



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. *Journal of Machine Learning Research*, 2024, 25 (251), pp.1-53. hal-04054021v4

HAL Id: hal-04054021

<https://hal.science/hal-04054021v4>

Submitted on 10 Sep 2024

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

PAUL.ESCANDE@MATH.UNIV-TOULOUSE.FR

*Institut de Mathématiques de Toulouse
UMR 5219, Université de Toulouse, CNRS
UPS, F-31062 Toulouse Cedex 9, France*

Editor: Joseph Salmon

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 high-level assumptions allowing to describe a regime that seems 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.

Keywords: empirical risk minimization, statistical learning, concentration inequalities, metric spaces, local Hölder-type error bound

1 Introduction

Let (\mathbb{X}, ρ, μ) be a metric probability space¹ and let (\mathcal{M}, ϑ) be another metric space. 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 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

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

Minimization (ERM) (Vapnik, 1991) 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. This work 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 (Valiant, 1984) with the important distinction that it proves bound on the estimation error $\mathcal{F}(\widehat{\phi}) - \inf \mathcal{F}$, also known as the excess risk, 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

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

which are the Fréchet means of μ and $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ respectively (Fréchet, 1948). 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 convex when the space \mathcal{M} are of non-positive curvature in the sense of Alexandrov (see Sturm, 2003). Otherwise, the risks might not be convex since the distance function might exhibit some concavity properties (see Ambrosio et al., 2008).

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 $\text{Cov}(\mu) = \mathbb{E} [XX^T]$, the covariance of a centered measure 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 an instance of (1) and (2) with $l(\phi, x) = -\langle \phi, x \rangle^2$, 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

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

with $V \in \mathbb{R}^n$ being the vector of outputs $V = (v_i)_{i=1}^n$ and $\Theta \in \mathbb{R}^{n \times m}$ is the design matrix with the vectors $(\theta(y_i)^T)_{i=1}^n$ as rows. The risks \mathcal{F} and $\widehat{\mathcal{F}}$ 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{F} and $\widehat{\mathcal{F}}$ 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 of Vapnik and Chervonenkis (1971, 1974). Two main classes of analyses can apparently be identified: the techniques using uniform convergence and the ones from algorithm stability. They are described in the following two paragraphs and will be linked to the question at stake in the third paragraph, that is 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 (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 Shalev-Shwartz et al., 2010, and reference therein).

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

1.2.1 UNIFORM CONVERGENCE

The techniques in the uniform convergence class can be tracked back to the 1970's with the works of Vapnik and Chervonenkis (1971, 1974); Vapnik (1991) and the PAC framework proposed in the 1980's by Valiant (1984). 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 Devroye et al., 1996; Shalev-Shwartz and Ben-David, 2014; Wainwright, 2019; Vershynin, 2018; Bach, 2024, to name a few).

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

$$\begin{aligned} \mathcal{F}(\widehat{\phi}) - \mathcal{F}_* &\leq \mathcal{F}(\widehat{\phi}) - \mathcal{F}(\phi_*) + \widehat{\mathcal{F}}(\phi_*) - \widehat{\mathcal{F}}(\widehat{\phi}) \quad (\text{since } \widehat{\phi} \in \arg \min \widehat{\mathcal{F}}) \\ &= \ell(\widehat{\phi}) - \ell(\phi_*) + \widehat{\ell}(\phi_*) - \widehat{\ell}(\widehat{\phi}) \quad (\text{since the terms in } \mathcal{R} \text{ cancel out}) \\ &\leq 2 \sup_{\phi \in \mathcal{M}} \Delta(\phi) \end{aligned} \tag{5}$$

where $\Delta(\phi) = |\ell(\phi) - \widehat{\ell}(\phi)|$. The uniform bound is taken to remove the dependency between $\widehat{\phi}$ and $\widehat{\ell}$ that was preventing the $(l(\widehat{\phi}, X_i))_{i=1}^n$ to be independent. This allows to use standard concentration tools such as the McDiarmid's inequality (McDiarmid, 1998) (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

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

with high probability at least $1 - \delta$. The downside of this uniform bound is to control $\mathbb{E} \left[\sup_{\phi \in \mathcal{M}} \Delta(\phi) \right]$ (that is, 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 latter can further be bounded by the Rademacher complexity or the Gaussian complexity of the function class $\{(l(\phi, x_i))_{i=1}^n \mid \phi \in \mathcal{M}\}$ (Bartlett and Mendelson, 2002). Note that when the values of l are discrete, the Rademacher complexity can be linked to the VC dimension of the class of functions (Bousquet et al., 2003). 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 (Wainwright, 2019; Shalev-Shwartz and Ben-David, 2014). Note that the use of Rademacher complexities can be improved by considering localized versions (Bartlett et al., 2005; Koltchinskii, 2006).

Second, when the cost l exhibits some smoothness properties (for example Lipschitz) with respect to the variable ϕ , the quantity $\mathbb{E} \left[\sup_{\phi \in \mathcal{M}} \Delta(\phi) \right]$ can alternatively be bounded using chaining techniques such as Dudley's chaining or the Talagrand's generic chaining (Talagrand, 2014; Vershynin, 2018). This leads to an upper-bound dependent on a measure of complexity of \mathcal{M} , such as 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 (see Ahidar-Coutrix et al., 2020; Schötz, 2019). In the case $s > 1$, the rate is customarily said to be *cursed by the dimension*.

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 (see Mendelson, 2008).

1.2.2 ALGORITHM STABILITY

Another body of literature studies the generalization of learning algorithms through stability. These ideas can be tracked back to the 1970's as well with Rogers and Wagner (1978) and later popularized in the founding article of Bousquet and Elisseeff (2002). The notions involved in these works will be introduced here for the case of the empirical minimizers, although they can be applied to the broader setting of learning algorithms.²

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 independent copy of X_1 . The random 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 denoted by $\tilde{\mathcal{J}}$ and its minimizer by $\tilde{\phi}$. It is important to remark at this point that the risk $\tilde{\mathcal{J}}$ 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 thought as exchanging the first component.

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

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

because the roles of Y and X_1 can be switched as being independent copies. This decomposition emphasizes the link between the estimation error and the *stability* of the empirical minimizer, that is 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 such as the hypothesis stability (Rogers and Wagner, 1978; Kearns and Ron, 1997) or the uniform one (Bousquet and Elisseeff, 2002). 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 $x \in \mathbb{X}$ and all $(\mathbf{X}, Y) \in \mathbb{X}^{n+1}$. This stability can be leveraged to directly bound the expected estimation error as $\mathbb{E}_{\mathbf{X}} \left[\mathcal{J}(\hat{\phi}) - \mathcal{J}_* \right] \leq \epsilon$, and obtain (see Bousquet and Elisseeff, 2002, Theorem 12) that

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

2. More 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 \mathcal{J} or an approximation of it obtained using any algorithm (for example gradient descent).

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) cannot be cursed by the dimension since it cannot depend exponentially on the complexity of the space \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} (Kontorovich, 2014; Maurer and Pontil, 2021).

Recently, a series of papers (Feldman and Vondrak, 2018, 2019; Bousquet et al., 2020) improved (8) to

$$\mathcal{F}(\widehat{\phi}) - \mathcal{F}_* \leq \epsilon \log(n) \log\left(\frac{1}{\delta}\right) + M \sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)}, \quad (9)$$

holding 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 these 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 (Kutin and Niyogi, 2002; Rakhlin et al., 2005).

1.2.3 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.

As a preliminary, it can be mentioned that when the minimizers of \mathcal{F} and $\widehat{\mathcal{F}}$ have explicit expressions with a clear dependency on the measure μ and the sample $(X_i)_{i=1}^n$, their structure, if any, could be leveraged to directly use concentration tools. This could happen for example when \mathcal{F} and $\widehat{\mathcal{F}}$ are quadratic, which is however restrictive for modern machine learning applications.

The first observation towards providing an indirect analysis, is that when an 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

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

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 (Vershynin, 2018),

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 $c > 0$ (Jin et al., 2019). 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 $\sqrt{n^{-1} \log(\delta^{-1})}$ would slow down the rate nevertheless. The same slow rate is obtained with the uniform convergence (6).

This gap is known in the literature (see Shalev-Shwartz et al., 2009, 2010; Klochkov and Zhivotovskiy, 2021) 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, such as (10), 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 $\epsilon = O(n^{-1})$ (see Shalev-Shwartz et al., 2010). Yet necessary, this improvement is still vain since the rate in (9) would now be controlled by the sampling error in $\sqrt{n^{-1} \log(\delta^{-1})}$. Shalev-Shwartz et al. (2009) 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 by Klochkov and Zhivotovskiy (2021) 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 $c > 0$ 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 in Section 1.1.

Probably one of the first work attempting to understand the fundamental mechanisms driving the concentration of the empirical minimizers is from van de Geer and Wainwright (2017) which was building on a former work of Chatterjee (2014). They aim at bounding, with the notation of this paper, the excess risk $\ell(\widehat{\phi}) - \ell(\phi_0) + \mathcal{R}(\widehat{\phi})$ where $\phi_0 = \arg \min \ell$ is the minimizer of the non-penalized version of (1). It is again assumed that ϕ_0 exists and is unique and that the penalty \mathcal{R} is convex. This question is different from the one at stake in this work, since (3) only seeks to provide bounds for the random part, that is between $\widehat{\phi}$ and ϕ_* . Relying on a curvature condition on the objective function around its minimizers, the analysis then uses uniform convergence techniques to provide the concentration bounds on the excess risk towards the expected excess risk. The guarantees on $\|\widehat{\phi} - \phi_0\|$ are then derived in the case where $\|\widehat{\phi} - \phi_0\|^2 = \ell(\widehat{\phi}) - \ell(\phi_0)$, which excludes many applications.

Recently, two works by Schötz (2019) and Ahidar-Coutrix et al. (2020) have explored the concentration of $\vartheta(\widehat{\phi}, \phi_*)$ using uniform convergence techniques too. Both assume a growth of \mathcal{F} around its minimizers of type (10) without assuming the strong convexity of l . In (Ahidar-Coutrix et al., 2020), the 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 (Schötz, 2019), the analysis is more general: \mathcal{M} can be different from \mathbb{X} and \mathcal{M} , \mathbb{X} and l can be 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 a first significant understanding on the question, the bounds obtained depend on the entropy of the space \mathcal{M} , which might lead to rates suffering from the curse of dimensionality. Indeed, in the example of the estimation of barycenters in non-positively curved spaces, the rates of Schötz (2019); Ahidar-Coutrix et al. (2020) might be cursed by the dimension, and therefore suboptimal, since it is known that parametric rates hold regardless of the entropy of \mathcal{M} (see Le Gouic et al., 2022).

1.3 Contributions

The review of the literature indicates that available results are insufficient to provide optimal concentration rates for $\vartheta(\arg \min \widehat{\mathcal{F}}; \arg \min \mathcal{F})$ for a wide variety of estimation problems, such as those described in Section 1.1. This work attempts to fill that gap by providing a set of high-level assumptions from which optimal concentration rates can be proven. These assumptions fit most of the 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 are not exponentially cursed by 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 space (\mathbb{X}, μ) is defined by $\|X\|_{L^p} = \mathbb{E}[|X|^p]^{1/p}$. This definition can be extended to the case $p = +\infty$ with the essential supremum $\|X\|_{L^\infty} = \text{ess sup } |X|$.

For $q \geq 1$, the q -exponential Orlicz norm of a random variable X on the probability space (\mathbb{X}, 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 $\|X\|_{L^p} \leq \|X\|_{\psi_q} p^{1/q}$. These properties can be shown to be equivalent (see Vershynin, 2018, for more details). Importantly, $\|X^2\|_{\psi_1} = \|X\|_{\psi_2}^2$ and $\|XY\|_{\psi_1} \leq \|X\|_{\psi_2} \|Y\|_{\psi_2}$ (see Vershynin, 2018, Lemma 2.7.6 and 2.7.7).

The p -th norm of vectors in \mathbb{R}^d will be written using the usual notation $\|\cdot\|_p$. 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, $\| \|X\|_2 \|_{\psi_q} \leq \sqrt{d} \|X\|_{\psi_q}$ (see Vershynin, 2018, Section 3.4).

Following Boucheron et al. (2013, 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 S if its logarithmic moment generating function satisfies

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

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

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

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

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

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, that is $\mathbf{X} = (X_1, \dots, X_n)$ for random vectors and $\mathbf{x} = (x_1, \dots, x_n)$ for vectors. Let Y be an independent copy of X_1 . The random vector \mathbf{X}' is defined by $\mathbf{X}' = (Y, X_2, \dots, X_n)$. 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 and its minimizers depend 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

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

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, in the sense that $\vartheta(\mathcal{E}_1; \mathcal{E}_2) \leq \vartheta(\mathcal{E}_1; \mathcal{E}_3) + \vartheta(\mathcal{E}_3; \mathcal{E}_2)$ for all $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ non-empty subsets of \mathcal{M} . The reverse triangle inequality holds in the form $|\vartheta(\mathcal{E}_1; \mathcal{E}_2) - \vartheta(\mathcal{E}_3; \mathcal{E}_2)| \leq \max(\vartheta(\mathcal{E}_1; \mathcal{E}_3), \vartheta(\mathcal{E}_3; \mathcal{E}_1))$.

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, for example $[\mathcal{F} \geq m] = \{\phi \in \mathcal{M} \mid \mathcal{F}(\phi) \geq m\}$ 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 the paper of Sturm (2003).

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, that is 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 μ^n -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, as functions of \mathbf{x} , are all assumed to be measurable.*

On a metric space (\mathcal{M}, ϑ) satisfying the Heine–Borel property (that is for which any closed and bounded subspace is compact), at least one minimizer exists 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.

The measurability assumption of the minimizers is made upfront to avoid technical considerations about their existence, which will be clear in usual applications. It can be shown, under mild conditions that there will be at least one measurable minimizer in $\widehat{\mathcal{S}}$ (see Hess, 1996). Moreover, the measurability assumptions can be weakened using outer expectations (see Van Der Vaart and Wellner, 1997).

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 $\phi \in \mathcal{M}$,*

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

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 $\mathcal{J}_0 = +\infty$. 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 (Bochnak et al., 2013), the bound (15) holds on any compact subset (Łojasiewicz, 1993, 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 (Basu et al., 2007, 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 (Bierstone and Milman, 1988, Theorem 6.4).

This assumption appears in many communities under different names such as local Hölder-type bound in functional analysis (Li, 2013), Łojasiewicz inequality in real algebraic geometry (Łojasiewicz, 1993) or Łojasiewicz error bound inequality in optimization (Bolte et al., 2017). In the particular case of Fréchet means where $\mathbb{X} = \mathcal{M}$ and $l = \rho^2$, the Assumption 2 is known as variance inequality (Sturm, 2003). It is also called a margin condition in the works of Barrio et al. (2007) and of van de Geer and Wainwright (2017) or a low noise assumption in the one of Ahidar-Coutrix et al. (2020).

In the ERM literature, the Bernstein condition is often presented as central to derive fast concentration rates (see for example Koltchinskii, 2006; Bartlett and Mendelson, 2006; Klochkov and Zhivotovskiy, 2021) which assumes that there exists a constant c such that for all $\phi \in \mathcal{M}$, there exists a $\phi_* \in \mathcal{S}$ such that $\mathbb{E}[(l(\phi, X) - l(\phi_*, X))^2] \leq c(\mathcal{F}(\phi) - \mathcal{F}_*)$. This condition is in particular verified when \mathcal{F} satisfies the Assumption 2 for $\beta = 2$, $\mathcal{J}_0 = +\infty$ and that $l(\cdot, X)$ is uniformly Lipschitz.

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

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

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 (Bolte et al., 2017). 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{F}_* so that $\widehat{\phi}$ would fall into the level-set $[\mathcal{F} \leq \mathcal{F}_0]$.

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 helpful. Following (5), if $\sup_{\phi \in \mathcal{M}} \Delta(\phi) \leq t$ holds with high probability, for example when the Rademacher complexity of the function class or the entropy of \mathcal{M} are bounded, then setting $t = (\mathcal{F}_0 - \mathcal{F}_*)/2$ 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 (Mendelson, 2014, 2018) might be used to prove that Assumption 3 holds. Finally, the convergence to the basin might also be proved considering approaches like the one of Schötz (2019).

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

Let $\mathbf{X} \in \mathbb{X}^n$ and $\mathbf{y} \in \text{supp}(\mu)^n$. Let $\widetilde{\mathcal{F}}$, respectively $\widetilde{\mathcal{S}}$, refer to the empirical risk and the empirical set of minimizers associated to \mathbf{y} . Let $\widehat{\tau}$ be the random variable defined by

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

The constant $\widehat{\tau}$ quantifies the behavior of $\widehat{\mathcal{F}}$, around its minimizers, regarding only the minimizers (falling into the basin $[\mathcal{F} \leq \mathcal{F}_0]$) of all possible empirical risks.

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*

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

Remark 2. *Similarly to Remark 1, it might only be possible to compute a lower bound to $\widehat{\tau}$. In this assumption, the lower bound however has to be sufficiently large to be greater 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 9.*

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

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

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 $\widehat{\mathcal{F}}$ regarding only the minimizers of $\widetilde{\mathcal{F}}$.

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 Lojasiewicz constants of \mathcal{F} and $\widehat{\mathcal{F}}$.

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 $a : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_+$ 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 function a should further satisfy $\|a(X, Y)\|_{\psi_1} < +\infty$ for independent $X, Y \sim \mu$.

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 $\mathbb{P}(\widehat{\mathcal{S}} \not\subseteq \Upsilon) \leq \iota(n)$.

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 $a(x, y)$ with the integrability condition that $a(X, Y)$ should be sub-exponential. In the remaining of the paper, the notation $\|a\|_{\psi_1}$ will be a shorthand to the quantity $\|a(X, Y)\|_{\psi_1}$ for independent $X, Y \sim \mu$. Note that $\|a\|_{\psi_1}$ is oftentimes called the sub-exponential diameter of the space (\mathbb{X}, a, μ) (Maurer and Pontil, 2021).

Remark that when $l(\cdot, x)$ satisfy a Hölder condition, that is 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 (Schötz, 2019). In the Fréchet mean setting, that is 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 (Sturm, 2003, 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.

3. A 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 $a(x, y) = 0 \implies x = y$.

Theorem 5. *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 $b : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}^+$ 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 (for example McDiarmid, 1998; Kutin, 2002; Kutin and Niyogi, 2002; Combes, 2015; Warnke, 2016). This result is a combination of ideas from Maurer and Pontil (2021) and Combes (2015) 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

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

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 the components of \mathbf{X} are perturbed. Let $\mathbf{X}, \mathbf{Y} \in \mathbb{X}^n$. Conditionally on the event $\mathbf{X} \in \mathcal{A}, \mathbf{Y} \in \mathcal{A}$, the minimizers in $\widehat{\mathcal{S}}$ and $\widetilde{\mathcal{S}}$ fall into the level set $[\mathcal{F} \leq \mathcal{F}_0]$. Furthermore, conditionally on this event, the Lojasiewicz constant $\widehat{\tau} \geq \frac{1}{2}\tau$ so that, for a minimizer $\widetilde{\phi} \in \widetilde{\mathcal{S}}$,

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

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

Lemma 6. *For $\beta > \alpha$ and conditionally on the event $\mathbf{X} \in \mathcal{A}$, $\mathbf{Y} \in \mathcal{A}$, it holds that*

$$\vartheta(\tilde{\mathcal{S}}; \hat{\mathcal{S}})^{\beta-\alpha} \leq \frac{2}{\tau} \frac{1}{n} \sum_{i=1}^n a(X_i, Y_i).$$

This notion of stability of the empirical minimizers is, in essence, similar to the notions of stability of Bousquet and Elisseeff (2002). Since the probability p_n is controlled by Assumptions 3, 4 and 5, the stability happens with high-probability, and echoes the work of Kutin and Niyogi (2002).

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

Lemma 7. *Under Assumptions 1–5 with the further condition that $p_n \leq 3/4$ then*

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

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 Y and X_1 in the expectation. This step allows to leverage the stability of the empirical minimizers of Lemma 6. However, due to the conditioning on $\mathbf{X} \in \mathcal{A}$, the random variables Y and X_i 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 of Bousquet and Elisseeff (2002); Feldman and Vondrak (2018, 2019); Bousquet et al. (2020). For these reasons, the proof of this lemma is postponed to Section 5.3.

The next corollary is an instance of Lemma 7 for particular values of α and β .

Corollary 8. *When $\beta = 2$ and $\alpha = 1$, the bound of Lemma 7 reads*

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

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

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

4.2 Concentration in Probability

The main results of this paper are now stated.

Theorem 9. *Suppose Assumptions 1–5 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, for any $\delta > 0$,*

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

with probability at least $1 - p_n - \delta$ where $C_n = L^{\frac{1}{Q}} \left(K^{\frac{q}{\beta}} + 2^{2+\frac{1}{Q}} n^{1-\frac{1}{Q}} \right)$ and where the constants c_1 , K and L are defined in Lemma 7.

Proof Let $f(\mathbf{X}) = \vartheta(\widehat{\mathcal{S}}; \mathcal{S})^q$. The McDiarmid's inequality of Theorem 5 is applied to bound $f(\mathbf{X}) - \mathbb{E}[f | \mathbf{X} \in \mathcal{A}]$ with high probability.

The conditional expectation will be bounded using Lemma 7 as

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

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 5. 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 $\mathbf{x}, \mathbf{y} \in \mathcal{A}$,

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

using the triangle inequality combined with Lemma 6. Therefore, the choice $b = \left(\frac{2}{n\tau} a\right)^{\frac{1}{Q}}$ satisfies the assumptions of Theorem 5 so that

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

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 $\mathbb{E}[f(\mathbf{X}) | \mathbf{X} \in \mathcal{A}]$ together with the fact that $\frac{q}{2\beta} < \frac{1}{2}$ proves the claimed result. \blacksquare

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

Corollary 10. *When $\beta = 2$ and $\alpha = 1$, the bound of Theorem 9 reads*

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

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 results. As long as β and α are independent of the complexity of \mathcal{M} , the rate of convergence will not depend exponentially on the dimension of \mathcal{M} . 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{F} and $\widehat{\mathcal{F}}$ 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 (Schötz, 2019) might be able to provide a usable rate.

Arguably the most studied case is the one where \mathcal{M} is a Hilbert space and the losses $l(\cdot, x)$ are bounded, ω -strongly convex and K -Lipschitz uniformly for μ -almost every $x \in \mathbb{X}$. This case was discussed in Section 1.2. A combination of (10) and the Proposition 2.1 of Klochkov and Zhivotovskiy (2021) was showing that, for all $\delta > 0$, the rate

$$\vartheta(\widehat{\phi}, \phi_*)^2 \leq c \frac{K^2 \log(n)}{\omega^2 n} \log\left(\frac{1}{\delta}\right), \quad (20)$$

was holding with high probability $1 - \delta$. Such functions guarantee that the Assumptions 1–5 hold for $\tau = \omega$, $\beta = 2$, $\alpha = 1$, $\|a\|_{\psi_1} = 2K$, $\mathcal{F}_0 = +\infty$ and $\eta = \iota = \kappa = 0$. In view of Corollary 10, a bound

$$\vartheta(\widehat{\phi}, \phi_*)^2 \leq c \frac{K^2}{\omega^2 n} \left(1 + \log\left(\frac{1}{\delta}\right) + \frac{1}{n} \log\left(\frac{1}{\delta}\right)^2 \right), \quad (21)$$

holding with high probability $1 - \delta$ is obtained for all $\delta > 0$. Note that the constants $c > 0$ in the two bounds are different. Up to the logarithm factor, the bound (20) is a parametric, while (21) gives a true parametric rate. The bounds also differ in their nature, (20) is a 0-concentration bound while (21) is a mean-concentration one. Finally, the bound (21) contains an extra term $n^{-1} \log(\delta^{-1})^2$ which inherits from the sub-exponential nature of $\|a\|_{\psi_1}$. With the further assumption that the losses are K -Lipschitz, it is possible to use a standard McDiarmid inequality, since in this case Lemma 6 would give $\vartheta(\widehat{\mathcal{S}}; \widetilde{\mathcal{S}}) \leq Kn^{-1} \sum_{i=1}^n \mathbb{1}_{X_i \neq Y_i}$. This strategy would remove the extra term $n^{-1} \log(\delta^{-1})^2$ of (21).

The case $\beta = 2$ and $\alpha = 1$ is particularly important in modern statistical learning because: (i) the losses l have bounded sub-gradients in order to use first-order descent algorithms and (ii) the risks are designed to have a strong local convexity to accelerate the convergence of optimization algorithms. The reason why the case $\beta = 2$ is often encountered can also be justified in view of a Taylor expansion of the risk (assuming that \mathcal{M} is a Riemannian manifold, proper smoothness of \mathcal{F} and the uniqueness of its minimizers). Sufficiently close to the minimizers, the quadratic growth driven by the Hessian of \mathcal{F} , if not vanishing, dominates the higher-order terms. This leads to local quadratic error bounds. Estimation

problems where $\beta \neq 2$ and $\alpha \neq 1$ can be constructed but are usually obtained by designing pathological settings. Such examples can be found, for instance, in the literature of power Fréchet means in Hadamard spaces (Schötz, 2023; Yun and Park, 2023). Moreover, in the literature of Fréchet means on positively curved spaces, it is possible to design ad-hoc measures μ to obtain flatter local error bounds, which is often called smeariness. An example on the sphere was found by Eltzner and Huckemann (2019) where the Hessian of the risk vanishes at the minimizer, leading to higher order polynomials of the Taylor expansion to drive the growth around the minimizer.

Finally, the optimality of the rates Theorem 9 is difficult to assess because of its level of generality. However, its application to various scenarii provides the optimal parametric rates even for complex underlying spaces \mathcal{M} .

4.3 Rates in Expectation

When $p_n = 0$, the Lemma 7 used to prove Theorem 9 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 7. 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

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \right] \leq \mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \mid \mathbf{X} \in \mathcal{A} \right] + \mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^r \mathbf{1}_{\mathbf{X} \notin \mathcal{A}} \right], \quad (22)$$

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 9.

5.1 General Concentration Results

The first result is adapted from the work of Maurer and Pontil (2021), a generalization of McDiarmid’s inequality.

Definition 11 (Maurer and Pontil (2021, Definition 1)). *Let $f : \mathbb{X}^n \rightarrow \mathbb{R}$, a $\mathbf{x} \in \mathbb{X}^n$ and a random $\mathbf{X} \in \mathbb{X}^n$ with its components sampled independently. Let $\mathbf{Y} \in \mathbb{X}^n$ an independent copy of \mathbf{X} . The k -th centered conditional version of f is the random variable*

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

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 (Maurer and Pontil, 2021) 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 $\|f_k(\mathbf{X})\|_\psi : \mathbb{X}^n \rightarrow \mathbb{R}_+$ 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 \mathbf{X}, \mathbf{Y} are independent copies then $\|f_k(\mathbf{X})\|_\psi = \|f_k(\mathbf{Y})\|_\psi$ and $\|f_k(\mathbf{X})\|_\psi(\mathbf{X})$ is independent and identically distributed to $\|f_k(\mathbf{X})\|_\psi(\mathbf{Y})$.

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

Theorem 12 (McDiarmid's inequality (Maurer and Pontil, 2021, 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 $\delta > 0$,*

$$f(\mathbf{X}) - \mathbb{E}f \leq e \left(2\sigma \sqrt{\log \left(\frac{1}{\delta} \right)} + S \log \left(\frac{1}{\delta} \right) \right), \quad (23)$$

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

Proof Since this result is not exactly stated in this form, a short proof is provided in the Appendix A. ■

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

5.2 Proof of Theorem 5

The Theorem 5 can now be proved using Theorem 12.

Proof Let $m = \mathbb{E}[f(\mathbf{X}) | \mathbf{X} \in \mathcal{B}]$ and $t > 0$. 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 12 cannot be directly applied to $\mathbb{P}(f(\mathbf{X}) - m, \mathbf{X} \in \mathcal{B})$ since the components of \mathbf{X} are no longer independent conditionally on $\mathbf{X} \in \mathcal{B}$.

Similarly to (Combes, 2015), 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

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

for all $\mathbf{x} \in \mathbb{X}^n$. This extension was also used in the Kirszbraun's theorem (Kirszbraun, 1934; Hiriart-Urruty, 1980) to extend the range of Lipschitz functions.

Lemma 13. *The extension function \bar{f} satisfies*

- $\bar{f}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{B}$,
- $|\bar{f}(\mathbf{x}) - \bar{f}(\mathbf{y})| \leq B(\mathbf{x}, \mathbf{y})$ 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} ,

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

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

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

By definition of $\bar{f}(\mathbf{x}) = \inf_{\mathbf{y} \in \mathcal{B}} f(\mathbf{y}) + B(\mathbf{x}, \mathbf{y}) \leq \mathbb{E}[f(\mathbf{Y}) + B(\mathbf{x}, \mathbf{Y}) \mid \mathbf{Y} \in \mathcal{B}]$, where $\mathbf{Y} \in \mathbb{X}^n$ is independent of \mathbf{X} , so that

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

Therefore,

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

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

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

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 $p \leq 3/4$, it holds that $p^{1/2}(1-p)^{-1/2} \leq 2\sqrt{p}$, 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 12 is used on \bar{f} since

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

so that $\|\bar{f}_k(\mathbf{X})(\mathbf{x})\|_{\psi_1} \leq \|\mathbb{E}[b(X_k, X'_k) \mid X_k]\|_{\psi_1} \leq \|b\|_{\psi_1}$ by (Maurer and Pontil, 2021, 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 $S = \|b\|_{\psi_1}$. ■

Remark 14. In Theorem 5, the function b has to be a pseudometric while Theorem 12 does not assume any further structure on b . It is not clear whether Theorem 5 holds dropping the pseudometric assumption.

5.3 Proof of Lemma 7

The goal of this section is to prove the Lemma 7 providing the bound on the conditional expectation $\mathbb{E}[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A}]$.

To start this section, it is important to describe some properties of \mathcal{A} . First, the subset \mathcal{A} , as the empirical risk $\widehat{\mathcal{J}}$, 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 $I(\mathbf{X}) = \{x \in \mathbb{X} \mid (x, X_2, \dots, X_n) \in \mathcal{A}\}$. The subset \mathcal{A} and its section satisfy the following properties

$$\begin{aligned} \{\mathbf{X} \in \mathcal{A}\} \cap \{\mathbf{X}' \in \mathcal{A}\} &= \{\mathbf{X} \in \mathcal{A}\} \cap \{Y \in I(\mathbf{X})\}, \\ \{\mathbf{X} \in \mathcal{A}\} \cap \{\mathbf{X}' \notin \mathcal{A}\} &= \{\mathbf{X} \in \mathcal{A}\} \cap \{Y \notin I(\mathbf{X})\}, \end{aligned} \quad (24)$$

recalling that $\mathbf{X}' = (Y, X_2, \dots, X_n)$ where Y is an independent 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 33. 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}' (that is \mathbf{X} with one sample point swapped) fall outside \mathcal{A} is small.

Proof [Proof of Lemma 7] The proof starts with the usual decomposition for any $\widehat{\phi} \in \widehat{\mathcal{S}}$

$$\begin{aligned} \tau\vartheta(\widehat{\phi}, \mathcal{S})^\beta &\leq \mathcal{J}(\widehat{\phi}) - \mathcal{J}_* && \text{(using Assumption 2 and } \mathbf{X} \in \mathcal{A}\text{)} \\ &= \mathbb{E}_Y[l(\widehat{\phi}, Y)] + \mathcal{R}(\widehat{\phi}) + \widehat{\mathcal{J}}_* - \widehat{\mathcal{J}}_* - \mathcal{J}_* \\ &= \mathbb{E}_Y[l(\widehat{\phi}, Y)] - \widehat{\ell}(\widehat{\phi}) + \widehat{\mathcal{J}}_* - \mathcal{J}_* \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_Y[l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_i)] + \widehat{\mathcal{J}}_* - \mathcal{J}_*. \end{aligned}$$

Remark that,

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

since \mathcal{A} is invariant under permutations of the components,

The standard argument of (7) cannot be directly applied here since Y and X_1 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 X_1 and (ii) $Y \notin I(\mathbf{X})$ where Y and X_1 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

$$\begin{aligned} \tau\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\beta \mid \mathbf{X} \in \mathcal{A} \right] &\leq E_1 + E_2, \quad \text{with } E_i = \mathbb{E} \left[\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} e_i \mid \mathbf{X} \in \mathcal{A} \right], \\ \text{and } e_1 &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1) \right) \mathbf{1}_{Y \in I(\mathbf{X})} \right], \\ e_2 &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1) \right) \mathbf{1}_{Y \notin I(\mathbf{X})} \right] + \widehat{\mathcal{J}}_* - \mathcal{J}_*. \end{aligned}$$

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

$$E_1 \leq \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\sup_{\widehat{\phi} \in \widehat{\mathcal{S}}} (l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1)) \mathbb{1}_{Y \in I(\mathbf{X})} \mid \mathbf{X} \in \mathcal{A} \right] \right].$$

The next step of the proof consists in bringing out the stability of the minimizers of the empirical risk in E_1 . Let $\widetilde{\ell}$, $\widetilde{\mathcal{F}}$ and $\widetilde{\mathcal{S}}$ be the empirical risks and the minimizers associated to $\mathbf{X}' = (Y, X_1, \dots, X_n)$. 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

$$\mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\widetilde{\mathcal{F}}_* \mathbb{1}_{Y \in I(\mathbf{X})} \right] \mathbb{1}_{\mathbf{X} \in \mathcal{A}} \right] = \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_Y \left[\widehat{\mathcal{F}}_* \mathbb{1}_{Y \in I(\mathbf{X})} \right] \mathbb{1}_{\mathbf{X} \in \mathcal{A}} \right]. \quad (\text{see Lemma 30})$$

This leads to the bound

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

with an argument similar to Lemma 6. Note that the order of integration has been interchanged by Tonelli's theorem. Therefore, using Lemma 6 shows that

$$\begin{aligned} E_1 &\leq \frac{1}{1-p_n} \mathbb{E} \left[a(Y, X_1) (2\tau^{-1}n^{-1}a(Y, X_1))^{\frac{\alpha}{\beta-\alpha}} \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A}} \right] \\ &\leq \frac{(2\tau^{-1}n^{-1})^{\frac{\alpha}{\beta-\alpha}}}{1-p_n} \left\| a^{\frac{\beta}{\beta-\alpha}} \right\|_{L^2} \sqrt{\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A})}, \end{aligned} \quad (25)$$

using Cauchy-Schwartz inequality. By definition of p_n ,

$$\frac{\sqrt{\mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \in \mathcal{A})}}{1-p_n} \leq \frac{1}{\sqrt{1-p_n}} \leq 2, \quad (26)$$

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

$$\left\| a^{\frac{\beta}{\beta-\alpha}} \right\|_{L^2} = \left\| a \right\|_{L^{\frac{2\beta}{\beta-\alpha}}}^{\frac{\beta}{\beta-\alpha}} \leq \left(\frac{2\beta}{\beta-\alpha} \|a\|_{\psi_1} \right)^{\frac{\beta}{\beta-\alpha}}. \quad (27)$$

Combining (25), (26) and (27) 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_1 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 31, 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. \blacksquare

The bound on 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 paragraph, the main results of Section 4 are applied on the examples presented in the introduction.

6.1 Barycenters in Hadamard Spaces

First introduced by Fréchet (1948), Fréchet means as defined in (4) have found numerous applications in geometry (Sturm, 2003; Villani, 2003), statistics (Pennec, 2006; Pelletier, 2005), data science (Beg et al., 2005; Peyré and Cuturi, 2019; Bigot et al., 2019a) and more. Note that the behavior of the empirical counterpart has been studied in specific scenarii for example in Riemannian manifold (Bhattacharya and Patrangenaru, 2003), in Wasserstein spaces (Ahidar-Coutrix et al., 2020; Le Gouic et al., 2022) or recently in non-positively curved metric spaces (Brunel and Serres, 2024). This section is dedicated to Fréchet means defined on Hadamard spaces.

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

Definition 15. *A Hadamard space is a complete geodesic metric space (\mathbb{X}, ρ) fulfilling*

$$\rho(\phi, x)^2 - \rho(\psi, x)^2 - \rho(\phi, y)^2 + \rho(\psi, y)^2 \leq 2\rho(x, y)\rho(\phi, \psi), \quad (28)$$

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

The property (28) characterizes CAT(0)-spaces (Berg and Nikolaev, 2008, Corollary 3) (Sturm, 2003, Proposition 2.4), that is 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 was proposed by Bacák (2014).

The computation of the barycenter 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 $\mathbb{E}_{X \sim \mu}[\rho(\phi, X)^2] < +\infty$ (Sturm, 2003, Proposition 4.3). Furthermore, Sturm (2003, Proposition 4.4) showed that Assumption 2 holds with $\mathcal{F}_0 = +\infty$, $\tau = 1$ 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 16. *Assuming that $\|\rho\|_{\psi_1} < +\infty$, then this example satisfies Assumptions 1–5 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 17. *Assuming that $\|\rho\|_{\psi_1} < +\infty$ then, for any $\delta > 0$*

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

with probability at least $1 - \delta$ and

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

Similar results are obtained by Le Gouic et al. (2022) and Brunel and Serres (2024).

Remark 18. *In the Euclidean case, the quantity $\|\rho\|_{\psi_1}$ can be bounded by*

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

(see Preliminaries) 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 by Lei (2020, Proposition 5.3)

Remark 19. *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 17 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 $\widehat{\phi} = (1/n) \sum_{i=1}^n X_i$. The concentration property of sub-exponential variables gives that for any $\delta > 0$*

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

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

$$\mathbb{E} \left[\|\phi_* - \widehat{\phi}\|_2^2 \right] = \frac{1}{n} \mathbb{E} \left[\|X - \mathbb{E}[X]\|_2^2 \right] \leq \frac{16}{n} \|X\|_{\psi_1}^2 = \frac{16d}{n} \|X\|_{\psi_1}^2. \quad (29)$$

The concentration in probability and expectation obtained via Corollary 17 are therefore optimal up to universal constants.

6.2 Largest Eigenvector of Covariance Matrices

Principal components analysis (PCA) (Pearson, 1901; Hotelling, 1933) 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 $(x_i)_{i=1}^n$ sampled from the unknown probability measure μ in \mathbb{R}^d . The covariance matrix $\text{Cov}(\mu)$ can thus be estimated only from the samples represented by the discrete measure $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$. One important question concerns the stability of the subspaces of $\text{Cov}(\mu_n)$ with respect to those of $\text{Cov}(\mu)$.

The case of the leading direction of variability, which is the largest eigenvector, will be analyzed in this section.

Let \mathbb{R}^d endowed with the Euclidean distance and equipped with a centered probability measure μ . The covariance matrix of μ is defined as $\text{Cov}(\mu) = \mathbb{E}_{X \sim \mu} [XX^T]$.

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

$$\phi_* \in \arg \max_{\phi \in \mathbb{S}^{d-1}} \langle \text{Cov}(\mu) \phi, \phi \rangle. \quad (30)$$

This problem fits in the formalism of (1) with $\mathcal{M} = \mathbb{S}^{d-1}$ endowed with the great circle distance $\vartheta(\phi, \psi) = \arccos \langle \phi, \psi \rangle_{\mathbb{R}^d}$, $\mathbb{X} = \mathbb{R}^d$ and $l(\phi, x) = -\langle \phi, x \rangle^2$. 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 20. *Assuming $\|X\|_{\psi_2} < +\infty$ then this example satisfies Assumptions 1–5 with*

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

with

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

$$c = \frac{(\lambda_1 - \lambda_2)^2}{32e^2(d+1)^2\|Y\|_{\psi_2}^4 + 8e(d+1)\|Y\|_{\psi_2}^2(\lambda_1 - \lambda_2)}.$$

Proof See Appendix B.2. ■

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

Corollary 21. *Assuming $\|X\|_{\psi_2} < +\infty$ then, for any $\delta > 0$*

$$\vartheta(\widehat{\mathcal{S}}; \mathcal{S}) \leq 8L \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log \left(\frac{1}{\delta} \right)} + \frac{1}{n} \log \left(\frac{1}{\delta} \right) \right) + C_1 d^{\frac{1}{4}} e^{-\frac{c}{4}n},$$

with probability at least $1 - e^{-cn} - \delta$ and

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^2 \right] \leq 2^6 L^2 n^{-1} + C_2 d^{\frac{1}{2}} e^{-\frac{c}{2}n},$$

with the constant c defined in Lemma 20 and where

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

Remark 22. *The usual way to obtain bounds between ϕ_* and $\widehat{\phi}$ is to use matrix perturbation tools such as the Davis-Kahan Theorem (Yu et al., 2015, Theorem 2) which gives*

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

The deviation of $\text{Cov}(\mu_n)$ from $\text{Cov}(\mu)$ can then be bounded using, for example, the matrix Bernstein inequality of Lemma 34 giving, for all $t > 0$,

$$\mathbb{P}\left(\inf_{\phi_* \in \mathcal{S}} \|\widehat{\phi} - \phi_*\|_2 > t\right) \leq 2d \exp\left(\frac{-nt^2}{2(L^2 + Lt)}\right).$$

On the other hand, the bound of Corollary 21 implies that

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

for all $t > 0$, which are similar, up to universal constants.

Remark 23. 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 general framework developed in this paper also allows to cover eigenvector associated to eigenvalues with multiplicity larger than 1. 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 independently 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 $g_* : \mathbb{Y} \rightarrow \mathbb{V}$ such that $V = g_*(Y) + \epsilon$ where ϵ is the noise of the model represented by the marginal $\mu_{\mathbb{V}}$. For instance, the noiseless case $\epsilon = 0$ is obtained when $\mu_{\mathbb{V}}$ is the Dirac in $g_*(X)$ while the additive Gaussian noise one is obtained when μ is the convolution between $\mu_{\mathbb{Y}}$ and the centered Gaussian distribution.

The best predictor, in the least-square sense, is defined as the minimizer of

$$g_0 \in \arg \min_{g: \mathbb{Y} \rightarrow \mathbb{V}} \mathbb{E}_{\mu}[(V - g(Y))^2], \quad (31)$$

where g is picked in the set of measurable functions from $\mathbb{Y} \rightarrow \mathbb{V}$. When Y and V are square-integrable, the optimum g_0 , also called the Bayes predictor, is defined by $g_0(Y) = \mathbb{E}[V | Y]$. Even with this closed form, this problem is intractable for a computer due to the infinite dimension of the space of measurable functions from $\mathbb{Y} \rightarrow \mathbb{V}$.

This issue is usually circumvented by parameterizing the space of solutions as $\mathcal{H} = \{g | g(y) = \langle \phi, \theta(y) \rangle, \phi \in \mathbb{R}^m\}$ where $\phi \in \mathbb{R}^m$ and $\theta : \mathbb{Y} \rightarrow \mathbb{R}^m$ are the feature vectors and solving

$$\phi_0 \in \arg \min_{\phi \in \mathbb{R}^m} \ell(\phi) = \mathbb{E}_{\mu}[(V - \langle \theta(Y), \phi \rangle)^2] \quad (32)$$

in place of (31). The class \mathcal{H} has to be large enough to approximate well the Bayes predictor. This question is fundamental but out of the scope of this paper (see Pinkus, 1985; Tsybakov, 2009). Still, the problem (32) cannot be solved since the joint distribution

μ is not known and only samples drawn from it are accessible. A direct minimization of $\hat{\ell}$, the empirical version of (32), might however result in overfitting, so that the penalized empirical minimization problem is solved

$$\hat{\phi} \in \arg \min_{\phi \in \mathbb{R}^m} \hat{\mathcal{F}}(\phi) = \frac{1}{2n} \|\Theta\phi - V\|_2^2 + \lambda \|\phi\|_1, \quad (33)$$

where $V \in \mathbb{R}^n$ being the vector of outputs $V = (v_i)_{i=1}^n$ and $\Theta \in \mathbb{R}^{n \times m}$ is the design matrix with the vectors $(\theta(y_i)^T)_{i=1}^n$ as rows. The problem (33) is referred to as LASSO (Least Absolute Shrinkage and Selection Operator) by Tibshirani (1996). The ℓ_1 regularization further promotes the sparsity of the minimizers $\hat{\phi}$ of $\hat{\mathcal{F}}$. It has been thoroughly studied in relation to the problem of sparse signal recovery (see for example Candes et al., 2006; Donoho, 2006; Wainwright, 2019, for a comprehensive treatment and bibliographic details). In particular, a fundamental question in sparse recovery related to the content of this paper is to understand the convergence of the minimizers of the empirical LASSO (33) towards those of the approximated Bayes predictor (32). High probability recovery guarantees are for example obtained assuming that ϕ_0 is sparse and that the design matrix Θ satisfies the so-called restricted eigenvalue condition (Raskutti et al., 2010; Rudelson and Zhou, 2012; Wainwright, 2019). This condition imply a quadratic growth of $\hat{\ell}$ around its minimizers in directions that are ‘‘sufficiently sparse’’. Koltchinskii (2009) carried another analysis on the excess risk $\ell(\hat{\phi}) - \ell(\phi_0)$ by introducing the intermediate problem

$$\phi_* \in \arg \min_{\phi \in \mathbb{R}^m} \mathcal{F}(\phi) = \ell(\phi) + \lambda \|\phi\|_1, \quad (34)$$

and by separately studying the bias $\ell(\phi_*) - \ell(\phi_0)$ and the random error $\ell(\hat{\phi}) - \ell(\phi_*)$. Finally, the excess risk between $\ell(\hat{\phi}) - \ell(\phi_0)$ is obtained by balancing both errors. While the bias part is bounded assuming some sparsity of the solution ϕ_0 and a local strong convexity of l , the random error is bounded using empirical process theory. Note that asymptotic convergence between (33) and (34) has also been studied by Knight and Fu (2000).

In this section, the random error between $\hat{\mathcal{S}}$ and \mathcal{S} is studied using the formalism of (1) with $\mathbb{X} = \mathbb{Y} \times \mathbb{V}$, $l : \mathbb{X} \rightarrow \mathbb{R}_+$ defined by $l(\phi, (y, v)) = \frac{1}{2}(\langle \phi, \theta(y) \rangle - v)^2$, $\mathcal{M} = \mathbb{R}^m$ endowed with the Euclidean distance and $\mathcal{R} = \lambda \|\cdot\|_1$.

The next result shows that all the assumptions required to apply Lemma 7 and Theorem 9 are satisfied. Proving the local error bound for \mathcal{F} and $\hat{\mathcal{F}}$ is the most challenging part. Since the function \mathcal{F} and $\hat{\mathcal{F}}$ are convex piecewise polynomials of degree 2, they both satisfy a local error bound of Assumption 2 with $\beta = 2$ (Bolte et al., 2017, Corollary 9) for any $\mathcal{J}_0 > \mathcal{J}_*$ with an adequate $\tau(\mathcal{J}_0)$. This powerful result however does not specify the values of the constants \mathcal{J}_0 and τ which are important to derive proper bounds and to check Assumption 4.

The following analysis will rely on the polyhedral geometry of the ℓ_1 ball for which error bounds analyses are available through the Hoffman’s lemma (Hoffman, 1952; Güler, 2010).

Definition 24 (Hoffman’s error bound). *Let P be a polyhedron defined by $P = \{x \in \mathbb{R}^q : Ax \leq a\}$ for some $A \in \mathbb{R}^{p \times q}$ and $a \in \mathbb{R}^p$. Let $O = \{x \in \mathbb{R}^q : Ex = e\}$ for some $E \in \mathbb{R}^{r \times q}$ and $e \in \mathbb{R}^r$. Assume that $O \cap P \neq \emptyset$. Then there exists a constant $H \geq 0$, depending only on the pair (O, P) such that*

$$\inf_{y \in O \cap P} \|x - y\| \leq H \|Ex - e\|, \quad \forall x \in P.$$

Alternatively, the Hoffman constant of the pair (O, P) can be written as

$$H = \max_{\substack{I \subseteq \{1, \dots, r+p\} \\ C_I \text{ full row rank}}} \frac{1}{\sigma_{|I|}(C_I)}, \quad (35)$$

where $C_I \in \mathbb{R}^{|I| \times \ell}$ refers to the submatrix built from the rows of indexes in I of the matrix $C = [E^T, A^T]^T \in \mathbb{R}^{(r+p) \times q}$ (Pena et al., 2021).

The polyhedral geometry of the ℓ_1 ball is finally exploited using the following result and will be key to prove Assumptions 2 and 4.

Lemma 25 (Bolte et al. (2017, 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 minimizers ϕ_* of \mathcal{L} satisfy $\|\phi_*\|_1 \leq R_*$. Fix $R > R_*$, then*

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

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

$$\tau = (4H^2 [1 + \lambda R + (R\|A\| + \|b\|)(4R\|A\| + \|b\|)])^{-1}, \quad (36)$$

and $H = H(A, \lambda) > 0$ is the Hoffman's constant (35) of the matrix

$$C(A, \lambda) = \begin{pmatrix} E & -1_{\mathbb{R}^{2^q \times 1}} \\ 0_{\mathbb{R}^{1 \times q}} & 1 \\ A & 0_{\mathbb{R}^{p \times 1}} \\ 0_{\mathbb{R}^{1 \times q}} & \lambda \end{pmatrix} \in \mathbb{R}^{(2^q+1+p+1) \times (q+1)}, \quad (37)$$

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

Lemma 26. *Assuming that $\|\theta(Y)\|_{\psi_2} < +\infty$ and $\|V\|_{\psi_2} < +\infty$ then this example satisfies Assumptions 1–5 with*

\mathcal{J}_0	β	τ	α	a	$\ a\ _{\psi_1}$	η	κ	ι
$\mathbb{E}[V^2]$	2	(38)	1	(49)	(39)	$2e^{-c_1 n}$	$2me^{-c_2 n}$	$2e^{-c_1 n}$

where the constants are defined with $A = \mathbb{E}[\theta(Y)\theta(Y)^T]^{\frac{1}{2}} \in \mathbb{R}^{m \times m}$ as

$$\tau = \left(4H^2 \left[1 + \mathcal{J}_0 + \left(\mathcal{J}_0 \|A\| \lambda^{-1} + \sqrt{\mathcal{J}_0} \right) \left(4\mathcal{J}_0 \|A\| \lambda^{-1} + \sqrt{\mathcal{J}_0} \right) \right] \right)^{-1}, \quad (38)$$

with H being the Hoffman constant of the matrix $C(A, \lambda)$ defined in (37) and

$$\|a\|_{\psi_1} = 4Rm \|\theta(Y)\|_{\psi_2}^2 + 4\sqrt{m} \|\theta(Y)\|_{\psi_2} \|V\|_{\psi_2}. \quad (39)$$

The constants c_1 and c_2 are defined by

$$\begin{aligned} c_1 &= \frac{\|V\|_2^4}{2(e^2 \|V\|_{\psi_2}^4 + e \|V\|_{\psi_2}^2 \|V\|_2^2)}, \\ c_2 &= \frac{c^2}{2(e^2(m+1)^2 \|\theta(Y)\|_{\psi_2}^4 + e(m+1) \|\theta(Y)\|_{\psi_2}^2 c)}, \\ c &= (\sqrt{2} - 1) \min \left(\frac{1}{\sqrt{2}H^2}, \|A\|^2 \right). \end{aligned}$$

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

Corollary 27. *Assuming that $\|\theta(Y)\|_{\psi_2} < +\infty$ and $\|V\|_{\psi_2} < +\infty$ then, for any $\delta > 0$*

$$\vartheta(\widehat{\mathcal{S}}; \mathcal{S}) \leq 8L \left(n^{-\frac{1}{2}} + 2\sqrt{\frac{1}{n} \log\left(\frac{1}{\delta}\right)} + \frac{1}{n} \log\left(\frac{1}{\delta}\right) \right) + 14Lp_n^{\frac{1}{4}},$$

with probability at least $1 - p_n - \delta$ and

$$\mathbb{E} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^2 \right] \leq L^2 2^6 n^{-1} + C_2 p_n^{-\frac{1}{2}},$$

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 Appendix B.3. ■

Koltchinskii (2009) obtained bounds either on the excess risk or on the ℓ_2 distance between the functions $\langle \theta(\cdot), \phi_* \rangle$ and $\langle \theta(\cdot), \widehat{\phi} \rangle$. The result of Corollary 27 shows that concentration bounds on the estimation error between (34) and (33) in terms of distance between their minimizers are also accessible. This additional granularity needed a local error bound that is proved via an intricate analysis of the geometry of the LASSO problem. The recent result of Lemma 25 was leveraged to obtain a local error bound involving the Hoffman's constant of a matrix defined by the regularization parameter and the covariance matrix of the feature vectors (37). The most challenging part of Lemma 26 was to prove the Assumption 4, that is the concentration of $\widehat{\tau}$ towards τ . This step involves the concentration of the Hoffman constant (35) of the matrix $C \left(\left(\frac{1}{n} \sum_{i=1}^n \theta(y_i) \theta(y_i)^T \right)^{\frac{1}{2}}, \lambda \right)$ towards the one of $C \left(\mathbb{E} [\theta(Y) \theta(Y)^T]^{\frac{1}{2}}, \lambda \right)$ (37). These technicalities essentially boil down to the concentration of the sample covariance matrix to its exact counterpart which can be established with tools such as Theorem 35.

Even though the present result sheds new light on the concentration of the solutions of the empirical LASSO problem, it still needs to be connected to the underlying problem of finding bounds on the distance between ϕ_0 and $\widehat{\phi}$. Answering this type of question requires the introduction of hypotheses on the support of the solutions ϕ_0 and to incorporate this support in the definition of the Hoffman constant. Furthermore, the dependency of the Hoffman constants of $C(A, \lambda)$ with respect to λ has to be properly understood in order to balance the bias and variance terms. From the author's perspective, there is still a substantial amount of work to accurately address the question of sparse recovery using Corollary 27.

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 the proposed analysis that cannot be applied verbatim. Deeper investigations

are thus 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 in good textbooks (such as Villani, 2003; Ambrosio et al., 2008; Ambrosio and Gigli, 2013).

6.4.1 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 $\mathbb{E}_{X \sim \phi}[h(X, y)^2] < +\infty$ 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

$$W_2^2(\phi, \psi) = \inf_{\pi \in \Pi_{\phi, \psi}} \mathbb{E}_{(X, Y) \sim \pi}[h(X, Y)^2]. \quad (40)$$

This section focuses on the case where H is a Hilbert space endowed with the distance h associated to the inner product, that is the spaces $(\mathcal{M}, \vartheta) = (\mathbb{X}, \rho) = (\mathcal{P}_2(H), W_2)$. The Wasserstein space \mathcal{M} is known to be geodesic and with non-negative curvature (see Ambrosio et al., 2008, Section 7.3). Let μ be a probability measure on $\mathcal{P}_2(\mathcal{M})$, its barycenters are defined as the Fréchet means

$$\phi_* \in \arg \min_{\phi \in \mathcal{P}_2(\mathcal{M})} \mathcal{F}(\phi) = \int W_2^2(\phi, \psi) d\mu(\psi), \quad (41)$$

that is 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 (see Agueh and Carlier, 2011; Le Gouic and Loubes, 2017), for example when (H, h) is a separable locally compact geodesic space.

6.4.2 THE DIFFICULTIES OF NON-NEGATIVE CURVATURE

As observed by Ambrosio et al. (2008, Theorem 7.3.2), the barycenter function (41) and its empirical counterpart are not strongly convex. This is due to the positive curvature of $(\mathcal{P}_2(H), W_2)$, 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 $\mathcal{J} = +\infty$ 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 $\mathcal{J} = +\infty$ 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 hold, for all measures $\mu \in \mathcal{P}_2(\mathcal{M})$, since this would imply that \mathcal{M} has non-positive curvature (Sturm, 2003, 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 can be unbounded (Stromme, 2020, 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} .

6.4.3 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 Ahidar-Coutrix et al. (2020); Le Gouic et al. (2022) 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 (see Le Gouic et al., 2022, Theorem 7). The distance induced by the metric will be denoted $\|\cdot\|_{\phi_*}$. Second they show (in their 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), that is $\phi = \nabla \zeta_{\phi_* \rightarrow \phi} \# \phi_*$, the function \mathcal{F} satisfies the global error bound of Assumption 2 with $\beta = 2$, $\tau = \frac{1}{\beta - \alpha - 1}$ and $\mathcal{F}_0 = +\infty$. 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 by Bigot et al. (2019b); Kroshnin (2018) to seek a regularized barycenter by adding an entropic regularization to the Wasserstein variance function. Bigot et al. (2019b) proved the existence and uniqueness of the regularized barycenters as well as convergence of their empirical counterpart. Recently, Carlier et al. (2021) derived regularity properties of the entropic-barycenters and established central limit theorems.

Let Ω be a bounded subset of \mathbb{R}^d , of positive Lebesgue measure and endowed with the Euclidean distance. Define $\mathcal{M} = \mathcal{P}_2^{ac}(\Omega)$ 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 up to a factor $\frac{1}{2}$ with the ℓ^1 norm on $\mathcal{P}_2^{ac}(\Omega)$, that is $\vartheta = \frac{1}{2} \|\cdot - \cdot\|_1$. 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 by Bigot et al. (2019b), 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

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

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 Carlier et al., 2021, Proposition 2.1).

While the Wasserstein distance squared is not geodesically convex (see Section 6.4), it is convex with respect to the linear structure in $\mathcal{P}_2(\Omega)$, that is 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 Proposition 7.17 Santambrogio, 2015). This leads to convex functions ℓ and $\widehat{\ell}$. 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 $\widehat{\mathcal{F}}$ so that Assumptions 2 and 4 hold.

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 28. *Assuming that Ω is bounded and is of positive Lebesgue measure, this example satisfies Assumptions 1–5 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 Appendix B.4. ■

Corollary 29. *Under the assumption of Lemma 28, for any $\delta > 0$*

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

with probability at least $1 - \delta$ and

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

As a comparison, Bigot et al. (2019b) obtained a bound of kind $\mathbb{E}[\text{KL}(\widehat{\phi}|\phi_*)^2] \leq cn^{-1}$, 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 29 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, 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 9 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 9.

Even though 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 Lojasiewicz constant in Assumption 4. Additionally to the technical questions raised in the core of the paper, Theorem 9 leaves aside the case $\beta = \alpha$ which might be of great interest for applications where the cost function l has, for example, a Lipschitz continuous gradient.

Acknowledgments and Disclosure of Funding

The author would like to thank the anonymous reviewers and the editor, Joseph Salmon, for their time and effort in handling the paper. Their careful review and constructive comments have greatly improved the quality of this manuscript.

Appendix A. Additional Proofs

A.1 General Concentration Results

Proof [Proof of Theorem 12] The Theorem 4 of Maurer and Pontil (2021) is not exactly stated in the form of (23). More precisely, it states that

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

for all $t \geq 0$, which is slightly weaker than (23). However, a closer inspection of the proofs shows that

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

for $t \leq (eS^{-1})$. This shows that $f(\mathbf{X}) \in \text{sub}\Gamma_+(2e^2\sigma^2, eS)$ which implies the tail (23), see Section 2. ■

A.2 Supporting Results for the Proof of Lemma 7

Lemma 30. *In the setting of the proof of Lemma 7,*

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

Proof This result will be clear, as long as the order of integration can be exchanged, because $\mathbb{1}_{\mathbf{X} \in \mathcal{S}}$ and $\mathbb{1}_{\mathbf{X}' \in \mathcal{S}}$ have the same distribution. This will be achieved checking that

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

The quantity is indeed bounded by

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

The third term is finite since, by picking any $\phi_* \in \mathcal{S}$,

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

together with the assumption that $\mathbb{E}_X [|l(\phi_*, X)|] < +\infty$ and the fact that $\mathcal{R}(\phi_*) < +\infty$ since ϕ_* is a minimizer of \mathcal{J} . 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

$$\begin{aligned} \mathbb{E}_{\mathbf{X}, Y} \left[\left(\inf_{\phi_* \in \mathcal{S}} \hat{\mathcal{J}}(\phi_*) - \hat{\mathcal{J}}_* \right) \mathbb{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{S}} \right] & \leq \mathbb{E}_{\mathbf{X}, Y} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{E}_Z [a(Z, X_i)] \vartheta(\hat{\mathcal{S}}; \mathcal{S})^\alpha \mathbb{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{S}} \right] \\ & = \mathbb{E}_{\mathbf{X}, Y} \left[\mathbb{E}_Z [a(Z, X_i)] \vartheta(\hat{\mathcal{S}}; \mathcal{S})^\alpha \mathbb{1}_{\mathbf{X}, \mathbf{X}' \in \mathcal{S}} \right] \\ & \leq \|a\|_{L^2} \|\vartheta(\hat{\mathcal{S}}; \mathcal{S})^\alpha \mathbb{1}_{\mathbf{X} \in \mathcal{S}}\|_{L^2}, \end{aligned}$$

by an application of Cauchy-Schwartz inequality. The latter quantity is finite since $\|a\|_{L^2} \leq 2\|a\|_{\psi_1}$ and using the bounds on the moments of $\vartheta(\hat{\mathcal{S}}; \mathcal{S})$ of Lemma 32.

Therefore, in virtue of Fubini's theorem,

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

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

Lemma 31. *The term E_2 from the proof of Lemma 7 can be bounded by*

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

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

$$\begin{aligned} e_2 &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1) \right) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] + \widehat{\mathcal{J}}_* - \mathcal{J}_* \\ &= e_{2,1} + e_{2,2}, \end{aligned}$$

with

$$\begin{aligned} e_{2,1} &= \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - l(\widehat{\phi}, X_1) - l(\phi_*, Y) + l(\phi_*, X_1) \right) \mathbb{1}_{Y \notin I(\mathbf{X})} \right], \\ e_{2,2} &= \mathbb{E}_Y \left[\left(l(\phi_*, Y) - l(\phi_*, X_1) \right) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] + \widehat{\mathcal{J}}_* - \mathcal{J}_*. \end{aligned}$$

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

$$\begin{aligned} e_{2,1} &\leq \mathbb{E}_Y \left[\left(l(\widehat{\phi}, Y) - l(\phi_*, Y) + l(\phi_*, X_i) - l(\widehat{\phi}, X_i) \right) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] \\ &\leq \mathbb{E}_Y \left[a(Y, X_1) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] \vartheta(\widehat{\phi}, \phi_*)^\alpha \quad (\text{Assumption 5}) \\ &\leq \mathbb{E}_Y \left[a(Y, X_1) \mathbb{1}_{Y \notin I(\mathbf{X})} \right] 2^{\max(0, \alpha-1)} \left(\vartheta(\widehat{\phi}, \mathcal{S})^\alpha + \text{diam}(\mathcal{S})^\alpha \right). \end{aligned}$$

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

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

using Tonelli's theorem to swap the order of integration and (24). The term

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

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

$$\begin{aligned} \mathbb{E}_{\mathbf{X}, Y} \left[\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha a(Y, X_1) \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}} \right] &\leq \|a\|_{L^4} \|\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^\alpha \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}}\|_{L^4} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \\ &\leq \|a\|_{L^4} \|\vartheta(\widehat{\mathcal{S}}; \mathcal{S})^{\beta-\alpha} \mathbb{1}_{\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A}}\|_{L^{\frac{\beta-\alpha}{\max(\frac{4\alpha}{\beta-\alpha}, 1)}}}^{\frac{\alpha}{\beta-\alpha}} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \\ &\leq \|a\|_{L^4} \left(\tau^{-1} \|a\|_{L^{\max(\frac{4\alpha}{\beta-\alpha}, 1)}} \right)^{\frac{\alpha}{\beta-\alpha}} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2} \quad (\text{Lemma 32}) \\ &\leq 4\|a\|_{\psi_1} \left(\tau^{-1} \max\left(\frac{4\alpha}{\beta-\alpha}, 1\right) \|a\|_{\psi_1} \right)^{\frac{\alpha}{\beta-\alpha}} \mathbb{P}(\mathbf{X} \in \mathcal{A}, \mathbf{X}' \notin \mathcal{A})^{1/2}. \end{aligned}$$

On the other hand, the Lemma 33 allows to show that

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

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

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

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

$$e_{2,2} \leq \mathbb{E}_Y [(l(\phi_*, Y) - l(\phi_*, X_1)) \mathbf{1}_{Y \notin I(\mathbf{X})}] + \widehat{\mathcal{J}}(\phi_*) - \mathcal{J}_*.$$

Since \mathcal{A} is invariant under permutations of the components of \mathbf{X} ,

$$\mathbb{E} [\widehat{\mathcal{J}}(\phi_*) | \mathbf{X} \in \mathcal{A}] = \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\widehat{\mathcal{J}}(\phi_*) | \mathbf{X} \in \mathcal{A}] = \mathbb{E} [l(\phi_*, X_1) | \mathbf{X} \in \mathcal{A}].$$

Therefore,

$$E_{2,2} \leq \frac{1}{1 - p_n} \mathbb{E}_{\mathbf{X}} [\mathbb{E}_Y [(l(\phi_*, X_1) - l(\phi_*, Y)) \mathbf{1}_{Y \in I(\mathbf{X})}] \mathbf{1}_{\mathbf{X} \in \mathcal{A}}].$$

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

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

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. ■

Lemma 32. *Let \mathcal{A} be the subset defined by (18) and assume Assumptions 1–5 hold, then conditionally on \mathcal{A}*

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

Furthermore, for all $p \geq 1$

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

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

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

Therefore,

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

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

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

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

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

Lemma 33. *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 independent 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 $g(\mathbf{Z}) = \mathbb{E}[\mathbf{1}_{\mathbf{X}' \notin \mathcal{A}} | \mathbf{Z}]$. Then

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

by the FKG inequality (Fortuin et al., 1971). Therefore $\text{Cov}(\mathbf{1}_{\mathbf{X} \in \mathcal{A}}, \mathbf{1}_{\mathbf{X}' \notin \mathcal{A}}) \leq 0$, which gives the desired result. \blacksquare

Appendix B. Additional Proofs for Applications

B.1 Concentration of Sample Covariance Matrices

Lemma 34. *Let $Z \in \mathbb{R}^d$ be a sub-Gaussian random vector and let $(Z_i)_{i=1}^n$ be n independent copies of Z . Let $A = \mathbb{E}[ZZ^T]$ and $\hat{A} = \frac{1}{n} \sum_{i=1}^n Z_i Z_i^T$ then*

$$\mathbb{P}\left(\|A - \hat{A}\| > t\right) \leq 2d \exp\left(-\frac{nt^2}{2\left(e^2(d+1)^2\|Z\|_{\psi_2}^4 + e(d+1)\|Z\|_{\psi_2}^2 t\right)}\right).$$

Proof This result is a consequence of matrix Bernstein inequalities (Tropp, 2012, Theorem 6.2) stated in the next theorem.

Theorem 35. *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 that 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 $\sigma^2 = \|\sum_{i=1}^n B_i^2\|$ and where λ_{\max} stands for the algebraically largest eigenvalue.

This result will be used by setting $S_i = \frac{1}{n}(Z_i Z_i^T - \mathbb{E}[ZZ^T])$. 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 $\hat{A} - A$ is equal to its largest magnitude eigenvalue that is $\|\hat{A} - A\| = \max(-\lambda_{\min}(\hat{A} - A), \lambda_{\max}(\hat{A} - A))$ (here λ_{\min} stands for the algebraically smallest 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

$$\begin{aligned} \|S_i\| &= \sup_{v \in \mathbb{S}^{d-1}} |\langle S_i v, v \rangle| = \frac{1}{n} \sup_{v \in \mathbb{S}^{d-1}} |\langle Z_i, v \rangle^2 - \mathbb{E}[\langle Z, v \rangle^2]| \\ &\leq \frac{1}{n} \max\left(\|Z_i\|_2^2, \sup_{v \in \mathbb{S}^{d-1}} \mathbb{E}[\langle Z, v \rangle^2]\right) \\ &\leq \frac{1}{n} \max(\|Z_i\|_2^2, \|Z\|_{\psi_2}^2), \end{aligned}$$

since $\mathbb{E}[\langle Z, v \rangle^2] \leq \|\langle Z, v \rangle\|_{\psi_1} = \|\langle Z, v \rangle\|_{\psi_2}^2$ and using the definition of the sub-Gaussian norm of vectors. This implies that

$$\begin{aligned} \|\|S_i\|\|_{\psi_1} &\leq \frac{1}{n} (\|\|Z_i\|_2\|_{\psi_1}^2 + \|Z\|_{\psi_2}^2) \\ &= \frac{1}{n} (\|\|Z_i\|_2\|_{\psi_2}^2 + \|Z\|_{\psi_2}^2) \\ &\leq \frac{d+1}{n} \|Z\|_{\psi_2}^2, \end{aligned}$$

leading to $\mathbb{E}[\|S_i\|^p] \leq \frac{p!}{2} \left(e^{\frac{d+1}{n}} \|Z\|_{\psi_2}^2 \right)^p$. The Theorem 35 can now be applied to $\lambda_{\max}(\widehat{A} - A)$ and $-\lambda_{\min}(\widehat{A} - A) = \lambda_{\max}(A - \widehat{A})$ with $R = e^{\frac{d+1}{n}} \|Z_i\|_{\psi_2}^2$ and $B_i = e^{\frac{d+1}{n}} \|Z_i\|_{\psi_2}^2 \text{Id}_d$ with the corresponding $\sigma^2 = \frac{1}{n} e^2 (d+1)^2 \|Z_i\|_{\psi_2}^4$. Applying the union bound to these estimates completes the proof. \blacksquare

B.2 Largest Eigenvector of Covariance Matrices

Proof [Proof of Lemma 20] Prior to the proof, general facts on \mathbb{S}^{d-1} as a Riemannian manifold are introduced.

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 $\mathbf{g}_\phi(u, v) = \langle u, v \rangle_{\mathbb{R}^d}$, the Euclidean scalar product, for all $u, v \in \mathcal{T}_\phi \mathbb{S}^{d-1}$. The norm of vectors in $\mathcal{T}_\phi \mathbb{S}^{d-1}$ will be denoted by $\|\cdot\|_\phi = \mathbf{g}_\phi(\cdot, \cdot)$. In particular, here $\|\cdot\|_\phi = \|\cdot\|_2$ for all $\phi \in \mathbb{S}^{d-1}$ so that the dependency on ϕ might be dropped. 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 $\text{Cov}(\mu)$.

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, so that $\vartheta(\psi, \mathcal{S}) \leq \pi/2$ 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 $\|\xi\|_2 = \|\xi\|_\phi = \vartheta(\psi, \phi) \leq \frac{\pi}{2}$. The bound is obtained using

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

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

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

The last inequality is obtained using Lemma 36 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 $\mathbb{P}(\mathcal{F}(\hat{\phi}) > \mathcal{F}_0) = 0$.

Assumption 4 Let $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_m$ be the eigenvalues of $\hat{A} = \frac{1}{n} \sum_{i=1}^n X_i X_i^T$. The derivation of Assumption 2 can also be carried on $\hat{\mathcal{F}}$ so that

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

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

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

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

When $2\|A - \hat{A}\| \leq \frac{1}{2}(\lambda_1 - \lambda_2)$ the eigengap $\hat{\lambda}_1 - \hat{\lambda}_2 \geq \frac{1}{2}(\lambda_1 - \lambda_2)$. This event happens with high probability since by Lemma 34:

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

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

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

using $(a^2 - b^2) = (a - b)(a + b)$ for the first equality and where 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

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

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

$$\|a\|_{\psi_1} \leq 4\|X\|_{\psi_1}^2 = 4\|X\|_{\psi_2}^2.$$

Finally, using the fact that the sub-Gaussian norm $\|X\|_{\psi_2} \leq \sqrt{d}\|X\|_{\psi_2}$, 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 36. *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 Since v is orthogonal to u_1 ,

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

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 [Proof of Lemma 26]

The proof consists in checking the assumptions one by one.

Assumption 1 Since the functions \mathcal{F} and $\widehat{\mathcal{F}}$ are convex and coercive, each of them admits at least one minimizer.

Assumption 2 In order to apply Lemma 25, the function \mathcal{F} has to be written in the form $\mathcal{F}(\phi) = \mathcal{L}(\phi) + c$ where \mathcal{L} is of the form $\mathcal{L}(\phi) = \frac{1}{2} \|A\phi - b\|_2^2 + \lambda \|\phi\|_1$ and c is a constant term. Reminding that $A = \mathbb{E} [\theta(Y)\theta(Y)^T]^{\frac{1}{2}}$, it can be shown that

$$\begin{aligned} \ell(\phi) &= \frac{1}{2} \langle A\phi, A\phi \rangle - \langle \phi, \mathbb{E}_\mu[\theta(Y)V] \rangle + \frac{1}{2} \mathbb{E}[V^2] \\ &= \frac{1}{2} \langle A(\phi - \phi_0), A(\phi - \phi_0) \rangle + \frac{1}{2} \mathbb{E}[V^2] - \frac{1}{2} \|A\phi_0\|_2^2, \end{aligned}$$

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

$$\min_{\phi \in \mathbb{R}^m} \ell(\phi). \quad (42)$$

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

$$\mathcal{L}(\phi) = \frac{1}{2} \|A\phi - b\|_2^2 + \lambda \|\phi\|_1,$$

setting $b = A\phi_0$. Remark at this point that since ϕ_0 is the solution of (42)

$$0 \leq \ell(\phi_0) = \frac{1}{2} \mathbb{E}[V^2] - \frac{1}{2} \|A\phi_0\|_2^2,$$

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

Finally, observe that $\lambda \|\phi_*\|_1 \leq \mathcal{F}(\phi_*) \leq \mathcal{F}(0) = \ell(0) = \frac{1}{2} \mathbb{E}[V^2]$. Thus, setting $\mathcal{J}_0 = \mathbb{E}[V^2]$ and $R = \mathbb{E}[V^2]/\lambda$ ensures that $\mathcal{F}(\phi) \leq \mathcal{J}_0$ implies that $\|\phi\|_1 \leq R$.

The Lemma 25 can now be applied with these matrix A and value R leading to a bound

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

as long as $\|\phi\|_1 \leq R$ for τ with expression given in (36). Using furthermore the fact that $\|b\|_2^2 \leq \mathbb{E}[V^2]$,

$$\begin{aligned} \tau &\geq \left(4H^2 \left[1 + \lambda R + \left(R\|A\| + \sqrt{\mathbb{E}[V^2]} \right) \left(4R\|A\| + \sqrt{\mathbb{E}[V^2]} \right) \right] \right)^{-1} \\ &= \left(4H^2 \left[1 + \lambda R + (R\|A\| + \sqrt{\lambda R})(4R\|A\| + \sqrt{\lambda R}) \right] \right)^{-1}. \end{aligned} \quad (43)$$

This lower-bound will be used in place of τ . This choice is not optimal, since it was used that $\|b\|_2^2 \leq \mathbb{E}[V^2]$, but it will ease the proof of Assumption 4.

Assumption 3 Similarly to the proof of Assumption 2, $\widehat{\mathcal{F}}(\widehat{\phi}) \leq \widehat{\mathcal{F}}(0) = \frac{1}{2} \frac{1}{n} \sum_{i=1}^n V_i^2$ so that $\mathcal{F}(\widehat{\phi}) \leq \mathcal{J}_0$ as soon as $\frac{1}{n} \sum_{i=1}^n V_i^2 \leq 2\mathbb{E}[V^2]$. This event happens with high probability since using the Bernstein inequality of Theorem 12,

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n V_i^2 - \mathbb{E}[V^2] \right| > t \right) \leq 2 \exp \left(- \frac{nt^2}{2(e^2 \|V^2\|_{\psi_1}^2 + e \|V^2\|_{\psi_1} t)} \right), \quad (44)$$

for all $t > 0$, because the random vector V^2 is sub-exponential with $\|V^2\|_{\psi_1} = \|V\|_{\psi_2}^2$. Choosing $t = \mathbb{E}[V^2]$ 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

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

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

In this proof, the matrix A will also refer to $\mathbb{E}[\theta(Y)\theta(Y)^T]^{\frac{1}{2}}$ and \hat{A} to its empirical counterpart.

Note that $\tilde{\tau}$ is defined in a similar fashion as (43) by

$$\tilde{\tau} = \left(4\hat{H}^2 \left[1 + \mathcal{F}_0 + \left(\mathcal{F}_0\|\hat{A}\|\lambda^{-1} + \sqrt{\mathcal{F}_0}\right) \left(4\mathcal{F}_0\|\hat{A}\|\lambda^{-1} + \sqrt{\mathcal{F}_0}\right)\right]\right)^{-1}, \quad (45)$$

so that the high-probability bound will be obtained by proving that $\hat{H}^2 \leq 2^{1/2}H^2$ and $\|\hat{A}\| \leq 2^{1/4}\|A\|$ with high probability. If all these inequalities hold, then

$$\begin{aligned} \tilde{\tau} &\geq \left(4\sqrt{2}H^2 \left[1 + \mathcal{F}_0 + \left(2^{1/4}\mathcal{F}_0\|A\|\lambda^{-1} + \sqrt{\mathcal{F}_0}\right) \left(2^{1/4}4\sqrt{\mathcal{F}_0}\|A\|\lambda^{-1} + \sqrt{\mathcal{F}_0}\right)\right]\right)^{-1} \\ &\geq \left(4\sqrt{2}H^2 \left[\sqrt{2}(1 + \mathcal{F}_0) + \sqrt{2} \left(\mathcal{F}_0\|A\|\lambda^{-1} + \sqrt{\mathcal{F}_0}\right) \left(4\mathcal{F}_0\|A\|\lambda^{-1} + \sqrt{\mathcal{F}_0}\right)\right]\right)^{-1} \\ &\geq \frac{1}{2}\tau. \end{aligned}$$

- Proving that $\hat{H}^2 \leq 2^{1/2}H^2$. Using definition of H , the Hoffman's constant of the matrix $C(A, \lambda)$ as defined in (37), and under the assumption that it is positive, it can be proved that

$$\lambda_{|I|}(C(A, \lambda)_I^T C(A, \lambda)_I) = \sigma_{|I|}^2(C(A, \lambda)_I) \geq \frac{1}{H^2}, \quad (46)$$

for all subsets 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 ,

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

by the interlacing property of eigenvalues. Further computations show that

$$r_{I,i} \leq \left\| \begin{pmatrix} AA - \hat{A}\hat{A} & 0 \\ 0 & 0 \end{pmatrix} \right\| = \left\| \mathbb{E}[\theta(Y)\theta(Y)^T] - \frac{1}{n} \sum_{i=1}^n \theta(y_i)\theta(y_i)^T \right\|.$$

In particular, under the event $\|AA - \hat{A}\hat{A}\| \leq \frac{1}{\epsilon H^2}$,

$$\begin{aligned} \lambda_{|I|}(C(\hat{A}, \lambda)_I^T C(\hat{A}, \lambda)_I) &\geq \lambda_{|I|}(C(A, \lambda)_I^T C(A, \lambda)_I) - \frac{1}{\epsilon H} \\ &\geq \left(1 - \frac{1}{\epsilon}\right) \lambda_{|I|}(C(A, \lambda)_I^T C(A, \lambda)_I), \end{aligned} \quad (47)$$

using (46). This implies that under this event, the submatrices $C(\hat{A}, \lambda)_I$ share the row rank of their counterpart $C(A, \lambda)_I$. Therefore, the set of constraints in the maximization problem (35) are equal. The equation (47) furthermore implies that $\hat{H}^2 \leq \frac{\epsilon}{\epsilon-1}H^2$. 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 $\|\widehat{A}\widehat{A}\| \leq 2^{1/2}\|AA\|$ which will also be proved using the Weyl's inequality as

$$\left| \|AA\| - \|\widehat{A}\widehat{A}\| \right| \leq \|AA - \widehat{A}\widehat{A}\|.$$

The expected bound happens under the event $\|AA - \widehat{A}\widehat{A}\| \leq (2^{1/2} - 1)\|AA\|$.

Finally, these computations show that

$$\begin{aligned} \mathbb{P}\left(\tilde{\tau} < \frac{1}{2}\tau\right) &\leq \mathbb{P}\left(\left\|\mathbb{E}[\theta(Y)\theta(Y)^T] - \frac{1}{n}\sum_{i=1}^n \theta(y_i)\theta(y_i)^T\right\| \geq c\right) \\ &\leq 2m \exp\left(-n \frac{c^2}{2(e^2(m+1)^2\|\theta(Y)\|_{\psi_2}^4 + e(m+1)\|\theta(Y)\|_{\psi_2}^2 c)}\right), \end{aligned} \quad (48)$$

with $c = \min\left(\frac{\sqrt{2}-1}{\sqrt{2}}\frac{1}{H^2}, (\sqrt{2}-1)\|AA\|\right)$ using Lemma 34. The assumption is therefore satisfied with $\kappa(n) = 2me^{-c_2n}$ with c_2 defined in (48).

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 $R = \mathbb{E}[V^2]/\lambda$ 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 $(a^2 - b^2) = (a - b)(a + b)$ and the Cauchy-Schwartz inequality. The function a defined by

$$a(x, x') = (\|\theta(y)\|_2 (2R\|\theta(y)\|_2 + 2|v|) + \|\theta(y')\|_2 (2R\|\theta(y')\|_2 + 2|v'|)) \mathbb{1}_{x \neq x'}, \quad (49)$$

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

$$\begin{aligned} \|a\|_{\psi_1} &\leq 4 \|\|\theta(Y)\|_2 (2R\|\theta(Y)\|_2 + 2|V|)\|_{\psi_1} \\ &\leq 4R \|\|\theta(Y)\|_2^2\|_{\psi_1} + 4 \|\|\theta(Y)\|_2 |V|\|_{\psi_1} \\ &\leq 4R \|\|\theta(Y)\|_2^2\|_{\psi_2}^2 + 4 \|\|\theta(Y)\|_2\|_{\psi_2} \|V\|_{\psi_2} \\ &\leq 4Rm \|\|\theta(Y)\|_{\psi_2}^2 + 4\sqrt{m} \|\|\theta(Y)\|_{\psi_2}\|V\|_{\psi_2}. \end{aligned} \quad (50)$$

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 $\frac{1}{n} \sum_{i=1}^n V_i \leq 2\mathbb{E}[V^2]$, the norm $\|\phi\|_2 \leq R$. This event holds with probability at least $1 - \iota(n)$ with $\iota(n) = \eta(n)$, see (44) in the proof of Assumption 3. ■

Proof [Proof of Corollary 27] The proof of the first part is a direct application of Theorem 9 with the constants of Lemma 26. The second part requires a little more work.

The total expectation can be split accordingly to (22) where the first term can be bounded using Lemma 7. The second term is bounded as follows. Due to the optimality of ϕ_* and $\hat{\phi}$ one has

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

see the proof of Lemma 26. Therefore,

$$\begin{aligned} \mathbb{E} \left[\vartheta(\hat{\mathcal{S}}; \mathcal{S})^2 \mathbf{1}_{\mathbf{x} \notin \mathcal{S}} \right] &\leq \frac{1}{2} \mathbb{E}[V^2]^2 \lambda^{-2} p_n + \frac{1}{2} \mathbb{E}[V^4 \mathbf{1}_{\mathbf{x} \notin \mathcal{S}}] \\ &\leq \frac{1}{2} \mathbb{E}[V^2]^2 \lambda^{-2} p_n + \frac{1}{2} \mathbb{E}[V^8]^{\frac{1}{2}} \sqrt{p_n} \\ &\leq (2\|V\|_{\psi_2}^2 + 2^5 \|V\|_{\psi_2}^4) \lambda^{-2} \sqrt{p_n}. \end{aligned}$$
■

B.4 Entropic-Wasserstein Barycenters

Proof [Proof of Lemma 28] The proof consists 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 measures $\mu \in \mathcal{P}_2(\mathcal{M})$. Using Proposition 2.1 of Carlier et al. (2021) allows to conclude that ϕ_* and $\hat{\phi}$ exist and are unique.

Assumption 2 The function $W_2^2(\cdot, \psi)$ is convex from Proposition 7.17 of Santambrogio (2015) 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{R} and $\hat{\mathcal{R}}$ 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 $\vartheta(\phi, \psi) \leq \sqrt{\frac{1}{2} \text{KL}(\phi|\psi)}$. 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

$$2\vartheta(\phi, \psi)^2 \leq \mathcal{F}(\phi) - \mathcal{F}_*,$$

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

Assumption 3 Since Assumption (2) is verified with $\mathcal{F}_0 = +\infty$, 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 Villani, 2003, 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 (Villani, 2003),

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

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,

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

Assuming without loss of generality that the point $0_{\mathbb{R}^d} \in \Omega$, the Kantorovich ζ are $2 \operatorname{diam}(\Omega)$ -Lipschitz so that

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

where the last inequality is obtained through the transport inequality $W_1(\phi_0, \phi_1) \leq \operatorname{diam}(\Omega)\vartheta(\phi_0, \phi_1)$ (see Gibbs and Su, 2002, 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

- Martial Agueh and Guillaume Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- 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.
- 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.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows: in Metric Spaces and in the Space of Probability Measures*. Springer, 2008.
- Miroslav Bacák. Computing medians and means in Hadamard spaces. *SIAM Journal on Optimization*, 24(3):1542–1566, 2014.
- Francis Bach. *Learning Theory from First Principles*. The MIT Press, 2024.
- Eustasio Del Barrio, Paul Deheuvels, and Sara van de Geer. *Lectures on Empirical Processes: Theory and Statistical Applications*. European Mathematical Society, 2007.
- Peter Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Peter Bartlett and Shahar Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.

- Peter Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *The Annals of Statistics*, 33(4):1487–1537, 2005.
- Saugata Basu, Richard Pollack, and Marie-Françoise Coste-Roy. *Algorithms in Real Algebraic Geometry*. Springer, 2007.
- Faisal Beg, Michael 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.
- Ira Berg and Igor Nikolaev. Quasilinearization and curvature of Aleksandrov spaces. *Geometriae Dedicata*, 133(1):195–218, 2008.
- 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.
- Edward Bierstone and Pierre Milman. Semianalytic and subanalytic sets. *Publications Mathématiques de l’IHÉS*, 67:5–42, 1988.
- 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, 2019a.
- Jérémie Bigot, Elsa Cazelles, and Nicolas Papadakis. Penalization of barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 51(3):2261–2285, 2019b.
- Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry*, volume 36. Springer, 2013.
- Jérôme Bolte, Trong Phong Nguyen, Juan Peypouquet, and Bruce Suter. From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming*, 165(2):471–507, 2017.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2:499–526, 2002.
- Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to statistical learning theory. In *Advanced Lectures on Machine Learning*, pages 169–207. Springer, 2003.
- Olivier Bousquet, Yegor Klochkov, and Nikita Zhivotovskiy. Sharper bounds for uniformly stable algorithms. In *Conference on Learning Theory*, pages 610–626, 2020.
- Victor-Emmanuel Brunel and Jordan Serres. Concentration of empirical barycenters in metric spaces. In *International Conference on Algorithmic Learning Theory*, pages 337–361, 2024.

- Emmanuel Candes, Justin Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- 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.
- Sourav Chatterjee. A new perspective on least squares under convex constraint. *The Annals of Statistics*, 42(6):2340–2381, 2014.
- Richard Combes. An extension of McDiarmid’s inequality. *arXiv preprint arXiv:1511.05240*, 2015.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996.
- David Donoho. For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics*, 59(6):797–829, 2006.
- Benjamin Eltzner and Stephan Huckemann. A smeary central limit theorem for manifolds with application to high-dimensional spheres. *The Annals of Statistics*, 47(6):3360–3381, 2019.
- Vitaly Feldman and Jan Vondrak. Generalization bounds for uniformly stable algorithms. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- 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, 2019.
- Cees Fortuin, Pieter Kasteleyn, and Jean Ginibre. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22(2):89–103, 1971.
- Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l’Institut Henri Poincaré*, 10:215–310, 1948.
- Alison Gibbs and Francis Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.
- Osman Güler. *Foundations of Optimization*, volume 258. Springer, 2010.
- Christian Hess. Epi-convergence of sequences of normal integrands and strong consistency of the maximum likelihood estimator. *The Annals of Statistics*, 24(3):1298–1315, 1996.
- Jean-Baptiste Hiriart-Urruty. Extension of Lipschitz functions. *Journal of Mathematical Analysis and Applications*, 77(2):539–554, 1980.
- Alan Hoffman. On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards*, 49(4):263–265, 1952.

- Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6):417, 1933.
- Chi Jin, Praneeth Netrapalli, Rong Ge, Sham Kakade, and Michael Jordan. A short note on concentration inequalities for random vectors with subgaussian norm. arXiv preprint arXiv:1902.03736, 2019.
- Michael Kearns and Dana Ron. Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. In *Conference on Computational Learning Theory*, pages 152–162, 1997.
- Mojzesz Kirszbraun. Über die zusammenziehende und lipschitzsche transformationen. *Fundamenta Mathematicae*, 22(1):77–108, 1934.
- Yegor Klochkov and Nikita Zhivotovskiy. Stability and deviation optimal risk bounds with convergence rate $O(1/n)$. In *Advances in Neural Information Processing Systems*, volume 34, pages 5065–5076, 2021.
- Keith Knight and Wenjiang Fu. Asymptotics for Lasso-type estimators. *The Annals of Statistics*, pages 1356–1378, 2000.
- Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- Vladimir Koltchinskii. Sparsity in penalized empirical risk minimization. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 45(1):7–57, 2009.
- Aryeh Kontorovich. Concentration in unbounded metric spaces and algorithmic stability. In *International Conference on Machine Learning*, pages 28–36, 2014.
- Alexey Kroshnin. Fréchet barycenters in the Monge-Kantorovich spaces. *Journal of Convex Analysis*, 15:1371–1395, 2018.
- Samuel Kutin. Extensions to McDiarmid’s inequality when differences are bounded with high probability. Technical report, Dept. Comput. Sci., Univ. Chicago, Chicago, IL, USA, 2002.
- Samuel Kutin and Partha Niyogi. Almost-everywhere algorithmic stability and generalization error. In *Conference on Uncertainty in Artificial Intelligence*, pages 275–282, 2002.
- Thibaut Le Gouic and Jean-Michel Loubes. Existence and consistency of Wasserstein barycenters. *Probability Theory and Related Fields*, 168(3):901–917, 2017.
- Thibaut Le Gouic, Quentin Paris, Philippe Rigollet, and Austin Stromme. Fast convergence of empirical barycenters in Alexandrov spaces and the Wasserstein space. *Journal of the European Mathematical Society*, 2022.
- Jing Lei. Convergence and concentration of empirical measures under Wasserstein distance in unbounded functional spaces. *Bernoulli*, 26(1):767–798, 2020.

- Guoyin Li. Global error bounds for piecewise convex polynomials. *Mathematical Programming*, 137(1):37–64, 2013.
- Stanislas Lojasiewicz. Sur la géométrie semi-et sous-analytique. *Annales de l'Institut Fourier*, 43(5):1575–1595, 1993.
- Andreas Maurer and Massimiliano Pontil. Concentration inequalities under sub-Gaussian and sub-exponential conditions. *Advances in Neural Information Processing Systems*, 34, 2021.
- Colin McDiarmid. Concentration. In *Probabilistic Methods for Algorithmic Discrete Mathematics*, pages 195–248. Springer, 1998.
- Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Transactions on Information Theory*, 54(8):3797–3803, 2008.
- Shahar Mendelson. Learning without concentration. In *Conference on Learning Theory*, pages 25–39, 2014.
- Shahar Mendelson. Learning without concentration for general loss functions. *Probability Theory and Related Fields*, 171(1):459–502, 2018.
- Karl Pearson. 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.
- Bruno Pelletier. Informative barycentres in statistics. *Annals of the Institute of Statistical Mathematics*, 57(4):767–780, 2005.
- Javier Pena, Juan Vera, and Luis Zuluaga. New characterizations of Hoffman constants for systems of linear constraints. *Mathematical Programming*, 187:79–109, 2021.
- Xavier Pennec. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Allan Pinkus. *N-widths in Approximation Theory*. Springer, 1985.
- Alexander Rakhlin, Sayan Mukherjee, and Tomaso Poggio. Stability results in learning theory. *Analysis and Applications*, 3(04):397–417, 2005.
- Garvesh Raskutti, Martin Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- William Rogers and Terry Wagner. A finite sample distribution-free performance bound for local discrimination rules. *The Annals of Statistics*, pages 506–514, 1978.
- Mark Rudelson and Shuheng Zhou. Reconstruction from anisotropic random measurements. In *Conference on Learning Theory*, pages 10–1, 2012.

- Filippo Santambrogio. *Optimal Transport for Applied Mathematicians*. Birkäuser, 2015.
- Christof Schötz. Convergence rates for the generalized Fréchet mean via the quadruple inequality. *Electronic Journal of Statistics*, 13(2):4280–4345, 2019.
- Christof Schötz. Variance inequalities for transformed Fréchet means in Hadamard spaces. *arXiv preprint arXiv:2310.13668*, 2023.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *Conference on Learning Theory*, 2009.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11:2635–2670, 2010.
- Austin Stromme. *Wasserstein Barycenters: Statistics and Optimization*. PhD thesis, Massachusetts Institute of Technology, 2020.
- Karl-Theodor Sturm. Probability measures on metric spaces of nonpositive curvature. In *Heat Kernels and Analysis on Manifolds, Graphs, and Metric Spaces*, volume 338, pages 357–390. American Mathematical Society, 2003.
- Michel Talagrand. *Upper and Lower bounds for Stochastic Processes*, volume 60. Springer, 2014.
- Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Joel Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12:389–434, 2012.
- Alexandre Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.
- Leslie Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- Sara van de Geer and Martin Wainwright. On concentration for (regularized) empirical risk minimization. *Sankhya A*, 79:159–200, 2017.
- Aad Van Der Vaart and Jon Wellner. *Weak Convergence and Empirical Processes: with Applications to Statistics*. Springer, 1997.
- Vladimir Vapnik. Principles of risk minimization for learning theory. In *Advances in Neural Information Processing Systems*, volume 4, 1991.
- Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264, 1971.

- Vladimir Vapnik and Alexey Chervonenkis. *Theory of Pattern Recognition*. Nauka, 1974.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47. Cambridge University Press, 2018.
- Cédric Villani. *Topics in Optimal Transportation*. Number 58 in Grad. Stud. Math. American Mathematical Society, 2003.
- Martin Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- Lutz Warnke. On the method of typical bounded differences. *Combinatorics, Probability and Computing*, 25(2):269–299, 2016.
- Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.
- Ho Yun and Byeong Park. Exponential concentration for geometric-median-of-means in non-positive curvature spaces. *Bernoulli*, 29(4):2927–2960, 2023.