

Regularization, sparse recovery, and median-of-means tournaments *

Gábor Lugosi^{†‡§} Shahar Mendelson[¶]

April 22, 2018

Abstract

We introduce a regularized risk minimization procedure for regression function estimation. The procedure is based on median-of-means tournaments, introduced by the authors in [11] and achieves near optimal accuracy and confidence under general conditions, including heavy-tailed predictor and response variables. It outperforms standard regularized empirical risk minimization procedures such as LASSO or SLOPE in heavy-tailed problems.

2010 Mathematics Subject Classification: 62J02, 62G08, 60G25.

1 Introduction

1.1 Empirical risk minimization, regularization

Regression function estimation is a fundamental problem in statistics and machine learning. In the most standard formulation of the problem, (X, Y) is a pair of random variables in which X , taking values in some general measurable space \mathcal{X} , represents the observation (or feature vector) and one would like to approximate the unknown real value Y by a function of X . In other words, one is interested in finding a function $f : \mathcal{X} \rightarrow \mathbb{R}$ such that $f(X)$ is “close” to Y . As the vast majority of the literature, we measure the quality of f by the *risk*

$$R(f) = \mathbb{E}(f(X) - Y)^2 ,$$

which is well defined whenever $f(X)$ and Y are square integrable, assumed throughout the paper. Clearly, the best possible function is the *regression function* $m(X) = \mathbb{E}(Y|X)$.

However, in statistical problems, the joint distribution of (X, Y) is unknown and the regression function is impossible to compute. Instead, a sample $\mathcal{D}_N = ((X_1, Y_1), \dots, (X_N, Y_N))$ of independent copies of the pair (X, Y) is available (such that \mathcal{D}_N and the pair (X, Y) are independent).

*Gábor Lugosi was supported by the Spanish Ministry of Economy and Competitiveness, Grant MTM2015-67304-P and FEDER, EU. Shahar Mendelson was supported in part by the Israel Science Foundation.

[†]Department of Economics and Business, Pompeu Fabra University, Barcelona, Spain, gabor.lugosi@upf.edu

[‡]ICREA, Pg. Lluís Companys 23, 08010 Barcelona, Spain

[§]Barcelona Graduate School of Economics

[¶]Mathematical Sciences Institute, The Australian National University, and Department of Mathematics, Technion, I.I.T, shahar.mendelson@anu.edu.au

A popular and thoroughly studied approach is to select a function \widehat{f}_N from a fixed class \mathcal{F} of functions. Formally, a *learning procedure* is a map $\Phi : (\mathcal{X} \times \mathbb{R})^N \rightarrow \mathcal{F}$ that assigns to each sample $\mathcal{D}_N = (X_i, Y_i)_{i=1}^N$ a (random) function $\Phi(\mathcal{D}_N) = \widehat{f}_N$. If the class \mathcal{F} is sufficiently “large,” then it is reasonable to expect that the best function in the class

$$f^* = \operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}(f(X) - Y)^2$$

has an acceptable performance, and the standard assumption is that the minimum is attained and $f^* \in \mathcal{F}$ is unique. We assume that \mathcal{F} is a closed and convex subset of $L_2(\mu)$ —where μ denotes the distribution of X —, guaranteeing the existence and uniqueness of f^* .

The quality of a learning procedure is typically measured by the *mean squared error*, which is the conditional expectation

$$\|\widehat{f}_N - f^*\|_{L_2}^2 = \mathbb{E}((\widehat{f}_N(X) - f^*(X))^2 | \mathcal{D}_N),$$

where, for $q \geq 1$, we use the notation

$$\|f - g\|_{L_q} = (\mathbb{E}|f(X) - g(X)|^q)^{1/q} \quad \text{and also} \quad \|f - Y\|_{L_q} = (\mathbb{E}|f(X) - Y|^q)^{1/q}.$$

A closely related, though not equivalent, measure of performance is the *excess risk*, defined by the conditional expectation

$$R(\widehat{f}_N) - R(f^*) = \mathbb{E}((\widehat{f}_N(X) - Y)^2 | \mathcal{D}_N) - \mathbb{E}(f^*(X) - Y)^2.$$

The goal of a statistical learning problem is to find a learning procedure that achieves a good *accuracy* with a high *confidence*. In particular, for $r > 0$ and $\delta \in (0, 1)$, we say that a procedure performs with accuracy parameter r with confidence $1 - \delta$ in the class \mathcal{F} (e.g., for the mean squared error) if

$$\mathbb{P}\left(R(\widehat{f}_N) - R(f^*) \leq r^2\right) \geq 1 - \delta$$

(sometimes one only considers $\mathbb{P}\left(\|\widehat{f}_N - f^*\|_{L_2} \leq r\right) \geq 1 - \delta$).

High accuracy and high confidence (i.e., small r and small δ) in the given class are obviously conflicting requirements. The achievable tradeoff has been thoroughly studied and it is fairly well understood. We refer the reader to Lecué and Mendelson [8], Lugosi and Mendelson [11] for recent accounts.

The most standard approach for a learning procedure is *empirical risk minimization* (ERM), also known as *least squares regression* in which

$$\widehat{f}_N \in \operatorname{argmin}_{f \in \mathcal{F}} \sum_{i=1}^N (f(X_i) - Y_i)^2$$

(where we assume that the minimum is achieved). One may show (see, e.g., Lecué and Mendelson [8]) that unless the function class and target are sub-Gaussian¹, then empirical risk minimization is far from achieving the optimal accuracy/confidence tradeoff. The reason for

¹Here, Sub-Gaussian means that the ψ_2 and the L_2 norms are equivalent in $\mathcal{F} \cup \{0\}$; that is, there is a constant L such that for any $f, h \in \mathcal{F} \cup \{0\}$ and any $p \geq 2$, $\|f - h\|_{L_p} \leq L\sqrt{p}\|f - h\|_{L_2}$, and that the same holds for any $Y - f(X)$ for any $f \in \mathcal{F}$.

the suboptimal behaviour of empirical risk minimization is that outliers distort the empirical means unless the problem is very close to being Gaussian. Thankfully, learning procedures that can tackle heavy-tailed problems exist, as it was recently pointed out by Lugosi and Mendelson [11] with the introduction of the *median-of-means tournament*.

A common problem that all learning procedures encounter is that of overfitting, which occurs when the underlying class is too big relative to the (random) information at the learner’s disposal. A standard way of dealing with learning problems involving classes that are too large is giving priority to functions in the class according to some prior belief of “simplicity”. For example, in regularized risk minimization, one selects a norm Ψ defined on a vector space E containing \mathcal{F} . A small value of $\Psi(f)$ is interpreted as simplicity and simple functions are given priority by way of adding a penalty term to the empirical risk that is proportional to $\Psi(f)$. In particular, for some *regularization parameter* $\lambda > 0$, a regularized risk minimizer selects

$$\hat{f}_N \in \operatorname{argmin}_{f \in \mathcal{F}} \left(\frac{1}{N} \sum_{i=1}^N (f(X_i) - Y_i)^2 + \lambda \Psi(f) \right) ,$$

and the term $\Psi(f)$ is sometimes called the *penalty*.

Just as the tournament procedure from [11] outperforms empirical risk minimization (in fact, the tournament procedure attains the optimal tradeoff between accuracy and confidence under minimal assumptions), the regularized tournament which we present here, outperforms regularized risk minimization. Since regularized procedures require the minimization of a functional that has the empirical mean as a component, they suffer from the same disadvantages as empirical risk minimization. Therefore, the accuracy/confidence tradeoff exhibited by regularized risk minimization is suboptimal once one leaves the sub-Gaussian realm, and deteriorates further if the problem is more heavy-tailed. In contrast, we show that the regularized tournament attains the optimal accuracy/confidence tradeoff under rather minimal conditions and, in particular, in heavy-tailed problems.

The paper is organized as follows. In Section 2 we introduce a new regularized “tournament” procedure in a quite general framework and illustrate how it works on an important specific case, the tournament LASSO (see Section 2.1). In Section 3 the main general performance bound is presented for the regularized tournament procedure under certain specific choice of the parameters of the procedure. The proof of the main result is detailed in Section 4. Finally, in Section 5 two examples are worked out. The first is a “tournament” version of LASSO (introduced in Section 2.1) and the second is the tournament SLOPE, a generalized version of tournament LASSO.

2 The procedure

Let us now describe the regularized tournament procedure. Recall that the learner is given a closed and convex class of functions $\mathcal{F} \subset L_2(\mu)$ and a regularization function Ψ which is assumed to be a norm on $\operatorname{span}(\mathcal{F})$.

An underlying assumption is that \mathcal{F} can be naturally decomposed to a hierarchy of subclasses, that is, there is a finite decreasing collection of subsets of \mathcal{F} , denoted by $(\mathcal{F}_\ell)_{\ell=1}^K$, such that

$$\mathcal{F} = \mathcal{F}_1 \supset \cdots \supset \mathcal{F}_K .$$

The idea is that the sets \mathcal{F}_ℓ capture some notion of ‘complexity’: the larger ℓ is, the simpler the functions in \mathcal{F}_ℓ are.

An important example of a hierarchy, discussed below in detail, is based on the notion of sparsity in \mathbb{R}^d relative to a fixed orthonormal basis $(e_i)_{i=1}^d$: the vector $t \in \mathbb{R}^d$ is s -sparse if its representation in the basis $(e_i)_{i=1}^d$, $\sum_{i=1}^d t_i e_i$, has at most s nonzero coefficients. For the class of linear functionals $\mathcal{F} = \{\langle t, \cdot \rangle : t \in \mathbb{R}^d\}$ one may set $\mathcal{F}_\ell = \{\langle t, \cdot \rangle : t \text{ is } d/2^{\ell-1} \text{ sparse}\}$.

Given a hierarchy $(\mathcal{F}_\ell)_{\ell=1}^K$, the first, second and third phases of the procedure are performed on each one of the subclasses \mathcal{F}_ℓ separately and the procedure returns subclasses $\mathcal{H}_\ell \subset \mathcal{F}_\ell$ consisting of functions whose statistical performance is good enough to be considered as candidates for selection. The fourth and final phase compares the candidates selected from each class \mathcal{F}_ℓ and selects one of them, as we describe below.

The first three stages of the procedure use independent data. In order to accommodate this, one needs to split the available data into three independent parts. For simplicity of the presentation, we assume that these parts have equal size, each containing N samples. (Thus, the total sample size is $3N$ rather than N but this change of convention only affects the constants in the bounds that we do not make explicit in any case).

The first phase: constructing the ‘referee’ in \mathcal{F}_ℓ

Like all tournaments, this one also requires a ‘referee’, which is there to decide whether a match will be allowed to take place. More formally, the goal of the first phase is to get a data-dependent estimate of $L_2(\mu)$ distances between elements in \mathcal{F}_ℓ . Its output is a (data-dependent) function $\mathcal{DO}_\ell : \mathcal{F}_\ell \times \mathcal{F}_\ell \rightarrow \{0, 1\}$, which takes the value 1 when two functions in \mathcal{F}_ℓ are far-enough. As we explain in what follows, the statistical match between f and h is allowed to take place only if $\mathcal{DO}_\ell(f, h) = 1$.

This first step uses unlabeled samples only.

Definition 2.1. *Suppose we are given a sample $(X_i)_{i=1}^N$. Fix an appropriately chosen positive integer n_1 . For each $\ell = 1, \dots, K$,*

- (1) *fix well-chosen parameters ρ_ℓ and $r_{\ell,1}$;*
- (2) *split $(X_i)_{i=1}^N$ to n_1 disjoint blocks (I_j) of equal size, denoted by $m_1 = N/n_1$;*
- (3) *for every $f, h \in \mathcal{F}_\ell$ let*

$$v_j = \frac{1}{m_1} \sum_{i \in I_j} |f(X_i) - h(X_i)|,$$

and set $d_\ell(f, h)$ to be the median of $\{v_j : 1 \leq j \leq n_1\}$;

- (4) *set $\mathcal{DO}_\ell(f, h) = 1$ if either $\Psi(f - h) \geq \rho_\ell$ or if $\Psi(f - h) < \rho_\ell$ and $d_\ell(f, h) \geq r_{\ell,1}$.*

The second phase: ℓ -elimination

This first phase is used as input to the second phase. The latter is a comparison of the statistical performance of any two functions f and h in \mathcal{F}_ℓ using a sample $(X_i, Y_i)_{i=N+1}^{2N}$. The

idea is to define ‘statistical matches’ between any two functions $f, h \in \mathcal{F}_\ell$, with each match designed to determine which function out of the two is more suited for our needs. Because the ‘matches’ are based on a random sample, a reliable comparison is impossible when the two functions are ‘too close’. This is where the output of the first phase comes into the frame. The random binary function \mathcal{DO}_ℓ tells the learner when f and h are too close (when $\mathcal{DO}_\ell(f, h) = 0$) and in such cases, the outcome of the statistical match between f and h is useless.

More accurately, the second phase is defined as follows:

Definition 2.2. *Suppose we are given a sample $(X_i, Y_i)_{i=N+1}^{2N}$ and, for each $\ell = 1, \dots, K$, the output \mathcal{DO}_ℓ of the first phase, Then*

- (1) *let λ_ℓ and $n_{\ell,2}$ be well-chosen parameters, and set ρ_ℓ as in the first phase;*
- (2) *split $\{N + 1, \dots, 2N\}$ to $n_{\ell,2}$ coordinate blocks (I_j) of equal size, denoted by $m_{\ell,2}$;*
- (3) *for any $f, h \in \mathcal{F}_\ell$ set $f \succ h$ if*

$$\frac{1}{m_{\ell,2}} \sum_{i \in I_j} ((h(X_i) - Y_i)^2 - (f(X_i) - Y_i)^2) + \lambda_\ell (\Psi(h) - \Psi(f)) > 0$$

for the majority of the blocks I_j ;

- (4) *denote by*

$$\mathcal{H}'_\ell = \{f \in \mathcal{F}_\ell : f \succ h \text{ for every } h \in \mathcal{F}_\ell \text{ such that } \mathcal{DO}_\ell(f, h) = 1\}.$$

Remark 2.3. *If $\mathcal{DO}_\ell(f, h) = 0$ for every $h \in \mathcal{F}_\ell$, then $f \in \mathcal{H}'_\ell$ trivially. We use the same convention throughout this article: for a function to be eliminated from the tournament it must actually lose a match that has taken place (i.e., to a function h such that $\mathcal{DO}_\ell(f, h) = 1$).*

The output of the second phase is the sequence of classes $\mathcal{H}'_\ell \subset \mathcal{F}_\ell$. Each \mathcal{H}'_ℓ consists of the set of functions that were superior to all the other ‘competitors’ in the ℓ -tournament, at least in ‘matches’ that were considered reliable by \mathcal{DO}_ℓ . We will show that if $f^* \in \mathcal{F}_\ell$ then functions in \mathcal{H}'_ℓ are ‘close’ in $L_2(\mu)$ to f^* . Also, as will become clearer in what follows, if at some stage ℓ the set \mathcal{H}'_ℓ is empty, one may terminate the procedure and not continue to the $\ell + 1$ stage.

The third phase: ℓ -champions league

The third phase in the procedure is a further selection process, this time conducted among the set of ‘winners’ of the ℓ -elimination tournament.

Definition 2.4. *Suppose we are given $(X_i, Y_i)_{i=2N+1}^{3N}$. For each $\ell = 1, \dots, K$, let \mathcal{H}'_ℓ be the outcome of the second phase. Then*

- (1) let $n_{\ell,2}$ be as above, and set $r_{\ell,3}$ to be a well-chosen parameter;
- (2) split $\{2N + 1, \dots, 3N\}$ to $n_{\ell,2}$ disjoint blocks (I_j) of size $m_{\ell,2} = N/n_{\ell,2}$;
- (3) for any $f, h \in \mathcal{H}'_\ell$ set $f \gg h$ if

$$\frac{2}{m_{\ell,2}} \sum_{i \in I_j} (h(X_i) - f(X_i)) \cdot (f(X_i) - Y_i) \geq -r_{\ell,3}^2$$

for a majority of the blocks I_j . (Observe that it is possible that $f \gg h$ and $h \gg f$ at the same time);

- (4) let

$$\mathcal{H}_\ell = \{f \in \mathcal{H}'_\ell : f \gg h \text{ for every } h \in \mathcal{H}'_\ell\}.$$

This further selection process improves the outcome of the ℓ -elimination phase: instead of just ensuring that the functions are close to f^* , the ℓ -champions league selects functions whose risk is close to the risk of f^* .

The fourth phase: naming a winner

The outcome of the first three phases results in data-dependent choices $\mathcal{H}_\ell \subset \mathcal{F}_\ell$ (and obviously at this point there is no way of knowing that the sets \mathcal{H}_ℓ are nonempty). All the functions in each one of the \mathcal{H}_ℓ 's are in some sense the best estimators one can find within \mathcal{F}_ℓ , taking into account its size and the data at the learner's disposal. Naturally, it is reasonable to expect that the larger \mathcal{F}_ℓ is, the larger the error one will incur by selecting an estimator in it. Therefore, in the final stage the aim is to find the smallest class in the hierarchy in which a good estimator still exists.

Definition 2.5. Given the classes $(\mathcal{H}_\ell)_{\ell=1}^K$, let ℓ_1 be the largest integer ℓ such that $\bigcap_{j \leq \ell} \mathcal{H}_j \neq \emptyset$. Select \hat{f} to be any function in $\bigcap_{j \leq \ell_1} \mathcal{H}_j$. If \mathcal{H}_1 is empty, one may output an arbitrary function, say $\hat{f} = 0$.

The first three phases are the key components of the procedure. Out of the three, the first one is an adaptation of the *distance oracle* used in [11] and which had been introduced in [15]; the third component is essentially the same as the champions league stage in the tournament procedure from [11]. The fourth phase is a version of Lepski's method of intersecting confidence intervals (see, e.g. [4]).

The truly new component is the second phase. Its analysis combines ideas from [9] (which focused on regularized risk minimization in 'sparse' problems) and from [11]. As it is the main novelty in this article we present it in detail and only sketch the arguments needed in the analysis of the other components.

Naturally, at this point there are no guarantees that this procedure performs well, let alone that it is close to optimal. That requires some assumptions on the class \mathcal{F} , the hierarchy

$(\mathcal{F}_\ell)_{\ell=1}^K$ and the regularization function Ψ . Moreover, the parameters that each phase requires as inputs have to be specified for the procedure to make any sense. All these issues are explored in what follows. Before we dive into technicalities, and to give the reader a feeling of how the regularized tournament looks like in a familiar situation, let us describe the tournament version of the LASSO.

2.1 The tournament LASSO

Consider the following standard setup: Let X be an isotropic random vector in \mathbb{R}^d (that is, for every $t \in \mathbb{R}^d$, $\mathbb{E} \langle t, X \rangle^2 = \|t\|_2^2$, where $\|\cdot\|_2$ denotes the Euclidean norm in \mathbb{R}^d). Let Y be the unknown target random variable and set t_0 to be the minimizer in \mathbb{R}^d of the risk functional $t \rightarrow \mathbb{E}(\langle X, t \rangle - Y)^2$. For the sake of simplicity we assume that $Y = \langle t_0, X \rangle + W$ for W that is mean-zero, square integrable and independent of X .

In sparse recovery problems one believes that t_0 is supported on at most s coordinates with respect to the standard basis in \mathbb{R}^d —or at least it is well approximated by an s -sparse vector—, but one does not know that for certain. The LASSO procedure (introduced in [19]) selects $\hat{t} \in \mathbb{R}^d$ that minimizes the regularized empirical squared-loss functional

$$t \rightarrow \frac{1}{N} \sum_{i=1}^N (\langle t, X_i \rangle - Y_i)^2 + \lambda \|t\|_1$$

for a well chosen regularization parameter λ , and $\|t\|_1 = \sum_{i=1}^d |t_i|$ is the ℓ_1 -norm of t .

The problem with the LASSO is that when either the class members $\langle t, X \rangle$ or the target Y are heavy-tailed, the tradeoff between the accuracy with which the LASSO performs and the confidence with which that accuracy is attained is far from optimal. That suboptimal tradeoff is what the tournament LASSO aims to remedy.

For the time being we assume that for every $t \in \mathbb{R}^d$ and for every $p \leq c \log d$, $\|\langle X, t \rangle\|_{L_p} \leq L\sqrt{p} \|\langle X, t \rangle\|_{L_2} = L\sqrt{p} \|t\|_2$ (where the last inequality holds because X is isotropic). The random variable W is assumed to be square-integrable.

The natural class of functions in this context is $\mathcal{F} = \{\langle t, \cdot \rangle : t \in \mathbb{R}^d\}$ and the hierarchy is given in terms of the sparsity of t : we set, for $\ell = 1, \dots, K = \lceil \log_2 d \rceil + 1$,

$$\mathcal{F}_\ell = \{\langle t, \cdot \rangle : t \text{ is } d/2^{\ell-1} \text{ sparse}\} .$$

Let us go through the four phases of the tournament LASSO.

The first phase

For the definition of the ‘referee’ \mathcal{DO}_ℓ one requires to specify three parameters: ρ_ℓ , $r_{\ell,1}$ and n_1 . First, let $n_1 \sim_L N$, meaning that the cardinality of each block is a constant, depending only on L and which we denote by m_1 . Thus, $n_1 = N/m_1$.

In what follows, $C(L), C_0(L), \dots, C_4(L)$ denote appropriately chosen constants whose value depends only on L . (The precise form may be extracted from the analysis but it is of secondary importance for our purpose.)

Next, set $s = d/2^{\ell-1}$. If $s \geq C(L)N/\log(ed/N)$ then the class \mathcal{F}_ℓ for the corresponding ℓ is too large and no useful statistical information can be derived from the sample. In that case,

we set $r_{\ell,1}, \rho_\ell = \infty$, and in particular, $\mathcal{DO}_\ell \equiv 0$. Otherwise, when $s < C_0(L)N/\log(ed/N)$, set

$$r_\ell = C_1(L)\|W\|_{L_2} \sqrt{\frac{s}{d} \log\left(\frac{ed}{s}\right)} = C_2(L)\|W\|_{L_2} \sqrt{\frac{d\ell}{2^\ell N}}, \quad (2.1)$$

and

$$\rho_\ell = C_3(L)\|W\|_{L_2} \frac{s}{\sqrt{N}} \sqrt{\log\left(\frac{ed}{s}\right)} = C_4(L)\|W\|_{L_2} \frac{d\sqrt{\ell}}{2^\ell \sqrt{N}}. \quad (2.2)$$

Set $r_{\ell,1}$ to be any number larger than r_ℓ .

The first phase is performed as follows: given $(X_i)_{i=1}^N$, then for any t_1, t_2 supported on at most $d/2^{\ell-1}$ coordinates, let

- For $1 \leq j \leq n_1$ (which is proportional to N), set $v_j = \frac{1}{m_1} \sum_{i \in I_j} |\langle t_1 - t_2, X_i \rangle|$ and define $d_\ell(t_1, t_2)$ to be a median of $\{v_j : 1 \leq j \leq n_1\}$.
- Set $\mathcal{DO}_\ell(t_1, t_2) = 1$ if either

$$\|t_1 - t_2\|_1 \geq \rho_\ell \quad \text{or} \quad \|t_1 - t_2\|_1 < \rho_\ell \quad \text{and} \quad d_\ell(t_1, t_2) \geq r_{\ell,1}.$$

Remark 2.6. *Since X is isotropic, the $L_2(\mu)$ distance coincides with the ℓ_2 distance. Thus, \mathcal{DO}_ℓ can be chosen as a deterministic function: $\mathcal{DO}_\ell(t_1, t_2) = 1$ if $\|t_1 - t_2\|_1 \geq \rho_\ell$ or if $\|t_1 - t_2\|_1 < \rho_\ell$ and $\|t_1 - t_2\|_2 \geq r_{\ell,1}$. However, when X is not isotropic, the first phase of the procedure is truly required, which is why we ignored the simpler option that is available in the isotropic case.*

The ℓ -elimination phase

The elimination phase in the tournament LASSO requires two parameters, λ_ℓ and $n_{\ell,2}$ as inputs, as well as the outcomes of \mathcal{DO}_ℓ obtained in the first phase. Recalling the choices of $r_{\ell,1}$ from (2.1) and ρ_ℓ from (2.2), we set

$$r = c_1(L)r_{\ell,1} \quad \lambda_\ell = c_2(L) \frac{r^2}{\rho_\ell^2}, \quad \text{and} \quad n_{\ell,2} = c_3(L)N \min \left\{ \frac{r^2}{\|W\|_{L_2}^2}, 1 \right\}, \quad (2.3)$$

where c_1, c_2 and c_3 are constants that depend only on L . Let (I_j) be the natural partition on $\{1, \dots, N\}$ to $n_{\ell,2}$ disjoint blocks of equal cardinality, denoted by $m_{\ell,2}$. It follows that given the sample $(X_i, Y_i)_{i=N+1}^{2N}$ and $t_1, t_2 \in \mathcal{F}_\ell$, we set $t_1 \succ t_2$ if

$$\frac{1}{m_{\ell,2}} \sum_{i \in I_j} (\langle X_i, t_2 \rangle - Y_i)^2 - (\langle X_i, t_1 \rangle - Y_i)^2 + \lambda_\ell (\|t_2\| - \|t_1\|) > 0$$

for a majority of the blocks I_j . Therefore, \mathcal{H}'_ℓ consists of all $t_1 \in \mathbb{R}^d$ that satisfy $t_1 \succ t_2$ for any $t_2 \in \mathcal{F}_\ell$ such that $\mathcal{DO}_\ell(t_1, t_2) = 1$.

The ℓ -champions league phase

The ℓ -champions league phase receives as input the set \mathcal{H}'_ℓ produced in the second phase. Let $n_{\ell,2}$ and r be as in (2.3) and set $r_{\ell,3} = c(L)r$. Given $(X_i, Y_i)_{i=2N+1}^{3N}$ and $t_1, t_2 \in \mathcal{H}'_\ell$ then $t_1 \gg t_2$ if

$$\frac{2}{m_{\ell,2}} \sum_{i \in I_j} \langle X_i, t_2 - t_1 \rangle \cdot (\langle X_i, t_1 \rangle - Y_i) \geq -r_{\ell,3}^2,$$

and $\mathcal{H}_\ell = \{t_1 \in \mathcal{H}'_\ell : t_1 \gg t_2 \text{ for every } t_2 \in \mathcal{H}'_\ell\}$.

The fourth phase

Given the classes $(\mathcal{H}_\ell)_{\ell=1}^K$, ℓ_1 is the largest integer ℓ such that $\bigcap_{j \leq \ell} \mathcal{H}_j \neq \emptyset$. We select \hat{t} to be any element of $\bigcap_{j \leq \ell_1} \mathcal{H}_j$.

The following theorem, proved in Section 5.1, summarizes the performance of the tournament LASSO.

Theorem 2.7. *For $L \geq 1$ there are constants c_0, \dots, c_5 that depend only on L such that the following holds. Let X be an isotropic random vector in \mathbb{R}^d . Let $Y = \langle t_0, X \rangle + W$ where $t_0 \in \mathbb{R}^d$ and W is mean-zero, square-integrable, and independent of X . Assume that for every $t \in \mathbb{R}^d$ and any $1 \leq p \leq c \log d$, $\|\langle t, X \rangle\|_{L_p} \leq L\sqrt{p} \|\langle t, X \rangle\|_{L_2}$. Assume further that there is $v \in \mathbb{R}^d$ that is s -sparse such that*

$$\|t_0 - v\|_1 \leq c_1(L) \|W\|_{L_2} \cdot s \sqrt{\frac{\log(ed/s)}{N}}.$$

If $N \geq c_2(L)s \log(ed/s)$, and

$$\hat{r} \geq c_3(L) \|W\|_{L_2} \sqrt{\frac{s}{N} \log\left(\frac{ed}{s}\right)}$$

then with probability at least

$$1 - 2 \exp\left(-c_4(L)N \min\left\{1, \left(\frac{\hat{r}}{\|W\|_{L_2}}\right)^2\right\}\right), \quad (2.4)$$

we have

$$\begin{aligned} \|\hat{t} - t_0\|_2 &\leq c_5(L)\hat{r}, & \|\hat{t} - t_0\|_1 &\leq c_5(L)\|W\|_{L_2} s \sqrt{\frac{\log(ed/s)}{N}}, \\ \text{and } \mathbb{E}(\langle \hat{t}, X \rangle - Y)^2 &\leq \mathbb{E}(\langle t_0, X \rangle - Y)^2 + c_5(L)\hat{r}^2. \end{aligned} \quad (2.5)$$

Theorem 2.7 shows that the tournament LASSO attains the optimal accuracy/confidence tradeoff even though $f^*(X) - Y$ can be heavy-tailed. In fact, the estimate is what one would expect in the most friendly of scenarios: if X were a sub-Gaussian random vector (i.e., linear forms exhibiting a $\psi_2 - L_2$ norm equivalence with constant L rather than an $L_p - L_2$ moment equivalence going only up to $p \sim \log d$), and $f^*(X) - Y$ were a Gaussian random variable, independent of X . The LASSO does not come close to such an accuracy/confidence tradeoff under the weak moment assumption of Theorem 2.7.

Note that $r = c_3 \|W\|_{L_2} \sqrt{\frac{s}{N} \log\left(\frac{ed}{s}\right)}$ is the best accuracy parameter one can hope for even if the learner knows that t_0 is s -sparse and X and W are Gaussian. The difference between the performance of the tournament LASSO and the standard LASSO can be seen in the confidence with which this accuracy is attained. In the situation described in Theorem 2.7, the standard LASSO performs with that accuracy parameter only with constant confidence, because all that we assume on W is that it is square-integrable. In contrast, the tournament LASSO attains the accuracy (2.5) with the optimal exponential probability estimate (2.4).

The tournament LASSO does not require prior information on the degree of sparsity of t_0 to be carried out, but its success does depend on having a large-enough sample and on that \hat{r} is in the right range. The former is a constraint that any recovery procedure faces while the latter is easily achieved by running the procedure for $\hat{r}_j = r_0/2^j$ for a large initial value of r_0 , followed by estimating the risk of the obtained function on independent data. The straightforward details are omitted.

The one item that does require extra attention is that an upper estimate on $\|W\|_{L_2}$ is used in the choice of parameters of the tournament LASSO. At times one is simply given that information; this is often the case in signal processing problems, where the nature of the ‘noise’ is known to the learner. If not, one may use the data-dependent procedure from [17] which holds for more general noise models: it leads to upper and lower estimates on $\|f^*(X) - Y\|_{L_2}$ that are sharp up to absolute multiplicative constants and under minimal assumptions.

Of course, proving Theorem 2.7 requires some work, and the choice of parameters used in the first three phases has to be clarified. We explain the choice in the general case in the next two sections and return to the example of the tournament LASSO in Section 5.

3 The main result

In the general setup we study, we merely assume a rather weak fourth-moment assumption. More precisely, we work under the following conditions.

Assumption 3.1. *Let $\mathcal{F} \subset L_2(\mu)$ be a locally compact, convex class of functions. Let $Y \in L_2$ and assume that, for some constant $L > 0$,*

- *for every $f, h \in \mathcal{F}$, $\|f - h\|_{L_4} \leq L\|f - h\|_{L_2}$;*
- *$\|f^* - Y\|_{L_4} \leq \sigma_4$ for a known value σ_4 .*

Remark 3.1. *The condition that $\|f^* - Y\|_{L_4} \leq \sigma_4$ may easily be replaced by a combination of two assumptions: that for every $f \in \mathcal{F}$, $\|f - Y\|_{L_4} \leq L\|f - Y\|_{L_2}$; and that $\|f^* - Y\|_{L_2} \leq \sigma$ for some known constant $\sigma > 0$. Also, in the case of independent additive noise, that is, when $Y = f_0(X) + W$ where $f_0 \in \mathcal{F}$ and W that is mean-zero, square-integrable and independent of X , the assumption that $\|f^* - Y\|_{L_4} \leq \sigma_4$ may be replaced by the weaker one, that $\|W\|_{L_2} \leq \sigma$ for a known constant σ .*

The necessary modifications to the proofs are straightforward and we do not explore this observation further. Also, as noted previously, we refer the reader to [17] for a data-dependent procedure of estimating $\|f^ - Y\|_{L_2}$ which may be easily modified to an estimate on $\|f^* - Y\|_{L_4}$. Since that is not the main focus of this paper we do not pursue it further and instead assume that the learner has access to σ_4 or to σ .*

3.1 Complexity parameters of a class

The choice of parameters $r_{\ell,1}$, $r_{\ell,3}$ and ρ_ℓ depends on a certain notion of “complexity” of the underlying class. More accurately, there are geometric parameters that measure the metric entropy of “localizations” of the class as well as the oscillation of various natural random processes indexed by those localizations.

The complexity is measured in terms of four parameters, depending both on the class \mathcal{F} and the distribution of (X, Y) . The four play an essential role in describing the optimal performance of learning procedures and for detailed discussion on the meaning we refer to Mendelson [13, 14] and Lugosi and Mendelson [11].

Before we define the four parameters we need some notation. Denote the unit ball in $L_2(\mu)$ by $D = \{f : \|f\|_{L_2} \leq 1\}$ and let $S = \{f : \|f\|_{L_2} = 1\}$ be the unit sphere. For $h \in L_2(\mu)$ and $r > 0$, we write $D_h(r) = \{f : \|f - h\|_{L_2} \leq r\}$. In a similar fashion for the norm Ψ used as a regularization function, let $\mathcal{B} = \{f : \Psi(f) \leq 1\}$, set $\rho\mathcal{B} = \{f : \Psi(f) \leq \rho\}$ and $\mathcal{B}_h(\rho) = \{f : \Psi(f - h) \leq \rho\}$.

In what follows we make two important modifications to the definitions of the complexity parameters used in [13, 14, 11]. First, just like in the above-mentioned articles, we are interested in “localized” classes. However, because regularized procedures are affected by two norms, Ψ and $L_2(\mu)$, the localization has to be with respect to both of them. Therefore, the “localization” of \mathcal{F} , centred in h and of radii $\rho, r > 0$ is defined by

$$\mathcal{F}_{h,\rho,r} = (\mathcal{F} - h) \cap (\rho\mathcal{B} \cap rD) = \{f - h : f \in \mathcal{F}, \Psi(f - h) \leq \rho, \|f - h\|_{L_2} \leq r\}.$$

The second minor modification is that each complexity parameter is associated with the ‘worse case’ centre $h \in \mathcal{F}'$ for some fixed $\mathcal{F}' \subset \mathcal{F}$, and not necessarily with the whole of \mathcal{F} .

Two of the four parameters are defined using the notion of *packing numbers*.

Definition 3.2. *Given a set $H \subset L_2(\mu)$ and $\varepsilon > 0$, denote the ε -packing number of H by $\mathcal{M}(H, \varepsilon D)$. In other words, $\mathcal{M}(H, \varepsilon D)$ is the maximal cardinality of a subset $\{h_1, \dots, h_m\} \subset H$, for which $\|h_i - h_j\|_{L_2} \geq \varepsilon$ for every $i \neq j$.*

The first relevant parameter $\lambda_{\mathbb{Q}}$ is defined as follows with appropriate numerical constants κ and η :

Definition 3.3. *Fix $\rho > 0$ and $h \in \mathcal{F}$. For $\kappa, \eta > 0$, set*

$$\lambda_{\mathbb{Q}}(\kappa, \eta, h, \rho) = \inf\{r : \log \mathcal{M}(\mathcal{F}_{h,\rho,r}, \eta r D) \leq \kappa^2 N\}. \quad (3.1)$$

For $\mathcal{F}' \subset \mathcal{F}$ let

$$\lambda_{\mathbb{Q}}(\kappa, \eta, \rho) = \sup_{h \in \mathcal{F}'} \lambda_{\mathbb{Q}}(\kappa, \eta, h, \rho).$$

While κ and η are adjustable parameters, we are mainly interested in the behaviour of $\lambda_{\mathbb{Q}}$ as a function of ρ . The way one selects ρ is clarified later.

The reason behind the choice of $\lambda_{\mathbb{Q}}$ comes from high-dimensional geometry: if the class $\mathcal{F}_{h,\rho,r} \subset L_2(\mu)$ is ‘less complex’ than an N -dimensional Euclidean ball of radius $\sim r$, then it can be covered by $\exp(cN)$ balls of radius $\sim r$. Thus, $\lambda_{\mathbb{Q}}$ captures the smallest radius for which there is still a chance that $\mathcal{F}_{h,\rho,r}$ resembles in some sense a subset of a ball of radius $\sim r$ in \mathbb{R}^N . The hope is that for such a choice of radius r , at least part of the metric

structure of $\mathcal{F}_{h,\rho,r}$ is reflected in a typical random set $\{(u(X_i))_{i=1}^N : u \in \mathcal{F}_{h,\rho,r}\}$. In more statistical terms, the hope is that a sample of cardinality N provides the learner with enough information to “separate” class members that are far enough, and the distance at which a sample of cardinality N no longer suffices is $\lambda_{\mathbb{Q}}$.

The next parameter, denoted by $\lambda_{\mathbb{M}}$, is also defined in terms of the packing numbers of the localization $\mathcal{F}_{h,\rho,r}$, though at a different scaling than $\lambda_{\mathbb{Q}}$.

Definition 3.4. Fix $h \in \mathcal{F}$ and $\rho > 0$. Let $\kappa > 0$, $0 < \eta < 1$, and define

$$\lambda_{\mathbb{M}}(\kappa, \eta, h, \rho) = \inf\{r : \log \mathcal{M}(\mathcal{F}_{h,\rho,r}, \eta r D) \leq \kappa^2 N r^2\}. \quad (3.2)$$

Also, for $\mathcal{F}' \subset \mathcal{F}$ let

$$\lambda_{\mathbb{M}}(\kappa, \eta, \rho) = \sup_{h \in \mathcal{F}'} \lambda_{\mathbb{M}}(\kappa, \eta, h, \rho).$$

For the remaining two complexity parameters, let $(\varepsilon_i)_{i=1}^N$ be independent, symmetric $\{-1, 1\}$ -valued random variables that are independent of $(X_i, Y_i)_{i=1}^N$.

Definition 3.5. Fix $h \in \mathcal{F}$ and $\rho > 0$. For $\kappa > 0$ let

$$r_E(\kappa, h, \rho) = \inf\left\{r : \mathbb{E} \sup_{u \in \mathcal{F}_{h,\rho,r}} \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i u(X_i) \right| \leq \kappa \sqrt{N} r\right\}, \quad (3.3)$$

and for $\mathcal{F}' \subset \mathcal{F}$ set $r_E(\kappa, \rho) = \sup_{h \in \mathcal{F}'} r_E(\kappa, h, \rho)$.

The idea behind r_E is similar to the one behind $\lambda_{\mathbb{Q}}$. $r_E(\kappa, h, \rho)$ is the smallest radius for which the expected supremum of the process $u \rightarrow \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i u(X_i)$ indexed by $\mathcal{F}_{h,\rho,r}$ exceeds the expectation of the supremum of the Bernoulli process in \mathbb{R}^N , $v \rightarrow \sum_{i=1}^N \varepsilon_i v_i$, indexed by a Euclidean ball of radius r . Thus, the point is to identify when a typical random set $\{N^{-1/2}(u(X_i))_{i=1}^N : u \in \mathcal{F}_{h,\rho,r}\}$ is richer than Euclidean ball of radius $\sim r$ in \mathbb{R}^N , where ‘richer’ is measured in terms of the Bernoulli mean-width of a set rather than by its metric entropy.

Definition 3.6. Fix $h \in \mathcal{F}$ and $\rho > 0$. For $\kappa > 0$, set $\bar{r}_{\mathbb{M}}(\kappa, h, \rho)$ to be

$$\bar{r}_{\mathbb{M}}(\kappa, h, \rho) = \inf\left\{r : \mathbb{E} \sup_{u \in \mathcal{F}_{h,\rho,r}} \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i u(X_i) \cdot (h(X_i) - Y_i) \right| \leq \kappa \sqrt{N} r^2\right\}. \quad (3.4)$$

For $\sigma > 0$ put $\mathcal{F}_Y^{(\sigma)} = \{f \in \mathcal{F}' : \|f(X) - Y\|_{L_2} \leq \sigma\}$ and let $\tilde{r}_{\mathbb{M}}(\kappa, \sigma, \rho) = \sup_{h \in \mathcal{F}_Y^{(\sigma)}} \bar{r}_{\mathbb{M}}(\kappa, h, \rho)$.

Finally, suppose that the distribution of (X, Y) is such that $\|Y - f^*(X)\|_{L_4} \leq \sigma_4$ for a known constant $\sigma_4 > 0$. The “complexity” of \mathcal{F} relative to centres in \mathcal{F}' and radius ρ is

$$r^*(\mathcal{F}, \mathcal{F}', \rho) = \max\{\lambda_{\mathbb{Q}}(c_1, c_2, \rho), \lambda_{\mathbb{M}}(c_1/\sigma_4, c_2, \rho), r_E(c_1, \rho), \tilde{r}_{\mathbb{M}}(c_1, \sigma_4, \rho)\}. \quad (3.5)$$

Here c_1, c_2 are appropriate positive numerical constants. (“Appropriate” means that $r^*(\mathcal{F}, \mathcal{F}', \rho)$ satisfies Propositions 4.1, 4.4 and 4.7 below). The existence of such constants is proved in [11] when $\mathcal{F}' = \mathcal{F}$, i.e., when any function in \mathcal{F} is a ‘legal choice’ of a centre, and under Assumption 3.1. In that case, the constants depend only on the value of L .

When \mathcal{F} and \mathcal{F}' are clear from the context, we simply write $r^*(\rho)$ for $r^*(\mathcal{F}, \mathcal{F}', \rho)$.

3.2 Properties of the hierarchy

After the complexity parameters are set in place, let us identify the conditions on the hierarchy $(\mathcal{F}_\ell)_{\ell=1}^K$ that are needed in the analysis of the regularized tournament.

Recall that \mathcal{F} is a (convex) subset of a normed space (E, Ψ) ; E is also a subspace of $L_2(\mu)$, though the norms Ψ and $\|\cdot\|_{L_2(\mu)}$ may have nothing to do with each other. Let B_{Ψ^*} and S_{Ψ^*} denote the unit ball and unit sphere in the dual space to (E, Ψ) , respectively. Therefore, B_{Ψ^*} consists of all the linear functionals $z \in E^*$ for which $\sup_{\{x \in E: \Psi(x)=1\}} |z(x)| \leq 1$. A linear functional $z^* \in S_{\Psi^*}$ is a norming functional for $f \in E$ if $z^*(f) = \Psi(f)$.

Definition 3.7. Let $\Gamma_f(\rho) \subset S_{\Psi^*}$ be the collection of functionals that are norming for some $v \in \mathcal{B}_f(\rho/20)$. Set

$$\Delta_\ell(\rho, r) = \inf_{f \in \mathcal{F}_\ell} \inf_h \sup_{z \in \Gamma_f(\rho)} z(h - f) ,$$

where the inner infimum is taken in the set

$$\{h \in \mathcal{F} : \Psi(h - f) = \rho \text{ and } \|h - f\|_{L_2} \leq r\} . \quad (3.6)$$

Let us examine $\Delta_\ell(\rho_\ell, r_\ell)$ and explain its meaning for some fixed values $\rho, r > 0$. Note that $\Delta_\ell(\rho, r) \leq \rho$. Indeed, $\Gamma_f(\rho) \subset S_{\Psi^*}$ and if $z \in S_{\Psi^*}$ and $\Psi(h - f) \leq \rho$ then

$$|z(h - f)| \leq \Psi^*(z) \cdot \Psi(h - f) \leq \rho .$$

The interesting situation is when one can ensure a reverse inequality, that is, that $\Delta_\ell(\rho, r)$ is proportional to ρ , say $\Delta_\ell(\rho, r) \geq (4/5)\rho$. Such a lower estimate on Δ_ℓ implies the following. Let $f \in \mathcal{F}_\ell$ and $h \in \mathcal{F}$ for which $\Psi(h - f) = \rho$ and $\|f - h\|_{L_2} \leq r$. It follows that there is some $z \in S_{\Psi^*}$ and $v \in \mathcal{B}_f(\rho/20)$ such that z is norming for v and $z(h) - z(f) \geq \Delta_\ell(\rho, r)$. Therefore,

$$\begin{aligned} \Psi(h) - \Psi(f) &= \Psi(h) - \Psi(v + (f - v)) \geq \Psi(h) - \Psi(v) - \Psi(f - v) \\ &\geq z(h) - z(v) - \Psi(f - v) \geq z(h) - z(f) - 2\Psi(f - v) \\ &\geq \Delta_\ell(\rho, r) - \rho/10 \geq 3\rho/5 . \end{aligned}$$

A lower bound on the term $\Psi(h) - \Psi(f)$ plays an essential role in the study of the elimination phase of the regularized tournament, when one has to compare

$$\frac{1}{m_{\ell,2}} \sum_{i \in I_j} (f(X_i) - Y_i)^2 + \lambda_\ell \Psi(f) \quad \text{and} \quad \frac{1}{m_{\ell,2}} \sum_{i \in I_j} (h(X_i) - Y_i)^2 + \lambda_\ell \Psi(h) .$$

Obviously, ensuring that $\Delta_\ell(\rho, r) \geq (4/5)\rho$ becomes simpler when the set $\Gamma_f(\rho)$ is large. In the extreme case, when $\rho > 30\Psi(f)$, then $\mathcal{B}_f(\rho/20)$ contains a nontrivial Ψ -ball around 0; thus, $\Gamma_f(\rho) = S_{\Psi^*}$ and $\Delta_\ell(\rho, r) = \rho$. The other extreme is if ρ is very small and one is left only with the functionals that are norming for f itself. Intuitively, the right choice of ρ_ℓ is the smallest one for which, for $r_\ell = r^*(\mathcal{F}, \mathcal{F}_\ell, \rho_\ell)$, one has $\Delta_\ell(\rho_\ell, r_\ell) \geq 4\rho_\ell/5$.

Definition 3.8. The sequence $(\mathcal{F}_\ell, \rho_\ell)_{\ell=1}^K$ is compatible if

- (1) $\mathcal{F} = \mathcal{F}_1 \supset \mathcal{F}_2 \supset \dots \supset \mathcal{F}_K$ is a finite hierarchy;

(2) $(\rho_\ell)_{\ell=1}^K$ is decreasing and $r_\ell = r^*(\mathcal{F}, \mathcal{F}_\ell, \rho_\ell)$;

(3) for every $1 \leq \ell \leq K$, $\Delta_\ell(\rho_\ell, r_\ell) \geq 4\rho_\ell/5$.

We allow the choice of $\rho_\ell = r_\ell = \infty$ for $\ell = 1, \dots, \ell_0$. In such cases the compatibility condition is to be verified from $\ell_0 + 1$ onward.

We are now ready to specify the parameters used in the definition of a regularized tournament.

- Let α , β , m_1 , θ_1 and θ_2 be well chosen constants that depend only on the norm equivalence constant L from Assumption 3.1, and assume that one has access to the value σ_4 from that assumption.
- Assume further that $(\mathcal{F}_\ell, \rho_\ell)_{\ell=1}^K$ is a compatible sequence, let r_ℓ be as in Definition 3.8 and set $\widehat{r}_\ell > r_\ell$ to be decreasing.

We set the following choice of parameters:

First phase: Let ρ_ℓ as above, and set

$$n_1 = \frac{N}{m_1} \quad \text{and} \quad r_{\ell,1} = \beta \widehat{r}_\ell.$$

ℓ -elimination phase: Set

$$n_{\ell,2} = \theta_1 N \min \left\{ 1, \frac{\widehat{r}_\ell^2}{\sigma_4^2} \right\} \quad \text{and} \quad \lambda_\ell = \theta_2 \frac{\widehat{r}_\ell^2}{\rho_\ell}.$$

ℓ -champions league: Set $r_{\ell,3} = c(\beta/\alpha)\widehat{r}_\ell$ for a suitable absolute constant c ($c = 1/\sqrt{5}$ would do).

With these choices set in place, let us formulate the main result of this article.

Theorem 3.9. *For $L \geq 1$ there exist constants c_0, \dots, c_2 that depend only on L and for which the following holds. Let (\mathcal{F}, X, Y) satisfy Assumption 3.1 and set $(\mathcal{F}_\ell)_{\ell=1}^K$ be a compatible sequence of the class \mathcal{F} . If ℓ^* is the largest index ℓ such that $f^* \in \mathcal{F}_\ell$, then with probability at least*

$$1 - 2 \sum_{\ell=1}^{\ell^*} \exp \left(-c_0(L) N \min \left\{ 1, \frac{\widehat{r}_\ell^2}{\sigma_4^2} \right\} \right),$$

we have

$$\Psi(\widehat{h} - f^*) \leq \rho_{\ell^*}, \quad \|\widehat{h} - f^*\|_{L_2} \leq c_1(L)\widehat{r}_{\ell^*}, \quad \text{and} \quad R(\widehat{h}) - R(f^*) \leq c_2(L)\widehat{r}_{\ell^*}^2.$$

3.3 Discussion

The main message of Theorem 3.9 is that the regularized tournament procedure achieves essentially the best performance that can be expected even under strong assumptions of sub-Gaussian distributions. The regularized procedure yields (almost) the optimal accuracy-confidence tradeoff for any accuracy parameter $r \geq r_{\ell^*}$: it behaves as if it “knew” the location of f^* in the hierarchy without actually knowing it. Indeed, the accuracy/confidence tradeoff established in Theorem 3.9 is essentially the best possible for any learning procedure taking values only in \mathcal{F}_{ℓ^*} .

We emphasize that Theorem 3.9 is quite general though finding the adequate parameters of the procedure requires additional work. For the “tournament” version of LASSO and SLOPE we work out the details in Section 5 under certain assumptions (such as isotropic design vector X and approximately sparse linear regression function) for illustration. Some of these assumption may be weakened but we prefer to keep the presentation as simple as possible.

Related work. The sensitivity of empirical risk minimization (or least squares regression) to heavy-tailed distributions has been pointed out and several proposals of robust regression function estimates have been made that avoid this sensitivity. We refer to Audibert and Catoni [1], Hsu and Sabato [5], Lerasle and Oliveira [10], Minsker [18], Brownlees, Joly, and Lugosi [3], Lugosi and Mendelson [11] for a sample of the literature. This paper mostly builds upon the methodology of median-of-means tournaments, developed in [11], (see also Lugosi and Mendelson [12]). Here we extend this methodology to the analysis of regularized robust risk minimization similarly to how the paper of Lecué and Mendelson [8] analyzes standard regularized risk minimization. The analysis of the LASSO and SLOPE procedures of [8] was extended and generalized by Bellec, Lecué, and Tsybakov [2]. In an independent parallel work to ours, and building on the arguments developed in [11], Lecué and Lerasle [6] point out a connection of median-of-means tournaments to Le Cam’s estimators, develop a version of LASSO—the so-called MOM-LASSO—and prove a performance bound quite similar to Theorem 2.7.

4 Analyzing the four phases

The first phase - the ℓ -distance oracle

At each stage $\ell = 1, \dots, K$ of the procedure, one initially uses a modification of the distance oracle from [11].

The ℓ -distance oracle is a data-dependent procedure that provides information on the distances between functions. It is used for any pair $f, h \in \mathcal{F}$, and aims at determining if $\Psi(f - h) \geq \rho_\ell$, or, if $\Psi(f - h) \leq \rho_\ell$, whether $\|f - h\|_{L_2} \geq r_{\ell,1}$. Note that Ψ is a known norm and therefore, $\Psi(f - h)$ is known for any pair $f, h \in \mathcal{F}$ but $\|f - h\|_{L_2}$ needs to be (crudely) estimated.

Recall that we work under Assumption 3.1, and that \mathcal{DO}_ℓ is defined as follows: we split $\{1, \dots, N\}$ to $n = n_1$ disjoint blocks $(I_j)_{j=1}^n$, each one of cardinality $m = N/n$. For a sample $\mathcal{C}_1 = (X_i)_{i=1}^N$ and functions f and h , let $w = (|f(X_i) - h(X_i)|)_{i=1}^N$ and set

$$\Phi_{\mathcal{C}_1}(f, h) = \text{Med}_m(w) ,$$

where $\text{Med}_m(w)$ is a median of the n values $\frac{1}{m} \sum_{i \in I_j} |(f - h)(X_i)|$.

The behaviour of Φ described below has been established in Mendelson [15] (see also Lugosi and Mendelson [11]):

Proposition 4.1. *Let \mathcal{F} satisfy Assumption 3.1. There exist constants m , $0 < \alpha < 1 < \beta$, and κ, η and c , all of them depending only on L , for which the following holds.*

- For $1 \leq \ell \leq K$ and $\rho > 0$, let $r > r^*(\mathcal{F}, \mathcal{F}_\ell, \rho)$, where r^* is defined relative to the constants κ and η .
- Let $n = N/m$ and fix $f \in \mathcal{F}_\ell$.

Then, with probability at least $1 - 2\exp(-cN)$, for any $h \in \mathcal{F}$ that satisfies $\Psi(f, h) \leq \rho$,

(1) if $\Phi_{\mathcal{C}_1}(f, h) \geq \beta r$ then

$$\beta^{-1}\Phi_{\mathcal{C}_1}(f, h) \leq \|f - h\|_{L_2} \leq \alpha^{-1}\Phi_{\mathcal{C}_1}(f, h);$$

(2) if $\Phi_{\mathcal{C}_1}(f, h) < \beta r$ then $\|f - h\|_{L_2} \leq (\beta/\alpha)r$.

The proof of Proposition 4.1 is a direct outcome of Proposition 3.2 from [11], applied to the set $\mathcal{F} \cap \mathcal{B}_f(\rho)$ for a fixed centre $f \in \mathcal{F}_\ell$ and we shall not present it here.

Based on Proposition 4.1, our choice of parameters in the first phase is clear. Recall that for $\widehat{r}_\ell > r_\ell = r^*(\mathcal{F}, \mathcal{F}_\ell, \rho_\ell)$ and $f_1, f_2 \in \mathcal{F}$ we set $\mathcal{DO}_\ell(f_1, f_2) = 1$ if either $\Psi(f_1 - f_2) > \rho_\ell$ or if $\Psi(f_1 - f_2) \leq \rho_\ell$ and $\Phi_{\mathcal{C}_1}(f_1, f_2) \geq \beta\widehat{r}_\ell$.

Thanks to Proposition 4.1, it follows that for a fixed centre $f \in \mathcal{F}_\ell$ (which is selected as f^* in what follows), with probability at least $1 - 2\exp(-cN)$ if $h \in \mathcal{F}$, and $\mathcal{DO}_\ell(f, h) = 0$ then $\Psi(h - f) \leq \rho_\ell$ and $\|f - h\|_{L_2} \leq \widehat{r}_\ell$.

Remark 4.2. *Although Proposition 4.1 is formulated for a designated single centre f , it is straightforward to extend it to any centre in \mathcal{F} and obtain a uniform distance oracle that holds for any pair $f, h \in \mathcal{F}$.*

The second phase— ℓ -elimination

Fix $1 \leq \ell \leq K$ and let $f, h \in \mathcal{F}$. Recall the definition of a regularized match between f and h : first, the ℓ -distance oracle defined above uses the first part of the sample $\mathcal{C}_1 = (X_i, Y_i)_{i=1}^N$ to determine the value of $\mathcal{DO}_\ell(f, h)$. If $\mathcal{DO}_\ell(f, h) = 0$, the match is abandoned.

Each match that is allowed to take place by the ℓ -distance oracle is played using the second part of the sample, $\mathcal{C}_2 = (X_i, Y_i)_{i=N+1}^{2N}$. The sub-sample is partitioned to $n = n_{\ell,2}$ blocks $(I_j)_{j=1}^n$ of cardinality $m = N/n$ where recall that n is chosen as

$$\theta_1(L)N \min \left\{ 1, \frac{\widehat{r}_\ell^2}{\sigma_4^2} \right\}. \quad (4.1)$$

for a well-chosen constant θ_1 that depends only on the equivalence constant L from Assumption 3.1. We also set

$$\lambda_\ell = \theta_2(L) \frac{\widehat{r}_\ell^2}{\rho_\ell},$$

with the choices of both constants θ_1 and θ_2 specified below.

The key definition in the elimination stage is the choice of a winner in a ‘statistical match’ between two functions in \mathcal{F}_ℓ .

Definition 4.3. *The function f defeats h (denoted by $f \succ h$) if*

$$\frac{1}{m} \sum_{i \in I_j} ((h(X_i) - Y_i)^2 - (f(X_i) - Y_i)^2) + \lambda_\ell(\Psi(h) - \Psi(f)) > 0$$

on a majority of the blocks I_j .

The set of winners \mathcal{H}'_ℓ of the ℓ -th elimination round consists of all the functions in \mathcal{F}_ℓ that have not lost a single match against a function in \mathcal{F}_ℓ .

The idea behind the elimination round is to ‘exclude’ functions that are far from f^* (without knowing the identity of f^* , of course). To that end, it suffices to show that, if $f^* \in \mathcal{F}_\ell$, then with high probability, f^* wins all the matches it takes part in. Indeed, that implies that \mathcal{H}'_ℓ is nonempty, and that all the matches between f^* and any $h \in \mathcal{H}'_\ell$ must have been abandoned; therefore, $\mathcal{DO}_\ell(f^*, h) = 0$, that is,

$$\Psi(f^* - h) \leq \rho_\ell \quad \text{and} \quad \|f^* - h\|_{L_2} \leq (\beta/\alpha)\widehat{r}_\ell. \quad (4.2)$$

The next theorem describes the outcome of the ℓ -elimination phase. Its proof may be found in Section 4.1.

Proposition 4.4. *Using the notation above, if $f^* \in \mathcal{F}_\ell$ then, with probability at least*

$$1 - 2 \exp\left(-c_0 N \min\left\{1, \frac{\widehat{r}_\ell^2}{\sigma_4^2}\right\}\right),$$

f^ wins all the matches it participates in. In particular, on that event, if $h \in \mathcal{H}'_\ell$, then (4.2) holds.*

4.1 Proof of Proposition 4.4—highlights

To explain why this elimination phase preforms well even when \mathcal{F} is very large, define, for each block I_j ($j = 1, \dots, n$),

$$B_{h,f}^\lambda(j) = \frac{1}{m} \sum_{i \in I_j} ((h(X_i) - Y_i)^2 - (f(X_i) - Y_i)^2) + \lambda(\Psi(h) - \Psi(f)). \quad (4.3)$$

Note that the regularized empirical excess risk of h on block I_j is $B_{h,f^*}^\lambda(j)$.

Consider the ℓ -th stage of the regularized tournament. The assertion of Proposition 4.4 is that, if $f^* \in \mathcal{F}_\ell$, then it is a winner of all the elimination phase matches it participates in. Hence, Proposition 4.4 is proved once we ensure that for the right choice of $\lambda = \lambda_\ell$, with high probability, if $h \in \mathcal{F}$ and $\mathcal{DO}_\ell(f^*, h) = 1$ then $B_{h,f^*}^\lambda(j)$ is positive for most of the blocks I_j .

To that end, observe that

$$\begin{aligned} & \frac{1}{m} \sum_{i \in I_j} ((h(X_i) - Y_i)^2 - (f^*(X_i) - Y_i)^2) \\ &= \frac{1}{m} \sum_{i \in I_j} (h - f^*)^2(X_i) + \frac{2}{m} \sum_{i \in I_j} (h - f^*)(X_i) \cdot (f^*(X_i) - Y_i), \end{aligned}$$

which is the natural decomposition of the empirical excess risk functional into its quadratic and multiplier components. Setting

$$\mathbb{Q}_{h,f}(j) = \frac{1}{m} \sum_{i \in I_j} (h - f)^2(X_i) \quad \text{and} \quad \mathbb{M}_{h,f}(j) = \frac{2}{m} \sum_{i \in I_j} (h - f)(X_i) \cdot (f(X_i) - Y_i),$$

we have

$$\begin{aligned} B_{h,f^*}^\lambda(j) &= \frac{1}{m} \sum_{i \in I_j} ((h(X_i) - Y_i)^2 - (f^*(X_i) - Y_i)^2) + \lambda(\Psi(h) - \Psi(f^*)) \\ &= \mathbb{Q}_{h,f^*}(j) + \mathbb{M}_{h,f^*}(j) + \lambda(\Psi(h) - \Psi(f^*)). \end{aligned}$$

The first observation we require is a version of a deterministic result from [9, Theorem 3.2] (see the appendix for the proof).

Lemma 4.5. *Let $f^* \in \mathcal{F}_\ell$ and $h \in \mathcal{F}$ for which either $\Psi(h - f^*) = \rho$, or $\Psi(h - f^*) < \rho$ and $\|h - f^*\|_{L_2} \geq r$. Assume that $\Delta_\ell(\rho, r) \geq 4\rho/5$, that λ satisfies*

$$\frac{C}{2} \cdot \frac{r^2}{\rho} \leq \lambda \leq \frac{3C}{4} \cdot \frac{r^2}{\rho} \quad (4.4)$$

for some $C > 0$. Assume further that

$$\mathbb{M}_{h,f^*}(j) \geq -(C/4) \max \{ \|h - f^*\|_{L_2}^2, r^2 \}, \quad (4.5)$$

and if also $\|h - f^*\|_{L_2} \geq r$ then

$$\mathbb{Q}_{h,f^*}(j) \geq C \|h - f^*\|_{L_2}^2. \quad (4.6)$$

Then

$$\mathbb{Q}_{h,f^*}(j) + \mathbb{M}_{h,f^*}(j) + \lambda(\Psi(h) - \Psi(f^*)) > 0.$$

Moreover, if $f \in \mathcal{F}$ such that $f = f^* + \alpha(h - f^*)$ for $\alpha > 1$ then also

$$\mathbb{Q}_{f,f^*}(j) + \mathbb{M}_{f,f^*}(j) + \lambda(\Psi(f) - \Psi(f^*)) > 0.$$

Thanks to Lemma 4.5 (which is proved in the appendix), all that is required to prove Proposition 4.4 is to verify that with the requested probability, (4.5) and (4.6) hold uniformly in h and on a majority of the blocks I_j , provided that $f^* \in \mathcal{F}_\ell$; thus f^* defeats any function in $h \in \mathcal{F}$ that satisfies either $\Psi(h - f^*) \geq \rho$ or $\Psi(h - f^*) \leq \rho$ and $\|f^* - h\|_{L_2} \leq r$.

The fact that we have the required control over coordinate blocks is formulated in the following lemma. Its proof may be found in the appendix.

Lemma 4.6. *There exists an absolute constant c and a constant $C_1 = C_1(L, \tau)$ for which the following holds. Let $f^* \in \mathcal{F}_\ell$. For $0 < \tau < 1$, with probability at least $1 - 2 \exp(-c\tau^2 n)$, for every $h \in \mathcal{B}_{f^*}(\rho)$ that satisfies $\|f - f^*\|_{L_2} \geq \widehat{r}_\ell$, we have*

$$|\{j : \mathbb{Q}_{f,f^*}(j) \geq C_1 \|f - f^*\|_{L_2}^2\}| \geq (1 - \tau) n$$

and

$$\left| \left\{ j : \mathbb{M}_{f,f^*}(j) \leq -\frac{C_1}{4} \|f - f^*\|_{L_2}^2 \right\} \right| \leq \tau n.$$

Moreover, for every $h \in \mathcal{B}_{f^*}(\rho)$ that satisfies $\|f - f^*\|_{L_2} \leq \widehat{r}_\ell$, we have

$$\left| \left\{ j : \mathbb{M}_{f,f^*}(j) \leq -\frac{C_1}{4} r^2 \right\} \right| \leq \tau n.$$

It follows from Lemma 4.6 that if $\tau < 1/4$ then with probability at least $1 - 2\exp(-c\tau^2 n)$, for every h as in Lemma 4.5, conditions (4.5) and (4.6) hold for $C = C_1$ and $r = \widehat{r}_\ell$ on the majority of the blocks I_j . Hence, setting $\tau = 1/10$, on an event with probability at least $1 - 2\exp(-cn)$, f^* wins all the matches it participates in and that are allowed to take place by the ℓ -distance oracle, as we require.

The third phase— ℓ -champions league

Once Proposition 4.4 is established, we turn to the third phase, aimed at selecting a set of “winners”. In order to do that, we run the champions league tournament defined in Definition 2.4, performed in each one of the sets \mathcal{H}'_ℓ . (This is the same procedure used in [11].) The crucial point is that if $f^* \in \mathcal{F}_\ell$ then \mathcal{H}'_ℓ satisfies the necessary conditions for a champions league tournament: that $f^* \in \mathcal{H}'_\ell$ and all the functions $h \in \mathcal{H}'_\ell$ have a mean-squared error at most $\sim \widehat{r}_\ell$.

Recall that the ℓ -champions league consists of matches that use the third part of the sample $\mathcal{C}_3 = (X_i, Y_i)_{i=2N+1}^{3N}$. Let $(I_j)_{j=1}^n$ be the partition of $\{2N+1, \dots, 3N\}$ to n blocks, for the same value of $n = n_{\ell,2}$ as in the ℓ -elimination phase, given in (4.1). Also, for α and β as in Proposition 4.1, set $c = \beta/\alpha$ and for $f, h \in \mathcal{H}'_\ell$, let $\Psi_{h,f} = (h(X) - f(X))(f(X) - Y)$.

Recall that the function f wins its home match against h (denoted by $f \gg h$) if

$$\frac{2}{m} \sum_{i \in I_j} \Psi_{h,f}(X_i, Y_i) \geq -(2c\widehat{r}_\ell)^2/10$$

on more than $n/2$ of the blocks I_j , and the set of winners \mathcal{H}_ℓ consists of all the “champions” in \mathcal{H}'_ℓ that win all of their home matches.

The outcome of the ℓ -champions league phase is as follows:

Proposition 4.7. *Let \mathcal{H}'_ℓ as above. With probability at least*

$$1 - 2 \exp \left(-c_0 N \min \left\{ 1, \frac{\widehat{r}_\ell^2}{\sigma_4^2} \right\} \right)$$

with respect to $(X_i, Y_i)_{i=2N+1}^{3N}$, the set of winners \mathcal{H}_ℓ contains f^ , and if $h \in \mathcal{H}_\ell$ then*

$$R(h) - R(f^*) \leq 16c^2 \widehat{r}_\ell^2 .$$

Proposition 4.7 is an immediate outcome of Proposition 3.8 from [11] for $H = \mathcal{H}'_\ell$ and using the fact that $f^* \in \mathcal{H}'_\ell$ and that if $h \in \mathcal{H}'_\ell$ then $\|h - f^*\|_{L_2} \leq (\beta/\alpha)\widehat{r}_\ell$.

Combining all these observations, Corollary 4.8 describes the outcome of the first three phases in the regularized tournament procedure, for each member of the hierarchy.

Corollary 4.8. *For $L \geq 1$ there exists a constant c that depends only on L for which the following holds. Using the above notation, if $f^* \in \mathcal{F}_\ell$, then with probability at least*

$$1 - 2 \exp \left(-c_1(L) N \min \left\{ 1, \frac{\widehat{r}_\ell^2}{\sigma_4^2} \right\} \right)$$

with respect to $(X_i, Y_i)_{i=1}^{3N}$, the set of winners \mathcal{H}_ℓ satisfies:

- $f^* \in \mathcal{H}_\ell$, and
- for any $h \in \mathcal{H}_\ell$,

$$\Psi(h - f^*) \leq \rho_\ell, \quad \|h - f^*\|_{L_2} \leq (\beta/\alpha)\widehat{r}_\ell, \quad \text{and} \quad R(h) - R(f^*) \leq 16(\beta/\alpha)^2 \widehat{r}_\ell^2 .$$

Selection of a final winner

Once the first three phases are completed, their outcome is K sets \mathcal{H}_ℓ consisting of stage winners. Of course, some of these sets may be empty. However, for those indices ℓ for which $f^* \in \mathcal{F}_\ell$, with high probability the set \mathcal{H}_ℓ contains at least f^* . In order to select the final “winner”, let ℓ_1 be the largest integer $1 \leq \ell \leq K$ for which $\bigcap_{j \leq \ell} \mathcal{H}_j \neq \emptyset$ and that the procedure returns any $\widehat{h} \in \bigcap_{j \leq \ell_1} \mathcal{H}_j$.

Clearly, on a high-probability event, $\ell_1 \geq \ell^*$. On this event, the selected function \widehat{h} belongs to \mathcal{H}_{ℓ^*} . Moreover, recalling that $(\widehat{r}_\ell)_{\ell=1}^K$ is decreasing,

$$\Psi(\widehat{h} - f^*) \leq \widehat{r}_{\ell^*}, \quad \|\widehat{h} - f^*\|_{L_2} \leq (\beta/\alpha)\widehat{r}_{\ell^*} \quad \text{and} \quad R(\widehat{h}) - R(f^*) \leq c(\beta/\alpha)^2 \widehat{r}_{\ell^*}^2,$$

which completes the proof of Theorem 3.9. ■

5 Examples

In what follows we present two examples: a tournament version of LASSO, and also, a tournament version of another popular sparse recovery procedure—SLOPE.

Here we use a stronger moment assumption than Assumption 3.1 because that allows us to give explicit estimates on the parameters ρ_ℓ and r_ℓ .

Assumption 5.1. *Let X be is an isotropic random vector in \mathbb{R}^d . Assume that there are constant c_1 and C such that for every $t \in \mathbb{R}^d$ and every $1 \leq p \leq c_1 \log d$,*

$$\|\langle t, X \rangle\|_{L_p} \leq C\sqrt{p}\|\langle t, X \rangle\|_{L_2} = C\sqrt{p}\|t\|_2; \tag{5.1}$$

the value of the constant c_1 is given in Theorem 5.2.

Assume further that $\|f^(X) - Y\|_{L_4} \leq \sigma_4$ for a known constant σ_4 .*

In other words, Assumption 5.1 means that linear forms satisfy a sub-Gaussian moment growth, but only up to a rather low exponent—logarithmic in the dimension of the underlying space. This moment assumption is a sufficient and almost necessary condition for the celebrated basis pursuit procedure to have a unique minimizer (see [7]), and as such, it is a natural assumption when studying such sparsity-driven bounds. Note that even with Assumption 5.1 replacing Assumption 3.1, the fact that the ‘noise’ $\xi = f^*(X) - Y$ may only be in L_4 means that there is not hope that

$$\sup_{t \in T} \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \xi_i \langle t, X_i \rangle - \mathbb{E} \xi \langle t, X \rangle \right|,$$

exhibits a fast tail decay, even in the extreme case when $|T| = 1$. This indicates why regularized risk minimization can only perform with a rather weak accuracy/confidence tradeoff in such situations.

On the other hand, (5.1) suffices to obtain bounds on the *expectation* of empirical and multiplier processes, as long as the indexing set has enough symmetries, and a suitable bound on the expectation suffices for the analysis of regularized tournaments.

Definition 5.1. Given a vector $x = (x_i)_{i=1}^n$, let $(x_i^*)_{i=1}^n$ be the non-increasing rearrangement of $(|x_i|)_{i=1}^n$.

The normed space $(\mathbb{R}^d, \|\cdot\|)$ is K -unconditional with respect to the basis $\{e_1, \dots, e_d\}$ if for every $x \in \mathbb{R}^d$ and every permutation of $\{1, \dots, n\}$,

$$\left\| \sum_{i=1}^d x_i e_i \right\| \leq K \left\| \sum_{i=1}^d x_{\pi(i)} e_i \right\|,$$

and if $y \in \mathbb{R}^d$ and $x_i^* \leq y_i^*$ for $1 \leq i \leq d$ then

$$\left\| \sum_{i=1}^d x_i e_i \right\| \leq K \left\| \sum_{i=1}^d y_i e_i \right\|.$$

There are many natural examples of K -unconditional spaces, most notably, all the ℓ_p spaces. Moreover, if $(v_i^*)_{i=1}^d$ denotes the nonincreasing rearrangement of $(|v_i|)_{i=1}^d$, then the norm $\|z\| = \sup_{v \in V} \sum_{i=1}^n v_i^* z_i^*$ is 1-unconditional. In fact, if $V \subset \mathbb{R}^d$ is closed under coordinate permutations and reflections (sign-changes), then $\|\cdot\| = \sup_{v \in V} |\langle \cdot, v \rangle|$ is 1-unconditional in the sense of Definition 5.1.

The following fact has been established in [16]:

Theorem 5.2. *There exists an absolute constant c_1 and for $K \geq 1$, $L \geq 1$ and $q_0 > 2$ there exists a constant c_2 that depends only on K , L and q_0 for which the following holds. Consider*

- $V \subset \mathbb{R}^d$ for which the norm $\|\cdot\| = \sup_{v \in V} |\langle v, \cdot \rangle|$ is K -unconditional with respect to the basis $\{e_1, \dots, e_d\}$,
- $\xi \in L_{q_0}$ for some $q_0 > 2$,
- an isotropic random vector $X \in \mathbb{R}^d$ that satisfies

$$\max_{1 \leq j \leq d} \sup_{1 \leq p \leq c_1 \log d} \frac{\|\langle X, e_j \rangle\|_{L_p}}{\sqrt{p}} \leq L.$$

If $(X_i, \xi_i)_{i=1}^N$ are independent copies of (X, ξ) then

$$\mathbb{E} \sup_{v \in V} \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N (\xi_i \langle X_i, v \rangle - \mathbb{E} \xi \langle X, v \rangle) \right| \leq c_2 \|\xi\|_{L_{q_0}} \ell_*(V),$$

where $\ell_*(V) = \mathbb{E} \sup_{v \in V} \sum_{i=1}^d g_i v_i$ and $G = (g_i)_{i=1}^d$ is a standard Gaussian vector in \mathbb{R}^d .

Therefore, as long as V is sufficiently symmetric and linear forms exhibit a sub-Gaussian moment growth up to $p \sim \log d$, the expectations of empirical and multiplier processes indexed by V behave as if X were the standard Gaussian vector and ξ were independent of X . In the cases we are interested in the indexing sets have enough symmetries, and since $\xi \in L_4$, the conditions of Theorem 5.2 hold for $q_0 = 4$.

5.1 The tournament LASSO

In this section we prove Theorem 2.7, the performance bound of the “tournament LASSO” procedure.

The proof of Theorem 2.7 follows from Theorem 3.9, combined with explicit estimates on the parameters ρ_ℓ , $r_{\ell,1}$ and $r_{\ell,2}$, as we now show.

Note that for any $h = \langle t_0, \cdot \rangle$ and every $\rho, r > 0$,

$$\mathcal{F}_{h,\rho,r} = \{ \langle t, \cdot \rangle \in \mathbb{R}^d : t \in \rho B_1^d \cap r B_2^d \} ,$$

where B_p^d is the unit ball of the normed space $(\mathbb{R}^d, \|\cdot\|_p)$.

Hence, for a fixed radius ρ the parameters r_E and \bar{r}_M are defined using the fixed-point conditions

$$\mathbb{E} \sup_{t \in \rho B_1^d \cap r B_2^d} \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i \langle t, X_i \rangle \right| \leq \kappa \sqrt{N} r \quad (5.2)$$

and

$$\mathbb{E} \sup_{t \in \rho B_1^d \cap r B_2^d} \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i \xi_i \langle t, X_i \rangle \right| \leq \kappa \sqrt{N} r^2 \quad (5.3)$$

respectively. The indexing set $V_{\rho,r} = \rho B_1^d \cap r B_2^d$ is invariant under coordinate permutations and sign reflections, and therefore satisfies the conditions of Theorem 5.2. Hence, an upper bound on r_E follows if

$$\mathbb{E} \sup_{t \in \rho B_1^d \cap r B_2^d} \sum_{i=1}^N g_i t_i \leq \kappa \sqrt{N} r , \quad (5.4)$$

while for an upper estimate on \bar{r}_M it suffices to ensure that

$$\|\xi\|_{L_4} \mathbb{E} \sup_{t \in \rho B_1^d \cap r B_2^d} \sum_{i=1}^N g_i t_i \leq \kappa \sqrt{N} r^2 . \quad (5.5)$$

Equations (5.4) and (5.5) cannot be improved; they are tight bounds on (5.2) and (5.3) when, for example, $X = (g_1, \dots, g_d)$ and ξ is a Gaussian variable that is independent of X .

The added value in (5.4) and (5.5) is that if r satisfies these inequalities then, necessarily, $\max\{\lambda_Q, \lambda_M\} \leq r$ (for a well chosen constant κ). This is an immediate consequence of Sudakov’s inequality, which implies that for some absolute constant $c > 0$, for any $T \subset \mathbb{R}^d$ and any $\varepsilon > 0$,

$$\varepsilon \sqrt{\log \mathcal{M}(T, \varepsilon B_2^d)} \leq c \mathbb{E} \sup_{t \in T} \sum_{i=1}^d g_i t_i \equiv \ell_*(T) .$$

Thus, when applied to the definition on λ_Q one obtains

$$\log \mathcal{M}(\rho B_1^d \cap r B_2^d, \eta r B_2^d) \leq \frac{\ell_*^2(\rho B_1^d \cap r B_2^d)}{(\eta r)^2} \leq \kappa^2 N ,$$

that is, it suffices that

$$\ell_*(\rho B_1^d \cap r B_2^d) \leq \eta \kappa \sqrt{N} r ,$$

which is precisely the type of condition in (5.4).

With those estimates in hand, we are now able to explain the choice of parameters ρ_ℓ , $r_{\ell,1}$, $r_{\ell,2}$ and λ_ℓ for the tournament LASSO. This requires several observations that have been established in [9].

First, the requirement that $\Delta(\rho, r) \geq 4\rho/5$ forces some constraint on the choice of ρ and r . To simplify things, assume that $t_0 = \operatorname{argmin}_{t \in \mathbb{R}^d} \mathbb{E}(\langle X, t \rangle - Y)^2$ is supported on $I \subset \{1, \dots, d\}$ and that $|I| \leq s$. Recall that by the definition of $\Delta(\rho, r)$, the fact that $\Psi(t) = \|t\|_1$ and since X is isotropic, it suffices to consider vectors $t \in \mathbb{R}^d$ for which $\|t - t_0\|_1 = \rho$ and $\|t - t_0\|_2 \leq r$. For such t ,

$$\|t\|_1 - \|t_0\|_1 = \sum_{i \in I^c} |t_i| + \sum_{i \in I} (|t_i| - |(t_0)_i|) \geq \sum_{i \in I^c} |t_i| - \sum_{i \in I} |t_i - (t_0)_i|,$$

and since $|I| \leq s$,

$$\sum_{i \in I} |t_i - (t_0)_i| \leq \sqrt{|I|} \|t - t_0\|_2 \leq \sqrt{sr}.$$

Therefore,

$$\sum_{i \in I^c} |t_i| = \sum_{i \in I^c} |t_i - (t_0)_i| = \sum_{i=1}^n |t_i - (t_0)_i| - \sum_{i \in I} |t_i - (t_0)_i| \geq \rho - \sqrt{sr}.$$

On the other hand, there is a functional z that is norming for both t_0 and $P_{I^c}t = \sum_{i \in I^c} t_i e_i$; hence,

$$\begin{aligned} z(t - t_0) &\geq z(P_{I^c}(t - t_0)) - \sum_{i \in I} |t_i - (t_0)_i| = \sum_{i \in I^c} |t_i - (t_0)_i| - \sum_{i \in I} |t_i - (t_0)_i| \\ &\geq \rho - 2\sqrt{sr} \geq \frac{4\rho}{5} \end{aligned}$$

as long as $s \lesssim (\rho/r)^2$.

This shows that, as long as the ratio ρ/r is larger than the square-root of the degree of sparsity of vectors we are interested in, $\Delta(\rho, r) \geq (4/5)\rho$ as our procedure requires. A similar observation is true if t_0 is not sparse, but rather well approximated by an s -sparse vector (see [9] for a detailed argument).

Set $k = (\rho/r)^2$ and assume without loss of generality that k is an integer. We also restrict ourselves to values $1 \leq k \leq d$, intuitively because the above implies that $(\rho/r)^2$ should capture the degree of sparsity. Recall that

$$\ell_*(\rho B_1^d \cap r B_2^d) = r \ell_*(\sqrt{k} B_1^d \cap B_2^d) \leq Cr \sqrt{k \log(ed/k)} = C\rho \sqrt{\log(edr^2/\rho^2)}$$

(see, e.g. [9] for the standard proof). Hence, (5.4) becomes

$$C\rho \sqrt{\log(edr^2/\rho^2)} \leq \kappa \sqrt{Nr}, \quad (5.6)$$

while (5.5) implies

$$\|\xi\|_{L_4} \cdot C\rho \sqrt{\log(edr^2/\rho^2)} \leq \kappa \sqrt{Nr^2}. \quad (5.7)$$

We consider only the case $N \leq Cd$, which is the more interesting range in sparse recovery—when the number of given linear measurements is significantly smaller than the dimension of the underlying space. An argument following the same path may be used when $N \geq Cd$ and we omit it.

It follows from a rather tedious computation that (5.6) holds provided that

$$r \geq c \frac{\rho}{\kappa \sqrt{N}} \sqrt{\log \left(\frac{cd}{\kappa N} \right)}, \quad (5.8)$$

and it follows from (5.7) that

$$r^2 \geq c \rho \frac{\|\xi\|_{L_4}}{\sqrt{N}} \sqrt{\log \left(c \frac{\|\xi\|_{L_4} d}{\sqrt{N} \rho} \right)} \quad (5.9)$$

as long as $\|\xi\|_{L_4} d / \sqrt{N} \rho \geq c'$.

Using the constraint that $\rho/r \geq c\sqrt{s}$, it is evident from (5.8) that $s \leq c(L)N / \log(ed/N)$, and that

$$\frac{1}{s} \geq c_1(L) \frac{r^2}{\rho^2} \gtrsim \frac{1}{\rho} \cdot \frac{\|\xi\|_{L_4}}{\sqrt{N}} \sqrt{\log \left(e \frac{\|\xi\|_{L_4} d}{\sqrt{N} \rho} \right)}.$$

Therefore, to have a ‘legal’ choice of ρ and r , we must have

$$N \geq c_2(L) s \log \left(\frac{ed}{s} \right),$$

and

$$\rho \geq c_3(L) \frac{s}{\sqrt{N}} \|\xi\|_{L_4} \cdot \sqrt{\log \left(\frac{ed}{s} \right)}.$$

This naturally leads to the choices made in Section 2.1: set

$$\mathcal{F}_\ell = \{t : \exists v, |\text{supp}(v)| \leq d/2^{\ell-1}, \|t - v\|_1 \leq \rho_\ell\}$$

to be the set of vectors that are ‘well-approximated’ by $d/2^{\ell-1}$ sparse vectors. For every ℓ let $s = d/2^{\ell-1}$; if $s \geq c(L)N / \log \left(\frac{ed}{N} \right)$, set $\rho_\ell = r_\ell = \infty$. If the reverse inequality holds, set

$$\rho_\ell = c(L) \frac{d}{2^\ell \sqrt{N}} \|\xi\|_{L_4} \sqrt{\log(e2^\ell)} \sim_L \frac{d\sqrt{\ell}}{2^\ell \sqrt{N}} \|\xi\|_{L_4},$$

and for that choice ρ_ℓ , the required value of r_ℓ is

$$r_\ell \geq c(L) \|\xi\|_{L_4} \sqrt{\frac{s}{N} \log \left(\frac{ed}{s} \right)} \sim_L \|\xi\|_{L_4} \sqrt{\frac{d\ell}{2^\ell N}}.$$

Finally, let $\hat{r}_\ell \geq r_\ell$ and recall that $\lambda_\ell \sim_L \hat{r}_\ell^2 / \rho_\ell$. Applying Theorem 3.9, these choices complete the proof of Theorem 2.7. \blacksquare

5.2 The tournament SLOPE

As a second example, we present and analyze a “tournament” version of the regularized risk minimization procedure SLOPE. SLOPE is defined using a set of non-increasing weights $(\beta_i)_{i=1}^d$. The corresponding norm is

$$\Psi(z) = \sum_{i=1}^d \beta_i z_i^* ,$$

where as always, $(z_i^*)_{i=1}^d$ denotes the non-increasing rearrangement of $(|z_i|)_{i=1}^d$. Clearly, SLOPE is a generalized version of LASSO, as the latter is given by the choice $\beta_i = 1$ for $1 \leq i \leq d$.

Just like the LASSO, most of the known results on the performance of SLOPE hold only when both the random vector X and the target Y have well behaved tails.

The tournament SLOPE we present below is defined for the penalty $\Psi(z) = \sum_{i=1}^d \beta_i z_i^*$, where $\beta_i \leq C\sqrt{\log(ed/i)}$. We obtain the following performance bound:

Theorem 5.3. *For $L \geq 1$ there exist constants c_0, \dots, c_5 that depend only on L and for which the following holds. Let X satisfy Assumption 5.1, set $t_0 = \operatorname{argmin}_{t \in \mathbb{R}^d} \mathbb{E}(Y - \langle t, X \rangle)^2$ and assume that $\|Y - \langle t_0, X \rangle\|_{L_4} \leq \sigma$. Assume further that there is v that is s -sparse such that*

$$\|t_0 - v\|_1 \leq c_1(L)\sigma \cdot \frac{s \log(ed/s)}{\sqrt{N}} .$$

If $N \geq c_2(L)s \log(ed/s)$ and

$$\hat{r} \geq c_3(L)\sigma \sqrt{\frac{s}{N} \log\left(\frac{ed}{s}\right)} ,$$

then with probability at least

$$1 - 2 \exp\left(-c_4(L)N \min\left\{1, \left(\frac{\hat{r}}{\sigma}\right)^2\right\}\right) ,$$

the tournament SLOPE produces \hat{t} that satisfies

$$\|\hat{t} - t_0\|_2 \leq c_5(L)\sigma \sqrt{\frac{s}{N} \log\left(\frac{ed}{s}\right)} ,$$

$$\Psi(\hat{t} - t_0) = \sum_{i=1}^d (\hat{t} - t_0)_i^* \sqrt{\log(ed/i)} \leq c_5(L)\sigma \frac{s}{\sqrt{N}} \log\left(\frac{ed}{s}\right) ,$$

and

$$\mathbb{E}\left(\langle \hat{t}, X \rangle - Y\right)^2 | (X_i, Y_i)_{i=1}^N \leq \mathbb{E}(\langle t_0, X \rangle - Y)^2 + c_5(L)\hat{r}^2 .$$

The estimate corresponds to the optimal accuracy/confidence tradeoff any procedure can attain even if the learner knows that t_0 is s -sparse. Moreover, in the heavy-tailed situations we study here, the performance of SLOPE is significantly weaker than in Theorem 5.3.

The argument we use here is similar to the one used for the tournament LASSO, and so we skip most of the details.

In tournament SLOPE one selects $\beta_i \leq c_0 \sqrt{\log(ed/i)}$ and therefore the corresponding indexing set is contained in

$$V_{\rho,r} = \rho \mathcal{B} \cap r B_2^d = \left\{ v \in \mathbb{R}^d : \|v\|_2 \leq r \text{ and } \sum_{i=1}^d v_i^* \sqrt{\log(ed/i)} \leq \rho/c_0 \right\}.$$

Because $V_{\rho,r}$ has enough symmetries, one may apply Theorem 5.2, leading to an upper bound on r_E when

$$\mathbb{E} \sup_{v \in V_{\rho,r}} \sum_{i=1}^d g_i v_i \leq \kappa \sqrt{N} r. \quad (5.10)$$

Also, to estimate r_M it suffices to ensure that

$$\|\xi\|_{L_4} \mathbb{E} \sup_{v \in V_{\rho,r}} \sum_{i=1}^d g_i v_i \leq \kappa \sqrt{N} r^2. \quad (5.11)$$

Next, one may verify (see Lemma 4.3 in [9]) that if we set $B_s = \sum_{i \leq s} \beta_i / \sqrt{i}$ and if $B_s \lesssim r/\rho$, then $\Delta(\rho, r) \geq (4/5)\rho$ for centres that are ‘well approximated’ by s -sparse vectors. Also, for our choice of β_i , $B_s \lesssim C \sqrt{s \log(ed/s)}$. Hence, for a fixed degree of sparsity $1 \leq s \leq d$, one has the constraint that

$$\frac{r}{\rho} \geq C_1 \sqrt{s \log(ed/s)} \quad (5.12)$$

for a constant C_1 that depends only on c_0 .

Following the same path used for the tournament LASSO let

$$\mathcal{F}_\ell = \{t : \exists v, |\text{supp}(v)| \leq d/2^{\ell-1}, \Psi(t - v) \leq \rho_\ell\}.$$

For every $1 \leq \ell \leq \log_2 K$ let $s = d/2^{\ell-1}$. There is a nontrivial choice of ρ and r only when $s \lesssim_L N/\log(ed/N)$; otherwise, $\rho = r = \infty$ as one would expect. When $s \lesssim_L N/\log(ed/N)$, we follow the computation in [9] and set

$$\rho_\ell \sim_L \|\xi\|_{L_4} \frac{s}{\sqrt{N}} \log\left(\frac{ed}{s}\right) \sim_L \|\xi\|_{L_4} \frac{d\ell}{2^\ell \sqrt{N}},$$

and

$$r_\ell \sim_L \|\xi\|_{L_4} \sqrt{\frac{s}{N} \log\left(\frac{ed}{s}\right)}.$$

Finally, fix $\widehat{r}_\ell \geq r_\ell$ and set $r_{\ell,1}, r_{\ell,2} \sim_L \widehat{r}_\ell$, $\lambda_\ell \sim_L \widehat{r}_\ell^2/\rho$. Applying Theorem 3.9 for these choices completes the proof of Theorem 5.3. \blacksquare

A Additional proofs

The proofs of Lemma 4.5 and Lemma 4.6 are, in fact, the same as in [9] and [11], respectively. The minor modifications to the original proofs are presented in this appendix solely for the sake of completeness and not in full detail.

Proof of Lemma 4.5

The proof of Lemma 4.5 follows the same path as that of Theorem 3.2 in [9]. Let us begin by examining

$$(*) = \mathbb{Q}_{f,f^*}(j) + \mathbb{M}_{f,f^*}(j) + \lambda(\Psi(f) - \Psi(f^*))$$

in the set $\{f \in \mathcal{F} : \Psi(f - f^*) = \rho\}$. If $\Psi(f - f^*) = \rho$ one should consider two cases. First, if $\|f - f^*\|_{L_2} \geq r$ then by the triangle inequality for Ψ , and since $\mathbb{Q}_{f,f^*}(j) \geq C\|f - f^*\|_{L_2}^2$ and $\mathbb{M}_{f,f^*}(j) \geq -(C/4)\|f - f^*\|_{L_2}^2$, we have

$$\begin{aligned} (*) &\geq C\|f - f^*\|_{L_2}^2 - \frac{C}{4}\|f - f^*\|_{L_2}^2 - \lambda\Psi(f - f^*) \\ &\geq \frac{3C}{4}\|f - f^*\|_{L_2}^2 - \lambda\rho \geq \frac{3C}{4}r^2 - \lambda\rho > 0, \end{aligned} \quad (\text{A.1})$$

provided that

$$\lambda \leq \frac{3C}{4} \cdot \frac{r^2}{\rho}. \quad (\text{A.2})$$

If, on the other hand, $\|f - f^*\|_{L_2} \leq r$, then $\mathbb{Q}_{f,f^*}(j) \geq 0$ and $\mathbb{M}_{f,f^*}(j) \geq -(C/4)r^2$. Therefore,

$$(*) \geq -\frac{C}{4}r^2 + \lambda(\Psi(f) - \Psi(f^*)).$$

Fix $v \in \mathcal{B}_{f^*}(\rho/20)$ and write $f^* = u + v$; thus $\Psi(u) \leq \rho/20$. Set z to be a linear functional that is norming for v and observe that for any $f \in E$,

$$\begin{aligned} \Psi(f) - \Psi(f^*) &\geq \Psi(f) - \Psi(v) - \Psi(u) \geq z(f - v) - \Psi(u) \geq z(f - f^*) - 2\Psi(u) \\ &\geq z(f - f^*) - \frac{\rho}{10}. \end{aligned} \quad (\text{A.3})$$

Hence, if $f^* \in \mathcal{F}_\ell$ and $f \in \mathcal{F} \cap \mathcal{B}_{f^*}(\rho) \cap D_{f^*}(r)$ then optimizing the choices of v and of z , $z(f - f^*) \geq \Delta_\ell(\rho, r)$; thus

$$\Psi(f) - \Psi(f^*) \geq \Delta_\ell(\rho, r) - \frac{\rho}{10} \geq \frac{7}{10}\rho. \quad (\text{A.4})$$

And, if

$$\lambda \geq \frac{C}{2} \cdot \frac{r^2}{\rho}, \quad (\text{A.5})$$

we have that

$$(*) \geq -\frac{C}{4}r^2 + \lambda \cdot \frac{7}{10}\rho > 0.$$

In other words, if λ is chosen to satisfy both (A.2) and (A.5), $f \in \mathcal{F}$ and $\Psi(f - f^*) = \rho$, it follows that

$$\mathbb{Q}_{f,f^*}(j) + \mathbb{M}_{f,f^*}(j) + \lambda(\Psi(f) - \Psi(f^*)) > 0.$$

Next, if $\Psi(f - f^*) > \rho$, there are $\theta \in (0, 1)$ and $h \in \mathcal{F}$ that satisfy

$$\Psi(h - f^*) = \rho \quad \text{and} \quad \theta(f - f^*) = h - f^*.$$

If $\|h - f^*\|_{L_2} \geq r$, then by the triangle inequality for Ψ followed by (A.1),

$$\begin{aligned} (*) &\geq \frac{1}{\theta^2} \mathbb{Q}_{h,f^*}(j) + \frac{1}{\theta} (\mathbb{M}_{h,f^*}(j) - \lambda \Psi(h - f^*)) \\ &\geq \frac{1}{\theta} (\mathbb{Q}_{h,f^*}(j) + \mathbb{M}_{h,f^*}(j) - \lambda \Psi(h - f^*)) > 0. \end{aligned}$$

If, on the other hand, $\|h - f^*\|_{L_2} \leq r$, then

$$\begin{aligned} (*) &\geq \frac{1}{\theta} \mathbb{M}_{h,f^*}(j) + \lambda(z(f - f^*) - 2\Psi(u)) \\ &\geq \frac{1}{\theta} (\mathbb{M}_{h,f^*}(j) + \lambda(z(h - f^*) - 2\theta\Psi(u))) \\ &\geq \frac{1}{\theta} (\mathbb{M}_{h,f^*}(j) + \lambda(z(h - f^*) - 2\Psi(u))) > 0, \end{aligned}$$

because $0 \leq \theta < 1$ and using (A.3).

Now, all that remains is to control $f \in \mathcal{F} \cap \mathcal{B}_{f^*}(\rho)$ and show that if $\|f - f^*\|_{L_2} \geq r$, then

$$\mathbb{Q}_{f,f^*}(j) + \mathbb{M}_{f,f^*}(j) + \lambda(\Psi(f) - \Psi(f^*)) > 0.$$

This follows from (A.1). ■

Proof of Lemma 4.6

The first part of Lemma 4.6 is identical to Lemma 5.1 from [11], with the trivial modification that the constant $-C/4$ replaces $-3C/4$ used in [11]. The second part of Lemma 4.6 was not needed in [11], but its proof follows the same path as Lemma 5.1 from [11].

Set $r = \hat{r}_\ell$ and fix $f \in \mathcal{F}$ that satisfies $\|f - f^*\|_{L_2} \leq r$. Recall that $m = N/n$ and that $\sqrt{n/N} \leq \sqrt{\theta}r/\sigma_4$ for a well-chosen constant θ that depends only on L and τ . Set $U = (f - f^*)(X) \cdot (f^*(X) - Y)$ and observe that

$$\mathbb{M}_{f,f^*} = \frac{1}{m} \sum_{i=1}^m U_i.$$

It follows from the convexity of \mathcal{F} that $\mathbb{E}U = \mathbb{E}(f - f^*)(X) \cdot (f^*(X) - Y) \geq 0$; therefore, $\mathbb{M}_{f,f^*} \geq \mathbb{M}_{f,f^*} - \mathbb{E}\mathbb{M}_{f,f^*}$. Also,

$$Pr(|\mathbb{M}_{f,f^*} - \mathbb{E}\mathbb{M}_{f,f^*}| > t) \leq t^{-1} \mathbb{E}|\mathbb{M}_{f,f^*} - \mathbb{E}\mathbb{M}_{f,f^*}|,$$

and by a straightforward symmetrization argument,

$$\mathbb{E}|\mathbb{M}_{f,f^*} - \mathbb{E}\mathbb{M}_{f,f^*}| \leq 2\mathbb{E} \left| \frac{1}{m} \sum_{i=1}^m \varepsilon_i U_i \right| \leq \frac{2}{\sqrt{m}} (\mathbb{E}|U|^2)^{1/2}.$$

Applying Assumption 3.1, it is evident that

$$(\mathbb{E}|U|^2)^{1/2} \leq \|f^*(X) - Y\|_{L_4} \cdot \|f - f^*\|_{L_4} \leq Lr\sigma_4,$$

and thus

$$Pr(|\mathbb{M}_{f,f^*} - \mathbb{E}\mathbb{M}_{f,f^*}| > t) \leq \frac{2Lr\sigma_4}{t\sqrt{m}} = \frac{2L\sigma_4 r \sqrt{n}}{t\sqrt{N}} \leq \frac{\tau}{3},$$

where we use the fact that $\sqrt{n/N} \leq \sqrt{\theta}r/\sigma_4$ and select $t = Cr^2/8$ and $\theta = \theta(\tau, L)$. Therefore,

$$Pr(\mathbb{M}_{f,f^*} \leq -(C/8)r^2) \leq \frac{\tau}{3},$$

and with probability at least $1 - 2\exp(-c\tau^2n)$,

$$|\{j : \mathbb{M}_{f,f^*}(j) \geq -(C/8)r^2\}| \geq (1 - \tau/2)n. \quad (\text{A.6})$$

The rest of the argument is identical to the proof of Lemma 5.1 from [11]: let \mathcal{H} be a maximal separated subset of $\mathcal{F} \cap \mathcal{B}_{f^*}(\rho) \cap D_{f^*}(r)$ with respect to the L_2 norm, of cardinality $\exp(c\tau^2n/2)$, and with the following property: for any $f \in \mathcal{F} \cap \mathcal{B}_{f^*}(\rho) \cap D_{f^*}(r)$ there is $h \in \mathcal{H}$ for which

$$\|f - h\|_{L_2} \leq \varepsilon \quad \text{and} \quad \mathbb{E}(f^*(X) - Y)(f(X) - h(X)) \geq 0; \quad (\text{A.7})$$

here ε denotes the mesh width of the net. The existence of such a separated set is established in [11] (see Lemma 5.3), and one may show that ε is a small proportion of r .

By (A.6), we have that with probability at least $1 - 2\exp(-c\tau^2n/2)$, for every $h \in \mathcal{H}$

$$|\{j : \mathbb{M}_{h,f^*}(j) \geq -(C/8)r^2\}| \geq (1 - \tau/2)n. \quad (\text{A.8})$$

For every $f \in \mathcal{F} \cap \mathcal{B}_{f^*}(\rho) \cap D_{f^*}(r)$ let $\pi f \in \mathcal{H}$ be as in (A.7), and at the heart of the proof of Lemma 5.4 in [11] is that with probability at least $1 - 2\exp(-c_1\tau^2n)$,

$$\sup_{f \in \mathcal{F} \cap \mathcal{B}_{f^*}(\rho) \cap D_{f^*}(r)} |\{j : \mathbb{M}_{f,f^*}(j) - \mathbb{M}_{\pi f,f^*}(j) \leq -(C/8)r^2\}| \leq \frac{\tau n}{2}. \quad (\text{A.9})$$

Combining (A.8) and (A.9), there is an event of probability at least $1 - 2\exp(-c_2\tau^2n)$ on which for any $f \in \mathcal{F} \cap \mathcal{B}_{f^*}(\rho) \cap D_{f^*}(r)$ there is a set of coordinate blocks $(I_j)_{j \in J}$, of cardinality $|J| \geq (1 - \tau)n$ and for $j \in J$,

$$\mathbb{M}_{f,f^*}(j) \geq \mathbb{M}_{\pi f,f^*}(j) + (\mathbb{M}_{f,f^*}(j) - \mathbb{M}_{\pi f,f^*}(j)) \geq -\frac{C}{4}r^2. \quad \blacksquare$$

Acknowledgements

We thank the referees for valuable suggestions that helped us improve the presentation significantly.

References

- [1] J.-Y. Audibert and O. Catoni. Robust linear least squares regression. *The Annals of Statistics*, 39:2766–2794, 2011.
- [2] P. Bellec, G. Lecué, and A. Tsybakov. Slope meets lasso: improved oracle bounds and optimality. *arXiv preprint arXiv:1605.08651*, 2016.
- [3] C. Brownlees, E. Joly, and G. Lugosi. Empirical risk minimization for heavy-tailed losses. *Annals of Statistics*, 43:2507–2536, 2015.

- [4] A. Goldenshluger and A. Nemirovski. On spatially adaptive estimation of nonparametric regression. *Math. Methods Statist.*, 6(2):135–170, 1997.
- [5] D. Hsu and S. Sabato. Approximate loss minimization with heavy tails. *Computing Research Repository*, abs/1307.1827, 2013.
- [6] G. Lecué and M. Lerasle. Learning from MOM’s principles. *arXiv preprint arXiv:1701.01961*, 2017.
- [7] G. Lecué and S. Mendelson. Sparse recovery under weak moment assumptions. *Journal of the European Mathematical Society*, 19(3):881–904, 2017.
- [8] G. Lecué and S. Mendelson. Learning subgaussian classes: Upper and minimax bounds. In S. Boucheron and N. Vayatis, editors, *Topics in Learning Theory*. Societe Mathematique de France, 2018, to appear.
- [9] Guillaume Lecué and Shahar Mendelson. Regularization and the small-ball method I: Sparse recovery. *Ann. Statist.*, 46(2):611–641, 2018.
- [10] M. Lerasle and R.I. Oliveira. Robust empirical mean estimators. *manuscript*, 2012.
- [11] G. Lugosi and S. Mendelson. Risk minimization by median-of-means tournaments. *Journal of the European Mathematical Society*, 2018, to appear.
- [12] G. Lugosi and S. Mendelson. Sub-gaussian estimators of the mean of a random vector. *Annals of Statistics*, 2018, to appear.
- [13] S. Mendelson. Learning without concentration. *Journal of the ACM*, 62:21, 2015.
- [14] S. Mendelson. Local vs. global parameters—breaking the Gaussian complexity barrier. *Annals of Statistics*, 45(5):1835–1862, 2017.
- [15] S. Mendelson. On aggregation for heavy-tailed classes. *Probability Theory and Related Fields*, 168(3):641–674, 2017.
- [16] S. Mendelson. On multiplier processes under weak moment assumptions. *Geometric Aspects of Functional Analysis, Lecture notes in Mathematics*, 2169:301–318, 2017.
- [17] S. Mendelson. An optimal unrestricted learning procedure. *arXiv preprint arXiv:1707.05342, v2*, 2017.
- [18] S. Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21:2308–2335, 2015.
- [19] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B.*, 58:267–288, 1996.