The Annals of Statistics
2020, Vol. 48, No. 2, 906–931
https://doi.org/10.1214/19-AOS1828

ROBUST MACHINE LEARNING BY MEDIAN-OF-MEANS: THEORY AND PRACTICE

BY GUILLAUME LECUÉ[1] AND MATTHIEU LERASLE[2]

[1]*ENSAE, École Polytechnique, guillaume.lecue@ensae.fr*
[2]*Mathematics Department, CNRS, University Paris Sud Orsay, matthieu.lerasle@math.u-psud.fr*

Median-of-means (MOM) based procedures have been recently introduced in learning theory (Lugosi and Mendelson (2019); Lecué and Lerasle (2017)). These estimators outperform classical least-squares estimators when data are heavy-tailed and/or are corrupted. None of these procedures can be implemented, which is the major issue of current MOM procedures (*Ann. Statist.* **47** (2019) 783–794).

In this paper, we introduce minmax MOM estimators and show that they achieve the same sub-Gaussian deviation bounds as the alternatives (Lugosi and Mendelson (2019); Lecué and Lerasle (2017)), both in small and high-dimensional statistics. In particular, these estimators are efficient under moments assumptions on data that may have been corrupted by a few outliers.

Besides these theoretical guarantees, the definition of minmax MOM estimators suggests simple and systematic modifications of standard algorithms used to approximate least-squares estimators and their regularized versions. As a proof of concept, we perform an extensive simulation study of these algorithms for robust versions of the LASSO.

1. Introduction. Consider the least-squares regression problem where, given a dataset $(X_i, Y_i)_{i\in\{1,\dots,N\}}$ of points in $\mathcal{X}\times\mathbb{R}$ and a new input $X\in\mathcal{X}$, one wants to predict the associated real valued output $Y\in\mathbb{R}$. A classical approach is to consider (X,Y) as a random variable with values in $\mathcal{X}\times\mathbb{R}$ and, given a set F of functions $f:\mathcal{X}\to\mathbb{R}$, to look for the oracle in F, which is defined by

$$f^*\in\operatorname*{argmin}_{f\in F} P\big(Y-f(X)\big)^2.$$

To estimate f^*, we have a dataset $(X_i, Y_i)_{i\in\{1,\dots,N\}}$ for which there exists a partition $\{1,\dots,N\}=\mathcal{O}\cup\mathcal{I}$ such that data $(X_i, Y_i)_{i\in\mathcal{I}}$ are *inliers* or *informative* and data $(X_i, Y_i)_{i\in\mathcal{O}}$ are "outliers" in the sense that *nothing* is assumed on these data. On inliers, one grants independence and finiteness of some moments, allowing for "heavy-tailed" data. Moreover, departing from the independent and identically distributed (i.i.d.) setup, we also allow inliers to have different distributions than (X,Y). We assume that, for all $i\in\mathcal{I}$ and all $f\in F$,

$$\mathbb{E}\big[\big(Y_i-f^*(X_i)\big)\big(f-f^*\big)(X_i)\big]=\mathbb{E}\big[\big(Y-f^*(X)\big)\big(f-f^*\big)(X)\big],$$

$$\mathbb{E}\big[\big(f-f^*\big)^2(X_i)\big]=\mathbb{E}\big[\big(f-f^*\big)^2(X)\big].$$

These assumptions imply that the distribution P of (X,Y) and the distribution P_i of (X_i, Y_i) for $i\in\mathcal{I}$ induce the same L^2-geometry on $F-f^*=\{f-f^*: f\in F\}$ and, therefore, in particular, that the oracles w.r.t. P and P_i for any $i\in\mathcal{I}$ are the same. Of course, the sets $\mathcal{O}$ and $\mathcal{I}$ are unknown to the statistician.

Regression problems with possibly heavy-tailed inliers cannot be handled by classical least-squares estimators, which are particular instances of empirical risk minimizers (ERM)

Received July 2018; revised February 2019.
MSC2010 subject classifications. Primary 60K35, 62G08; secondary 62C20, 62G05, 62G20.
Key words and phrases. Empirical processes, high-dimensional statistics.

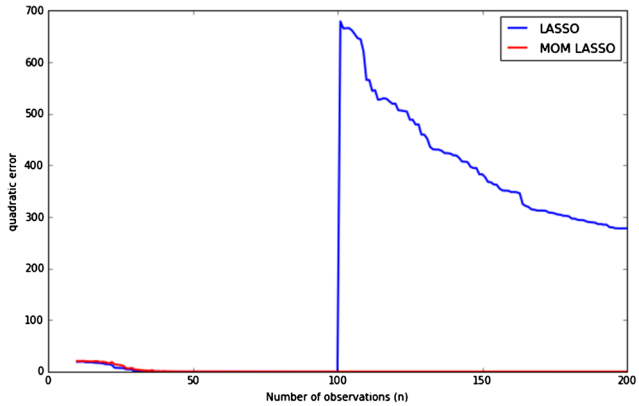

FIG. 1. *Estimation error of the LASSO* (*blue curve*) *and MOM LASSO* (*red curve*) *after one outliers was added at observation* 100.

of Vapnik [65]. Least-squares estimators have sub-Gaussian deviations under stronger assumptions, such as boundedness [44] or sub-Gaussian [37] assumptions on the noise and the design. In this paper, the main hypothesis is the *small ball assumption* of [32, 48] which says that $L^2(P)$ and $L_1(P)$ norms are equivalent over $F - f^*$; see Section 3.1 for details. Although sometimes restrictive [26, 56], this assumption does not involve high moment conditions unnecessary for the problem to make sense.

Least-squares estimators and their regularized versions are also useless in corrupted environments. This has been known for a long time and can easily be checked in practice. Figure 1, for example, shows estimation bounds of the LASSO [58] on a dataset containing a single outlier in the outputs.

These restrictions of least-squares estimators gave rise in the 1960s to the theory of robust statistics of John Tukey [59, 60], Peter Huber [27, 28] and Frank Hampel [24, 25]. The most classical alternatives to least-squares estimators are M-estimators, which are ERM based on loss functions $\ell_f(X, Y)$ less sensitive to outliers than the square loss, such as a truncated version of the square loss. The idea is that, while $(Y_i - f(X_i))^2$ can be very large for some outliers data and influence all the empirical mean $N^{-1}\sum_{i=1}^N (Y_i - f(X_i))^2$, the influence of these anomalies will be asymptotically null if $\ell_f(X_i, Y_i)$ is bounded. Recent works study deviation properties of M-estimators: [2, 21, 22, 69] considered the Huber-loss in linear regression with heavy-tailed noise and sub-Gaussian design. They obtain minimax optimal deviation bounds in this setting. The limitation on the design is not surprising: it is well known that M-estimators using loss functions such as Huber or L_1 loss are not robust to outliers in the inputs X_i. This problem is called the"leverage points problem" [29]. In a slightly different approach than M-estimation, [6] proposed a minmax estimator based on losses introduced in [18] in a least-squares regression framework and prove optimal sub-Gaussian bounds under a L_2 assumption on the noise and a L_4/L_2 assumption on the design, which is close to the assumptions we grant on inliers.

This paper focuses on Median-of-means (MOM) [1, 30, 53], which provide alternatives to M-estimators. MOM estimators of the real valued expectation $\mathbb{E}[Z]$ are built as follows: the dataset $Z_1, \dots, Z_N$ is partitioned into blocks $(Z_i)_{i\in B_k}$, $k = 1, \dots, K$ of the same cardinality. The MOM estimator is the median of the K empirical means constructed on each block:

$$\mathrm{MOM}_K(Z) = \mathrm{median}\left\{\frac{1}{|B_k|}\sum_{i\in B_k} Z_i, k = 1, \dots, K\right\}.$$

Sub-Gaussian properties of these estimators can be found in [20, 40].

As in [35, 43], MOM estimators are used to estimate real valued increments of square risks $P[(Y-f(X))^2-(Y-g(X))^2]$, where $f,g\in F$. This construction does not require a notion of median in dimension larger than 1, contrary to "geometric median-of-means" approach presented in [50, 51]. In [35, 43], each $f\in F$ receives a score which is the $L^2(P)$-diameter $\Delta(f)$ of the set $\mathcal{B}(f)$, where $g\in\mathcal{B}(f)$ if $\mathrm{MOM}_K(\ell_f-\ell_g)<0$. The approach of [35, 43] requires therefore an evaluation of the diameter of the sets $\mathcal{B}(f)$ for all $f\in F$, which makes the procedure impossible to implement.

This paper presents an alternative to [35, 43] which relies on the following minmax formulation. By linearity of P, f^* is a solution of

$$f^*\in\operatorname*{argmin}_{f\in F}\sup_{g\in F}P\big[\big(Y-f(X)\big)^2-\big(Y-g(X)\big)^2\big].$$

Replacing the real valued means $P[(Y-f(X))^2-(Y-g(X))^2]$ in this equation by their MOM estimators produces the minmax MOM estimators of f^* which are rigorously introduced in Section 2.3. Compared with [35, 43], minmax MOM estimators do not require an estimation of L^2-distances between elements in F and are therefore simpler to define. Minmax strategies have also been considered in [6] and [9, 10]. The idea of building estimators of f^* from estimators of increments goes back to seminal works by Le Cam [33, 34] and was further developed by Birgé with the T-estimators [14]. In Le Cam and Birgé's works, the authors used "robust tests" to compare densities f and g and deduce from these an alternative to the nonrobust maximum likelihood estimators. Baraud [8] showed that robust tests could be obtained by estimating the difference of Hellinger risks of f and g and used a variational formula to build these new tests. Finally, Baraud, Birgé and Sart [10] used Baraud's estimators of increments in a minmax procedure to build ρ-estimators.

The first aim of this paper is to show that minmax MOM estimators satisfy the same sub-Gaussian deviation bounds as other MOM estimators [35, 42]. The analysis of minmax MOM estimators is conceptually and technically simpler: an adaptation of Lemmas 5.1 and 5.5 in [43] or Lemmas 2 and 3 [35] is sufficient to prove sub-Gaussian bound for minmax MOM estimators while a robust estimation (based on MOM estimates) of the $L^2(P)$-metric was required in [35, 42].

Another advantage of the minmax MOM approach lies in the Lepski-step (see Theorem 2), which selects adaptively the number K of blocks. This step is way easier to implement and to study than the one presented in [35], as only one confidence region is sufficient to grant adaptation with respect to the excess risk, the regularization and L^2 norms. Recall that, in corrupted environments, a data-driven choice of K has to be performed since K must be larger than twice the (unknown) number of outliers. Note that the idea of aggregating estimators built on blocks of data and selecting the number of blocks by Lepski's method was already present in Birgé [13], proof of Theorem 1. It was also used in [20] to build "multiple-δ" sub-Gaussian estimators of univariate means.

In our opinion, the most interesting feature of the minmax formulation is that it suggests a generic method to modify descent algorithms designed to approximate ERM and their regularized versions and make them efficient even if run on corrupted datasets. Let us give a rough presentation of a "MOM version" of descent algorithms: at each time-step t, all empirical means $P_{B_k}(Y-f_t(X))^2$ for $k=1,\dots,K$ are evaluated and one computes the index $k_{\mathrm{med}}\in[K]$ of the block such that

$$P_{B_{k_{\mathrm{med}}}}\big(Y-f_t(X)\big)^2=\mathrm{med}\big\{P_{B_k}\big(Y-f_t(X)\big)^2,k=1,\dots,K\big\}.$$

The descent direction is the opposite gradient $-\nabla(f\to P_{B_{k_{\mathrm{med}}}}(Y-f(X))^2)_{|f=f_t}$. This descent algorithm can be turned into a descent-ascent algorithm approximating minmax MOM estimators. Section 5 presents several examples of modifications of classical algorithms.

In practice, these basic algorithms perform poorly when applied on a *fixed* partition of the dataset. However, empirical performance are improved when the partition is chosen uniformly at random at each descent step of the algorithm; cf. Section 6.2. In particular, the shuffling step prevents the algorithms to converge to local minimaxima. Besides, randomized algorithms define a notion of depth of data: each time a data belongs to the median block, its "score" is incremented by 1. The higher the final score is, the deeper is the data. This notion of depth is based on the risk function which is natural in a learning framework and should probably be investigated more carefully in future works. It also suggests an empirical definition of outliers and, therefore, an outliers detection algorithm. This by-product is presented in Section 6.2.

The paper is organized as follows. Section 2 introduces the framework and presents the minmax MOM estimator, Section 3 details the main theoretical results. These are illustrated in Section 4 on some classical problems of machine learning. Many robust versions of standard optimization algorithms are presented in Section 5. An extensive simulation study illustrating our results is performed in Section 6. Proofs of the main results, complementary theorems showing minimax optimality of our bounds are postponed to the Supplementary Material [36].

2. Setting. Let $\mathcal{X}$ denote a measurable space. Let $(X_i, Y_i)_{i\in\{1,\dots,N\}}$, (X, Y) denote random variables taking values in $\mathcal{X}\times\mathbb{R}$. Let P denote the distribution of (X, Y) and, for $i \in \{1,\dots,N\}$, let P_i denote the distribution of (X_i, Y_i). Let F denote a convex class of functions $f : \mathcal{X} \to \mathbb{R}$ and suppose that $\mathbb{E}[Y^2] < \infty$. For any $Q \in \{P, (P_i)_{i\in[N]}\}$ and any $p \geq 1$, let L_Q^p denote the set of functions f such that the norm $\|f\|_{L_Q^p} = (Q|f|^p)^{1/p}$, where $Qg = \mathbb{E}_{Z\sim Q}[g(Z)]$. Assume that $F \subset L_P^2$. For any $(x, y) \in \mathcal{X}\times\mathbb{R}$, let $\ell_f(x, y) = (y - f(x))^2$ denote the square loss and let f^* denote an oracle

$$f^* \in \operatorname*{argmin}_{f\in F} P\ell_f \quad \text{where } \forall g \in L_P^1,\ Pg = \mathbb{E}\big[g(X, Y)\big]. \tag{1}$$

Let $R(f) = P\ell_f$ denote the risk. The goal is to build estimators $\hat{f}$ satisfying: with probability at least $1-\delta$,

$$R(\hat{f}) \leq \min_{f\in F} R(f) + r_N^{(1)} \quad \text{and} \quad \big\|\hat{f} - f^*\big\|_{L_P^2} \leq r_N^{(2)}.$$

The residue $r_N^{(1)}$ of the oracle inequality, the estimation rate $r_N^{(2)}$ and the confidence level δ should be as small as possible. Oracle inequalities provide risk bounds for the estimation the regression function $\overline{f}(x) = \mathbb{E}[Y|X = x]$: $R(\hat{f}) \leq R(f^*) + r_N^{(1)}$ is equivalent to

$$\|\overline{f} - \hat{f}\|_{L_P^2}^2 \leq \|\overline{f} - f^*\|_{L_P^2}^2 + r_N^{(1)}.$$

Finally, let $\|\cdot\|$ be a norm defined on the span of F; $\|\cdot\|$ will be used as a regularization norm to induce some low-dimensional structure or some regularity, such as the ℓ_1 or SLOPE norm (see Section 4).

2.1. *Minmaximization.* The oracle f^* is solution of the minmax problem:

$$f^* \in \operatorname*{argmin}_{f\in F} P\ell_f = \operatorname*{argmin}_{f\in F} \sup_{g\in F} P(\ell_f - \ell_g). \tag{2}$$

Any estimator of real valued expectations $P\ell_f$ or $P(\ell_f - \ell_g)$ can be plugged in (2) to obtain estimators of f^*. Plugging the empirical means (in both the min and the minmax problems) yields the classical ERM over F, for example. In general, plugging nonlinear (robust or not) estimators of the mean in the minmax problem or in the min problem in (2) does not yield the

same estimator of f^* though. The main advantage of the minmax formulation is that it allows to bound the risk of the estimator using the complexity of F around f^*. This "localization" idea is central to derive optimal (fast) rates for the ERM [15, 31, 44] and cannot be used directly when empirical means are simply replaced by nonlinear estimators of the mean in a minimization formulation.

2.2. *MOM estimators.* Let K denote an integer smaller than $N/2$ and let $B_1, \dots, B_K$ denote a partition of $[N] = \{1, \dots, N\}$ into blocks of equal size N/K (w.l.o.g. we assume that K divides N). For all functions $\mathcal{L} : \mathcal{X} \times \mathbb{R} \to \mathbb{R}$ and $k \in [K] = \{1, \dots, K\}$, let $P_{B_k}\mathcal{L} = |B_k|^{-1} \sum_{i \in B_k} \mathcal{L}(X_i, Y_i)$.

For all $\alpha \in (0, 1)$ and real numbers $x_1, \dots, x_K$, denote by $\mathcal{Q}_\alpha(x_1, \dots, x_K)$ the set of α-quantiles of $\{x_1, \dots, x_K\}$:

$$\big\{u \in \mathbb{R} : \big|\{k \in [K] : x_k \geq u\}\big| \geq (1-\alpha)\ell, \big|\{k \in [K] : x_k \leq u\}\big| \geq \alpha\ell\big\}$$

and let $Q_\alpha(x)$ denote any point in $\mathcal{Q}_\alpha(x_1, \dots, x_K)$. For $x = (x_1, \dots, x_K) \in \mathbb{R}^K$ and $t \in \mathbb{R}$, we say that $Q_\alpha(x) \geq t$ when there exists $J \subset [K]$ such that $|J| \geq (1-\alpha)K$ and for all $k \in J, x_k \geq t$; we write $Q_\alpha(x) \leq t$ if there exists $J \subset [K]$ such that $|J| \geq \alpha K$ and for all $k \in J, x_k \leq t$.

Let $y = (y_1, \dots, y_K) \in \mathbb{R}^K$. We write $Q_{1/2}(x-y) \leq Q_{3/4}(x) - Q_{1/4}(y)$ when there exist $u, l \in \mathbb{R}$ such that $Q_{1/2}(x-y) \leq u - l$, $Q_{3/4}(x) \leq u$ and $Q_{1/4}(y) \geq l$.

DEFINITION 1. Let $\alpha \in (0, 1)$, $K \in [N]$. For any $\mathcal{L} : \mathcal{X} \times \mathbb{R} \to \mathbb{R}$ the *α-quantile on K blocks of $\mathcal{L}$* is $Q_{\alpha,K}(\mathcal{L}) = Q_\alpha((P_{B_k}\mathcal{L})_{k\in[K]})$. In particular, the Median-of-Means (MOM) of $\mathcal{L}$ on K blocks is defined as $\mathrm{MOM}_K(\mathcal{L}) = Q_{1/2,K}(\mathcal{L})$. For all $f, g \in F$, the *MOM estimator on K blocks of the loss increment from g to f* is defined by

$$T_K(g, f) = \mathrm{MOM}_K(\ell_f - \ell_g)$$

and, for a given regularization parameter $\lambda \geq 0$, its regularized version is

$$T_{K,\lambda}(g, f) = \mathrm{MOM}_K(\ell_f - \ell_g) + \lambda\big(\|f\| - \|g\|\big).$$

2.3. *Minmax MOM estimators.* Minmax MOM estimators are obtained by replacing the unknown expectations $P(\ell_f - \ell_g)$ in (2) by their MOM estimators.

DEFINITION 2. For any $K \in [N/2]$, let

$$\hat{f}_K \in \operatorname*{argmin}_{f\in F} \max_{g\in F} T_K(g, f) \quad \text{and} \quad \hat{f}_{K,\lambda} \in \operatorname*{argmin}_{f\in F} \max_{g\in F} T_{K,\lambda}(g, f). \tag{3}$$

We shall provide results for $\hat{f}_{K,\lambda}$ only in the main text. The estimators $\hat{f}_K$ are studied in the Supplementary Material [36] in Section 7.

REMARK 1 ($K = 1$ and ERM). If one chooses $K = 1$, then for all $f, g \in F$, $T_K(g, f) = P_N(\ell_f - \ell_g)$ and it is straightforward to check that $\hat{f}_K$ and $\hat{f}_{K,\lambda}$ are respectively the Empirical risk Minimization (ERM) and its regularized version (RERM).

3. Assumptions and main results. Denote by $\{\mathcal{O}, \mathcal{I}\}$ a partition of $[N]$ and by $|\mathcal{O}|$ the cardinality of $\mathcal{O}$. On $(X_i, Y_i)_{i\in\mathcal{O}}$, *no assumptions* is granted, these data are *outliers*. They may not be independent, nor independent from the remaining data (not even random). (X_i, Y_i), $i \in \mathcal{I}$ are called *inliers* or *informative* data. They are hereafter assumed *independent*. The sets $\mathcal{O}, \mathcal{I}$ are unknown.

3.1. *Assumptions*. The main assumptions involve first and second moments of the functions in F and Y under the distributions $P, (P_i)_{i\in\mathcal{I}}$.

ASSUMPTION 1. For all $f \in F$ and all $i \in \mathcal{I}$,

$$P_i(f-f^*)^2 = P(f-f^*)^2 \quad \text{and} \quad P_i[\zeta(f-f^*)] = P[\zeta(f-f^*)],$$

where $\zeta(x,y) = (y - f^*(x))$ for all $x \in \mathcal{X}$ and $y \in \mathbb{R}$.

Assumption 1 holds in the i.i.d. framework, with $\mathcal{I} = [N]$ but it covers also other cases where inliers follow different distributions (see, for instance, multimodal datasets such as in [45] or heteroscedastic noise [4]). It is also possible to weaken Assumption 1 such as in [35]. The second assumption bounds the correlation between $\zeta_i = Y_i - f^*(X_i)$ and the shifted class $F - f^*$.

ASSUMPTION 2. There exists $\theta_m > 0$ such that, for any $i \in \mathcal{I}$ and $f \in F$,

$$\operatorname{var}(\zeta_i(f-f^*)(X_i)) \le \theta_m^2 \|f-f^*\|_{L_P^2}^2.$$

Assumption 2 holds when data are i.i.d. and $Y - f^*(X)$ has uniformly bounded L^2-moments conditionally to X. This last assumption holds when $Y - f^*(X)$ is independent of X and has a L^2-moment bounded by θ_m. Assumption 2 also holds if, for all $i \in \mathcal{I}$, $\|\zeta\|_{L_{P_i}^4} \le \theta_2 < \infty$—where $\zeta(x,y) = y - f^*(x)$ for all $x \in \mathcal{X}$ and $y \in \mathbb{R}$—and, for every $f \in F$, $\|f-f^*\|_{L_{P_i}^4} \le \theta_1 \|f-f^*\|_{L_P^2}$. Actually, in this case,

$$\sqrt{\operatorname{var}(\zeta_i(f-f^*)(X_i))} \le \|\zeta\|_{L_{P_i}^4} \|f-f^*\|_{L_{P_i}^4} \le \theta_1\theta_2 \|f-f^*\|_{L_P^2},$$

so Assumption 2 holds for $\theta_m = \theta_1\theta_2$. The third assumption states that the norms L_P^2 and L_P^1 are equivalent over $F - f^*$.

ASSUMPTION 3. There exists $\theta_0 \ge 1$ such that for all $f \in F$ and $i \in \mathcal{I}$,

$$\|f-f^*\|_{L_P^2} \le \theta_0 \|f-f^*\|_{L_{P_i}^1}.$$

Under Assumption 1, $\|f-f^*\|_{L_{P_i}^1} \le \|f-f^*\|_{L_{P_i}^2} = \|f-f^*\|_{L_P^2}$ for all $f \in F$ and $i \in \mathcal{I}$, hence, Assumptions 1 and 3 imply that the norms $L_P^1, L_P^2, L_{P_i}^2, L_{P_i}^1, i \in \mathcal{I}$ are equivalent over $F - f^*$. Assumption 3 is equivalent to the small ball property (cf. [32, 48]); see Proposition 1 in [35].

3.2. *Complexity measures*. For all $\rho, r \ge 0$, let

$$B(f^*,\rho) = \{f \in F : \|f-f^*\| \le \rho\},$$
$$B_2(f^*,r) = \{f \in F : \|f-f^*\|_{L_P^2} \le r\}.$$

DEFINITION 3. Let $(\epsilon_i)_{i\in[N]}$ be independent random variables uniformly distributed in $\{-1,1\}$, independent from $(X_i,Y_i)_{i=1}^N$. For all $f \in F$, $r > 0$ and $\rho \in (0,+\infty]$, let

$$B_{\mathrm{reg}}(f,\rho,r) = \{g \in F : \|g-f\|_{L_P^2} \le r, \|g-f\| \le \rho\}.$$

Let $\zeta_i = Y_i - f^*(X_i)$ for all $i \in \mathcal{I}$ and for $\gamma_Q, \gamma_M > 0$ define $r_Q(\rho,\gamma_Q)$ as

$$\inf\left\{r > 0 : \sup_{J\subset\mathcal{I},|J|\ge N/2} \mathbb{E} \sup_{f\in B_{\mathrm{reg}}(f^*,\rho,r)} \left|\frac{1}{|J|}\sum_{i\in J}\epsilon_i(f-f^*)(X_i)\right| \le \gamma_Q r\right\},$$

and $r_M(\rho,\gamma_M)$ as

$$\inf\Big\{r>0:\sup_{J\subset\mathcal{I},|J|\geq N/2}\mathbb{E}\sup_{f\in B_{\mathrm{reg}}(f^*,\rho,r)}\Big|\frac{1}{|J|}\sum_{i\in J}\epsilon_i\zeta_i(f-f^*)(X_i)\Big|\leq\gamma_M r^2\Big\}.$$

Let $\rho\to r(\rho,\gamma_Q,\gamma_M)$ be a continuous and nondecreasing function such that for every $\rho>0$, $r(\rho)=r(\rho,\gamma_Q,\gamma_M)\geq\max\{r_Q(\rho,\gamma_Q),r_M(\rho,\gamma_M)\}$.

It follows from Lemma 2.3 in [37] that r_M and r_Q are continuous and nondecreasing functions, that depend on f^*. According to [37], for appropriate choice of γ_Q, γ_M, $r(\rho)=\max(r_M(\rho,\gamma_M),r_Q(\rho,\gamma_Q))$ is the minimax rate of convergence over $B(f^*,\rho)$. Note also that r_Q and r_M are well defined when $|\mathcal{I}|\geq N/2$, meaning that at least half data should be informative.

3.3. *The sparsity equation.* Risk bounds follow from upper bounds on $T_{K,\lambda}(f,f^*)$ for functions f far from f^* either in L_P^2-norm or for the regularization norm $\|\cdot\|$. Let $f\in F$ and let $\rho=\|f-f^*\|$. When $\|f-f^*\|_{L_P^2}$ is small, $T_{K,\lambda}$ has to be bounded from above by $\lambda(\|f^*\|-\|f\|)$. To bound $\|f^*\|-\|f\|$ from below, introduce the subdifferentials of $\|\cdot\|$. Let $(E^*,\|\cdot\|^*)$ be the dual normed space of $(E,\|\cdot\|)$ and for all $f\in F$, let

$$\big(\partial\|\cdot\|\big)_f=\big\{z^*\in E^*:\forall h\in E,\|f+h\|\geq\|f\|+z^*(h)\big\}.$$

For any $\rho>0$, let H_ρ denote the set of functions "close" to f^* in L_P^2 and at distance ρ from f^* in regularization norm and let $\Gamma_{f^*}(\rho)$ denote the set of subdifferentials of all vectors close to f^*:

$$\Gamma_{f^*}(\rho)=\bigcup_{f\in F:\|f-f^*\|\leq\rho/20}\big(\partial\|\cdot\|\big)_f$$

and $H_\rho=\{f\in F:\|f-f^*\|=\rho$ and $\|f-f^*\|_{L_P^2}\leq r(\rho)\}$. If there exists f^{**} such that $\|f^*-f^{**}\|\leq\rho/20$ and $(\partial\|\cdot\|)_{f^{**}}$ is almost all the unit dual sphere, then $\|f\|-\|f^{**}\|$ is large for any $f\in H_\rho$ so $\|f\|-\|f^*\|\geq\|f\|-\|f^{**}\|-\|f^*-f^{**}\|$ is large as well. Formally, for all $\rho>0$, let

$$\Delta(\rho)=\inf_{f\in H_\rho}\sup_{z^*\in\Gamma_{f^*}(\rho)}z^*\big(f-f^*\big).$$

The sparsity equation, introduced in [39], quantifies these notions of "large."

DEFINITION 4. A radius $\rho>0$ is said to satisfy the *sparsity equation* when $\Delta(\rho)\geq 4\rho/5$.

If ρ^* satisfies the sparsity equation, so do all $\rho\geq\rho^*$. Let

$$\rho^*=\inf\Big(\rho>0:\Delta(\rho)\geq\frac{4\rho}{5}\Big).$$

If $\rho\geq 20\|f^*\|$, then $0\in\Gamma_{f^*}(\rho)$. Moreover, $(\partial\|\cdot\|)_0$ is the unit ball of $(E^*,\|\cdot\|^*)$, so $\Delta(\rho)=\rho$. This implies that any $\rho\geq 20\|f^*\|$ satisfies the sparsity equation. This simple observation can be used to get "complexity-dependent rates of convergence" [38].

3.4. *Main results.* The first results give risk bounds for $\widehat{f}_{K,\lambda}$. Similar bounds have been obtained for other MOM estimators [35, 42].

THEOREM 1. *Grant Assumptions 1, 2 and 3 and let r_Q, r_M denote the complexity functions introduced in Definition 3. Assume that $N \geq 384\theta_0^2$ and $|\mathcal{O}| \leq N/(768\theta_0^2)$. Let ρ^* be solution to the sparsity equation from Definition 4. Let $\epsilon = 1/(833\theta_0^2)$ and $r^2(\cdot)$ is defined in Definition 3 for $\gamma_Q = (384\theta_0)^{-1}$ and $\gamma_M = \epsilon/192$. Let K^* denote the smallest integer such that*

$$K^* \geq \frac{N\epsilon^2}{384\theta_m^2} r^2(\rho^*).$$

For any $K \geq K^*$, define the radius ρ_K and the regularization parameter as

$$r^2(\rho_K) = \frac{384\theta_m^2}{\epsilon^2}\frac{K}{N} \quad \text{and} \quad \lambda = \frac{16\epsilon r^2(\rho_K)}{\rho_K}.$$

Then, for all $K \in [\max(K^*, 8|\mathcal{O}|), N/(96\theta_0^2)]$, with probability larger than $1 - 4\exp(-7K/9216)$, the estimator $\widehat{f}_{K,\lambda}$ defined in Section 2.3 satisfies

$$\|\widehat{f}_{K,\lambda} - f^*\| \leq 2\rho_K, \qquad \|\widehat{f}_{K,\lambda} - f^*\|_{L_P^2} \leq r(2\rho_K),$$

$$R(\widehat{f}_{K,\lambda}) \leq R(f^*) + (1 + 52\epsilon)r^2(2\rho_K).$$

The function r is used to define the regularization parameter in Theorem 1, so it cannot depend on f^*. When r_M, r_Q depend on f^*, r should be a computable upper bound independent from f^*. The best rates of estimation and prediction that follow from Theorem 1 are obtained for $K = K^*$ when $|\mathcal{O}| \leq K^*/8 \sim Nr^2(\rho^*)$. In that case, it is proved in Section 4 on two examples that the rate ρ_{K^*} and the residue $r(2\rho_{K^*})$ are minimax optimal. In a setup where data only induce the same L^2 metric as P and may have been corrupted by up to $K^*/8 \sim Nr^2(\rho^*)$ outliers, Theorem 1 shows that our estimators achieve the sub-Gaussian deviations bounds of the ERM when data are i.i.d. with a noise ζ independent of X and both X and ζ have Gaussian distributions (see Section 8 in the Supplementary Material [36]).

3.4.1. *Adaptive choice of K.* In Theorem 1, all rates depend on K, which has to be larger than the number of outliers and $Nr^2(\rho^*)$. These quantities are unknown in general, for instance, $Nr^2(\rho^*) \sim s\log(ed/s)$ in high-dimensional statistics where s is the unknown sparsity parameter. This section presents an adaptive choice of K inspired from Lepski's method that allows to bypass this issue. However, this construction requires the knowledge of constants θ_0 and θ_m (see Section 6.1 for a fully data driven choice of K in practice).

For all $J \in [K]$, $\lambda > 0$, $f \in F$ and $c_{ad} > 0$, let

$$\mathcal{C}_{J,\lambda}(f) = \sup_{g \in F} T_{J,\lambda}(g, f) \quad \text{and} \quad \hat{R}_{J,c_{ad}} = \left\{ f \in F : \mathcal{C}_{J,\lambda}(f) \leq \frac{c_{ad}}{\theta_0^2} r^2(\rho_J) \right\}.$$

Let $\hat{K}_{c_{ad}} = \inf\{K \in [1, N/(96\theta_0^2)] : \bigcap_{J=K}^{N/(96\theta_0^2)} \hat{R}_{J,c_{ad}} \neq \varnothing\}$ and

$$\widehat{f}_{c_{ad}} \in \bigcap_{J=\hat{K}_{c_{ad}}}^{N/(96\theta_0^2)} \hat{R}_{J,c_{ad}}. \tag{4}$$

The following theorem gives risk bounds for these estimators. Bounds in regularization and L_P^2 norms have been proved for Le Cam test estimators in [35]. To the best of our knowledge, adaptive bounds in excess risk have never been proved before.

THEOREM 2. *Grant the assumptions of Theorem 1. Choose* $c_{ad} = 18/833$ *in* (4) *and let* $\epsilon = (833\theta_0^2)^{-1}$. *For any* $K \in [\max(K^*, 8|\mathcal{O}|), N/(96\theta_0^2)]$, *with probability larger than*

$$1 - 4\exp(-K/2304) = 1 - 4\exp\bigl(-\epsilon^2 N r^2(\rho_K)/884736\bigr)$$

one has $\|\widehat{f}_{c_{ad}} - f^*\| \le 2\rho_K$, $\|\widehat{f}_{c_{ad}} - f^*\|_{L_P^2} \le r(2\rho_K)$ *and*

$$R(\widehat{f}_{c_{ad}}) \le R\bigl(f^*\bigr) + (1 + 52\epsilon) r^2(2\rho_K).$$

In particular, for $K = K^*$, *we have* $r(2\rho_{K^*}) = \max(r(2\rho^*), \sqrt{|\mathcal{O}|/N})$.

Theorem 2 shows that $\widehat{f}_{c_{ad}}$ achieves similar performance as $\widehat{f}_{K,\lambda}$ simultaneously for all K from K^* to $O(N)$. For $K = K^*$, these rates match the optimal minimax rates of convergence; see Section 4. The main difference with Theorem 1 is that the knowledge of K^* and $|\mathcal{O}|$ is not necessary to design $\widehat{f}_{c_{ad}}$. This is very useful in applications where these quantities are typically unknown. Moreover, both the construction and the analysis are much simpler for $\widehat{f}_{c_{ad}}$ than the adaptive estimator in [35] since they are based on the analysis of confidence regions for $\mathcal{C}_{J,\lambda}$ only, instead of multiple criteria in [35].

REMARK 2 (Deviation parameter). Note that $r(\cdot)$ can be any continuous, nondecreasing function such that $r(\rho) \ge \max(r_Q(\rho, \gamma_Q), r_M(\rho, \gamma_M))$. In particular, if $r_* : \rho \to \max(r_Q(\rho, \gamma_Q), r_M(\rho, \gamma_M))$ is continuous, as it is clearly nondecreasing, then for every $x > 0$, $r(\rho) = \max(r_Q(\rho, \gamma_Q), r_M(\rho, \gamma_M)) + x/N$ is another nondecreasing upper bound. Therefore, one can derive results similar to Theorem 2 but with an extra confidence parameter: for all $x > 0$, with probability at least $1 - 4\exp(-c_0 N r_*^2(\rho_{K^*}) + c_0 x)$,

$$\|\widehat{f}_{c_{ad}} - f^*\| \le 2\rho_K, \qquad \|\widehat{f}_{c_{ad}} - f^*\|_{L_P^2} \le r_*(2\rho_K) + \frac{x}{N}$$

$$R(\widehat{f}_{c_{ad}}) \le R\bigl(f^*\bigr) + (1 + 52\epsilon)\biggl(r^2(2\rho_K) + \frac{x}{N}\biggr).$$

In that case, $\widehat{f}_{c_{ad}}$ depends on x since $\lambda = 16\epsilon(r_*(\rho_K) + x/N)/\rho_K$.

4. Examples of applications. This section presents two examples of regularization in high-dimensional statistics: the ℓ_1 and the SLOPE norms.

4.1. *The LASSO.* The LASSO is obtained when $F = \{\langle t, \cdot\rangle : t \in \mathbb{R}^d\}$ and the regularization function is the ℓ_1-norm:

$$\hat{t} \in \operatorname*{argmin}_{t\in\mathbb{R}^d}\biggl(\frac{1}{N}\sum_{i=1}^N \bigl(\langle t, X_i\rangle - Y_i\bigr)^2 + \lambda\|t\|_1\biggr) \quad \text{where } \|t\|_1 = \sum_{i=1}^d |t_i|.$$

Even if recent advances show some limitations of LASSO [54, 62, 67], it remains the benchmark estimator in high-dimensional statistics because a high-dimensional parameter space does not significantly affect its performance as long as t^* is sparse. One can refer to [12, 41, 47, 52, 61, 63, 64] for estimation and sparse oracle inequalities, [7, 46, 68] for support recovery results; more results and references on LASSO can be found in [17, 31].

4.2. *SLOPE.* SLOPE is an estimator introduced in [16, 57]. The class F is still $F = \{\langle t, \cdot\rangle : t \in \mathbb{R}^d\}$ and the regularization function is defined for parameters $\beta_1 \ge \beta_2 \ge \cdots \ge \beta_d > 0$ by $\|t\|_{\mathrm{SLOPE}} = \sum_{i=1}^d \beta_i t_i^\sharp$, where $(t_i^\sharp)_{i=1}^d$ denotes the nonincreasing rearrangement of $(|t_i|)_{i=1}^d$. SLOPE norm is a weighted ℓ_1-norm that coincide with ℓ_1-norm when $(\beta_1, \ldots, \beta_d) = (1, \ldots, 1)$.

4.3. *Classical results for LASSO and SLOPE.* Typical results for LASSO and SLOPE have been obtained when data are i.i.d. with sub-Gaussian design X and, most of the time, sub-Gaussian noise ζ as well.

DEFINITION 5. Let ℓ_2^d be a d-dimensional inner product space and let X be a random variable with values in ℓ_2^d. We say that X is isotropic when for every $t \in \ell_2^d$, $\|\langle X,t\rangle\|_{L_P^2} = \|t\|_{\ell_2^d}^2$ and it is L-sub-Gaussian if for every $t \in \ell_2^d$ and every $p \geq 2$, $\|\langle X,t\rangle\|_{L_P^p} \leq L\sqrt{p}\|\langle X,t\rangle\|_{L_P^2}$.

The covariance structure of an isotropic random variable coincides with the inner product in ℓ_2^d. If X is a L-sub-Gaussian random vector, the L_P^p norms of all linear forms do not grow faster than the L_P^p norm of a Gaussian variable. When dealing with the LASSO and SLOPE, the natural Euclidean structure is used in $\mathbb{R}^d$.

ASSUMPTION 4.

1. Data are i.i.d. (in particular, $|\mathcal{I}| = N$ and $|\mathcal{O}| = 0$, that is, there is no outlier),
2. X is isotropic and L-sub-Gaussian,
3. for $f^* = \langle t^*, \cdot\rangle$, $\xi = Y - f^*(X) \in L_P^{q_0}$ for some $q_0 > 2$.

Assumption 4 requires a L^{q_0} for $q_0 > 2$ moment on the noise. LASSO and SLOPE still achieve optimal rates of convergence under this assumption but with a severely deteriorated probability estimate.

THEOREM 3 (Theorem 1.4 in [39]). *Grant Assumption 4. Let $s \in [d]$. Assume that $N \geq c_1 s\log(ed/s)$ and that there is some $v \in \mathbb{R}^d$ supported on at most s coordinates for which $\|t^* - v\|_1 \leq c_2\|\xi\|_{L_P^{q_0}} s\sqrt{\log(ed)/N}$. The Lasso estimator $\hat{t}$ with regularization parameter $\lambda = c_3\|\xi\|_{L_P^{q_0}}\sqrt{\log(ed)/N}$ is such that with probability at least*

$$1 - \frac{c_4\log^{q_0} N}{N^{q_0/2-1}} - 2\exp\bigl(-c_5 s\log(ed/s)\bigr) \tag{5}$$

for every $1 \leq p \leq 2$

$$\bigl\|\hat{t} - t^*\bigr\|_p \leq c_6\|\xi\|_{L_P^{q_0}} s^{1/p}\sqrt{\frac{\log(ed)}{N}}.$$

The constants $(c_j)_{j=1}^6$ depend only on L and q_0.

Theorem 3 shows that LASSO achieves its optimal rate (cf. [12]) if t^* is close to a sparse vector and the noise ζ may be heavy tailed and may not be independent from X. On the other hand, the dataset cannot contain outliers and the data should be i.i.d. with sub-Gaussian design matrix X.

Turning to SLOPE, recall the following result for the regularization norm $\Psi(t) = \sum_{j=1}^d \beta_j t_j^\sharp$ when $\beta_j = C\sqrt{\log(ed/j)}$.

THEOREM 4 (Theorem 1.6 in [39]). *Consider the SLOPE under Assumption 4. Assume that $N \geq c_1 s\log(ed/s)$ and that there is $v \in \mathbb{R}^d$ such that $|\operatorname{supp}(v)| \leq s$ and $\Psi(t^* - v) \leq c_2\|\xi\|_{L_P^{q_0}} s\log(ed/s)/\sqrt{N}$. The SLOPE estimator with $\lambda = c_3\|\xi\|_{L_P^{q_0}}/\sqrt{N}$ satisfies, with probability at least (5),*

$$\Psi\bigl(\hat{t} - t^*\bigr) \leq c_4\|\xi\|_{L_P^{q_0}}\frac{s}{\sqrt{N}}\log\biggl(\frac{ed}{s}\biggr), \qquad \bigl\|\hat{t} - t^*\bigr\|_2^2 \leq c_5\|\xi\|_{L_P^{q_0}}^2\frac{s}{N}\log\biggl(\frac{ed}{s}\biggr).$$

The constants $(c_j)_{j=1}^5$ depend only on L and q_0.

4.4. *Minmax MOM LASSO and SLOPE.* In this section, Theorem 2 is applied to the set F of linear functionals indexed by $\mathbb{R}^d$ with regularization functions being either the ℓ_1 or the SLOPE norm. The aim is to show that the results from Section 4.3 hold and are sometimes even improved by MOM versions of LASSO and SLOPE under weaker assumptions and with a better probability deviation. Start with the new set of assumptions.

ASSUMPTION 5. Denote by $(e_j)_{j=1}^d$ the canonical basis of $\mathbb{R}^d$ and assume

1. $(X, Y), (X_i, Y_i)_{i\in\mathcal{I}}$ are i.i.d.
2. X is isotropic and for every $t \in \mathbb{R}^d$, $p \in [C_0 \log(ed)]$ and $j \in [d]$, $\|\langle X, e_j\rangle\|_{L_P^p} \le L\sqrt{p}\|\langle X, e_j\rangle\|_{L_P^2}$,
3. $\xi = Y - \langle t^*, X\rangle \in L_P^{q_0}$ for some $q_0 > 2$.
4. there exists θ_0 such that for all $t \in \mathbb{R}^d$, $\|\langle X, t\rangle\|_{L_P^2} \le \theta_0\|\langle X, t\rangle\|_{L_P^1}$,
5. there exists θ_m such that $\operatorname{var}(\xi\langle X, t\rangle) \le \theta_m\|\langle X, t\rangle\|_{L_P^2}$.

In order to apply Theorem 2, we have to compute the fixed-point functions $r_Q(\cdot)$, $r_M(\cdot)$ and solve the sparsity equation in both cases. To compute the fixed point functions, recall the definition of Gaussian mean widths: for a set $V \subset \mathbb{R}^d$, the Gaussian mean width of V is defined as

$$\ell^*(V) = \mathbb{E}\sup_{v\in V}\langle G, v\rangle \quad \text{where } G \sim \mathcal{N}_d(0, I_d). \tag{6}$$

The dual norm of the ℓ_1^d-norm is the ℓ_∞^d-norm which is 1-unconditional with respect to the canonical basis of $\mathbb{R}^d$ [49], Definition 1.4. Therefore, [49], Theorem 1.6, applies under the following assumption.

ASSUMPTION 6. There exist constants $q_0 > 2$, C_0 and L such that $\xi \in L_P^{q_0}$, X is isotropic and for every $j \in [d]$ and $1 \le p \le C_0 \log d$, $\|\langle X, e_j\rangle\|_{L_P^p} \le L\sqrt{p}\|\langle X, e_j\rangle\|_{L_P^2}$.

Under Assumption 6, if $\sigma = \|\xi\|_{L_P^{q_0}}$, [49], Theorem 1.6, shows that for $\zeta_i = Y_i - \langle X_i, t^*\rangle$ and for every $\rho > 0$,

$$\mathbb{E}\sup_{v\in\rho B_1^d\cap r B_2^d}\left|\sum_{i\in[N]}\epsilon_i\langle v, X_i\rangle\right| \le c_2\sqrt{N}\ell^*(\rho B_1^d \cap r B_2^d),$$

$$\mathbb{E}\sup_{v\in\rho B_1^d\cap r B_2^d}\left|\sum_{i\in[N]}\epsilon_i\zeta_i\langle v, X_i\rangle\right| \le c_2\sigma\sqrt{N}\ell^*(\rho B_1^d \cap r B_2^d).$$

Local Gaussian mean widths $\ell^*(\rho B_1^d \cap r B_2^d)$ are bounded from above in [39], Lemma 5.3, and computations of $r_M(\cdot)$ and $r_Q(\cdot)$ follow

$$r_M^2(\rho) \lesssim_{L,q_0,\gamma_M} \begin{cases} \sigma^2\dfrac{d}{N} & \text{if } \rho^2 N \ge \sigma^2 d^2, \\ \rho\sigma\sqrt{\dfrac{1}{N}\log\left(\dfrac{e\sigma d}{\rho\sqrt{N}}\right)} & \text{otherwise,} \end{cases}$$

$$r_Q^2(\rho) \begin{cases} = 0 & \text{if } N \gtrsim_{L,\gamma_Q} d, \\ \lesssim_{L,\gamma_Q} \dfrac{\rho^2}{N}\log\left(\dfrac{c(L,\gamma_Q)d}{N}\right) & \text{otherwise.} \end{cases}$$

Therefore, one can take

$$r^2(\rho) \sim_{L,q_0,\gamma_Q,\gamma_M} \begin{cases} \max\left(\rho\sigma\sqrt{\frac{1}{N}\log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)}, \frac{\sigma^2 d}{N}\right) & \text{if } N \gtrsim_L d, \\ \max\left(\rho\sigma\sqrt{\frac{1}{N}\log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)}, \frac{\rho^2}{N}\log\left(\frac{d}{N}\right)\right) & \text{otherwise.} \end{cases} \tag{7}$$

Now we turn to a solution of the sparsity equation for the ℓ_1^d-norm. This equation has been solved in [39], Lemma 4.2, we recall this result.

LEMMA 1. *If there exists* $v \in \mathbb{R}^d$ *such that* $v \in t^* + (\rho/20)B_1^d$ *and* $|\operatorname{supp}(v)| \le c\rho^2/r^2(\rho)$, *then*

$$\Delta(\rho) = \inf_{h \in \rho S_1^{d-1} \cap r(\rho) B_2^d} \sup_{g \in \Gamma_{t^*}(\rho)} \langle h, g - t^* \rangle \ge \frac{4\rho}{5},$$

where S_1^{d-1} *is the unit sphere of the* ℓ_1^d*-norm and* B_2^d *is the unit Euclidean ball in* $\mathbb{R}^d$.

As a consequence, if $N \gtrsim s\log(ed/s)$ and if there exists a s-sparse vector in $t^* + (\rho/20)B_1^d$, Lemma 1 and the choice of $r(\cdot)$ in (7) imply that for $\sigma = \|\xi\|_{L^{q_0}}$,

$$\rho^* \sim_{L,q_0} \sigma s\sqrt{\frac{1}{N}\log\left(\frac{ed}{s}\right)} \quad \text{and} \quad r^2(\rho^*) \sim \frac{\sigma^2 s}{N}\log\left(\frac{ed}{s}\right)$$

then ρ^* satisfies the sparsity equation and $r^2(\rho^*)$ is the rate of convergence of the LASSO for $\lambda \sim r^2(\rho^*)/\rho^* \sim \|\xi\|_{L_P^{q_0}}\sqrt{\log(ed/s)/N}$. This choice of λ requires to know the sparsity parameter s. That is the reason why we either need to choose a larger value for the $r(\cdot)$ function as in [39]—this results in the suboptimal $\sqrt{\log(ed)/N}$ rates of convergence from Theorem 3—or to use an adaptation step as in Section 3.4.1. This results in the better minimax rate $\sqrt{\log(ed/s)/N}$. Finally, one needs to compute the radii ρ_K and $\lambda \sim r^2(\rho_K)/\rho_K$. Let $K \in [N]$ and $\sigma = \|\xi\|_{L^{q_0}}$. The equation $K = cr(\rho_K)^2 N$ is solved by

$$\rho_K \sim_{L,q_0} \frac{K}{\sigma}\sqrt{\frac{1}{N}\log^{-1}\left(\frac{\sigma^2 d}{K}\right)} \tag{8}$$

for the $r(\cdot)$ function defined in (7). Therefore,

$$\lambda \sim \frac{r^2(\rho_K)}{\rho_K} \sim_{L,q_0} \sigma\sqrt{\frac{1}{N}\log\left(\frac{e\sigma d}{\rho_K\sqrt{N}}\right)} \sim_{L,q_0} \sigma\sqrt{\frac{1}{N}\log\left(\frac{e\sigma^2 d}{K}\right)}. \tag{9}$$

The regularization parameter depends on the $L_P^{q_0}$-norm of ξ. This parameter is unknown in practice. Nevertheless, it can be replaced by an estimator in the regularization parameter as in [23], Sections 5.4 and 5.6.2.

The following result follows from Theorem 2 together with the previous computation of ρ^*, $r_Q(\cdot)$, $r_M(\cdot)$, $r(\cdot)$ and λ.

THEOREM 5. *Grant Assumption 5. Let* $s \in [d]$. *Assume that* $N \ge c_1 s\log(ed/s)$ *and that there is some* $v \in \mathbb{R}^d$ *supported on at most* s *coordinates for which* $\|t^* - v\|_1 \le c_2\|\xi\|_{L^{q_0}} s\sqrt{\log(ed)/N}$. *Assume that* $|\mathcal{I}| \ge N/2$ *and* $|\mathcal{O}| \le c_3 s\log(ed/s)$. *The MOM-LASSO estimator* $\hat{t}$ *with the adaptively chosen number of blocks* K *(and* λ*) from Section 3.4.1 satisfies, with probability at least* $1 - c_4\exp(-c_5 s\log(ed/s))$, *for every* $1 \le p \le 2$,

$$\|\hat{t} - t^*\|_p \le c_6\|\xi\|_{L_{q_0}} s^{1/p}\sqrt{\frac{1}{N}\log\left(\frac{ed}{s}\right)},$$

where $(c_j)_{j=1}^6$ *depends only on* θ_0, θ_m *and* q_0.

PROOF. It follows from Theorem 2, the computation of $r(\rho_K)$ from (7) and ρ_K in (8) that with probability at least $1-c_0\exp(-cr(\rho_K)^2N/\overline{C})$, $\|\hat{t}-t^*\|_1 \le \rho_{K^*}$ and $\|\hat{t}-t^*\|_2 \lesssim r(\rho_K)$. The result follows since $\rho_{K^*} \sim \rho^* \sim_{L,q_0} \sigma s\sqrt{\frac{1}{N}\log(\frac{ed}{s})}$ and $\|v\|_p \le \|v\|_1^{-1+2/p}\|v\|_2^{2-2/p}$ for all $v\in\mathbb{R}^d$ and $1\le p\le 2$. $\square$

Theoretical properties of MOM LASSO (cf. Theorem 5) outperform those of LASSO (cf. Theorem 3) in several ways:

- Estimation rates achieved by MOM-LASSO are the actual minimax rates $s\log(ed/s)/N$ (see [11]), while classical LASSO estimators achieve the rate $s\log(ed)/N$. This improvement is possible thanks to the adaptation step in MOM-LASSO.
- the probability deviation in (5) is polynomial – $1/N^{(q_0/2-1)}$, whereas it is exponentially small for MOM LASSO. Exponential rates for LASSO hold only if ξ is sub-Gaussian ($\|\xi\|_{L_p} \le C\sqrt{p}\|\xi\|_{L_2}$ for all $p\ge 2$).
- MOM LASSO is insensitive to data corruption by up to $s\log(ed/s)$ outliers while only one outlier can be responsible of a dramatic breakdown of the performance of LASSO (cf. Figure 1).
- Assumptions on X are weaker for MOM LASSO than for LASSO. In the LASSO case, we assume that X is sub-Gaussian whereas for the MOM LASSO we assume that the coordinates of X have $C_0\log(ed)$ sub-Gaussian moments and that X satisfies a L^2/L^1 equivalence assumption.

Let us now turn to the study of a "minmax MOM version" of the SLOPE estimator. The computation of the fixed point functions $r_Q(\cdot)$ and $r_M(\cdot)$ rely on [49], Theorem 1.6, and the computation from [39]. Again, the SLOPE norm has a dual norm which is 1-unconditional with respect to the canonical basis of $\mathbb{R}^d$, [49], Definition 1.4. Therefore, it follows from [49], Theorem 1.6, that under Assumption 6, one has

$$\mathbb{E}\sup_{v\in\rho\mathcal{B}\cap rB_2^d}\left|\sum_{i\in[N]}\epsilon_i\langle v,X_i\rangle\right| \le c_2\sqrt{N}\ell^*(\rho\mathcal{B}\cap rB_2^d),$$

$$\mathbb{E}\sup_{v\in\rho\mathcal{B}\cap rB_2^d}\left|\sum_{i\in[N]}\epsilon_i\zeta_i\langle v,X_i\rangle\right| \le c_2\sigma\sqrt{N}\ell^*(\rho\mathcal{B}\cap rB_2^d),$$

where $\mathcal{B}$ is the unit ball of the SLOPE norm and $\zeta_i = Y_i - \langle X_i, t^*\rangle$. Local Gaussian mean widths $\ell^*(\rho\mathcal{B}\cap rB_2^d)$ are bounded from above in [39], Lemma 5.3: $\ell^*(\rho\mathcal{B}\cap rB_2^d) \lesssim \min\{C\rho, \sqrt{d}r\}$ when $\beta_j = C\sqrt{\log(ed)/j}$ for all $j\in[d]$ and computations of $r_M(\cdot)$ and $r_Q(\cdot)$ follow

$$r_Q^2(\rho) \lesssim_L \begin{cases} 0 & \text{if } N \gtrsim_L d,\\ \dfrac{\rho^2}{N} & \text{otherwise,}\end{cases} \quad \text{and}$$

$$r_M^2(\rho) \lesssim_{L,q,\delta} \begin{cases} \|\xi\|_{L_q}^2\dfrac{d}{N} & \text{if } \rho^2N \gtrsim_{L,q,\delta} \|\xi\|_{L_q}^2 d^2,\\ \|\xi\|_{L_q}\dfrac{\rho}{\sqrt{N}} & \text{otherwise.}\end{cases}$$

The sparsity equation has been solved in [39], Lemma 4.3.

LEMMA 2. *Let $1\le s\le d$ and set $\mathcal{B}_s = \sum_{j\le s}\beta_j/\sqrt{j}$. If t^* is $\rho/20$ approximated (relative to the SLOPE norm) by an s-sparse vector and if $40\mathcal{B}_s \le \rho/r(\rho)$ then $\Delta(\rho)\ge 4\rho/5$.*

For $\beta_j \le C\sqrt{\log(ed/j)}$, $\mathcal{B}_s = \sum_{j\le s}\beta_j/\sqrt{j} \lesssim C\sqrt{s\log(ed/s)}$. The condition $\mathcal{B}_s \lesssim \rho/r(\rho)$ holds when $N \gtrsim_{L,q_0} s\log(ed/s)$, $\rho \gtrsim_{L,q_0} \|\xi\|_{L_q}\frac{s}{\sqrt{N}}\log(\frac{ed}{s})$. Lemma 2 implies that $\Delta(\rho) \ge 4\rho/5$ when there is an s-sparse vector in $t^* + (\rho/20)B_\Psi$. Therefore, Theorem 1 applies for $\lambda \sim r^2(\rho)/\rho \sim_{L,q,\delta} \|\xi\|_{L_q}/\sqrt{N}$.

The final ingredient is to compute the ρ_K solution to $K = cr(\rho_K)^2 N$. It is solved for $\rho_K \sim K/(\sigma\sqrt{N})$ and, therefore, $\lambda \sim r^2(\rho_K)/\rho_K \sim_{L,q,\delta} \|\xi\|_{L_q}/\sqrt{N}$.

The following result follows from Theorem 2 together with the previous computations of ρ^*, ρ_K, $r_Q(\cdot)$, $r_M(\cdot)$ and $r(\cdot)$. The proof, similar to Theorem 5, is omitted.

THEOREM 6. *Grant Assumption 5. Let $s \in [d]$. Assume that $N \ge c_1 s\log(ed/s)$ and that there is $v \in \mathbb{R}^d$ such that $|\operatorname{supp}(v)| \le s$ and $\Psi(t^* - v) \le c_2\|\xi\|_{L_q} s\log(ed/s)/\sqrt{N}$. Assume that $|\mathcal{I}| \ge N/2$ and $|\mathcal{O}| \le c_3 s\log(ed/s)$. The MOM-SLOPE estimator $\hat{t}$ with the adaptive number of blocks K from Section 3.4.1 satisfies, with probability at least $1 - c_4\exp(-c_5 s\log(ed/s))$,*

$$\|\hat{t} - t^*\|_2^2 \le c_6\|\xi\|_{L_{q_0}}^2 \frac{s}{N}\log\left(\frac{ed}{s}\right),$$

where $(c_j)_{j=1}^6$ depends only on θ_0, θ_m and q_0.

MOM-SLOPE has the same advantages upon SLOPE as MOM-LASSO upon LASSO. These improvements, listed below Theorem 5 are not repeated. The only difference is that SLOPE, unlike LASSO, already achieves the minimax rate $s\log(ed/s)/N$ whereas, without an extra adaptation step as in [11], the LASSO is not known to achieve a rate better than $s\log(ed)/N$.

5. Algorithms for minmax MOM LASSO. The aim of this section is to show that there is a systematic way to transform classical descent based algorithms (such as Newton or gradient descent algorithm, or proximal gradient descent algorithms, etc.) into robust ones using MOM approach. This section provides several examples of such modifications.

These algorithms are tested in high-dimensional frameworks. In this setup, there exists an important number of algorithms approximating LASSO. The aim of this section is to show that there is a natural modification of these algorithms that makes them more robust to outliers. The choice of hyper-parameters like the number of blocks or the regularization parameter cannot be done via classical Cross-Validation (CV) because of possible outliers in the test sets. CV procedures are also adapted using MOM's principle in Section 6. We also advocate for using random blocks at every iterations of the algorithms, to bypasses a problem of "local saddle points" we have identified. A byproduct of the latter approach is a definition of depth adapted to the learning task and, therefore, of an outliers detection algorithm. This material and a simulation study are given in Section 6 of the Supplementary Material [36].

5.1. *From algorithms for LASSO to MOM LASSO.* Each algorithm designed for the LASSO can be transformed into a robust algorithm for the minmax MOM estimator. Recall that minmax MOM LASSO estimator is

$$\hat{t}_{K,\lambda} \in \operatorname*{argmin}_{t\in\mathbb{R}^d} \sup_{t'\in\mathbb{R}^d} T_{K,\lambda}(t', t), \tag{10}$$

where $T_{K,\lambda}(t', t) = \mathrm{MOM}_K(\ell_t - \ell_{t'}) + \lambda(\|t\|_1 - \|t'\|_1)$, $\mathrm{MOM}_K(\ell_t - \ell_{t'})$ is a median of the set of real numbers $\{P_{B_1}(\ell_t - \ell_{t'}), \dots, P_{B_K}(\ell_t - \ell_{t'})\}$ and for all $k \in [K]$,

$$P_{B_k}(\ell_t - \ell_{t'}) = \frac{1}{|B_k|}\sum_{i\in B_k}\big(Y_i - \langle X_i, t\rangle\big)^2 - \big(Y_i - \langle X_i, t'\rangle\big)^2.$$

A natural idea to implement (10) is to consider algorithms based on a sequence of alternating descents (in t) and ascents (in t') steps with possible proximal/projection steps and for various choices of step sizes. A key issue here is that $t \to T_{K,\lambda}(t_0', t)$ (resp., $t' \to T_{K,\lambda}(t', t_0)$), for some given $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$, may not be convex (resp., concave). Nevertheless, one can still compute the steepest descent by assuming that the index in $[K]$ of the block achieving the median in $\mathrm{MOM}_K(\ell_{t_0} - \ell_{t_0'})$ remains constant on a convex open set containing (t_0, t_0'), for almost all (t_0, t_0'). The median is set as the minimal value of the median interval.

ASSUMPTION 7. Almost surely (with respect to $(X_i, Y_i)_{i=1}^N$) for almost all $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$ (with respect to the Lebesgue measure on $\mathbb{R}^d \times \mathbb{R}^d$), there exists a convex open set B containing (t_0, t_0') and $k \in [K]$ such that for all $(t, t') \in B$, $P_{B_k}(\ell_t - \ell_{t'}) \in \mathrm{MOM}_K(\ell_t - \ell_{t'})$.

Under Assumption 7, for almost all couples $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$, $t \to T_{K,\lambda}(t_0', t)$ is "locally convex" and $t' \to T_{K,\lambda}(t', t_0)$ is "locally concave." Therefore, for k such that $P_{B_k}(\ell_{t_0} - \ell_{t_0'}) \in \mathrm{MOM}_K(\ell_{t_0} - \ell_{t_0'})$,

$$\nabla_t \mathrm{MOM}_K(\ell_t - \ell_{t_0'})_{|t=t_0} = -2\big(X^{(k)}\big)^\top \big(Y^{(k)} - X^{(k)} t_0\big), \tag{11}$$

where $Y^{(k)} = (Y_i)_{i \in B_k}$ and $X^{(k)}$ is the $|B_k| \times d$ matrix with rows given by $X_i^\top$ for $i \in B_k$. The integer $k \in [K]$ is the index of the median of K real numbers $P_{B_1}(\ell_t - \ell_{t'}), \ldots, P_{B_K}(\ell_t - \ell_{t'})$, which is straightforward to compute. The gradient $-2(X^{(k)})^\top (Y^{(k)} - X^{(k)} t_0)$ in (11) depends on t_0' only through the index k.

REMARK 3 (Block gradient descent). Algorithms developed for the minmax estimator using steepest descent steps such as (11) are special instances of Block Gradient Descent (BGD). The major difference with standard BGD (which takes sequentially all blocks), is that the index of the block is chosen here as $P_{B_k}(\ell_{t_0} - \ell_{t_0'}) \in \mathrm{MOM}_K(\ell_{t_0} - \ell_{t_0'})$. In particular, we expect blocks corrupted by outliers to be avoided which is not the case in the classical BGD. Moreover, choosing the "descent/ascent" block k using its centrality, we also expect $P_{B_k}(\ell_{t_0} - \ell_{t_0'})$ to be close to the objective function $P(\ell_{t_0} - \ell_{t_0'})$. This should make every descent (resp., ascent) steps particularly efficient.

REMARK 4 (Map-reduce). The algorithms presented in this section particularly fits the map-reduce paradigm [19], where data are spread out in a cluster of servers and are therefore naturally split into blocks. Our procedures use for *mapper* a mean and for *reducer* a median. This makes our algorithms easily scalable into the big data framework even when some servers have crashed down (making blocks of outliers data). The median identifies the correct block of data onto which one should make a descent or an ascent and leaves aside servers which have crashed down.

REMARK 5 (Normalization). In the i.i.d. setup, the design matrix $\mathbb{X}$ (i.e., the $N \times d$ matrix with row vectors $X_1, \ldots, X_N$) is normalized to make ℓ_2^N-norms of the columns equal to one. In a corrupted setup, one row of $\mathbb{X}$ may be corrupted and normalizing each column of $\mathbb{X}$ would corrupt the entire matrix $\mathbb{X}$. We therefore do not normalize the design matrix in the following.

5.2. *Subgradient descent algorithm.* LASSO is solution of the minimization problem $\min_{t \in \mathbb{R}^d} \psi(t)$ where ψ is defined for all $t \in \mathbb{R}^d$ by $\psi(t) = \|\mathbb{Y} - \mathbb{X}t\|_2^2 + \lambda \|t\|_1$ with $\mathbb{Y} = (Y_i)_{i=1}^N$ and $\mathbb{X}$ is the $N \times d$ matrix with row vectors $X_1, \ldots, X_N$. LASSO can be approximated

input : $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$: initial point, $\epsilon > 0$: a stopping parameter, $(\eta_p)_p$, $(\beta_p)_p$: two step size sequences

output: approximated solution to the min–max problem (10)

1 **while** $\|t_{p+1} - t_p\|_2 \geq \epsilon$ **or** $\|t'_{p+1} - t'_p\|_2 \geq \epsilon$ **do**

2 find $k \in [K]$ such that $P_{B_k}(\ell_{t_p} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_p} - \ell_{t'_p})$

3 $$t_{p+1} = t_p + 2\eta_p \mathbb{X}_k^\top (\mathbb{Y}_k - \mathbb{X}_k t_p) - \lambda \eta_p \operatorname{sign}(t_p)$$

4 find $k \in [K]$ such that $P_{B_k}(\ell_{t_{p+1}} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_{p+1}} - \ell_{t'_p})$

5 $$t'_{p+1} = t'_p + 2\beta_p \mathbb{X}_k^\top (\mathbb{Y}_k - \mathbb{X}_k t'_p) - \lambda \beta_p \operatorname{sign}(t'_p)$$

6 **end**

7 **Return** (t_p, t'_p)

Algorithm 1: A "minmax MOM version" of the subgradient descent

by a subgradient descent procedure: given $t_0 \in \mathbb{R}^d$ and step sizes $(\gamma_p)_p$ (i.e., $\gamma_p > 0$ and $(\gamma_p)_p$ decreases), at step p we update

$$t_{p+1} = t_p - \gamma_p \partial\psi(t_p), \tag{12}$$

where $\partial\psi(t_p)$ is a subgradient of ψ at t_p like $\partial\psi(t_p) = -2\mathbb{X}^\top(\mathbb{Y} - \mathbb{X}t_p) + \lambda \operatorname{sign}(t_p)$ where $\operatorname{sign}(t_p)$ is the vector of signs of the coordinates of t_p with the convention $\operatorname{sign}(0) = 0$. The subgradient descent algorithm (12) can be turned into an alternating subgradient ascent/descent algorithm for the min–max estimator (10): let

$$\mathbb{Y}_k = (Y_i)_{i \in B_k} \quad \text{and} \quad \mathbb{X}_k = \big(X_i^\top\big)_{i \in B_k} \in \mathbb{R}^{|B_k| \times d}. \tag{13}$$

The key insight in Algorithm 1 are steps 2 and 4 where the blocks number have been chosen by the median operator. Those steps are expected (1) to remove outliers from the descent/ascent directions, (2) to improve the accuracy of the latter directions.

A classical choice of step size γ_p in (12) is $\gamma_p = 1/L$ where $L = \|\mathbb{X}\|_{S_\infty}^2$ ($\|\mathbb{X}\|_{S_\infty}$ is the operator norm of $\mathbb{X}$). Another possible choice follows from the Armijo–Goldstein condition with the following backtracking line search: γ is decreased geometrically while the Armijo–Goldstein condition is not satisfied

$$\textbf{while} \quad \psi\big(t_p + \gamma_\ell \partial\psi(t_p)\big) > \psi(t_p) + \delta\gamma_\ell \big\|\partial\psi(t_p)\big\|_2^2 \quad \textbf{do} \quad \gamma_{\ell+1} = \rho\gamma_\ell$$

for some given $\rho \in (0, 1)$, $\delta = 10^{-4}$ and initial point $\gamma_0 = 1$.

Of course, the same choices of step size cannot be made for $(\eta_p)_p$ and $(\beta_p)_p$ in Algorithm 1 because $\mathbb{X}$ may be corrupted but it can be adapted. In the first case, one can take $\eta_p = 1/\|\mathbb{X}_k\|_{S_\infty}^2$ where $k \in [K]$ is the index defined in line 2 of Algorithm 1 and $\beta_p = 1/\|\mathbb{X}_k\|_{S_\infty}^2$ where $k \in [K]$ is the index defined in line 4 of Algorithm 1. In the other backtracking line search case, the Armijo–Goldstein condition adapted for Algorithm 1 reads like

$$\textbf{while} \quad \psi_k\big(t_p + \gamma_\ell \partial\psi_k(t_p)\big) > \psi_k(t_p) + \delta\gamma_\ell \big\|\partial\psi_k(t_p)\big\|_2^2 \quad \textbf{do} \quad \eta_{\ell+1} = \rho\eta_\ell,$$

where $\psi_k(t) = \|\mathbb{Y}_k - \mathbb{X}_k t\|_2^2 + \lambda\|t\|_1$ where $k \in [K]$ is defined in line 2 of Algorithm 1 and, for β_p, with $k \in [K]$ defined in line 4 of Algorithm 1.

input : $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$: initial point, $\epsilon > 0$: a stopping parameter, $(\eta_k)_k$, $(\beta_k)_k$: two step size sequences
output: approximated solution to the min–max problem (10)

while $\|t_{p+1} - t_p\|_2 \geq \epsilon$ **or** $\|t'_{p+1} - t'_p\|_2 \geq \epsilon$ **do**
  find $k \in [K]$ such that $P_{B_k}(\ell_{t_p} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_p} - \ell_{t'_p})$
  $t_{p+1} = \mathrm{prox}_{\lambda\|\cdot\|_1}(t_p + 2\eta_k \mathbb{X}_k^\top(\mathbb{Y}_k - \mathbb{X}_k t_p))$
  find $k \in [K]$ such that $P_{B_k}(\ell_{t_{p+1}} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_{p+1}} - \ell_{t'_p})$
  $t'_{p+1} = \mathrm{prox}_{\lambda\|\cdot\|_1}(t'_p + 2\beta_k \mathbb{X}_k^\top(\mathbb{Y}_k - \mathbb{X}_k t'_p))$
end
Return (t_p, t'_p)

Algorithm 2: A "minmax MOM version" of ISTA

5.3. *Proximal gradient descent algorithms*. This section provides MOM versions of ISTA (Iterative Shrinkage-Thresholding Algorithm) and its accelerated version FISTA. ISTA and FISTA are proximal gradient descent where the objective function $\psi(t) = f(t) + g(t)$ with $f(t) = \|\mathbb{Y} - \mathbb{X}t\|_2^2$ (convex and differentiable) and $g(t) = \lambda\|t\|_1$ (convex). ISTA alternates between a descent in the direction of the gradient of f and a projection through the proximal operator of g, which, for the ℓ_1-norm, is the soft-thresholding:

$$t_{p+1} = \operatorname*{prox}_{\lambda\|\cdot\|_1}\big(t_p + 2\gamma_p \mathbb{X}^\top(\mathbb{Y} - \mathbb{X}t_p)\big), \tag{14}$$

where $\mathrm{prox}_{\lambda\|\cdot\|_1}(t) = (\mathrm{sign}(t_j)\max(|t_j| - \lambda, 0))_{j=1}^d$ for all $t = (t_j)_{j=1}^d \in \mathbb{R}^d$.

A natural "MOM version" for ISTA is in Algorithm 2 given by the following alternating method where the step sizes sequences $(\eta_p)_p$ and $(\beta_p)_p$ may be chosen according to the remarks below Algorithm 1 or chosen a posteriori.

5.4. *Douglas–Racheford/ADMM*. This section presents the Alternating Direction Method of Multipliers (ADMM) algorithm. It is also a splitting algorithm which reads as follows in the LASSO case: at step p,

$$\begin{aligned} t_{p+1} &= \big(\mathbb{X}^\top\mathbb{X} + \rho I_{d\times d}\big)^{-1}\big(\mathbb{X}^\top\mathbb{Y} + \rho z_p - u_p\big), \\ z_{p+1} &= \operatorname*{prox}_{\lambda\|\cdot\|_1}(t_{p+1} + u_p/\rho), \\ u_{p+1} &= u_p + \rho(t_{p+1} - z_{p+1}), \end{aligned}$$

where ρ is a tuning parameter. ADMM algorithm returns t_p after a stopping criteria is met. In Algorithm 3, we provide a MOM version of this algorithm.

6. Simulations study. This section provides an extensive simulation study based on algorithms of Section 5. In particular, their robustness and their convergence properties are illustrated on simulated data. The algorithms depend on hyperparameters that need to be tuned. Due to possible corruption, classical approaches relying on test samples cannot be trusted. The section starts therefore by introducing a robust CV procedure based on MOM principle.

6.1. *Adaptive choice of hyperparameters via MOM V-fold CV.* MOM's principles can be combined with the idea of multiple splitting into training/test datasets in cross-validation.

input : $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$: initial point, $\epsilon > 0$: a stopping parameter, ρ: a parameter
output: approximated solution to the min–max problem (10)

while $\|t_{p+1} - t_p\|_2 \geq \epsilon$ **or** $\|t'_{p+1} - t'_p\|_2 \geq \epsilon$ **do**

find $k \in [K]$ such that $P_{B_k}(\ell_{t_p} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_p} - \ell_{t'_p})$

$$t_{p+1} = \big(\mathbb{X}_k^\top \mathbb{X}_k + \rho I_{d\times d}\big)^{-1}\big(\mathbb{X}_k^\top \mathbb{Y}_k + \rho z_p - u_p\big),$$
$$z_{p+1} = \operatorname*{prox}_{\lambda\|\cdot\|_1}(t_{p+1} + u_p/\rho),$$
$$u_{p+1} = u_p + \rho(t_{p+1} - z_{p+1}),$$

find $k \in [K]$ such that $P_{B_k}(\ell_{t_{p+1}} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_{p+1}} - \ell_{t'_p})$

$$t'_{p+1} = \big(\mathbb{X}_k^\top \mathbb{X}_k + \rho I_{d\times d}\big)^{-1}\big(\mathbb{X}_k^\top \mathbb{Y}_k + \rho z'_p - u'_p\big),$$
$$z'_{p+1} = \operatorname*{prox}_{\lambda\|\cdot\|_1}\big(t'_{p+1} + u'_p/\rho\big),$$
$$u'_{p+1} = u'_p + \rho\big(t'_{p+1} - z'_{p+1}\big),$$

end
Return (t_p, t'_p)

Algorithm 3: A "minmax MOM version" of ADMM

Let $V \in [N]$ be such that N can be divided by V. Let also $\mathcal{G}_K \subset [N]$ and $\mathcal{G}_\lambda \subset (0, 1]$. The aim is to select an optimal number of blocks and regularization parameter within both grids. The dataset is split into V disjoints blocks $\mathcal{D}_1, \ldots, \mathcal{D}_V$. For each $v \in [V]$, $\bigcup_{u\neq v} \mathcal{D}_u$ is used to train a family of estimators

$$\big(\hat{f}^{(v)}_{K,\lambda} : K \in \mathcal{G}_K, \lambda \in \mathcal{G}_\lambda\big). \tag{15}$$

The remaining $\mathcal{D}_v$ of the dataset is used to test the performance of each estimator in the family (15). Using these notation, we define a MOM version of the cross-validation procedure.

DEFINITION 6. The *Median of Means V-fold CV* associated to the estimators (15) is $\hat{f}_{\hat{K},\hat{\lambda}}$ where $(\hat{K}, \hat{\lambda})$ is a minimizer of

$$(K, \lambda) \in \mathcal{G}_K \times \mathcal{G}_\lambda \to \mathrm{MomCv}_V(K, \lambda) = Q_{1/2}\big(\mathrm{MOM}^{(v)}_{K'}(\ell_{\hat{f}^{(v)}_{K,\lambda}})_{v\in[V]}\big),$$

where, for all $v \in [V]$ and $f \in F$,

$$\mathrm{MOM}^{(v)}_{K'}(\ell_f) = \mathrm{MOM}_{K'}(P_{B_1^{(v)}}\ell_f, \ldots, P_{B_{K'}^{(v)}}\ell_f) \tag{16}$$

and $B_1^{(v)} \cup \cdots \cup B_{K'}^{(v)}$ is a partition of the test set $\mathcal{D}_v$ into K' blocks where $K' \in [N/V]$ such that K' divides N/V.

The difference with standard V-fold CV is that empirical means in classical V-fold CV are replaced by MOM estimators in (16). Moreover, the mean over all V splits in the classical V-fold CV is replaced by a median.

The choice of V raises the same issues for MOM CV as for classical V-fold CV [3, 5]. In the simulations, we use $V = 5$. The construction of MOM-CV requires to choose another parameter: K', the number of blocks used to build MOM criteria (16) over the test set. One

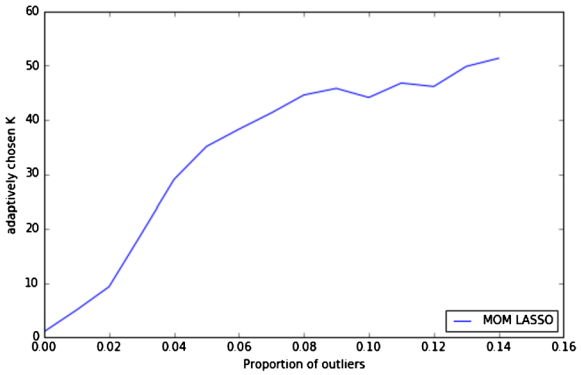

FIG. 2. *Adaptively chosen number of blocks K for the minmax MOM LASSO.*

can choose $K' = K/V$ to make only one split of $\mathcal{D}$ into K blocks and use, for each round, $(V-1)K/V$ blocks to build estimators (15) and K/V blocks to test them.

In Figure 2, hyperparameters K (i.e., the number of blocks) and λ (i.e., the regularization parameter) have been chosen for MOM LASSO estimators via MOM V-fold CV. Only the evolution of $\hat{K}$ in function of the proportion of outliers has been depicted (the choice of the adaptively chosen regularization parameter is more erratic and may first require a more deeper understanding of CV in the classical i.i.d. before the study of MOM CV in the $\mathcal{O} \cup \mathcal{I}$ framework). The adaptive $\hat{K}$ grows with the number of outliers as expected, since the number of blocks has to be at least twice the number of outliers. In particular, when there are no outliers in the dataset, MOMCV selects $K = 1$ so minmax MOM LASSO is the LASSO. The algorithm learns that splitting the database is useless in the absence of outliers: LASSO is the best choice among all minmax MOM LASSO estimators for $K \in [N/2]$.

REMARK 6. Median of Means V-fold CV introduced in Definition 6 aims at testing the performance of estimators on a possibly corrupted test set. This is done by excluding outliers from the test set thanks to the median operator. However, there are situations, for instance in image recognition, where the test set is corrupted but still we expect estimators to perform well even on these corrupted data in the test set. This is a classical robustness issue in Deep Learning [66]. Indeed, deep learning methods are known to fail if a small Gaussian noise is added to images even with a small variance undetectable by human eyes. Even though minmax MOM estimators introduced in this paper have been initially designed to be robust to outliers in the train set, one can use classical tricks to be also robust to corruption in the test set by training minmax MOM estimators onto an augmented database: in practice, given a (clean or not) dataset $(X_i, Y_i)_{i=1}^N$, one can construct an augmented dataset where each data (X_i, Y_i) is replicated m times with an added Gaussian noise: $(X_i + Z_{i1}, Y_i), \ldots, (X_i + Z_{im}, Y_i)$—where $(Z_{ij} : 1 \le i \le N, 1 \le j \le m)$ are i.i.d. Gaussian variable, and then a minmax MOM estimator can be trained onto the dataset

$$(X_i + Z_{ij}, Y_i), \quad i = 1, \ldots, n, \, j = 1, \ldots, m.$$

By doing so, we expect the minmax MOM estimator to improve its robustness performance evaluated on a corrupted test set.

6.2. *Saddle-point, random blocks, outliers detection and depth.* The aim of this section is to show some advantages of choosing randomly the blocks at every (descent and ascent) steps of the algorithm and how this modified version works on the example of ADMM. As a byproduct, it is possible to define an outliers detection algorithm.

Let us first explain a problem of "local saddle point" in the case of fixed blocks. Minmax MOM estimators are based on the observation that the oracle f^* is solution to the minmax problem $f^* \in \operatorname{argmin}_{f\in F} \sup_{g\in F} P(\ell_f - \ell_g)$. Likewise, f^* is solution of the maxmin problem: $f^* \in \operatorname{argmax}_{g\in F} \inf_{f\in F} P(\ell_f - \ell_g)$. One can also define the maxmin MOM estimator

$$\hat{g}_{K,\lambda} \in \operatorname*{argmax}_{g\in F} \inf_{f\in F} T_{K,\lambda}(g, f). \tag{17}$$

Following the proofs of Section 6 in the Supplementary Material [36], one can prove the same results for $\hat{g}_{K,\lambda}$ and $\hat{f}_{K,\lambda}$ (see Section 7 in the Supplementary Material for a proof in small dimension). However, $\hat{g}_{K,\lambda}$ and $\hat{f}_{K,\lambda}$ may differ since, in general

$$\operatorname*{argmin}_{f\in F} \sup_{g\in F} T_{K,\lambda}(g, f) \neq \operatorname*{argmax}_{g\in F} \inf_{f\in F} T_{K,\lambda}(g, f). \tag{18}$$

In other words, the duality gap may not be null. Since $T_{K,\lambda}(g, f) = -T_{K,\lambda}(f, g)$ for all $f, g \in F$, (18) holds if and only if

$$\inf_{f\in F} \sup_{g\in F} T_{K,\lambda}(f, g) = 0.$$

In that case, $\hat{f}$ is a *saddle-point* estimator and minmax and maxmin estimators are equal. The left-hand side of Figure 3 shows a simulation where this happens. The choice of fixed blocks $B_1, \ldots, B_K$ may result in a problem of "*local saddle points*" and the algorithms remain close to suboptimal local saddle points. To see this, consider the vector case (i.e., for $F = \{f(\cdot) = \langle \cdot, t\rangle : t \in \mathbb{R}^d\}$ and introduce, for all $k \in [K]$,

$$\mathcal{C}_k = \big\{(t, t') \in \mathbb{R}^d \times \mathbb{R}^d : \operatorname{MOM}_K(\ell_t - \ell_{t'}) = P_{B_k}(\ell_t - \ell_{t'})\big\}. \tag{19}$$

The problem is that, if a cell $\mathcal{C}_k$ contains a saddle-point of $(t, t') \to P_{B_k}(\ell_t - \ell_{t'}) + \lambda(\|t\|_1 - \|t'\|_1)$ the algorithms gets stuck in that cell instead of looking for "better saddle-point" in other cells.

To overcome this issue, the partition is chosen at random at every descent and ascent steps of the algorithms, so the decomposition into cells $\mathcal{C}_1, \ldots, \mathcal{C}_K$ changes at every steps. As an example, we develop the ADMM procedure with a random choice of blocks in Algorithm 4.

In Figure 3, both MOM LASSO via ADMM with fixed and changing blocks are run. Both the objective function and the estimation error of MOM LASSO jump with fixed blocks. These jumps correspond to a change of cell number. The algorithm converges to local saddle-points before jumping to other cells, thanks to the regularization of the ℓ_1-norm. On the other hand, the algorithms with changing blocks do not suffer this drawback. Figure 3 shows that the estimation error converges faster and more smoothly for changing blocks. The objective

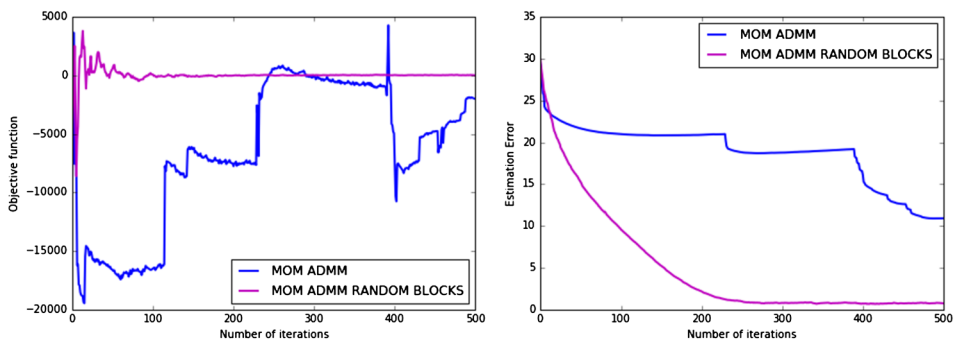

FIG. 3. *Fixed blocks against random blocks.*

input : $(t_0, t_0') \in \mathbb{R}^d \times \mathbb{R}^d$: initial point, $\epsilon > 0$: a stopping parameter, ρ: parameter
output: approximated solution to the min–max problem (10)

while $\|t_{p+1} - t_p\|_2 \geq \epsilon$ **or** $\|t'_{p+1} - t'_p\|_2 \geq \epsilon$ **do**
  Build an equipartition $B_1, \ldots, B_K$ of $[N]$ at random.
  Find $k \in [K]$ such that $P_{B_k}(\ell_{t_p} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_p} - \ell_{t'_p})$
  $t_{p+1} = (\mathbb{X}_k^\top \mathbb{X}_k + \rho I_{d\times d})^{-1}(\mathbb{X}_k^\top \mathbb{Y}_k + \rho z_p - u_p)$
  $z_{p+1} = \mathrm{prox}_{\lambda\|\cdot\|_1}(t_{p+1} + u_p/\rho)$
  $u_{p+1} = u_p + \rho(t_{p+1} - z_{p+1})$
  Build an equipartition $B_1, \ldots, B_K$ of $[N]$ at random.
  Find $k \in [K]$ such that $P_{B_k}(\ell_{t_{p+1}} - \ell_{t'_p}) = \mathrm{MOM}_K(\ell_{t_{p+1}} - \ell_{t'_p})$
  $t'_{p+1} = (\mathbb{X}_k^\top \mathbb{X}_k + \rho I_{d\times d})^{-1}(\mathbb{X}_k^\top \mathbb{Y}_k + \rho z'_p - u'_p)$
  $z'_{p+1} = \mathrm{prox}_{\lambda\|\cdot\|_1}(t'_{p+1} + u'_p/\rho)$
  $u'_{p+1} = u'_p + \rho(t'_{p+1} - z'_{p+1})$
end
Return (t_p, t'_p)

Algorithm 4: minmax MOM ADMM with changing random blocks

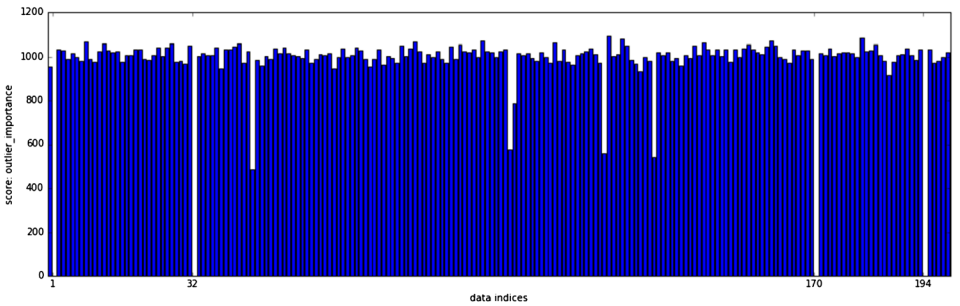

FIG. 4. *Outliers detection algorithm. The dataset has been corrupted by* 4 *outliers at number* 1, 32, 170 *and* 194. *The score of the outliers is* 0: *they have not been selected even once.*

function of MOM ADMM with changing blocks converges to zero so the duality gap converges to zero. This gives a natural stopping criterion and shows that minmax and maxmin MOM LASSO are solution of a saddle point problem even though the objective function is not convex-concave.

A byproduct is an *outliers detection procedure*. Count the number of times each data is selected in the selected median blocks of steps 3 and 7 of Algorithm 4. At the end of the algorithm (for instance, Algorithm 4), every data ends up with a score revealing its centrality for the learning task. Aggressive outliers are likely to corrupt their respective blocks and should therefore not be selected at steps 3 and 7 of Algorithm 4. With fixed blocks, informative data cannot be distinguished from outliers lying in the same block, therefore, this outliers detection algorithm only makes sense when blocks are changing at every steps. Figure 4 shows performance of this strategy on synthetic data (cf. Section 6.3 for more details on the simulations). Outliers (data 1, 32, 170 and 194) end up with a null score.

6.3. *Simulations setup for the figures.* All codes are available at [55] and can be used to reproduce the figures. Many other simulations and algorithms can be found in [55].

6.3.1. *Data generating process and corruption by outliers.* The algorithms introduced in Section 5 are tested on datasets corrupted by outliers of various forms in [55]. The basic set of informative data is called $\mathcal{D}_1$. The outliers are named $\mathcal{D}_2, \mathcal{D}_3, \mathcal{D}_4$ and $\mathcal{D}_5$. These data are merged and shuffled in the dataset $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3 \cup \mathcal{D}_4 \cup \mathcal{D}_5$ given to the algorithm.

1. The set $\mathcal{D}_1$ of inliers contains N_{good} i.i.d. data (X_i, Y_i) with common distribution

$$Y = \langle X, t^* \rangle + \xi, \tag{20}$$

where $t^* \in \mathbb{R}^d$, $X \sim \mathcal{N}(0, I_{d\times d})$ and $\xi \sim \mathcal{N}(0, \sigma^2)$ is independent of X.

2. $\mathcal{D}_2$ is a dataset of N_{bad-1} outliers (X_i, Y_i) such that $Y_i = 1$ and $X_i = (1)_{j=1}^d$
3. $\mathcal{D}_3$ is a dataset of N_{bad-2} outliers (X_i, Y_i) such that $Y_i = 10{,}000$ and $X_i = (1)_{j=1}^d$
4. $\mathcal{D}_4$ is a dataset of N_{bad-3} outliers (X_i, Y_i) where Y_i is a $0-1$-Bernoulli random variable and X_i is uniformly distributed over $[0, 1]^d$
5. $\mathcal{D}_5$ is also a set of outliers that have been generated according to a linear model (20), with the same target vector t^* and a different choice of design X and noise ξ. The design $X \sim \mathcal{N}(0, \Sigma)$ with $\Sigma = (\rho^{|i-j|})_{1\le i,j\le d}$ and ξ is a heavy-tailed noise distributed according to a Student distribution with various degrees of freedom.

The different types of outliers $\mathcal{D}_j$, $j = 2, 3, 4, 5$ are useless to learn the oracle t^* some are not independent nor random as in $\mathcal{D}_2$ and $\mathcal{D}_3$.

6.3.2. *Simulations setup for the figures.* Let us now precise the parameters of the simulations in Figure 1 and Figure 2: the number of observations is $N = 200$, the number of features is $d = 500$, $t^* \in \mathbb{R}^d$ has sparsity $s = 10$ and support chosen at random, with nonzero coordinates t_j^* being either equal to 10, -10 or decreasing according to $\exp(-j/10)$. Informative data $\mathcal{D}_1$, described in Section 6.3.1, have variance $\sigma = 1$. This dataset is increasingly corrupted with outliers in $\mathcal{D}_3$.

The proportion of outliers are $0, 1/100, 2/100, \ldots, 15/100$. The ADMM algorithm is run with adaptive λ chosen by V-fold CV with $V = 5$ for the LASSO. Then MOM ADMM is run with adaptive K and λ chosen by MOM CV with $V = 5$ and $K' = \max(\text{grid}_K)/V$ where $\text{grid}_K = \{1, 4, \ldots, 115/4\}$ and $\text{grid}_\lambda = \{0, 10, 20, \ldots, 100\}/\sqrt{100}$ are the search grids used to select the best K and λ during the CV and MOM CV steps. The number of iterations of ADMM and MOM ADMM is 200. Simulations have been run 70 times and the averaged values of the estimation error and adaptive $\hat{K}$ have been reported in Figure 1, Figure 5 and Figure 2. The ℓ_2 estimation error of LASSO increases roughly from 0 when there is no outliers and stabilize at 550 right after a single outliers enters the dataset. The value 550 comes from the fact that $Y = 10{,}000$ and $X = (1)_{j=1}^{500}$ satisfy that the vector t with minimal ℓ_1^d norm among all the solutions t of $Y = \langle X, t\rangle$ is $t^{**} = (20)_{j=1}^{500}$, and $\|t^{**} - t^*\|_2$ is approximately 550. This means that LASSO is trying to fit a model on the single outlier instead of solving the linear problem associated with the 200 other informative data. A single outliers is therefore completely misleading the LASSO.

For Figure 3, we have run similar experiments with $N = 200$, $d = 300$, $s = 20$, $\sigma = 1$, $K = 10$, the number of iterations was 500 and the regularization parameter was $1/\sqrt{N}$.

For Figure 4, we took $N = 200$, $d = 500$, $s = 20$, $\sigma = 1$, the number of outliers is $|\mathcal{O}| = 4$ and the outliers are of the form $Y = 10{,}000$ and $X = (1)_{j=1}^d$, $K = 10$, the number of iterations is 5.000 and $\lambda = 1/\sqrt{200}$.

Acknowledgments. The first author was supported by Labex ECODEC (ANR—11-LABEX-0047), Maison des Sciences de l'Homme (MSH—17MA01) and a donation of NVIDIA Corporation.

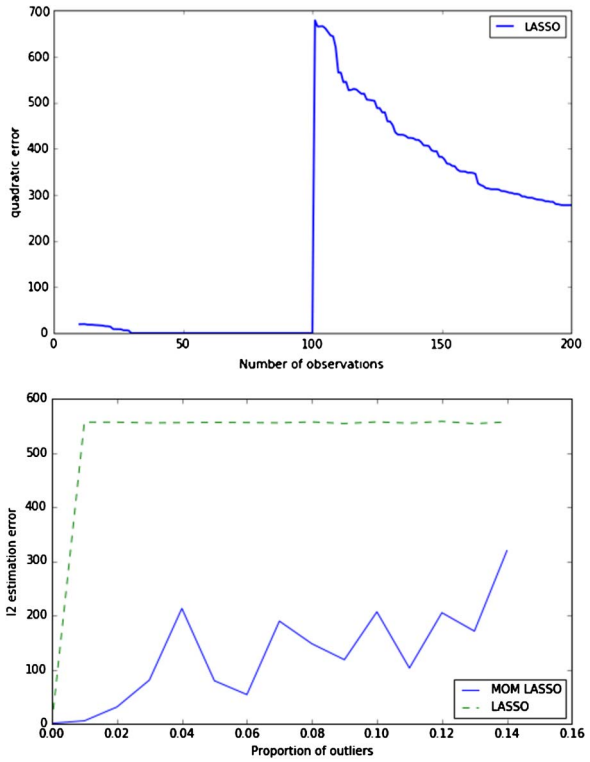

FIG. 5. *Estimation error versus proportion of outliers for LASSO and the minmax MOM LASSO.*

SUPPLEMENTARY MATERIAL

Supplementary material to "Estimation bounds and sharp oracle inequalities of regularized procedures with Lipschitz loss functions" (DOI: 10.1214/19-AOS1828SUPP; .pdf). Section 6 gives the proof of the main results. These main results focus on the regularized version of the MOM estimates of the increments presented in this Introduction that are well suited for high-dimensional learning frameworks. We complete these results in Section 7, providing results for the basic estimators without regularization in small dimension. Finally, Section 8 provides minimax optimality results for our procedures.

REFERENCES

[1] ALON, N., MATIAS, Y. and SZEGEDY, M. (1999). The space complexity of approximating the frequency moments. *J. Comput. System Sci.* **58** 137–147. MR1688610 https://doi.org/10.1006/jcss.1997.1545

[2] ALQUIER, P., COTTET, V. and LECUÉ, G. (2019). Estimation bounds and sharp oracle inequalities of regularized procedures with Lipschitz loss functions. *Ann. Statist.* **47** 2117–2144. MR3953446 https://doi.org/10.1214/18-AOS1742

[3] ARLOT, S. and CELISSE, A. (2010). A survey of cross-validation procedures for model selection. *Stat. Surv.* **4** 40–79. MR2602303 https://doi.org/10.1214/09-SS054

[4] ARLOT, S. and CELISSE, A. (2011). Segmentation of the mean of heteroscedastic data via cross-validation. *Stat. Comput.* **21** 613–632. MR2826696 https://doi.org/10.1007/s11222-010-9196-x

[5] ARLOT, S. and LERASLE, M. (2016). Choice of V for V-fold cross-validation in least-squares density estimation. *J. Mach. Learn. Res.* **17** Paper No. 208, 50. MR3595142

[6] AUDIBERT, J.-Y. and CATONI, O. (2011). Robust linear least squares regression. *Ann. Statist.* **39** 2766–2794. MR2906886 https://doi.org/10.1214/11-AOS918

[7] BACH, F. R. (2010). Structured sparsity-inducing norms through submodular functions. In *Advances in Neural Information Processing Systems* 23: 24*th Annual Conference on Neural Information Processing Systems* 2010. *Proceedings of a Meeting Held* 6–9 *December* 2010 118–126, Vancouver, BC.

[8] BARAUD, Y. (2011). Estimator selection with respect to Hellinger-type risks. *Probab. Theory Related Fields* **151** 353–401. MR2834722 https://doi.org/10.1007/s00440-010-0302-y

[9] BARAUD, Y. and BIRGÉ, L. (2016). Rho-estimators for shape restricted density estimation. *Stochastic Process. Appl.* **126** 3888–3912. MR3565484 https://doi.org/10.1016/j.spa.2016.04.013

[10] BARAUD, Y., BIRGÉ, L. and SART, M. (2017). A new method for estimation and model selection: ρ-estimation. *Invent. Math.* **207** 425–517. MR3595933 https://doi.org/10.1007/s00222-016-0673-5

[11] BELLEC, P. C., LECUÉ, G. and TSYBAKOV, A. B. (2016). Slope meets lasso: Improved oracle bounds and optimality. Technical report, CREST, CNRS, Université Paris Saclay.

[12] BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. MR2533469 https://doi.org/10.1214/08-AOS620

[13] BIRGÉ, L. (1984). Stabilité et instabilité du risque minimax pour des variables indépendantes équidistribuées. *Ann. Inst. Henri Poincaré Probab. Stat.* **20** 201–223. MR0762855

[14] BIRGÉ, L. (2006). Model selection via testing: An alternative to (penalized) maximum likelihood estimators. *Ann. Inst. Henri Poincaré Probab. Stat.* **42** 273–325. MR2219712 https://doi.org/10.1016/j.anihpb.2005.04.004

[15] BLANCHARD, G., BOUSQUET, O. and MASSART, P. (2008). Statistical performance of support vector machines. *Ann. Statist.* **36** 489–531. MR2396805 https://doi.org/10.1214/009053607000000839

[16] BOGDAN, M., VAN DEN BERG, E., SABATTI, C., SU, W. and CANDÈS, E. J. (2015). SLOPE—adaptive variable selection via convex optimization. *Ann. Appl. Stat.* **9** 1103–1140. MR3418717 https://doi.org/10.1214/15-AOAS842

[17] BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data*: *Methods*, *Theory and Applications*. *Springer Series in Statistics*. Springer, Heidelberg. MR2807761 https://doi.org/10.1007/978-3-642-20192-9

[18] CATONI, O. (2012). Challenging the empirical mean and empirical variance: A deviation study. *Ann. Inst. Henri Poincaré Probab. Stat.* **48** 1148–1185. MR3052407 https://doi.org/10.1214/11-AIHP454

[19] DEAN, J. and GHEMAWAT, S. (2010). Mapreduce: A flexible data processing tool. *Commun. ACM* **53** 72–77.

[20] DEVROYE, L., LERASLE, M., LUGOSI, G. and OLIVEIRA, R. I. (2016). Sub-Gaussian mean estimators. *Ann. Statist.* **44** 2695–2725. MR3576558 https://doi.org/10.1214/16-AOS1440

[21] ELSENER, A. and VAN DE GEER, S. (2018). Robust low-rank matrix estimation. *Ann. Statist.* **46** 3481–3509. MR3852659 https://doi.org/10.1214/17-AOS1666

[22] FAN, J., LI, Q. and WANG, Y. (2017). Estimation of high dimensional mean regression in the absence of symmetry and light tail assumptions. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 247–265. MR3597972 https://doi.org/10.1111/rssb.12166

[23] GIRAUD, C. (2015). *Introduction to High-Dimensional Statistics*. *Monographs on Statistics and Applied Probability* **139**. CRC Press, Boca Raton, FL. MR3307991

[24] HAMPEL, F. R. (1971). A general qualitative definition of robustness. *Ann. Math. Stat.* **42** 1887–1896. MR0301858 https://doi.org/10.1214/aoms/1177693054

[25] HAMPEL, F. R. (1974). The influence curve and its role in robust estimation. *J. Amer. Statist. Assoc.* **69** 383–393. MR0362657

[26] HAN, Q. and WELLNER, J. A. (2017). A sharp multiplier inequality with applications to heavy-tailed regression problems. Available at arXiv:1706.02410.

[27] HUBER, P. J. (1964). Robust estimation of a location parameter. *Ann. Math. Stat.* **35** 73–101. MR0161415 https://doi.org/10.1214/aoms/1177703732

[28] HUBER, P. J. (1967). The behavior of maximum likelihood estimates under nonstandard conditions. In *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability* (*Berkeley*, *Calif.*, 1965/66), *Vol. I*: *Statistics* 221–233. Univ. California Press, Berkeley, CA. MR0216620

[29] HUBER, P. J. and RONCHETTI, E. M. (2009). *Robust Statistics*, 2nd ed. *Wiley Series in Probability and Statistics*. Wiley, Hoboken, NJ. MR2488795 https://doi.org/10.1002/9780470434697

[30] JERRUM, M. R., VALIANT, L. G. and VAZIRANI, V. V. (1986). Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.* **43** 169–188. MR0855970 https://doi.org/10.1016/0304-3975(86)90174-X

[31] KOLTCHINSKII, V. (2011). *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. *Lecture Notes in Math.* **2033**. Springer, Heidelberg. MR2829871 https://doi.org/10.1007/978-3-642-22147-7

[32] KOLTCHINSKII, V. and MENDELSON, S. (2015). Bounding the smallest singular value of a random matrix without concentration. *Int. Math. Res. Not. IMRN* **23** 12991–13008. MR3431642 https://doi.org/10.1093/imrn/rnv096

[33] LE CAM, L. (1986). *Asymptotic Methods in Statistical Decision Theory*. *Springer Series in Statistics*. Springer, New York. MR0856411 https://doi.org/10.1007/978-1-4612-4946-7

[34] LeCam, L. (1973). Convergence of estimates under dimensionality restrictions. *Ann. Statist.* **1** 38–53. MR0334381
[35] Lecué, G. and Lerasle, M. (2017). Learning from mom's principle: Le Cam's approach. Technical report, CNRS, ENSAE, Paris-sud, 2017. To appear in Stochastic Processes and their applications.
[36] Lecué, G. and Lerasle, M. (2020). Supplement to "Robust machine learning by median-of-means: Theory and practice." https://doi.org/10.1214/19-AOS1828SUPP.
[37] Lecué, G. and Mendelson, S. (2013). Learning subgaussian classes: Upper and minimax bounds Technical report, CNRS, Ecole polytechnique and Technion.
[38] Lecué, G. and Mendelson, S. (2017). Regularization and the small-ball method II: Complexity dependent error rates. *J. Mach. Learn. Res.* **18** Paper No. 146, 48. MR3763780
[39] Lecué, G. and Mendelson, S. (2018). Regularization and the small-ball method I: Sparse recovery. *Ann. Statist.* **46** 611–641. MR3782379 https://doi.org/10.1214/17-AOS1562
[40] Lerasle, M. and Oliveira, R. I. (2011). Robust empirical mean estimators. Available at arXiv:1112.3914.
[41] Lounici, K. (2008). Sup-norm convergence rate and sign concentration property of Lasso and Dantzig estimators. *Electron. J. Stat.* **2** 90–102. MR2386087 https://doi.org/10.1214/08-EJS177
[42] Lugosi, G. and Mendelson, S. (2019). Regularization, sparse recovery, and median-of-means tournaments. *Bernoulli* **25** 2075–2106. MR3961241 https://doi.org/10.3150/18-BEJ1046
[43] Lugosi, G. and Mendelson, S. (2019). Risk minimization by median-of-means tournaments. *J. Eur. Math. Soc.* (*JEMS*). To appear.
[44] Massart, P. and Nédélec, É. (2006). Risk bounds for statistical learning. *Ann. Statist.* **34** 2326–2366. MR2291502 https://doi.org/10.1214/009053606000000786
[45] Massias, M., Fercoq, O., Gramfort, A. and Salmon, J. (2017). Generalized concomitant multi-task lasso for sparse multimodal regression.
[46] Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.* **34** 1436–1462. MR2278363 https://doi.org/10.1214/009053606000000281
[47] Meinshausen, N. and Yu, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.* **37** 246–270. MR2488351 https://doi.org/10.1214/07-AOS582
[48] Mendelson, S. (2014). Learning without concentration. In *Proceedings of the 27th annual conference on Learning Theory COLT14*, 25–39.
[49] Mendelson, S. (2016). On multiplier processes under weak moment assumptions Technical report, Technion.
[50] Minsker, S. (2015). Geometric median and robust estimation in Banach spaces. *Bernoulli* **21** 2308–2335. MR3378468 https://doi.org/10.3150/14-BEJ645
[51] Minsker, S. and Strawn, N. (2019). Distributed statistical estimation and rates of convergence in normal approximation. *Electron. J. Statist.* **13** 5213–5252. MR4043072
[52] Negahban, S. N., Ravikumar, P., Wainwright, M. J. and Yu, B. (2012). A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statist. Sci.* **27** 538–557. MR3025133 https://doi.org/10.1214/12-STS400
[53] Nemirovsky, A. S. and Yudin, D. B. (1983). *Problem Complexity and Method Efficiency in Optimization. A Wiley-Interscience Publication*. Wiley, New York. MR0702836
[54] Nickl, R. and van de Geer, S. (2013). Confidence sets in sparse regression. *Ann. Statist.* **41** 2852–2876. MR3161450 https://doi.org/10.1214/13-AOS1170
[55] Notebook. Available at https://github.com/lecueguillaume/MOMpower.
[56] Saumard, A. (2018). On optimality of empirical risk minimization in linear aggregation. *Bernoulli* **24** 2176–2203. MR3757527 https://doi.org/10.3150/17-BEJ925
[57] Su, W. and Candès, E. (2016). SLOPE is adaptive to unknown sparsity and asymptotically minimax. *Ann. Statist.* **44** 1038–1068. MR3485953 https://doi.org/10.1214/15-AOS1397
[58] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. MR1379242
[59] Tukey, J. W. (1960). A survey of sampling from contaminated distributions. In *Contributions to Probability and Statistics* 448–485. Stanford Univ. Press, Stanford, CA. MR0120720
[60] Tukey, J. W. (1962). The future of data analysis. *Ann. Math. Stat.* **33** 1–67. MR0133937 https://doi.org/10.1214/aoms/1177704711
[61] van de Geer, S. (2014). Weakly decomposable regularization penalties and structured sparsity. *Scand. J. Stat.* **41** 72–86. MR3181133 https://doi.org/10.1111/sjos.12032
[62] van de Geer, S., Bühlmann, P., Ritov, Y. and Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Statist.* **42** 1166–1202. MR3224285 https://doi.org/10.1214/14-AOS1221

[63] VAN DE GEER, S. A. (2007). The deterministic lasso. Technical report, ETH Zürich. Available at http://www.stat.math.ethz.ch/~geer/lasso.pdf.

[64] VAN DE GEER, S. A. (2008). High-dimensional generalized linear models and the lasso. *Ann. Statist.* **36** 614–645. MR2396809 https://doi.org/10.1214/009053607000000929

[65] VAPNIK, V. N. (1998). *Statistical Learning Theory*. *Adaptive and Learning Systems for Signal Processing*, *Communications*, *and Control*. Wiley, New York. MR1641250

[66] WANG, T. E., GU, Y., MEHTA, D., ZHAO, X. and BERNAL, E. A. (2018). Towards robust deep neural networks. Preprint. Available at arXiv:1810.11726.

[67] ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 217–242. MR3153940 https://doi.org/10.1111/rssb.12026

[68] ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *J. Mach. Learn. Res.* **7** 2541–2563. MR2274449

[69] ZHOU, W.-X., SUN, Q. and FAN, J. (2017). Adaptive huber regression: Optimality and phase transition. Preprint. Available at arXiv:1706.06991.