

Robust classification via MOM minimization

G. Lecué and M. Lerasle and T. Mathieu

July 3, 2018

Abstract

We present an extension of Vapnik’s classical empirical risk minimizer (ERM) where the empirical risk is replaced by a median-of-means (MOM) estimator, the new estimators are called MOM minimizers. While ERM is sensitive to corruption of the dataset for many classical loss functions used in classification, we show that MOM minimizers behave well in theory, in the sense that it achieves Vapnik’s (slow) rates of convergence under weak assumptions: data are only required to have a finite second moment and some outliers may also have corrupted the dataset.

We propose an algorithm inspired by MOM minimizers. These algorithms can be analyzed using arguments quite similar to those used for Stochastic Block Gradient descent. As a proof of concept, we show how to modify a proof of consistency for a descent algorithm to prove consistency of its MOM version. As MOM algorithms perform a smart subsampling, our procedure can also help to reduce substantially time computations and memory resources when applied to non linear algorithms. These empirical performances are illustrated on both simulated and real datasets.

1 Introduction

Machine learning has experienced unprecedented growth in recent years with a major societal impact exceeding that it has had in science for several decades. The article is set within the framework of robust learning theory, in the sense that only moment hypotheses are guaranteed on the data and the dataset may contain outliers. Robust learning theory has received particular attention in recent years, the aim sometimes being to construct reliable procedures under weaker assumptions than strong stochastic ones (such as the Gaussian character of noise in regression) and sometimes to detect outliers. This interest can be appreciated, for example, by the challenges recently posted on “kaggle”, the most popular data science competition platform. The 1.5 million dollars problem “Passenger Screening Algorithm Challenge” involves the discovery of terrorist activity from 3D images. The “NIPS 2017: Defense Against Adversarial Attack” consists in building algorithms robust to adversarial data.

Our approach extends Vapnik’s *empirical risk minimization* (ERM) (cf., [66]). The ERM is defined as a minimizer of the empirical risk $f \rightarrow (1/N) \sum_{i=1}^N I_{\{Y_i \neq f(X_i)\}}$ over a class \mathcal{F} of functions from \mathcal{X} to $\{\pm 1\}$:

$$\hat{f}_{\text{ERM}} \in \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N I\{Y_i \neq f(X_i)\} \quad (1)$$

where $I\{Y_i \neq f(X_i)\} = 1$ if $Y_i \neq f(X_i)$ and 0 otherwise. This estimator is fairly insensitive to both heavy-tailed data and to the presence of a small proportion of outliers in the dataset (see Theorem 1 below). The reason is that the 0 – 1 loss $\ell_f(x, y) = I_{\{y \neq f(x)\}}$ is bounded, which grants concentration no matter the distribution of X and a small number of data cannot really impact the empirical mean performance. However, in practice, we find that algorithms supposed to approach ERM (usually following a convex relaxation principle) have poor performance in the presence of a few outliers. Figure 1 illustrates this problem on a toy example on which most of the data would be well separated by a linear classifier like Perceptron [61] or logistic classifier, but adding some anomalies completely flaws these algorithms.




Figure 1: Scatter plot of the toy dataset, the color of the points gives their class. The background color gives the linear separation provided by the perceptron (left) and the logistic regression (right) trained on this corrupted dataset.

The example in Figure 1 is representative of a general problem. Algorithms approaching ERM do not directly solve the minimization problem (1), which is NP-hard even for simple function classes like half-spaces indicators [26, 23]. Instead, the algorithm seeks to approach a convex relaxation of the problem, in which the considered loss (e.g. hinge or logistic loss) is not necessarily bounded. The advantage of this relaxation is that the function to be minimized is convex and that the minimum can easily be approached using for example gradient descent algorithms. On the other hand, this function becomes unbounded and can be diverted by a single outlier. Minimizing this corrupted criterion then leads to a bad classifier choice.

In learning theory, most alternatives to ERM manage the problem of outliers and heavy tail distributions for outputs only. These solutions are based on the pioneering work of John Tukey [64, 65], Peter Huber [32, 33] and Frank Hampel [27, 28], replacing the square loss by a robust alternative like Huber loss or Tuckey’s biweight loss. These methods do not allow to treat the case where the inputs are with heavy tails or corrupted, which is a classical problem in robust statistics also known as the “*leverage point problem*”, see [34]. In this article, we address this question by considering an alternative to M-estimators, called median-of-means. Several estimators based on this principle have recently been proposed in the literature [56, 42, 47, 48, 49, 53, 43]. To our knowledge, these articles use the small ball hypothesis [41, 54] to treat problems of least square regression or Lipschitzian loss regression. This assumption is restrictive in some classic functional regression frameworks [63, 29] or for problems such as the construction of recommendation system where inputs are sampled in the canonical basis and therefore do not satisfy a small ball condition. In this article, we propose a natural estimator based on the MOM principle, which we call MOM minimizer, and which we study without the small ball hypothesis. Instead we assume an a priori bound on the L^2 -norm of learning functions. We can identify mainly two streams of hypothesis in Learning theory: 1) boundedness with respect to some norm of the class F of functions and the output Y , the typical example is the boundedness in L_∞ assumption or 2) norm equivalence assumption over the class F (or, more precisely, on the shifted class $F - f^* = \{f - f^* : f \in F\}$ where f^* is the oracle in F) and Y , the typical example being the subgaussian assumption, i.e. $\|f - f^*\|_{\psi_2} \leq L \|f - f^*\|_{L_2}, \forall f \in F$. The small ball assumption is a norm equivalence assumption between the L_1 and L_2 norms and is concerned with the second type of assumptions. Our approach deals with the first type of assumption since we only assume boundedness in L_2 which can be seen as an significant improvement upon a L_∞ boundedness assumption even though the rates of convergence obtained for our estimators turn out to be minimax [18] in the absence of a margin condition [51] even under a L_∞ assumption.

The estimation of the expectation of a univariate variable by median-of-means (MOM) [2, 37, 59] is done as follows: given a partition of the dataset into blocks of the same size, the empirical mean is constructed on each block and the MOM estimator is obtained by taking the median of these empirical means (see Section 2.2 for details). These estimators are naturally resistant to the presence of a few outliers in the dataset: if the number of these outliers does not exceed half the number of blocks, more than half of these blocks are made of “clean” data and the median is a reliable estimator.

In this article, we also study a family of algorithms approaching the MOM minimizer. The general idea we wish to promote is to use the MOM principle within algorithms originally intended for the evaluation of ERM associated to convex loss functions. In Section 4, we present the modification of gradient descent algorithms following this philosophy. The general principle of this iterative algorithm is as follows: at iteration t , a dataset equipartition B_1, \dots, B_K is selected uniformly at random and the most central block B_{med} is determined according to the following formula

$$\sum_{i \in B_{\text{med}}} \ell_{f_t}(X_i, Y_i) = \text{median} \left(\sum_{i \in B_k} \ell_{f_t}(X_i, Y_i) : k = 1, \dots, K \right). \quad (2)$$

The next iteration f_{t+1} is then produced by taking from f_t a step down in the direction opposite to the gradient of $f \rightarrow \sum_{i \in B_{\text{med}}} \ell_f(X_i, Y_i)$ at f_t , cf. Algorithm 1. The underlying heuristic is that the data in the selected block B_{med} are safe for estimating the mean at the point, in the sense that empirical risk $|B_{\text{med}}|^{-1} \sum_{i \in B_{\text{med}}} \ell_{f_t}(X_i, Y_i)$ is a subgaussian estimator of $\mathbb{E} \ell_f(X_i, Y_i)$, cf. [19] and that data indexed by B_{med} should not be outliers. The consistency of the algorithm is studied in Section 4.2. Theorem 3 provides an example showing how to modify the convergence proof of SGD algorithms.

In Section 5, the practical performances of this algorithm are illustrated on several simulations, involving in particular different loss functions. These simulations illustrate not surprisingly the gain of robustness that there is to use these algorithms in their MOM version rather than in their traditional version, as can for example be appreciated on the toy-example of Figure 1, cf. Figure 5. MOM estimators are compared to different learning algorithms on real datasets that can be modeled by heavy tailed data, obtaining in each case performances comparable to the best of these benchmarks.

Another advantage of our procedure is that it works on blocks of data. This can improve speed of execution and reduce memory requirements, which can be decisive on massive datasets and/or when one wishes to use non-linear algorithms as in Section 4.3. This principle of dividing the dataset to calculate estimators more quickly and then aggregating them is a powerful tool in statistics and machine learning [38]. Among others, one can mention bagging methods [12] or subbagging —a variant of bagging where the bootstrap is replaced by subsampling— [15]. These methods are considered difficult to study theoretically in general and their analysis is often limited to obtaining asymptotic guarantees. By contrast, the theoretical tools for non-asymptotic risk analysis of MOM minimizer have already essentially been developed and the consistency of the algorithm can also be derived from standard arguments. Finally, subsampling by the central block B_{med} ensures robustness properties that cannot be guaranteed by traditional alternatives.

The algorithm provides an empirical notion of data depth: data providing good risk estimates of $f \rightarrow \mathbb{E} \ell_f(X, Y)$ are likely to be selected many times in the central block B_{med} along the descent, while outliers will be systematically ignored. This notion of depth, based on the risk function, is very natural for prediction problems. It is complemented by an outliers detection procedure: data that are selected a number of times below a predetermined threshold are classified outliers. This procedure is evaluated on our toy example of Figure 1. Data represented by the dots in the top right corner all end with a null score (see Figure 9 below). The procedure is then tested on a real dataset on which the conclusions are more interesting. On this experiment, according to the theoretical bounds in Theorem 2, MOM minimizer’s prediction qualities are all the better when the number of blocks is small, cf. Figure 8. On the other hand, outlier detection is best when the number of blocks is large, cf. Figure 10. Outlier resistance and anomaly detection tasks can therefore both be handled using the MOM principle, but the main hyper-parameter K – the number of blocks of data – for setting this method must be chosen carefully according to the objective. A number of blocks as small as possible (about twice the number of outliers) will give the best

predictions, while large values of this number of blocks will accelerate the detection of anomalies. Note that it is essential for outliers detection to use different (for instance, random) partitions at each step of the descent to avoid giving the same score to an outlier and to all the data in the same block containing it.

Detecting outliers is usually performed in machine learning via some unsupervised preprocessing algorithm that detects outliers outside a bulk of data, see for example [35, 30, 17, 58] or other algorithms like DBSCAN [7] or isolation forest [46]. These algorithms assume elliptical symmetry of the data, a solution for skewed data can also be found in [36]. These unsupervised preprocessing removes outliers in advance, i.e. before starting any learning task. As expected, these strategies work well in the toy example from Figure 1. There are several cases where it will fail though. First, as explained in [34], this strategy classifies data independently of the risk, it is likely to remove from the dataset outlier coming from heavy-tailed distribution, yielding biased estimators. Moreover, a small group of misclassified data inside a bulk won't be detected. Our notion of depth, based on the risk, seems more adapted to the learning task than any preprocessing procedure blind to the risk.

The paper is organized as follows. Section 2 presents the classification problem, the ERM and its MOM versions and gathers the assumptions granted for the main results. Section 3 presents theoretical risk bounds for the ERM and MOM minimizers on corrupted datasets. Section 4 deals with theoretical results on the algorithm computing MOM minimizers. We present the algorithm, study its convergence and provide theoretical bounds on its complexity. Section 5 shows empirical performance of our estimators in both simulated and real datasets. Proofs of the main results are postponed to Section 6 where we also added heuristics on the practical choice of the hyper-parameters.

2 Setting

2.1 Empirical risk minimization for binary classification

Consider the supervised binary classification problem, where one observes a sample $(X_1, Y_1), \dots, (X_N, Y_N)$ taking values in $\mathcal{X} \times \mathcal{Y}$. The set \mathcal{X} is a measurable space and $\mathcal{Y} = \{-1, 1\}$. The goal is to build a classifier—that is, a measurable map $f : \mathcal{X} \rightarrow \mathcal{Y}$ —such that, for any new observation (X, Y) , $f(X)$ is a good prediction for Y . For any classifier f , let

$$\ell_f^{0-1}(x, y) = I\{y \neq f(x)\}, \quad R^{0-1}(f) = P\ell_f^{0-1} = \mathbb{P}_{(X, Y) \sim P}(Y \neq f(X)) .$$

The 0–1 risk $R^{0-1}(\cdot)$ is a standard measure of the quality of a classifier. Following Vapnik [67], a popular way to build estimators is to replace the unknown measure P in the definition of the risk by the empirical measure P_N defined for any real valued function g by $P_N g = N^{-1} \sum_{i=1}^N g(X_i, Y_i)$ and minimize the empirical risk. The *empirical risk minimizer* for the 0–1 loss on a class \mathcal{F} of classifiers is $\hat{f}_{\text{ERM}}^{0-1} \in \operatorname{argmin}_{f \in \mathcal{F}} \{P_N \ell_f^{0-1}\}$.

The main issue with $\hat{f}_{\text{ERM}}^{0-1}$ is that it cannot be computed efficiently in general. One source of computational complexity is that both \mathcal{F} and the 0–1 loss function are non-convex. This is why various convex relaxations of the 0–1 loss have been introduced since the very beginning of machine learning. These proceed in two steps. First, \mathcal{F} should be replaced by a convex set F of functions taking values in \mathbb{R} . Then one builds an alternative loss function ℓ for ℓ^{0-1} defined for all $f \in F$. The new function ℓ should be convex and put less weight on those $f \in F$ such that $f(X_i)Y_i > 0$. Classical examples include the *hinge loss* $\ell_f^{\text{hinge}}(x, y) = (1 - yf(x))_+$, or the *logistic loss* $\ell_f^{\text{logistic}}(x, y) = \log(1 + e^{-yf(x)})$. A couple (F, ℓ) such that F is a convex set of real valued functions and ℓ is a convex function (i.e. for all $y \in \{-1, 1\}$ and $x \in \mathcal{X}$, $f \in F \rightarrow \ell_f(x, y)$ is convex) such that $\ell_f(x, y) < \ell_f(x, -y)$ whenever $yf(x) > 0$ will be called a convex relaxation of $(\mathcal{F}, \ell^{0-1})$. Given a convex relaxation (F, ℓ) of $(\mathcal{F}, \ell^{0-1})$, one can define the associated empirical risk minimizer (ERM) by

$$\hat{f}_{\text{ERM}} \in \operatorname{argmin}_{f \in F} P_N \ell_f . \quad (3)$$

Note that \hat{f}_{ERM} does not build a classifier. To deduce, a classification rule from \hat{f}_{ERM} one can simply consider its sign function defined for all $x \in \mathcal{X}$ by $\operatorname{sign}(\hat{f}_{\text{ERM}}(x)) = 2(I\{\hat{f}_{\text{ERM}}(x) \geq 0\} - 1/2)$. The

procedure \widehat{f}_{ERM} is solution of a convex optimization problem that can therefore be approximated using a descent algorithm. We refer for example to [14] for a recent overview of this topic and Section 4 for more examples.

2.2 Corrupted datasets

In this paper, we consider a framework where the dataset may have been corrupted by *outliers* (or anomalies). There are several definitions of outliers in the literature, here, we assume that the dataset is divided into two parts. The first part is the set of inliers, indexed by \mathcal{I} , data $(X_i, Y_i)_{i \in \mathcal{I}}$ are hereafter always assumed to be independent and identically distributed (i.i.d.) with common distribution P . The second one is the set of outliers, indexed by $\mathcal{O} \subset [N]$ which has cardinality $|\mathcal{O}|$. Nothing is assumed on these data which may not be independent, have distributions P_i totally different from P , with no moment at all, etc... In particular, this framework is sufficiently general to cover the case where outliers are i.i.d. with distribution $Q \neq P$ as in the ϵ -contamination model [34, 16, 25, 21].

Our first result shows that the rate of convergence of $\widehat{f}_{\text{ERM}}^{0-1}$ is not affected by this corruption as long as $|\mathcal{O}|$ does not exceed $N \times (\text{rate of convergence})$ see Theorem 1 and the remark afterward. However, it is easy to remark that, when the number N of data is finite as it is always the case in practice, even one aggressive outliers may yield disastrous breakdown of the empirical mean's statistical performance. Consequently, even if $\widehat{f}_{\text{ERM}}^{0-1}$ behaves correctly, its proxy \widehat{f}_{ERM} defined in (3) for a convex relaxation (F, ℓ) can have disastrous statistical performances, particularly when F and ℓ are unbounded, cf. Figure 5 for an illustration.

To bypass this problem, we consider in this paper an alternative to the empirical mean called *median-of-means* [2, 37, 59]. Let $K \leq N$ denote an integer and let B_1, \dots, B_K denote a partition of $\{1, \dots, N\}$ into bins B_k of equal size $|B_k| = N/K$. If K doesn't divide N , one can always drop a few data. For any function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ and any non-empty subset $B \subset \{1, \dots, N\}$, define the empirical mean on B by $P_B f = |B|^{-1} \sum_{i \in B} f(X_i, Y_i)$. The median-of-means (MOM) estimator of Pf is defined as the empirical median of the empirical means on the blocks B_k

$$\text{MOM}_K(f) = \text{median} \{P_{B_k} f : k = 1, \dots, K\} .$$

As the classical Huber's estimator [32], MOM estimators interpolate between the unbiased but non robust empirical mean and the robust but biased median. In particular, when applied to loss functions, these new estimators of the risk $P\ell_f, f \in F$ suggest to define the following alternative to Vapnik's ERM, called MOM minimizers

$$\widehat{f}_{\text{MOM}, K} \in \underset{f \in F}{\text{argmin}} \text{MOM}_K(\ell_f) . \quad (4)$$

From a theoretical point of view, we will prove that, when the number $|\mathcal{O}|$ of outliers is smaller than $N \times (\text{rate of convergence})$, $\widehat{f}_{\text{MOM}, K}$ performs well under 2-moments assumptions on F and ℓ . To illustrate our main assumptions and theoretical results, we will regularly use the following classical example.

Example 1 (linear classification.). *Let $\mathcal{X} = \mathbb{R}^p$ and let $\|\cdot\|_2$ denote the classical Euclidean norm on \mathbb{R}^p . Let F denote a set of linear functions*

$$F = \{f_t : x \mapsto \langle x, t \rangle : \|t\|_2 \leq \Gamma\} .$$

Let ℓ denote either the hinge loss or the logistic loss defined respectively for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and $f \in F$ by

$$\ell_f^{\text{hinge}}(x, y) = (1 - yf(x))_+, \quad \ell_f^{\text{logistic}}(x, y) = \log(1 + e^{-yf(x)}) .$$

2.3 Main assumptions

As already mentioned, data are divided into two groups, a subset $\{(X_i, Y_i) : i \in \mathcal{O}\}$ made of outliers (on which we will make no assumption) and the remaining data $\{(X_i, Y_i) : i \in \mathcal{I}\}$ contains all data that bring

information on the target/oracle

$$f^* \in \operatorname{argmin}_{f \in F} P\ell_f .$$

Data indexed by \mathcal{I} are therefore called *inliers* or *informative data*. To keep the presentation as simple as possible, inliers are assumed to be i.i.d. distributed according to P although this assumption could be relaxed as in [42, 43]. Finally, note that the $\mathcal{O} \cup \mathcal{I} = \{1, \dots, N\}$ partition of the dataset is of course unknown from the statistician.

Let us now turn to the set of assumptions we will use to study MOM minimizers procedures. For any measure Q and any function f for which it makes sense, denote by $Qf = \int f dQ$. Denote also, for all $q \geq 1$, by L^q the set of real valued functions f such that $\int |f|^q dP < \infty$ and, for any $f \in L^q$, by

$$\|f\|_{L^q} = \left(\int |f|^q dP \right)^{1/q} .$$

Our first assumption is an L^2 -assumption on the functions in F .

Assumption 1. *For all $f \in F$, we have $\|f\|_{L^2} \leq \theta_2$.*

Of course, Assumption 1 is granted if F is a set of classifiers. It also holds for the linear class of functions from Example 1 as long as $P\|X\|_2^2 < \infty$ with $\theta_2 = \Gamma(P\|X\|_2^2)^{1/2}$.

The second assumption deals with the complexity of the class F . This complexity appears in the upper bound of the risk. It is defined using only informative data. Let

$$\mathcal{K} = \{k \in \{1, \dots, K\} : B_k \cap \mathcal{O} = \emptyset\} \quad \text{and} \quad \mathcal{J} = \cup_{k \in \mathcal{K}} B_k .$$

Definition 1. *Let \mathcal{G} denote a class of functions $f : \mathcal{X} \rightarrow \mathbb{R}$ and let $(\epsilon_i)_{i \in \mathcal{I}}$ denote i.i.d. Rademacher random variables independent from $(X_i, Y_i)_{i \in \mathcal{I}}$. The Rademacher complexity of \mathcal{G} is defined by*

$$\mathcal{R}(\mathcal{G}) = \max_{A \in \{\mathcal{I}, \mathcal{J}\}} \mathbb{E} \left[\sup_{f \in \mathcal{G}} \sum_{i \in A} \epsilon_i f(X_i) \right] .$$

The Rademacher complexity is a standard measure of complexity in classification problems [5]. It can be upper bounded, depending on the situation, by the VC dimension or the Dudley's entropy integral or the Gaussian mean width of the class F , see for example [9, 39, 5, 10, 18] for a presentation of these classical bounds. Our second assumption is simply that the Rademacher complexity of the class F is finite.

Assumption 2. *The Rademacher complexity of F is finite, $\mathcal{R}(F) < \infty$.*

Assumption 2 holds in the linear classification example under Assumption 1 since it follows from Cauchy-Schwarz inequality that $\mathcal{R}(F) \leq \theta_2 \sqrt{|\mathcal{I}|p}$. Finally, our last assumption is that the loss function ℓ considered is Lipschitz in the following sense.

Assumption 3. *The loss function ℓ satisfies for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and all $f, f' \in F$,*

$$|\ell_f(x, y) - \ell_{f'}(x, y)| \leq L|f(x) - f'(x)| .$$

Assumption 3 holds for classical convex relaxation of the 0 – 1 loss such as hinge loss ℓ^{hinge} or logistic loss ℓ^{logistic} as in Example 1. In these examples, the constant L can be chosen equal to 1.

3 Theoretical guarantees

Our first result follows Vapnik’s original risk bound for the ERM and shows that $\hat{f}_{\text{ERM}}^{0-1}$ is insensitive to the presence of outliers in the dataset. Moreover, it quantifies this robustness property since Vapnik’s rate of convergence is still achieved by $\hat{f}_{\text{ERM}}^{0-1}$ when there are less than (number of observations) times (Vapnik’s rate of convergence) outliers.

Theorem 1. *Let \mathcal{F} denote a collection of classifiers. Let $\mathcal{L}_{\mathcal{F}}^{0-1} = \{\ell_f^{0-1} - \ell_{f^*}^{0-1} : f \in \mathcal{F}\}$ be the family of excess loss functions indexed by \mathcal{F} where $f^* \in \operatorname{argmin}_{f \in \mathcal{F}} R^{0-1}(f)$. For all $K > 0$, with probability at least $1 - e^{-K}$, we have*

$$R^{0-1}(\hat{f}_{\text{ERM}}^{0-1}) - \inf_{f \in \mathcal{F}} R^{0-1}(f) \leq \frac{2\mathcal{R}(\mathcal{L}_{\mathcal{F}}^{0-1})}{N} + \sqrt{\frac{K}{2|\mathcal{I}|}} + \frac{2|\mathcal{O}|}{N}.$$

Theorem 1 is proved in Section 6.1. It is an adaptation of Vapnik’s proof of the excess risk bounds satisfied by $\hat{f}_{\text{ERM}}^{0-1}$ in the presence of outliers.

Remark 1. *In the last result, one can easily bound the excess risk using $\mathcal{R}(\mathcal{F})$ instead of $\mathcal{R}(\mathcal{L}_{\mathcal{F}}^{0-1})$ since*

$$\mathcal{R}(\mathcal{L}_{\mathcal{F}}^{0-1}) = \max_{A \in \{\mathcal{I}, \mathcal{J}\}} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \sum_{i \in A} \epsilon_i (f(X_i) - f^*(X_i)) \right] = \mathcal{R}(\mathcal{F}).$$

The final bound is of similar flavor: for all $K > 0$, with probability at least $1 - e^{-K}$, we have

$$R^{0-1}(\hat{f}_{\text{ERM}}^{0-1}) - \inf_{f \in \mathcal{F}} R^{0-1}(f) \lesssim \max \left(\frac{\mathcal{R}(\mathcal{F})}{N}, \sqrt{\frac{K}{|\mathcal{I}|}}, \frac{|\mathcal{O}|}{N} \right). \quad (5)$$

Remark 2. *When \mathcal{F} is the class of all linear classifiers, that is when $\mathcal{F} = \{\operatorname{sgn}(\langle t, \cdot \rangle) : t \in \mathbb{R}^p\}$, one has $\mathcal{R}(\mathcal{F}) \leq \sqrt{|\mathcal{I}|}p$ (see Theorem 3.4 in [9]). Therefore, when $|\mathcal{I}| \geq N/2$, Theorem 1 implies that for all $1 \leq K \leq p$, with probability at least $1 - \exp(-K)$,*

$$R^{0-1}(\hat{f}_{\text{ERM}}^{0-1}) - \inf_{f \in \mathcal{F}} R^{0-1}(f) \lesssim \max(\sqrt{p/N}, |\mathcal{O}|/N).$$

As a consequence, when the number of outliers is such that $|\mathcal{O}| \lesssim N \times \sqrt{p/N}$, Vapnik’s classical “slow” rate of convergence $\sqrt{p/N}$ is still achieved by the ERM even if $|\mathcal{O}|$ outliers have polluted the dataset. The interested reader can also check that “fast rates” p/N could also be achieved by the ERM in the presence of outliers if $|\mathcal{O}| \lesssim p$ and when the so-called strong margin assumption holds (see, [9, 9]). Note also that the previous remark also holds if F is a class with VC dimension p beyond the case of indicators of half spaces.

The conclusion of Theorem 1 can be misleading in practice. Indeed, theoretical performance of the ERM for the 0 – 1 loss function are not downgraded by outliers, but its proxies based on convex relaxation (F, ℓ) of $(\mathcal{F}, \ell^{0-1})$ are. This can be seen on the toy example in Figure 1 and in Figure 5 from Section 5. In this work, we propose a robust surrogate, based on MOM estimators of the risk and defined in (2), to the natural empirical risk estimation of the risk which works for unbounded loss functions. In the next result, we prove that the MOM minimizer $\hat{f}_{\text{MOM},K}$ defined as

$$\hat{f}_{\text{MOM},K} \in \operatorname{argmin}_{f \in F} \text{MOM}_K(\ell_f) \quad (6)$$

satisfies an excess risk bound under weak assumptions introduced in Section 2.

Theorem 2. *Grant Assumptions 1, 2 and 3. Assume that $N > K > 4|\mathcal{O}|$ and let $\Delta = 1/4 - |\mathcal{O}|/K$. Then, with probability larger than $1 - 2 \exp(-2\Delta^2 K)$, we have*

$$R(\hat{f}_{MOM,K}) \leq \inf_{f \in F} R(f) + 4L \max \left(\frac{4\mathcal{R}(F)}{N}, 2\theta_2 \sqrt{\frac{K}{N}} \right) .$$

Theorem 2 is proved in Section 6.2. Compared to Theorem 1, $\hat{f}_{MOM,K}$ achieves the same rate $(\mathcal{R}(F)/N) \vee (\sqrt{K/N})$ under the same conditions on the number of outliers with the same exponential control of the probability as for the ERM f_{ERM}^{0-1} . The main difference is that the loss function may be unbounded, which is often the case in practice. Moreover, unlike classical analysis of ERM obtained by minimizing an empirical risk associated with a convex surrogate loss function, we only need a second moment assumption on the class F .

These theoretical improvements have already been noticed in previous works [56, 20, 42, 47, 48, 49, 53, 43]. Contrary to tournaments of [49], Le Cam MOM estimators of [42] or minmax MOM estimators [43], Theorem 2 does not require the small ball assumption on F but only shows “slow rates” of convergence. These slow rates are minimax optimal in the absence of a margin or Bernstein assumption [6, 51]. Proof of Theorem 2 does not enable fast rates to be obtained. Indeed, the non-linearity of the median excludes the possibility of using localization techniques leading to these fast rates. However, we show in the simulation study (cf. left side picture of Figure 14) that fast rates seem to be reached by the MOM minimizer.

The main advantage of our approach is its simplicity, we just have to replace empirical means by their MOM alternative in the definition of the ERM. Moreover, as expected, this simple alternative to ERM yields a systematically way to modify algorithms designed for approximating the ERM. The resulting “MOM version” of these algorithms are more robustness and faster than their original “ERM version”. Before illustrating these facts on simulations, let us describe algorithms approximating MOM minimizers.

4 Computation of MOM minimizers

In this section, we present first a generic algorithm to provide a MOM version of descent algorithms minimizing the empirical risk. We show how classical convergence proofs can be adapted to these algorithms.

4.1 MOM algorithms

The general idea is that any descent algorithms such as gradient descent, Newton method, alternate gradient descent, etc. (cf. [57, 14, 11, 4]) can easily be turned into a robust MOM-version. To illustrate this idea, a basic gradient descent is analyzed in the sequel.

The choice of blocks greatly influences the practical performance of the algorithm. In particular, a recurring flaw is that iterations tend to get stuck in local minima, which greatly slows the convergence of the algorithm. To overcome this default and improve the stability of the procedure, we modify the partition at each iteration by drawing it uniformly at random, cf. step 2 of Algorithm 1.

Let \mathcal{S}_N denote the set of permutations of $\{1, \dots, N\}$. For each $\sigma \in \mathcal{S}_N$, let $B_0(\sigma) \cup \dots \cup B_{K-1}(\sigma) = \{1, \dots, N\}$ denote an equipartition of $\{1, \dots, N\}$ defined for all $j \in \llbracket 0, K-1 \rrbracket$ by

$$B_j(\sigma) = \{\sigma(Kj+1), \dots, \sigma(K(j+1))\} = \sigma(\{Kj+1, \dots, K(j+1)\}) .$$

To simplify the presentation, let us assume the class F to be parametrized $F = \{f_u : u \in \mathbb{R}^p\}$, for some $p \in \mathbb{N}^*$. Let’s assume that the function $u \mapsto f_u$ is as regular as needed and convex (a typical example is $f_u(x) = \langle u, x \rangle$ for all $x \in \mathbb{R}^p$). Denote by $\nabla_u \ell_{f_u}$ the gradient or subgradient of $u \mapsto f_u$ in $u \in \mathbb{R}^p$. The step-sizes sequence is denoted by $(\eta_t)_{t \geq 0}$ and satisfies the classical conditions: $\sum_{t=1}^{\infty} \eta_t = \infty$ and $\sum_{t=1}^{\infty} \eta_t^2 < \infty$. Iterations will go on until a stopping time $T \in \mathbb{N}^*$ has been achieved. With these notations, a generic MOM version of a gradient descent algorithm (with random choice of blocks) is detailed in Algorithm 1.

<p>input : $u_0 \in \mathbb{R}^p$, $K \in \llbracket 3, N/2 \rrbracket$, $T \in \mathbb{N}^*$ and $(\eta_t)_{t \in \{0, \dots, T-1\}} \in \mathbb{R}_+^T$</p> <p>output: approximated solution to the minimization problem (4)</p> <pre> 1 for $t = 0, \dots, T - 1$ do 2 choose a permutation at random: $\sigma_t \sim \text{Unif}(\mathcal{S}_N)$, 3 build a partition of the dataset: $B_0(\sigma_t), \dots, B_{K-1}(\sigma_t)$, 4 find a median block: $k_{med}(t)$ s.t. $\text{MOM}_K(\ell_{f_{u_t}}) = P_{B_{k_{med}(t)}(\sigma_t)}(\ell_{f_{u_t}})$, 5 do a descent step $u_{t+1} = u_t - \eta_t \sum_{i \in B_{k_{med}(t)}(\sigma_t)} \nabla_{u_t} \ell_{f_{u_t}}(X_i, Y_i).$ 6 end 7 Return u_T </pre>
--

Algorithm 1: MOM gradient descent algorithm.

Remark 3 (MOM gradient descent algorithm and stochastic block gradient descent). *Algorithm 1 can be seen as a stochastic block gradient descent (SBGD) algorithm minimizing the function $t \rightarrow \mathbb{E} \ell_t(X, Y)$ using a given dataset. The main difference with the classical SBGD is that the choice of the block along which the gradient direction is performed is chosen according to a centrality measure computed thanks to the median operator in step 4 of Algorithm 1.*

In Section 5, we use MOM ideas (as in the generic Algorithm 1) to construct MOM versions for various classical algorithms such as Perceptron, Logistic Regression, Kernel Logistic Regression, SGD Classifiers or Multi-layer Perceptron.

4.2 Convergence properties

We denote by $\mathcal{D}_N = \{(X_1, Y_1), \dots, (X_N, Y_N)\}$ the dataset. We consider the following set of assumptions:

Assumption 4. *There exist $L > 0$ and a sequence of decreasing positive numbers $(\eta_t)_{t \geq 1}$, such that*

H-1 *for all $u \in \mathbb{R}^p$, and P -almost all $x, y \in \mathcal{X} \times \{-1, 1\}$, $\|\nabla_u \ell_{f_u}(x, y)\|_2 \leq L$,*

H-2 *for almost all datasets $\mathcal{D}_N = \{(X_i, Y_i) : i = 1, \dots, N\}$, there exists a unique minimum $u_{min} = \text{argmin}_{u \in \mathbb{R}^p} \mathbb{E}_\sigma[\text{MOM}_K(\ell_{f_u}) | \mathcal{D}_N]$ where \mathbb{E}_σ denotes the expectation with respect to σ , the uniformly distributed permutation used to construct the blocks of data,*

H-3 $\sum_{t=1}^{\infty} \eta_t^2 < \infty$,

H-4 $\sum_{t=1}^{\infty} \eta_t = \infty$,

H-5 *for almost all datasets $\mathcal{D}_N = \{(X_i, Y_i) : i = 1, \dots, N\}$, for Lebesgue almost all $u \in \mathbb{R}^p$, for all $\varepsilon > 0$, $\inf_{\|u_{min} - u\|_2 > \varepsilon} (u_{min} - u)^T \mathbb{E}_\sigma[\nabla_u \text{MOM}_K(\ell_{f_u}) | \mathcal{D}_N] > 0$.*

H-6 *for almost all datasets $\mathcal{D}_N = \{(X_i, Y_i) : i = 1, \dots, N\}$ and Lebesgue almost all $u \in \mathbb{R}^p$, there exists an open convex set B containing u such that for any equipartition of $\{1, \dots, N\}$ into K blocks B_1, \dots, B_K there exists $k_{med} \in \{1, \dots, K\}$ such that for all $v \in B$, $P_{B_{k_{med}}}(\ell_{f_v}) \in \text{MOM}_K(\ell_{f_v})$.*

Recall that the dataset $\mathcal{D}_N = \{(X_i, Y_i) : i = 1, \dots, N\}$ may contain outlier, which might be deterministic. We use the terminology “for almost all datasets $\mathcal{D}_N = \{(X_i, Y_i) : i = 1, \dots, N\}$ such assumption holds” when the latter holds for $P^{|\mathcal{I}|}$ -almost all informative data $(X_i, Y_i)_{i \in \mathcal{I}}$ and for every outliers $(X_i, Y_i)_{i \in \mathcal{O}} \in (\mathcal{X} \times \mathcal{Y})^{|\mathcal{O}|}$.

Theorem 3 (Convergence of the algorithm). *Grant Assumptions **H-1** to **H-6**, then the MOM gradient descent algorithm 1 $(u_t)_{t \in \mathbb{N}}$ converges: for almost all \mathcal{D}_N , we have*

$$\|u_T - u_{min}\| \xrightarrow[T \rightarrow \infty]{a.s.} 0$$

(where the a.s. convergence is with respect to the randomness on the choice of data partitions at each step).

Note that Theorem 3 insures the convergence of the MOM version of the gradient descent algorithm with random blocks as defined in Algorithm 1 toward a minimizer of the conditional expectation $u \rightarrow \mathbb{E}_\sigma[\text{MOM}_K(\ell_{f_u})|\mathcal{D}_N]$. It does not prove its convergence toward a minimizer of the MOM criteria $u \rightarrow \text{MOM}_K(\ell_{f_u})$. Nevertheless, we can expect a strong concentration of $\text{MOM}_K(\ell_{f_u})$ around $\mathbb{E}_\sigma[\text{MOM}_K(\ell_{f_u})|\mathcal{D}_N]$ by concentration of random permutations, see [52].

Remark 4. [MOM gradient descent, SGD and assumption **H-6**] All assumptions in Theorem 3 are the classical assumption used to prove the convergence of the classical SGD except for **H-6** which is specific to MOM algorithms.

Indeed, the picture of the MOM gradient descent algorithm is pretty simple and depicted in Figure 2. At every step t , the median operator makes a partition of \mathbb{R}^p into K cells $\mathcal{C}_k(t) = \{u \in \mathbb{R}^p : \text{MOM}_K(\ell_{f_u}) = P_{B_k} \ell_{f_u}\}$ for $k = 1, \dots, K$ – this partition changes at every step because the blocks B_1, \dots, B_K are chosen randomly at the beginning of every step according to the random partition σ_t . We want every iteration u_t of the MOM algorithm to be in the interior of a cell and not on a frontier in order to differentiate the objective function $u \rightarrow \text{MOM}_K(\ell_{f_u})$ at u_t . This is indeed the case under **H-6**, given that in that case, there is an open neighbor B of u_t such that for all $v \in B$, $\text{MOM}_K(\ell_{f_v}) = P_{B_k} \ell_{f_v}$ where the index $k = k_{med}$ of the block is common to every $v \in B$. Therefore, to differentiate the objective function $u \rightarrow \text{MOM}_K(\ell_{f_u})$ at u_t one just needs to differentiate $u \rightarrow P_{B_k} \ell_{f_u}$ at u_t . The objective function to minimize is differentiable almost everywhere under **H-6** and a gradient of the objective function is given by $\nabla(u \rightarrow P_{B_k} \ell_{f_u})|_{u=u_t}$.




Figure 2: Partition of \mathbb{R}^p at step t by the median operator and iteration number $t - 2, t - 1, t$ and $t + 1$ of the MOM gradient descent algorithm. Under Assumption **H-6**, there is a natural descent direction given at step t by $-\nabla_u(u \rightarrow P_{B_{k_{med}(t)}(\sigma_t)}(\ell_{f_u}))|_{u=u_t}$.

Remark 5 (Assumption **H-5**). *The assumption **H-5** illustrate why we used a random shuffle of the data at each step of the algorithm, the principle is that the MOM risk in general may present local minima and*

hence the algorithm without the shuffle would stop in a local minimum. By using the random shuffle at each step, it seems that the problem has a unique minimum, at least experiment seems to say so. For example, in figure 3 we see the plot of the loss function when we change one of the parameters (the intercept) of a linear classifier for the classification problem depicted in figure 4.

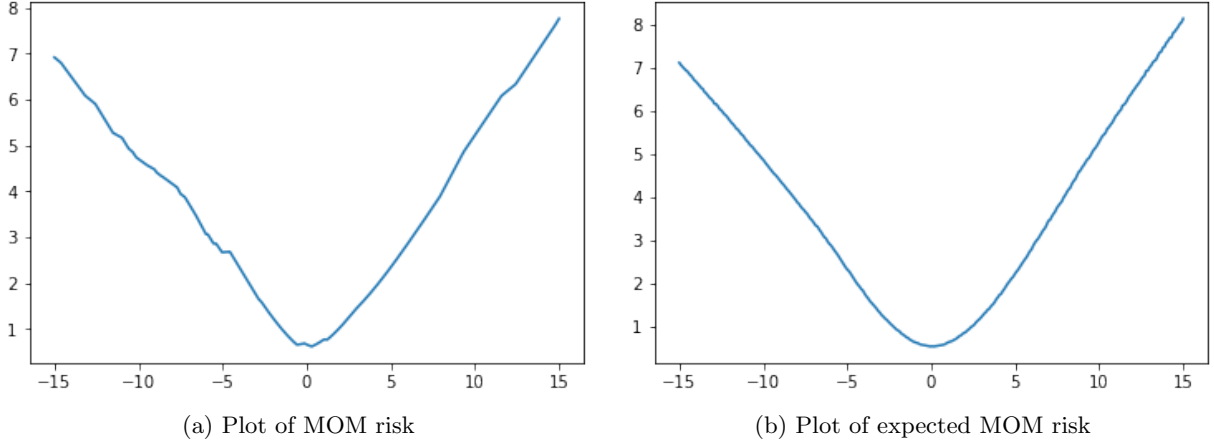


Figure 3: Plot of the two loss functions the MOM risk $MOM_K(\ell_{f_b})$ and the expected MOM risk $\mathbb{E}[MOM_K(\ell_{f_b})|\mathcal{D}_N]$ (estimated using Monte Carlo method with 300 samples), as a function of b where ℓ is the logistic loss and f_b is the linear classifier $f_b(x) = bx_1 + x_2 + x_3, \forall x = (x_j)_{j=1}^3 \in \mathbb{R}^3$ and $b \in \mathbb{R}$.

*Remark that assumption **H-5** is not convexity, and indeed in figure 3 the expected loss is not convex but it verifies assumption **H-5** nonetheless.*

Example 2. Let us give a relatively general example where Assumption **H-6** is satisfied. Let $B_1 \cup \dots \cup B_K = \{1, \dots, N\}$ be an equipartition and let ψ be defined for all $x = (x_i)_{i=1}^N \in \mathbb{R}^N$ by,

$$\psi_p(x) = MOM_K(f_p(x)) = \text{median} \left(\frac{K}{N} \sum_{i \in B_k} f_p(x_i), 1 \leq k \leq K \right) = P_{B(K/2)(p)}(f_p),$$

where for all blocks $B \subset \{1, \dots, N\}$, $P_B f_p = |B|^{-1} \sum_{i \in B} f_p(x_i)$ and the blocks $B(k)(p), k = 1, \dots, K$ are rearranged blocks defined such that $P_{B(1)(p)}(f_p) \geq \dots \geq P_{B(K)(p)}(f_p)$. Proposition 1 shows that Assumption **H-6** is satisfied in several situations.

Before stating the proposition, recall that the hazard function of an absolutely continuous real-valued random variable Z is defined for all $x \in \mathbb{R}$ by $h(x) = f(x)/(1 - F(x))$ where f is the density function of Z and F is its cumulative distribution function.

Proposition 1. Let X_1, \dots, X_N be N real-valued random variables and let $(f_p)_{p \in \mathbb{R}^d}$ be a family of functions with values in \mathbb{R} , if for all $x \in \mathbb{R}$, the function $p \mapsto f_p(x)$ is Lipschitz and the probability distribution of a mean of a block $P_{B(k)(p)}(f_p)$ is a random variable (recall that the choice of block is random) with a bounded hazard function, then with probability 1, Assumption **H-6** is satisfied, in particular, the partial derivative of $p \mapsto \psi_{f_p}((X_i)_{i=1}^N) = MOM_K(f_p((X_i)_{i=1}^N))$ with respect to the j^{th} coordinate is given for almost all X_1, \dots, X_N by

$$\partial_j \psi_{f_p}((X_i)_{i=1}^N) = \frac{K}{N} \sum_{i \in B(K/2)(p)} \partial_j f_p(X_i)$$

where ∂_j denote the derivative with respect to the j^{th} coordinate of p .

Proposition 1 is proved in Section 6. The function h is bounded for a large panel of probability distributions. For example, an exponential variable with parameter λ has a constant hazard function equal to λ ; a Cauchy random variable has also a bounded hazard function. Intuitively, this property of derivation of the median of means is more easily verified when the distribution of $P_{B(k)(p)}(f)$ is heavy tailed (i.e. the mean blocks are far from one another) and in fact we can verify that the hazard function of the Gaussian is unbounded and hence this proposition cannot be used on Gaussian random variables.

4.3 Complexity of MOM risk minimization algorithms

Let $C(m)$ be the computational complexity of a single ERM gradient descent update step on a dataset of size m and let $L(m)$ be the computational complexity of the evaluation of the loss function on a dataset containing m data. Here the computational complexity is simply the number of basic operations needed to perform a task [3].

For each epoch, we begin by computing the ‘‘MOM empirical risk’’. We perform K times N/K evaluations of the loss function, then we sort the K means of these blocks of loss to finally get the median. The complexity of this step is then $O(KL(N/K) + K \ln(K))$, assuming that the sort algorithm is in $O(K \ln(K))$ (like *quick sort* [31]). Then we do the ERM gradient step on a sample of size N/K . Hence, the time complexity of this algorithm is

$$O(T(KL(N/K) + K \ln(K) + C(N/K))).$$

Example 3 (Linear complexity ERM algorithms). *For example, if the ERM gradient step and the loss function evaluation have linear complexity – like Perceptron or Logistic Regression – the complexity of the MOM algorithm is $O(T(N + K \log(K)))$ against $O(TN)$ for the ERM algorithm. Therefore, the two complexities are of the same order and the only advantage of MOM algorithms lies in their robustness to outliers.*

Example 4 (Super-linear complexity ERM algorithms). *If on the other hand the complexity is more than linear as for Kernel Logistic Regression (KLR), taking into account the matrix multiplications whose complexity can be found in [24], the complexity of the MOM version of KLR, due to the additional need of the computation of the kernel matrix, is $O(N^2 + T(N^2/K + K \log(K) + (N/K)^{2.373}))$ against $O(TN^{2.373})$ for the ERM version. MOM versions of KLR are therefore faster than the classical version of KLR on top of being more robust. This advantage comes from the fact that MOM algorithms work on blocks of data instead on the entire dataset at every step. More informations about Kernel Logistic Regression can be found in [62] for example.*

In this last example, the complexity comes in part from the evaluation of the kernel matrix that can be computationally expensive. Following the idea that MOM algorithms are performing ERM algorithm restricted to a wisely chosen block of data, then one can modify our generic strategy in this particular example to reduce drastically its complexity. The idea here is that we only need to construct the kernel matrix on the median block. The resulting algorithm, called Fast KLR MOM is described in Figure 2.

In Figure 2, we compute only the block kernel matrices, denoted by N^1, \dots, N^k and constructed from the samples in the block B_k . We also denote by N_i^k the i^{th} row in N^k .

input : $\alpha_0 \in \mathbb{R}^p$, $K \in \llbracket 3, N/2 \rrbracket$, $T \in \mathbb{N}^*$, $(\eta_t)_{t \in \{0, \dots, T-1\}} \in \mathbb{R}_+^T$, $\beta \in \mathbb{R}_+^*$, $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ a positive definite kernel and a bloc decomposition B_1, \dots, B_K of $\{1, \dots, N\}$.

output: approximated solution to the minimization problem (4) among the KLR classifiers

- 1 Construct the bloc Kernel matrices $N^k = (\kappa(X_i, X_j))_{i, j \in B_k}$ for $1 \leq k \leq K$,
- 2 **for** $t = 0, \dots, T - 1$ **do**
- 3 find a median block: $k_{med}(t)$ s.t. $\text{MOM}_K(\ell_{f_{\alpha_t}}) = P_{B_{k_{med}(t)}}(\ell_{f_{\alpha_t}})$ with
- 4
$$P_{B_k}(\ell_{f_{\alpha_t}}) = \frac{1}{|B_k|} \sum_{i \in B_k} \ln(1 + e^{-N_i^k \alpha_t^k Y_i}) + \beta \sum_{k=1}^K (\alpha_t^k)^T N^k \alpha_t^k,$$
- 5 where α_t^k is the vector in $\mathbb{R}^{|B_k|}$ made of the coordinates of α_t with indices in B_k .
Do an IRLS descent step for KLR with weight matrix $W_{k_{med}(t)}$, design matrix $X_{k_{med}(t)}$ and labels $y_{k_{med}(t)}$ on $B_{k_{med}(t)}$

$$\alpha_{t+1}^{k_{med}(t)} = \alpha_t^{k_{med}(t)}(1 - \eta_t) + \eta_t (X_{k_{med}(t)}^T W_{k_{med}(t)} X_{k_{med}(t)})^{-1} X_{k_{med}(t)}^T W_{k_{med}(t)} y_{k_{med}(t)}.$$

$$\alpha_{t+1}^k = \alpha_t^k(1 - \eta_t), \quad \forall k \neq k_{med}(t).$$
- 6 **end**
- 7 **Return** $\alpha_T, N_{k_{med}(T)}$

Algorithm 2: Description of Fast KLR MOM algorithm.

There are several drawbacks in the approach of Algorithm 2. First, the blocks are fixed at the beginning of the algorithm; therefore the algorithm needs a bigger dataset to work well and it may converge to a local minimum. Nonetheless, from the complexity point of view, this algorithm will be much faster than both the classical KLR and MOM KLR (see below for a computation of its complexity) which is important given the growing use of kernel methods on very large databases for example in biology. The choice of K should ultimately realize a trade-off between complexity and performance (in term of accuracy for example) when dealing with big databases containing few outliers.

Example 5 (Complexity of Fast KLR-MOM algorithm). *The complexity of Fast KLR-MOM is $O(N^2/K + T(N^2/K + K \log(K) + (N/K)^{2.373}))$ against $O(TN^{2.373})$ for the ERM version.*

5 Implementation and Simulations

5.1 Basic results on a toy dataset

The toy model we consider models outliers due to human or machine errors we would like to ignore in our learning process. It is also a dataset corrupted to make linear classifiers fail. The dataset is a 2D dataset constituted of three “labeled Gaussian distribution”. Two informative Gaussians $\mathcal{N}((-1, -1), 1.4I_2)$ and $\mathcal{N}((1, 1), 1.4I_2)$ with label respectively 1 and -1 and one outliers Gaussian $\mathcal{N}((24, 8), 0.1I_2)$ with label 1. In other words, the distribution of informative data is given by $\mathcal{L}(X|Y = 1) = \mathcal{N}((-1, -1), 1.4I_2)$, $\mathcal{L}(X|Y = -1) = \mathcal{N}((1, 1), 1.4I_2)$ and $\mathbb{P}(Y = 1) = \mathbb{P}(Y = -1) = 1/2$. Outliers data have distribution given by $Y = 1$ a.s. and $X \sim \mathcal{N}((24, 8), 0.1I_2)$.

The algorithms we study are the MOM adaptations of Perceptron, Logistic Regression and Kernel Logistic Regression.

Based on our theoretical results, we know that the number of blocks K has to be larger than 4 times the number of outliers for our procedure to be on the safe side. The value $K = 120$ is therefore used in all subsequent applications of MOM algorithms on the toy dataset except when told otherwise. To quantify performance, we compute the miss-classification error on a clean dataset made of data distributed like the informative data.




Figure 4: Scatter plot of 630 samples from the training dataset (600 informative data, 30 outliers), the color of the points correspond to their labels.

For Kernel Logistic Regression, we study here a linear kernel because outliers in this dataset are clearly adversarial when dealing with linear classifiers. The algorithm can also use more sophisticated kernels, a comparison of the MOM algorithms with similar ERM algorithms is represented in figure 5, the ERM algorithms are taken from the python library scikit-learn [60] with their default parameters.

Figure 5 illustrates resistance to outliers of MOM’s algorithms compared to their classical version.




Figure 5: Scatter plot of 500 samples from the test dataset (500 informative data), the color of the points correspond to their labels and the background color correspond to the prediction. The score in the title of each subfigure is the accuracy of the algorithm.

These first results are completed in Figure 6 where we computed accuracy on several run of the algorithms. These results confirm the visual impression of our first experiment.

Finally, we illustrate our results regarding complexities of the algorithms on a simulated example. MOM algorithms have been computed together with state-of-the art algorithms from scikit-learn [60] (we use Random forest, SVM classifier as well as SGD classifier optimizing Huber loss which entail a robustness in Y but not in X , see [34, Chapter 7]) on a simulated dataset composed of two Gaussian blobs $\mathcal{N}((-1, -1), 1.4I_2)$ and $\mathcal{N}((1, 1), 1.4I_2)$ with label respectively 1 and -1 . We sample 20000 points for the training dataset and 20000 for the test dataset. The parameters used in the algorithms are those for which we obtained the optimal accuracy, (this accuracy is illustrated in the next section). Time of training plus time of evaluation on the test dataset are gathered in Figure 7.

Not surprisingly, very efficient versions of linear algorithms from Python’s library are extremely fast (results are sometimes provided before we even charged the dataset in some experiments). The performance




Figure 6: Comparison of the MOM algorithms and their counterpart with the boxplots of the accuracy on the test dataset from 50 runs of the algorithms on 50 sample of the training/test toy dataset (one run for each dataset sampled).

Algorithm	Perceptron MOM	Log. Reg. MOM	KLR MOM	Fast KLR MOM
Time (s)	1.06	1.05	13.6	1.2
Algorithm	Rand. Forest	SVM	SGD Hub. loss	
Time (s)	0.21	9.0	0.0078	

Figure 7: Time of different algorithms on a simulated dataset .

of our algorithm are nevertheless acceptable in general (around 5 times longer than random forest for example). The important fact here is that non linear algorithms such as SVM take much more time to provide a result. FAST KLR MOM is able to reduce substantially the execution time of SVM with comparable predictive performance.

5.2 Applications on real datasets

We used the HTRU2 dataset, also studied in [50], that is provided by the UCI Machine Learning Repository. The goal is to detect pulsars (a rare type of Neutron star) based on radio emission detectable on earth from which features are extracted to gives us this dataset. The problem is that most of the signal comes from noise and not pulsar, the goal is then to classify pulsar against noise, using the 17 898 points in the dataset.

The accuracy of different algorithms is obtained using on several runs of the algorithms each using 4/5 of the datasets for training and 1/5 for testing algorithms. Boxplots presenting performance of various algorithms are displayed in Figure 8. To improve performance, RBF kernel was used both for KLR MOM and Fast KLR MOM.

5.3 Outlier detection with MOM algorithms

When we run MOM version of a descent algorithm, we select at each step a block of data points realizing the median of a set of “local/block empirical risk” at the current iteration of the algorithm. The number of times a point is selected by the algorithm can be used as a depth function measuring reliability of the data. Note that this definition of depth of a data point has the advantage of taking into account the learning task we want to solve, that is the loss ℓ and the class F . It means that outliers are considered w.r.t. the problem we want to solve and not w.r.t. some a priori notion of centrality of points in \mathbb{R}^d unrelated with the problem considered at the beginning.

We apply this idea on the toy dataset with the Logistic Regression MOM algorithm. Results are gathered in a sorted histogram given in Figure 9. Red bars represent outliers in the original datasets.




Figure 8: Comparison of the MOM algorithms and common algorithms with the boxplots and the medians of the accuracy $\frac{1}{n} \sum_{i=1}^n \mathbb{1}\{\hat{f}(x_i) = y_i\}$ on the test dataset from 50 runs of the algorithms on 100 sample of a 4/5 cut of the dataset HTRU2 (one run is trained on a sample of 4/5 of the dataset and tested on the remaining 1/5)

Quite remarkably, outliers are in fact those data that have been used the smallest number of times. The method targets a very specific type of outliers, those disturbing the classification task at hand. If there was a point very far away from the bulk of data but in the half-space of its label, it wouldn't be detected.

This detection algorithm doesn't scale well when the dataset gets bigger as a large number of iterations is necessary to choose each point a fair number of times. For bigger datasets, we suggest to adapt usual outlier detection algorithms [1]. We emphasize that clustering techniques and K-Means are rather easy to adapt in a MOM algorithm and detect points far from the bulk of data. This technique might greatly improve usual K-Means as MOM K-Means is more robust.

Let us now analyze the effect of K on the outlier detection task. The histogram of the 1000 smaller counts of points of HTRU2 dataset as K gets bigger is plotted in Figure 10.

It appears from Figure 10 that K measures the sensitivity of the algorithm. Severe outliers (as in the toy example) are detected for small K while mild outliers are only discovered as K gets bigger.

It seems therefore that the optimal choice of K in MOM depends on the task one is interested in. For classification, K should be as small as possible to get better risk bounds (but it still should be larger than the number of outliers) whereas for detecting outliers we may want to choose K much larger to even detect an outlier, (but it should also be small enough for the underlying classification to perform correctly). As a proof of concept, for Pulsar database, we got optimal results choosing $K = 10$ for classification whereas we only detect a significant amount of outliers when K is around 1000.

6 Proofs

6.1 Proof of Theorem 1

We adapt Vapnik's classical analysis [66] of excess risk bound of ERM to a dataset corrupted by outliers. We first recall that $f^* \in \operatorname{argmin}_{f \in \mathcal{F}} R^{0-1}(f)$ and for all $f \in \mathcal{F}$, the excess loss function of f is $\mathcal{L}_f^{0-1} = \ell_f^{0-1} - \ell_{f^*}^{0-1}$. For simplicity we denote $\hat{f} = \hat{f}_{\text{ERM}}^{0-1}$ and for all $f \in \mathcal{F}$, $\mathcal{L}_f^{0-1} = \mathcal{L}_f$ and $R(f) = R^{0-1}(f)$.

It follows from the definition of the ERM that $P_N \mathcal{L}_{\hat{f}} \leq 0$. Therefore, if we denote by $P_{\mathcal{I}}$ (resp. $P_{\mathcal{O}}$)



Figure 9: Sorted Histogram of the score (number of times a data belongs to the selected median block) of each points in a Logistic Regression MOM algorithm on a toy dataset. Red is an outlier and blue is an informative sample. $K = 120$ and $T = 2000$.

the empirical measure supported on $\{(X_i, Y_i) : i \in \mathcal{I}\}$ (resp. $\{(X_i, Y_i) : i \in \mathcal{O}\}$), we have

$$\begin{aligned} R(\hat{f}) - R(f^*) &= (P - P_N)\mathcal{L}_{\hat{f}} + P_N\mathcal{L}_{\hat{f}} \leq (P - P_N)\mathcal{L}_{\hat{f}} = \frac{|\mathcal{I}|}{N}(P - P_{\mathcal{I}})\mathcal{L}_{\hat{f}} + \frac{|\mathcal{O}|}{N}(P - P_{\mathcal{O}})\mathcal{L}_{\hat{f}} \\ &\leq \frac{|\mathcal{I}|}{N} \sup_{f \in \mathcal{F}} (P - P_{\mathcal{I}})\mathcal{L}_f + \frac{2|\mathcal{O}|}{N} \end{aligned}$$

because $|\mathcal{L}_{\hat{f}}| \leq 1$ a.s.. Then, by the bounded difference inequality [10, Theorem 6.2], since all $f \in \mathcal{F}$ satisfies $-1 \leq \mathcal{L}_f \leq 1$, one has, for any $x > 0$,

$$\mathbb{P} \left(\sup_{f \in \mathcal{F}} (P - P_{\mathcal{I}})\mathcal{L}_f \geq \mathbb{E}[\sup_{f \in \mathcal{F}} (P - P_{\mathcal{I}})\mathcal{L}_f] + x \right) \leq e^{-2|\mathcal{I}|x^2} .$$

Furthermore, by the symmetrization argument (cf. Chapter 4 in [45]),

$$\mathbb{E}[\sup_{f \in \mathcal{F}} (P - P_{\mathcal{I}})\mathcal{L}_f] \leq 2 \frac{\mathcal{R}(\mathcal{L}_{\mathcal{F}})}{|\mathcal{I}|} .$$

Therefore, for any $x > 0$, with probability larger than $1 - e^{-2|\mathcal{I}|x^2}$,

$$R(\hat{f}) - R(f^*) \leq \frac{2\mathcal{R}(\mathcal{L}_{\mathcal{F}})}{N} + x + \frac{2|\mathcal{O}|}{N} .$$

The proof is completed by choosing $x = \sqrt{K/(2|\mathcal{I}|)}$.

6.2 Proof of Theorem 2

Let $f^* \in \operatorname{argmin}_{f \in \mathcal{F}} Pl_f$. By definition, one has $\operatorname{MOM}_K(\ell_{\hat{f}_{\operatorname{MOM},K}}) \leq \operatorname{MOM}_K(\ell_{f^*})$, therefore,

$$R(\hat{f}_{\operatorname{MOM},K}) - R(f^*) \leq Pl_{\hat{f}_{\operatorname{MOM},K}} - \operatorname{MOM}_K(\ell_{\hat{f}_{\operatorname{MOM},K}}) - (Pl_{f^*} - \operatorname{MOM}_K(\ell_{f^*})) . \quad (7)$$

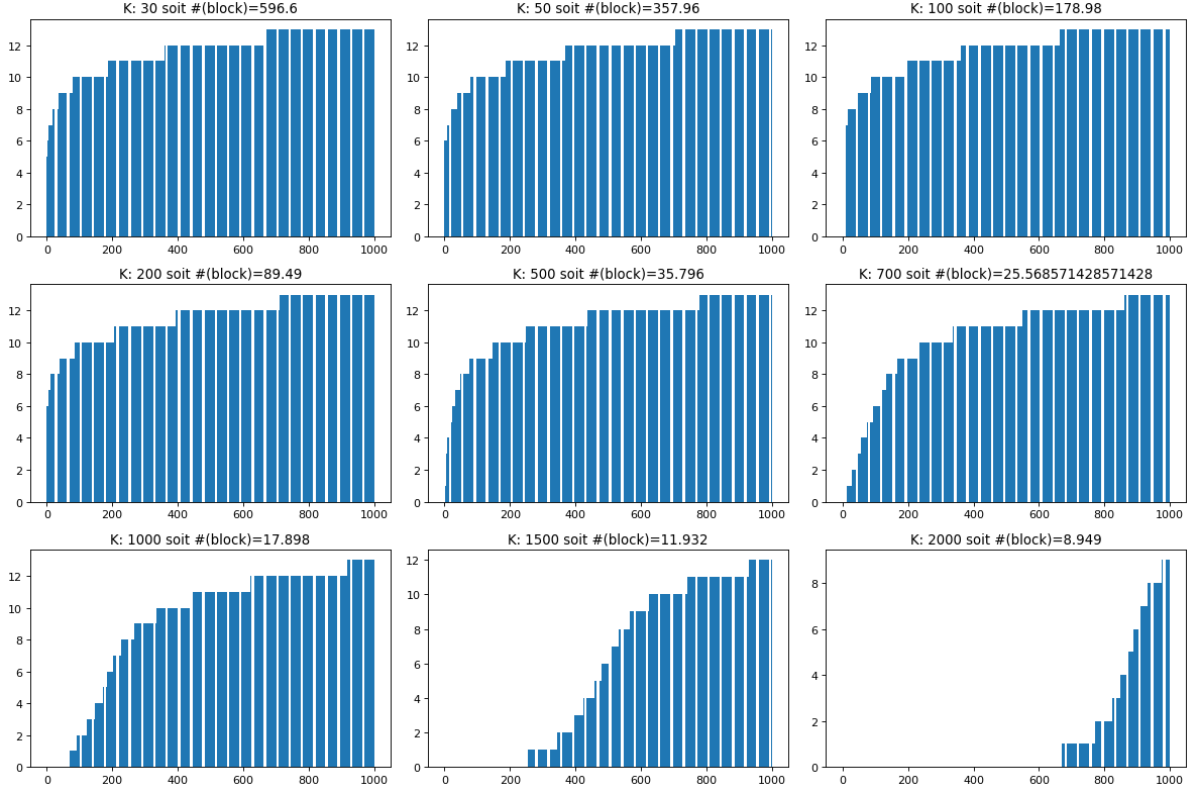


Figure 10: Sorted Histogram on the score (number of times a data is selected in a median block) of each points in a Logistic Regression MOM algorithm on the pulsar dataset for various values of K and $T = 20 \times K$ (only the 1000 smaller counts among the 17898 sample of the pulsar dataset are represented).

Let us now control the two expressions in the right-hand side of (7). Let $x > 0$. We have

$$\mathbb{P}[P\ell_{f^*} - \text{MOM}_K(\ell_{f^*}) > x] = \mathbb{P}\left[\sum_{k=1}^K I(P\ell_{f^*} - P_{B_k}\ell_{f^*} > x) \geq \frac{K}{2}\right] = \sum_{k=K/2}^K \binom{K}{k} p^k (1-p)^{K-k} \leq p^{K/2} 2^K$$

where $p = \mathbb{P}[P\ell_{f^*} - P_{B_k}\ell_{f^*} > x]$. Using Markov inequality together with $\text{var}(\ell_{f^*}) \leq 2L^2 \mathbb{E}(f^*(X))^2 \leq 2L^2\theta^2$, we obtain

$$\mathbb{P}[P\ell_{f^*} - \text{MOM}_K(\ell_{f^*}) > x] \leq \left(\frac{4\text{var}(\ell_{f^*})K}{Nx^2}\right)^{K/2} \leq \left(\frac{8L^2\theta^2K}{Nx^2}\right)^{K/2} = \exp(-K/2)$$

when $x = 2L\theta\sqrt{2eK/N}$.

Now, for any $x > 0$, one has $\sup_{f \in F} \text{MOM}_K(P\ell_f - \ell_f) > x$ iff

$$\sup_{f \in F} \sum_{k=1}^K I\{(P - P_{B_k})\ell_f > x\} \geq \frac{K}{2}. \quad (8)$$

Let us now control the probability that (8) holds via an adaptation of the small ball method [41, 55]. Let $x > 0$ and let $\phi(t) = (t-1)I\{1 \leq t \leq 2\} + I\{t \geq 2\}$ be defined for all $t \in \mathbb{R}$. As $\phi(t) \geq I\{t \geq 2\}$, one

has

$$\begin{aligned} & \sup_{f \in F} \sum_{k=1}^K I \{(P - P_{B_k})\ell_f > x\} \\ & \leq \sup_{f \in F} \sum_{k \in \mathcal{K}} \mathbb{E}[\phi(2(P - P_{B_k})\ell_f/x)] + |\mathcal{O}| + \sup_{f \in F} \sum_{k \in \mathcal{K}} (\phi(2(P - P_{B_k})\ell_f/x) - \mathbb{E}[\phi(2(P - P_{B_k})\ell_f/x)]) \end{aligned}$$

where we recall that $\mathcal{K} = \{k \in \{1, \dots, K\} : B_k \cap \mathcal{O} = \emptyset\}$.

Since, $\phi(t) \leq I\{t \geq 1\}$ and for all $f \in F$, $\text{Var}(\ell_f) \leq 2L^2\mathbb{E}f(X)^2 \leq 2L^2\theta_2^2$, we have for all $f \in F$ and $k \in \mathcal{K}$,

$$\mathbb{E}[\phi(2(P - P_{B_k})\ell_f/x)] \leq \mathbb{P}\left((P - P_{B_k})\ell_f \geq \frac{x}{2}\right) \leq \frac{4\text{Var}(\ell_f)}{x^2|B_k|} \leq \frac{8L^2\theta_2^2K}{x^2N}.$$

One has therefore

$$\begin{aligned} & \sup_{f \in F} \sum_{k=1}^K I \{(P - P_{B_k})\ell_f > x\} \\ & \leq K \left(\frac{8L^2\theta_2^2K}{x^2N} + \frac{|\mathcal{O}|}{K} + \sup_{f \in F} \frac{1}{K} \sum_{k \in \mathcal{K}} \left(\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right) - \mathbb{E}\left[\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right)\right] \right) \right). \end{aligned}$$

As $0 \leq \phi(\cdot) \leq 1$, by the bounded-difference inequality, for any $y > 0$, with probability larger than $1 - e^{-2y^2K}$,

$$\begin{aligned} & \sup_{f \in F} \frac{1}{K} \sum_{k \in \mathcal{K}} \left(\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right) - \mathbb{E}\left[\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right)\right] \right) \\ & \leq \mathbb{E} \left[\sup_{f \in F} \frac{1}{K} \sum_{k \in \mathcal{K}} \left(\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right) - \mathbb{E}\left[\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right)\right] \right) \right] + y. \end{aligned}$$

Now, by the symmetrization inequality,

$$\begin{aligned} & \mathbb{E} \left[\sup_{f \in F} \frac{1}{K} \sum_{k \in \mathcal{K}} \left(\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right) - \mathbb{E}\left[\phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right)\right] \right) \right] \\ & \leq 2\mathbb{E} \left[\sup_{f \in F} \frac{1}{K} \sum_{k \in \mathcal{K}} \epsilon_k \phi\left(\frac{2(P - P_{B_k})\ell_f}{x}\right) \right]. \end{aligned}$$

Since ϕ is 1-Lipschitz and $\phi(0) = 0$, by the contraction principle (see [45, Chapter 4] or more precisely equation (2.1) in [40]),

$$\mathbb{E} \left[\sup_{f \in F} \frac{1}{K} \sum_{k \in \mathcal{K}} \epsilon_k \phi\left(\frac{(P - P_{B_k})\ell_f}{x}\right) \right] \leq \mathbb{E} \left[\sup_{f \in F} \frac{1}{xK} \sum_{k \in \mathcal{K}} \epsilon_k (P - P_{B_k})\ell_f \right].$$

By the symmetrization principle,

$$\mathbb{E} \left[\sup_{f \in F} \frac{2}{xK} \sum_{k \in \mathcal{K}} \epsilon_k (P - P_{B_k})\ell_f \right] \leq \frac{2}{xN} \mathbb{E} \left[\sup_{f \in F} \sum_{i \in \mathcal{J}} \epsilon_i \ell_f(X_i, Y_i) \right].$$

Finally, since ℓ is L -Lipschitz, by the contraction principle (see equation (2.1) in [40]),

$$\mathbb{E} \left[\sup_{f \in F} \sum_{i \in \mathcal{J}} \epsilon_i \ell_f(X_i, Y_i) \right] \leq 2L\mathcal{R}(F).$$

Thus, for any $y > 0$, with probability larger than $1 - \exp(-2y^2K)$,

$$\sup_{f \in F} \sum_{k=1}^K I \{(P - P_{B_k})\ell_f > x\} \leq K \left(\frac{8L^2\theta_2^2K}{x^2N} + \frac{|\mathcal{O}|}{K} + y + \frac{4L\mathcal{R}(F)}{xN} \right).$$

Let $\Delta = 1/4 - |\mathcal{O}|/K$ and let $y = \Delta$ and $x = 8L \max\left(\theta_2\sqrt{K/N}, 4\mathcal{R}(F)/N\right)$ so

$$\mathbb{P} \left(\sup_{f \in F} \sum_{k=1}^K I \{(P - P_{B_k})\ell_f > x\} < \frac{K}{2} \right) \geq 1 - e^{-\Delta^2K/8}.$$

Going back to (8), this means that

$$\mathbb{P} \left(\sup_{f \in F} \text{MOM}_K(\ell_f - P\ell_f) \leq 4L \max \left(\theta_2\sqrt{\frac{K}{N}}, \frac{4\mathcal{R}(F)}{N} \right) \right) \geq 1 - \exp(-2\Delta^2K).$$

Plugging this result in (7) concludes the proof of the theorem.

6.3 Proof of Theorem 3

First, hypothesis **H-6** implies that, for any $t \geq 1$, there exists an open convex neighbor B of u_{t-1} such that for all $u \in B$, $P_{B_{k_{med}}(\sigma_{t-1})}(\ell_{f_u}) \in \text{MOM}_K(\ell_{f_u})$. In particular therefore, for any $u \in B$,

$$\frac{K}{N} \sum_{b \in B_{k_{med}}(\sigma_{t-1})} \nabla_u \ell_{f_u}(X_b, Y_b) = \nabla_u \text{MOM}_K(f_u). \quad (9)$$

From here, the rest of the proof is identical to the proof of consistency of the SGD. As a proof of concept, we reproduce here the arguments of [8]. In what follows, all results are given conditionally on \mathcal{D}_N . We set

$$G_{t-1} = \nabla_u \text{MOM}_K(f_u)|_{u=u_{t-1}} = \frac{K}{N} \sum_{i \in B_{k_{med}}(\sigma_{t-1})} \nabla_u \ell_{f_u}(X_i, Y_i)|_{u=u_{t-1}}.$$

By definition of the update of the algorithm, we have

$$\begin{aligned} \|u_t - u_{min}\|_2^2 &= \|u_{t-1} - u_{min} - \eta_{t-1}G_{t-1}\|_2^2 \\ &= \|u_{t-1} - u_{min}\|_2^2 - 2\eta_{t-1}\langle u_{t-1} - u_{min}, G_{t-1} \rangle + \eta_{t-1}^2 \|G_{t-1}\|_2^2. \end{aligned}$$

Let F_t be the σ -algebra spanned by $\sigma_1, \dots, \sigma_{t-2}$, we have

$$\mathbb{E} \left[\|u_t - u_{min}\|_2^2 | F_t \right] = \|u_{t-1} - u_{min}\|_2^2 - 2\eta_{t-1}\langle u_{t-1} - u_{min}, \mathbb{E}[G_{t-1} | F_t] \rangle + \eta_{t-1}^2 \mathbb{E} \left[\|G_{t-1}\|_2^2 | F_t \right].$$

Then, we use Jensen's inequality, we have

$$\begin{aligned} \mathbb{E} \left[\|u_t - u_{min}\|_2^2 | F_t \right] &\leq \|u_{t-1} - u_{min}\|_2^2 - 2\eta_{t-1} \frac{K}{N} \sum_{i \in B_{k_{med}}(\sigma_{t-1})} \langle u_{t-1} - u_{min}, \mathbb{E} \left[\nabla_{u_{t-1}} \ell_{f_{u_{t-1}}}(X_i, Y_i) | F_t \right] \rangle \\ &\quad + \eta_{t-1}^2 \frac{K}{N} \sum_{i \in B_{k_{med}}(\sigma_{t-1})} \mathbb{E} \left[\left\| \nabla_{u_{t-1}} \ell_{f_{u_{t-1}}}(X_i, Y_i) \right\|_2^2 | F_t \right]. \quad (10) \end{aligned}$$

By convexity of $u \mapsto \ell_{f_u}(x, y)$, for any x, y , we have

$$\langle u_{t-1} - u_{min}, \nabla_{u_{t-1}} \ell_{f_{u_{t-1}}}(X_i, Y_i) \rangle \geq \ell_{f_{u_{t-1}}}(X_i, Y_i) - \ell_{f_{u_{min}}}(X_i, Y_i)$$

Therefore,

$$\frac{K}{N} \sum_{i \in B_{k_{med}}(\sigma_{t-1})} \langle u_{t-1} - u_{min}, \mathbb{E} \left[\nabla_{u_{t-1}} \ell_{f_{u_{t-1}}}(X_i, Y_i) \middle| F_t \right] \rangle \geq \mathbb{E} \left[\text{MOM}_K(\ell_{f_{u_{t-1}}}) - P_{B_{k_{med}}(\sigma_{t-1})} \ell_{f_{u_{min}}} \middle| F_t \right]$$

Hence, using hypothesis **H-1** and **H-5**, we get

$$\mathbb{E} \left[\|u_t - u_{min}\|_2^2 \middle| F_t \right] \leq \|u_{t-1} - u_{min}\|_2^2 + \eta_{t-1}^2 L^2. \quad (11)$$

Now, defining $h(t) = \|u_t - u_{min}\|_2^2$ and $\delta_t = \mathbb{1}\{\mathbb{E}[h(t) - h(t-1) | F_t] > 0\}$, it follows from (11) that

$$\mathbb{E}[\delta(t)(h(t) - h(t-1))] = \mathbb{E}[\delta(t)\mathbb{E}[h(t) - h(t-1) | F_t]] \leq \eta_t^2 L^2.$$

Using hypothesis **H-3**, we have that

$$\sum_{t \geq 1} \mathbb{E}[\delta(t)(h(t) - h(t-1))] < \infty.$$

From Doob's first martingale convergence theorem (see for example [22, Theorem 5.2.8]), there exists $h_\infty \in \mathbb{R}_+$ such that

$$h(t) \xrightarrow[t \rightarrow \infty]{a.s.} h_\infty.$$

We plug this convergence in the inequality (10) to get

$$\begin{aligned} 2\eta_t(u_{t-1} - u_{min})^T \frac{K}{n} \mathbb{E} \left[\sum_{i \in B_{k_{med}}(\sigma_{t-1})} \nabla_{u_{t-1}} \ell_{f_{u_{t-1}}}(X_i, Y_i) \middle| F_t \right] &\leq h(t-1) - \mathbb{E}[h(t) | F_t] + \eta_t^2 L^2 \\ &\leq h(t-1) + \eta_t^2 L^2. \end{aligned}$$

Then, using hypothesis **H-3**, we have

$$\sum_{t=1}^{\infty} \eta_t(u_{t-1} - u_{min})^T \frac{K}{N} \mathbb{E} \left[\sum_{i \in B_{k_{med}}(\sigma_{t-1})} \nabla_{u_{t-1}} \ell_{f_{u_{t-1}}}(X_i, Y_i) \middle| F_t \right] < \infty.$$

From the hypotheses **H-4** and **H-5**, we can then conclude that

$$\|u_{t-1} - u_{min}\| \xrightarrow[t \rightarrow \infty]{a.s.} 0.$$

6.4 Proof of Proposition 1

We denote by $B(1)(f), \dots, B(K)(f)$ the blocks such that the corresponding block empirical mean $P_{B(i)(f)}(f(X_1^n))$ are sorted: $P_{B(1)(f)}(f(X_1^n)) \geq \dots \geq P_{B(K)(f)}(f(X_1^n))$. To simplify the notations we denote by $K/2 = \lceil K/2 \rceil$ that correspond to the rank of the median when K is odd, we suppose K odd.

Let us take $\epsilon \in \mathbb{R}^N$, we show that for all ϵ with $\|\epsilon\|$ sufficiently small, we have $B_{(K/2)}(p) = B_{(K/2)}(p+t\epsilon)$ for all $t \in [0, 1]$.

To show that, it is sufficient to check that we have

$$\forall 1 \leq k \leq K-1, \forall t \in [0, 1], \quad P_{B(k)(p)}(f_{p+t\epsilon}) - P_{B(k+1)(p)}(f_{p+t\epsilon}) \geq 0.$$

We decompose this difference in three parts,

$$\begin{aligned} P_{B(k)(p)}(f_{p+t\epsilon}) - P_{B(k+1)(p)}(f_{p+t\epsilon}) &\geq P_{B(k)(p)}(f_p) - P_{B(k+1)(p)}(f_p) \\ &\quad - \left| P_{B(k)(p)}(f_p) - P_{B(k)(p)}(f_{p+t\epsilon}) \right| \\ &\quad - \left| P_{B(k+1)(p)}(f_{p+t\epsilon}) - P_{B(k+1)(p)}(f_p) \right| \end{aligned}$$

The two last terms are controlled by the Lipschitz property of $p \mapsto f_p$,

$$\forall t \in [0, 1], \quad P_{B(k)(p)}(f_{p+t\varepsilon}) - P_{B(k+1)(p)}(f_{p+t\varepsilon}) \geq P_{B(k)(p)}(f_p) - P_{B(k+1)(p)}(f_p) - 2tL\|\varepsilon\|_2.$$

Then, we use order statistics and interquantile range to control the first term.

To do that, we use Renyi representation to say that the order statistics of the blocks can be expressed as functions of the order statistics of exponential variables. Let E_1, \dots, E_K be K i.i.d sample with a standard Exponential law, then using the function $U(t) = F_p^{\leftarrow}(1 - 1/t)$ where F_p is the c.d.f of the random variable $P_{B_k}(f_p)$ to make the link between the sample $P_{B(1)(p)}, \dots, P_{B(K)(p)}$ and the sample E_1, \dots, E_K , we have

$$\forall 1 \leq k \leq K, \quad P_{B(k)(p)}(f_p) \sim U \left(\exp \left(\sum_{i=k}^K \frac{E(i)}{i} \right) \right),$$

Then, because the hazard function is the inverse of the derivative of U and is bounded by $1/m_{N,K}$, we have by the mean value inequality,

$$U \left(e^{\sum_{i=k}^K \frac{E(i)}{i}} \right) - U \left(e^{\sum_{i=k+1}^K \frac{E(i)}{i}} \right) \geq \frac{m_{N,K}E}{k}$$

hence, for all $\varepsilon' > 0$

$$\begin{aligned} \mathbb{P} \left(P_{B(k)(p)}(f_p) - P_{B(k+1)(p)}(f_p) \leq \varepsilon' \right) &\leq \mathbb{P} \left(\frac{m_{N,K}E}{k} \leq \varepsilon' \right) \\ &= 1 - e^{-\frac{k\varepsilon'}{m_{N,K}}}. \end{aligned}$$

which we can rewrite as: for all $\delta \in (0, 1)$, with probability greater than δ ,

$$P_{B(k)(p)}(f_p) - P_{B(k+1)(p)}(f_p) \geq \frac{m \ln(1/\delta)}{k}.$$

Finally, with probability greater than δ ,

$$\forall t \in [0, 1], \quad P_{B(k)(p)}(f_{p+t\varepsilon}) - P_{B(k+1)(p)}(f_{p+t\varepsilon}) \geq \frac{m_{N,K} \ln(1/\delta)}{k} - 2Lt\|\varepsilon\|_2.$$

Which implies that with probability greater than $\exp(-ktL\|\varepsilon\|_2/(2m_{N,K}))$,

$$\forall t \in [0, 1], \quad P_{B(k)(p)}(f_{p+t\varepsilon}) - P_{B(k+1)(p)}(f_{p+t\varepsilon}) \geq 0.$$

Then, taking the union bound for $1 \leq k \leq K - 1$, we have

$$\mathbb{P} \left(\forall t \in [0, 1], \quad B_{(K/2)(p)} = B_{(K/2)(p+t\varepsilon)} \right) \geq 1 - \sum_{k=1}^{K-1} \exp(-ktL\|\varepsilon\|_2/(2m_{N,K}))$$

We want to compute the partial derivatives of $\psi_{f(p)}$, to do that we denote by $e_1, \dots, e_n \in \mathbb{R}^d$ the canonical basis of \mathbb{R}^d and we define $\varepsilon_n^j = \frac{e_j}{2^n} m_{N,K}$. We denote by A_n the event

$$A_n^j = \left\{ \forall t \in [0, 1], \quad \forall n' \geq n, \quad B_{(K/2)(p)} = B_{(K/2)(p+t\varepsilon_{n'}^j)} \right\}$$

and we study the limiting event $\Omega^j = \overline{\lim}_{n \rightarrow \infty} A_n^j$.

First, let us note that for all $1 \leq j \leq d$, the sequence of set $(A_n^j)_n$ is non-increasing, hence

$$\Omega^j = \overline{\lim}_{n \rightarrow \infty} A_n^j = \underline{\lim}_{n \rightarrow \infty} A_n^j = (\overline{\lim}_{n \rightarrow \infty} (A_n^j)^c)^c,$$

then, for all $1 \leq j \leq d$, we can study the $\overline{\lim}_{n \rightarrow \infty} (A_n^j)^c$ with Borel-Cantelli Lemma. Indeed, we have

$$\begin{aligned} \mathbb{P}((A_n^j)^c) &\leq \sum_{k=1}^K 1 - \exp(-ktL \|\varepsilon_n^j\|_2 / (2m_{N,K})) \\ &\leq \sum_{k=1}^{K-1} \frac{ktL \|\varepsilon_n^j\|_2}{2m_{N,K}} \\ &= \frac{1}{2^n} \frac{K(K-1)tL}{4} \end{aligned}$$

Hence, the series $\sum_n \mathbb{P}((A_n^j)^c)$ converges and by Borel Cantelli Lemma, $\mathbb{P}(\overline{\lim}_{n \rightarrow \infty} (A_n^j)^c) = 0$, then for all $1 \leq i \leq d$, $\mathbb{P}(\Omega^j) = 1$.

Rewriting that we have that

$$\forall \omega \in \Omega^j, \quad \exists N \geq 1 : \quad \forall n \geq N, \quad \omega \in A_n^j.$$

Hence,

$$\exists N \geq 1 : \quad \forall n \geq N, \quad \forall t \in [0, 1], \quad B_{(K/2)}(p) = B_{(K/2)}(p + t\varepsilon_n^j),$$

Hence,

$$\exists N \geq 1 : \quad \forall n \geq N, \quad \forall t \in [0, 1], \quad B_{(K/2)}(p) = B_{(K/2)}(p + t\varepsilon_n^j),$$

which implies that for all $1 \leq j \leq d$, $\forall n \geq N$,

$$\begin{aligned} \partial_j \psi_{f_p}(X) &= \lim_{t \rightarrow 0} \frac{\psi_{f_{p+t\varepsilon_n^j}}(X) - \psi_{f_p}(X)}{t} \\ &= \lim_{t \rightarrow 0} \frac{P_{B(K/2)(p)}(f_{p+t\varepsilon_n^j}) - P_{B(K/2)(p)}(f_p)}{t} \\ &= \frac{1}{n/K} \lim_{t \rightarrow 0} \sum_{i \in B(K/2)(p)} \frac{f_{p+t\varepsilon_n^j}(X_i) - f_p(X_i)}{t} \\ &= \frac{1}{n/K} \sum_{i \in B(K/2)(p)} \partial_j f_p(X_i) \end{aligned}$$

7 Annex

7.1 Choice of the number of blocks

Let us study the behaviour of our algorithms when the number of blocks changes. We plot the accuracy as a function of K averaged on 50 runs to have a good idea of the evolution of the performance with respect to K , the result is represented in figure 11.

There is a clear separation around $2|\mathcal{O}| = 60$ that is consistent with the theory. On the other hand the accuracy doesn't decrease when K gets bigger one would expect. This may be due to the symmetry of the dataset. If we run the same experiment on the real dataset, we get a much more regular plot, see Figure 12.

Figure 12 confirms our predictions on clean datasets, the accuracy getting better as K gets smaller (the MOM minimizer is the ERM when $K = 1$ and ERM is optimal in the i.i.d. setup, [44]). This may be due to the small number of outliers in this dataset.




Figure 11: Plot of the accuracy on the toy dataset of Logistic Regression MOM as a function of K .

7.2 Illustration of convergence rate

In this section, we estimate the rate of convergence of the MOM risk minimization algorithm Logistic Regression on two databases (see figure 13). The first dataset is composed of points located on two interlaced half-circle with a Gaussian noise of standard deviation 0.3, the two "moons" are each of a different class. We assume that these moons don't satisfy the margin property (we checked that the rate was slow for ERM algorithms, using the vanilla logistic regression). The second dataset is composed of two Gaussians $\mathcal{N}((-1, -1), 1.4^2 I_2)$ and $\mathcal{N}((1, 1), 1.4^2 I_2)$ with respective label 1 and 0, we can prove that this dataset verifies the margin property needed to obtain fast rate in ERM

There are no outliers in the datasets because we only want to test the rate of convergence. To illustrate the rates of convergence of our algorithms, we plot the curve $\log \left(\left| \hat{R}^{0-1}(\hat{f}_K) - \hat{R}^{0-1}(f^*) \right| \right)$ as a function of $\log(n)$ where the risk is estimated by Monte-Carlo. The figure obtained for Logistic Regression MOM is represented in figure 14. It seems that MOM minimizers can achieve fast rates of convergence even if we did not prove them.

Remark 6. *We used random blocks sampled at each iteration for this application because it is the algorithm that we described earlier but even if we use one partition of blocks for the whole algorithm (as in the theory we developed) we obtain nonetheless fast rate for the Gaussians dataset.*

7.3 Comparison with robust algorithms based on M-estimators.

In this section we compare the algorithm Logistic Regression MOM with two other algorithms based on M-estimators, these algorithms are studied on the toy dataset presented in Section 5.

One algorithm is a gradient on the Huber estimation of the loss function, it follows the same reasoning as MOM risk minimization and minimizes $\mathbb{E}[l_f(X, Y)]$ using as a proxy the Huber estimator for this quantity.




Figure 12: Plot of the accuracy on HTRU2 dataset of Logistic Regression MOM as a function of K .




Figure 13: Scatter plot of the two dataset used in this section, the color represent the class of the points.

The Huber estimator is then defined as a M -estimator, denoted here $\hat{\mu}_f$, solution of

$$\sum_{i=1}^n \psi_c(\hat{\mu}_f - l(f(X_i), Y_i)) = 0$$

where $\psi_c = \max(-c, \min(c, x))$ is the Huber function, $c > 0$. Using this definition of $\hat{\mu}_f$, it is then easy to compute the gradient $\nabla \hat{\mu}_f$ and then use a gradient descent algorithm. The theory behind this algorithm is studied further in [13].

The second algorithm uses a "redescending" loss function, in short we do ERM with a bounded loss function. Here we use Tukey biweight loss function rescaled by MADN scale estimator and IRLS algorithm to optimize the empirical risk.

Figure 15 shows that all algorithms perform similarly on this easy, low dimensional dataset. The situation is quite different in higher dimension. In Figure 16 we used a 200 dimensional dataset and the algorithm using a redescending loss function does not perform well. This may be due to local minima in which the algorithm gets stuck, as local minima are multiplied when the dimension gets higher. The other algorithms don't suffer this drawback since they use a "projection by the loss function" that makes the problem one dimensional.

The algorithm using redescending loss functions is a simple gradient descent that has linear complexity. The Huber gradient algorithm estimates at each iteration a Huber estimator of location. The complexity

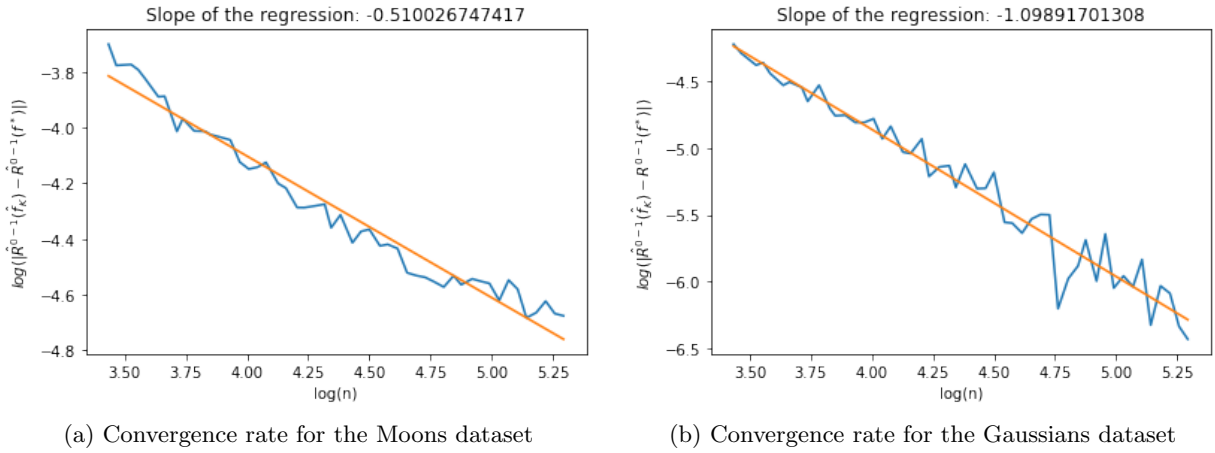


Figure 14: Plot of the logarithm of the excess risk as a function of $\log(n)$ in two cases: (a) where the margin assumption does not hold and (b) where the margin assumption holds. A linear regression is fitted on the curve, its slope is printed at the top of each figure revealing a slow $n^{-0.51}$ rate of convergence in case (a) and a fast $n^{-1.1}$ in case (b).

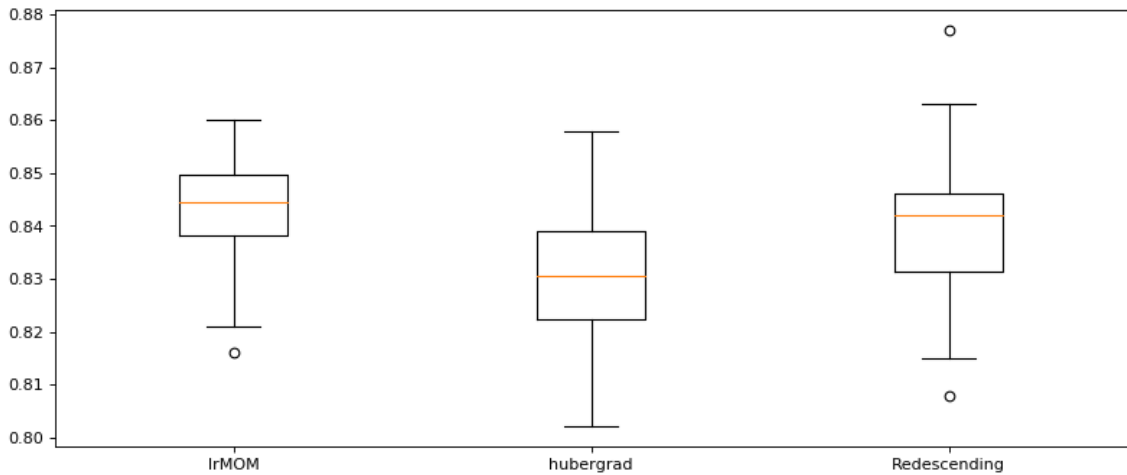


Figure 15: Boxplot of the accuracy obtained on 50 training/test run (1000 training sample, 2% corruption) of each algorithms on a 2-dimensional toy dataset.




Figure 16: Boxplot of the accuracy obtained on 50 training/test run (2000 training sample, 2% corruption) of each algorithms on a 200-dimensional toy dataset.

of this estimator depends on the algorithm used but for most M-estimators a commonly used algorithm is an iteratively reweighted algorithm whose complexity is linear in the sample size. In practice we can nonetheless notice a great complexity of the Huber estimator in some cases where data are not well spread. In most cases, Logistic Regression MOM is the fastest among these three algorithms and the gradient Huber is the slowest, even though logistic regression may need a lot more iterations than the other algorithms.

References

- [1] Charu C. Aggarwal. *Outlier Analysis*. Springer Publishing Company, Incorporated, 2013.
- [2] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *J. Comput. System Sci.*, 58(1, part 2):137–147, 1999. Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996).
- [3] Sanjeev Arora and Boaz Barak. *Computational complexity*. Cambridge University Press, Cambridge, 2009. A modern approach.
- [4] Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Structured sparsity through convex optimization. *Statist. Sci.*, 27(4):450–468, 2012.
- [5] Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: risk bounds and structural results. *J. Mach. Learn. Res.*, 3(Spec. Issue Comput. Learn. Theory):463–482, 2002.
- [6] Peter L. Bartlett and Shahar Mendelson. Empirical minimization. *Probab. Theory Related Fields*, 135(3):311–334, 2006.
- [7] Derya Birant and Alp Kut. St-dbscan: An algorithm for clustering spatial-temporal data. *Data & Knowledge Engineering*, 60(1):208–221, 2007.
- [8] Léon Bottou. On-line learning in neural networks. chapter On-line Learning and Stochastic Approximations, pages 9–42. Cambridge University Press, New York, NY, USA, 1998.
- [9] Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: a survey of some recent advances. *ESAIM Probab. Stat.*, 9:323–375, 2005.
- [10] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. Concentration inequalities: A nonasymptotic theory of independence, 2013.
- [11] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004.
- [12] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [13] Christian Brownlees, Emilien Joly, and Gábor Lugosi. Empirical risk minimization for heavy-tailed losses. *Ann. Statist.*, 43(6):2507–2536, 12 2015.
- [14] S. Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015.

- [15] Peter Bühlmann and Bin Yu. Analyzing bagging. *The Annals of Statistics*, 30(4):927–961, 2002.
- [16] Mengjie Chen, Chao Gao, and Zhao Ren. Robust covariance and scatter matrix estimation under huber’s contamination model. Technical report, University of Chicago and University of Pittsburgh, 2017. Preprint available on arXiv:1506.00691.
- [17] Christophe CROUX and Catherine DEHON. Robust linear discriminant analysis using s-estimators. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 29(3):473–493, 2001.
- [18] Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*, volume 31 of *Applications of Mathematics*. Springer, corrected 2nd edition, 1997. missing.
- [19] Luc Devroye, Matthieu Lerasle, Gabor Lugosi, and Roberto I. Oliveira. Sub-Gaussian mean estimators. *Ann. Statist.*, 44(6):2695–2725, 2016.
- [20] Luc Devroye, Matthieu Lerasle, Gabor Lugosi, and Roberto I. Oliveira. Sub-gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016.
- [21] D. Donoho and A. Montanari. Variance breakdown of huber (m)-estimators: $n/p \rightarrow m \in (1, +\infty)$. Technical report, Stanford University, 2015. Preprint available on arXiv:1503.02106.
- [22] Rick Durrett. *Probability: Theory and examples*, 2011.
- [23] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM J. Comput.*, 41(6):1558–1590, 2012.
- [24] François Le Gall. Powers of tensors and fast matrix multiplication. *CoRR*, abs/1401.7714, 2014.
- [25] Chao Gao. Robust regression via multivariate regression depth. Technical report, University of Chicago, 2017. Preprint available on arXiv:1702.04656.
- [26] Venkatesan Guruswami and Prasad Raghavendra. Hardness of learning halfspaces with noise. *SIAM J. Comput.*, 39(2):742–765, 2009.
- [27] Frank R Hampel. A general qualitative definition of robustness. *The Annals of Mathematical Statistics*, pages 1887–1896, 1971.
- [28] Frank R Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974.
- [29] Qiyang Han and Jon A. Wellner. A sharp multiplier inequality with applications to heavy-tailed regression problems. *arXiv:1706.02410*, 2017.
- [30] Xuming He and Wing K Fung. High breakdown estimation for multiple populations with applications to discriminant analysis. *Journal of Multivariate Analysis*, 72(2):151 – 162, 2000.
- [31] Charles AR Hoare. Quicksort. *The Computer Journal*, 5(1):10–16, 1962.
- [32] Peter J. Huber. Robust estimation of a location parameter. *Ann. Math. Statist.*, 35:73–101, 1964.
- [33] Peter J Huber. The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 221–233. Berkeley, CA, 1967.
- [34] Peter J Huber and Elvezio M Ronchetti. *Robust statistics; 2nd ed.* Wiley Series in Probability and Statistics. Wiley, Hoboken, NJ, 2009.
- [35] Mia Hubert and Katrien Van Driessen. Fast and robust discriminant analysis. *Computational Statistics & Data Analysis*, 45(2):301 – 320, 2004.
- [36] Mia Hubert and Stephan Van Der Veeken. Robust classification for skewed data. *Adv. Data Anal. Classif.*, 4(4):239–254, December 2010.
- [37] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.*, 43(2-3):169–188, 1986.
- [38] Michael I. Jordan. On statistics, computation and scalability. *Bernoulli*, 19(4):1378–1390, 2013.
- [39] Vladimir Koltchinskii. *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].
- [40] Vladimir Koltchinskii. Oracle inequalities in empirical risk minimization and sparse recovery problems. 2011.
- [41] Vladimir Koltchinskii and Shahar Mendelson. Bounding the smallest singular value of a random matrix without concentration. *Int. Math. Res. Not. IMRN*, (23):12991–13008, 2015.
- [42] G. Lecué and M. Lerasle. Learning from mom’s principle : Le cam’s approach. Technical report, CNRS, ENSAE, Paris-sud. Preprint available on arXiv:1701.01961.
- [43] G. Lecué and M. Lerasle. Robust machine learning by median-of-means : theory and practice. Technical report, CNRS, ENSAE, Paris-sud. Preprint available on arXiv:1711.10306.

- [44] Guillaume Lecué and Shahar Mendelson. Learning subgaussian classes: Upper and minimax bounds. Technical report, CNRS, Ecole polytechnique and Technion, 2013.
- [45] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*. Classics in Mathematics. Springer-Verlag, Berlin, 2011. Isoperimetry and processes, Reprint of the 1991 edition.
- [46] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*, pages 413–422. IEEE, 2008.
- [47] Gabor Lugosi and Shahar Mendelson. Regularization, sparse recovery, and median-of-means tournaments. *Preprint available on arXiv:1701.04112*.
- [48] Gabor Lugosi and Shahar Mendelson. Sub-gaussian estimators of the mean of a random vector. *Preprint available on arXiv:1702.00482*.
- [49] Gabor Lugosi and Shahar Mendelson. Risk minimization by median-of-means tournaments. *Preprint available on ArXiv:1608.00757*, 2017.
- [50] R. J. Lyon, B. W. Stappers, S. Cooper, and J. D. Brooke, J. M. Knowles. Fifty Years of Pulsar Candidate Selection: From simple filters to a new principled real-time classification approach. *MNRAS*, 000:000–000, 2015.
- [51] Enno Mammen and Alexandre B. Tsybakov. Smooth discrimination analysis. *Ann. Statist.*, 27(6):1808–1829, 1999.
- [52] COLIN McDIARMID. Concentration for independent permutations. *Combinatorics, Probability and Computing*, 11(2):163–178, 2002.
- [53] Shahar Mendelson. An optimal unrestricted learning procedure. *Preprint available on arXiv:1707.05342*.
- [54] Shahar Mendelson. Learning without concentration. In *Proceedings of the 27th annual conference on Learning Theory COLT14*, pages pp 25–39. 2014.
- [55] Shahar Mendelson. Learning without concentration. *J. ACM*, 62(3):Art. 21, 25, 2015.
- [56] Stanislav Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4):2308–2335, 2015.
- [57] Eric Moulines and Francis R Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pages 451–459, 2011.
- [58] Ernest Fokoue Necla Gunduz. Robust classification of high dimension low sample size data. 2015.
- [59] A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization*. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1983. Translated from the Russian and with a preface by E. R. Dawson, Wiley-Interscience Series in Discrete Mathematics.
- [60] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [61] Frank Rosenblatt. *The perceptron: A theory of statistical separability in cognitive systems*. Cornell Aeronautical Laboratory, Inc., Rep. No. VG-1196-G-1. U.S. Department of Commerce, Office of Technical Services, PB 151247, 1958.
- [62] Volker Roth. Probabilistic discriminative kernel classifiers for multi-class problems, 2001.
- [63] Adrien Saumard. On optimality of empirical risk minimization in linear aggregation. *Bernoulli*, 24(3):2176–2203, 2018.
- [64] John W Tukey. A survey of sampling from contaminated distributions. *Contributions to probability and statistics*, 2:448–485, 1960.
- [65] John W Tukey. The future of data analysis. *The annals of mathematical statistics*, 33(1):1–67, 1962.
- [66] Vladimir N. Vapnik. *Statistical learning theory*. Adaptive and Learning Systems for Signal Processing, Communications, and Control. John Wiley & Sons, Inc., New York, 1998. A Wiley-Interscience Publication.
- [67] Vladimir N. Vapnik. *The nature of statistical learning theory*. Statistics for Engineering and Information Science. Springer-Verlag, New York, second edition, 2000.