

# Exploring Local Norms in Exp-concave Statistical Learning

**Nikita Puchkin**

*HSE University and IITP RAS, Russian Federation*

NPUCHKIN@HSE.RU

**Nikita Zhivotovskiy**

*UC Berkeley, Department of Statistics, USA*

ZHIVOTOVSKIY@BERKELEY.EDU

**Editors:** Gergely Neu and Lorenzo Rosasco

## Abstract

We consider the problem of stochastic convex optimization with exp-concave losses using Empirical Risk Minimization in a convex class. Answering a question raised in several prior works, we provide a

$$O\left(\frac{d}{n} + \frac{1}{n} \log\left(\frac{1}{\delta}\right)\right)$$

excess risk bound valid for a wide class of bounded exp-concave losses, where  $d$  is the dimension of the convex reference set,  $n$  is the sample size, and  $\delta$  is the confidence level. Our result is based on a unified geometric assumption on the gradient of losses and the notion of local norms.

## 1. Problem statement

We consider a stochastic convex optimization problem, that is, minimization of the objective

$$F(w) = Pf(w, Z)$$

over a compact *convex* set  $\mathcal{W} \subset \mathbb{R}^d$ . Here  $Z$  is a random element with a distribution  $P$  supported on a set  $Z$ , and  $Pf(w, Z)$  stands for the expectation of  $f(w, Z)$  with respect to  $Z \sim P$ . Since  $P$  is usually unknown, the optimizer

$$w^* \in \operatorname{argmin}_{w \in \mathcal{W}} F(w) \tag{1}$$

cannot be computed explicitly. Instead, a learner has an access to a sample  $S_n = \{Z_i : 1 \leq i \leq n\}$  with independent elements distributed according to  $P$ . The most natural and popular strategy in this case is to replace the expectation with respect to  $Z \sim P$  by the sample mean and study an empirical risk minimizer (ERM)

$$\hat{w} \in \operatorname{argmin}_{w \in \mathcal{W}} \hat{F}(w), \tag{2}$$

where

$$\hat{F}(w) = P_n f(w, Z) = \frac{1}{n} \sum_{i=1}^n f(w, Z_i) \quad \text{for all } w \in \mathcal{W}. \tag{3}$$

Here and further in the paper,  $P_n$  stands for the empirical measure, corresponding to the sample  $S_n$ . We are interested in tight high probability bounds on the excess risk. Our main result will imply a bound of the form

$$F(\hat{w}) - F(w^*) = O\left(\frac{d}{n} + \frac{1}{n} \log\left(\frac{1}{\delta}\right)\right),$$

for a general class of bounded exp-concave losses.

**Contributions.** Our contribution combines two ideas: specific assumptions on the loss function and a careful probabilistic analysis utilizing empirical process techniques. We propose a general assumption on loss functions that applies to most commonly used exp-concave losses in the bounded setup. Our main result, Theorem 3.4, addresses the question raised explicitly in a series of papers (Mahdavi, Zhang, and Jin, 2015; Koren and Levy, 2015; Mehta, 2017; Gonen and Shalev-Shwartz, 2018; Yang, Li, and Zhang, 2018) by further exploring the concept of local norms. Despite the presence of established empirical process techniques in related contexts, our approach handles the loss functions exhibiting data-dependent curvature. Compared to the existing localized analysis for the losses with (data-independent) curvature (Mendelson, 2002; Sridharan et al., 2008; Koltchinskii, 2011), we do not assume any linear structure, so the common approach based on computing the local complexity of the linear class is not applicable in our situation. Similarly, some more recent applications of empirical process techniques to exp-concave losses (Vijaykumar, 2021) do not fully exploit the local geometric structure of the reference set and thus lead to suboptimal bounds in our setup. We discuss this in more detail in Section 4.

**Notation.** Throughout the paper, the notation  $f \lesssim g$  or  $g \gtrsim f$  means that for some universal constant  $c > 0$  we have  $f \leq cg$ . For any set  $\mathcal{W} \subseteq \mathbb{R}^d$ ,  $\text{Int}(\mathcal{W})$  denotes its interior. For a symmetric positive semidefinite matrix  $A \in \mathbb{R}^{d \times d}$ ,  $\|\cdot\|_A$  stands for a seminorm, induced by  $A$ :

$$\|w\|_A = \sqrt{w^\top A w}, \quad \text{for all } w \in \mathcal{W}.$$

Finally, for any two positive semidefinite matrices  $A$  and  $B$  of the same size the inequalities  $A \preceq B$  and  $B \succeq A$  mean that  $(B - A)$  is positive semidefinite.

**Structure of the paper.** The rest of the paper is organized as follows. In Section 2, we make a brief overview of related works. In Section 3, we introduce our geometric assumption, present our main result, Theorem 3.4, and discuss its implications. Section 4 is devoted to the proof of Theorem 3.4. It heavily relies on a technical statement about the supremum of a certain empirical process, Lemma 4.2. We provide its proof in Section 5. Proofs of auxiliary results appear in Appendix.

## 2. Related work

Theoretical properties of the empirical risk minimizer  $\hat{w}$ , especially its generalization ability, were extensively studied in the contexts of stochastic convex optimization and statistical learning theory. Given the extensive body of relevant literature, we review only a subset of the most relevant papers. Early results in this direction include the analysis of linear regression with squared loss via local Rademacher complexities (see, e.g., (Mendelson, 2002) and (Koltchinskii, 2011, Section 5.1)). In particular, Koltchinskii (2011, Section 5.1, Example 1) shows that, in the problem of bounded linear regression in  $\mathbb{R}^d$ , the excess squared risk of a ERM scales as  $O(d/n)$  with high probability. In contrast to the standard results on learning general VC-type classes under various margin conditions (see, e.g., (Bartlett, Bousquet, and Mendelson, 2005; Massart and Nédélec, 2006; Zhivotovskiy and Hanneke, 2018), and (Koltchinskii, 2011, Section 5.1, Example 3)), the linear regression specific rate  $O(d/n)$  avoids an extra  $\log n$  factor and exhibits tighter bounds on generalization ability of ERM. The reason for that is a linear structure of the reference set, which allows to bound its local Rademacher complexity using the standard Cauchy-Schwarz or Fenchel inequalities (see (Koltchinskii, 2011, Proposition 3.2) and (Kakade et al., 2008, Theorem 1)).

Many researchers have further explored the possibility of achieving an excess risk bound better than  $O(d \log n/n)$  for problems with specific local structures in the reference set. Recent investigations in stochastic approximation lead to numerous tight in-expectation  $O(d/n)$  bounds on the excess risk of particular optimization algorithms (see, e.g., (Moulines and Bach, 2011; Bach and Moulines, 2013; Woodworth and Srebro, 2021) to mention a few). The current paper focuses on high probability bounds, which typically require a more technical analysis. In a study by Sridharan, Shalev-Shwartz, and Srebro (2008), the authors extended the analysis of local Rademacher complexities to the case of regularized strongly convex losses,  $f(w, z) = \ell(y, w^\top x) + R(w)$ , where  $z = (x, y)$ ,  $\ell$  is convex with respect to its second argument for any  $y$ , and  $R(w)$  is strongly convex. Similar to an example in (Koltchinskii, 2011, Section 5.1, Example 1), the key to the proof was a bound on the local Rademacher complexity of a linear class, as demonstrated in Lemma 7 in (Sridharan et al., 2008). The PAC-Bayesian localization technique was introduced by Audibert and Catoni (2010, 2011) to obtain high-probability bounds on the squared excess risk of ERM in the bounded linear regression problem. These authors achieved a rate of  $O((d + \log(1/\delta))/n)$ . This approach has also been used in (Zhang, 2006) in the context of density estimation. Related to our study, Klochkov and Zhivotovskiy (2021) proved a dimension-free excess risk bound of  $O(\log(n) \log(1/\delta)/n)$  with high probability for arbitrary strongly convex and Lipschitz losses, using the uniform stability of ERM. We also would like to mention that excess risk bounds can be derived from regret bounds for online learning algorithms via the online-to-batch conversion. We refer to (Kakade and Tewari, 2008; Gaillard and Wintenberger, 2018; Van der Hoeven et al., 2022) for related high probability upper bounds. However, in our context, the online-to-batch approach yields risk bounds with additional logarithmic factors, which makes it unsuitable for our purposes.

This paper focuses on stochastic convex optimization problems using a broad class of exp-concave losses. Our goal is to provide tight high probability bounds  $O((d + \log(1/\delta))/n)$  on the excess risk of ERM under mild assumptions on the function  $f(w, z)$ . In the context of exp-concave losses, this question was directly addressed in a series of papers and turned out to be quite challenging. It appears that one also needs to make some additional assumptions on the loss function that allows to capture the local geometry of the reference set. Several works have achieved the desired  $O(d/n)$  rate in expectation (e.g., (Koren and Levy, 2015; Gonen and Shalev-Shwartz, 2018; Zhang and Zhou, 2019)), but transforming this guarantee into a high probability bound results in suboptimal  $O((d \log n + \log(1/\delta))/n)$  guarantees. Several groups of authors (Mahdavi, Zhang, and Jin, 2015; Koren and Levy, 2015; Mehta, 2017; Gonen and Shalev-Shwartz, 2018; Yang, Li, and Zhang, 2018) have questioned the possibility of achieving the optimal  $O((d + \log(1/\delta))/n)$  high probability excess risk bound. The best bound in this direction is given by the confidence boosting approach suggested by Mehta (2017), which replaces ERM with an aggregate of independently trained ERMs, resulting in a bound of  $O(d \log(1/\delta)/n)$ . We additionally refer to the recent papers by Mourtada and Gaïffas (2022) and Bilodeau, Foster, and Roy (2021) where a detailed survey of relevant results is provided.

The cornerstone of the analysis of empirical risk minimizers in learning problems with the squared loss (or more general strongly convex losses) is localization. Certain fixed point equations appeared in influential works (Van de Geer, 1990; Birgé and Massart, 1993; Yang and Barron, 1999) and then became familiar to the machine learning community thanks to the papers (Koltchinskii, 2006; Bartlett, Bousquet, and Mendelson, 2005). We refer to the monographs (Van de Geer, 2000; Koltchinskii, 2011) and the paper (Mendelson, 2002) for examples of application of local entropy and local Rademacher complexity to the analysis of linear regression with squared loss. Finally, we

note that the analysis of the local structure of the reference set in the context of density estimation has its roots in the classical works of LeCam (1973) and Birgé (1983). For related upper bounds, we refer to Section 32.2.3 in the forthcoming book by Polyanskiy and Wu (2023).

### 3. Main results and examples

In this section, we present our main result. Throughout the paper, we assume that the loss function  $f(w, z)$  satisfies the following condition.

**Assumption 3.1** *For any training sample  $Z_1, \dots, Z_n$  there is positive semidefinite matrix  $H \in \mathbb{R}^{d \times d}$ , which may depend on the sample, such that*

- *the expectation of  $H$  is finite;*
- *$\widehat{F}(w)$  given by (3) is  $\sigma$ -strongly convex on  $\mathcal{W}$  with respect to the seminorm, induced by  $H$ ; that is, for any  $u, v \in \mathcal{W}$  and  $\alpha \in [0, 1]$ , the next inequality holds almost surely:*

$$\widehat{F}(\alpha u + (1 - \alpha)v) \leq \alpha \widehat{F}(u) + (1 - \alpha) \widehat{F}(v) - \frac{\sigma \alpha (1 - \alpha)}{2} \|u - v\|_H^2; \quad (4)$$

- *for any  $u, v \in \mathcal{W}$ , the empirical  $L_2$ -distance between  $f(u, Z)$  and  $f(v, Z)$  satisfies the inequality*

$$P_n(f(u, Z) - f(v, Z))^2 \leq L^2 \|u - v\|_H^2.$$

Here,  $\sigma$  and  $L$  are positive constants that do not depend on  $Z_1, \dots, Z_n$ .

When  $H$  is the identity matrix, we simply say that  $\widehat{F}$  is a strongly convex function. Similarly, when  $H$  is the identity matrix, the second part of this assumption corresponds to the standard Lipschitz property. Though Assumption 3.1 may appear similar to the commonly used assumptions of strong convexity and Lipschitzness, the introduction of the data-dependent matrix  $H$  broadens the range of applications of our model. In particular, it includes regularized objectives, considered in (Sridharan et al., 2008), and generalized linear models (see, e.g., (Gonen and Shalev-Shwartz, 2018)), where the loss function has a form

$$f(w, z) = \ell(w^\top x, y). \quad (5)$$

Here  $z = (x, y)$  is a feature-label pair,  $w$  and  $x$  belong to bounded subsets of  $\mathbb{R}^d$ , and  $\ell$  is  $L$ -Lipschitz and  $\sigma$ -strongly convex with respect to the first argument for all  $y$ . In this case,  $\widehat{F}(w)$  is  $\sigma$ -strongly convex with respect to the seminorm  $\|\cdot\|_H$ , induced by  $H = (X_1 X_1^\top + \dots + X_n X_n^\top)/n$  (see Example 3.6). Our setup covers both of the most commonly used generalized linear models: linear regression, which corresponds to  $f(w, z) = (y - w^\top x)^2$ , and logistic regression, which corresponds to  $f(w, z) = \log(1 + e^{-y w^\top x})$  with  $y \in \{-1, 1\}$ .

Under mild differentiability assumptions,  $\widehat{F}$  demonstrates exponential concavity, as shown in the following result.

**Proposition 3.2** *Suppose that  $\widehat{F}$  satisfies Assumption 3.1 and let the loss function  $f(w, z)$  be twice differentiable with respect to the first argument for all  $z \in Z$ . Then  $\widehat{F}$  is  $(\sigma/L^2)$ -exponentially concave on  $\mathcal{W}$ . That is,  $e^{-\sigma \widehat{F}/L^2}$  is concave on  $\mathcal{W}$ .*

However, our assumptions are slightly stronger than exp-concavity alone. In fact, our assumption implies the boundedness of the loss.

**Proposition 3.3** *Under Assumption 3.1, we have  $|f(w, Z) - f(w^*, Z)| \leq 4L^2/\sigma$  for all  $w \in \mathcal{W}$  and almost surely with respect to  $Z$ .*

Observe that this result does not necessarily imply the boundedness of the set  $\mathcal{W}$ . This distinguishes our assumptions from the more restrictive strongly convex Lipschitz assumption, which implies both the boundedness of the loss and  $\mathcal{W}$ . In fact, the proof of Proposition 3.3 yields that

$$\|w - \hat{w}\|_{\mathbf{H}}^2 \leq \frac{4}{\sigma} \left( \hat{F}(w) - \hat{F}(\hat{w}) \right) \leq \frac{32L^2}{\sigma^2} \quad \text{for all } w \in \mathcal{W}.$$

If  $\det(\mathbf{H}) \neq 0$ , then  $\mathcal{W}$  is bounded. However, in general, the matrix  $\mathbf{H}$  is not assumed to be nondegenerate. Hence,  $\mathcal{W}$  may potentially be unbounded under our assumption. The proofs of Propositions 3.3 and 3.2 are deferred to Appendix A. After listing these basic properties, we are ready to present our main result.

**Theorem 3.4** *Assume that  $\mathcal{W} \subset \mathbb{R}^d$  is a compact convex set. Under Assumption 3.1, for any  $\delta \in (0, 1)$ , it holds that, with probability at least  $1 - \delta$ , the ERM estimator (2) satisfies*

$$F(\hat{w}) - \min_{w \in \mathcal{W}} F(w) \lesssim \frac{L^2(d + \log(1/\delta))}{\sigma n}.$$

Theorem 3.4 answers a natural question raised in several previous papers. The most relevant to our result are the papers of [Koren and Levy \(2015\)](#); [Gonen and Shalev-Shwartz \(2018\)](#), where the authors obtained optimal in-expectation bounds of order  $O(d/n)$ , and of [Mahdavi et al. \(2015\)](#); [Mehta \(2017\)](#); [Yang et al. \(2018\)](#), where the authors proved  $O((d \log n + \log(1/\delta))/n)$  high probability upper bounds. We will now discuss the differences in assumptions made on the loss functions. In ([Koren and Levy, 2015](#); [Yang et al., 2018](#)), the authors required the gradient  $\partial f(w, z)/\partial w$  be Lipschitz on  $\mathcal{W}$ . [Koren and Levy \(2015\)](#) also asked if the smoothness of the gradient is necessary for obtaining fast rates for ERM. A partial answer was given in ([Gonen and Shalev-Shwartz, 2018](#)), where the authors considered generalized linear models with the losses of the form (5), where  $\ell(\cdot, y)$  is Lipschitz and strongly convex with respect to its first argument for all  $y$ 's. Under these assumptions, the authors showed that the expected excess risk of ERM can decay as fast as  $O(d/n)$ . Our result strictly improves the one of [Gonen and Shalev-Shwartz \(2018\)](#). As it was discussed earlier, our Assumption 3.1 is more general than the requirements of [Gonen and Shalev-Shwartz \(2018\)](#). At the same time, we derive a rate optimal  $O((d + \log(1/\delta))/n)$  high probability excess risk bound. Concerning the papers ([Mahdavi et al., 2015](#); [Mehta, 2017](#)), the authors work under more general assumptions than Assumption 3.1. Namely, they only assume that the loss  $f(w, z)$  is bounded, Lipschitz and  $\sigma$ -exponentially concave with respect to  $w \in \mathcal{W}$  for all  $z \in \mathcal{Z}$ . It is not clear, if it is possible to get rid of extra  $\log n$  factors in their high probability bounds. However, we would like to note that our setup covers all interesting bounded loss functions arising in exp-concave learning with convex reference sets.

Theorem 3.4 has several implications. We discuss them in the examples below.

**Example 3.5 (Linear regression with squared loss)** *Fix  $R > 0$ . Assume that the observations  $Z_1, \dots, Z_n$  have a form  $Z_i = (X_i, Y_i)$ ,  $1 \leq i \leq n$ , where  $X_i \in \mathbb{R}^d$  and  $Y_i \in \mathbb{R}$  have distributions*

supported on  $\{x \in \mathbb{R}^d : \|x\| \leq 1\}$  and  $[-R, R]$ , respectively. Consider  $\mathcal{W} = \{w \in \mathbb{R}^d : \|w\| \leq R\}$  and the quadratic loss function  $f(w, z) = (y - w^\top x)^2$ . Note that  $f(w, z)$  fulfils Assumption 3.1 with the matrix  $\mathbf{H} = (X_1 X_1^\top + \dots + X_n X_n^\top)/n$  and the constants  $\sigma = 2$ ,  $L = 4R$ . According to Theorem 3.4, for any  $\delta \in (0, 1)$ , an empirical risk minimizer  $\hat{w}$  satisfies the inequality

$$F(\hat{w}) - F(w^*) \lesssim \frac{R^2(d + \log(1/\delta))}{n}$$

with probability at least  $1 - \delta$ . This agrees with the upper bound of [Koltchinskii \(2011, Section 5.1, Example 1\)](#), based on local Rademacher complexities, and the lower bound of [Shamir \(2015, Theorem 1\)](#), who showed that for any estimator  $\tilde{w} \in \mathcal{W}$  there exists a data distribution over  $\mathcal{B}(0, 1) \times [-R, R]$  such that  $\mathbb{E}F(\tilde{w}) - F(w^*) \gtrsim R^2 d/n$ .

**Example 3.6 (Logistic regression)** Again, assume that the observations  $Z_1, \dots, Z_n$  have a form  $Z_i = (X_i, Y_i)$ ,  $1 \leq i \leq n$ , where  $X_i \in \mathbb{R}^d$  and  $Y_i \in \{-1, 1\}$ . Fix  $R > 0$  and let the distribution of  $X_i$ 's be supported on  $\mathcal{B}(0, R) = \{x \in \mathbb{R}^d : \|x\| \leq R\}$ . Consider  $\mathcal{W} = \{w \in \mathbb{R}^d : \|w\| \leq B\}$  and the logistic loss

$$f(w, z) = \log(1 + e^{-y w^\top x}).$$

Direct calculations show that the map  $u \mapsto \log(1 + e^u)$  is 1-Lipschitz and  $e^{-BR}$ -strongly convex on  $[-BR, BR]$ . Hence, the logistic loss meets Assumption 3.1 with the matrix

$$\mathbf{H} = \frac{1}{n} \sum_{i=1}^n X_i X_i^\top$$

and the constants  $\sigma = e^{-BR}$ ,  $L = 1$ . Theorem 3.4 implies that, for any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , it holds that

$$F(\hat{w}) - F(w^*) \lesssim \frac{e^{BR}(d + \log(1/\delta))}{n}. \quad (6)$$

We refer to [\(Hazan, Koren, and Levy, 2014\)](#), where the authors discuss the inavoidability of the  $e^{BR}$  factor for ERM. We note that the bound (6) has not appeared in the literature before.

Furthermore, our bound covers strongly convex and Lipschitz losses, regardless of any linear structure in the loss function. In [\(Klochkov and Zhivotovskiy, 2021\)](#), the authors used the uniform stability argument to prove that, if  $f(\cdot, z)$  is  $L$ -Lipschitz and  $\sigma$ -strongly convex for all  $z \in \mathcal{Z}$ , then, for any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , it holds that

$$F(\hat{w}) - F(w^*) \lesssim \frac{L^2 \log n \log(1/\delta)}{\sigma n}. \quad (7)$$

Taking together their result and Theorem 3.4, we get the following bound.

**Corollary 3.7 (Strongly convex and Lipschitz losses)** Assume that, for any  $z \in \mathcal{Z}$ , the loss function  $f(\cdot, z)$  is  $L$ -Lipschitz and  $\sigma$ -strongly convex on  $\mathcal{W} \subseteq \mathbb{R}^d$ , which is a compact convex set. Then, for any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , it holds that

$$F(\hat{w}) - F(w^*) \lesssim \frac{L^2 \min\{d + \log(1/\delta), \log n \log(1/\delta)\}}{\sigma n}.$$

This gives the best known high probability upper bound in this standard setup. Recent results on algorithmic stability (Bousquet et al., 2020; Klochkov and Zhivotovskiy, 2021) leave the question if the  $\log n$  term can be removed in the bounds such as (7). Our Corollary (3.7) shows that if this factor cannot be removed, then any demonstration of this must occur in the regime  $d + \log(1/\delta) \gtrsim \log n \log(1/\delta)$ .

**On logarithmic factors and improper learning.** We focus on the sharp analysis of ERM in a convex reference set  $\mathcal{W}$ . Nevertheless, recent work has demonstrated that the use of so-called *improper* estimators — estimators that predict values outside of the reference set  $\mathcal{W}$  — can lead to significant improvements in the dependence on some parameters. These estimators are typically based on aggregating the functions within certain *non-convex* reference sets constructed by truncating the original functions that correspond to  $\mathcal{W}$ . Such improvements are known in both logistic regression (Kakade and Ng, 2004; Foster et al., 2018; Mourtada and Gaïffas, 2022) and regression with the squared loss (Forster and Warmuth, 2002; Vaškevičius and Zhivotovskiy, 2023; Mourtada et al., 2021). However, it should be noted that the improper learning setup poses a considerable challenge for the localized analysis utilized in this paper. Specifically, the lower bound of Rakhlin, Sridharan, and Tsybakov (2017, Theorem 6) indicates that, in the context of bounded regression with a non-convex reference set, the lower bound of  $\Omega(\log(n)/n)$  cannot be surpassed even with the use of improper learners. At present, there is no known analysis that can yield high probability excess risk bounds of order  $O((d + \log(1/\delta))/n)$  for any improper learner in non-convex reference sets.

#### 4. Proof of Theorem 3.4

Our proof is making most of the standard arguments on the Laplace transform of *shifted* or *offset* empirical processes. A similar approach was exploited by many authors Wegkamp (2003); Lecué and Rigollet (2014); Liang, Rakhlin, and Sridharan (2015); Zhivotovskiy and Hanneke (2018); Kanade, Rebeschini, and Vaškevičius (2022) in the statistical setup, though their analysis is specific to strongly convex losses or binary losses under additional probabilistic assumptions. In (Vijaykumar, 2021), the author generalized the approach of Liang et al. (2015) to study the properties of ERM under exp-concave losses, though their analysis only focuses on getting the  $O(1/n)$  rate of convergence and fails to capture the local structure of the reference set. In our context their result would give an additional  $\log n$  factor, which is known to be achievable using an  $\varepsilon$ -net argument (Mehta, 2017). The key point of our analysis is based on relating Assumption 3.1 with the localized covering numbers of the set  $\mathcal{W}$  with respect to the data-dependent semi-norm induced by  $H$ . Two papers that are worth mentioning on the topic of online learning and exp-concave losses, which utilize sequential offset processes, are (Rakhlin and Sridharan, 2015) and (Bilodeau et al., 2020).

**Lemma 4.1** *Under Assumption 3.1, let  $\Phi : \mathbb{R} \rightarrow \mathbb{R}$  be a convex monotonously increasing function. Then, it holds that*

$$\mathbb{E}\Phi(F(\hat{w}) - F(w^*)) \leq \mathbb{E}\mathbb{E}_\varepsilon \Phi \left( 4 \sup_{w \in \mathcal{W}} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_H^2 \right] \right),$$

where  $\varepsilon_1, \dots, \varepsilon_n$  are i.i.d. Rademacher random variables.

The proof of Lemma 4.1 is postponed to Appendix A.3. From now on, we focus our attention on the shifted multiplier process of the form  $P_n \varepsilon(f(w, Z) - f(w^*, Z)) - \sigma \|w - w^*\|_{\mathbb{H}}^2/8$ ,  $w \in \mathcal{W}$ , and consider exponential moments of its supremum. In Section 5, we prove our main technical bound.

**Lemma 4.2** *Let  $f(w, z)$  satisfy Assumption 3.1. Set  $\lambda = \sigma n/(32eL^2)$ . Then, for all  $Z_1, \dots, Z_n$ , it holds that*

$$\mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}} \left[ P_n \varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \leq e + e^{3d} + 14e^{2048(1+e)^2 d/e}.$$

Let us briefly describe the idea of the proof of Lemma 4.2. We use the standard peeling argument and represent the whole set  $\mathcal{W}$  as the union

$$\mathcal{W} = \mathcal{W}[0, r] \cup \left( \bigcup_{k=0}^{\infty} \mathcal{W}[2^k r, 2^{k+1} r] \right),$$

where  $r > 0$  is a fixed number and, for any  $b \geq a \geq 0$ ,  $\mathcal{W}[a, b]$  is defined as

$$\mathcal{W}[a, b] = \{w \in \mathcal{W} : a \leq \|w - w^*\|_{\mathbb{H}} \leq b\}.$$

After this step, we essentially reduce the problem to the analysis of localized offset process, that is, the process of the form

$$P_n \varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2,$$

where  $w$  runs over a local set  $\mathcal{W}[0, r]$ , rather than over the whole set  $\mathcal{W}$ . However, the peeling argument alone is not sufficient to avoid unnecessary  $\log n$  factors, and one should be careful at this point. In (Mehta, 2017, Theorem 1) and (Liang et al., 2015, Lemma 6, Lemma 7), the authors use the  $\varepsilon$ -net argument, which leads to a suboptimal result in our situation. The main difference of our approach from the previous works is hidden in the proof Lemma 5.1, where we combine the localized analysis with chaining and Talagrand's inequality to derive a tight upper bound on the exponential moment

$$\mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon(f(w, Z) - f(w^*, Z)) \right\}.$$

Given that we consider the geometry of the loss in defining  $\mathcal{W}[0, r]$ , deriving an upper bound on its covering number is relatively straightforward. We provide this bound in Lemma 5.2. The key point of this approach, which sets it apart and enables improvement over existing results, is the decision not to use the naive  $\varepsilon$ -net argument.

Lemma 4.1 and Lemma 4.2 immediately imply that

$$\log \mathbb{E} \exp \{ \lambda (F(\hat{w}) - F(w^*)) \} \lesssim d \quad \text{with } \lambda = \frac{\sigma n}{128eL^2}.$$

Hence, for any  $\delta \in (0, 1)$ , the following inequality holds, with probability at least  $1 - \delta$ :

$$F(\hat{w}) - F(w^*) \lesssim \frac{L^2(d + \log(1/\delta))}{\sigma n}.$$

The claim follows. ■



## 5. Proof of Lemma 4.2

Our proof relies on the classical peeling argument. For any  $b \geq a \geq 0$  define

$$\mathcal{W}[a, b] = \{w \in \mathcal{W} : a \leq \|w - w^*\|_{\mathbb{H}} \leq b\}.$$

Let  $r > 0$  be a positive real number to be specified later. Since the function  $\Phi(x) = e^{\lambda x}$  takes only positive values, the expectation of interest does not exceed

$$\begin{aligned} & \mathbb{E}_{\varepsilon} \exp \left\{ \lambda \sup_{w \in \mathcal{W}} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \\ & \leq \mathbb{E}_{\varepsilon} \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, r]} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \\ & \quad + \sum_{k=0}^{\infty} \mathbb{E}_{\varepsilon} \exp \left\{ \lambda \sup_{w \in \mathcal{W}[2^k r, 2^{k+1} r]} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\}. \end{aligned} \quad (8)$$

Thus, we have bounded the exponential moment of the supremum of the empirical process by the sum of exponential moments of localized processes. The main ingredient in the proof of Lemma 4.2 is Lemma 5.1, which provides an upper bound on the exponential moment of the supremum of a localized set without introducing additional logarithmic factors.

**Lemma 5.1** *Let  $f(w, Z)$  satisfy Assumption 3.1. Let*

$$B = \sup_{w \in \mathcal{W}[0, r]} \operatorname{esssup}_{z \in \mathcal{Z}} |f(w, z) - f(w^*, z)|.$$

*Then, for any  $\lambda \geq 0$ , it holds that*

$$\begin{aligned} & \log \mathbb{E}_{\varepsilon} \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon (f(w, Z) - f(w^*, Z)) \right\} \\ & \leq 64\lambda Lr \sqrt{\frac{d}{n}} + \frac{B^2 \lambda^2 e^{B\lambda/n}}{2n} \left( \frac{128Lr}{B} \sqrt{\frac{d}{n}} + \frac{L^2 r^2}{B^2} \right). \end{aligned}$$

*If, in addition,  $\lambda \leq n/B$ , then*

$$\log \mathbb{E}_{\varepsilon} \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon (f(w, Z) - f(w^*, Z)) \right\} \leq 64(1+e)\lambda Lr \sqrt{\frac{d}{n}} + \frac{e\lambda^2 L^2 r^2}{2n}.$$

We defer the proof of Lemma 5.1 to Section 5.1 and focus on finishing the proof of Lemma 4.2. Note that Proposition 3.3 implies that  $B \leq 4L^2/\sigma$  under Assumption 3.1. Since we consider  $\lambda = \sigma n/(32eL^2)$ , the conditions of the second part of Lemma 5.1 are satisfied. Applying it to the

first term in the right-hand side of (8), we obtain that

$$\begin{aligned}
 & \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0,r]} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \\
 & \leq \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0,r]} P_n \varepsilon (f(w, Z) - f(w^*, Z)) \right\} \\
 & \leq \exp \left\{ 64(1+e) \lambda L r \sqrt{\frac{d}{n}} + \frac{e \lambda^2 L^2 r^2}{2n} \right\}.
 \end{aligned} \tag{9}$$

Next, fix a non-negative integer  $k$  and consider

$$\mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[2^k r, 2^{k+1} r]} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\}.$$

In contrast to (9), dropping the negative term leads to a suboptimal bound in this situation. At the same time, the following simple observation allows us to obtain a tight bound. By the definition of  $\mathcal{W}[2^k r, 2^{k+1} r]$ , we have  $\|w - w^*\|_{\mathbb{H}} \geq 2^k r$  for any  $w \in \mathcal{W}[2^k r, 2^{k+1} r]$ . Therefore,

$$\begin{aligned}
 & \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[2^k r, 2^{k+1} r]} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \\
 & \leq \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[2^k r, 2^{k+1} r]} [P_n \varepsilon (f(w, Z) - f(w^*, Z))] - 4^{k-1} \sigma \lambda r^2 / 2 \right\} \\
 & \leq \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, 2^{k+1} r]} [P_n \varepsilon (f(w, Z) - f(w^*, Z))] - 4^{k-1} \sigma \lambda r^2 / 2 \right\}.
 \end{aligned}$$

Applying Lemma 5.1 again, we obtain that

$$\begin{aligned}
 & \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, 2^{k+1} r]} [P_n \varepsilon (f(w, Z) - f(w^*, Z))] - 4^{k-1} \sigma \lambda r^2 / 2 \right\} \\
 & \leq \exp \left\{ 64(1+e) 2^{k+1} \lambda L r \sqrt{\frac{d}{n}} + \frac{e \cdot 4^{k+1} \lambda^2 L^2 r^2}{2n} - 4^{k-1} \sigma \lambda r^2 / 2 \right\} \\
 & = \exp \left\{ 128(1+e) 2^k \lambda L r \sqrt{\frac{d}{n}} + \frac{2e \cdot 4^k \lambda^2 L^2 r^2}{n} - 4^{k-1} \sigma \lambda r^2 / 2 \right\}.
 \end{aligned} \tag{10}$$

Taking (8), (9), and (10) together, we get that

$$\begin{aligned}
 & \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}} \left[ P_n \varepsilon (f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \\
 & \leq \exp \left\{ 64(1+e) \lambda L r \sqrt{\frac{d}{n}} + \frac{e \lambda^2 L^2 r^2}{2n} \right\} \\
 & \quad + \sum_{k=0}^{\infty} \exp \left\{ 128(1+e) 2^k \lambda L r \sqrt{\frac{d}{n}} + \frac{2e \cdot 4^k \lambda^2 L^2 r^2}{n} - 4^{k-1} \sigma \lambda r^2 / 2 \right\}.
 \end{aligned}$$

The remainder of the proof relies on standard computations and on making the appropriate choice for the parameters  $\lambda$  and  $r$ . Let us recall that  $\lambda = \sigma n / (32eL^2)$ . It is easy to observe that such a choice ensures

$$\frac{2e \cdot 4^k \lambda^2 L^2 r^2}{n} = \frac{4^k \sigma \lambda r^2}{16}.$$

Then, it holds that

$$\begin{aligned} & \exp \left\{ 64(1+e)\lambda Lr \sqrt{\frac{d}{n}} + \frac{e\lambda^2 L^2 r^2}{2n} \right\} \\ & + \sum_{k=0}^{\infty} \exp \left\{ 128(1+e) 2^k \lambda Lr \sqrt{\frac{d}{n}} + \frac{2e \cdot 4^k \lambda^2 L^2 r^2}{n} - 4^{k-1} \sigma \lambda r^2 / 2 \right\} \\ & \leq \exp \left\{ 64(1+e)\lambda Lr \sqrt{\frac{d}{n}} + \frac{\sigma \lambda r^2}{64} \right\} + \sum_{k=0}^{\infty} \exp \left\{ 128(1+e) 2^k \lambda Lr \sqrt{\frac{d}{n}} - 4^{k-2} \sigma \lambda r^2 \right\}. \end{aligned}$$

Choosing

$$r = \frac{L}{\sigma} \sqrt{\frac{d}{n}},$$

we obtain that

$$64(1+e)\lambda Lr \sqrt{\frac{d}{n}} + \frac{\sigma \lambda r^2}{64} = \frac{64(1+e)\lambda L^2 d}{\sigma n} + \frac{\lambda L^2 d}{64\sigma n} = \frac{2(1+e)d}{e} + \frac{d}{2048e} < 3d,$$

and, for any  $k \geq 0$ , it holds that

$$128(1+e) 2^k \lambda Lr \sqrt{\frac{d}{n}} - 4^{k-2} \sigma \lambda r^2 \leq \frac{\lambda L^2 d}{\sigma n} (128(1+e) 2^k - 4^{k-2}).$$

Thus, we conclude that

$$\begin{aligned} & \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}} \left[ P_n \varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \\ & \leq e^{3d} + \sum_{k=0}^{\infty} \exp \left\{ \frac{\lambda L^2 d}{\sigma n} (128(1+e) 2^k - 4^{k-2}) \right\}, \end{aligned} \quad (11)$$

where  $\lambda = \sigma n / (32eL^2)$ . It remains to bound the sum

$$\sum_{k=0}^{\infty} \exp \left\{ \frac{\lambda L^2 d}{\sigma n} (128(1+e) 2^k - 4^{k-2}) \right\}.$$

Note that, if  $k \geq 14$ , then  $128(1+e) 2^k < 2^{k+9} \leq 4^{k-2}/2$ . Hence,

$$\begin{aligned} & \sum_{k=0}^{\infty} \exp \left\{ \frac{\lambda L^2 d}{\sigma n} (128(1+e) 2^k - 4^{k-2}) \right\} \\ & \leq \sum_{k=0}^{13} \exp \left\{ \frac{\lambda L^2 d}{\sigma n} (128(1+e) 2^k - 4^{k-2}) \right\} + \sum_{k=14}^{\infty} \exp \left\{ -\frac{4^k \lambda L^2 d}{32\sigma n} \right\}. \end{aligned} \quad (12)$$

Applying the standard bound for the quadratic function

$$128(1+e)2^k - 4^{k-2} = \left(128(1+e)x - \frac{x^2}{16}\right) \Big|_{x=2^k} \leq 256^2(1+e)^2,$$

we deduce that the first term in the right-hand side does not exceed

$$14 \exp \left\{ \frac{256^2(1+e)^2 \lambda L^2 d}{\sigma n} \right\} = 14 \exp \left\{ \frac{2048(1+e)^2 d}{e} \right\}. \quad (13)$$

For the second one, we have

$$\begin{aligned} \sum_{k=14}^{\infty} \exp \left\{ -\frac{4^k \lambda L^2 d}{32 \sigma n} \right\} &= \sum_{k=0}^{\infty} \exp \left\{ -\frac{4^{k+14} d}{32^2 e} \right\} \leq \sum_{k=0}^{\infty} \exp \left\{ -\frac{4^{14} k d}{32^2 e} \right\} \\ &\leq \sum_{k=0}^{\infty} \exp \left\{ -\frac{4^9 k d}{e} \right\} \leq e. \end{aligned} \quad (14)$$

The inequalities (11), (12), (13), and (14) imply that

$$\mathbb{E}_{\varepsilon} \exp \left\{ \lambda \sup_{w \in \mathcal{W}} \left[ P_n \varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_{\mathbb{H}}^2 \right] \right\} \leq e + e^{3d} + 14e^{2048(1+e)^2 d/e}.$$

The claim follows. ■

### 5.1. Proof of Lemma 5.1

We start with an upper bound on the expected value of the supremum. Let  $\mathcal{F}[0, r] = \{f(w, z) - f(w^*, z) : w \in \mathcal{W}[0, r]\}$ . According to Assumption 3.1, we have

$$P_n(f(u, Z) - f(v, Z))^2 \leq L^2 \|u - v\|_{\mathbb{H}}^2 \quad \text{for all } u, v \in \mathcal{W}.$$

This yields that, for any  $\gamma > 0$ , a  $(\gamma/L)$ -net in  $\mathcal{W}[0, r]$  with respect to the seminorm  $\|\cdot\|_{\mathbb{H}}$  induces a  $\gamma$ -net in  $\mathcal{F}[0, r]$  with respect to the empirical  $L_2$ -norm. Indeed, let  $\mathcal{W}_{\gamma/L}[0, r]$  be a  $(\gamma/L)$ -net in  $\mathcal{W}[0, r]$ . Consider any  $g \in \mathcal{F}[0, r]$ . By the definition of  $\mathcal{F}[0, r]$ , there exists  $w \in \mathcal{W}[0, r]$ , such that  $g(z) = f(w, z)$ . Let  $w_{\gamma/L}$  be the closest to  $w$  element of  $\mathcal{W}_{\gamma/L}[0, r]$ . By the definition of  $\mathcal{W}_{\gamma/L}[0, r]$ , we have  $\|w - w_{\gamma/L}\|_{\mathbb{H}} \leq \gamma/L$ . Then, it holds that

$$\begin{aligned} &\left| \sqrt{P_n(f(w, Z) - f(w^*, Z))^2} - \sqrt{P_n(f(w_{\gamma/L}, Z) - f(w^*, Z))^2} \right| \\ &\leq \sqrt{P_n(f(w, Z) - f(w_{\gamma/L}, Z))^2} \leq L \|w - w_{\gamma/L}\|_{\mathbb{H}} \leq \gamma. \end{aligned}$$

Hence, the covering number of  $\mathcal{F}[0, r]$  with respect to the empirical  $L_2$ -norm does not exceed the one of  $\mathcal{W}[0, r]$  with respect to the seminorm  $\|\cdot\|_{\mathbb{H}}$ . Applying the standard chaining technique, we obtain that

$$\mathbb{E}_{\varepsilon} \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon(f(w, Z) - f(w^*, Z)) \leq \frac{12}{\sqrt{n}} \int_0^{Lr} \sqrt{\log \mathcal{N}(\mathcal{W}[0, r], \|\cdot\|_{\mathbb{H}}, \gamma/L)} d\gamma. \quad (15)$$

We use the following lemma to bound the covering number of  $\mathcal{W}[0, r]$ .

**Lemma 5.2** For any  $u \in (0, r]$ , it holds that  $\mathcal{N}(\mathcal{W}[0, r], \|\cdot\|_{\mathbb{H}}, u) \leq (6r/u)^d$ .

The proof of Lemma 5.2 is moved to Section A.4. Substituting the bound of Lemma 5.2 into the Dudley integral (15), we obtain that

$$\mathbb{E}_\varepsilon \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon(f(w, Z) - f(w^*, Z)) \leq \frac{12}{\sqrt{n}} \int_0^{Lr} \sqrt{d \log(6Lr/\gamma)} d\gamma.$$

Substituting  $\gamma$  with  $6Lre^{-u}$ , we get

$$\int_0^{Lr} \sqrt{\log(6Lr/\gamma)} d\gamma = 6Lr \int_{\log 6}^{+\infty} \sqrt{ue^{-u}} du \leq 6Lr \int_0^{+\infty} \sqrt{ue^{-u}} du = 6Lr \Gamma(3/2) = 3Lr\sqrt{\pi}.$$

Therefore, we have

$$\mathbb{E}_\varepsilon \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon(f(w, Z) - f(w^*, Z)) \leq 12\sqrt{\frac{d}{n}} \cdot 3Lr\sqrt{\pi} \leq 64Lr\sqrt{\frac{d}{n}}. \quad (16)$$

Since  $|\varepsilon(f(w, Z) - f(w^*, Z))| \leq B$  almost surely, we can apply Talagrand's concentration inequality for the supremum of the empirical process to

$$\sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon(f(w, Z) - f(w^*, Z)).$$

We use an upper bound on the exponential moment of the supremum, which follows from (Bousquet, 2002, Theorems 2.1 and 2.3):

$$\begin{aligned} & \log \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, r]} \sum_{i=1}^n \varepsilon_i(f(w, Z_i) - f(w^*, Z_i)) \right\} \\ & \leq \lambda \mathbb{E}_\varepsilon \sup_{w \in \mathcal{W}[0, r]} \sum_{i=1}^n \varepsilon_i(f(w, Z_i) - f(w^*, Z_i)) + V_n(e^{B\lambda} - B\lambda - 1) \quad \text{for all } \lambda \geq 0, \end{aligned} \quad (17)$$

where

$$\begin{aligned} V_n &= \frac{2}{B} \mathbb{E}_\varepsilon \sup_{w \in \mathcal{W}[0, r]} \sum_{i=1}^n \varepsilon_i(f(w, Z) - f(w^*, Z)) + \frac{1}{B^2} \sup_{w \in \mathcal{W}[0, r]} \sum_{i=1}^n (f(w, Z) - f(w^*, Z))^2 \\ &\leq \frac{2}{B} \mathbb{E}_\varepsilon \sup_{w \in \mathcal{W}[0, r]} \sum_{i=1}^n \varepsilon_i(f(w, Z) - f(w^*, Z)) + \frac{nL^2r^2}{B^2}. \end{aligned} \quad (18)$$

Combining the inequalities (16), (17), (18) with the bound  $e^x - x - 1 \leq \frac{x^2 e^x}{2}$ , which holds for all  $x \geq 0$ , we obtain for any  $\lambda \geq 0$ ,

$$\begin{aligned} & \log \mathbb{E}_\varepsilon \exp \left\{ \lambda \sup_{w \in \mathcal{W}[0, r]} P_n \varepsilon(f(w, Z) - f(w^*, Z)) \right\} \\ & \leq 64\lambda Lr \sqrt{\frac{d}{n}} + \frac{B^2 \lambda^2 e^{B\lambda/n}}{2n} \left( \frac{128Lr}{B} \sqrt{\frac{d}{n}} + \frac{L^2 r^2}{B^2} \right). \end{aligned}$$

If  $\lambda \leq n/B$ , then

$$64\lambda Lr\sqrt{\frac{d}{n}} + \frac{B^2\lambda^2 e^{B\lambda/n}}{2n} \left( \frac{128Lr}{B}\sqrt{\frac{d}{n}} + \frac{L^2r^2}{B^2} \right) \leq 64(1+e)\lambda Lr\sqrt{\frac{d}{n}} + \frac{e\lambda^2 L^2r^2}{2n}.$$

The claim follows. ■

## 6. Conclusion

In this work, we have established a tight high-probability upper bound on the excess risk of Empirical Risk Minimization (ERM), specifically  $O(d/n + \log(1/\delta)/n)$ , for a subclass of exp-concave losses. The key aspect of our analysis is the convexity of the domain  $\mathcal{W}$ , though it does not necessitate boundedness. Nevertheless, it is important for the loss function  $f(w, Z_i)$  to remain bounded for every  $w \in \mathcal{W}$  and all  $i \in \{1, \dots, n\}$ . This assumption proves instrumental in Lemma 5.1, when we are applying Talagrand’s inequality for the suprema of empirical processes. The counterpart to Talagrand’s inequality in the unbounded case, as indicated in (Adamczak, 2008), might introduce an extra  $\log n$  factor. Currently, it remains unclear whether a similar  $O(d/n + \log(1/\delta)/n)$  bound can be achieved in the unbounded scenario. While applications like logistic regression can mitigate the unbounded case to the bounded one with clipping (see (Foster et al., 2018)), this results in non-convex classes, making our localized analysis inapplicable. Moreover, as noted in Section 3, the additional  $\log n$  factor is sometimes unavoidable when dealing with a non-convex reference set (see (Rakhlin et al., 2017, Theorem 6)). Understanding the geometry of non-convex classes, particularly in relation to proper and improper learners, appears to be a promising direction for future research.

## Acknowledgments

The article was prepared within the framework of the HSE University Basic Research Program. Nikita Puchkin is a Young Russian Mathematics award winner and would like to thank its sponsors and jury.

## References

- Radoslaw Adamczak. A tail inequality for suprema of unbounded empirical processes with applications to Markov chains. *Electronic Journal of Probability*, 13:1000 – 1034, 2008.
- Jean-Yves Audibert and Olivier Catoni. Linear regression through PAC-Bayesian truncation. *arXiv preprint arXiv:1010.0072*, 2010.
- Jean-Yves Audibert and Olivier Catoni. Robust linear least squares regression. *The Annals of Statistics*, 39(5):2766–2794, 2011.
- Francis Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate  $O(1/n)$ . *Advances in Neural Information Processing Systems*, 26, 2013.
- Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.

- Blair Bilodeau, Dylan Foster, and Daniel Roy. Tight bounds on minimax regret under logarithmic loss via self-concordance. In *International Conference on Machine Learning*, pages 919–929, 2020.
- Blair Bilodeau, Dylan Foster, and Daniel Roy. Minimax rates for conditional density estimation via empirical entropy. *arXiv preprint arXiv:2109.10461*, 2021.
- Lucien Birgé. Approximation dans les espaces métriques et théorie de l’estimation. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65:181–237, 1983.
- Lucien Birgé and Pascal Massart. Rates of convergence for minimum contrast estimators. *Probability Theory and Related Fields*, 97:113–150, 1993.
- Olivier Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *Comptes Rendus Mathématique. Académie des Sciences. Paris*, 334(6):495–500, 2002.
- Olivier Bousquet, Yegor Klochkov, and Nikita Zhivotovskiy. Sharper bounds for uniformly stable algorithms. In *Conference on Learning Theory*, pages 610–626, 2020.
- Jürgen Forster and Manfred K Warmuth. Relative expected instantaneous loss bounds. *Journal of Computer and System Sciences*, 64(1):76–102, 2002.
- Dylan J Foster, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan. Logistic regression: The importance of being improper. In *Conference On Learning Theory*, pages 167–208. PMLR, 2018.
- Pierre Gaillard and Olivier Wintenberger. Efficient online algorithms for fast-rate regret bounds under sparsity. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Alon Gonen and Shai Shalev-Shwartz. Average stability is invariant to data preconditioning. Implications to exp-concave empirical risk minimization. *Journal of Machine Learning Research*, 18(222):1–13, 2018.
- Elad Hazan, Tomer Koren, and Kfir Y. Levy. Logistic regression: Tight bounds for stochastic and online optimization. In *Proceedings of The 27th Conference on Learning Theory*, volume 35, pages 197–209, 2014.
- Sham M Kakade and Andrew Ng. Online bounds for bayesian algorithms. *Advances in Neural Information Processing Systems*, 17, 2004.
- Sham M. Kakade and Ambuj Tewari. On the generalization ability of online strongly convex programming algorithms. In *Advances in Neural Information Processing Systems*, volume 21, 2008.
- Sham M. Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in Neural Information Processing Systems*, volume 21, 2008.
- Varun Kanade, Patrick Rebeschini, and Tomas Vaškevičius. Exponential tail local Rademacher complexity risk bounds without the Bernstein condition. *arXiv preprint arXiv:2202.11461*, 2022.

- 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.
- Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593 – 2656, 2006.
- Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: Ecole d’Eté de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.
- Tomer Koren and Kfir Levy. Fast rates for exp-concave empirical risk minimization. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- Lucien LeCam. Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, pages 38–53, 1973.
- Guillaume Lecué and Philippe Rigollet. Optimal learning with  $Q$ -aggregation. *The Annals of Statistics*, 42(1):211 – 224, 2014.
- Tengyuan Liang, Alexander Rakhlin, and Karthik Sridharan. Learning with square loss: Localization through offset Rademacher complexity. In *Proceedings of The 28th Conference on Learning Theory*, volume 40, pages 1260–1285, 2015.
- Mehrdad Mahdavi, Lijun Zhang, and Rong Jin. Lower and upper bounds on the generalization of stochastic exponentially concave optimization. In *Conference on Learning Theory*, pages 1305–1320, 2015.
- Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- Nishant Mehta. Fast rates with high probability in exp-concave statistical learning. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 1085–1093, 2017.
- Shahar Mendelson. Geometric parameters of kernel machines. In *International Conference on Computational Learning Theory*, pages 29–43. Springer, 2002.
- Eric Moulines and Francis Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, volume 24, 2011.
- Jaouad Mourtada and Stéphane Gaïffas. An improper estimator with optimal excess risk in misspecified density estimation and logistic regression. *Journal of Machine Learning Research*, 23(31):1–49, 2022.
- Jaouad Mourtada, Tomas Vaškevičius, and Nikita Zhivotovskiy. Distribution-free robust linear regression. *Mathematical Statistics and Learning*, 4(3-4):253–292, 2021.
- Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning*. Cambridge University Press, 2023.



- Alexander Rakhlin and Karthik Sridharan. Sequential probability assignment with binary alphabets and large classes of experts. *arXiv preprint arXiv:1501.07340*, 2015.
- Alexander Rakhlin, Karthik Sridharan, and Alexandre B. Tsybakov. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23(2):789–824, 2017. ISSN 1350-7265.
- Ohad Shamir. The sample complexity of learning linear predictors with the squared loss. *Journal of Machine Learning Research*, 16(108):3475–3486, 2015.
- Karthik Sridharan, Shai Shalev-Shwartz, and Nathan Srebro. Fast rates for regularized objectives. In *Advances in Neural Information Processing Systems*, volume 21, 2008.
- Sara Van de Geer. Estimating a regression function. *The Annals of Statistics*, pages 907–924, 1990.
- Sara Van de Geer. *Empirical Processes in M-estimation*. Cambridge University Press, 2000.
- Dirk Van der Hoeven, Nikita Zhivotovskiy, and Nicolò Cesa-Bianchi. A regret-variance trade-off in online learning. *arXiv preprint arXiv:2206.02656*, 2022.
- Tomas Vaškevičius and Nikita Zhivotovskiy. Suboptimality of constrained least squares and improvements via non-linear predictors. *Bernoulli*, 29(1):473–495, 2023.
- Suhas Vijaykumar. Localization, convexity, and star aggregation. In *Advances in Neural Information Processing Systems*, volume 34, pages 4570–4581, 2021.
- Marten Wegkamp. Model selection in nonparametric regression. *The Annals of Statistics*, 31(1):252–273, 2003.
- Blake E. Woodworth and Nathan Srebro. An even more optimal stochastic optimization algorithm: Minibatching and interpolation learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 7333–7345, 2021.
- Tianbao Yang, Zhe Li, and Lijun Zhang. A simple analysis for exp-concave empirical minimization with arbitrary convex regularizer. In *International Conference on Artificial Intelligence and Statistics*, pages 445–453, 2018.
- Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, pages 1564–1599, 1999.
- Lijun Zhang and Zhi-Hua Zhou. Stochastic approximation of smooth and strongly convex functions: Beyond the  $O(1/t)$  convergence rate. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99, pages 3160–3179, 2019.
- Tong Zhang. Information-theoretic upper and lower bounds for statistical estimation. *IEEE Transactions on Information Theory*, 52(4):1307–1321, 2006.
- Nikita Zhivotovskiy and Steve Hanneke. Localization of VC classes: Beyond local Rademacher complexities. *Theoretical Computer Science*, 742:27–49, 2018.

## Appendix A. Proofs of auxiliary results

### A.1. Proof of Proposition 3.2

It is enough to check that

$$\nabla^2 e^{-\sigma \widehat{F}(w)/L^2} \preceq \mathbf{O}$$

for all  $w$  from the interior of  $\mathcal{W}$ . Here  $\mathbf{O}$  is a matrix with zero entries. This is equivalent to

$$\frac{\sigma}{L^2} \nabla \widehat{F}(w) \nabla \widehat{F}(w)^\top \preceq \nabla^2 \widehat{F}(w) \quad \text{for all } w \in \text{Int}(\mathcal{W}).$$

The strong convexity of  $\widehat{F}$  with respect to the seminorm  $\|\cdot\|_{\mathbf{H}}$  implies that the function  $\widehat{F}(w) - \sigma \|w\|_{\mathbf{H}}^2$  is convex on  $\mathbf{H}$ . The last is equivalent to the inequality

$$\nabla^2 \left( \widehat{F}(w) - \sigma \|w\|_{\mathbf{H}}^2 \right) \succeq \mathbf{O} \quad \text{for all } w \in \text{Int}(\mathcal{W})$$

or simply  $\nabla^2 \widehat{F}(w) \succeq \sigma \mathbf{H}$  for all  $w \in \text{Int}(\mathcal{W})$ . On the other hand, it holds that

$$\left( \widehat{F}(u) - \widehat{F}(v) \right)^2 \leq P_n (f(u, Z) - f(v, Z))^2 \leq L^2 \|u - v\|_{\mathbf{H}}^2 \quad \text{for all } u, v \in \mathcal{W}.$$

This yields that

$$\nabla \widehat{F}(w) \nabla \widehat{F}(w)^\top \preceq L^2 \mathbf{H} \quad \text{for all } w \in \text{Int}(\mathcal{W}).$$

Hence, for any  $w$  in the interior of  $\mathcal{W}$ , we have

$$\frac{\sigma}{L^2} \nabla \widehat{F}(w) \nabla \widehat{F}(w)^\top \preceq \sigma \mathbf{H} \preceq \nabla^2 \widehat{F}(w).$$

The claim follows. ■

### A.2. Proof of Proposition 3.3

If Assumption 3.1 holds, then Jensen's inequality immediately implies that

$$\left( \widehat{F}(u) - \widehat{F}(v) \right)^2 \leq P_n (f(u, Z) - f(v, Z))^2 \leq L^2 \|u - v\|_{\mathbf{H}}^2 \quad \text{for all } u, v \in \mathcal{W}.$$

Thus,  $\widehat{F}$  is  $L$ -Lipschitz on  $\mathcal{W}$  with respect to the seminorm  $\|\cdot\|_{\mathbf{H}}$ . On the other hand, the strong convexity of  $\widehat{F}$  yields that

$$\widehat{F}(\widehat{w}) \leq \widehat{F} \left( \frac{\widehat{w} + w}{2} \right) \leq \frac{1}{2} \widehat{F}(\widehat{w}) + \frac{1}{2} \widehat{F}(w) - \frac{\sigma}{8} \|\widehat{w} - w\|_{\mathbf{H}}^2 \quad \text{for all } w \in \mathcal{W},$$

where the first inequality follows from the convexity of  $\mathcal{W}$  and the definition of  $\widehat{w}$ . Hence,

$$\widehat{F}(w) \geq \widehat{F}(\widehat{w}) + \frac{\sigma}{4} \|\widehat{w} - w\|_{\mathbf{H}}^2 \quad \text{for all } w \in \mathcal{W}.$$

As a result, we obtain that

$$\left( \widehat{F}(w) - \widehat{F}(\widehat{w}) \right)^2 \leq L^2 \|\widehat{w} - w\|_{\mathbf{H}}^2 \leq \frac{4L^2}{\sigma} \left( \widehat{F}(w) - \widehat{F}(\widehat{w}) \right) \quad \text{for all } w \in \mathcal{W}.$$

Consequently,

$$\left| \widehat{F}(w) - \widehat{F}(w^*) \right| \leq \max \left\{ \widehat{F}(w) - \widehat{F}(\widehat{w}), \widehat{F}(w^*) - \widehat{F}(\widehat{w}) \right\} \leq \frac{4L^2}{\sigma} \quad \text{for all } w \in \mathcal{W}. \quad (19)$$

We show that  $|f(w, Z) - f(w^*, Z)| \leq 4L^2/\sigma$  almost surely. Assume the opposite. Let  $Z_{>} \subseteq Z$  be a set of all such  $z \in Z$  that  $f(w, z) - f(w^*, z) > 4L^2/\sigma$ . If  $Z_{>}$  has a positive probability measure, then we have  $\widehat{F}(w) - \widehat{F}(w^*) > 4L^2/\sigma$  on  $Z_{>}^{\otimes n}$ , which contradicts (19). Hence,  $f(w, Z) - f(w^*, Z) \leq 4L^2/\sigma$  almost surely. Similarly, we can prove that  $f(w^*, Z) - f(w, Z) \leq 4L^2/\sigma$  almost surely.  $\blacksquare$

### A.3. Proof of Lemma 4.1

Note that

$$\begin{aligned} F(\widehat{w}) - F(w^*) &= 2(P - P_n)(f(\widehat{w}, Z) - f(w^*, Z)) + P(f(w^*, Z) - f(\widehat{w}, Z)) \\ &\quad + 2P_n(f(\widehat{w}, Z) - f(w^*, Z)). \end{aligned}$$

According to Assumption 3.1,  $\widehat{F}$  is strongly convex with respect to the seminorm, induced by the matrix  $H$ . Taking into account the convexity of  $\mathcal{W}$  and the equation (4), we obtain that

$$\widehat{F}(\widehat{w}) \leq \widehat{F}\left(\frac{\widehat{w} + w}{2}\right) \leq \frac{1}{2}\widehat{F}(\widehat{w}) + \frac{1}{2}\widehat{F}(w) - \frac{\sigma}{8}\|\widehat{w} - w\|_H^2 \quad \text{for all } w \in \mathcal{W}.$$

Therefore, it holds that

$$\widehat{F}(w) \geq \widehat{F}(\widehat{w}) + \frac{\sigma}{4}\|\widehat{w} - w\|_H^2 \quad \text{for all } w \in \mathcal{W}. \quad (20)$$

Similarly, we deduce that the following inequality holds for all  $w \in \mathcal{W}$ :

$$F(w^*) = \mathbb{E}\widehat{F}(w^*) \leq \mathbb{E}\widehat{F}\left(\frac{w^* + w}{2}\right) \leq \frac{1}{2}\mathbb{E}\widehat{F}(w^*) + \frac{1}{2}\mathbb{E}\widehat{F}(w) - \frac{\sigma}{8}\mathbb{E}\|w^* - w\|_H^2.$$

Consequently, we have

$$F(w) \geq F(w^*) + \frac{\sigma}{4}\mathbb{E}\|w^* - w\|_H^2 \quad \text{for all } w \in \mathcal{W}. \quad (21)$$

Observe that  $\mathbb{E}\|w^* - w\|_H^2 = \|w^* - w\|_{\mathbb{E}H}^2$ . Then the inequalities (20) and (21) imply that

$$P(f(w^*, Z) - f(\widehat{w}, Z)) \leq -\frac{\sigma}{4}\|\widehat{w} - w^*\|_{\mathbb{E}H}^2 \quad \text{and} \quad P_n(f(\widehat{w}, Z) - f(w^*, Z)) \leq -\frac{\sigma}{4}\|\widehat{w} - w^*\|_{\mathbb{E}H}^2.$$

Thus, it holds that

$$F(\widehat{w}) - F(w^*) \leq 2(P - P_n)(f(\widehat{w}, Z) - f(w^*, Z)) - \frac{\sigma}{4}\|\widehat{w} - w^*\|_{\mathbb{E}H}^2 - \frac{\sigma}{2}\|\widehat{w} - w^*\|_{\mathbb{E}H}^2.$$

The last expression does not exceed

$$2 \sup_{w \in \mathcal{W}} \left[ (P - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{\mathbb{E}H}^2 + \|w - w^*\|_H^2) \right].$$

We are ready to apply the symmetrization argument. Let  $Z'_1, \dots, Z'_n$  be independent copies of  $Z_1, \dots, Z_n$ , and let  $H' = H(Z'_1, \dots, Z'_n)$  be the corresponding matrix from Assumption 3.1. Then

$$\begin{aligned} & \mathbb{E}\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ (P - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{\mathbb{E}H}^2 + \|w - w^*\|_H^2) \right]\right) \\ &= \mathbb{E}\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ \mathbb{E}'(P'_n - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\mathbb{E}'\|w - w^*\|_{H'}^2 + \|w - w^*\|_H^2) \right]\right), \end{aligned}$$

where  $P'_n$  denotes the expectation with respect to the uniform measure on  $Z'_1, \dots, Z'_n$  and  $\mathbb{E}'$  stands for the expectation with respect to  $Z'_1, \dots, Z'_n$ . Using Jensen's inequality, we obtain that

$$\begin{aligned} & \mathbb{E}\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ (P - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{\mathbb{E}H}^2 + \|w - w^*\|_H^2) \right]\right) \\ & \leq \mathbb{E}\mathbb{E}'\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ (P'_n - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{H'}^2 + \|w - w^*\|_H^2) \right]\right). \end{aligned}$$

Let  $\varepsilon_1, \dots, \varepsilon_n$  be i.i.d. Rademacher random variables. Note that, whatever the realization of  $\varepsilon_1, \dots, \varepsilon_n$  is,

$$(P'_n - P_n)\varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{H'}^2 + \|w - w^*\|_H^2)$$

has the same distribution as

$$(P'_n - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{H'}^2 + \|w - w^*\|_H^2).$$

Thus,

$$\begin{aligned} & \mathbb{E}\mathbb{E}'\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ (P'_n - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{H'}^2 + \|w - w^*\|_H^2) \right]\right) \\ &= \mathbb{E}\mathbb{E}'\mathbb{E}_\varepsilon\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ (P'_n - P_n)\varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{H'}^2 + \|w - w^*\|_H^2) \right]\right) \\ & \leq \mathbb{E}\mathbb{E}_\varepsilon\Phi\left(4 \sup_{w \in \mathcal{W}} \left[ P_n\varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_H^2 \right]\right), \end{aligned}$$

where the last line is due to the triangle inequality. Hence, we conclude that

$$\begin{aligned} & \mathbb{E}\Phi(F(\hat{w}) - F(w^*)) \\ & \leq \mathbb{E}\mathbb{E}_\varepsilon\Phi\left(2 \sup_{w \in \mathcal{W}} \left[ (P - P_n)(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} (\|w - w^*\|_{\mathbb{E}H}^2 + \|w - w^*\|_H^2) \right]\right) \\ & \leq \mathbb{E}\mathbb{E}_\varepsilon\Phi\left(4 \sup_{w \in \mathcal{W}} \left[ P_n\varepsilon(f(w, Z) - f(w^*, Z)) - \frac{\sigma}{8} \|w - w^*\|_H^2 \right]\right). \end{aligned}$$

The claim follows. ■

#### A.4. Proof of Lemma 5.2

Let  $H = \sum_{j=1}^{\text{rank}(H)} \lambda_j v_j v_j^\top$  be the eigendecomposition of  $H$ . Here we assume that  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{\text{rank}(H)} > 0$ , and the vectors  $v_1, \dots, v_{\text{rank}(H)}$  form an orthonormal system. Define the Moore-Penrose pseudoinverse of  $H$

$$H^\dagger = \sum_{j=1}^{\text{rank}(H)} \lambda_j^{-1} v_j v_j^\top$$

and the square roots

$$H^{1/2} = \sum_{j=1}^{\text{rank}(H)} \sqrt{\lambda_j} v_j v_j^\top \quad \text{and} \quad (H^\dagger)^{1/2} = \sum_{j=1}^{\text{rank}(H)} \frac{1}{\sqrt{\lambda_j}} v_j v_j^\top.$$

It is easy to observe that

$$\Pi = H^{1/2} (H^\dagger)^{1/2} = \sum_{j=1}^{\text{rank}(H)} v_j v_j^\top.$$

is nothing but the orthogonal projector onto the image of  $H^{1/2}$ . Let  $\{a_1, \dots, a_N\}$  be the smallest  $u$ -net of the Euclidean ball  $\mathcal{B}(0, r) \subset \mathbb{R}^d$  with respect to the Euclidean norm. It is known that  $N \leq (6r/u)^d$ . We show that  $\{w^* + (H^\dagger)^{1/2} a_1, \dots, w^* + (H^\dagger)^{1/2} a_N\}$  is a  $u$ -net of  $\mathcal{W}[0, r]$  with respect to the seminorm  $\|\cdot\|_H$ . For this purpose, fix any  $w \in \mathcal{W}[0, r]$  and let

$$a = H^{1/2}(w - w^*).$$

By the definition of  $\Pi$  and  $a$ ,  $\Pi a = a$ . Find the closest to  $a$  element amongst  $\{a_1, \dots, a_N\}$  (with respect to the Euclidean distance). Without loss of generality, assume that it is  $a_1$ . Denote  $w^* + (H^\dagger)^{1/2} a_1$  by  $w_1$ . Then, it holds that

$$\begin{aligned} \|w - w_1\|_H &= \|H^{1/2}(w - w_1)\| = \|a - H^{1/2}(w_1 - w^*)\| \\ &= \left\| a - H^{1/2} (H^\dagger)^{1/2} a_1 \right\| = \|a - \Pi a_1\| = \|\Pi a - \Pi a_1\| \leq \|a - a_1\| \leq u. \end{aligned}$$

Hence,  $\{w^* + (H^\dagger)^{1/2} a_1, \dots, w^* + (H^\dagger)^{1/2} a_N\}$  is a  $u$ -net of  $\mathcal{W}[0, r]$  with respect to  $\|\cdot\|_H$ . The claim follows. ■