi_1) & \leq 4 \operatorname{diam}(\Omega) W_1(\phi_0, \phi_1) \\ & \leq 4 \operatorname{diam}(\Omega)^2 \|\phi_0 - \phi_1\|_1 \end{aligned}$$

where the last inequality is obtained through the transport inequality $W_1(\phi_0, \phi_1) \leq \operatorname{diam}(\Omega) \|\phi_0 - \phi_1\|_1$ see [31, Theorem 4]. Therefore, the assumption is verified with $\Upsilon = \mathcal{M}$, $a = 4 \operatorname{diam}(\Omega)^2$, $\|a\|_{\psi_1} = 4 \operatorname{diam}(\Omega)^2$ and $\iota = 0$.

□

References

- [1] Martial Agueh and Guillaume Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- [2] Adil Ahidar-Coutrix, Thibaut Le Gouic, and Quentin Paris. Convergence rates for empirical barycenters in metric spaces: curvature, convexity and extendable geodesics. *Probability theory and related fields*, 177(1):323–368, 2020.
- [3] Luigi Ambrosio and Nicola Gigli. A user’s guide to optimal transport. In *Modelling and optimisation of flows on networks*, pages 1–155. Springer, 2013.
- [4] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.
- [5] Miroslav Bacák. Computing medians and means in Hadamard spaces. *SIAM journal on Optimization*, 24(3):1542–1566, 2014.
- [6] Francis Bach. Learning theory from first principles. *Draft of a book, version of Sept, 6:2021*, 2021.
- [7] Eustasio Del Barrio, Paul Deheuvels, and Sara A. Geer. *Lectures on Empirical Processes: Theory and Statistical Applications*. European Mathematical Society, 2007.
- [8] Peter L Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

- [9] Saugata Basu, Richard Pollack, and Marie-Françoise Coste-Roy. *Algorithms in Real Algebraic Geometry*. Springer Science & Business Media. Google-Books-ID: 8Hyp1SLis4IC.
- [10] Amir Beck and Shimrit Shtern. Linearly convergent away-step conditional gradient for non-strongly convex functions. *Mathematical Programming*, 164(1):1–27, 2017.
- [11] M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005.
- [12] Ira D Berg and Igor G Nikolaev. Quasilinearization and curvature of Aleksandrov spaces. *Geometriae Dedicata*, 133(1):195–218, 2008.
- [13] Rabi Bhattacharya and Vic Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds. *The Annals of Statistics*, 31(1):1–29, 2003.
- [14] Edward Bierstone and Pierre D Milman. Semianalytic and subanalytic sets. *Publications Mathématiques de l’IHÉS*, 67:5–42, 1988.
- [15] Jérémie Bigot, Elsa Cazelles, and Nicolas Papadakis. Data-driven regularization of Wasserstein barycenters with an application to multivariate density registration. *Information and Inference: A Journal of the IMA*, 8(4):719–755, 2019.
- [16] Jérémie Bigot, Elsa Cazelles, and Nicolas Papadakis. Penalization of barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 51(3):2261–2285, 2019.
- [17] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real algebraic geometry*, volume 36. Springer Science & Business Media, 2013.
- [18] Jérôme Bolte, Trong Phong Nguyen, Juan Peypouquet, and Bruce W Suter. From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming*, 165(2):471–507, 2017.
- [19] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [20] Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to statistical learning theory. In *Summer school on machine learning*, pages 169–207. Springer, 2003.
- [21] Olivier Bousquet and André Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526, 2002.
- [22] Olivier Bousquet, Yegor Klochkov, and Nikita Zhivotovskiy. Sharper bounds for uniformly stable algorithms. In *Conference on Learning Theory*, pages 610–626. PMLR, 2020.
- [23] Victor-Emmanuel Brunel and Jordan Serres. Concentration of empirical barycenters in metric spaces. *arXiv preprint arXiv:2303.01144*, 2023.
- [24] Guillaume Carlier, Katharina Eichinger, and Alexey Kroshnin. Entropic-Wasserstein barycenters: PDE characterization, regularity, and CLT. *SIAM Journal on Mathematical Analysis*, 53(5):5880–5914, 2021.

- [25] Richard Combes. An extension of McDiarmid’s inequality. *arXiv preprint arXiv:1511.05240*, 2015.
- [26] Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 1996.
- [27] Vitaly Feldman and Jan Vondrak. Generalization bounds for uniformly stable algorithms. *Advances in Neural Information Processing Systems*, 31, 2018.
- [28] Vitaly Feldman and Jan Vondrak. High probability generalization bounds for uniformly stable algorithms with nearly optimal rate. In *Conference on Learning Theory*, pages 1270–1279. PMLR, 2019.
- [29] Cees M Fortuin, Pieter W Kasteleyn, and Jean Ginibre. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22(2):89–103, 1971.
- [30] Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. In *Annales de l’institut Henri Poincaré*, volume 10, pages 215–310, 1948.
- [31] Alison L Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *International statistical review*, 70(3):419–435, 2002.
- [32] Osman Güler. *Foundations of optimization*, volume 258. Springer Science & Business Media, 2010.
- [33] J-B Hiriart-Urruty. Extension of Lipschitz functions. *Journal of Mathematical Analysis and Applications*, 77(2):539–554, 1980.
- [34] Alan J Hoffman. On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards*, 49(4):263–265, 1952.
- [35] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- [36] Chi Jin, Praneeth Netrapalli, Rong Ge, Sham M Kakade, and Michael I Jordan. A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*, 2019.
- [37] Michael Kearns and Dana Ron. Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. In *Proceedings of the tenth annual conference on Computational learning theory*, pages 152–162, 1997.
- [38] Mojzesz Kirszbraun. Über die zusammenziehende und lipschitzsche transformationen. *Fundamenta Mathematicae*, 22(1):77–108, 1934.
- [39] Yegor Klochkov and Nikita Zhivotovskiy. Stability and deviation optimal risk bounds with convergence rate $O(1/n)$. *Advances in Neural Information Processing Systems*, 34:5065–5076, 2021.
- [40] Keith Knight and Wenjiang Fu. Asymptotics for lasso-type estimators. *Annals of statistics*, pages 1356–1378, 2000.

- [41] Aryeh Kontorovich. Concentration in unbounded metric spaces and algorithmic stability. In *International Conference on Machine Learning*, pages 28–36. PMLR, 2014.
- [42] Alexey Kroshnin. Fréchet barycenters in the Monge-Kantorovich spaces. *Journal of Convex Analysis*, 15:1371–1395, 2018.
- [43] Samuel Kutin. Extensions to McDiarmid’s inequality when differences are bounded with high probability. *Dept. Comput. Sci., Univ. Chicago, Chicago, IL, USA, Tech. Rep. TR-2002-04*, 2002.
- [44] Samuel Kutin and Partha Niyogi. Almost-everywhere algorithmic stability and generalization error. *Proc. 18th Conference on Uncertainty in Artificial Intelligence*, pages 275–282, 2002.
- [45] Thibaut Le Gouic and Jean-Michel Loubes. Existence and consistency of Wasserstein barycenters. *Probability Theory and Related Fields*, 168(3):901–917, 2017.
- [46] Thibaut Le Gouic, Quentin Paris, Philippe Rigollet, and Austin J Stromme. Fast convergence of empirical barycenters in Alexandrov spaces and the Wasserstein space. *Journal of the European Mathematical Society*, 2022.
- [47] Jing Lei et al. Convergence and concentration of empirical measures under Wasserstein distance in unbounded functional spaces. *Bernoulli*, 26(1):767–798, 2020.
- [48] Guoyin Li. Global error bounds for piecewise convex polynomials. *Mathematical programming*, 137(1):37–64, 2013.
- [49] Stanislas Łojasiewicz. Sur la géométrie semi-et sous-analytique. In *Annales de l’institut Fourier*, volume 43, pages 1575–1595, 1993.
- [50] Andreas Maurer and Massimiliano Pontil. Concentration inequalities under sub-Gaussian and sub-exponential conditions. *Advances in Neural Information Processing Systems*, 34, 2021.
- [51] Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pages 195–248. Springer, 1998.
- [52] Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Transactions on Information Theory*, 54(8):3797–3803, 2008.
- [53] Shahar Mendelson. Learning without concentration. In *Conference on Learning Theory*, pages 25–39. PMLR, 2014.
- [54] Shahar Mendelson. Learning without concentration for general loss functions. *Probability Theory and Related Fields*, 171(1):459–502, 2018.
- [55] Karl Pearson. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11):559–572, 1901.
- [56] Bruno Pelletier. Informative barycentres in statistics. *Annals of the Institute of Statistical Mathematics*, 57(4):767–780, 2005.

- [57] Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- [58] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [59] Alexander Rakhlin, Sayan Mukherjee, and Tomaso Poggio. Stability results in learning theory. *Analysis and Applications*, 3(04):397–417, 2005.
- [60] William H Rogers and Terry J Wagner. A finite sample distribution-free performance bound for local discrimination rules. *The Annals of Statistics*, pages 506–514, 1978.
- [61] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.
- [62] Christof Schötz. Convergence rates for the generalized Fréchet mean via the quadruple inequality. *Electronic Journal of Statistics*, 13(2):4280–4345, 2019.
- [63] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [64] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *COLT*, volume 2, page 5, 2009.
- [65] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670, 2010.
- [66] Austin James Stromme. *Wasserstein barycenters: statistics and optimization*. PhD thesis, Massachusetts Institute of Technology, 2020.
- [67] Karl-Theodor Sturm. Probability measures on metric spaces of nonpositive curvature. volume 338, page 357. American Mathematical Soc., 2003.
- [68] Michel Talagrand. *Upper and lower bounds for stochastic processes*, volume 60. Springer, 2014.
- [69] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [70] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12:389–434, 2012.
- [71] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [72] Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991.
- [73] Vladimir Vapnik and Alexey Chervonenkis. Theory of pattern recognition, 1974.

- [74] VN Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264, 1971.
- [75] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [76] Cédric Villani. *Topics in optimal transportation*. Number 58. American Mathematical Soc., 2003.
- [77] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [78] Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.
- [79] Krzysztof Zająkowski. Norms of sub-exponential random vectors. *Statistics & Probability Letters*, 152:147–152, 2019.