



Robust statistical learning with Lipschitz and convex loss functions

Geoffrey Chinot¹ · Guillaume Lecué¹  · Matthieu Lerasle¹

Received: 3 April 2019 / Revised: 17 June 2019
© Springer-Verlag GmbH Germany, part of Springer Nature 2019

Abstract

We obtain estimation and excess risk bounds for Empirical Risk Minimizers (ERM) and minmax Median-Of-Means (MOM) estimators based on loss functions that are both Lipschitz and convex. Results for the ERM are derived under weak assumptions on the outputs and subgaussian assumptions on the design as in Alquier et al. (Estimation bounds and sharp oracle inequalities of regularized procedures with Lipschitz loss functions. [arXiv:1702.01402](https://arxiv.org/abs/1702.01402), 2017). The difference with Alquier et al. (2017) is that the global Bernstein condition of this paper is relaxed here into a local assumption. We also obtain estimation and excess risk bounds for minmax MOM estimators under similar assumptions on the output and only moment assumptions on the design. Moreover, the dataset may also contain outliers in both inputs and outputs variables without deteriorating the performance of the minmax MOM estimators. Unlike alternatives based on MOM's principle (Lecué and Lerasle in *Ann Stat*, 2017; Lugosi and Mendelson in *JEMS*, 2016), the analysis of minmax MOM estimators is not based on the small ball assumption (SBA) of Koltchinskii and Mendelson (*Int Math Res Not IMRN* 23:12991–13008, 2015). In particular, the basic example of non parametric statistics where the learning class is the linear span of localized bases, that does not satisfy SBA (Saumard in *Bernoulli* 24(3):2176–2203, 2018) can now be handled. Finally, minmax MOM estimators are analysed in a setting where the local Bernstein condition is also dropped out. It is shown to achieve excess risk bounds with exponentially large probability under minimal assumptions insuring only the existence of all objects.

Mathematics Subject Classification 62G35

✉ Guillaume Lecué
lecueguillaume@gmail.com

Geoffrey Chinot
geoffrechinot@gmail.com

Matthieu Lerasle
lerasle@gmail.com

¹ Palaiseau, France

1 Introduction

In this paper, we study learning problems where the loss function is simultaneously Lipschitz and convex. This situation happens in classical examples such as quantile, Huber and L_1 regression or logistic and hinge classification [42]. As the Lipschitz property allows to make only weak assumptions on the outputs, these losses have been quite popular in robust statistics [17]. Empirical risk minimizers (ERM) based on Lipschitz losses such as the Huber loss have received recently an important attention [2,15,45].

Based on a dataset $\{(X_i, Y_i) : i = 1, \dots, N\}$ of points in $\mathcal{X} \times \mathcal{Y}$, a class F of functions and a risk function $R(\cdot)$ defined on F , the statistician want to estimate an oracle $f^* \in \operatorname{argmin}_{f \in F} R(f)$ or to predict an output Y at least as good as $f^*(X)$. The risk function $R(\cdot)$ is often defined as the expectation of a loss function $\ell : (f, x, y) \in F \times \mathcal{X} \times \mathcal{Y} \rightarrow \ell_f(x, y) \in \mathbb{R}$ with respect to the unknown distribution P of a random variable $(X, Y) \in \mathcal{X} \times \mathcal{Y}$: $R(f) = \mathbb{E} \ell_f(X, Y)$. Hereafter, the risk is assumed to have this form for a loss function ℓ such that, for any (f, x, y) , $\ell_f(x, y) = \bar{\ell}(f(x), y)$, for some function $\bar{\ell} : \bar{\mathcal{Y}} \times \mathcal{Y} \rightarrow \mathbb{R}$, where the set $\bar{\mathcal{Y}}$ is a convex set containing all possible values of $f(x)$. The loss function ℓ is said Lipschitz and convex when the following assumption holds.

Assumption 1 There exists $L > 0$ such that, for any $y \in \mathcal{Y}$, $\bar{\ell}(\cdot, y)$ is L -Lipschitz and convex.

Many classical loss functions satisfy Assumption 1 and we recall some of them below.

- The **logistic loss** defined, for any $u \in \bar{\mathcal{Y}} = \mathbb{R}$ and $y \in \mathcal{Y} = \{-1, 1\}$, by $\bar{\ell}(u, y) = \log(1 + \exp(-yu))$ satisfies Assumption 1 with $L = 1$.
- The **hinge loss** defined, for any $u \in \bar{\mathcal{Y}} = \mathbb{R}$ and $y \in \mathcal{Y} = \{-1, 1\}$, by $\bar{\ell}(u, y) = \max(1 - uy, 0)$ satisfies Assumption 1 with $L = 1$.
- The **Huber loss** defined, for any $\delta > 0$, $u, y \in \mathcal{Y} = \bar{\mathcal{Y}} = \mathbb{R}$, by

$$\bar{\ell}(u, y) = \begin{cases} \frac{1}{2}(y - u)^2 & \text{if } |u - y| \leq \delta \\ \delta|y - u| - \frac{\delta^2}{2} & \text{if } |u - y| > \delta \end{cases},$$

satisfies Assumption 1 with $L = \delta$.

- The **quantile loss** is defined, for any $\tau \in (0, 1)$, $u, y \in \mathcal{Y} = \bar{\mathcal{Y}} = \mathbb{R}$, by $\bar{\ell}(u, y) = \rho_\tau(u - y)$ where, for any $z \in \mathbb{R}$, $\rho_\tau(z) = z(\tau - I\{z \leq 0\})$. It satisfies Assumption 1 with $L = 1$. For $\tau = 1/2$, the quantile loss is the L_1 loss.

All along the paper, the following assumption is also granted.

Assumption 2 The class F is convex.

When (X, Y) and the data $((X_i, Y_i))_{i=1}^N$ are independent and identically distributed (i.i.d.), for any $f \in F$, the empirical risk $R_N(f) = (1/N) \sum_{i=1}^N \ell_f(X_i, Y_i)$ is a natural estimator of $R(f)$. The empirical risk minimizers (ERM) [43] obtained by minimizing $f \in F \rightarrow R_N(f)$ are expected to be close to the oracle f^* . This procedure and its

regularized versions have been extensively studied in learning theory [20]. When the loss is both convex and Lipschitz, results have been obtained in practice [4,12] and theory [42]. Risk bounds with exponential deviation inequalities for the ERM can be obtained under weak assumptions on the outputs Y , but stronger assumptions on the design X . Moreover, fast rates of convergence [41] can only be obtained under margin type assumptions such as the Bernstein condition [8,42].

The Lipschitz assumption and global Bernstein conditions (that hold over the entire F as in [2]) imply boundedness in L_2 -norm of the class F , see the discussion preceding Assumption 4 for details. This boundedness is not satisfied in linear regression with unbounded design so the results of [2] don't apply to this basic example such as linear regression with a Gaussian design. To bypass this restriction, the global condition is relaxed into a "local" one as in [15,42], see Assumption 4 below.

The main constraint in our results on ERM is the assumption on the design. This constraint can be relaxed by considering alternative estimators based on the "median-of-means" (MOM) principle of [1,9,18,37] and the minmax procedure of [3,5]. The resulting minmax MOM estimators have been introduced in [24] for least-squares regression as an alternative to other MOM based procedures [23,29–31]. In the case of convex and Lipschitz loss functions, these estimators satisfy the following properties (1) as the ERM, they are efficient under weak assumptions on the noise, (2) they achieve optimal rates of convergence under weak stochastic assumptions on the design and (3) the rates are not downgraded by the presence of some outliers in the dataset.

These improvements of MOM estimators upon ERM are not surprising. For univariate mean estimation, rate optimal sub-Gaussian deviation bounds can be shown under minimal L_2 moment assumptions for MOM estimators [14] while the empirical mean needs each data to have sub-Gaussian tails to achieve such bounds [13]. In least-squares regression, MOM-based estimators [23,24,29–31] inherit these properties, whereas the ERM has downgraded statistical properties under moment assumptions (see Proposition 1.5 in [25]). Furthermore, MOM procedures are resistant to outliers: results hold in the " $\mathcal{O} \cup \mathcal{I}$ " framework of [23,24], where inliers or informative data (indexed by \mathcal{I}) only satisfy weak moments assumptions and the dataset may contain outliers (indexed by \mathcal{O}) on which no assumption is made, see Sect. 3. This robustness, that almost comes for free from a technical point of view is another important advantage of MOM estimators compared to ERM in practice. Figure 1¹ illustrates this fact, showing that statistical performance of the standard logistic regression are strongly affected by a single corrupted observation, while the minmax MOM estimator maintains good statistical performance even with 5% of corrupted data.

Compared to [24,29], considering convex-Lipschitz losses instead of the square loss allows to simplify simultaneously some assumptions and the presentation of the results for MOM estimators: L_2 -assumptions on the noise in [24,29] can be removed and complexity parameters driving risk of ERM and MOM estimators only involve a single stochastic linear process, see Eqs. (3) and (6) below. Also, contrary to the analysis in least-squares regression, the small ball assumption [22,33] is not required here. Recall that this assumption states that there are absolute constants κ and β such that, for all $f \in F$, $\mathbb{P}[|f(X) - f^*(X)| \geq \kappa \|f - f^*\|_{L_2}] \geq \beta$. It is interesting as it

¹ All figures can be reproduced from the code available at <https://github.com/lecueguillaume/MOMpower>.




Fig. 1 MOM Logistic Regression VS Logistic regression from Sklearn ($p = 50$ and $N = 1000$)

involves only moments of order 1 and 2 of the functions in F . However, it does not hold with absolute constants in classical frameworks such as histograms, see [16,39] and Sect. 5.

Finally, minmax MOM estimators are studied in a framework where the Bernstein condition is dropped out. In this setting, they are shown to achieve an oracle inequality with exponentially large probability (see Sect. 4). The results are slightly weaker in this relaxed setting: the excess risk is bounded but not the L_2 risk and the rates of convergence are “slow” in $1/\sqrt{N}$ in general. Fast rates of convergence in $1/N$ can still be recovered from this general result if a local Bernstein type condition is satisfied though, see Sect. 4 for details. This last result shows that minmax MOM estimators can be safely used with Lipschitz and convex losses, assuming only that inliers data are independent with enough finite moments to give sense to the results.

To approximate minmax MOM estimators, an algorithm inspired from [24,26] is also proposed. Asymptotic convergence of this algorithm has been proved in [26] under strong assumptions, but, to the best of our knowledge, convergence rates have not been established. Nevertheless, the simulation study presented in Sect. 7 shows that it has good robustness performances.

The paper is organized as follows. Optimal results for the ERM are presented in Sect. 2. Minmax MOM estimators are introduced and analysed in Sect. 3 under a local Bernstein condition and in Sect. 4 without the Bernstein condition. A discussion of the main assumptions is provided in Sect. 5. Section 6 presents the theoretical limits of the ERM compared to the minmax MOM estimators. Finally, Sect. 7 provides a simulation study where a natural algorithm associated to the minmax MOM estimator for logistic loss is presented. The proofs of the main theorems are gathered in Sects. A, B and C.

Notations Let \mathcal{X}, \mathcal{Y} be measurable spaces and let $\bar{\mathcal{Y}}$ denote a convex set $\bar{\mathcal{Y}} \supset \mathcal{Y}$. Let F be a class of measurable functions $f : \mathcal{X} \rightarrow \bar{\mathcal{Y}}$ and let $(X, Y) \in \mathcal{X} \times \mathcal{Y}$ be a random variable with distribution P . Let μ denote the marginal distribution of X . For any probability measure Q on $\mathcal{X} \times \mathcal{Y}$, and any function $g \in L_1(Q)$, let $Qg = \int g(x, y)dQ(x, y)$. Let $\ell : F \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}, (f, x, y) \mapsto \ell_f(x, y)$ denote a loss function measuring the error made when predicting y by $f(x)$. It is always assumed that there exists a function $\bar{\ell} : \bar{\mathcal{Y}} \times \mathcal{Y} \rightarrow \mathbb{R}$ such that, for any $(f, x, y) \in F \times \mathcal{X} \times \mathcal{Y}, \bar{\ell}(f(x), y) = \ell_f(x, y)$. Let $R(f) = P\ell_f = \mathbb{E}\ell_f(X, Y)$ for f in F denote the risk and let $\mathcal{L}_f = \ell_f - \ell_{f^*}$ denote the excess loss. If $F \subset L_1(P) := L_1$ and Assumption 1 holds, an equivalent risk can be defined even if $Y \notin L_1$. Actually, for any $f_0 \in F, \ell_f - \ell_{f_0} \in L_1$ so one can define $R(f) = P(\ell_f - \ell_{f_0})$. W.l.o.g. the set of risk minimizers is assumed to be reduced to a singleton $\operatorname{argmin}_{f \in F} R(f) = \{f^*\}$. f^* is called the oracle as $f^*(X)$ provides the prediction of Y with minimal risk among functions in F . For any f and $p > 0$, let $\|f\|_{L_p} = (P|f|^p)^{1/p}$, for any $r \geq 0$, let $rB_{L_2} = \{f \in F : \|f\|_{L_2} \leq r\}$ and $rS_{L_2} = \{f \in F : \|f\|_{L_2} = r\}$. For any set H for which it makes sense, $H + f^* = \{h + f^* \text{ s.t } h \in H\}, H - f^* = \{h - f^* \text{ s.t } h \in H\}$. For any real numbers a, b , we write $a \lesssim b$ when there exists a positive constant c such that $a \leq cb$, when $a \lesssim b$ and $b \lesssim a$, we write $a \asymp b$.

2 ERM in the sub-Gaussian framework

This section studies the ERM, improving some results from [2]. In particular, the global Bernstein condition in [2] is relaxed into a local hypothesis following [42]. All along this section, data $(X_i, Y_i)_{i=1}^N$ are independent and identically distributed with common distribution P . The ERM is defined for $f \in F \rightarrow P_N\ell_f = (1/N) \sum_{i=1}^N \ell_f(X_i, Y_i)$ by

$$\hat{f}^{ERM} = \operatorname{argmin}_{f \in F} P_N\ell_f. \tag{1}$$

The results for the ERM are shown under a sub-Gaussian assumption on the class $F - F$ with respect to the distribution of X . This result is the benchmark for the following minmax MOM estimators.

Definition 1 Let $B \geq 1$. F is called B -sub-Gaussian (with respect to X) when for all $f \in F$ and all $\lambda > 1$

$$\mathbb{E} \exp(\lambda|f(X)|/\|f\|_{L_2}) \leq \exp(\lambda^2 B^2/2).$$

Assumption 3 The class $F - F$ is B -sub-Gaussian with respect to X , where $F - F = \{f_1 - f_2 : f_1, f_2 \in F\}$.

Under this sub-Gaussian assumption, statistical complexity can be measured via Gaussian mean-widths.

Definition 2 Let $H \subset L_2$. Let $(G_h)_{h \in H}$ be the canonical centered Gaussian process indexed by H (in particular, the covariance structure of $(G_h)_{h \in H}$ is given by

$(\mathbb{E}(G_{h_1} - G_{h_2})^2)^{1/2} = (\mathbb{E}(h_1(X) - h_2(X))^2)^{1/2}$ for all $h_1, h_2 \in H$). The **Gaussian mean-width** of H is $w(H) = \mathbb{E} \sup_{h \in H} G_h$.

The complexity parameter driving the performance of \hat{f}^{ERM} is presented in the following definition.

Definition 3 The **complexity parameter** is defined as

$$r_2(\theta) \geq \inf\{r > 0 : 32Lw((F - f^*) \cap rB_{L_2}) \leq \theta r^2 \sqrt{N}\}$$

where $L > 0$ is the Lipschitz constant from Assumption 1.

Let $A > 0$. In [8], the class F is called $(1, A)$ -Bernstein if, for all $f \in F$, $P\mathcal{L}_f^2 \leq AP\mathcal{L}_f$. Under Assumption 1, F is $(1, AL^2)$ -Bernstein if the following stronger assumption is satisfied

$$\|f - f^*\|_{L_2}^2 \leq AP\mathcal{L}_f. \quad (2)$$

This stronger version was used, for example in [2] to study ERM. However, under Assumption 1, Eq (2) implies that

$$\|f - f^*\|_{L_2}^2 \leq AP\mathcal{L}_f \leq AL\|f - f^*\|_{L_1} \leq AL\|f - f^*\|_{L_2}.$$

Therefore, $\|f - f^*\|_{L_2} \leq AL$ for any $f \in F$. The class F is bounded in L^2 -norm, which is restrictive as this assumption is not verified by the class of linear functions for example. To bypass this issue, the following condition is introduced.

Assumption 4 There exists a constant $A > 0$ such that, for all $f \in F$ satisfying $\|f - f^*\|_{L_2} = r_2(1/(2A))$, we have $\|f - f^*\|_{L_2}^2 \leq AP\mathcal{L}_f$.

In Assumption 4, Bernstein condition is granted in a L_2 -sphere centered in f^* only. Outside of this sphere, there is no restriction on the excess loss. From the previous remark, it is clear that we necessarily have $r_2(1/(2A)) \leq AL$ (as long as there exists some $f \in F$ such that $\|f - f^*\|_2 \geq r_2(1/(2A))$). This relaxed assumption is satisfied for many Lipschitz-convex loss functions under moment assumptions and weak assumptions on the noise as it will be checked in Sect. 5. The following theorem is the main result of this section.

Theorem 1 Grant Assumptions 1, 2, 3 and 4, \hat{f}^{ERM} defined in (1) satisfies, with probability larger than

$$1 - 2 \exp\left(-C \frac{Nr_2^2(1/(2A))}{(AL)^2}\right), \quad (3)$$

$$\|\hat{f}^{ERM} - f^*\|_{L_2}^2 \leq r_2^2(1/(2A)) \text{ and } P\mathcal{L}_{\hat{f}^{ERM}} \leq \frac{r_2^2(1/(2A))}{2A}, \quad (4)$$

where C is an absolute constant.

Theorem 1 is proved in Sect. A.1. It shows deviation bounds both in L_2 norm and for the excess risk, which are both minimax optimal as proved in [2]. As in [2], a similar result can be derived if the sub-Gaussian Assumption 3 is replaced by a boundedness in L_∞ assumption. An extension of Theorem 1 can be shown, where Assumption 4 is replaced by the following hypothesis: there exists κ such that for all $f \in F$ in a L_2 -sphere centered in f^* , $\|f - f^*\|_{L_2}^{2\kappa} \leq AP\mathcal{L}_f$. The case $\kappa = 1$ is the most classical and its analysis contains all the ingredients for the study of the general case with any parameter $\kappa \geq 1$. More general Bernstein conditions can also be considered as in [42, Chapter 7]. These extensions are left to the interested reader.

Notice that none of the Assumptions 1, 2, 3 and 4 involve the output Y directly. All assumptions on Y are done through the oracle f^* . Yet, as will become transparent in the applications in Sect. 5, some assumptions on the distributions of Y are required to check the assumptions of Theorem 1. These assumptions are not very restrictive though and Lipschitz losses have been quite popular in robust statistics for this reason.

3 Minmax MOM estimators

This section presents and studies minmax MOM estimators, comparing them to ERM. We relax the sub-Gaussian assumption on the class $F - F$ and the i.i.d assumption on the data $(X_i, Y_i)_{i=1}^N$.

3.1 The estimators

The framework of this section is a relaxed version of the i.i.d. setup considered in Sect. 2. Following [23,24], there exists a partition $\mathcal{O} \cup \mathcal{I}$ of $\{1, \dots, N\}$ in two subsets unknown to the statistician. No assumption is granted on the set of “outliers” $(X_i, Y_i)_{i \in \mathcal{O}}$. “Inliers”, $(X_i, Y_i)_{i \in \mathcal{I}}$, are only assumed to satisfy the following assumption. For all $i \in \mathcal{I}$, (X_i, Y_i) has distribution P_i , X_i has distribution μ_i and for any $p > 0$ and any function g for which it makes sense $\|g\|_{L_p(\mu_i)} = (P_i|g|^p)^{1/p}$.

Assumption 5 $(X_i, Y_i)_{i \in \mathcal{I}}$ are independent and, for any $i \in \mathcal{I}$, $\|f - f^*\|_{L_2} = \|f - f^*\|_{L_2(\mu_i)}$ and $P_i\mathcal{L}_f = P\mathcal{L}_f$.

Assumption 5 holds in the i.i.d case but it covers other situations where informative data $(X_i, Y_i)_{i \in \mathcal{I}}$ may have different distributions. Typically, when F is the class of linear functions on \mathbb{R}^d , $F = \{\langle t, \cdot \rangle, t \in \mathbb{R}^d\}$ and $(X_i)_{i \in \mathcal{I}}$ are vectors with independent coordinates $(X_{i,j})_{j=1, \dots, d}$, then Assumption 5 is met if the coordinates $(X_{i,j})_{i \in \mathcal{I}}$ have the same first and second moments for all $j = 1, \dots, d$.

Recall the definition of MOM estimators of univariate means. Let $(B_k)_{k=1, \dots, K}$ denote a partition of $\{1, \dots, N\}$ into blocks B_k of equal size N/K (if N is not a multiple of K , just remove some data). For any function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ and $k \in \{1, \dots, K\}$, let $P_{B_k}f = (K/N) \sum_{i \in B_k} f(X_i, Y_i)$. MOM estimator is the median of these empirical means:

$$\text{MOM}_K(f) = \text{Med}(P_{B_1}f, \dots, P_{B_K}f).$$

The estimator $\text{MOM}_K(f)$ achieves rate optimal sub-Gaussian deviation bounds, assuming only that $Pf^2 < \infty$, see for example [14]. The number K is a tuning parameter. The larger K , the more outliers are allowed. When $K = 1$, $\text{MOM}_K(f)$ is the empirical mean, when $K = N$, the empirical median.

Following [24], remark that the oracle is also solution of the following minmax problem:

$$f^* \in \underset{f \in F}{\operatorname{argmin}} P\ell_f = \underset{f \in F}{\operatorname{argmin}} \sup_{g \in F} P(\ell_f - \ell_g).$$

Minmax MOM estimators are obtained by plugging MOM estimators of the unknown expectations $P(\ell_f - \ell_g)$ in this formula:

$$\hat{f} \in \underset{f \in F}{\operatorname{argmin}} \sup_{g \in F} \text{MOM}_K(\ell_f - \ell_g). \quad (5)$$

The minmax MOM construction can be applied systematically as an alternative to ERM. For instance, it yields a robust version of logistic classifiers. The minmax MOM estimator with $K = 1$ is the ERM.

The linearity of the empirical process P_N is important to use localisation technics and derive “fast rates” of convergence for ERM [21], improving “slow rates” derived with the approach of [44], see [41] for details on “fast and slow rates”. The idea of the minmax reformulation comes from [3], where this strategy allows to overcome the lack of linearity of some alternative robust mean estimators. [23] introduced minmax MOM estimators to least-squares regression.

3.2 Theoretical results

3.2.1 Setting

The assumptions required for the study of estimator (5) are essentially those of Sect. 2 except for Assumption 3 which is relaxed into Assumption 5. Instead of Gaussian mean width, the complexity parameter is expressed as a fixed point of local Rademacher complexities [6,7,10]. Let $(\sigma_i)_{i=1,\dots,N}$ denote i.i.d. Rademacher random variables (uniformly distributed on $\{-1, 1\}$), independent from $(X_i, Y_i)_{i \in \mathcal{I}}$. Let

$$\tilde{r}_2(\gamma) \geq \inf \left\{ r > 0, \forall J \subset \mathcal{I} : |J| \geq \frac{N}{2}, \right. \\ \left. \mathbb{E} \sup_{f \in F: \|f - f^*\|_{L_2} \leq r} \left| \sum_{i \in J} \sigma_i (f - f^*)(X_i) \right| \leq r^2 |J| \gamma \right\}. \quad (6)$$

The outputs do not appear in the complexity parameter. This is an interesting feature of Lipschitz losses. It is necessary to adapt the Bernstein assumption to this framework.

Assumption 6 There exists a constant $A > 0$ such that for all $f \in F$ if $\|f - f^*\|_{L_2}^2 = C_{K,r}$ then $\|f - f^*\|_{L_2}^2 \leq AP\mathcal{L}_f$ where

$$C_{K,r} = \max \left(\tilde{r}_2^2(1/(575AL)), 864A^2L^2\frac{K}{N} \right). \tag{7}$$

Assumptions 6 and 4 have a similar flavor as both require the Bernstein condition in a L_2 -sphere centered in f^* with radius given by the rate of convergence of the associated estimator (see Theorems 1 and 2). For $K \leq (\tilde{r}_2^2(575/(AL))N)/(846A^2L^2)$ the sphere $\{f \in F : \|f - f^*\|_{L_2} = \sqrt{C_{K,r}}\}$ is a L_2 -sphere centered in f^* of radius $\tilde{r}_2(575/(AL))$ which can be of order $1/\sqrt{N}$ (see Sect. 3.2.3). As a consequence, Assumption 6 holds in examples where the small ball assumption does not (see discussion after Assumption 9).

3.2.2 Main results

We are now in position to state the main result regarding the statistical properties of estimator (5) under a local Bernstein condition.

Theorem 2 Grant Assumptions 1, 2, 5 and 6 and assume that $|\mathcal{O}| \leq 3N/7$. Let $\gamma = 1/(575AL)$ and $K \in [7|\mathcal{O}|/3, N]$. The minmax MOM estimator \hat{f} defined in (5) satisfies, with probability at least

$$1 - \exp(-K/2016), \tag{8}$$

$$\|\hat{f} - f^*\|_{L_2}^2 \leq C_{K,r} \quad \text{and} \quad P\mathcal{L}_{\hat{f}} \leq \frac{2}{3A}C_{K,r}. \tag{9}$$

Suppose that $K = \tilde{r}_2^2(\gamma)N$, which is possible as long as $|\mathcal{O}| \lesssim N\tilde{r}_2^2(\gamma)$. The deviation bound is then of order $\tilde{r}_2^2(\gamma)$ and the probability estimate $1 - \exp(-N\tilde{r}_2^2(\gamma)/2016)$. Therefore, minmax MOM estimators achieve the same statistical bound with the same deviation as the ERM as long as $\tilde{r}_2^2(\gamma)$ and $r_2(\theta)$ are of the same order. Using generic chaining [40], this comparison is true under Assumption 3. It can also be shown under weaker moment assumption, see [35] or the example of Sect. 3.2.3.

When $\tilde{r}_2^2(\gamma) \asymp r_2(\theta)$, the bounds are rate optimal as shown in [2]. This is why these bounds are called rate optimal sub-Gaussian deviation bounds. While these hold for ERM in the i.i.d. setup with sub-Gaussian design in the absence of outliers (see Theorem 1), they hold for minmax MOM estimators in a setup where inliers may not be i.i.d., nor have sub-Gaussian design and up to $N\tilde{r}_2^2(\gamma)$ outliers may have contaminated the dataset.

This section is concluded by presenting an estimator achieving (5) simultaneously for all K . For all $K \in \{1, \dots, N\}$ and $f \in F$, define $T_K(f) = \sup_{g \in F} \text{MOM}_K(\ell_f - \ell_g)$ and let

$$\hat{R}_K = \{g \in F : T_K(g) \leq (1/3A)C_{K,r}\}. \tag{10}$$

Now, building on the Lepskii's method, define a data-driven number of blocks

$$\hat{K} = \inf \left(K \in \{1, \dots, N\} : \bigcap_{J=K}^N \hat{R}_J \neq \emptyset \right) \quad (11)$$

and let \tilde{f} be such that

$$\tilde{f} \in \bigcap_{J=\hat{K}}^N \hat{R}_J. \quad (12)$$

Theorem 3 *Grant Assumptions 1, 2, 5 and 6 and assume that $|\mathcal{O}| \leq 3N/7$. Let $\gamma = 1/(575AL)$. The estimator \tilde{f} defined in (12) is such that for all $K \in [7|\mathcal{O}|/3, N]$, with probability at least*

$$1 - 4 \exp(-K/2016),$$

$$\|\tilde{f} - f^*\|_{L_2}^2 \leq C_{K,r} \text{ and } P\mathcal{L}_{\tilde{f}} \leq \frac{2}{3A} C_{K,r}.$$

Theorem 3 states that \tilde{f} achieves the results of Theorem 2 simultaneously for all $K \geq 7|\mathcal{O}|/3$. This extension is useful as the number $|\mathcal{O}|$ of outliers is typically unknown in practice. However, contrary to \hat{f} , the estimator \tilde{f} requires the knowledge of A and $\tilde{r}(\gamma)$. These parameters allow to build confidence regions for f^* , which is necessary to apply Lepski's method. Similar limitations appear in least-squares regression [24] and even in the basic problem of univariate mean estimation. In this simpler problem, it can be shown that one can build sub-Gaussian estimators depending on the confidence level (through K) under only a second moment assumption. On the other hand, to build estimators achieving the same risk bounds simultaneously for all K , more informations on the distribution of the data are required, see [14, Theorem 3.2]. In particular, the knowledge of the variance, which allows to build confidence intervals for the unknown univariate mean, is sufficient. The necessity of extra-information to obtain adaptivity with respect to K is therefore not surprising here.

3.2.3 Some basic examples

The following example illustrates the optimality of the rates provided in Theorem 2 even under a simple L_2 -moment assumption.

Lemma 1 [19] *In the $\mathcal{O} \cup \mathcal{I}$ framework with $F = \{\langle t, \cdot \rangle : t \in \mathbb{R}^d\}$, we have $\tilde{r}_2^2(\gamma) \leq \text{Rank}(\Sigma)/(2\gamma^2 N)$, where $\Sigma = \mathbb{E}[XX^T]$ is the $d \times d$ covariance matrix of X .*

The proof of Lemma 1 is recalled in Sect. B for the sake of completeness. Lemma 1 grants only the existence of a second moment for X even though the rate obtained $\text{Rank}(\Sigma)/(2\gamma^2 N)$ is the same as the one we would get under a sub-Gaussian assumption given that $r_2(\theta) \sim \text{Rank}(\Sigma)/(2\theta^2 N)$. Moreover, Sect. 5 shows that Assumptions 4

and 6 are satisfied when $F = \{ \langle t, \cdot \rangle : t \in \mathbb{R}^d \}$ and X is a vector with i.i.d. entries having only a few finite moments. Theorem 2 applies therefore in this setting and the Min-max MOM estimator (5) achieves the optimal fast rate of convergence $\text{Rank}(\Sigma)/N$. This shows that when the model is the entire space \mathbb{R}^d , the results for the ERM from Theorem 1 obtained under a sub-Gaussian assumption is the same as the one for the minmax MOM from Theorem 2 under only weak moment assumption.

However, Lemma 1 does not describe a typical situation. Having $\tilde{r}_2^2(\gamma)$ of the same order as $r_2^2(\theta)$ under only a second moment assumption is mainly happening on large models such as the entire space \mathbb{R}^d . For smaller size models such as the B_1^d -ball (the unit ball of the ℓ_1^d -norm), the picture is different: $\tilde{r}_2^2(\gamma)$ should be bigger than $r_2^2(\theta)$ unless X has enough moment. To make this statement simple let us consider the case $N = 1$. In that case, we have $r_2^2(\theta) \sim \sqrt{\log d}$. Let us now describe $\tilde{r}_2^2(\gamma)$ under various moment assumptions on X to see when $\tilde{r}_2^2(\gamma)$ compares with $\sqrt{\log d}$.

Let $X = (x_j)_{j=1}^d$ be a random vector. It follows from Equation (3.1) in [36] that

$$\mathbb{E} \sup_{t \in B_1^d \cap r B_2^d} |\langle t, X \rangle| \lesssim \begin{cases} r \mathbb{E} \left(\sum_{j=1}^d x_j^2 \right)^{1/2} & \text{if } r \leq 1/\sqrt{d} \\ r \mathbb{E} \left(\sum_{j=1}^k (x_j^*)^2 \right)^{1/2} & \text{if } 1/\sqrt{d} \leq r \leq 1 \\ \mathbb{E} \max_{j=1, \dots, d} |x_j| & \text{if } r \geq 1. \end{cases}$$

where $k = \lceil 1/r^2 \rceil$ and $x_1^* \geq \dots \geq x_d^*$ is a non-increasing rearrangement of the absolute values of the coordinates $x_j, j = 1, \dots, d$ of X . Assume that x_1, \dots, x_d are i.i.d. distributed like x such that $\mathbb{E}x = 0$ and $\mathbb{E}x^2 = 1$. Assume that x has $2p$ moments, for $p \geq 1$ and let c_0 be such that $\mathbb{E}[x^{2p}] \leq c_0$. Then, using Jensen's inequality, we obtain

$$\left(\mathbb{E} \left[\frac{1}{k} \sum_{j=1}^k (x_j^*)^2 \right] \right)^p \leq \frac{1}{k} \sum_{j=1}^k \mathbb{E}[(x_j^*)^{2p}] \leq \frac{1}{k} \sum_{j=1}^d \mathbb{E}[x_j^{2p}] \leq \frac{c_0 d}{k}.$$

It follows that

$$\begin{aligned} \mathbb{E} \left[\left(\frac{1}{k} \sum_{j=1}^k (x_j^*)^2 \right)^{1/2} \right] &\leq \left(\mathbb{E} \left[\frac{1}{k} \sum_{j=1}^k (x_j^*)^2 \right] \right)^{1/2} \\ &\leq \left(\frac{c_0 d}{k} \right)^{1/(2p)} \leq \left(\frac{c_0 d}{k} \right)^{1/(2p)} \leq (c_0 d r^2)^{1/(2p)}. \end{aligned}$$

Hence,

$$\mathbb{E} \sup_{t \in B_1^d \cap r B_2^d} |\langle t, X \rangle| \lesssim \begin{cases} r\sqrt{d} & \text{if } r \leq 1/\sqrt{d} \\ (c_0 d r^2)^{1/(2p)} & \text{if } 1/\sqrt{d} \leq r \leq 1 \\ (c_0 d)^{1/(2p)} & \text{if } r \geq 1. \end{cases}$$

Assume that $f^* = \langle t^*, \cdot \rangle$, with $t^* = 0$. Then,

$$\tilde{r}_2(\gamma) = \inf \left\{ r > 0 : \mathbb{E} \sup_{t \in B_1^d \cap r B_2^d} |\langle t, X \rangle| \leq \gamma r^2 \right\} \lesssim (c_0 d)^{1/(4p)} / \sqrt{\gamma}.$$

In particular,

$$\begin{cases} \tilde{r}_2(\gamma) \asymp 1 & \text{when } p \geq \log(c_0 d), \\ 1 \lesssim r_2(\gamma) \lesssim \sqrt{\log d} & \text{when } \log(c_0 d) / \log \log d \leq p \leq \log(c_0 d), \\ \tilde{r}_2(\gamma) \asymp d^{1/(4p)} & \text{when } p \leq \log(c_0 d) / \log \log d. \end{cases}$$

Let us now show that these estimates are sharp by considering $x = \epsilon(1 + R\eta)$ where ϵ is a Rademacher variable, η is a Bernoulli variable (independent of ϵ) with mean $\delta = 1/d$ and $R = d^{1/(4p)}$. We have $\mathbb{E}x = 0$ and $\mathbb{E}x^2 = 1 + R\delta \leq 2$ because $R\delta \leq 1$ when $p \geq 1$. Let $x_j = \epsilon_j(1 + R\eta_j)$, $j = 1, \dots, d$ be i.i.d. copies of x . We have

$$\begin{aligned} & \mathbb{E} \max_{j=1, \dots, d} |x_j| \\ & \geq (1 + R) \mathbb{P} \left[\max_{j=1, \dots, d} |x_j| \geq 1 + R \right] \\ & = (1 + R) (1 - \mathbb{P}[\eta_j = 0, \forall j = 1, \dots, d]) \\ & = (1 + R)(1 - (1 - \delta)^d) \geq (1 + R)e^{-1} \gtrsim d^{1/(4p)}. \end{aligned}$$

As a consequence, for all $1 \leq r \lesssim d^{1/(8p)}$,

$$\mathbb{E} \sup_{t \in B_1^d \cap r B_2^d} |\langle t, X \rangle| = \mathbb{E} \max_{j=1, \dots, d} |x_j| > r^2$$

and so $\tilde{r}_2^2(\gamma) \gtrsim d^{1/(4p)}$. As a consequence, under only a L_{2p} moment assumption one cannot have $\tilde{r}_2^2(\gamma)$ better than $d^{1/(4p)}$.

As a consequence, $\tilde{r}_2(\gamma)$ can be much larger than $r_2(\theta)$ when x has less than $\log(c_0 d) / \log(\log d)$ moments, for instance, $\tilde{r}_2^2(\gamma)$ can be of the order of $d^{1/8}$ when x has only 2 moments. This picture is different from the one given by Lemma 1 where we were able to get equivalence between $\tilde{r}_2(\gamma)$ and $r_2(\theta)$ only under a second moment assumption.

4 Relaxing the Bernstein condition

This section shows that minmax MOM estimators satisfy sharp oracle inequalities with exponentially large deviation under minimal stochastic assumptions insuring the existence of all objects. These results are slightly weaker than those of the previous section: the L_2 risk is not controlled and only slow rates of convergence hold in this relaxed setting. However, the bounds are sufficiently precise to imply fast rates

of convergence for the excess risk as in Theorems 2 if a slightly stronger Bernstein condition holds.

Given that data may not have the same distribution as (X, Y) , the following relaxed version of Assumption 5 is introduced.

Assumption 7 $(X_i, Y_i)_{i \in \mathcal{I}}$ are independent and for all $i \in \mathcal{I}$, (X_i, Y_i) has distribution P_i , X_i has distribution μ_i . For any $i \in \mathcal{I}$, $F \subset L_2(\mu_i)$ and $P_i \mathcal{L}_f = P \mathcal{L}_f$ for all $f \in F$.

When Assumption 6 does not necessary hold, the localization argument has to be modified. Instead of the L_2 -norm, the excess risk $f \in F \rightarrow P \mathcal{L}_f$ is used to define neighborhoods around f^* . The associated complexity is then defined for all $\gamma > 0$ and $K \in \{1, \dots, N\}$ by

$$\bar{r}_2(\gamma) \geq \inf \left\{ r > 0 : \max \left(\frac{E(r)}{\gamma}, \sqrt{1536} V_K(r) \right) \leq r^2 \right\} \tag{13}$$

where

$$E(r) = \sup_{J \subset \mathcal{I}: |J| \geq N/2} \mathbb{E} \sup_{f \in F: P \mathcal{L}_f \leq r^2} \left| \frac{1}{|J|} \sum_{i \in J} \sigma_i(f - f^*)(X_i) \right|,$$

$$\text{and } V_K(r) = \max_{i \in \mathcal{I}} \sup_{f \in F: P \mathcal{L}_f \leq r^2} \sqrt{\text{Var}_{P_i}(\mathcal{L}_f)} \sqrt{\frac{K}{N}}.$$

There are two important differences between $\bar{r}_2(\gamma)$ on one side and $r_2(\theta)$ in Definition 2 or $\tilde{r}_2(\gamma)$ in (6) on the other side. The first one is the extra variance term $V_K(r)$. Under the Bernstein condition, this term is negligible in front of the “expectation term” $E(r)$ see [8]. In the general setting considered here, the variance term is handled in the complexity parameter. The second important consequence is that \bar{r}_2 is a fixed point of the complexity of F localized around f^* with respect to the excess risk rather than with respect to the L_2 -norm. An important consequence is that this quantity is harder to compute in practical examples. As a consequence, the results of this section are more of theoretical importance.

Theorem 4 Grant Assumptions 1, 2, 7 and assume that $|\mathcal{O}| \leq 3N/7$. Let $\gamma = 1/(768L)$ and $K \in [7|\mathcal{O}|/3, N]$. The minmax MOM estimator \hat{f} defined in (5) satisfies, with probability at least $1 - \exp(-K/2016)$, $P \mathcal{L}_{\hat{f}} \leq \bar{r}_2^2(\gamma)$.

Recall that Assumptions 1 and 2 are only meaning that the loss function is convex and Lipschitz and that the class F is convex. Assumption 7 says that inliers are independent and define the same excess risk as (X, Y) over F . In particular, Theorem 4 holds, as Theorem 2, without assumptions on the outliers $(X_i, Y_i)_{i \in \mathcal{O}}$ and with weak assumptions on the outputs $(Y_i)_{i \in \mathcal{I}}$ of the inliers (we remark that excess loss function $f \rightarrow P \mathcal{L}_f$ is well-defined under no assumption on Y —even if $Y \notin L_1$ —because $|\mathcal{L}_f| \leq L|f - f^*|$). Moreover, the excess risk bound holds with exponentially large

probability without assuming sub-Gaussian design, a small ball hypothesis or a Bernstein condition. This generality can be achieved by combining MOM estimators with convex-Lipschitz loss functions.

The following result discuss relationships between Theorems 2 and 4. Introduce the following modification of the Bernstein condition.

Assumption 8 Let $\gamma = 1/(768L)$. There exists a constant $A > 0$ such that for all $f \in F$ if $P\mathcal{L}_f = C'_{K,r}$ then $\|f - f^*\|_{L_2}^2 \leq AP\mathcal{L}_f$ where, for $\tilde{r}_2(\gamma)$ defined in (6),

$$C'_{K,r} = \max\left(\frac{\tilde{r}_2^2(\gamma/A)}{A}, \frac{1536AL^2K}{N}\right).$$

Assumption 8 is slightly stronger than Assumption 4 since the L_2 -metric to define the sphere is replaced by the excess risk metric. If Assumption 8 holds then Theorem 5 implies the same statistical bounds for (5) as Theorem 2 up to constants, as shown by the following result.

Theorem 5 Grant Assumptions 1, 2, 7 and assume that $|\mathcal{O}| \leq 3N/7$. Assume that the local Bernstein condition Assumption 8 holds. Let $\gamma = 1/(768L)$ and $K \in [7|\mathcal{O}|/3, N]$. The minmax MOM estimator \hat{f} defined in (5) satisfies, with probability at least $1 - \exp(-K/2016)$,

$$\begin{aligned} \|\hat{f} - f^*\|_{L_2}^2 &\leq \max\left(\tilde{r}_2^2(\gamma/A), \frac{1536L^2A^2K}{N}\right) \\ \text{and } P\mathcal{L}_{\hat{f}} &\leq \max\left(\frac{\tilde{r}_2^2(\gamma/A)}{A}, \frac{1536L^2AK}{N}\right). \end{aligned}$$

Proof First, $V_K(r) \leq LV'_K(r)$ for all $r > 0$ where $V'_K(r) = \sqrt{K/N} \max_{i \in \mathcal{I}} \sup_{f \in F: P\mathcal{L}_f \leq r^2} \|f - f^*\|_{L_2(\mu_i)}$. Moreover, $r \rightarrow E(r)/r^2$ and $r \rightarrow V'_K(r)/r^2$ are non-increasing, therefore by Assumption 8 and the definition of $\tilde{r}_2(\gamma)$, $V'_K(r)$,

$$\frac{1}{\gamma} E\left(\frac{\tilde{r}_2(\gamma/A)}{\sqrt{A}}\right) \leq \frac{\tilde{r}_2^2(\gamma/A)}{A} \quad \text{and} \quad \sqrt{1536}V'_K\left(\sqrt{1536}L\sqrt{\frac{AK}{N}}\right) \leq \frac{1536A^2LK}{N}.$$

Hence, $\tilde{r}_2^2(\gamma) \leq \max(\tilde{r}_2^2(\gamma/A)/A, 1536L^2A(K/N))$. □

5 Bernstein’s assumption

This section shows that the local Bernstein condition holds for various loss functions and design X . In Assumption 4 and 6, the comparizon between $P\mathcal{L}_f$ and $\|f - f^*\|_{L_2}^2$ is only required on a L_2 -sphere. In this section, we prove that the local Bernstein assumption can be verified over the entire L_2 -ball and not only on the sphere under mild moment conditions. The class $F - \{f^*\}$ is assumed to satisfy a “ $L_{2+\varepsilon}/L_2$ -norm equivalence assumption”, for $\varepsilon > 0$.

Assumption 9 Let $\varepsilon > 0$. There exists $C' > 0$ such that for all $f \in F$, $\|f - f^*\|_{L_{2+\varepsilon}} \leq C' \|f - f^*\|_{L_2}$

Assumption 9 is a “ $L_{2+\varepsilon}/L_2$ ” norm equivalence assumption over $F - \{f^*\}$. A “ L_4/L_2 ” norm equivalence assumption over $F - \{f^*\}$ has been used for the study of MOM estimators (see [29]). Examples of distributions satisfying Assumption 9 can be found in [33,34].

There are situations where the constant C' depends on the dimension d of the model. In that case, the results in [24,29] provide sub-optimal statistical upper bounds. For instance, if X is uniformly distributed on $[0, 1]$ and $F = \{\sum_{j=1}^d \alpha_j I_{A_j} : (\alpha_j)_{j=1}^d \in \mathbb{R}^d\}$ where I_{A_j} is the indicator of $A_j = [(j - 1)/d, j/d]$ then for all $f \in F$, $\|f - f^*\|_{L_{2+\varepsilon}} \leq d^{\varepsilon/(4+2\varepsilon)} \|f - f^*\|_{L_2}$ so $C' = d^{\varepsilon/(4+2\varepsilon)}$. This dependence with respect to the dimension d is inevitable. For instance, in [24,29], a L_4/L_2 norm equivalence is required. In this case, $C' = d^{1/4}$ which ultimately yields sub-optimal rates in this example. On the other hand, as will become clear in this section, the rates given in Theorem 2 or Theorem 3 are not deteriorated in this example. This improvement is possible since the Bernstein condition is only required in a neighborhood of f^* .

5.1 Quantile loss

The proof is based on [15, Lemma 2.2] and is postponed to Sect. C.1. Recall that $\ell_f(x, y) = (y - f(x))(\tau - I\{y - f(x) \leq 0\})$.

Assumption 10 Let C' be the constant defined in Assumption 9. There exist $\alpha > 0$ and $r > 0$ such that, for all $x \in \mathcal{X}$ and for all z in \mathbb{R} such that $|z - f^*(x)| \leq r(\sqrt{2}C')^{(2+\varepsilon)/\varepsilon}$, we have $f_{Y|X=x}(z) \geq \alpha$, where $f_{Y|X=x}$ is the conditional density function of Y given $X = x$.

Theorem 6 Grant Assumptions 9 (with constant C') and 10 (with parameter r and α). Then, for all $f \in F$ satisfying $\|f - f^*\|_{L_2} \leq r$, $\|f - f^*\|_{L_2}^2 \leq (4/\alpha)P\mathcal{L}_f$.

Consider the example from Sect. 3.2.3, assume that $K \lesssim \text{Rank}(\Sigma)$ and let $r^2 = C_{K,r} = \tilde{r}_2^2(\gamma) \leq \text{Rank}(\Sigma)/(2\gamma^2N)$. If $C' = d^{\varepsilon/(4+2\varepsilon)}$, Assumption 10 holds for r and an associated $\alpha \gtrsim 1$ as long as $d^{1/2}\sqrt{\text{Rank}(\Sigma)/N} \lesssim 1$ and, for all $x \in \mathcal{X}$ and for all z in \mathbb{R} such that $|z - f^*(x)| \lesssim 1$, $f_{Y|X=x}(z) \gtrsim 1$. As $\text{Rank}(\Sigma) \leq d$, the first condition reduces to $N \gtrsim d^2$. In this situation, the rates given in Theorems 2 and 3 are still $\text{Rank}(\Sigma)/N$. This gives a partial answer, in our setting, to the issue raised in [39] regarding results based on the small ball method.

5.2 Huber loss

Consider the Huber loss function defined, for all $f \in F$, $x \in \mathcal{X}$ and $y \in \mathbb{R}$, by $\ell_f(x, y) = \rho_H(y - f(x))$ where $\rho_H(t) = t^2/2$ if $|t| \leq \delta$ and $\rho_H(t) = \delta|t| - \delta^2/2$ otherwise. Introduce the following assumption.

Assumption 11 Let C' be the constant defined in Assumption 9. There exist $\alpha > 0$ and $r > 0$ such that for all $x \in \mathcal{X}$ and all z in \mathbb{R} such that $|z - f^*(x)| \leq (\sqrt{2}C')^{(2+\varepsilon)/\varepsilon}r$,

$F_{Y|X=x}(z + \delta) - F_{Y|X=x}(z - \delta) \geq \alpha$, where $F_{Y|X=x}$ is the conditional cumulative function of Y given $X = x$.

Under this assumption and a “ $L_{2+\epsilon}/L_2$ ” assumption, the local Bernstein condition is proved to be satisfied in the following result whose proof is postponed to Sect. C.2.

Theorem 7 *Grant Assumptions 9 (with constant C') and 11 (with parameter r and α). Then, for all $f \in F$ satisfying $\|f - f^*\|_{L_2} \leq r$, there exists $\alpha > 0$ (given by Assumption 11) such that $\|f - f^*\|_{L_2}^2 \leq (4/\alpha)P\mathcal{L}_f$.*

5.3 Logistic classification

In this section we consider the logistic loss function.

Assumption 12 There exists $c_0 > 0$ such that

$$\mathbb{P}(|f^*(X)| \leq c_0) \geq 1 - \frac{1}{(2C')^{(4+2\epsilon)/\epsilon}}.$$

where C' is defined in Assumption 9.

The following result is proved in Sect. C.3.

Theorem 8 *Grant Assumptions 9 and 13. Then, for all $r > 0$ and all $f \in F$ such that $\|f - f^*\|_{L_2} \leq r$,*

$$P\mathcal{L}_f \geq \frac{e^{-c_0 - r(2C')^{(2+\epsilon)/\epsilon}}}{2(1 + e^{c_0 + r(2C')^{(2+\epsilon)/\epsilon}})^2} \|f - f^*\|_{L_2}^2.$$

The proof is postponed to Sect. C.3. As for the Huber Loss and the Hinge Loss, the rates of convergence are not deteriorated when C' may depend on the dimension as long as $r \times (C')^{(2+\epsilon)/\epsilon}$ is smaller than some absolute constant.

5.4 Hinge loss

In this section, we show that the local Bernstein condition holds for various design X for the Hinge loss function. We obtain the result under the assumption that the oracle f^* is actually the Bayes rules which is the function minimizing the risk $f \mapsto R(f)$ over all measurable functions from \mathcal{X} to \mathbb{R} . Recall that, under this assumption, $f^*(x) = \text{sign}(2\eta(x) - 1)$ where $\eta(X) = \mathbb{P}(Y = 1|X)$. In that case, the Bernstein condition (see [8]) coincides with the margin assumption (see [32,41]).

Assumption 13 Let C' be the constant defined in Assumption 9. There exist $\alpha > 0$ and $0 < r \leq (\sqrt{2}C')^{-(2+\epsilon)/\epsilon}$ such that for all $x \in \mathcal{X}$, for all $z \in \mathbb{R}$, $|z - f^*(x)| \leq (\sqrt{2}C')^{(2+\epsilon)/\epsilon} r$

$$\min(\eta(x), 1 - \eta(x), |1 - 2\eta(x)|) \geq \alpha.$$

Assumption 13 is also local and has the same flavor as Assumptions 10 and 11.

Theorem 9 *Grant Assumptions 9 (with constant C') and 13 (with parameter r and α). Assume that the oracle f^* is the Bayes estimator i.e. $f^*(x) = \text{sign}(2\eta(x) - 1)$ for all $x \in \mathcal{X}$. Then, for all $f \in F$ such that $\|f - f^*\|_{L_2} \leq r$, $\|f - f^*\|_{L_2}^2 \leq \frac{2}{\alpha} P\mathcal{L}_f$.*

The proof is postponed to Sect. C.4.

6 Comparison between ERM and minmax MOM

In this section, we show that robustness properties with respect to heavy-tailed data and to outliers of the minmax MOM estimator in Theorem 2 cannot be achieved by the ERM. We prove two lower bounds on the statistical risk of ERM. First, we show that ERM is not robust to contamination in the design X and second that ERM cannot achieve the optimal rate with a sub-Gaussian deviation under only moment assumptions.

We first show the absence of robustness of ERM w.r.t. contamination by even a single input variable. We consider the absolute loss function of linear functionals $\ell_t(x, y) = |y - \langle x, t \rangle|$. Let X_1, \dots, X_N denote i.i.d. Gaussian vectors, and suppose that there exists t^* such that $Y_i = \langle X_i, t^* \rangle, i = 1, \dots, N$. Assume that a vector $v \in \mathbb{R}^d$ was added to X_1 (and that this is the only corrupted data). Hence, we are given the dataset $(X_1 + v, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)$. Consider the ERM constructed on this dataset i.e $\hat{t}^{ERM} \in \text{argmin}_{t \in \mathbb{R}^d} P_N \ell_t$ where $P_N \ell_t = (1/N)|Y_1 - \langle X_1 + v, t \rangle| + (1/N) \sum_{i=2}^N |Y_i - \langle X_i, t \rangle|$. In this context, the following lower bound holds.

Proposition 1 *There exist absolute constants c_0 and c_2 such that the following holds. If the contamination vector v satisfies $|\langle v, t^* \rangle| \geq (1/2) \|v\|_2 \|t^*\|_2$, with $\|v\|_2 \geq c_2 N$, then with probability at least $1 - 4 \exp(-c_0 N)$, $\|\hat{t}^{ERM} - t^*\|_2 \geq \|t^*\|_2 / 4$.*

When $N \asymp d$, from Theorem 2 with $K \asymp d$, minmax MOM estimators yields, with probability at least $1 - 2 \exp(-cd)$, $\|\hat{t}^{MOM} - t^*\|_2 \lesssim d/N$ on the same dataset as the one used by ERM in Proposition 1. If $\|t^*\|_2 \gtrsim 1$ then the ERM is suboptimal compared with the minmax MOM estimator.

Proof To show that \hat{t}^{ERM} is outside $\mathcal{B} = B_2(t^*, (1/4) \|t^*\|_2) = \{t \in \mathbb{R}^d : \|t - t^*\|_2 \leq (1/4) \|t^*\|_2\}$, it is enough to show that $P_N \ell_0$ is smaller than the smallest value of $t \rightarrow P_N \ell_t$ over \mathcal{B} . It follows from Gaussian concentration that, with probability at least $1 - \exp(-c_0 N)$,

$$P_N \ell_0 = \frac{1}{N} \sum_{i=1}^N |\langle X_i, t^* \rangle| \leq \frac{3}{2} \sqrt{\frac{2}{\pi}} \|t^*\|_2. \tag{14}$$

Let us now bound from below the empirical loss function $t \rightarrow P_N \ell_t$ uniformly over all t in \mathcal{B} . First,

$$|\langle v, t \rangle| \geq |\langle v, t^* \rangle| - |\langle v, t - t^* \rangle| \geq \|v\|_2 (\|t^*\|_2 / 2 - \|t - t^*\|_2) \geq \|v\|_2 \|t^*\|_2 / 4. \tag{15}$$

Then, it follows from Borell-TIS inequality (see Theorem 7.1 in [27] or pages 56-57 in [28]) that with probability at least $1 - 2 \exp(-c_1 N)$, $\|X_1\|_2 \leq \mathbb{E} \|X_1\|_2 + \sqrt{2c_1 N} \leq c_3 \sqrt{N + d}$. Therefore, $\mathbb{P}(\Omega_1) \geq 1 - 2 \exp(-c_1 N)$, where

$$\Omega_1 = \{\forall t \in \mathbb{R}^d, \quad |\langle X_1, t - t^* \rangle| \leq c_3 \sqrt{N + d} \|t - t^*\|_2\}.$$

On Ω_1 , we have for all $t \in \mathcal{B}$ that $|\langle X_1, t^* - t \rangle| \leq c_3 \sqrt{N + d} \|t - t^*\|_2 \leq (c_3/4) \sqrt{N + d} \|t^*\|_2$. Therefore, using (15) and $\|v\|_2 \geq c_2 N$ for a large enough constant c_2 ,

$$\begin{aligned} P_N \ell_t &= \frac{1}{N} |\langle X_1, t^* - t \rangle - \langle v, t \rangle| + \frac{1}{N} \sum_{i=2}^N |\langle X_i, t^* - t \rangle| \\ &\geq \frac{1}{N} |\langle v, t \rangle| - \frac{1}{N} |\langle X_1, t^* - t \rangle| \\ &\geq \frac{1}{4N} \|v\|_2 \|t^*\|_2 - \frac{c_3 \|t^*\|_2}{4\sqrt{N}} > \frac{3}{2} \sqrt{\frac{2}{\pi}} \|t^*\|_2. \end{aligned} \tag{16}$$

□

It follows from Proposition 1 that ERM is not consistant when there is even a single outlier among the X_i . By comparison, the minmax MOM has optimal performance even when the dataset has been corrupted by up to d outliers when $N \gtrsim d$. This shows a first advantage of the minmax MOM approach.

Now, we prove a second advantage of the minmax MOM over the ERM by considering heavy-tailed design. We also consider the absolute L_1 -loss function as in the previous example and suppose that data are generated from a linear model in dimension $d = 1$: $Y = Xt^* + \zeta$ where X and ζ are independent mean zero random variables and $t^* \in \mathbb{R}$ (we choose $d = 1$ so that we have access to a canonical definition of median which simplifies the proof). Our aim is to show that if the design X has only a second moment then the ERM \hat{t}^{ERM} cannot achieve the optimal rate $\sqrt{x/N}$ with a sub-Gaussian deviation that is $1 - \exp(-c_0 x)$ as does the minmax MOM for all $x \in [1, N]$.

Proposition 2 *Let $N \geq 8000$ and $10 \leq x \leq N/800$. There exist X and ζ two symmetric and independent random variables such that $\mathbb{E}X^2 \in [1, 16]$, $\mathbb{E}\zeta^2 \leq 5x^2$ and, for any $t^* \in \mathbb{R}$ and $Y = Xt^* + \zeta$, we have $\{t^*\} = \operatorname{argmin}_{t \in \mathbb{R}} \mathbb{E}|Y - Xt|$. Let $(X_i, Y_i)_{i=1}^N$ be N i.i.d. copies of (X, Y) such that $Y = Xt^* + \zeta$ for some $t^* \in \mathbb{R}$. Let $\hat{t}^{ERM} \in \operatorname{argmin}_{t \in \mathbb{R}} \sum_{i=1}^N |Y_i - X_i t|$. Then, with probability at least $3/(5x)$,*

$$\sqrt{\mathbb{E}[(X(\hat{t}^{ERM} - t^*))^2]} \geq (1/5)\sqrt{x/N}.$$

Proof Let $\delta' = (1/8)\sqrt{x/(2N)}$ and let ζ be uniformly distributed over $[-x - 1/2 + \delta', -x] \cup [-\delta', \delta'] \cup [x, x + 1/2 - \delta']$. Let ϵ denote a Rademacher variable, let η be a Bernoulli variable with parameter $\delta = 1/(xN)$ and $R = 4/\sqrt{\delta} = 4\sqrt{xN}$. We

assume that ζ, ϵ and η are independent and let $X = \epsilon(1 + R\eta)$. Let $t^* \in \mathbb{R}$ and let $\mathcal{D}_N = (X_i, Y_i)_{i=1}^N$ be a dataset of N i.i.d. copies of (X, Y) , where $Y = Xt^* + \zeta$.

Since the median of ζ is 0, for all $u \in \mathbb{R}, \mathbb{E}|u - \zeta| \geq \mathbb{E}|\zeta|$ with equality iff $u = 0$. As a consequence, for all $t \in \mathbb{R}, \mathbb{E}|Y - Xt| = \mathbb{E}_X \mathbb{E}_\zeta |\zeta - X(t^* - t)| \geq \mathbb{E}|\zeta|$ and the only minimizer of $t \in \mathbb{R} \rightarrow \mathbb{E}|Y - Xt|$ is t^* . In other words, t^* is the oracle. For all $t \in \mathbb{R}$,

$$\mathbb{E}[(X(t - t^*))^2] = \mathbb{E}[X^2](t - t^*)^2 = (1 + 2R\delta + R^2\delta)(t - t^*)^2.$$

Since $R^2\delta = 16$ and $2R\delta \leq \sqrt{8}/100 \leq 1$, we have $(t - t^*)^2 \leq \mathbb{E}(X(t - t^*))^2 \leq 18(t - t^*)^2$, that is, the $L^2(\mu)$ -norm is equivalent to the absolute value.

Observe that $\hat{t}^{ERM} - t^*$ is solution of the minimization problem

$$(\hat{t}^{ERM} - t^*) \in \operatorname{argmin}_{u \in \mathbb{R}} \sum_{i=1}^N |X_i| \left| \frac{\zeta_i}{X_i} - u \right| = \operatorname{argmin}_{u \in \mathbb{R}} \mathbb{E}[|W - u| | \mathcal{D}_N].$$

Here, defining $\zeta'_i = \epsilon_i \zeta_i, i \in [N], W$ is a random variable such that

$$\mathbb{P}\left[W = \frac{\zeta'_i}{|X_i|} | \mathcal{D}_N\right] = \frac{|X_i|}{\sum_{i=1}^N |X_i|}.$$

Notice that, almost surely, all $\zeta'_i/|X_i|$ are different. In particular, $|\hat{t}^{ERM} - t^*|$ is the absolute value of the empirical median $|\operatorname{Median}(W)|$. Therefore, $|\hat{t}^{ERM} - t^*| \geq c_1\sqrt{x/N}$ when the median of W does not belong to $(-c_1\sqrt{x/N}, c_1\sqrt{x/N})$. This holds when $\mathbb{P}[W \leq -c_1\sqrt{x/N} | \mathcal{D}_N] > 1/2$ or $\mathbb{P}[W \geq c_1\sqrt{x/N} | \mathcal{D}_N] > 1/2$. Introduce the following sets

$$I_{\leq -x} := \{i \in [N] : \zeta'_i \leq -x\}, \\ I_{\delta'} := \{i \in [N] : |\zeta'_i| \leq \delta'\} \text{ and } I_{\geq x} := \{i \in [N] : \zeta'_i \geq x\}.$$

Define also the following events

$$\Omega_0 := \{|I_{\delta'}| \leq \sqrt{2xN}, \quad ||I_{\leq -x}| - |I_{\geq x}|| \leq \sqrt{2xN}\}, \\ \Omega_1 := \{\forall i \in I_{\delta'} : \eta_i = 0 \text{ and } |\{i \in [N] : \eta_i = 1\}| = 1\}.$$

By Hoeffding's inequality (see Chapter 2 in [11]), as $(\zeta'_i)_{i=1}^N$ is a family of i.i.d. random variables distributed like ζ_1 , with probability at least $1 - \exp(-x/4)$,

$$|I_{\delta'}| = \sum_{i=1}^N I(|\zeta'_i| \leq \delta') \leq N\mathbb{P}[|\zeta'_1| \leq \delta'] + \sqrt{\frac{xN}{2}} = 2\delta'N + \sqrt{\frac{xN}{2}} \leq \sqrt{2xN}.$$

Since $\mathbb{P}[\zeta'_1 \leq -x] = \mathbb{P}[\zeta'_1 \geq x]$, $I(\zeta'_i \leq -x) - I(\zeta'_i \geq x)$ are independent, centered random variables taking values in $[-1, 1]$. By Hoeffding's inequality, with probability at least $1 - 2 \exp(-x/2)$,

$$||I_{\leq -x}| - |I_{\geq x}|| = \left| \sum_{i=1}^N I(\zeta'_i \leq -x) - I(\zeta'_i \geq x) \right| \leq \sqrt{2xN}.$$

Using a union bound, we have $\mathbb{P}[\Omega_0] \geq 1 - 2 \exp(-x/2) - \exp(-x/4) \geq 1 - 1/(10x)$ when $x \geq 10$. Since the ζ'_i 's and the η_i 's are independent, on the event Ω_0 , we have

$$\begin{aligned} \mathbb{P}[\forall i \in I_{\delta'} : \eta_i = 0 | \mathcal{D}_N] &= (1 - \delta)^{|I_{\delta'}|} \geq (1 - \delta)^{\sqrt{2xN}} \geq 1 - 2\delta\sqrt{2xN} \\ &= 1 - \frac{2\sqrt{2}}{\sqrt{xN}} \geq 1 - \frac{1}{10x}. \end{aligned}$$

The last inequality holds since $x \leq N/800$. Moreover,

$$\mathbb{P}[|\{i \in [N] : \eta_i = 1\}| = 1] = N\delta(1 - \delta)^{N-1} \geq \frac{1 - 2/x}{x}.$$

When $x \geq 10$, this implies

$$\mathbb{P}[|\{i \in [N] : \eta_i = 1\}| = 1] \geq \frac{4}{5x}.$$

Therefore, on the event Ω_0 , $\mathbb{P}[\Omega_1 | \mathcal{D}_N] \geq 7/(10x)$ and so

$$\mathbb{P}[\Omega_0 \cap \Omega_1] = \mathbb{E}[\mathbf{1}_{\Omega_0} \mathbb{E}[\mathbf{1}_{\Omega_1} | \mathcal{D}_N]] \geq \frac{4}{5x} \left(1 - \frac{1}{10x}\right) \geq 3/(5x).$$

We want to show that $|\text{Median}(W)| \geq c_1\sqrt{x/N}$ on the event $\Omega_0 \cap \Omega_1$, for some well-chosen constant $c_1 > 0$. We have $\mathbb{P}[W \leq -c_1\sqrt{x/N} | \mathcal{D}_N] > 1/2$ if and only if

$$\sum_{i=1}^N I\left(\frac{\zeta'_i}{|X_i|} \leq -c_1\sqrt{x/N}\right) |X_i| \geq \frac{1}{2} \sum_{i=1}^N |X_i|.$$

In particular, if

$$\sum_{i \in I_{\leq -x}} |X_i| \geq \frac{1}{2} \sum_{i=1}^N |X_i| \text{ or } \sum_{i \in I_{\geq x}} |X_i| \geq \frac{1}{2} \sum_{i=1}^N |X_i| \tag{17}$$

then the median of W takes value in $\{\zeta'_i/|X_i| : i \in I_{\leq -x}\}$, resp. in $\{\zeta'_i/|X_i| : i \in I_{\geq x}\}$. Since, for all $i \in I_{\leq -x}$, $\zeta_i/|X_i| \leq -x/(1 + R) < -\delta'$ and for all $i \in I_{\geq x}$, $\zeta_i/|X_i| \geq x/(1 + R) > \delta'$, in these cases, $|\text{Median}(W)| \geq x/(1 + R) = x/(1 + 4\sqrt{xN}) \geq (1/5)\sqrt{x/N}$. Since $|\hat{t}^{ERM} - t^*| = |\text{Median}(W)|$, the proof is finished if (17) is proved.

Let us now prove that (17) holds on the event $\Omega_0 \cap \Omega_1$. On this event, only one η_i equals to 1. Therefore only one $|X_i|$ equals to $1 + R$ and all the others equal 1. Moreover, $\eta_i = 0$ for all $i \in I_{\delta'}$. Therefore, if $i^* \in [N]$ denotes the only index such that $\eta_{i^*} = 1$, then either $i^* \in I_{\leq -x}$ or $i^* \in I_{\geq x}$. If $i^* \in I_{\leq -x}$, on $\Omega_0 \cap \Omega_1$,

$$\begin{aligned} \sum_{i \in I_{\leq -x}} |X_i| &= |I_{\leq x}| - 1 + (1 + R) \\ &= |I_{\leq -x}| + 4\sqrt{xN} \geq |I_{\geq x}| - \sqrt{\frac{xN}{2}} + |I_{\delta'}| - \sqrt{2xN} + 4\sqrt{xN} \\ &\geq |I_{\geq x}| + |I_{\delta'}|. \end{aligned}$$

Moreover, all the $|X_i|$ equal 1 when $i \in I_{\geq x} \cup I_{\delta'}$. Therefore, $|I_{\geq x}| + |I_{\delta'}| = \sum_{i \in I_{\geq x} \cup I_{\delta'}} |X_i|$. Overall, $\sum_{i \in I_{\leq -x}} |X_i| \geq \sum_{i \in I_{\geq x} \cup I_{\delta'}} |X_i|$ which is equivalent to $\sum_{i \in I_{\leq -x}} |X_i| \geq (1/2) \sum_{i=1}^N |X_i|$. Likewise, if $i^* \in I_{\geq x}$ then $\sum_{i \in I_{\geq x}} |X_i| \geq (1/2) \sum_{i=1}^N |X_i|$. Therefore, on the event $\Omega_0 \cap \Omega_1$ (17) holds. \square

Proposition 2 shows that the distance between the ERM and t^* is larger than $(1/5)\sqrt{x/N}$ with probability at least $3/(5x)$. This probability is larger than $1 - \exp(-x/2016)$ for large values of x , which shows that the ERM does not have sub-Gaussian deviations. Let us now show that, using the same data and a number of blocks $K \asymp x$ (or using the adaptive estimator (12)), the minmax MOM estimators achieve the rate $\sqrt{x/N}$ with probability at least $1 - \exp(-x/2016)$. This will show a second advantage of minmax MOM estimators compared with ERM for heavy-tailed designs.

To apply Theorem 2 (or Theorem 3 for the adaptive estimator), we show that the local Bernstein condition is satisfied for the example of Proposition 2 and compute the complexity parameter $\tilde{r}_2(\gamma)$. We have

$$\begin{aligned} &\mathbb{E} \sup_{t \in \mathbb{R}: \mathbb{E}[(X(t-t^*))^2] \leq r^2} \left| \sum_{i=1}^N \sigma_i X_i(t - t^*) \right| \\ &= \frac{r}{\sqrt{\mathbb{E}X^2}} \mathbb{E} \left| \sum_{i=1}^N \sigma_i \epsilon_i (1 + R\eta_i) \right| \\ &\leq \frac{r\sqrt{2}}{17} \left(\mathbb{E} \left(\sum_{i=1}^N \sigma_i \epsilon_i \right)^2 + R^2 \mathbb{E} \left(\sum_{i=1}^N \sigma_i \epsilon_i \eta_i \right)^2 \right)^{1/2} \\ &\leq \frac{r\sqrt{2}}{17} \sqrt{N + R^2 N \delta} = r \sqrt{\frac{2N}{17}}. \end{aligned}$$

As a consequence, $\tilde{r}_2(\gamma) = (1/\gamma)\sqrt{2/(17N)}$ satisfies (6). We now prove Assumption 6 in this particular example. Let $K \in [N]$ be such that $K \geq 2(575)^2/(17 \times 865)$ so that, for $C_{K,r}$ is defined in (7) with $L = 1$ and A defined later,

$$C_{K,r} = A^2 \max \left(\frac{2(575)^2}{17N}, \frac{865K}{N} \right) = \frac{A^2 865K}{N},$$

Let $t \in \mathbb{R}$ be such that $\mathbb{E}[(X(t - t^*))^2] = C_{K,r}$. We have to show that $P\mathcal{L}_t \geq A\mathbb{E}[(X(t - t^*))^2]$ for some well chosen A and $P\mathcal{L}_t = \mathbb{E}[|Y - Xt| - |Y - Xt^*|]$. It follows from (39) that $P\mathcal{L}_t = \mathbb{E}[g(X, Xt) - g(X, Xt^*)]$ where $g : (x, a) \in \mathbb{R}^2 \mapsto \int \mathbf{1}_{y \geq a}(1 - F_{Y|X=x}(y))dy + (1/2)a$ and $F_{Y|X=x}$ is the cdf of Y given $X = x$. Therefore, if we denote by F the cdf of ζ , we have

$$\begin{aligned} P\mathcal{L}_t &= \mathbb{E} \int_{Xt}^{Xt^*} (1 - F_{Y|X=X}(y))dy \\ &= \frac{1 - \delta}{2} \int_{t-t^*}^0 (1 - F(y))dy + \frac{1 - \delta}{2} \int_{t^*-t}^0 (1 - F(y))dy \\ &\quad + \frac{\delta}{2} \int_{(1+R)(t-t^*)}^0 (1 - F(y))dy + \frac{\delta}{2} \int_{(1+R)(t^*-t)}^0 (1 - F(y))dy. \end{aligned}$$

Let us choose K such that $\sqrt{C_{K,r}} \leq \sqrt{\mathbb{E}X^2}\delta'$ (which holds for instance when $865A^2K \leq 17x/128$). In that case, $|t - t^*| \leq \delta'$ and so $(1 - F(y)) = (1/2 - y)$ for all $y \in [-|t - t^*|, |t - t^*|]$. We therefore have

$$\frac{1 - \delta}{2} \int_{t-t^*}^0 (1 - F(y))dy + \frac{1 - \delta}{2} \int_{t^*-t}^0 (1 - F(y))dy = \frac{(1 - \delta)(t - t^*)^2}{2}.$$

Moreover, since $(1 + R)|t - t^*| = (1 + 4\sqrt{xN})\sqrt{C_{K,r}/\mathbb{E}X^2} \leq 5\sqrt{xN}\delta' = 5x/(8\sqrt{2})$, we have

$$\frac{\delta}{2} \int_{(1+R)(t-t^*)}^0 (1 - F(y))dy + \frac{\delta}{2} \int_{(1+R)(t^*-t)}^0 (1 - F(y))dy \geq \frac{-10x\delta}{8\sqrt{2}} = \frac{-5}{4\sqrt{2}N}.$$

Assume that $(1/16)^2 865K \geq 18 * 40/\sqrt{2}$, so $|t - t^*| \geq C_{K,r}/18 \geq 40/(\sqrt{2}N)$. Then, we have

$$P\mathcal{L}_t \geq (1 - \delta)(t - t^*)^2/2 - 5/(4\sqrt{2}N) \geq (t - t^*)^2/16.$$

For $A = 1/(16 \times 18)$, this yields $P\mathcal{L}_t \geq A\mathbb{E}[X(t - t^*)^2]$, which concludes the proof of the Bernstein's assumption.

7 Simulation study

This section provides a short simulation study that illustrates our theoretical findings for the minmax MOM estimators. Let us consider the following setup: $X = (\xi_1, \dots, \xi_d)$, where $(\xi_j)_{j=1}^d$ are independent and identically distributed, with $\xi_1 \sim \mathcal{T}(5)$, and

$$\log \left(\frac{\mathbf{P}(Y = 1|X)}{\mathbf{P}(Y = -1|X)} \right) = \langle X, t^* \rangle + \epsilon$$

where $\epsilon \sim \mathcal{LN}(0, 1)$. Let $(X_i, Y_i)_{i=1}^N$ be i.i.d with the same distribution as (X, Y) . We study the minmax MOM estimator defined as:

$$\hat{t}_K^{\text{MOM}} \in \arg \min_{t \in \mathbb{R}^p} \sup_{\tilde{t} \in \mathbb{R}^p} \text{MOM}_K(\ell_t - \ell_{\tilde{t}}). \tag{18}$$

Following [24], a gradient ascent-descent step is performed on the empirical incremental risk $(t, \tilde{t}) \rightarrow P_{B_k}(\ell_t - \ell_{\tilde{t}})$ constructed on the block B_k of data realizing the median of the empirical incremental risk. Initial points $t_0 \in \mathbb{R}^d$ and $\tilde{t}_0 \in \mathbb{R}^d$ are taken at random. In logistic regression, the step sizes η and $\tilde{\eta}$ are usually chosen equal to $\|\mathbb{X}\mathbb{X}^\top\|_{\text{op}}/4N$, where \mathbb{X} is the $N \times d$ matrix with row vectors equal to $X_1^\top, \dots, X_N^\top$ and $\|\cdot\|_{\text{op}}$ denotes the operator norm. In a corrupted environment, this choice might lead to disastrous performance. This is why η and $\tilde{\eta}$ are computed at each iteration using only data in the median block: let B_k denote the median block at the current step, then one chooses $\eta = \tilde{\eta} = \|\mathbb{X}_{(k)}\mathbb{X}_{(k)}^\top\|_{\text{op}}/4|B_k|$ where $\mathbb{X}_{(k)}$ is the $|B_k| \times p$ matrix with rows given by X_i^\top for $i \in B_k$. In practice, K is chosen by robust cross-validation choice as in [24].

In a first approach and according to our theoretical results, the blocks are chosen at the beginning of the algorithm. As illustrated in Fig. 2, this first strategy has some limitations. To understand the problem, for all $k = 1, \dots, K$, let C_k denote the following set

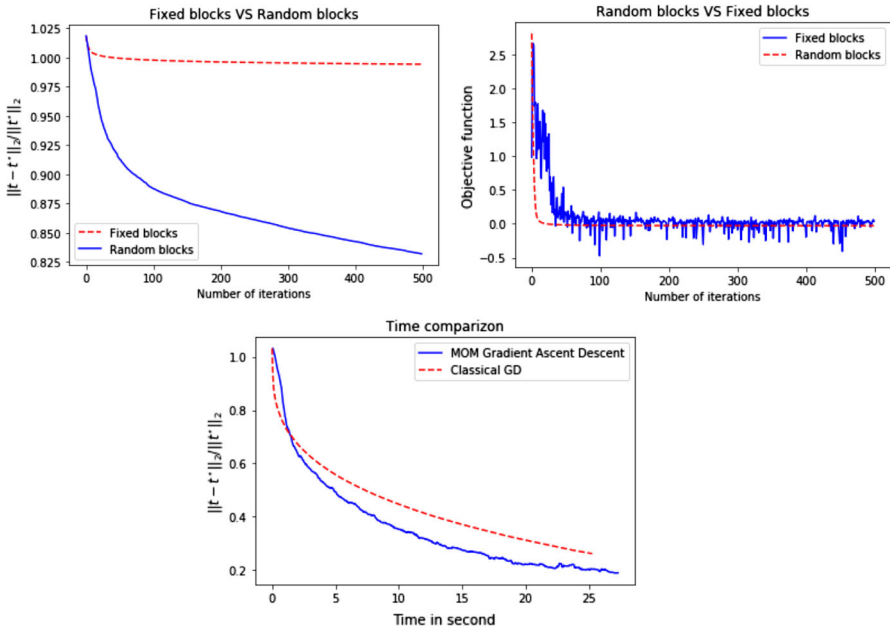


Fig. 2 Top left and right: comparison of the algorithm with fixed and changing blocks. Bottom: comparison of running time between classical gradient descent and Algorithm 1. In all simulation $N = 1000$, $p = 100$ and there is no outliers

$$C_k = \left\{ t \in \mathbb{R}^d : P_{B_k} \ell_t = \text{Median} \{ P_{B_1} \ell_t, \dots, P_{B_K} \ell_t \} \right\}.$$

If the minimum of $t \rightarrow P_{B_k} \ell_t$ lies in C_k , the algorithm typically converges to this minimum if one iteration enters C_k . As a consequence, when the minmax MOM estimator (18) lies in another cell, the algorithm does not converge to this estimator.

To bypass this issue, the partition is changed at every ascent/descent steps of the algorithm, it is chosen uniformly at random among all equipartition of the dataset. This alternative algorithm is described in Algorithm 1. In practice, changing the partition seems to widely accelerate the convergence (see Fig. 2).

Input: The number of block K , initial points t_0 and \tilde{t}_0 in \mathbb{R}^p and the stopping criterion $\epsilon > 0$
Output: An estimator of t^*

```

1 while  $\|t_i - \tilde{t}_i\|_2 \geq \epsilon$  do
2   Split the data into  $K$  disjoint blocks  $(B_k)_{k \in \{1, \dots, K\}}$  of equal sizes chosen at random:
    $B_1 \cup \dots \cup B_K = \{1, \dots, N\}$ .
3   Find  $k \in [K]$  such that  $\text{MOM}_K \ell_{t_i} = P_{B_k} \ell_{t_i}$ .
4   Compute  $\eta = \tilde{\eta} = \|\mathbb{X}_{(k)}^T \mathbb{X}_{(k)}\|_{op} / 4N$ .
5   Update  $t_{i+1} = t_i - \frac{1}{\eta} \nabla_t (P_{B_k} \ell_t)|_{t=t_i}$  and  $\tilde{t}_{i+1} = \tilde{t}_i - \frac{1}{\tilde{\eta}} \nabla_{\tilde{t}} (P_{B_k} \ell_{\tilde{t}})|_{\tilde{t}=\tilde{t}_i}$ .
6 end

```

Algorithm 1: Descent-ascent gradient method with blocks of data chosen at random at every steps.

Simulation results are gathered in Fig. 2. In these simulations, there is no outlier, $N = 1000$ and $d = 100$ with $(X_i, Y_i)_{i=1}^{1000}$ i.i.d with the same distribution as (X, Y) . Minmax MOM estimators (18) are compared with the Logistic Regression algorithm from the scikit-learn library of [38].

The upper pictures compare performance of MOM ascent/descent algorithms with fixed and changing blocks. These pictures give an example where the fixed block algorithm is stuck into local minima and another one where it does not converge. In both cases, the changing blocks version converges to t^* .

Running times of logistic regression (LR) and its MOM version (MOM LR) are compared in the lower picture of Fig. 2 in a dataset free from outliers. LR and MOM LR are coded with the same algorithm in this example, meaning that MOM gradient descent-ascent and simple gradient descent are performed with the same descent algorithm. As illustrated in Fig. 2, running each step of the gradient descent on one block only and not on the whole dataset accelerates the running time. The larger the dataset, the bigger the benefit is expected.

The resistance to outliers of logistic regression and its minmax MOM alternative are depicted in Fig. 1 in the introduction. We added an increasing number of outliers to the dataset. Outliers $\{(X_i, Y_i), i \in \mathcal{O}\}$ in this simulation are such that $X_i \sim \mathcal{LN}(0, 5)$ and $Y_i = -\text{sign}(\langle X_i, t \rangle + \epsilon_i)$, with $\epsilon_i \sim \epsilon$ as above. Figure 1 shows that logistic classification is misled by a single outlier while MOM version maintains reasonable performance with up to 50 outliers (i.e 5% of the database is corrupted).

A byproduct of Algorithm 1 is an outlier detection algorithm. Each data receives a score equal to the number of times it is selected in a median block in the random

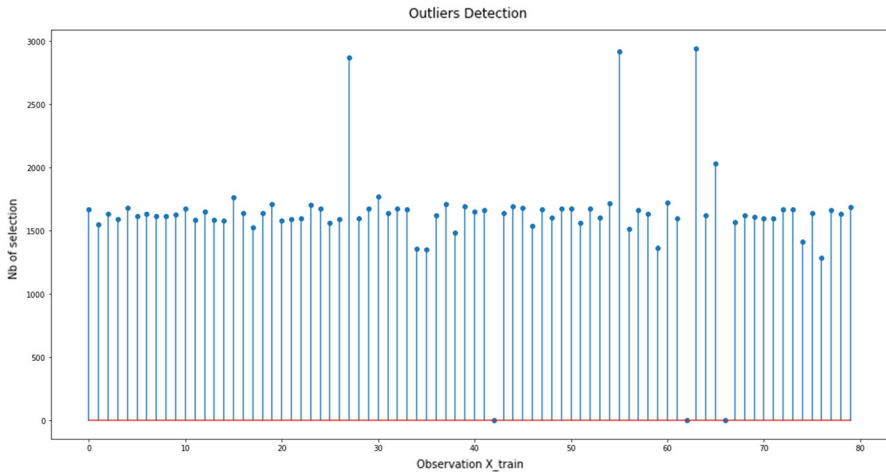


Fig. 3 Outliers Detection Procedure for $N = 100$, $p = 10$ and outliers are $i = 42, 62, 66$

choice of block version of the algorithm. The first iterations may be misleading: before convergence, the empirical loss at the current point may not reveal the centrality of the data because the current point may be far from t^* . Simulations are run with $N = 100$, $d = 10$ and 5000 iterations and therefore only the score obtained by each data in the last 4000 iterations are displayed. 3 outliers $(X_i, Y_i)_{i \in \{1,2,3\}}$ with $X_i = (10)_{j=1}^d$ and $Y_i = -\text{sign}\langle X_i, t \rangle$ have been introduced at number 42, 62 and 66. Figure 3 shows that these are not selected once.

8 Conclusion

The paper introduces a new homogeneity argument for learning problems with convex and Lipschitz losses. This argument allows to obtain estimation rates and oracle inequalities for ERM and minmax MOM estimators improving existing results. The ERM requires sub-Gaussian hypotheses on the class F with respect to the distribution of the design and a local Bernstein condition (see Theorem 1), both assumptions can be removed for minmax MOM estimators (see Theorem 5). The local Bernstein conditions provided in this article can be verified in several learning problems. In particular, it allows to derive optimal risk bounds in examples where analyses based on the small ball hypothesis fail. Minmax MOM estimators applied to convex and Lipschitz losses are efficient under weak assumptions on the outputs Y , under minimal L_2 assumptions on the class F with respect to the distribution of the design and the results are robust to the presence of few outliers in the dataset. A modification of these estimators can be implemented efficiently and confirm all these conclusions.

A Proof of Theorems 1, 2, 3 and 4

A.1 Proof of Theorem 1

The proof is splitted in two parts. First, we identify an event where the statistical behavior of the regularized estimator \hat{f}^{ERM} can be controlled. Then, we prove that this event holds with probability at least (3). Introduce $\theta = 1/(2A)$ and define the following event:

$$\Omega := \left\{ \forall f \in F \cap (f^* + r_2(\theta)B_{L_2}), \quad |(P - P_N)\mathcal{L}_f| \leq \theta r_2^2(\theta) \right\}$$

where θ is a parameter appearing in the definition of r_2 in Definition 3.

Proposition 3 *On the event Ω , one has*

$$\|\hat{f}^{ERM} - f^*\|_{L_2} \leq r_2(\theta) \quad \text{and} \quad P\mathcal{L}_{\hat{f}^{ERM}} \leq \theta r_2^2(\theta).$$

Proof By construction, \hat{f}^{ERM} satisfies $P_N\mathcal{L}_{\hat{f}^{ERM}} \leq 0$. Therefore, it is sufficient to show that, on Ω , if $\|f - f^*\|_{L_2} > r_2(\theta)$, then $P_N\mathcal{L}_f > 0$. Let $f \in F$ be such that $\|f - f^*\|_{L_2} > r_2(\theta)$. By convexity of F , there exists $f_0 \in F \cap (f^* + r_2(\theta)S_{L_2})$ and $\alpha > 1$ such that

$$f = f^* + \alpha(f_0 - f^*). \quad (19)$$

For all $i \in \{1, \dots, N\}$, let $\psi_i : \mathbb{R} \rightarrow \mathbb{R}$ be defined for all $u \in \mathbb{R}$ by

$$\psi_i(u) = \bar{\ell}(u + f^*(X_i), Y_i) - \bar{\ell}(f^*(X_i), Y_i). \quad (20)$$

The functions ψ_i are such that $\psi_i(0) = 0$, they are convex because $\bar{\ell}$ is, in particular $\alpha\psi_i(u) \leq \psi_i(\alpha u)$ for all $u \in \mathbb{R}$ and $\alpha \geq 1$ and $\psi_i(f(X_i) - f^*(X_i)) = \bar{\ell}(f(X_i), Y_i) - \bar{\ell}(f^*(X_i), Y_i)$ so that the following holds:

$$\begin{aligned} P_N\mathcal{L}_f &= \frac{1}{N} \sum_{i=1}^N \psi_i(f(X_i) - f^*(X_i)) = \frac{1}{N} \sum_{i=1}^N \psi_i(\alpha(f_0(X_i) - f^*(X_i))) \\ &\geq \frac{\alpha}{N} \sum_{i=1}^N \psi_i((f_0(X_i) - f^*(X_i))) = \alpha P_N\mathcal{L}_{f_0}. \end{aligned} \quad (21)$$

Until the end of the proof, the event Ω is assumed to hold. Since $f_0 \in F \cap (f^* + r_2(\theta)S_{L_2})$, $P_N\mathcal{L}_{f_0} \geq P\mathcal{L}_{f_0} - \theta r_2^2(\theta)$. Moreover, by Assumption 4, $P\mathcal{L}_{f_0} \geq A^{-1}\|f_0 - f^*\|_{L_2}^2 = A^{-1}r_2^2(\theta)$, thus

$$P_N\mathcal{L}_{f_0} \geq (A^{-1} - \theta)r_2^2(\theta). \quad (22)$$

From Eqs. (21) and (22), $P_N \mathcal{L}_f > 0$ since $A^{-1} > \theta$. Therefore, $\|\hat{f}^{ERM} - f^*\|_{L_2} \leq r_2^2(\theta)$. This proves the L_2 -bound.

Now, as $\|\hat{f}^{ERM} - f^*\|_{L_2} \leq r_2^2(\theta)$, $|(P - P_N)\mathcal{L}_{\hat{f}^{ERM}}| \leq \theta r_2^2(\theta)$. Since $P_N \mathcal{L}_{\hat{f}^{ERM}} \leq 0$,

$$P \mathcal{L}_{\hat{f}^{ERM}} = P_N \mathcal{L}_{\hat{f}^{ERM}} + (P - P_N)\mathcal{L}_{\hat{f}^{ERM}} \leq \theta r_2^2(\theta).$$

This shows the excess risk bound. □

Proposition 3 shows that \hat{f}^{ERM} has the risk bounds given in Theorem 1 on the event Ω . To show that Ω holds with probability (3), recall the following results from [2].

Lemma 2 [2] [Lemma 8.1] *Grant Assumptions 1 and 3 . Let $F' \subset F$ with finite L_2 -diameter $d_{L_2}(F')$. For every $u > 0$, with probability at least $1 - 2 \exp(-u^2)$,*

$$\sup_{f, g \in F'} |(P - P_N)(\mathcal{L}_f - \mathcal{L}_g)| \leq \frac{16L}{\sqrt{N}} (w(F') + u d_{L_2}(F')).$$

It follows from Lemma 2 that for any $u > 0$, with probability larger than $1 - 2 \exp(-u^2)$,

$$\begin{aligned} & \sup_{f \in F \cap (f^* + r_2(\theta) B_{L_2})} |(P - P_N)\mathcal{L}_f| \\ & \leq \sup_{f, g \in F \cap (f^* + r_2(\theta) B_{L_2})} |(P - P_N)(\mathcal{L}_f - \mathcal{L}_g)| \\ & \leq \frac{16L}{\sqrt{N}} (w((F - f^*) \cap r_2(\theta) B_{L_2}) + u d_{L_2}((F - f^*) \cap r_2(\theta) B_{L_2})) \end{aligned}$$

where $d_{L_2}((F - f^*) \cap r_2(\theta) B_{L_2}) \leq r_2(\theta)$. By definition of the complexity parameter (see Eq. (3)), for $u = \theta \sqrt{N} r_2(\theta) / (64L)$, with probability at least

$$1 - 2 \exp(-\theta^2 N r_2^2(\theta) / (16^3 L^2)), \tag{23}$$

for every f in $F \cap (f^* + r_2(\theta) B_{L_2})$,

$$|(P - P_N)\mathcal{L}_f| \leq \theta r_2^2(\theta). \tag{24}$$

Together with Proposition 3, this concludes the proof of Theorem 1.

A.2 Proof of Theorem 2

The proof is splitted in two parts. First, we identify an event Ω_K where the statistical properties of \hat{f} from Theorem 2 can be established. Next, we prove that this event

holds with probability (8). Let α, θ and γ be positive numbers to be chosen later. Define

$$C_{K,r} = \max \left(\frac{4L^2K}{\theta^2\alpha N}, \tilde{r}_2^2(\gamma) \right)$$

where the exact form of α, θ and γ are given in Eq. (33). Set the event Ω_K to be such that

$$\Omega_K = \left\{ \forall f \in F \cap \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right), \exists J \subset \{1, \dots, K\} : |J| > K/2 \right. \\ \left. \text{and } \forall k \in J, |(P_{B_k} - P)\mathcal{L}_f| \leq \theta C_{K,r} \right\}. \quad (25)$$

A.2.1 Deterministic argument

The goal of this section is to show that, on the event Ω_K , $\|\hat{f} - f^*\|_{L_2}^2 \leq C_{K,r}$ and $P\mathcal{L}_{\hat{f}} \leq 2\theta C_{K,r}$.

Lemma 3 *If there exists $\eta > 0$ such that*

$$\sup_{f \in F \setminus \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)} \text{MOM}_K(\ell_{f^*} - \ell_f) < -\eta \quad \text{and} \\ \sup_{f \in F \cap \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)} \text{MOM}_K(\ell_{f^*} - \ell_f) \leq \eta, \quad (26)$$

then $\|\hat{f} - f^*\|_{L_2}^2 \leq C_{K,r}$.

Proof Assume that (26) holds, then

$$\inf_{f \in F \setminus \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)} \text{MOM}_K[\ell_f - \ell_{f^*}] > \eta. \quad (27)$$

Moreover, if $T_K(f) = \sup_{g \in F} \text{MOM}_K[\ell_f - \ell_g]$ for all $f \in F$, then

$$T_K(f^*) = \sup_{f \in F \cap \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)} \text{MOM}_K[\ell_{f^*} - \ell_f] \\ \vee \sup_{f \in F \setminus \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)} \text{MOM}_K[\ell_{f^*} - \ell_f] \leq \eta. \quad (28)$$

By definition of \hat{f} and (28), $T_K(\hat{f}) \leq T_K(f^*) \leq \eta$. Moreover, by (27), any $f \in F \setminus \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)$ satisfies $T_K(f) \geq \text{MOM}_K[\ell_f - \ell_{f^*}] > \eta$. Therefore $\hat{f} \in F \cap \left(f^* + \sqrt{C_{K,r}} B_{L_2} \right)$. \square

Lemma 4 Grant Assumption 6 and assume that $\theta - A^{-1} < -\theta$. On the event Ω_K , (26) holds with $\eta = \theta C_{K,r}$.

Proof Let $f \in F$ be such that $\|f - f^*\|_{L_2} > C_{K,r}$. By convexity of F , there exists $f_0 \in F \cap (f^* + \sqrt{C_{K,r}}S_{L_2})$ and $\alpha > 1$ such that $f = f^* + \alpha(f_0 - f^*)$. For all $i \in \{1, \dots, N\}$, let $\psi_i : \mathbb{R} \rightarrow \mathbb{R}$ be defined for all $u \in \mathbb{R}$ by

$$\psi_i(u) = \bar{\ell}(u + f^*(X_i), Y_i) - \bar{\ell}(f^*(X_i), Y_i). \tag{29}$$

The functions ψ_i are convex because $\bar{\ell}$ is and such that $\psi_i(0) = 0$, so $\alpha\psi_i(u) \leq \psi_i(\alpha u)$ for all $u \in \mathbb{R}$ and $\alpha \geq 1$. As $\psi_i(f(X_i) - f^*(X_i)) = \bar{\ell}(f(X_i), Y_i) - \bar{\ell}(f^*(X_i), Y_i)$, for any block B_k ,

$$\begin{aligned} P_{B_k} \mathcal{L}_f &= \frac{1}{|B_k|} \sum_{i \in B_k} \psi_i(f(X_i) - f^*(X_i)) = \frac{1}{|B_k|} \sum_{i \in B_k} \psi_i(\alpha(f_0(X_i) - f^*(X_i))) \\ &\geq \frac{\alpha}{|B_k|} \sum_{i \in B_k} \psi_i((f_0(X_i) - f^*(X_i))) = \alpha P_{B_k} \mathcal{L}_{f_0}. \end{aligned} \tag{30}$$

As $f_0 \in F \cap (f^* + \sqrt{C_{K,r}}S_{L_2})$, on Ω_K , there are strictly more than $K/2$ blocks B_k where $P_{B_k} \mathcal{L}_{f_0} \geq P \mathcal{L}_{f_0} - \theta C_{K,r}$. Moreover, from Assumption 6, $P \mathcal{L}_{f_0} \geq A^{-1} \|f_0 - f^*\|_{L_2}^2 = A^{-1} C_{K,r}$. Therefore, on strictly more than $K/2$ blocks B_k ,

$$P_{B_k} \mathcal{L}_{f_0} \geq (A^{-1} - \theta) C_{K,r}. \tag{31}$$

From Eqs. (30) and (31), there are strictly more than $K/2$ blocks B_k where $P_{B_k} \mathcal{L}_f \geq (A^{-1} - \theta) C_{K,r}$. Therefore, on Ω_K , as $(\theta - A^{-1}) < -\theta$,

$$\sup_{f \in F \setminus (f^* + \sqrt{C_{K,r}}B_{L_2})} \text{MOM}_K(\ell_{f^*} - \ell_f) < (\theta - A^{-1}) C_{K,r} < -\theta C_{K,r}.$$

In addition, on the event Ω_K , for all $f \in F \cap (f^* + \sqrt{C_{K,r}}B_{L_2})$, there are strictly more than $K/2$ blocks B_k where $|(P_{B_k} - P)\mathcal{L}_f| \leq \theta C_{K,r}$. Therefore

$$\text{MOM}_K(\ell_{f^*} - \ell_f) \leq \theta C_{K,r} - P \mathcal{L}_f \leq \theta C_{K,r}.$$

□

Lemma 5 Grant Assumption 6 and assume that $\theta - A^{-1} < -\theta$. On the event Ω_K , $P \mathcal{L}_{\hat{f}} \leq 2\theta C_{K,r}$.

Proof Assume that Ω_K holds. From Lemmas 3 and 4, $\|\hat{f} - f^*\|_{L_2} \leq \sqrt{C_{K,r}}$. Therefore, on strictly more than $K/2$ blocks B_k , $P \mathcal{L}_{\hat{f}} \leq P_{B_k} \mathcal{L}_{\hat{f}} + \theta C_{K,r}$. In addition, by definition of \hat{f} and (28) (for $\eta = \theta C_{K,r}$),

$$\text{MOM}_K(\ell_{\hat{f}} - \ell_{f^*}) \leq \sup_{f \in F} \text{MOM}_K(\ell_{f^*} - \ell_f) \leq \theta C_{K,r}.$$

As a consequence, there exist at least $K/2$ blocks B_k where $P_{B_k} \mathcal{L}_{\hat{f}} \leq \theta C_{K,r}$. Therefore, there exists at least one block B_k where both $P \mathcal{L}_{\hat{f}} \leq P_{B_k} \mathcal{L}_{\hat{f}} + \theta C_{K,r}$ and $P_{B_k} \mathcal{L}_{\hat{f}} \leq \theta C_{K,r}$. Hence $P \mathcal{L}_{\hat{f}} \leq 2\theta C_{K,r}$. \square

A.2.2 Stochastic argument

This section shows that Ω_K holds with probability at least (8).

Proposition 4 *Grant Assumptions 1, 2, 5 and 6 and assume that $(1 - \beta)K \geq |\mathcal{O}|$. Let $x > 0$ and assume that $\beta(1 - \alpha - x - 8\gamma L/\theta) > 1/2$. Then Ω_K holds with probability larger than $1 - \exp(-x^2 \beta K/2)$.*

Proof Let $\mathcal{F} = F \cap (f^* + \sqrt{C_{K,r}} B_{L_2})$ and set $\phi : t \in \mathbb{R} \rightarrow I\{t \geq 2\} + (t - 1)I\{1 \leq t \leq 2\}$ so, for all $t \in \mathbb{R}$, $I\{t \geq 2\} \leq \phi(t) \leq I\{t \geq 1\}$. Let $W_k = ((X_i, Y_i))_{i \in B_k}$, $G_f(W_k) = (P_{B_k} - P)\mathcal{L}_f$. Let

$$z(f) = \sum_{k=1}^K I\{|G_f(W_k)| \leq \theta C_{K,r}\}.$$

Let \mathcal{K} denote the set of indices of blocks which have not been corrupted by outliers, $\mathcal{K} = \{k \in \{1, \dots, K\} : B_k \subset \mathcal{I}\}$ and let $f \in \mathcal{F}$. Basic algebraic manipulations show that

$$\begin{aligned} z(f) &\geq |\mathcal{K}| - \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \left(\phi(2\theta^{-1} C_{K,r}^{-1} |G_f(W_k)|) - \mathbb{E}\phi(2\theta^{-1} C_{K,r}^{-1} |G_f(W_k)|) \right) \\ &\quad - \sum_{k \in \mathcal{K}} \mathbb{E}\phi(2\theta^{-1} C_{K,r}^{-1} |G_f(W_k)|). \end{aligned}$$

By Assumptions 1 and 5, using that $C_{K,r}^2 \geq \|f - f^*\|_{L_2}^2 [(4L^2 K)/(\theta^2 \alpha N)]$,

$$\begin{aligned} &\mathbb{E}\phi(2\theta^{-1} C_{K,r}^{-1} |G_f(W_k)|) \\ &\leq \mathbb{P}\left(|G_f(W_k)| \geq \frac{\theta C_{K,r}}{2}\right) \leq \frac{4}{\theta^2 C_{K,r}^2} \mathbb{E}G_f(W_k)^2 = \frac{4}{\theta^2 C_{K,r}^2} \text{Var}(P_{B_k} \mathcal{L}_f) \\ &\leq \frac{4K^2}{\theta^2 C_{K,r}^2 N^2} \sum_{i \in B_k} \mathbb{E}[\mathcal{L}_f^2(X_i, Y_i)] \leq \frac{4L^2 K}{\theta^2 C_{K,r}^2 N} \|f - f^*\|_{L_2}^2 \leq \alpha. \end{aligned}$$

Therefore,

$$z(f) \geq |\mathcal{K}|(1 - \alpha) - \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \left(\phi(2\theta^{-1} C_{K,r}^{-1} |G_f(W_k)|) - \mathbb{E}\phi(2\theta^{-1} C_{K,r}^{-1} |G_f(W_k)|) \right). \tag{32}$$

Using Mc Diarmid’s inequality [11, Theorem 6.2], for all $x > 0$, with probability larger than $1 - \exp(-x^2|\mathcal{K}|/2)$,

$$\begin{aligned} & \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \left(\phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) - \mathbb{E}\phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) \right) \\ & \leq x|\mathcal{K}| + \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \left(\phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) - \mathbb{E}\phi(2\theta^{-1}\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) \right). \end{aligned}$$

Let $\epsilon_1, \dots, \epsilon_K$ denote independent Rademacher variables independent of the (X_i, Y_i) , $i \in \mathcal{I}$. By Giné-Zinn symmetrization argument,

$$\begin{aligned} & \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \left(\phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) - \mathbb{E}\phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) \right) \\ & \leq x|\mathcal{K}| + 2\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \epsilon_k \phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) \end{aligned}$$

As ϕ is 1-Lipschitz with $\phi(0) = 0$, using the contraction lemma [28, Chapter 4],

$$\begin{aligned} \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \epsilon_k \phi(2\theta^{-1}C_{K,r}^{-1}|G_f(W_k)|) & \leq 2\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \epsilon_k \frac{G_f(W_k)}{\theta C_{K,r}} \\ & = 2\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \epsilon_k \frac{(P_{B_k} - P)\mathcal{L}_f}{\theta C_{K,r}}. \end{aligned}$$

Let $(\sigma_i : i \in \cup_{k \in \mathcal{K}} B_k)$ be a family of independent Rademacher variables independent of $(\epsilon_k)_{k \in \mathcal{K}}$ and $(X_i, Y_i)_{i \in \mathcal{I}}$. It follows from the Giné-Zinn symmetrization argument that

$$\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{k \in \mathcal{K}} \epsilon_k \frac{(P_{B_k} - P)\mathcal{L}_f}{C_{K,r}} \leq 2\mathbb{E} \sup_{f \in \mathcal{F}} \frac{K}{N} \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{\mathcal{L}_f(X_i, Y_i)}{C_{K,r}}.$$

By the Lipschitz property of the loss, the contraction principle applies and

$$\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{\mathcal{L}_f(X_i, Y_i)}{C_{K,r}} \leq L\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{C_{K,r}}.$$

To bound from above the right-hand side in the last inequality, consider two cases 1) $C_{K,r} = \tilde{r}_2^2(\gamma)$ or 2) $C_{K,r} = 4L^2K/(\alpha\theta^2N)$. In the first case, by definition of the complexity parameter $\tilde{r}_2(\gamma)$ in (6),

$$\begin{aligned} & \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{C_{K,r}} \\ &= \mathbb{E} \sup_{f \in F: \|f - f^*\|_{L_2} \leq \tilde{r}_2(\gamma)} \frac{1}{\tilde{r}_2^2(\gamma)} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i (f - f^*)(X_i) \right| \\ &\leq \frac{\gamma |\mathcal{K}| N}{K}. \end{aligned}$$

In the second case,

$$\begin{aligned} & \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \frac{\sigma_i (f - f^*)(X_i)}{C_{K,r}} \\ &\leq \mathbb{E} \left[\sup_{\substack{f \in F: \\ \|f - f^*\|_{L_2} \leq \tilde{r}_2(\gamma)}} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \frac{\sigma_i (f - f^*)(X_i)}{\tilde{r}_2^2(\gamma)} \right| \right] \\ &\vee \sup_{\substack{f \in F: \\ \tilde{r}_2(\gamma) \leq \|f - f^*\|_{L_2} \leq \sqrt{\frac{4L^2K}{\alpha\theta^2N}}}} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{\frac{4L^2K}{\alpha\theta^2N}} \right|. \end{aligned}$$

Let $f \in F$ be such that $\tilde{r}_2(\gamma) \leq \|f - f^*\|_{L_2} \leq \sqrt{[4L^2K]/[\alpha\theta^2N]}$; by convexity of F , there exists $f_0 \in F$ such that $\|f_0 - f^*\|_{L_2} = \tilde{r}_2(\gamma)$ and $f = f^* + \alpha(f_0 - f^*)$ with $\alpha = \|f - f^*\|_{L_2} / \tilde{r}_2(\gamma) \geq 1$. Therefore,

$$\begin{aligned} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{\frac{4L^2K}{\alpha\theta^2N}} \right| &\leq \frac{1}{\tilde{r}_2(\gamma)} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{\|f - f^*\|_{L_2}} \right| \\ &= \frac{1}{\tilde{r}_2^2(\gamma)} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i (f_0 - f^*)(X_i) \right| \end{aligned}$$

and so

$$\begin{aligned} & \sup_{\substack{f \in F: \\ \tilde{r}_2(\gamma) \leq \|f - f^*\|_{L_2} \leq \sqrt{\frac{4L^2K}{\alpha\theta^2N}}}} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{\frac{4L^2K}{\alpha\theta^2N}} \right| \\ &\leq \frac{1}{\tilde{r}_2^2(\gamma)} \sup_{\|f - f^*\|_{L_2} = \tilde{r}_2(\gamma)} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i (f - f^*)(X_i) \right|. \end{aligned}$$

By definition of $\tilde{r}_2(\gamma)$, it follows that

$$\mathbb{E} \sup_{f \in \mathcal{F}} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i \frac{(f - f^*)(X_i)}{C_{K,r}} \right| \leq \frac{\gamma |\mathcal{K}| N}{K}.$$

Therefore, as $|\mathcal{K}| \geq K - |\mathcal{O}| \geq \beta K$, with probability larger than $1 - \exp(-x^2\beta K/2)$, for all $f \in F$ such that $\|f - f^*\|_{L_2} \leq \sqrt{C_{K,r}}$,

$$z(f) \geq |\mathcal{K}| \left(1 - \alpha - x - \frac{8\gamma L}{\theta} \right) > \frac{K}{2}. \tag{33}$$

□

A.2.3 End of the proof of Theorem 2

Theorem 2 follows from Lemmas 3, 4, 5 and Proposition 4 for the choice of constant

$$\theta = 1/(3A) \quad \alpha = 1/24, \quad x = 1/24, \quad \beta = 4/7 \text{ and } \gamma = 1/(575AL).$$

A.3 Proof of Theorem 3

Let $K \in [7|\mathcal{O}|/3, N]$ and consider the event Ω_K defined in (25). It follows from the proof of Lemmas 3 and 4 that $T_K(f^*) \leq \theta C_{K,r}$ on Ω_K . Setting $\theta = 1/(3A)$, on $\cap_{J=K}^N \Omega_J$, $f^* \in \hat{R}_J$ for all $J = K, \dots, N$, so $\cap_{J=K}^N \hat{R}_J \neq \emptyset$. By definition of \hat{K} , it follows that $\hat{K} \leq K$ and by definition of \tilde{f} , $\tilde{f} \in \hat{R}_K$ which means that $T_K(\tilde{f}) \leq \theta C_{K,r}$. It is proved in Lemmas 3 and 4 that on Ω_K , if $f \in F$ satisfies $\|f - f^*\|_{L_2} \geq \sqrt{C_{K,r}}$ then $T_K(f) > \theta C_{K,r}$. Therefore, $\|\tilde{f} - f^*\|_{L_2} \leq \sqrt{C_{K,r}}$. On Ω_K , since $\|\tilde{f} - f^*\|_{L_2} \leq \sqrt{C_{K,r}}$, $P\mathcal{L}_{\tilde{f}} \leq 2\theta C_{K,r}$. Hence, on $\cap_{J=K}^N \Omega_J$, the conclusions of Theorem 3 hold. Finally, by Proposition 4,

$$\mathbb{P} \left[\cap_{J=K}^N \Omega_J \right] \geq 1 - \sum_{J=K}^N \exp(-K/2016) \geq 1 - 4 \exp(-K/2016).$$

A.4 Proof of Theorem 4

The proof of Theorem 4 follows the same path as the one of Theorem 2. We only sketch the different arguments needed because of the localization by the excess loss and the lack of Bernstein condition.

Define the event Ω'_K in the same way as Ω_K in (25) where $C_{K,r}$ is replaced by $\bar{r}_2^2(\gamma)$ and the L_2 localization is replaced by the ‘‘excess loss localization’’:

$$\Omega'_K = \left\{ \forall f \in (\mathcal{L}_F)_{\bar{r}_2^2(\gamma)}, \exists J \subset \{1, \dots, K\} : |J| > K/2 \right. \\ \left. \text{and } \forall k \in J, |(P_{B_k} - P)\mathcal{L}_f| \leq (1/4)\bar{r}_2^2(\gamma) \right\} \tag{34}$$

where $(\mathcal{L}_F)_{\bar{r}_2^2(\gamma)} = \{f \in F : P\mathcal{L}_f \leq \bar{r}_2^2(\gamma)\}$. Our first goal is to show that on the event Ω'_K , $P\mathcal{L}_{\tilde{f}} \leq (1/4)\bar{r}_2^2(\gamma)$. We will then handle $\mathbb{P}[\Omega'_K]$.

Lemma 6 *Grant Assumptions 1 and 2. For every $r \geq 0$, the set $(\mathcal{L}_F)_r := \{f \in F : P\mathcal{L}_f \leq r\}$ is convex and relatively closed to F in $L_1(\mu)$. Moreover, if $f \in F$ is such that $P\mathcal{L}_f > r$ then there exists $f_0 \in F$ and $(P\mathcal{L}_f/r) \geq \alpha > 1$ such that $(f - f^*) = \alpha(f_0 - f^*)$ and $P\mathcal{L}_{f_0} = r$.*

Proof Let f and g be in $(\mathcal{L}_F)_r$ and $0 \leq \alpha \leq 1$. We have $\alpha f + (1 - \alpha)g \in F$ because F is convex and for all $x \in \mathcal{X}$ and $y \in \mathbb{R}$, using the convexity of $u \rightarrow \bar{\ell}(u + f^*(x), y)$, we have

$$\begin{aligned} & \ell_{\alpha f + (1-\alpha)g}(x, y) - \ell_{f^*}(x, y) \\ &= \bar{\ell}(\alpha(f - f^*)(x) + (1 - \alpha)(g - f^*)(x) + f^*(x), y) - \bar{\ell}(f^*(x), y) \\ &\leq \alpha(\bar{\ell}((f - f^*)(x) + f^*(x), y) - \bar{\ell}(f^*(x), y)) \\ &\quad + (1 - \alpha)(\bar{\ell}((g - f^*)(x) + f^*(x), y) - \bar{\ell}(f^*(x), y)) \\ &= \alpha(\ell_f - \ell_{f^*}) + (1 - \alpha)(\ell_g - \ell_{f^*}) \end{aligned}$$

and so $P\mathcal{L}_{\alpha f + (1-\alpha)g} \leq \alpha P\mathcal{L}_f + (1 - \alpha)P\mathcal{L}_g$. Given that $P\mathcal{L}_f, P\mathcal{L}_g \leq r$ we also have $P\mathcal{L}_{\alpha f + (1-\alpha)g} \leq r$. Therefore, $\alpha f + (1 - \alpha)g \in (\mathcal{L}_F)_r$ and $(\mathcal{L}_F)_r$ is convex.

For all $f, g \in F$, $|P\mathcal{L}_f - P\mathcal{L}_g| \leq \|f - g\|_{L_1(\mu)}$ so that $f \in F \rightarrow P\mathcal{L}_f$ is continuous onto F in $L_1(\mu)$ and therefore its level sets, such as $(\mathcal{L}_F)_r$, are relatively closed to F in $L_1(\mu)$.

Finally, let $f \in F$ be such that $P\mathcal{L}_f > r$. Define $\alpha_0 = \sup\{\alpha \geq 0 : f^* + \alpha(f - f^*) \in (\mathcal{L}_F)_r\}$. Note that $P\mathcal{L}_{f^* + \alpha(f - f^*)} \leq \alpha P\mathcal{L}_f = r$ for $\alpha = r/P\mathcal{L}_f$ so that $\alpha_0 \geq r/P\mathcal{L}_f$. Since $(\mathcal{L}_F)_r$ is relatively closed to F in $L_1(\mu)$, we have $f^* + \alpha_0(f - f^*) \in (\mathcal{L}_F)_r$ and in particular $\alpha_0 < 1$ otherwise, by convexity of $(\mathcal{L}_F)_r$, we would have $f \in (\mathcal{L}_F)_r$. Moreover, by maximality of α_0 , $f_0 = f^* + \alpha_0(f - f^*)$ is such that $P\mathcal{L}_{f_0} = r$ and the results follows for $\alpha = \alpha_0^{-1}$. \square

Lemma 7 *Grant Assumptions 1 and 2. On the event Ω'_K , $P\mathcal{L}_{\hat{f}} \leq \bar{r}_2^2(\gamma)$.*

Proof Let $f \in F$ be such that $P\mathcal{L}_f > \bar{r}_2^2(\gamma)$. It follows from Lemma 6 that there exists $\alpha \geq 1$ and $f_0 \in F$ such that $P\mathcal{L}_{f_0} = \bar{r}_2^2(\gamma)$ and $f - f^* = \alpha(f_0 - f^*)$. According to (30), we have for every $k \in \{1, \dots, K\}$, $P_{B_k}\mathcal{L}_f \geq \alpha P_{B_k}\mathcal{L}_{f_0}$. Since $f_0 \in (\mathcal{L}_F)_{\bar{r}_2^2(\gamma)}$, on the event Ω'_K , there are strictly more than $K/2$ blocks B_k such that $P_{B_k}\mathcal{L}_{f_0} \geq P\mathcal{L}_{f_0} - (1/4)\bar{r}_2^2(\gamma) = (3/4)\bar{r}_2^2(\gamma)$ and so $P_{B_k}\mathcal{L}_f \geq (3/4)\bar{r}_2^2(\gamma)$. As a consequence, we have

$$\sup_{f \in F \setminus (\mathcal{L}_F)_{\bar{r}_2^2(\gamma)}} \text{MOM}_K(\ell_{f^*} - \ell_f) \leq (-3/4)\bar{r}_2^2(\gamma). \quad (35)$$

Moreover, on the event Ω'_K , for all $f \in (\mathcal{L}_F)_{\bar{r}_2^2(\gamma)}$, there are strictly more than $K/2$ blocks B_k such that $P_{B_k}(-\mathcal{L}_f) \leq (1/4)\bar{r}_2^2(\gamma) - P\mathcal{L}_f \leq (1/4)\bar{r}_2^2(\gamma)$. Therefore,

$$\sup_{f \in (\mathcal{L}_F)_{\bar{r}_2^2(\gamma)}} \text{MOM}_K(\ell_{f^*} - \ell_f) \leq (1/4)\bar{r}_2^2(\gamma). \quad (36)$$

We conclude from (35) and (36) that $\sup_{f \in F} \text{MOM}_K(\ell_{f^*} - \ell_f) \leq (1/4)\bar{r}_2^2(\gamma)$ and that every $f \in F$ such that $P\mathcal{L}_f > \bar{r}_2^2(\gamma)$ satisfies $\text{MOM}_K(\ell_f - \ell_{f^*}) \geq (3/4)\bar{r}_2^2(\gamma)$. But, by definition of \hat{f} , we have

$$\text{MOM}_K(\ell_{\hat{f}} - \ell_{f^*}) \leq \sup_{f \in F} \text{MOM}_K(\ell_{f^*} - \ell_f) \leq (1/4)\bar{r}_2^2(\gamma).$$

Therefore, we necessarily have $P\mathcal{L}_{\hat{f}} \leq \bar{r}_2^2(\gamma)$. □

Now, we prove that Ω'_K is an exponentially large event using similar argument as in Proposition 4.

Proposition 5 *Grant Assumptions 1, 2 and 7 and assume that $(1 - \beta)K \geq |\mathcal{O}|$ and $\beta(1 - 1/12 - 32\gamma L) > 1/2$. Then Ω'_K holds with probability larger than $1 - \exp(-\beta K/1152)$.*

Sketch of proof The proof of Proposition 5 follows the same line as the one of Proposition 4. Let us precise the main differences. We set $\mathcal{F}' = (\mathcal{L}_F)_{\bar{r}_2^2(\gamma)}$ and for all $f \in \mathcal{F}'$, $z'(f) = \sum_{k=1}^K I\{|G_f(W_k)| \leq (1/4)\bar{r}_2^2(\gamma)\}$ where $G_f(W_k)$ is the same quantity as in the proof of Proposition 5. Let us consider the contraction ϕ introduced in Proposition 5. By definition of $\bar{r}_2^2(\gamma)$ and $V_K(\cdot)$, we have

$$\begin{aligned} & \mathbb{E}\phi(8(\bar{r}_2^2(\gamma))^{-1}|G_f(W_k)|) \\ & \leq \mathbb{P}\left(|G_f(W_k)| \geq \frac{\bar{r}_2^2(\gamma)}{8}\right) \leq \frac{64}{(\bar{r}_2^2(\gamma))^2} \mathbb{E}G_f(W_k)^2 = \frac{64}{(\bar{r}_2^2(\gamma))^2} \text{Var}(P_{B_k}\mathcal{L}_f) \\ & \leq \frac{64K^2}{(\bar{r}_2^2(\gamma))^2 N^2} \sum_{i \in B_k} \text{Var}_{P_i}(\mathcal{L}_f) \leq \frac{64K}{(\bar{r}_2^2(\gamma))^2 N} \sup\{\text{Var}_{P_i}(\mathcal{L}_f) : f \in \mathcal{F}', i \in \mathcal{I}\} \\ & \leq \frac{64K}{(\bar{r}_2^2(\gamma))^2 N} \sup\{\text{Var}_{P_i}(\mathcal{L}_f) : P\mathcal{L}_f \leq \bar{r}_2^2(\gamma), i \in \mathcal{I}\} \leq \frac{1}{24}. \end{aligned}$$

Using Mc Diarmid’s inequality, the Giné-Zinn symmetrization argument and the contraction lemma twice and the Lipschitz property of the loss function, such as in the proof of Proposition 4, we obtain with probability larger than $1 - \exp(-|\mathcal{K}|/1152)$, for all $f \in \mathcal{F}'$,

$$z(f) \geq |\mathcal{K}|(1 - 1/12) - \frac{32LK}{N} \mathbb{E} \sup_{f \in \mathcal{F}'} \frac{1}{\bar{r}_2^2(\gamma)} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i(f - f^*)(X_i) \right|. \tag{37}$$

Now, it remains to use the definition of $\bar{r}_2^2(\gamma)$ to bound the expected supremum in the right-hand side of (37) to get

$$\mathbb{E} \sup_{f \in \mathcal{F}'} \frac{1}{\bar{r}_2(\gamma)^2} \left| \sum_{i \in \cup_{k \in \mathcal{K}} B_k} \sigma_i (f - f^*)(X_i) \right| \leq \frac{\gamma |\mathcal{K}| N}{K}. \quad (38)$$

□

Proof of Theorem 4 The proof of Theorem 4 follows from Lemma 7 and Proposition 5 for $\beta = 4/7$ and $\gamma = 1/(768L)$. □

B Proof of Lemma 1

Proof We have

$$\frac{1}{\sqrt{N}} \mathbb{E} \sup_{f \in \mathcal{F}: \|f - f^*\|_{L_2} \leq r} \sum_{i=1}^N \sigma_i (f - f^*)(X_i) = \mathbb{E} \sup_{t \in \mathbb{R}^d: \mathbb{E}\langle t, X \rangle^2 \leq r^2} \langle t, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle.$$

Let $\Sigma = \mathbb{E} X^T X$ denote the covariance matrix of X and consider its SVD, $\Sigma = Q D Q^T$ where $Q = [Q_1 | \dots | Q_d] \in \mathbb{R}^{d \times d}$ is an orthogonal matrix and D is a diagonal $d \times d$ matrix with non-negative entries. For all $t \in \mathbb{R}^d$, we have $\mathbb{E}\langle X, t \rangle^2 = t^T \Sigma t = \sum_{j=1}^d d_j \langle t, Q_j \rangle^2$. Then

$$\begin{aligned} & \mathbb{E} \sup_{t \in \mathbb{R}^d: \sqrt{\mathbb{E}\langle t, X \rangle^2} \leq r} \langle t, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle \\ &= \mathbb{E} \sup_{t \in \mathbb{R}^d: \sqrt{\mathbb{E}\langle t, X \rangle^2} \leq r} \langle \sum_{j=1}^d \langle t, Q_j \rangle Q_j, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle \\ &= \mathbb{E} \sup_{t \in \mathbb{R}^d: \sqrt{\sum_{j=1}^d d_j \langle t, Q_j \rangle^2} \leq r} \sum_{j=1: d_j \neq 0}^d \sqrt{d_j} \langle t, Q_j \rangle \langle \frac{Q_j}{\sqrt{d_j}}, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle \\ &\leq r \mathbb{E} \sqrt{\sum_{j=1: d_j \neq 0}^d \langle \frac{Q_j}{\sqrt{d_j}}, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle^2} \leq r \sqrt{\mathbb{E} \sum_{j=1: d_j \neq 0}^d \langle \frac{Q_j}{\sqrt{d_j}}, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle^2}. \end{aligned}$$

Moreover, for any j such that $d_j \neq 0$,

$$\begin{aligned} & \mathbb{E} \langle \frac{Q_j}{\sqrt{d_j}}, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \rangle^2 \\ &= \mathbb{E} \frac{1}{N} \sum_{k,l=1}^N \sigma_l \sigma_k \langle \frac{Q_j}{\sqrt{d_j}}, X_k \rangle \langle \frac{Q_j}{\sqrt{d_j}}, X_l \rangle = \frac{1}{N} \sum_{k=1}^N \mathbb{E} \langle \frac{Q_j}{\sqrt{d_j}}, X_k \rangle^2 \end{aligned}$$

$$= \frac{1}{N} \sum_{k=1}^N \left(\frac{Q_j}{\sqrt{d_j}} \right)^T \mathbb{E} X_k^T X_k \left(\frac{Q_j}{\sqrt{d_j}} \right) = \frac{1}{N} \sum_{k=1}^N \left(\frac{Q_j}{\sqrt{d_j}} \right)^T \Sigma \left(\frac{Q_j}{\sqrt{d_j}} \right)$$

By orthonormality, $Q^T Q_j = e_j$ and $Q_j^T Q = e_j^T$, then, for any j such that $d_j \neq 0$,

$$\mathbb{E} \left\langle \frac{Q_j}{\sqrt{d_j}}, \frac{1}{\sqrt{N}} \sum_{i=1}^N \sigma_i X_i \right\rangle^2 = \frac{1}{N} \sum_{k=1}^N \frac{1}{d_j} e_j^T D e_j = 1.$$

Finally, we obtain

$$\frac{1}{\sqrt{N}} \mathbb{E} \sup_{f \in F: \|f - f^*\|_{L_2} \leq r} \sum_{i=1}^N \sigma_i (f - f^*)(X_i) \leq r \sqrt{\sum_{j=1}^d \mathbf{1}_{\{d_j \neq 0\}}} = r \sqrt{\text{Rank}(\Sigma)}$$

and therefore the fixed point $\tilde{r}_2(\gamma)$ is such that

$$\begin{aligned} \tilde{r}_2(\gamma) &= \inf \left\{ r > 0, \forall J \in \mathcal{I} : |J| \geq N/2, \right. \\ &\quad \left. \mathbb{E} \sup_{t \in \mathbb{R}^d: \sqrt{\mathbb{E}(t - t^*, X)^2} \leq r} \sum_{i \in J} \sigma_i \langle X_i, t - t^* \rangle \leq r^2 |J| \gamma \right\} \\ &\leq \inf \left\{ r > 0, \forall J \in \mathcal{I} : |J| \geq N/2, \quad r \sqrt{\text{Rank}(\Sigma)} \leq r^2 \sqrt{|J|} \gamma \right\} \\ &\leq \sqrt{\frac{\text{Rank}(\Sigma)}{2\gamma^2 N}}. \end{aligned}$$

□

C Proofs of the results of Sect. 5

We begin this section with a simple Lemma coming from the convexity of F .

Lemma 8 For any $f \in F$,

$$\lim_{t \rightarrow 0^+} \frac{R(f^* + t(f - f^*)) - R(f^*)}{t} \geq 0$$

where we recall that $R(f) = \mathbb{E}_{(X, Y) \sim P} [\ell_f(X, Y)]$.

Proof Let $t \in (0, 1)$. By convexity of F , $f^* + t(f - f^*) \in F$ and $R(f^* + t(f - f^*)) - R(f^*) \geq 0$ because f^* minimizes the risk over F . □

C.1 Proof of Theorem 6

Let $r > 0$. Let $f \in F$ be such that $\|f - f^*\|_{L_2} \leq r$. For all $x \in \mathcal{X}$ denote by $F_{Y|X=x}$ the conditional c.d.f. of Y given $X = x$. We have

$$\begin{aligned} \mathbb{E}\left[\ell_f(X, Y)|X = x\right] &= (\tau - 1) \int \mathbf{1}_{y \leq f(x)}(y - f(x))F_{Y|X=x}(\mathrm{d}y) \\ &\quad + \tau \int \mathbf{1}_{y > f(x)}(y - f(x))F_{Y|X=x}(\mathrm{d}y) \\ &= \int \mathbf{1}_{y > f(x)}(y - f(x))F_{Y|X=x}(\mathrm{d}y) \\ &\quad + (\tau - 1) \int \mathbf{1}_{\mathbb{R}}(y - f(x))F_{Y|X=x}(\mathrm{d}y). \end{aligned}$$

By Fubini's theorem,

$$\begin{aligned} \int \mathbf{1}_{z \geq f(x)}(1 - F_{Y|X=x}(z))\mathrm{d}z &= \int \mathbf{1}_{z \geq f(x)}\left(1 - \mathbb{P}(Y \leq z|X = x)\right)\mathrm{d}z \\ &= \int \mathbf{1}_{z \geq f(x)}\mathbb{E}[\mathbf{1}_{Y > z}|X = x]\mathrm{d}z \\ &= \int \int \mathbf{1}_{y > z \geq f(x)}f_{Y|X=x}(y)\mathrm{d}y\mathrm{d}z \\ &= \int \mathbf{1}_{y > f(x)}(y - f(x))f_{Y|X=x}(y)\mathrm{d}y \\ &= \int \mathbf{1}_{y > f(x)}(y - f(x))F_{Y|X=x}(\mathrm{d}y). \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}\left[\ell_f(X, Y)|X = x\right] &= \int \mathbf{1}_{y \geq f(x)}(1 - F_{Y|X=x}(y))\mathrm{d}y + (\tau - 1)\left(\int_{\mathbb{R}} yF_{Y|X=x}(\mathrm{d}y) - f(x)\right) \\ &= g(x, f(x)) + (\tau - 1) \int_{\mathbb{R}} yF_{Y|X=x}(\mathrm{d}y) \end{aligned}$$

where $g : (x, a) \in \mathcal{X} \times \mathbb{R} \rightarrow \int \mathbf{1}_{y \geq a}(1 - F_{Y|X=x}(y))\mathrm{d}y + (1 - \tau)a$. It follows that

$$P\mathcal{L}_f = \mathbb{E}[g(X, f(X)) - g(X, f^*(X))]. \quad (39)$$

Since for all $x \in \mathcal{X}$, $a \mapsto g(x, a)$ is twice differentiable, from a second order Taylor expansion we get

$$\begin{aligned}
 P\mathcal{L}_f &= \mathbb{E} \left[g(X, f(X)) - g(X, f^*(X)) \right] \\
 &= \mathbb{E} \left[\frac{\partial g(X, a)}{\partial a} (f^*(X))(f(X) - f^*(X)) \right] \\
 &\quad + \frac{1}{2} \int_{x \in \mathcal{X}} \frac{\partial^2 g(x, a)}{\partial a^2} (z_x)(f(x) - f^*(x))^2 dP_X(x)
 \end{aligned}$$

where for all $x \in \mathcal{X}$, z_x is some point in $[\min(f(x), f^*(x)), \max(f(x), f^*(x))]$. For the first order term, we have

$$\begin{aligned}
 &\mathbb{E} \left[\frac{\partial g(X, a)}{\partial a} (f^*(X))(f(X) - f^*(X)) \right] \\
 &= \mathbb{E} \lim_{t \rightarrow 0^+} \frac{g(X, f^*(X) + t(f(X) - f^*(X))) - g(X, f^*(X))}{t}.
 \end{aligned}$$

For all $x \in \mathcal{X}$, we have $[g(x, f^*(x) + t(f(x) - f^*(x))) - g(x, f^*(x))]/t \leq (2 - \tau)|f^*(x) - f(x)|$ which is integrable with respect to P_X . Thus, by the dominated convergence theorem, it is possible to interchange integral and limit and therefore using Lemma 8, we obtain

$$\begin{aligned}
 &\mathbb{E} \left[\frac{\partial g(X, a)}{\partial a} (f^*(X))(f(X) - f^*(X)) \right] \\
 &= \lim_{t \rightarrow 0^+} \mathbb{E} \frac{g(X, f^*(X) + t(f(X) - f^*(X))) - g(X, f^*(X))}{t} \\
 &= \lim_{t \rightarrow 0^+} \frac{R(f^* + t(f - f^*)) - R(f^*)}{t} \geq 0.
 \end{aligned}$$

Given that for all $x \in \mathcal{X}$, $\frac{\partial^2 g(x, a)}{\partial a^2} (z) = f_{Y|X=x}(z)$ for all $z \in \mathbb{R}$ it follows that

$$P\mathcal{L}_f \geq \frac{1}{2} \int_{x \in \mathcal{X}} f_{Y|X=x}(z_x)(f(x) - f^*(x))^2 dP_X(x).$$

Consider $A = \{x \in \mathcal{X}, |f(x) - f^*(x)| \leq (\sqrt{2}C')^{(2+\varepsilon)/\varepsilon} r\}$. Given that $\|f - f^*\|_{L_2} \leq r$, by Markov's inequality, $P(X \in A) \geq 1 - 1/(\sqrt{2}C')^{(4+2\varepsilon)/\varepsilon}$. From Assumption 10 we get

$$\frac{2P\mathcal{L}_f}{\alpha} \geq \mathbb{E}[I_A(X)(f(X) - f^*(X))^2] = \|f - f^*\|_{L_2}^2 - \mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2]. \tag{40}$$

By Holder and Markov's inequalities,

$$\begin{aligned}\mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2] &\leq (\mathbb{E}[I_{A^c}(X)])^{\varepsilon/(2+\varepsilon)} (\mathbb{E}[(f(X) - f^*(X))^{2+\varepsilon}])^{2/(2+\varepsilon)} \\ &\leq \frac{\|f - f^*\|_{L_{2+\varepsilon}}^2}{2(C')^2}.\end{aligned}$$

By Assumption 9, it follows that $\mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2] \leq \|f - f^*\|_{L_2}^2/2$ and we conclude with (40).

C.2 Proof of Theorem 7

Let $r > 0$. Let $f \in F$ be such that $\|f - f^*\|_{L_2} \leq r$. We have

$$\begin{aligned}P\mathcal{L}_f &= \mathbb{E}_X \mathbb{E} \left[\rho_H(Y - f(x)) - \rho_H(Y - f^*(x)) | X = x \right] \\ &= \mathbb{E} [g(X, f(X)) - g(X, f^*(X))]\end{aligned}$$

where $g : (x, a) \in \mathcal{X} \times \mathbb{R} = \mathbb{E}[\rho_H(Y - a) | X = x]$. Let $F_{Y|X=x}$ denote the c.d.f. of Y given $X = x$. Since for all $x \in \mathcal{X}$, $a \mapsto g(x, a)$ is twice differentiable in its second argument (see Lemma 2.1 in [15]), a second Taylor expansion yields

$$\begin{aligned}P\mathcal{L}_f &= \mathbb{E} \left[\frac{\partial g(X, a)}{\partial a} (f^*(X))(f(X) - f^*(X)) \right] \\ &\quad + \frac{1}{2} \int_{x \in \mathcal{X}} (f(x) - f^*(x))^2 \frac{\partial^2 g(x, a)}{\partial a^2} (z_x) dP_X(x)\end{aligned}$$

where for all $x \in \mathcal{X}$, z_x is some point in $[\min(f(x), f^*(x)), \max(f(x), f^*(x))]$. By Lemma 8, with the same reasoning as the one in Sect. C.1, we get

$$P\mathcal{L}_f \geq \frac{1}{2} \int_{x \in \mathcal{X}} (f(x) - f^*(x))^2 \frac{\partial^2 g(x, a)}{\partial a^2} (z_x) dP_X(x).$$

Moreover, for all $z \in \mathbb{R}$,

$$\frac{\partial^2 g(x, a)}{\partial a^2} (z) = F_{Y|X=x}(z + \delta) - F_{Y|X=x}(z - \delta).$$

Now, let $A = \{x \in \mathcal{X} : |f(x) - f^*(x)| \leq (\sqrt{2}C')^{(2+\varepsilon)/\varepsilon} r\}$. It follows from Assumption 10 that $P\mathcal{L}_f \geq (\alpha/2)\mathbb{E}[(f(X) - f^*(X))^2 I_A(X)]$. Since $\|f - f^*\|_{L_2} \leq r$, by Markov's inequality, $P(X \in A) \geq 1 - 1/(\sqrt{2}C')^{(4+2\varepsilon)/\varepsilon}$. By Holder and Markov's inequalities,

$$\begin{aligned} \mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2] &\leq (\mathbb{E}[I_{A^c}(X)])^{\varepsilon/(2+\varepsilon)} (\mathbb{E}[(f(X) - f^*(X))^{2+\varepsilon}])^{2/(2+\varepsilon)} \\ &\leq \frac{\|f - f^*\|_{L_{2+\varepsilon}}^2}{2(C')^2}. \end{aligned}$$

By Assumption 9, it follows that $\mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2] \leq \frac{\|f - f^*\|_{L_2}^2}{2}$, which concludes the proof.

C.3 Proof of Theorem 8

Let $r > 0$. Let $f \in F$ be such that $\|f - f^*\|_{L_2} \leq r$. Let $\eta(x) = P(Y = 1|X = x)$. Write first that $P\mathcal{L}_f = \mathbb{E} \left[g(X, f(X)) - g(X, f^*(X)) \right]$ where for all $x \in \mathcal{X}$ and $a \in \mathbb{R}$, $g(x, a) = \eta(x) \log(1 + \exp(-a)) + (1 - \eta(x)) \log(1 + \exp(a))$. From Lemma 8 and the same reasoning as in Sects. C.1 and C.2 we get

$$\begin{aligned} P\mathcal{L}_f &\geq \int_{x \in \mathcal{X}} \frac{\partial_2^2 g(x, a)}{\partial a^2}(z_x) \frac{(f(x) - f^*(x))^2}{2} dP_X(x) \\ &= \int_{x \in \mathcal{X}} \frac{e^{z_x}}{(1 + e^{z_x})^2} \frac{(f(x) - f^*(x))^2}{2} dP_X(x) \end{aligned}$$

for some $z_x \in [\min(f(x), f^*(x)), \max(f(x), f^*(x))]$. Now, let

$$A = \left\{ x \in \mathcal{X} : |f^*(x)| \leq c_0, |f(x) - f^*(x)| \leq (2C')^{(2+\varepsilon)/\varepsilon} r \right\}.$$

On the event A we have

$$P\mathcal{L}_f \geq \frac{e^{-c_0 - (2C')^{(2+\varepsilon)/\varepsilon} r}}{2(1 + e^{c_0 + (2C')^{(2+\varepsilon)/\varepsilon} r})^2} \mathbb{E}[I_A(X)(f(X) - f^*(X))^2]$$

Using the fact that $P(X \notin A) \leq P(|f^*(X)| > c_0) + P(|f(X) - f^*(X)| > (2C')^{(2+\varepsilon)/\varepsilon} r) \leq 2/(2C')^{(4+\varepsilon)/\varepsilon}$, we conclude with Assumption 9 and the same analysis as in the two previous proofs.

C.4 Proof of Theorem 9

Let $r > 0$ such that $r(\sqrt{2}C')^{(2+\varepsilon)/\varepsilon} \leq 1$. Let f be in F such that $\|f - f^*\|_{L_2} \leq r$. For all x in \mathcal{X} let us denote $\eta(x) = \mathbb{P}(Y = 1|X = x)$. It is easy to verify that the Bayes estimator (which is equal to the oracle) is defined as $f^*(x) = \text{sign}(2\eta(x) - 1)$. Consider the set $A = \{x \in \mathcal{X}, |f(x) - f^*(x)| \leq r(\sqrt{2}C')^{(2+\varepsilon)/\varepsilon}\}$. Since $\|f - f^*\|_{L_2} \leq r$, by Markov's inequality $\mathbb{P}(X \in A) \geq 1 - 1/(\sqrt{2}C')^{(4+2\varepsilon)/\varepsilon}$. Let x be in A . If $f^*(x) = -1$ (i.e $2\eta(x) \leq 1$) and $f(x) \leq f^*(x) = -1$ we obtain

$$\begin{aligned}\mathbb{E}[\ell_f(X, Y)|X = x] - \mathbb{E}[\ell_{f^*}(X, Y)|X = x] &= \eta(x)(1 - f(x)) - \eta(x)(1 - f^*(x)) \\ &\geq \eta(x)(f(x) - f^*(x))^2\end{aligned}$$

where we used the fact that on A , $|f(x) - f^*(x)| \leq r(\sqrt{2}C')^{(2+\varepsilon)/\varepsilon} \leq 1$. Using the same analysis for the other cases we get that

$$\begin{aligned}\mathbb{E}[\ell_f(X, Y)|X = x] - \mathbb{E}[\ell_{f^*}(X, Y)|X = x] \\ \geq \min(\eta(x), 1 - \eta(x), |1 - 2\eta(x)|)(f(x) - f^*(x))^2 \\ \geq \alpha(f(x) - f^*(x))^2\end{aligned}$$

Therefore,

$$\frac{P\mathcal{L}_f}{\alpha} \geq \mathbb{E}[I_A(X)(f(X) - f^*(X))^2] = \|f - f^*\|_{L_2}^2 - \mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2]. \quad (41)$$

By Holder and Markov's inequalities,

$$\begin{aligned}\mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2] &\leq (\mathbb{E}[I_{A^c}(X)])^{\varepsilon/(2+\varepsilon)} (\mathbb{E}[(f(X) - f^*(X))^{2+\varepsilon}])^{2/(2+\varepsilon)} \\ &\leq \frac{\|f - f^*\|_{L_{2+\varepsilon}}^2}{2(C')^2}.\end{aligned}$$

By Assumption 9, it follows that $\mathbb{E}[I_{A^c}(X)(f(X) - f^*(X))^2] \leq \frac{\|f - f^*\|_{L_2}^2}{2}$ and we conclude with (41).

References

- Alon, N., Matias, Y., Szegedy, M.: The space complexity of approximating the frequency moments. *J. Comput. Syst. Sci.* **58**(1, part 2), 137–147 (1999). Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)
- Alquier, P., Cottet, V., Lecué, G.: Estimation bounds and sharp oracle inequalities of regularized procedures with lipschitz loss functions (2017). [arXiv:1702.01402](https://arxiv.org/abs/1702.01402)
- Audibert, J.-Y., Catoni, O.: Robust linear least squares regression. *Ann. Stat.* **39**(5), 2766–2794 (2011)
- Bach, F., Jenatton, R., Mairal, J., Obozinski, G., et al.: Convex optimization with sparsity-inducing norms. *Optim. Mach. Learn.* **5**, 19–53 (2011)
- Baraud, Y., Birgé, L., Sart, M.: A new method for estimation and model selection: ρ -estimation. *Invent. Math.* **207**(2), 425–517 (2017)
- Bartlett, P.L., Bousquet, O., Mendelson, S.: Local Rademacher complexities. *Ann. Stat.* **33**(4), 1497–1537 (2005)
- Bartlett, P.L., Bousquet, O., Mendelson, S., et al.: Local rademacher complexities. *Ann. Stat.* **33**(4), 1497–1537 (2005)
- Bartlett, P.L., Mendelson, S.: Empirical minimization. *Probab. Theory Relat. Fields* **135**(3), 311–334 (2006)
- Birgé, L.: Stabilité et instabilité du risque minimax pour des variables indépendantes équidistribuées. *Ann. Inst. H. Poincaré Probab. Stat.* **20**(3), 201–223 (1984)
- Boucheron, S., Bousquet, O., Lugosi, G.: Theory of classification: a survey of some recent advances. *ESAIM Probab. Stat.* **9**, 323–375 (2005)

11. Boucheron, S., Lugosi, G., Massart, P.: Concentration Inequalities. Oxford University Press, Oxford (2013). A nonasymptotic theory of independence, with a foreword by Michel Ledoux
12. Bubeck, S.: Convex optimization: algorithms and complexity. *Found. Trends@ Mach. Learn.* **8**(3–4), 231–357 (2015)
13. Catoni, O.: Challenging the empirical mean and empirical variance: a deviation study. *Ann. Inst. Henri Poincaré Probab. Stat.* **48**(4), 1148–1185 (2012)
14. Devroye, L., Lerasle, M., Lugosi, G., Oliveira, R.I., et al.: Sub-Gaussian mean estimators. *Ann. Stat.* **44**(6), 2695–2725 (2016)
15. Elsener, A., van de Geer, S.: Robust low-rank matrix estimation (2016). [arXiv:1603.09071](https://arxiv.org/abs/1603.09071)
16. Han, Q., Wellner, J.A.: Convergence rates of least squares regression estimators with heavy-tailed errors. *Ann. Statist.* **47**(4), 2286–2319 (2019)
17. Huber, P.J., Ronchetti, E.: Robust statistics. In: *International Encyclopedia of Statistical Science*, pp. 1248–1251. Springer, New York (2011)
18. Jerrum, M.R., Valiant, L.G., Vazirani, V.V.: Random generation of combinatorial structures from a uniform distribution. *Theor. Comput. Sci.* **43**(2–3), 169–188 (1986)
19. Koltchinskii, V.: Local Rademacher complexities and oracle inequalities in risk minimization. *Ann. Stat.* **34**(6), 2593–2656 (2006)
20. Koltchinskii, V.: Empirical and Rademacher processes. In: *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*, pp. 17–32. Springer, New York (2011)
21. Koltchinskii, V.: Oracle inequalities in empirical risk minimization and sparse recovery problems, volume 2033 of *Lecture Notes in Mathematics*. Springer, Heidelberg (2011). Lectures from the 38th Probability Summer School held in Saint-Flour, 2008, École d'Été de Probabilités de Saint-Flour (Saint-Flour Probability Summer School)
22. Koltchinskii, V., Mendelson, S.: Bounding the smallest singular value of a random matrix without concentration. *Int. Math. Res. Not. IMRN* **23**, 12991–13008 (2015)
23. Lecué, G., Lerasle, M.: Learning from mom's principles: Le cam's approach. *Stochast. Process. Appl.* (2018). [arXiv:1701.01961](https://arxiv.org/abs/1701.01961)
24. Lecué, G., Lerasle, M.: Robust machine learning by median-of-means: theory and practice. *Ann. Stat.* (2017). [arXiv:1711.10306](https://arxiv.org/abs/1711.10306)
25. Lecué, G., Mendelson, S.: Performance of empirical risk minimization in linear aggregation. *Bernoulli* **22**(3), 1520–1534 (2016)
26. Lecué, G., Lerasle, M., Mathieu, T.: Robust classification via mom minimization (2018). [arXiv:1808.03106](https://arxiv.org/abs/1808.03106)
27. Ledoux, M.: The Concentration of Measure Phenomenon, Volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (2001)
28. Ledoux, M., Talagrand, M.: *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, New York (2013)
29. Lugosi, G., Mendelson, S.: Risk minimization by median-of-means tournaments. *J. Eur. Math. Soc.* (2019). [arXiv:1608.00757](https://arxiv.org/abs/1608.00757)
30. Lugosi, G., Mendelson, S.: Regularization, sparse recovery, and median-of-means tournaments (2017). [arXiv:1701.04112](https://arxiv.org/abs/1701.04112)
31. Lugosi, G., Mendelson, S.: Sub-gaussian estimators of the mean of a random vector (2017). To appear in *Ann. Stat.* [arXiv:1702.00482](https://arxiv.org/abs/1702.00482)
32. Mammen, E., Tsybakov, A.B.: Smooth discrimination analysis. *Ann. Stat.* **27**(6), 1808–1829 (1999)
33. Mendelson, S.: Learning without concentration. In: *Conference on Learning Theory*, pp. 25–39 (2014)
34. Mendelson, S.: Learning without concentration. *J. ACM* **62**(3), Art. 21, 25 (2015)
35. Mendelson, S.: On multiplier processes under weak moment assumptions. In: *Geometric Aspects of Functional Analysis*, Volume 2169 of *Lecture Notes in Math.*, pp. 301–318. Springer, Cham (2017)
36. Mendelson, S., Pajor, A., Tomczak-Jaegermann, N.: Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geom. Funct. Anal.* **17**(4), 1248–1282 (2007)
37. Nemirovsky, A.S., Yudin, D.B.: *Problem Complexity and Method Efficiency in Optimization*. A Wiley-Interscience Publication. Wiley, New York (1983). Translated from the Russian and with a preface by E. R. Dawson, Wiley-Interscience Series in Discrete Mathematics
38. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)

39. Saumard, A.: On optimality of empirical risk minimization in linear aggregation. *Bernoulli* **24**(3), 2176–2203 (2018)
40. Talagrand, M.: Upper and lower bounds for stochastic processes, volume 60 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics (Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics)*. Springer, Heidelberg (2014). Modern methods and classical problems
41. Tsybakov, A.B.: Optimal aggregation of classifiers in statistical learning. *Ann. Stat.* **32**(1), 135–166 (2004)
42. van de Geer, S.: *Estimation and Testing Under Sparsity*, Volume 2159 of *Lecture Notes in Mathematics*. Springer, Cham (2016). Lecture notes from the 45th Probability Summer School held in Saint-Flour, 2015, École d'Été de Probabilités de Saint-Flour (Saint-Flour Probability Summer School)
43. Vapnik, V.: *The Nature of Statistical Learning Theory*. Springer, New York (1995)
44. Vapnik, V.: *Statistical Learning Theory*, vol. 1. Wiley, New York (1998)
45. Zhou, W.-X., Bose, K., Fan, J., Liu, H.: A new perspective on robust m-estimation: finite sample theory and applications to dependence-adjusted multiple testing. *Ann. Stat.* **46**(5), 1904–1931 (2018). <https://doi.org/10.1214/17-AOS1606>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.