

Minimax regularization

Raphaël Deswarte^{1,3} and Guillaume Lecué^{2,4,5}

Abstract

Classical approach to regularization is to design norms enhancing smoothness or sparsity and then to use this norm or some power of this norm as a regularization function. The choice of the regularization function (for instance a power function) in terms of the norm is mostly dictated by computational purpose rather than theoretical considerations.

In this work, we design regularization functions that are motivated by theoretical arguments. To that end we introduce a concept of optimal regularization called “minimax regularization” and, as a proof of concept, we show how to construct such a regularization function for the ℓ_1^d norm for the random design setup. We develop a similar construction for the deterministic design setup. It appears that the resulting regularized procedures are different from the one used in the LASSO in both setups.

Keywords: Regularization, Gaussian mean width, empirical processes, regularization.

1 Introduction

Let (\mathcal{X}, μ) be a probability space and (X, Y) be a couple of random variables, in which X is distributed according to μ . One is given a sample of N independent couples $(X_i, Y_i)_{i=1..N}$ distributed according to the joint law of (X, Y) . On the basis of this sample, one tries to link X and Y by a random mapping \hat{f} with $\hat{f}(X)$ close (in L_2) to Y . This is the classical problem, in learning theory, of the prediction of an output Y from an input X given a i.i.d. copies of the couple (X, Y) .

To that end, one is given a class F of functions from \mathcal{X} to \mathbb{R} and the aim in learning theory is to mimic the *best* element in F for the prediction of Y by a function of X in F . We assume that F is closed and convex in $L_2(\mu)$ so that it exists a function f^* that minimizes the square loss in F :

$$f^* \in \operatorname{argmin}_{f \in F} \mathbb{E} (Y - f(X))^2. \quad (1.1)$$

This function is usually called the oracle (cf. [30]); it is the closest function in F to Y in L_2 . Now, the goal is to construct an estimator \hat{f} whose $L_2(\mu)$ distance to f^* is as small as possible using the dataset $\{(X_i, Y_i) : i = 1, \dots, N\}$. In the framework considered in this

¹CMAP. Ecole Polytechnique. 91120 Palaiseau. France

²(corresponding author) CNRS, CREST, ENSAE, 5 avenue Henry Le Chatelier. 91120 Palaiseau. France.

³Email: raphael.deswarte@polytechnique.edu

⁴Email: guillaume.lecue@ensae.fr

⁵Supported by the French National Research Agency (ANR) under the grant Labex Ecodec (ANR-11-LABEX-0047) and by the Maison des sciences de l’Homme under the grant 17MA01

paper, the excess risk of \hat{f} , which is the difference $\mathbb{E}(Y - \hat{f}(X))^2 - \mathbb{E}(Y - f^*(X))^2$, is actually equal to $\|\hat{f} - f^*\|_{L_2(\mu)}^2$ and so estimating f^* is equivalent to predicting Y ; thus we fall back on the original prediction problem by estimating f^* in $L_2(\mu)$.

One may therefore try to bound the quadratic error $\|f - f^*\|_{L_2(\mu)}$ either *in expectation* or *in deviation* with respect to the sample. In this work, we obtain upper bounds on the quadratic error that are valid in deviation, showing that the results are true for “most” samples rather than in average.

Given that we want to be close to a function f^* minimizing $f \rightarrow \mathbb{E}(Y - f(X))^2$ over F , a natural candidate for this problem is the Empirical Risk Minimizer (ERM) also known as the “least squares estimator”:

$$\hat{f}_{ERM} \in \operatorname{argmin}_{f \in F} \frac{1}{N} \sum_{i=1}^N (Y_i - f(X_i))^2. \quad (1.2)$$

Many works have been carried out for general classes (see, [20, 27, 40, 44]) or on the vectorial case (see [36] for the famous Stein paradox, and [14] for elements about the admissibility of the ERM).

It appears that when F is too large (for instance the whole $L_2(\mu)$ space), the ERM tends to “overfit”. The understanding of this phenomenon has led to the introduction of “regularization methods” which were originally used to smooth estimators in order to overcome the “overfitting phenomena”. Those procedures are nowadays used beyond their smoothing effect and, in particular, they are now extensively used in Statistics and learning theory for their “low-dimensional / sparsity inducing properties”. At a high level description, those methods make a trade-off between an “adequation to the data term” and a “regularization term” and their general form (for the quadratic loss) is

$$\hat{f} \in \operatorname{argmin}_{f \in F} \left(\frac{1}{N} \sum_{i=1}^N (Y_i - f(X_i))^2 + \Psi(f) \right) \quad (1.3)$$

where Ψ is a function usually called the *regularization function*.

The “adequation to the data term” can be constructed from any loss function; for the case of the quadratic loss, this term reads like $N^{-1} \sum_{i=1}^N (Y_i - f(X_i))^2$.

As for the regularization term Ψ , several choices are possible, enabling to smooth the estimator, or to force a low-dimensional structure. It depends thus on the a priori knowledge one has on the data (and in particular on f^*), and on computational issues.

A first option is the Tikhonov / ridge regularization:

$$\hat{f}_\lambda \in \operatorname{argmin}_{f \in \mathcal{H}} \left(\frac{1}{N} \sum_{i=1}^N (Y_i - f(X_i))^2 + \lambda \|f\|_{\mathcal{H}}^2 \right) \quad (1.4)$$

where $\|\cdot\|_{\mathcal{H}}$ is a Hilbert norm. One expects in this case that $\|f\|_{\mathcal{H}}$ reflects the smoothness of f (for instance $\|f\|_{\mathcal{H}} = f(0) + (\int_{\mathbb{R}} |f'(t)|^2 dt)^{1/2}$) and that the oracle function f^* has a small $\|f^*\|_{\mathcal{H}}$ norm.

In the finite but high dimensional vectorial case, one often wishes the estimator to be sparse, i.e., to have few non-zero components (for some well designed basis, cf. [26]). This

may come from the fact that the vector to be estimated is known in advance to be sparse; or that, in high-dimensional problems, it is computationally important not to have to manage a huge amount of non-zero coefficients.

Then, a natural way to address this question is to use a sparsity-inducing penalization: like the number of non-zero components of the vector, sometimes called the “ ℓ_0 -norm”. Even though it is theoretically appealing (cf. [18]), it proves to be computationally intractable (actually NP-Hard, in general, cf. [29]). But for geometric reasons, another regularization is efficient to induce sparsity: the ℓ_1 norm (which can be seen as the convex relaxation of the ℓ_0 norm on the unit ℓ_∞ -ball). The associated estimator is called the LASSO (“Least Absolute Shrinkage and Selection Operator”, [38]):

$$\hat{t}_\lambda \in \operatorname{argmin}_{t \in \mathbb{R}^d} \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2 + \lambda \|t\|_1 \right). \quad (1.5)$$

In [9], it is emphasized that Lasso leads generally to sparsity, though giving some counter-examples in which it doesn’t work well –in particular Lasso struggles when dealing with a data matrix with high correlations among the columns. To tackle this kind of issues, it is possible to “mix” regularizations: it is the principle of the “Elastic net method” ([47]), which penalty is a combination of the ℓ_1 and ℓ_2 norms.

All these methods rely on the choice of one (even two for the Elastic net) regularization parameter λ , fixed by the statistician on the basis of empirical methods such as cross-validation. It has to be chosen wisely in order to make the right trade-off between the two terms and thus to minimize the rate of convergence of the regularization procedure towards the oracle.

1.1 Regularization and Model Selection

At this point, it should be clear to the reader that choosing (or even designing) a specific regularization norm like $\|\cdot\|_{\mathcal{H}}$ or $\|\cdot\|_1$ depends only on an a priori knowledge we have. But once this choice has been made, why would someone use the square of this norm in one case (as for the Tikhonov / ridge regularization), or the norm itself (as for the LASSO) or some other power of this norm (cf. [33] for some examples) in other situations? In many cases, this choice is only made following some computational considerations.

In this work, we want to support the choice of regularization functions (given some norm) on theoretical arguments. To that end we will rely on the model selection theory and, in particular, on a key principle in model selection which is to design penalty functions that capture the “complexity” of a model in the most accurate way.

The right calibration of penalty functions has opened an important stream of researches since the work [2]. It has led many researchers to (re)think about the notion of complexity in statistics. In a nutshell, there are mainly three types of quantities that have been introduced to measure the statistical complexity of a statistical model: combinatorial (like the VC dimension, [45]), metric (like the entropy) and random (like Gaussian mean width and Rademacher complexities, [3, 19]). Penalty calibration has culminated with the notion of “minimal penalty” that are sharp penalty functions with exact constants (cf. for instance [7]) thanks to second order term analysis of the notion of complexity of a model.

In the present work, we want to put forward the idea that the “right” choice (from a theoretical point of view) of a regularization function may also follow from a careful study of the complexity of a specific family of models. To that end, the first argument is to look at regularization as a model selection problem for which one has to design a sharp penalty. This has been done for instance in Chapter 3.7 in [21]. In the particular case where one is given a norm $\|\cdot\|$ for regularization, then the associated regularized ERM is a penalized estimator associated to the sequence of embedded models $(m_r)_{r \geq 0}$ where for all $r \geq 0$, $m_r = \{f \in F : \|f\| \leq r\}$ and the right way to regularize is given by a function $\text{reg} : f \in F \rightarrow \text{pen}(m_{\|f\|})$ where $m_{\|f\|}$ is the smallest model in $(m_r)_{r \geq 0}$ containing f , see [7]. This idea is a baseline of this work.

Before diving into further details about the way we suggest to construct a regularization function, let us precise what we expect from a good procedure, in particular how we evaluate that a regularization function is the “right” one, at least from a theoretical point of view. We therefore need to introduce a concept of optimality for regularized estimators. Once again we rely on the basics of model selection theory.

Model selection procedures have been used originally to construct adaptive estimators. For the model selection problem we want to solve, this adaptivity problem reads like selecting the smallest model in the family $(\{f \in F : \|f\| \leq r\})_{r \geq 0}$ containing f^* which is obviously $\{f \in F : \|f\| \leq \|f^*\|\}$. Therefore, the adaptation problem we want to solve here is to construct a procedure which performance is as good as if we had been given the value $\|f^*\|$ in advance. In particular, an estimator achieving the minimax rate of convergence over the model $\{f \in F : \|f\| \leq \|f^*\|\}$ would solve this adaptation problem. In what follows, we design regularization functions in order to meet this requirement but before that we clarify the notion of minimax rate over a model for the type of deviation results we prove below.

To simplify the exposition, we will focus on a specific, though very classical and widely-used, framework: the vectorial case, i.e. when $F = \{\langle \cdot, t \rangle : t \in T\}$ is a class of linear functionals from \mathbb{R}^d to \mathbb{R} indexed by some subset $T \subset \mathbb{R}^d$, with Gaussian design, and Gaussian noise (with known variance σ^2).

Definition 1.1 *Let $T \subset \mathbb{R}^d$, X denote a standard Gaussian vector in \mathbb{R}^d and ξ be a centered real-valued Gaussian random variable with variance σ^2 , independent of X . For all $t^* \in T$, define the random variable $Y^{t^*} = \langle X, t^* \rangle + \xi$ and denote by $\mathcal{Y}^T := \{Y^{t^*} : t^* \in T\}$ the set of all such random variables.*

*Let \hat{t}_N be a statistics from $(\mathbb{R}^d \times \mathbb{R})^N$ to \mathbb{R}^d . Let $0 < \delta_N < 1$ and $\zeta_N > 0$. We say that \hat{t}_N **performs with accuracy ζ_N and confidence $1 - \delta_N$ relative to the set of targets \mathcal{Y}^T** , if for all $Y \in \mathcal{Y}^T$, with probability, w.r.t. to a sample $\mathcal{D} := \{(X_i, Y_i) : i = 1, \dots, N\}$ of i.i.d. copies of (X, Y) , at least $1 - \delta_N$, $\|\hat{t}_N - t^*\|_2^2 \leq \zeta_N$.*

*We say that \mathcal{R}_N is a **minimax rate of convergence over T for the confidence $1 - \delta_N$** if the two following statements hold:*

1. *there exists a statistics \hat{t}_N which performs with accuracy \mathcal{R}_N and confidence $1 - \delta_N$ relative to the set of targets \mathcal{Y}^T*
2. *there exists an absolute constant $g_0 > 0$ such that if \tilde{t}_N is a statistics which attains an accuracy ζ_N with confidence $1 - \delta_N$ relative to the set of targets \mathcal{Y}^T then necessarily $\zeta_N \geq g_0 \mathcal{R}_N$.*

In the following (cf. Theorem 1.3 below), we recall a result from [22] on the minimax rate of convergence over $T = \rho B_1^d$, the unit ℓ_1^d -ball of radius $\rho \geq 0$, for a constant confidence (i.e., for instance, $\delta_N = 1/4$). Note that classical minimax rates of convergence are usually given in expectation (cf. for instance [39]). The main difference here with Definition 1.1 is that it is given for deviation results: the minimax rate \mathcal{R}_N may depend on the confidence parameter δ_N (cf. [22]).

In the present work, we are interested in procedures achieving the minimax rate of convergence over the model $\{t \in \mathbb{R}^d : \|t\| \leq \|t^*\|\}$. This provides a natural way to introduce a notion of optimality for the problem of designing regularization functions.

Definition 1.2 *Let $\|\cdot\|$ be a norm on \mathbb{R}^d , $0 < \delta_N < 1$ and $T \subset \mathbb{R}^d$. Let us consider the following RERM for some function $\Psi : \mathbb{R}_+ \rightarrow \mathbb{R}$:*

$$\hat{t} \in \operatorname{argmin}_{t \in \mathbb{R}^d} \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2 + \Psi(\|t\|) \right)$$

constructed from a sample $\mathcal{D} := \{(X_i, Y_i) : i = 1, \dots, N\}$ of i.i.d. copies of (X, Y^{t^}) where $Y^{t^*} = \langle X, t^* \rangle + \xi$ with $X \sim \mathcal{N}(0, I_{d \times d})$, $\xi \sim \mathcal{N}(0, \sigma^2)$ is independent of X and $t^* \in \mathbb{R}^d$. We say that Ψ is a **minimax regularization function for the norm $\|\cdot\|$ and the confidence $1 - \delta_N$ over T** , if there exists an absolute constant $g_1 > 0$ such that for all $t^* \in T$, the RERM \hat{t} is such that with \mathbb{P}_{t^*} -probability at least $1 - \delta_N$, $\|\hat{t} - t^*\|_2^2 \leq g_1 \mathcal{R}_N^{\|t^*\|_1}$, where $\mathcal{R}_N^{\|t^*\|_1}$ is the minimax rate of convergence over $\{t \in \mathbb{R}^d : \|t\| \leq \|t^*\|\}$ and \mathbb{P}_{t^*} denotes the probability distribution of a N sample of i.i.d. copies of (X, Y^{t^*}) .*

The aim of this work is to show that one can design minimax regularization functions by finding the right notion of complexity of the sequence of embedded models $(\{t \in \mathbb{R}^d : \|t\| \leq r\})_{r \geq 0}$. Note however that there should be some situations where designing such an optimal regularization function would be impossible at some given confidence parameter δ_N . In particular, such a situation should happen when the Empirical risk minimization (ERM) procedure over the “true model” $\{t \in \mathbb{R}^d : \|t\| \leq \|t^*\|\}$ is not itself a minimax procedure over the model $\{t \in \mathbb{R}^d : \|t\| \leq \|t^*\|\}$. This happens for constant confidence bound (for instance, when $\delta_N = 1/4$) when there is a gap in Sudakov inequality (cf. [22] for more details). Nevertheless, in [22], it is proved that for high confidence bounds (that is when δ_N decays exponentially fast with the complexity of the model) ERM is always minimax over convex classes. It appears that for the case of ℓ_1^d -balls ERM is minimax for all confidence regime therefore this subtlety will not show up in this special case.

1.2 General approach provided in this paper

Let us now present our approach. As we mentioned before, we want to construct a regularization function depending on the complexity of the models $\{t \in \mathbb{R}^d : \|t\| \leq r\}$ for all $r \geq 0$. This leads to choose a RERM (in the vectorial case) having the following form:

$$\hat{t} \in \operatorname{argmin}_{t \in \mathbb{R}^d} \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2 + \operatorname{comp}(\|t\| B_{\|\cdot\|}) \right) \quad (1.6)$$

where $B_{\|\cdot\|}$ is the unit ball associated with the given regularization norm $\|\cdot\|$ and for all $t \in \mathbb{R}^d$, $\|t\|B_{\|\cdot\|} = \{u \in \mathbb{R}^d : \|u\| \leq \|t\|\}$. The key feature in (1.6) is the “complexity function” $r \geq 0 \rightarrow \text{comp}(rB_{\|\cdot\|})$ which aims at measuring with the best possible accuracy the complexity of the models $rB_{\|\cdot\|}$ for all $r \geq 0$ from a statistical point of view. Of course, finding the right notion of complexity is paramount in this approach.

To aim at optimality, in this paper we advocate complexities that are tailored for the specific statistical problem at stake. It turns out that the “right” choice of complexity, and hence, of regularization, is linked to the behavior of two empirical processes. Those two empirical processes are ultimately connected to the two sources of statistical complexities in the considered problem. When estimating t^* from the data $(X_i, Y_i)_{i=1}^N$ there are two statistical issues : 1)(an inverse problem) t^* is observed only through \mathbb{X} , where $\mathbb{X} \in \mathbb{R}^{N \times d}$ is the operator whose rows vectors are given by the X_i 's; 2)(noisy data) the observations have been corrupted by some noise ξ . The action of the operator \mathbb{X} on the models $rB_{\|\cdot\|}$ for all $r \geq 0$ plays a prominent role in our analysis. In particular, the size of the intersection of its kernel with the model is a natural minmax lower bound for any estimator since any two vectors in the kernel of \mathbb{X} and the model are indistinguishable. The effect of the “distortion” of the operator \mathbb{X} does not show up for small models (i.e. small values of r) because of the presence of the noise which blurs everything at small scales. But passing beyond some threshold for the signal-to-noise ratio r/σ , only the distortion of \mathbb{X} matters from a statistical point of view. This phenomenon occurs only when $N \lesssim d$, because that is the regime where \mathbb{X} has a non trivial kernel. On the contrary, in the low dimensional setup $d \lesssim N$, \mathbb{X} is well conditioned with high probability and therefore, there is no statistical complexity coming from the distortion of \mathbb{X} since in that regime there is no such distortion. Controlling the distortion of \mathbb{X} is a key issue in high-dimensional statistics. It is behind all classical properties like RIP (cf. [12]) or REC (cf. [6]) and it will play equally a key role in our analysis. In particular, the ℓ_2^d diameter of the intersection of \mathbb{X} with the model $rB_{\|\cdot\|}$ will appear explicitly in the optimal regularization. Given that \mathbb{X} is a standard Gaussian random matrix, this diameter will be the Gelfand width of $rB_{\|\cdot\|}$ in our example (cf. [31], Chapter 2 in [13] or [22] for more details on Gelfand widths and their role in signal processing and learning theory).

1.3 Overview of the paper; main results

As a proof of concept we present an example of the construction of a minimax regularization function in the popular set-up of regularization by the ℓ_1^d norm. Let us recall the statistical model we used in both Definition 1.2 and Definition 1.1: the Gaussian linear regression model with a Gaussian design

$$Y = \langle X, t^* \rangle + \xi \tag{1.7}$$

where $X \sim \mathcal{N}(0, I_{d \times d})$ and $\xi \sim \mathcal{N}(0, \sigma^2)$ are independent, centered Gaussian variables in \mathbb{R}^d and \mathbb{R} respectively. As written before, a dataset $(X_i, Y_i)_{i=1}^N$ of N i.i.d. copies of the couple (X, Y) is provided and one wants to use it to estimate t^* .

Note that we choose a Gaussian random design to make the exposition as simple as possible. The results can be extended to more general sub-Gaussian designs. Nonetheless, our goal is not to provide general results but to show that the approach we present allows

to achieve minimax regularization in some classical set-up. Moreover, we want to see the effect of the random design on the construction of a minimax regularization.

In the supplementary material 4, we consider a fixed design setup, in which one still has for all i : $Y_i = \langle X_i, t^* \rangle + \xi_i$ with $\xi_i \sim \mathcal{N}(0, \sigma^2)$ i.i.d., but with X_i deterministic, satisfying an “isomorphic property” on “compressible vectors”, equivalent to the RIP from [12]. We will see that under this property, the arguments and the results will be quite similar to the random design case.

We will not be interested in getting optimal or sharp numerical constants, and some of the inequalities and coefficients in the arguments will be rather loose from this point of view, but what actually matters is that the quantity will have the right order of magnitude w.r.t. N, d, σ and $\|t^*\|_1$.

As for our choice to consider regularizations that are functions of the ℓ_1^d norm, its motivation is that the ℓ_1^d -norm has been one of the most studied regularization norm since the beginning of high-dimensional statistics, in particular for the reasons presented in the Introduction. Moreover, as mentioned previously, the ERM over ℓ_1^d -balls is minimax for every confidence $1 - \delta_N$ (cf. [22]), this makes the construction of minimax regularization possible for different deviation parameters and this makes the exposition also simpler.

Let the choice of the ℓ_1^d -norm as a regularization norm be made once and for all. Now, the problem we want to solve is to construct a regularization function $\Psi : \mathbb{R} \rightarrow \mathbb{R}$ such that the regularized procedure

$$\hat{t}_\Psi \in \operatorname{argmin}_{t \in \mathbb{R}^d} \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2 + \Psi(\|t\|_1) \right) \quad (1.8)$$

achieves the minimax rate of convergence over $\|t^*\|_1 B_1^d$ given N i.i.d. data $(X_i, Y_i), i = 1, \dots, N$ distributed according to (1.7). And, we want \hat{t}_Ψ to satisfy that property whatever $t^* \in \mathbb{R}^d$ is.

Now, let us explain the strategy we use to design a minimax regularization function. We denote for all $\rho \geq 0$ and $r \geq 0$, $\rho B_1^d = \{t \in \mathbb{R}^d : \|t\|_1 \leq \rho\}$, and $r B_2^d = \{t \in \mathbb{R}^d : \|t\|_2 \leq r\}$. The starting point to our approach is that \hat{t}_Ψ minimizes $t \mapsto P_N \mathcal{L}_t^\Psi$ over \mathbb{R}^d , where, for every $t \in \mathbb{R}^d$,

$$P_N \mathcal{L}_t^\Psi := \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2 + \Psi(\|t\|_1) \right) - \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t^* \rangle)^2 + \Psi(\|t^*\|_1) \right),$$

in particular, $P_N \mathcal{L}_{\hat{t}_\Psi}^\Psi \leq P_N \mathcal{L}_{t^*}^\Psi = 0$. So if one shows that $P_N \mathcal{L}_t^\Psi > 0$ for all $\|t\|_1 \gtrsim \|t^*\|_1$ then this will prove that $\|\hat{t}_\Psi\|_1 \lesssim \|t^*\|_1$, proving that \hat{t}_Ψ belongs to the right model. This will be essentially the main step since for the correct choice of Ψ , we will show that the regularization has no effect within the right model and that the RERM \hat{t}^Ψ has essentially the same statistical behavior as the ERM in $\|t^*\|_1 B_1^d$ which is known to be minimax. We will therefore conclude that \hat{t}^Ψ can learn t^* at the minimax rate of convergence within the model $\|t^*\|_1 B_1^d$ without knowing in advance the radius $\|t^*\|_1$.

Using the quadratic / multiplier decomposition as in [22, 34, 35], one can write $P_N \mathcal{L}_t^\Psi$ as the sum of three terms: $P_N \mathcal{L}_t^\Psi = P_N \mathcal{Q}_{t-t^*} + P_N \mathcal{M}_{t-t^*} + \mathcal{R}_{t,t^*}$ where

- $P_N \mathcal{Q}_{t-t^*} := \sum_{i=1}^N (\langle X_i, t^* \rangle - \langle X_i, t \rangle)^2 / N$ is the “quadratic process”
- $P_N \mathcal{M}_{t-t^*} := 2 \sum_{i=1}^N (Y_i - \langle X_i, t^* \rangle) (\langle X_i, t^* \rangle - \langle X_i, t \rangle) / N$ is the “multiplier process”
- $\mathcal{R}_{t,t^*} := \Psi(\|t\|_1) - \Psi(\|t^*\|_1)$ is the regularization part.

The definition of our complexity will thus be a consequence of the study of the behavior of the quadratic and multiplier empirical processes indexed by $t \in \rho B_1^d$ for all $\rho \geq 0$. The two processes are associated to the two statistical complexities previously discussed: 1) the quadratic process can be written as $P_N \mathcal{Q}_{t-t^*} = \|\mathbb{X}(t - t^*)\|_2^2$ and is well behaved (i.e. of the order of $\|t - t^*\|_2^2$) when \mathbb{X} is well conditioned; 2) the multiplier process is measuring the statistical complexity coming from the noise $\xi = Y - \langle X, t^* \rangle$, $P_N \mathcal{M}_{t-t^*}$ is the empirical correlation between the noise and the model shifted by $\langle \cdot, t^* \rangle$. All the game is now to identify regions of the space \mathbb{R}^d where the statistical complexity come from the distortion of \mathbb{X} or from the noise. This drives the construction of the optimal regularization function Ψ .

In order to identify those regions, note that for every fixed $t \in \mathbb{R}^d$, the distribution of these two processes depend on $t - t^*$ only by its ℓ_2^d -norm $\|t - t^*\|_2$, in two different ways: $P_N \mathcal{Q}_{t-t^*}$ in a quadratic way, $P_N \mathcal{M}_{t-t^*}$ in a linear way. So it is natural to partition the model ρB_1^d into vectors with “small” ℓ_2 norm – i.e. the intersection of $\rho B_1^d \cap r B_2^d$ for an adequate radius r – and vectors of ρB_1^d with ℓ_2^d -norm larger than r . We will see that outside $r B_2^d$, with high probability the two processes are “well-behaved” and regularization is unnecessary; but inside $r B_2^d$ it is not the case, the operator \mathbb{X} may have a kernel and the noise is making the estimation hard: hence, this is where the regularization will be needed to keep control of the situation and this is precisely where the regularization function is designed. In that case, either the statistical complexity comes from the size of the intersection of the kernel of \mathbb{X} with ρB_1^d and therefore one needs to take $\Psi(\rho)$ of the order of this diameter (which appears to be equal to the Gelfand width of ρB_1^d to the square) or the statistical complexity comes from the noise and then $\Psi(\rho)$ is of the order of the oscillations of the multiplier process inside $\rho B_1^d \cap r B_2^d$.

The choice of the “adequate radius” r is of course paramount in our approach. It results from the right understanding of the two previously discussed sources of statistical complexities: the bigger these complexities, the bigger this radius (since, as we mentioned, outside $r B_2^d$ the processes are well-behaved). First, we want to identify the smallest ℓ_2^d radius $r_Q(\rho)$ above which \mathbb{X} is well-behaved in ρB_1^d , i.e. such that for every $t \in \rho B_1^d$, if $\|t - t^*\|_2 \geq r_Q(\rho)$, then $P_N \mathcal{Q}_{t-t^*} = \|\mathbb{X}(t - t^*)\|_2^2 \sim \|t - t^*\|_2^2$. Then, we need to identify the smallest ℓ_2^d -radius $r_M(\rho)$ above which the effect of the noise is below the signal intensity that is above which one can clearly identify if $t \neq t^*$ when $\|t - t^*\|_2 \geq r_M(\rho)$. To that end we want to make the oscillations of the multiplier process smaller than the one of the quadratic process, which is of the order of $\|t - t^*\|_2^2$ when $\|t - t^*\|_2 \geq r_Q(\rho)$. It will appear that, in our framework, the two radii obtained from the above trade-offs are solution of fixed point equations for all $\rho \geq 0$: for some absolute constants Q and η (to be chosen later):

- the “quadratic fixed point” is $r_Q(\rho) := \inf \left(r > 0 : \ell^*(\rho B_1^d \cap r B_2^d) = Q r \sqrt{N} \right)$
- the “multiplier fixed point” is $r_M(\rho) := \inf \left(r > 0 : \sigma \ell^*(\rho B_1^d \cap r B_2^d) = \eta r^2 \sqrt{N} \right)$

where $\ell^*(\rho B_1^d \cap r B_2^d)$ is the Gaussian mean width of the localized set $\rho B_1^d \cap r B_2^d$ defined as

$$\ell^*(\rho B_1^d \cap r B_2^d) = \mathbb{E} \sup_{t \in \rho B_1^d \cap r B_2^d} \langle G, t \rangle$$

where G is a standard Gaussian vector in \mathbb{R}^d . As our framework involves ‘‘Gaussian randomness’’ in both the design and the noise, it is not surprising that the Gaussian mean width arise when dealing with the control of the two processes. However, Gaussian mean widths appear in learning theory, statistics and signal processing way beyond the ‘‘full Gaussian framework’’ as considered here (see, for instance, [24]).

These two fixed points have been introduced in [22] and used later in [23] for the study of ERM and RERM. As their names suggest, the quadratic fixed point will be used to control the quadratic process, and the multiplier fixed point to control the multiplier process. Their general definitions use an inequality rather than an equality inside the infimum: $r_Q(\rho)$ is defined as $\inf \left(r > 0 : \ell^*(\rho B_1^d \cap r B_2^d) \leq Qr\sqrt{N} \right)$ and $r_M(\rho)$ as $\inf \left(r > 0 : \sigma \ell^*(\rho B_1^d \cap r B_2^d) \leq \eta r^2 \sqrt{N} \right)$. This allows to deal with infinite-dimensional set-ups in which the mapping $r \mapsto \ell^*(\mathcal{F} \cap r B_2)$ is not necessarily continuous. But in our case, this mapping is continuous and the infimum is attained in a point for which there is exact equality.

It appears that one can provide an explicit formulation for the two fixed point $r_Q(\rho)$ and $r_M(\rho)$ in many situations and, in particular, in the case of the ℓ_1^d -norm (cf. [22]): for some absolute constants $C_M^{(1)}, C_M^{(2)}, C_Q^{(1)}, C_Q^{(2)}$ and $\zeta < 1 < \zeta'$, for all ρ , there exists $C_M \in [C_M^{(1)}, C_M^{(2)}]$ such that:

$$r_M^2(\rho) = C_M \begin{cases} \frac{\sigma^2 d}{N} & \text{if } \rho^2 N \geq \sigma^2 d^2 \\ \rho \sigma \sqrt{\frac{1}{N} \log \left(\frac{e \sigma d}{\rho \sqrt{N}} \right)} & \text{if } \sigma^2 \log d \leq \rho^2 N \leq \sigma^2 d^2 \\ \rho \sigma \sqrt{\frac{\log(ed)}{N}} & \text{if } \rho^2 N \leq \sigma^2 \log d. \end{cases} \quad (1.9)$$

and there exists $C_Q \in [C_Q^{(1)}, C_Q^{(2)}]$ such that

$$r_Q^2(\rho) = C_Q \begin{cases} 0 & \text{if } N \geq \zeta' d \\ \frac{\rho^2}{N} \log \left(\frac{ed}{N} \right) & \text{if } N \leq \zeta d. \end{cases} \quad (1.10)$$

Note that when $\zeta d \leq N \leq \zeta' d$, $r_Q(\rho)$ decays from $(\rho^2/N) \log(ed/N)$ to 0 and one only has an upper estimate on $r_Q(\rho)$ given by $C_Q \rho^2/N$. We will therefore not consider this case in the following since it involves to deal with sharp estimates on the spectra of squared or approximatively squared Gaussian random matrices. Note also that C_M and C_Q may depend on ρ, N, d, σ but they are both controlled from above and below by absolute constants (independent of ρ, N, d and σ).

Now that we have a way to measure the statistical complexity of a model we need one more thing before turning to the effective construction of a minimax regularization for the ℓ_1^d -norm: we need to know the minimax rate of convergence over ℓ_1^d -ball ρB_1^d for all $\rho \geq 0$.

We will see below that one way to measure the statistical complexity of a model is closely related to its minimax rate. To that end, we summarize the main results in the constant deviation case $\delta_N = 1/4$ from section 4.1 in [22] in the following theorem.

Theorem 1.3 *Consider the Gaussian linear model with Gaussian design introduced in (1.7). Let $\rho > 0$. The minimax rate of convergence for constant confidence parameter $\delta_N = 1/4$ over ρB_1^d is achieved by the ERM and is given (up to absolute constants) by*

$$\min(r^2(\rho), \rho^2) \quad \text{where } r(\rho) = \max(r_Q(\rho), r_M(\rho)). \quad (1.11)$$

Up to multiplicative absolute constants, this rate is given for some $\zeta < 1 < \zeta'$,

1. when $N \leq \log d$, by ρ^2 ,
2. when $\log d \leq N \leq \zeta d$, by

$$\begin{cases} \rho^2 & \text{if } \rho^2 N \leq \sigma^2 \log d, \\ \rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{ed^2\sigma^2}{\rho^2 N}\right)} & \text{if } \sigma^2 \log d \leq \rho^2 N \leq \frac{\sigma^2 N^2}{\log(ed/N)}, \\ \frac{\rho^2}{N} \log\left(\frac{ed}{N}\right) & \text{if } \rho^2 N \geq \frac{\sigma^2 N^2}{\log(ed/N)} \end{cases}$$

3. when $N \geq \zeta' d$, by

$$\begin{cases} \rho^2 & \text{if } \rho^2 N \leq \sigma^2 \log d, \\ \rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{ed^2\sigma^2}{\rho^2 N}\right)} & \text{if } \sigma^2 \log d \leq \rho^2 N \leq \sigma^2 d^2, \\ \frac{\sigma^2 d}{N} & \text{if } \rho^2 N \geq \sigma^2 d^2. \end{cases}$$

In other words, for all $\rho \geq 0$ and $t^* \in \rho B_1^d$, the ERM $\hat{t}_\rho^{ERM} \in \operatorname{argmin}_{t \in \rho B_1^d} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2$, is such that, with probability at least $3/4$, $\|\hat{t}_\rho^{ERM} - t^*\|_2^2 \leq \min(r^2(\rho), \rho^2)$. Moreover, there are no estimator that can do uniformly better than the ERM \hat{t}_ρ^{ERM} over ρB_1^d when $N \notin (\zeta d, \zeta' d)$.

Note that we have decided to present the result in the constant deviation result (that is for $\delta_N = 1/4$) whereas it is actually true with a much better probability estimate in section 4.1 in [22]. We will also obtain our main results with an exponentially large deviation below.

As mentioned previously, when $N \in [\zeta d, \zeta' d]$, we only have an upper bound on $(r_Q(\rho))^2$ that does not match the minimax lower bound. As a consequence, the $N \sim d$ regime is not considered in Theorem 1.3. Notable is that the rate ρ^2 is the trivial rate obtained by taking the ℓ_2^d diameter of the model ρB_1^d which is simply 2ρ . Therefore, any statistics \tilde{t}_N (like the ERM \hat{t}_ρ^{ERM}) taking its values in ρB_1^d satisfies with probability 1, $\|\tilde{t}_N - t^*\|_2^2 \leq 4\rho^2$ for all $t^* \in \rho B_1^d$. This is a trivial bound that one can get for free as long as the radius ρ is known. However, for the construction of an optimal regularization function which can be seen as an adaptation to the radius $\|t^*\|_1$, which is therefore not known, this trivial bound is not

available. This will be an issue for designing a minimax regularization function when $\|t^*\|_1$ is unknown and small (actually smaller than $\sigma\sqrt{\log(ed)/N}$). Somehow the “signal-to-noise ratio” is too small for the models ρB_1^d with small ρ 's. Therefore, the trivial upper bound ρ^2 is optimal when ρ is known but in the other case we will have to pay the price due to the noise and there will be no way to achieve the trivial optimal ρ^2 bound for small ρ 's (except for the trivial estimator $\hat{t}_0 = 0$, see the discussion after Proposition 1.6). That is the reason why we will not be able to construct a minimax regularization function over the entire space \mathbb{R}^d but only for t^* such that $\|t^*\|_1 \gtrsim \sigma/\sqrt{\log(ed)/N}$. We will also show that such a construction of an optimal regularization function over the entire space \mathbb{R}^d is actually not possible at all later in Proposition 1.6.

Finally let us turn to the construction of a minimax regularization function for the ℓ_1^d -norm. To that end we will use the function $\rho \geq 0 \mapsto r^2(\rho) = \max(r_Q^2(\rho), r_M^2(\rho))$ as a sharp way to measure the complexity of the model ρB_1^d . The main result of this article is that this function is a minimax regularization function as introduced in Definition 1.2.

Theorem 1.4 *There are absolute constants $\eta, Q, \zeta, \zeta', \Delta_0, c_0$ such that the following holds. When $\zeta'd \geq N$ or $\zeta d \leq N$, a minimax regularization function for the ℓ_1^d -norm over $\mathbb{R}^d \setminus (\Delta_0 \sigma \sqrt{\log(ed)/N} B_1^d)$ for the confidence parameter $\delta_N = 1/4$ is given by the following function: for all $\rho > 0$,*

$$\Psi(\rho) = c_0 r^2(\rho)$$

where $r(\rho) = \max(r_Q(\rho), r_M(\rho))$ for

$$r_Q(\rho) = \inf \left(r > 0 : \ell^*(\rho B_1^d \cap r B_2^d) = Qr\sqrt{N} \right)$$

and

$$r_M(\rho) = \inf \left(r > 0 : \sigma \ell^*(\rho B_1^d \cap r B_2^d) = \eta r^2 \sqrt{N} \right)$$

denoting by $\ell^*(\rho B_1^d \cap r B_2^d)$ the Gaussian mean width of the localized sets $\rho B_1^d \cap r B_2^d$. In that case, the rate achieved by the RERM \hat{t}_Ψ is the minimax rate $r^2(\|t^*\|_1)$ when $\|t^*\|_1 \geq \Delta_0 \sigma \sqrt{\log(ed)/N}$.

The shape of the minimax regularization function $\rho \rightarrow \Psi(\rho) = c_0 r^2(\rho)$ is given in Figure 1 in the two cases $N \leq \zeta d$ (“high-dimensional statistics”) and $N \geq \zeta' d$ (“classical or low-dimensional statistics”).

The only difference between the two cases ($\zeta' d \geq N$ or $\zeta d \leq N$) appears for large radii ρ . The reason for that lies in the statistical complexity coming (or not) from the distortion of the operator \mathbb{X} . In the low-dimensional case, \mathbb{X} is such that (with high probability), $\|\mathbb{X}t\|_2 \sim \|t\|_2$ for all $t \in \mathbb{R}^d$. There is no distortion coming from \mathbb{X} . Somehow observing $\mathbb{X}t^*$ is the same as observing t^* itself, one just has to invert \mathbb{X} – this can be done because \mathbb{X} acts like an isomorphism on the entire space \mathbb{R}^d . Therefore, there is no statistical complexity coming from \mathbb{X} and so its associated complexity parameter $r_Q(\cdot)$ does not show up in the final complexity parameter $r(\cdot) = \max(r_M(\cdot), r_Q(\cdot))$. We therefore end up with $r(\cdot) = r_M(\cdot)$ in the low-dimensional case. In particular, for large radii ρ (for which, one has $\rho B_1^d \cap r_M(\rho) B_2^d = r_M(\rho) B_2^d$), we pay the worst rate of convergence in \mathbb{R}^d , which is $\sigma^2 d/N$




Figure 1: Shape of the graph of the minimax regularization function $\rho \rightarrow r^2(\rho)$ of the ℓ_1^d -norm for the cases $N \leq \zeta d$ (left) and $N \geq \zeta' d$ (right)

because learning over ρB_1^d for large values of ρ is as hard as learning over the entire space \mathbb{R}^d and the price for the latter is the rate $\sigma^2 d/N$.

The situation is totally different in the high-dimensional setup because in that case the operator $\mathbb{X} \in \mathbb{R}^{N \times d}$ has a non-trivial kernel; therefore, observing $\mathbb{X}t^*$ is totally different from observing t^* (for instance, imagining that $t^* \in \ker \mathbb{X}$). This adds to the statistical complexity of the problem of estimating t^* . In this regime, both the noise and the distortion effect of \mathbb{X} appear in the statistical complexity of the estimation problem; this means that both complexity parameter $r_Q(\cdot)$ and $r_M(\cdot)$ appear in the total complexity parameter $r(\cdot)$ and therefore in the ultimately designed minimax regularization function. For small values of ρ , the effect of the noise is predominant but for large values of ρ this is the effect of \mathbb{X} which is the main responsible of the statistical complexity. In particular, the ℓ_2^d -diameter of $\ker \mathbb{X} \cap \|t^*\|_1 B_1^d$ is important because there is of course no way to distinguish t^* from $t^* + h$ for all $h \in \mathbb{R}^d$ such that $\mathbb{X}t^* = \mathbb{X}(t^* + h)$ that is for all $h \in \ker \mathbb{X}$ such that $t^* + h \in \|t^*\|_1 B_1^d$. Hence, estimating t^* is at least as hard as estimating any point in $(t^* + \ker \mathbb{X}) \cap \|t^*\|_1 B_1^d$ and therefore, no estimator \tilde{t} can estimate t^* at a rate better than $\text{diam}(\ker \mathbb{X} \cap \|t^*\|_1 B_1^d, \ell_2^d)^2$. The latter quantity is itself lower bounded by the Gelfand's N -width of $\|t^*\|_1 B_1^d$ defined as

$$c_N(\|t^*\|_1 B_1^d) := \inf \left\{ \text{diam}(\ker \Gamma \cap \|t^*\|_1 B_1^d) : \Gamma \in \mathbb{R}^{N \times d} \right\} \sim \|t^*\| \min \left\{ 1, \sqrt{\frac{\log(ed/N)}{N}} \right\} \quad (1.12)$$

the latter result is due to Garanaev and Gluskin [17]. It appears that the Gelfand's N -width of $\|t^*\|_1 B_1^d$ are achieved (up to absolute constants and with high probability) by the kernel of standard $N \times d$ Gaussian matrices, which is exactly the case of the design matrix \mathbb{X} . Therefore, with high probability,

$$\text{diam}(\ker \mathbb{X} \cap \|t^*\|_1 B_1^d, \ell_2^d)^2 \sim c_N^2(\|t^*\|_1 B_1^d) \sim \frac{\|t^*\|_1}{N} \log\left(\frac{ed}{N}\right) \sim r_Q^2(\|t^*\|_1) \quad (1.13)$$

This is exactly the price we pay in $r_Q(\rho)$ when $\rho \geq \sigma\sqrt{N/\log(ed/N)}$. That is the reason why we take the regularization function $\Psi(\rho)$ of the order of the Gelfand's N -width of ρB_1^d

(to the square) for large radii ρ : it is the right concept of statistical complexity that shows up in this part of the space \mathbb{R}^d , where the statistical complexity coming from the distortion of \mathbb{X} becomes more important than the one due to the noise.

Remark 1.5 (Regularization function for the LASSO) *The LASSO is the RERM procedure obtained for a linear regularization function $\Psi(\rho) = \sigma\rho\sqrt{\log d/N}$ which is obtained by using a trivial upper bound on the complexity of the model ρB_1^d in (1.8):*

$$\text{comp}(\rho B_1^d) = r_M^2(\rho) \leq \frac{\sigma \ell^*(\rho B_1^d)}{\sqrt{N}} = \sigma\rho\sqrt{\frac{\log(ed)}{N}}. \quad (1.14)$$

This complexity is obtained by simply removing the localization (i.e. the intersection with ρB_2^d) in the multiplier process when computing $r_M(\cdot)$, and does not take $r_Q(\cdot)$ into account. This means that the distortion of the operator \mathbb{X} is supposed to have no effect on the statistical complexity of the problem. This is why estimation results for the LASSO deal only with the reconstruction of vectors which are sparse or almost sparse, i.e. for vectors belonging to the cone appearing in the RE or CC conditions, cf. [6]. Over this cone, the quadratic process behaves nicely (that is, the isomorphic property from Proposition 2.1 holds on this cone) or in other words, the operator \mathbb{X} is well-conditioned on the set of vectors we want to reconstruct, so that there is no statistical complexity coming from the distortion of this operator. So, as long as estimation of sparse or approximately sparse vectors is concerned, there is no need for the complexity function $r_Q(\cdot)$. That is why the regularization function used for the LASSO take into account only the fixed point $r_M(\cdot)$ associated to the statistical complexity due to the noise and not the one from the inverse problem. On the contrary, by taking $r_M(\cdot)$ and $r_Q(\cdot)$ into account, our regularization function allows us to deal with the full space \mathbb{R}^d (except for a small ℓ_1^d -ball centered in 0, cf. Proposition 1.6) and not only a cone.

Moreover, as said before, the way the regularization function is designed in (1.14) is sub-optimal because it uses a trivial upper bound on $r_M(\cdot)$ instead of using the exact formulation of $r_M(\cdot)$ as in (1.9). Contrary to the LASSO, this latter exact formulation takes into account, thanks to the localization, the fact that the regularization is not needed on the whole space –in some areas the random processes behave nicely whatever. The suboptimal approach for the LASSO is likely to be responsible for a loss in the rate of convergence achieved by the LASSO, which is $s\sigma^2 \log(ed)/N$ whereas the minimax rate is $\sigma^2 \log(ed/s)/N$ (cf. [5]). This is not a big loss, especially when $d \gg s$, but from a purely theoretical point of view the right way to regularize for the reconstruction of sparse vectors should be using $r_M^2(\|t\|_1)$ instead of $\sigma\|t\|_1\sqrt{\log(ed)/N}$ as it is the case for the LASSO. However, the resulting regularization function would be concave (cf. the right-hand side plot in Figure 1). Therefore, the small price paid from a theoretical point of view by using the trivial upper bound in (1.14) seems to be worth the computational gain obtained by using a convex regularization as does the LASSO.

Let us now turn to the adaptation problem in the ball ρB_1^d for $\rho \sim \sigma\sqrt{\log(ed)/N}$. We want to answer the following question: is it possible to construct a regularization function $\Psi(\cdot)$ so that the associated regularized procedure \hat{t}_Ψ is adaptive on the entire space \mathbb{R}^d ? Or (even stronger) is there any statistic that can be adaptive (in the sense that it achieves

the rate of the ERM on $\|t^*\|_1 B_1^d$ without knowing $\|t^*\|_1$ beforehand) on the entire space \mathbb{R}^d (this statistics may not be a regularized procedure)? It appears that the answer to this question is negative, which we prove in the following proposition.

However, one needs to be cautious with the next statement because there is a trivial estimator $\hat{t}_0 = 0$ such that for every $t^* \in \mathbb{R}^d$, with probability 1, $\|\hat{t}_0 - t^*\|_2^2 = \|t^*\|_2^2$ and therefore \hat{t}_0 is adaptive on ρB_1^d as long as the minimax rate over ρB_1^d is ρ^2 , which is the case for any $\rho \lesssim \sigma \sqrt{\log(ed)/N}$. Therefore, there exists a procedure adaptive on ρB_1^d when $\rho \sim \sigma \sqrt{\log(ed)/N}$. Moreover, according to Theorem 1.4, there exists a procedure adaptive on $\mathbb{R}^d \setminus \rho B_1^d$. But the question concerns the adaptation on the *entire space* \mathbb{R}^d *at the same time*.

The following statement shows that if \hat{t} is a procedure adaptive on $\mathbb{R}^d \setminus \rho B_1^d$ then it cannot be adaptive on ρB_1^d for $\rho \sim \sigma \sqrt{\log(ed)/N}$. Moreover, it also proves that adaptation on the entire space \mathbb{R}^d is not possible and that Theorem 1.4 is optimal given that the range of radii $[\Delta_0 \sigma \sqrt{\log(ed)/N}, +\infty)$ on which it is adaptive cannot be inflated (up to absolute constants). Before turning to the statement let us denote by \mathbb{P}_{t^*} the probability distribution of a N -sample $(X_i, Y_i)_{i=1}^N$ of i.i.d. copies of (X, Y) when (X, Y) is distributed according to (1.7).

Proposition 1.6 *Assume that $2d \geq \exp(544/225)$ and that there exists an absolute constant χ_1 such that the following holds. Let $\rho \leq 2\sigma \sqrt{(\log(2d))/(96N)}$ be such that $16\chi_1 r^2(\rho) \leq \rho^2$ and denote by $(e_j)_{j=1}^d$ the canonical basis of \mathbb{R}^d . Assume that \hat{t} is an estimator such that for every $t^* \in \{\pm \rho e_1, \dots, \pm \rho e_d\}$,*

$$\mathbb{P}_{t^*} [\|\hat{t} - t^*\|_2^2 \leq \chi_1 r^2(\rho)] \geq \frac{3}{4}.$$

Then, for every $t^ \in (\rho/2)B_1^d$,*

$$\mathbb{P}_{t^*} [\|\hat{t} - t^*\|_2^2 \geq \rho^2/16] \geq \frac{1}{2}. \quad (1.15)$$

The proof of Proposition 1.6 is given in Section 3. Note that the only property of the design X used to prove Proposition 1.6 is isotropicity. Since isotropicity does not tell much on the distortion properties of the design matrix \mathbb{X} , it means that Proposition 1.6 is only based on the statistical complexity coming from the noise. This is not a surprise given that Proposition 1.6 is a result for very small radii less than $\sim \sigma \sqrt{\log(ed)/N}$. At that scale, even if $\ker \mathbb{X}$ is in the worst possible position, i.e. $\text{diam}(\ker \mathbb{X} \cap \rho B_1^2, \ell_2^d)^2 = \text{diam}(\rho B_1^2, \ell_2^d)^2 = \rho^2$, we still have $\rho^2 \lesssim r_M^2(\rho)$. Hence, the distortion of \mathbb{X} does not play any role at this very small scale and therefore that is not a surprise that Proposition 1.6 is true for any isotropic design X .

Finally, let us rephrase Proposition 1.6 in other words. Proposition 1.6 shows that if a procedure can learn all vectors in $\{\pm \rho e_1, \dots, \pm \rho e_d\}$ at the minimax rate $r^2(\rho)$ then this estimator cannot learn any $t^* \in (\rho/2)B_1^d$ at the optimal minimax rate ρ^2 for confidence $1/4$. For instance, given that the result (1.15) holds for any $t^* \in (\rho/2)B_1^d$, in particular, for $t^* = 0$, it tells that \hat{t} cannot estimate $t^* = 0$ at a rate better than $\rho^2 \sim \sigma^2 \log(ed)/N$ whereas the minimax rate over $\rho^* B_1^d$ for $\rho^* = 0$ is obviously 0. Finally, note that the condition $16\chi_1 r^2(\rho) \leq \rho^2$ implies that $\rho \gtrsim \sigma \sqrt{\log(ed)/N}$ so that the phase transition

radius above which adaptation is possible but not below is of the order of $\sigma\sqrt{\log(ed)/N}$ which is the radius we have found in Theorem 1.4.

2 Proof of Theorem 1.4

Most of the proof consists in showing that with high probability, \hat{t} belongs to $t^* + \rho^* B_1^d$ where

$$\rho^* = \max \left(10, 8 \frac{C_M^{(2)2}}{(C_M^{(1)})^2 \eta} + 1 \right) \|t^*\|_1. \quad (2.1)$$

(once this goal is achieved, it is straightforward to show (again with high probability) that $\|\hat{t} - t^*\|_2^2$ is less than the minimax rate of convergence over $\rho^* B_1^d$).

To do so, we will prove that with high probability, any t outside $t^* + \rho^* B_1^d$ satisfies $P_N \mathcal{L}_t^\Psi > 0$ (whereas $P_N \mathcal{L}_t^\Psi \leq 0$). We partition $\mathbb{R}^d \setminus (t^* + \rho^* B_1^d)$ into shelves of the form $t^* + (2^{j+1} \rho^* B_1^d \setminus (2^j \rho^*) B_1^d)$, in which the regularization function remains mostly constant. We only need to study the smallest shells, i.e. for $k = 1, \dots, K_0$ for some well-chosen K_0 (K_0 is the smallest integer so that $2^{K_0-1} \rho^* B_1^d \cap r(2^{K_0-1} \rho^*) B_2^d = r(2^{K_0-1} \rho^*) B_2^d$), the part of \mathbb{R}^d for which $\|t\|_1 \geq 2^{K_0} \rho^*$ will be treated by an homogeneity argument.

On each of the smallest shelves, the argument is roughly and heuristically the following: we place ourselves on a high probability event on which random processes ($P_N \mathcal{M}, P_N \mathcal{Q}$, their supremum, their infimum,...) “behave nicely” (i.e. they both scale like $\|t - t^*\|_2^2$). Then, as $P_N \mathcal{M}_{t-t^*}$ is the only possibly negative term, it suffices to identify zones where $P_N \mathcal{Q}_{t-t^*} > P_N \mathcal{M}_{t-t^*}$ (then directly $P_N \mathcal{L}_t^\Psi > 0$), and compensate $|P_N \mathcal{M}_{t-t^*}|$ on the other part by using a penalty that is close to the supremum (on this other part) of $P_N \mathcal{M}_{t-t^*}$. As $P_N \mathcal{Q}_{t-t^*}$ grows quicker than $P_N \mathcal{M}_{t-t^*}$ with respect to $\|t - t^*\|_2$, the big zone where $P_N \mathcal{Q}_{t-t^*} > P_N \mathcal{M}_{t-t^*}$ will be the exterior of $t^* + r B_2^d$ for an adequate r (cf. Figure 2). This r must be such that any t in the exterior of this ball satisfies $P_N \mathcal{Q}_{t-t^*} \gtrsim \|t - t^*\|_2^2 \gtrsim P_N \mathcal{M}_{t-t^*}$, and we will see that the first inequality amounts to $r \geq r_Q(\rho)$, and the second one to $r \geq r_M(\rho)$. Next, the supremum of $P_N \mathcal{M}_{t-t^*}$ on $\rho B_1^d \cap r(\rho) B_2^d$ is less than $r_M(\rho) r(\rho) \leq r^2(\rho)$ for $r(\rho) = \max(r_M(\rho), r_Q(\rho))$. We therefore set the regularization function $\Psi(\rho)$ at level ρ to be proportional to the quantity $r^2(\rho)$ because it is this quantity measuring the amplitude of the oscillation of the multiplier process in $\rho B_1^d \cap r(\rho) B_2^d$.

As for its presentation, the proof of Theorem 1.4 is divided into two parts. The first part (Section 2.1) defines the event on which the two processes “behave nicely” and computes a lower bound on its probability. In the second part (Section 2.2) we will place ourselves on this event and carry out the deterministic geometric part of the argument.

2.1 Probabilistic control of the processes

Instead of controlling the two processes on shelves, we will control them on the full ℓ_1 balls, because it does not change the complexity, up to constants, and the very last step of the proof requires a control on the two processes on the full ℓ_1 ball $\rho^* B_1^d$.

2.1.1 Control of the quadratic process

This first section provides the classical analysis of the quadratic process based upon its isomorphic properties on the set of “almost sparse vectors”. Such a property holds in the optimal regime of observation (or the optimal size of the cone of “almost sparse vectors”), only in the sub-Gaussian case. It is the case we are considering here since we assumed that the design is a standard Gaussian random variable. This analysis borrows some ideas from the “isomorphic method” from [3] or the Restricted Isometry Property from [10] in the sub-Gaussian case. For the sake of completeness we recall here the argument from [22].

Proposition 2.1 *There are absolute constants C_1 and C'_1 such that the following holds. Let X_1, \dots, X_N be N i.i.d. standard Gaussian vectors in \mathbb{R}^d . Denote by Ω^* the event on which: for every $\rho \geq \rho^*$ and all $t \in t^* + \rho B_1^d$,*

$$\text{if } \|t - t^*\|_2 \geq r_Q(\rho) \text{ then } \frac{1}{2} \|t - t^*\|_2^2 \leq \frac{1}{N} \sum_{i=1}^N \langle X_i, t - t^* \rangle^2 \leq \frac{3}{2} \|t - t^*\|_2^2. \quad (2.2)$$

Then, one has $\mathbb{P}[\Omega^*] \geq 1 - 2 \exp(-C_1 Q^2 N)$ as long as $Q \leq C'_1$.

Proof. First note that for all $\rho > 0$, $r_Q(\rho) = \rho r_Q(1)$. Indeed, we have

$$\ell^*(\rho B_1^d \cap r B_2^d) = \ell^*(\rho(B_1^d \cap (r/\rho)B_2^d)) = \rho \ell^*(B_1^d \cap (r/\rho)B_2^d)$$

and so

$$\begin{aligned} r_Q(\rho) &= \inf\{r > 0 : \ell^*(\rho B_1^d \cap r B_2^d) = Q\sqrt{N}r\} \\ &= \inf\{r > 0 : \ell^*(B_1^d \cap (r/\rho)B_2^d) = Q\sqrt{N}(r/\rho)\} \\ &= \rho \inf\{r > 0 : \ell^*(B_1^d \cap r B_2^d) = Q\sqrt{N}r\} = \rho r_Q(1). \end{aligned} \quad (2.3)$$

For all $\rho > 0$, define the event $\Omega(\rho)$ on which one has for all $t \in t^* + \rho B_1^d$,

$$\text{if } \|t - t^*\|_2 \geq r_Q(\rho) \text{ then } \frac{1}{2} \|t - t^*\|_2^2 \leq \frac{1}{N} \sum_{i=1}^N \langle X_i, t - t^* \rangle^2 \leq \frac{3}{2} \|t - t^*\|_2^2.$$

Let us show that if $\Omega(\rho^*)$ holds then for any $\rho \geq \rho^*$, $\Omega(\rho)$ holds as well. Suppose that $\Omega(\rho^*)$ holds. Consider $t \in t^* + \rho B_1^d$ such that $\|t - t^*\|_2 > r_Q(\rho)$ and define

$$t' := t^* + (\rho^*/\rho)(t - t^*) \in t^* + \rho^* B_1^d.$$

It follows from (2.3) that $r_Q(\rho^*) = (\rho^*/\rho)r_Q(\rho)$. Thus $\|t' - t^*\|_2 = (\rho^*/\rho)\|t - t^*\|_2 > (\rho^*/\rho)r_Q(\rho) = r_Q(\rho^*)$, and since $\Omega(\rho^*)$ holds, it follows that $\|t' - t^*\|_2^2/2 \leq \sum_{i=1}^N \langle X_i, t' - t^* \rangle^2/N \leq 3\|t' - t^*\|_2^2/2$. This implies that $\|t - t^*\|_2^2/2 \leq \sum_{i=1}^N \langle X_i, t - t^* \rangle^2/N \leq 3\|t - t^*\|_2^2/2$ so $\Omega(\rho)$ holds.

As a conclusion, $\Omega^* = \Omega(\rho^*)$ and we can now lower bound the probability that this event holds.

Let us consider the class of linear functions

$$F = \left\{ \langle \cdot, t - t^* \rangle, t \in t^* + \rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1} \right\} = \left\{ \langle \cdot, t \rangle, t \in \rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1} \right\}.$$

We assume that F is non empty (if $F = \emptyset$ then the theorem is trivially satisfied). It follows from Theorem 1.12 in [28] that for any $x > 0$, with probability at least $1 - 2 \exp(-C_1 \min(x^2, x\sqrt{N}))$,

$$\sup_{f \in F} \left| \frac{1}{N} \sum_{i=1}^N f^2(X_i) - \mathbb{E} f^2(X) \right| \leq C_2 \left(\frac{\Delta \gamma}{\sqrt{N}} + \frac{\gamma^2}{N} + \frac{x \Delta^2}{\sqrt{N}} \right)$$

where Δ is the diameter in ψ_2 of F and γ is Talagrand's γ_2 functional of F w.r.t. ψ_2 . Note that since X is a standard Gaussian variable in \mathbb{R}^d , for any $t \in \mathbb{R}^d$, $\|\langle X, t \rangle\|_{\psi_2} = C_3 \|t\|_2$ for some absolute constant C_3 . It follows that

$$\Delta = 2 \sup_{t \in \rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1}} C_3 \|t\|_2 = 2C_3 r_Q(\rho^*) \text{ and } \gamma = \gamma_2(\rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1}, \ell_2^d).$$

Moreover, it follows from the Majorizing measure theorem (cf. Chapter 1 in [37]) that

$$\gamma_2(\rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1}, \ell_2^d) \leq C_4 \ell^*(\rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1})$$

Since F is non-empty, by Lemma 3.1 the right-hand side is equal to $C_4 \ell^*(\rho^* B_1^d \cap r_Q(\rho^*) B_2^{d-1})$ and so by definition of $r_Q(\rho^*)$, one has $\gamma \leq C_4 Q r_Q(\rho^*) \sqrt{N}$.

Since X is isotropic (i.e. for any $t \in \mathbb{R}^d$, $\mathbb{E} \langle X, t \rangle^2 = \|t\|_2^2$), we obtain for $x = Q\sqrt{N}$ that, with probability greater than $1 - 2 \exp(-C_1 \min(Q, Q^2)N)$, for any $t \in t^* + \rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1}$,

$$\left| \frac{1}{N} \sum_{i=1}^N \langle X_i, t - t^* \rangle^2 - \|t - t^*\|_2^2 \right| \leq (2C_2 C_3 C_4 Q + C_2 C_4^2 Q^2 + 4C_2 Q C_3^2) r_Q^2(\rho^*)$$

So, as long as long as: $Q \leq C'_1 := \min\{1, (12C_2 C_3 C_4)^{-1}, (\sqrt{6}C_2 C_4)^{-1}, (24C_2 C_3^2)^{-1}\}$, one has, with probability greater than $1 - 2 \exp(-C_1 Q^2 N)$, for all $t \in t^* + \rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1}$,

$$\left| \frac{1}{N} \sum_{i=1}^N \langle X_i, t - t^* \rangle^2 - \|t - t^*\|_2^2 \right| \leq \frac{r_Q^2(\rho^*)}{2} = \frac{1}{2} \|t - t^*\|_2^2. \quad (2.4)$$

In other words, the quadratic process satisfies an isomorphic property on the set $t^* + (\rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1})$. Now, it remains to extend this result to the set of vectors $t \in t^* + \rho^* B_1^d$ such that $\|t - t^*\|_2 \geq r_Q(\rho^*)$. Let t be such a vector and define $t' := t^* + (r_Q(\rho^*) / \|t - t^*\|)(t - t^*)$. Since $t' \in t^* + (\rho^* B_1^d \cap r_Q(\rho^*) S_2^{d-1})$, it satisfies the isomorphic property from (2.4) and so

$$\frac{1}{2} \|t - t^*\|_2^2 \leq \frac{1}{N} \sum_{i=1}^N \langle X_i, t - t^* \rangle^2 \leq \frac{3}{2} \|t - t^*\|_2^2$$

which corresponds exactly to the event $\Omega(\rho^*)$. Therefore, $\mathbb{P}[\Omega^*] = \mathbb{P}[\Omega(\rho^*)] \geq 1 - 2 \exp(-C_1 Q^2 N)$. \blacksquare

2.1.2 Control of the multiplier process

In this section, we provide a control of the multiplier process on several shelves of the space \mathbb{R}^d .

Define r_0 as the non-null solution to $\sigma\ell^*(rB_2^d) = \eta r^2\sqrt{N}$ – i.e. $r_0 = \sigma\ell^*(B_2^d)/(\eta\sqrt{N}) = (\sigma/\eta)\sqrt{d/N}$. Let ρ_0 be the smallest ρ such that ρB_1^d contains $r_0 B_2^d$ – i.e., $\rho_0 = r_0\sqrt{d}$. We can see that ρ_0 is such that $r_M(\rho) = r_M(\rho_0)$ for all $\rho \geq \rho_0$. Indeed, one first sees that $r_M(\rho_0) = r_0$, since $\sigma\ell^*(r_0 B_2^d \cap \rho_0 B_1^d) = \sigma\ell^*(r_0 B_2^d) = \eta r_0^2\sqrt{N}$, and $\sigma\ell^*(r_0 B_2^d \cap \rho_0 B_1^d) = \sigma r\ell^*(B_2^d) > \eta r^2\sqrt{N}$ for all $r < r_0$ and for $r > r_0$, $\sigma\ell^*(r_0 B_2^d \cap \rho_0 B_1^d) \leq \sigma r\ell^*(B_2^d) \leq \eta r^2\sqrt{N}$. This last argument also holds for $\rho \geq \rho_0$. In the latter case, if $r > r_0$ then

$$\sigma\ell^*(r_0 B_2^d \cap \rho B_1^d) \leq \frac{r}{r_0}\sigma\ell^*(r_0 B_2^d \cap \rho B_1^d) = \frac{r}{r_0}\sigma\ell^*(r_0 B_2^d \cap \rho_0 B_1^d) = \frac{r}{r_0}\eta r_0^2\sqrt{N} < \eta r^2\sqrt{N}$$

which means that $r_M(\rho) \leq r_0$. And as $\rho \geq \rho_0$, $r_M(\rho) \geq r_M(\rho_0) = r_0$. Therefore, for $\rho \geq \rho_0$, $r_M(\rho)$ is constant, equal to r_0 . And on $[0, \rho_0]$, r_M is non-decreasing: let $\rho' \leq \rho'' \leq \rho_0$, then $\sigma\ell^*(r_M(\rho')B_2^d \cap \rho''B_1^d) \geq \sigma\ell^*(r_M(\rho')B_2^d \cap \rho'B_1^d) = \eta r_M(\rho')^2\sqrt{N}$ so $r_M(\rho'') \geq r_M(\rho')$.

We denote $K_0 = \min\{k \in \mathbb{N} : 2^k \rho^* \geq 2\rho_0\}$: we will see later that K_0 is defined that way to be the number of the first “shell” such that $r_M(2^{K_0-1}\rho^*)B_2^d \subset 2^{K_0-1}\rho^*B_1^d$.

Proposition 2.2 *There exists an absolute constant C_5 such that the following holds. Let X_1, \dots, X_N be N i.i.d. standard Gaussian vectors in \mathbb{R}^d and ξ_1, \dots, ξ_N be N standard real-valued Gaussian variables independent of the X_i ’s. For all $k = 0, \dots, K_0$, denote by A_k the event on which, for every $t \in \mathbb{R}^d$ such that $\|t - t^*\|_1 \leq 2^k \rho^*$:*

$$|P_N \mathcal{M}_{t-t^*}| \leq \frac{1}{4} \max\left(r_M(2^k \rho^*)^2, \|t - t^*\|_2^2\right). \quad (2.5)$$

Then, for $\eta = 1/(16\sqrt{2})$, one has

$$\mathbb{P}\left[\bigcap_{k=0}^{K_0} A_k\right] \geq 1 - 2\exp(-C_5 N) - 40\exp\left(-\frac{C_M^{(1)} N r_M(\rho^*)^2}{1024 C_M^{(2)} \sigma^2}\right)$$

when $\rho^* \geq 4096 \log(2)\sigma / (C_M^{(1)}\sqrt{N})$.

Proof. We first work conditionally to the $\xi_i, i = 1, \dots, N$. Let $\rho > 0$ and define $T(\rho) := t^* + \rho B_1^d \cap r_M(\rho)B_2^d$. It follows from the Gaussian concentration inequality (cf. Borell’s inequality in [25]) that, for all $x > 0$, with probability greater than $1 - 2\exp(-x^2/2)$,

$$\left|\sup_{t \in T(\rho)} \sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle - \mathbb{E} \sup_{t \in T(\rho)} \sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle\right| \leq x\sigma(T(\rho))$$

where

$$\sigma(T(\rho)) = \sup_{t \in T(\rho)} \sqrt{\mathbb{E} \left(\sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle \right)^2}.$$

Conditionally to the ξ_i 's, the Gaussian process $\left(\sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle\right)_{t \in T(\rho)}$ has the same distribution as the Gaussian process $(\widehat{\sigma}_N \langle X_1, t - t^* \rangle)_{t \in T(\rho)}$ for $\widehat{\sigma}_N := \sqrt{\sum_{i=1}^N \xi_i^2}$. This yields

$$\mathbb{E} \left[\sup_{t \in T(\rho)} \sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle \right] = \widehat{\sigma}_N \ell^*(T(\rho) - t^*) = \widehat{\sigma}_N \ell^*(\rho B_1^d \cap r_M(\rho) B_2^d) = \widehat{\sigma}_N \frac{\eta \sqrt{N} r_M^2(\rho)}{\sigma}$$

and $\sigma(T(\rho)) = \sup_{t \in T(\rho)} \widehat{\sigma}_N \sqrt{\mathbb{E}[\langle X, t - t^* \rangle^2]} = \widehat{\sigma}_N \sup_{t \in T(\rho)} \|t - t^*\|_2 \leq \widehat{\sigma}_N r_M(\rho)$.

So, conditionally on $(\xi_i)_{i=1}^N$, for all $x > 0$, one has, with probability at least $1 - 2 \exp(-x^2/2)$,

$$\sup_{t \in T(\rho)} \left| \sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle \right| \leq \widehat{\sigma}_N \frac{\eta \sqrt{N} r_M^2(\rho)}{\sigma} + x \widehat{\sigma}_N r_M(\rho)$$

Thus, taking $x = \eta \sqrt{N} r_M(\rho) / \sigma$ in the previous statement, one gets, on an event which probability is at least $1 - 2 \exp(-\eta^2 N r_M^2(\rho) / (2\sigma^2))$,

$$\sup_{t \in t^* + \rho B_1^d \cap r_M(\rho) B_2^d} \left| \frac{1}{N} \sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle \right| \leq 2\eta \frac{\widehat{\sigma}_N}{\sigma \sqrt{N}} r_M^2(\rho). \quad (2.6)$$

It remains to prove, on the same event, the result for all $t \in t^* + \rho B_1^d$ such that $\|t - t^*\|_2 > r_M(\rho)$. Define $t' := t^* + (r_M(\rho) / \|t - t^*\|_2)(t - t^*)$. Since $t' \in T(\rho)$, it follows from (2.6) that (on the same event):

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i \langle X_i, \frac{r_M(\rho)}{\|t - t^*\|_2} (t - t^*) \rangle \right| = \left| \frac{1}{N} \sum_{i=1}^N \xi_i \langle X_i, t' - t^* \rangle \right| \leq 2\eta \frac{\widehat{\sigma}_N}{\sigma \sqrt{N}} r_M(\rho)^2$$

and since $r_M(\rho) \leq \|t - t^*\|_2$ one gets

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i \langle X_i, t - t^* \rangle \right| \leq 2\eta \frac{\widehat{\sigma}_N}{\sigma \sqrt{N}} \|t - t^*\| r_M(\rho) \leq 2\eta \frac{\widehat{\sigma}_N}{\sigma \sqrt{N}} \|t - t^*\|^2.$$

Hence, with probability (conditionally to the ξ_i) at least $1 - 2 \exp(-\eta^2 N r_M(\rho)^2 / (2\sigma^2))$, the multiplier process is controlled such that

$$\sup_{t \in t^* + \rho B_1^d} |P_N \mathcal{M}_{t-t^*}| \leq 4\eta \frac{\widehat{\sigma}_N}{\sigma \sqrt{N}} \max(r_M(\rho)^2, \|t - t^*\|_2^2). \quad (2.7)$$

A control of the probability measure of the event A_k follows by applying the previous result to $\rho = 2^k \rho^*$ when $\eta \leq 1/(16\sqrt{2})$ together with a control of the term $\widehat{\sigma}_N$. It follows from an union bound that, conditionally to the ξ_i , (2.7) is satisfied for all $\rho = 2^k \rho^*$, $k = 0, \dots, K_0$ on an event whose probability measure is larger than

$$1 - 2 \sum_{k=0}^{K_0} \exp\left(-\eta^2 N r_M(2^k \rho^*)^2 / (2\sigma^2)\right). \quad (2.8)$$

We handle the last term below thanks to Lemma 3.2.

Now, we handle the random variables ξ_1, \dots, ξ_N . It appears that only a control of the empirical variance term $\widehat{\sigma}_N/\sqrt{N}$ is needed to get a fully deterministic upper bound in the right-hand term of (2.7). It follows from Bernstein inequality for subexponential variables (cf. Theorem 1.2.7 in [13]) that with probability greater than $1 - 2\exp(-C_5N)$,

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i^2 - \sigma^2 \right| \leq \sigma^2,$$

which implies $\widehat{\sigma}_N/\sqrt{N} \leq \sqrt{2}\sigma$. Therefore, for $\eta = 1/(16\sqrt{2})$ we have

$$\mathbb{P} \left[4\eta\widehat{\sigma}_N/(\sqrt{N}\sigma) \leq 1/4 \right] \geq 1 - 2\exp(-C_5N). \quad (2.9)$$

Binding together (2.7), (2.8) and (2.9) gives

$$\mathbb{P} \left[\bigcap_{k=0}^{K_0} A_k \right] \geq 1 - 2\exp(-C_5N) - 2 \sum_{k=0}^{K_0} \exp \left(-\frac{Nr_M(2^k\rho^*)^2}{1024\sigma^2} \right).$$

Finally, Lemma 3.2 yields the following bound

$$\sum_{k=0}^{K_0} \exp \left(-\frac{Nr_M(2^k\rho^*)^2}{1024\sigma^2} \right) \leq \frac{10}{1 - \exp \left(-\frac{C_M^{(1)}\sqrt{N}\rho^*}{4096\sigma} \right)} \exp \left(-\frac{C_M^{(1)}Nr_M(\rho^*)^2}{C_M^{(2)}1024\sigma^2} \right)$$

and the result follows when $\rho^* \geq 4096 \log(2)\sigma/(C_M^{(1)}\sqrt{N})$, which implies that the denominator of the right-hand side is greater than 1/2. ■

2.1.3 Conclusion: construction of the event Ω_0

We define the event

$$\Omega_0 = \Omega^* \cap \bigcap_{k=0}^{K_0} A_k.$$

It follows from Proposition 2.1 and Proposition 2.2 (as well as Lemma 2.4 below), that, as long as $\rho^* \geq 4096 \log(2)\sigma/(C_M^{(1)}\sqrt{N})$,

$$\mathbb{P}[\Omega_0] \geq 1 - 4\exp(-C_6N) - 40 \exp \left(-\frac{C'_6Nr_M(\|t^*\|_1)^2}{\sigma^2} \right)$$

where C_6 and C'_6 are absolute constants.

2.2 Deterministic part of the proof

We first start with some few lemmas on the growth of $r_M(\cdot)$ and $r_Q(\cdot)$. We then construct a partition of \mathbb{R}^d depending on the behavior of function $r^2(\cdot)$ (in particular, its concavity for intermediate values is an issue; we solve it thanks to a peeling argument). We then turn to the main deterministic argument showing that \hat{t} belongs to a ℓ_1^d -ball of radius ρ^* around t^* . The latter holds on the event Ω_0 introduced in Section 2.1.3.

2.2.1 Two Lemmas on the growth of r_M and r_Q

Lemma 2.3 *Let $\rho > 0$ and $\phi = 4(C_M^{(2)})^2/(C_M^{(1)})^2$. If $\phi\rho \leq \rho_0 \min(1, \eta)$, then for any $\rho' \geq \phi\rho$,*

$$r_M^2(\rho') > 2r_M^2(\rho) \text{ and } r^2(\rho') > 2r^2(\rho).$$

Proof. Since $r_M(\cdot)$, $r_Q(\cdot)$ and $r(\cdot)$ are non-decreasing, we only have to prove the result for $\rho' = \phi\rho$. Recall that $C_M^{(2)} \geq C_M^{(1)}$ (so $\phi \geq 4$). First note that if $N \geq \zeta'd$ then $r_Q(\rho) = 0$ and so the second claim follows from the first one since $r(\rho) = r_M(\rho)$ in this case. And when $N \leq \zeta'd$, one has $r_Q^2(\phi\rho) = \phi^2 r_Q(\rho) \geq 16r_Q(\rho) > 2r_Q(\rho)$ (because in that case $r_Q(\rho)^2 > 0$). Therefore the second claim is a straightforward consequence of the first one. So it only remains to study the behavior of $r_M^2(\cdot)$. For $\rho < \phi\rho \leq \rho_0$, r_M is given by one of the two last expressions of (1.9).

First assume that $(\phi\rho)^2 N \leq \sigma^2 \log d$ then

$$r_M^2(\phi\rho) \geq C_M^{(1)} \phi\rho\sigma \sqrt{\log(ed/N)} > 2C_M^{(2)} \rho\sigma \sqrt{\log(ed/N)} \geq 2r_M^2(\rho).$$

Now, assume that $\sigma^2 \log d \leq \rho^2 N \leq (\phi\rho)^2 N$ then

$$r_M^2(\phi\rho) = C_M \phi\rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{\phi\rho\sqrt{N}}\right)}$$

for some $C_M \in [C_M^{(1)}, C_M^{(2)}]$. One has that for all $x \geq \phi e$, $\log(x/\phi) > \log(x)/\phi$, and, since $\rho_0 = \sigma d/(\eta\sqrt{N})$, the assumption $\phi\rho \leq \rho_0 \min(1, \eta)$ guarantees that $e\sigma d/\rho\sqrt{N} \geq \phi e$. Therefore, we have:

$$r_M^2(\phi\rho) > C_M^{(1)} \phi\rho\sigma \sqrt{\frac{1}{\phi N} \log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)} \geq 2C_M^{(2)} \rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)} \geq 2r_M^2(\rho).$$

Finally, when $\rho^2 N \leq \sigma^2 \log d \leq (\phi\rho)^2 N$, since r_M is increasing, it is clear considering the two previous cases that one has again $r_M^2(\phi\rho) > (\sqrt{\phi}C_M^{(1)}/C_M^{(2)})r_M^2(\rho)$, so $r_M^2(\phi\rho) > 2r_M^2(\rho)$. \blacksquare

Lemma 2.4 *Let $\nu > 0$. If $\nu \geq 1$ then $r_M(\nu\rho) \leq \sqrt{\nu}r_M(\rho)$ and $r(\nu\rho) \leq \nu r(\rho)$. If $\nu \leq 1$ then $r_M(\nu\rho) \geq \sqrt{\nu}r_M(\rho)$ and $r(\nu\rho) \geq \nu r(\rho)$.*

Proof. It is clear that $r_Q(\nu\rho) = \nu r_Q(\rho)$, because $\ell^*(\nu\rho B_1^d \cap \nu r_Q(\rho) B_2^d) = \nu \ell^*(\rho B_1^d \cap r_Q(\rho) B_2^d)$. As for $r_M(\nu\rho)$, for $\nu \geq 1$, one has $\sigma \ell^*(\nu\rho B_1^d \cap \sqrt{\nu}r_M(\rho) B_2^d) \leq \sigma \nu \ell^*(\rho B_1^d \cap r_M(\rho) B_2^d) = \nu \eta \sqrt{N} r_M(\rho)^2 = \eta \sqrt{N} (\sqrt{\nu}r_M(\rho))^2$. So $r_M(\nu\rho) \leq \sqrt{\nu}r_M(\rho)$.

As for the case $\nu \leq 1$, then $1/\nu \geq 1$, and it suffices to write that $r_M(\rho) = r_M((1/\nu)\nu\rho) \leq (1/\nu)r_M(\nu\rho)$ to get the result (and still $r_Q(\rho/\nu) = r_Q(\rho)/\nu$). \blacksquare

2.2.2 Partition of \mathbb{R}^d into three zones

Recall that $\rho^* = \max\left(10, 8(C_M^{(2)})^2 / ((C_M^{(1)})^2 \eta) + 1\right) \|t^*\|_1$, that ρ_0 is the smallest ρ such that $r_M(\rho) = r_M(\rho_0)$ for all $\rho \geq \rho_0$ and that $K_0 = \min\{k \in \mathbb{N} : 2^k \rho^* \geq 2\rho_0\}$.

Hereunder, K_0 is the number of shelves used to partition the intermediate ‘‘peeling zone’’ defined below. We use ρ^* and K_0 to construct a partition of \mathbb{R}^d into three main zones:

- the ‘‘central zone’’ $t^* + \rho^* B_1^d$,
- the intermediate ‘‘peeling zone’’ $\{t \in \mathbb{R}^d : \rho^* < \|t - t^*\|_1 \leq 2^{K_0} \rho^*\}$ (recall that $2^{K_0} \rho^* \simeq 2\rho_0$). This zone is considered only when $K_0 \geq 1$. We use a ‘‘peeling’’, i.e. a partition of this zone into K_0 sub-areas called the ‘‘shelves’’: $2^{k-1} \rho^* < \|t - t^*\|_1 \leq 2^k \rho^*$, for $k = 1, \dots, K_0$.
- the ‘‘exterior zone’’ $\{t \in \mathbb{R}^d : \|t - t^*\|_1 > 2^{K_0} \rho^*\}$, on which $r_M(\cdot)$ is constant.

Our main objective is now to show that, on the event Ω_0 , \hat{t} belongs to the central zone.




Figure 2: $P_N \mathcal{M}_{t-t^*}$ is dominated by $P_N \mathcal{Q}_{t-t^*}$ in (A) and by \mathcal{R}_{t,t^*} in (B). The regularization function $\Psi(\cdot)$ is designed in order to dominate $P_N \mathcal{M}_{t-t^*}$ in (B).

2.2.3 Locating \hat{t} in the central zone on the event Ω_0

To show that, on the event Ω_0 , \hat{t} belongs to the central zone, it is enough to show that any t outside this area satisfies $P_N \mathcal{L}_t^\Psi > 0$ where we recall that in our case,

$$P_N \mathcal{L}_t^\Psi := \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t \rangle)^2 + c_0 r^2 (\|t\|_1) \right) - \left(\frac{1}{N} \sum_{i=1}^N (Y_i - \langle X_i, t^* \rangle)^2 + c_0 r^2 (\|t^*\|_1) \right).$$

This will indeed prove that $\|\hat{t} - t^*\|_1 \leq \rho^*$ since $P_N \mathcal{L}_{\hat{t}}^\Psi \leq 0 = P_N \mathcal{L}_{t^*}^\Psi$.

Let $t \in \mathbb{R}^d$ be outside the central zone – i.e. $\|t - t^*\|_1 > \rho^*$. First note that t satisfies $\|t\|_1 \geq \|t - t^*\|_1 - \|t^*\|_1 \geq \rho^* - \|t^*\|_1 \geq 9\|t^*\|_1$. Therefore, we have $\mathcal{R}_{t', t^*} = \Psi(\|t\|_1) - \Psi(\|t^*\|_1) = c_0 r^2 (\|t\|_1) - c_0 r^2 (\|t^*\|_1) > 0$ and $(8/9)\|t\|_1 \leq \|t - t^*\|_1 \leq (10/9)\|t\|_1$.

From now on, we place ourselves on the event Ω_0 introduced in Section 2.1.3.

First suppose that $\rho^* < 2\rho_0$, then both the “intermediate peeling zone” and the “exterior zone” must be considered. For t in any of these two areas, one has $\mathcal{R}_{t', t^*} \geq 0$.

Let us begin by the “peeling zone”. Consider t and $k \geq 1$ such that

$$2^{k-1} \rho^* < \|t - t^*\|_1 \leq 2^k \rho^* \text{ and } \|t - t^*\|_1 \leq 2^{K_0} \rho^*.$$

We recall that $2^{K_0-1} \rho^* < 2\rho_0 \leq 2^{K_0} \rho^*$ and ρ_0 is the smallest radius such that $r_M(\rho) = r_M(\rho_0)$ for all $\rho \geq \rho_0$.

One possibility is that $\|t - t^*\|_2^2 \geq r^2(2^k \rho^*)$, then on Ω^* one has $P_N \mathcal{Q}_{t-t^*} \geq \|t - t^*\|_2^2/2$ and, on A_k , one has $|P_N \mathcal{M}_{t-t^*}| \leq \|t - t^*\|_2^2/4$. Hence, on Ω_0 , $P_N \mathcal{L}_t^\Psi \geq \|t - t^*\|_2^2/4 + \mathcal{R}_{t', t^*} > 0$.

Let us tackle now the other case: $\|t - t^*\|_2^2 < r^2(2^k \rho^*)$. We will show that in this situation, if c_0 is large enough, $\mathcal{R}_{t, t^*}/4 > |P_N \mathcal{M}_{t-t^*}|$.

As $\|t\|_1 > 4\|t^*\|_1 (C_M^{(2)})^2 / (C_M^{(1)})^2$ (t is not in the “central zone”) and $4\|t^*\|_1 (C_M^{(2)})^2 / (C_M^{(1)})^2 \leq \eta \rho_0$ (since $2\rho_0 \geq \rho^* > 8\|t^*\|_1 (C_M^{(2)})^2 / ((C_M^{(1)})^2 \eta)$), by Lemma 2.3 one has that $c_0 r^2 (\|t\|_1) \geq 2c_0 r^2 (\|t^*\|_1)$. Thus, $\mathcal{R}_{t, t^*} \geq c_0 r^2 (\|t\|_1)/2$. So, thanks to Lemma 2.4, and since

$$\frac{\|t\|_1}{2^k \rho^*} = \frac{\|t - t^*\|_1}{2^k \rho^*} \frac{\|t\|_1}{\|t - t^*\|_1} > \frac{1}{2} \cdot \frac{9}{10} > \frac{2}{5},$$

one has

$$\mathcal{R}_{t, t^*} \geq \frac{1}{2} c_0 r^2 (\|t\|_1) \geq \frac{1}{2} \cdot \left(\frac{\|t\|_1}{2^k \rho^*} \right)^2 c_0 r^2 (2^k \rho^*) > (2/25) c_0 r^2 (2^k \rho^*).$$

In addition, we have $\|t - t^*\|_1 \leq 2^k \rho^*$ and $\|t - t^*\|_2^2 < r^2(2^k \rho^*)$, hence, on the event A_k , $|P_N \mathcal{M}_{t-t^*}| \leq r^2(2^k \rho^*)/4$. As a consequence, for $c_0 > 13$ one has

$$\frac{\mathcal{R}_{t, t^*}}{4} - |P_N \mathcal{M}_{t-t^*}| > \frac{2}{25} \cdot c_0 r^2 (2^k \rho^*) \cdot \frac{1}{4} - \frac{1}{4} r^2 (2^k \rho^*) > 0. \quad (2.10)$$

A fortiori, $\mathcal{R}_{t', t^*} - |P_N \mathcal{M}_{t-t^*}| > 0$ and as $P_N \mathcal{Q}_{t-t^*} \geq 0$, this implies in particular that $P_N \mathcal{L}_t^\Psi > 0$.

To sum up, we proved that for all t in the peeling zone, we have

$$P_N \mathcal{L}_t^\Psi \geq P_N \mathcal{Q}_{t-t^*} - |P_N \mathcal{M}_{t-t^*}| + \mathcal{R}_{t', t^*}/4 > 0. \quad (2.11)$$

Let us now study the exterior zone. We will mainly use homogeneity arguments. Let $t \in \mathbb{R}^d$ be outside the ball $t^* + 2^{K_0} \rho^* B_1^d$. Define $t' \in t^* + 2^{K_0} \rho^* S_1^{d-1}$ such that $t - t^* = \alpha_t(t' - t^*)$ for some $\alpha_t \geq 1$. In particular, $\|t' - t^*\|_2 \geq 2r_0 > r_M(2^{K_0} \rho^*)$ so by Proposition 2.2, on the event A_{K_0} , one has $|P_N \mathcal{M}_{t'-t^*}| \leq \|t' - t^*\|_2^2/4$. We consider now two cases.

If $\|t' - t^*\|_2 \geq r_Q(\|t' - t^*\|_1)$, then by Lemma 2.1, on the event Ω^* , one has $P_N \mathcal{Q}_{t'-t^*} \geq \|t' - t^*\|_2^2/2$. So on the event Ω_0 , $P_N \mathcal{Q}_{t'-t^*} - |P_N \mathcal{M}_{t'-t^*}| > 0$ and

$$P_N \mathcal{Q}_{t-t^*} - |P_N \mathcal{M}_{t-t^*}| = \alpha_t^2 P_N \mathcal{Q}_{t'-t^*} - \alpha_t |P_N \mathcal{M}_{t'-t^*}| \geq \alpha_t (P_N \mathcal{Q}_{t'-t^*} - |P_N \mathcal{M}_{t'-t^*}|) > 0$$

therefore $P_N \mathcal{L}_t^\Psi > 0$.

On the contrary, if $\|t' - t^*\|_2 < r_Q(\|t' - t^*\|_1)$, as $\|t' - t^*\|_2 \geq 2r_0$, then $r_Q(\|t'\|_1) \geq 9r_Q(\|t' - t^*\|_1)/10 > 9r_0/5 = 9r_M(\|t'\|_1)/5 > r_M(\|t'\|_1)$ so $r(\|t'\|_1) = r_Q(\|t'\|_1)$. One has $\|t\|_1 = \|t^* + \alpha_t(t' - t^*)\|_1 \geq (4\alpha_t/5)\|t'\|_1$ since $\alpha_t \geq 1$ and $\|t' - t^*\|_1 \geq \rho^*$. So $r^2(\|t\|_1) \geq r_Q^2(\alpha_t 4\|t'\|_1/5) = 4^2 \alpha_t^2 r^2(\|t'\|_1)/5^2$. We have seen before that $r_Q(\|t'\|_1) \geq 9r_0/5 \geq 9r_M(\|t^*\|_1)/5$, and $r_Q(\|t'\|_1) \geq 9r_Q(\|t^*\|_1)$, so $r^2(\|t'\|_1) \geq 3r^2(\|t^*\|_1)$. As a consequence,

$$\begin{aligned} \mathcal{R}_{t,t^*} &= c_0 r^2(\|t\|_1) - c_0 r^2(\|t^*\|_1) \geq c_0 \left(\alpha_t^2 \frac{4^2}{5^2} r^2(\|t'\|_1) - (1/3) r^2(\|t'\|_1) \right) \\ &\geq \alpha_t \frac{c_0}{4} r^2(\|t'\|_1) \geq \alpha_t \mathcal{R}_{t',t^*}/4. \end{aligned}$$

Moreover, since $t - t^* = \alpha_t(t' - t^*)$, one has $P_N \mathcal{Q}_{t-t^*} = \alpha_t^2 P_N \mathcal{Q}_{t'-t^*}$ and $P_N \mathcal{M}_{t-t^*} = \alpha_t P_N \mathcal{M}_{t'-t^*}$. So, in the case $\|t' - t^*\|_2 < r_Q(\|t' - t^*\|_1)$, by (2.11) applied to t' ,

$$P_N \mathcal{L}_t^\Psi \geq \alpha_t (P_N \mathcal{Q}_{t'-t^*} - |P_N \mathcal{M}_{t'-t^*}| + \mathcal{R}_{t',t^*}/4) > 0.$$

Let us now consider the case $\rho^* > 2\rho_0$. In this situation, there is no need for the intermediate ‘‘peeling’’ zone. Let $t \in \mathbb{R}^d$ be outside the ball $t^* + \rho^* B_1^d$ and set $t' \in t^* + \rho^* S_1^{d-1}$ such that $t - t^* = \alpha_t(t' - t^*)$ with $\alpha_t \geq 1$. Then one can apply arguments similar to the peeling case (with $\|t' - t^*\|_1 = \rho^*$, but this time $k = K_0 = 0$), on the event $A_0 \cap \Omega^*$, to t' . If $\|t' - t^*\|_2 \geq r_Q(\|t' - t^*\|_1)$, then by Lemma 2.1, on the event Ω_0 one has $P_N \mathcal{Q}_{t'-t^*} \geq |P_N \mathcal{M}_{t'-t^*}|/2$. Conversely, if $\|t' - t^*\|_2 < r_Q(\|t' - t^*\|_1)$, then $r^2(\|t'\|_1) \geq 3r^2(\|t^*\|_1)$ so for c_0 big enough, in the same spirit as (2.10), one gets that $\mathcal{R}_{t',t^*}/4 - |P_N \mathcal{M}_{t'-t^*}| > 0$. So in both cases, $P_N \mathcal{Q}_{t-t^*} - |P_N \mathcal{M}_{t-t^*}| + \mathcal{R}_{t',t^*}/4 > 0$. The same argument as previously for the exterior zone, shows that $P_N \mathcal{L}_t^\Psi \geq \alpha_t (P_N \mathcal{Q}_{t'-t^*} - |P_N \mathcal{M}_{t'-t^*}| + \mathcal{R}_{t',t^*}/4)$, so $P_N \mathcal{L}_t^\Psi > 0$. As a conclusion, in the case $\rho^* > 2\rho_0$, we have for all $t \in \mathbb{R}^d$ satisfying $\|t - t^*\|_1 \geq \rho^* > 2\rho_0$ that $P_N \mathcal{L}_t^\Psi > 0$.

To sum up, on the event Ω_0 , any t outside the central zone satisfies $P_N \mathcal{L}_t^\Psi > 0$. Therefore, given that \hat{t} satisfies $P_N \mathcal{L}_{\hat{t}}^\Psi \leq 0$, we conclude that \hat{t} belongs to the central zone.

2.2.4 Conclusion of the proof of Theorem 1.4

On the event Ω_0 , $\hat{t} \in (t^* + \rho^* B_1^d)$. Hence, either $\|\hat{t} - t^*\|_2^2 \leq r^2(\rho^*)$ and the proof is over or $\|\hat{t} - t^*\|_2^2 > r^2(\rho^*)$. In the latter case, one has

$$P_N \mathcal{Q}_{\hat{t}-t^*} > \frac{1}{2} \|\hat{t} - t^*\|_2^2 \text{ and } |P_N \mathcal{M}_{\hat{t}-t^*}| < \frac{1}{4} \|\hat{t} - t^*\|_2^2$$

and so

$$0 \geq P_N \mathcal{L}_{\hat{t}}^\Psi \geq \frac{1}{4} \|\hat{t} - t^*\|_2^2 + c_0 r^2 (\|\hat{t}\|_1) - c_0 r^2 (\|t^*\|_1)$$

which implies $\|\hat{t} - t^*\|_2^2 \leq 4c_0 r^2 (\|t^*\|_1)$.

Thus, taking $\theta_0 = \max\left(100, (8(C_M^{(2)})^2 / (C_M^{(1)})^2 + 1)^2, 4c_0\right)$, with probability at least

$$\mathbb{P}[\Omega_0] \geq 1 - 4 \exp(-C_6 N) - 40 \exp(-C'_6 N r_M (\|t^*\|_1)^2 / \sigma^2)$$

one gets that in both cases, $R(\hat{t}) - R(t^*) = \|\hat{t} - t^*\|_2^2 \leq \theta_0 r^2 (\|t^*\|_1)$. Moreover, for $\|t^*\|_1 \geq \Delta_0 \sigma \sqrt{\log(ed)/N}$ for Δ_0 an absolute constant large enough, we have $\mathbb{P}[\Omega_0] \geq 3/4$. Given that $r^2 (\|t^*\|_1)$ is the minimax rate of convergence over $\|t^*\|_1 B_1^d$ (cf. Theorem 1.3), we conclude that $\Psi(\rho) = c_0 r^2(\rho)$ is indeed a minimax regularization function.

3 Technical material and proof of Proposition 1.6

3.1 Localization with balls and spheres

The next lemma shows that when the intersection is not trivial, localizing by intersecting an ℓ_1 ball with an ℓ_2 sphere or the corresponding full ℓ_2 ball is equivalent.

Lemma 3.1 *If $\rho \geq r$, then $\ell^*(\rho B_1^d \cap r S_2^{d-1}) = \ell^*(\rho B_1^d \cap r B_2^d)$. If $\rho < r$, then $\ell^*(\rho B_1^d \cap r B_2^d) = \ell^*(\rho B_1^d)$ and $\ell^*(\rho B_1^d \cap r S_2^{d-1}) = 0$.*

Proof. Since for any set $T \subset \mathbb{R}^d$, $\ell^*(T) = \ell^*(\text{conv}(T))$ (with $\text{conv}(T)$ denoting the convex hull of T), the result is a direct consequence of the fact that for $\rho < r$, $\text{conv}(\rho B_1^d \cap r B_2^d) = \rho B_1^d$ and $\text{conv}(\rho B_1^d \cap r S_2^{d-1}) = \emptyset$ (which are two obvious statements), and that if $\rho \geq r$, then $\text{conv}(\rho B_1^d \cap r S_2^{d-1}) = \text{conv}(\rho B_1^d \cap r B_2^d)$, which we prove now.

One inclusion is immediate, it remains to prove that $\text{conv}(\rho B_1^d \cap r B_2^d) \subset \text{conv}(\rho B_1^d \cap r S_2^{d-1})$. First,

$$\text{conv}(\rho B_1^d \cap r B_2^d) = \rho B_1^d \cap r B_2^d = \text{conv}\{t \in \mathbb{R}^d : \max(\|t\|_1/\rho, \|t\|_2/r) = 1\}$$

so it only remains to show that $\{t \in \mathbb{R}^d : \max(\|t\|_1/\rho, \|t\|_2/r) = 1\} \subset \text{conv}(\rho B_1^d \cap r S_2^{d-1})$.

First, remark that $\{t \in \mathbb{R}^d : \|t\|_1/\rho \leq 1, \|t\|_2/r = 1\}$ is included in $\rho B_1^d \cap r S_2^{d-1}$. Let us now consider the set $\{t \in \mathbb{R}^d : \|t\|_1/\rho = 1, \|t\|_2/r < 1\}$ and consider an element t in it. We denote by e_1, \dots, e_d the canonical basis of \mathbb{R}^d and we recall that each face of ρB_1^d is the convex hull of its vertices. So, since $t \in \rho S_1^{d-1}$, there exist $b_1 \in \{\rho e_1, -\rho e_1\}, b_2 \in \{\rho e_2, -\rho e_2\}, \dots, b_d \in \{\rho e_d, -\rho e_d\}$ such that t is in the convex hull of $\{b_1, \dots, b_d\}$: there exist non-negative coefficients μ_1, \dots, μ_d such that $\sum_{j=1}^d \mu_j (b_j - t) = 0$. Consider for each $j \in \{1, \dots, d\}$ the mapping

$$f_j : x \in [0, 1] \mapsto \|t + x(b_j - t)\|_2/r - \|t + x(b_j - t)\|_1/\rho$$

This mapping is continuous, $f_j(0) < 0$ (t is in ρS_1^{d-1} and $r B_2^d$ but not in $r S_2^{d-1}$), and $f_j(1) \geq 0$ (because $\|b_j\|_2 = \|b_j\|_1 = \rho$ and $\rho \geq r$). So there exists x_j in $(0, 1]$ such that $f_j(x_j) = 0$, which means that $t + x_j(b_j - t) \in \rho S_1^{d-1} \cap r S_2^{d-1}$. One has that

$$\sum_{j=1}^d \frac{\mu_j}{x_j} (t + x_j(b_j - t) - t) = \sum_{j=1}^d \frac{\mu_j}{x_j} x_j (b_j - t) = \sum_{j=1}^d \mu_j (b_j - t) = 0.$$

As a consequence, t is in the convex hull of the vectors $t + x_j(b_j - t) \in \rho S_1^{d-1} \cap r S_2^{d-1}$, $j \in \{1, \dots, d\}$ which achieves the proof. \blacksquare

3.2 Control of the probability estimate

Lemma 3.2 *Set $\eta = 1/(16\sqrt{2})$. For all $\nu > 0$, we have*

$$\sum_{k=0}^{K_0} \exp\left(-\nu r_M(2^k \rho^*)^2\right) \leq \frac{10}{1 - \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)} \exp\left(-\frac{C_M^{(1)}}{C_M^{(2)}} \nu r_M(\rho^*)^2\right).$$

Proof. First, the terms of the sum are non-increasing (remember that r_M is a non-decreasing function). So skipping the last terms will not change the order of magnitude:

$$\sum_{k=0}^{K_0} \exp\left(-\nu r_M(2^k \rho^*)^2\right) \leq \max\left(10 \exp\left(-\nu r_M(\rho^*)^2\right), 10 \sum_{k=0}^{K_0-9} \exp\left(-\nu r_M(2^k \rho^*)^2\right)\right).$$

One has $10 \exp\left(-\nu r_M(\rho^*)^2\right) > 10 \sum_{k=0}^{K_0-9} \exp\left(-\nu r_M(2^k \rho^*)^2\right)$ when $K_0 \leq 8$. Let us now assume that $K_0 \geq 9$. We study the sum $\sum_{k=0}^{K_0-9} \exp\left(-\nu r_M(2^k \rho^*)^2\right)$.

In order to get rid of the “range” $[C_M^{(1)}, C_M^{(2)}]$ in the definition of $r_M(\cdot)$, notice that $\sum_{k=0}^{K_0-9} \exp\left(-\nu r_M(2^k \rho^*)^2\right) \leq \sum_{k=0}^{K_0-9} a_k$ with:

$$a_k := \begin{cases} \exp\left(-\nu C_M^{(1)} 2^k \rho^* \sigma \sqrt{\frac{\log(ed)}{N}}\right) & \text{if } (2^k \rho^*)^2 N \leq \sigma^2 \log(d) \\ \exp\left(-\nu C_M^{(1)} 2^k \rho^* \sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{2^k \rho^* \sqrt{N}}\right)}\right) & \text{otherwise.} \end{cases}$$

We emphasize that the sum goes only up to $2^{K_0} \rho^*$, which excludes the constant third form of $r_M(\cdot)$.

Applying a second time the “range” $[C_M^{(1)}, C_M^{(2)}]$ in the bounds on r_M allows to bound a_0 in terms of $r_M(\rho^*)^2$: $a_0 \leq \exp\left(-\frac{C_M^{(1)}}{C_M^{(2)}} \nu r_M(\rho^*)^2\right)$. Therefore, in the following we will bound $\sum_{k=0}^{K_0-9} a_k$ with respect to a_0 , and then get back to $r_M(\rho^*)^2$.

We now prove that there exists α independent on k (but dependent on the other parameters) such that for any $k \leq K_0 - 9$, $a_{k+1}/a_k \leq \alpha < 1$.

If $(2^{k+1} \rho^*)^2 \leq \sigma^2 \log(d)/N$ then

$$\begin{aligned} \frac{a_{k+1}}{a_k} &= \exp\left(-\nu C_M^{(1)} 2^{k+1} \rho^* \sigma \sqrt{\frac{\log(ed)}{N}} + \nu C_M^{(1)} 2^k \rho^* \sigma \sqrt{\frac{\log(ed)}{N}}\right) \\ &= \exp\left(-\nu C_M^{(1)} 2^k \rho^* \sigma \sqrt{\frac{\log(ed)}{N}}\right) \leq \exp(-\nu C_M^{(1)} \rho^* \beta_1) \end{aligned}$$

with $\beta_1 = \sigma \sqrt{\log(ed)/N} \geq \sigma/\sqrt{N}$.

As for the case $(2^k \rho^*)^2 \leq \sigma^2 \log(d)/N < (2^{k+1} \rho^*)^2$ (which can only occur when $d > 1$), then one has $2^{k+1} \rho^* \leq 2\sigma\sqrt{\log d}/\sqrt{N}$ and so

$$\begin{aligned} \frac{a_{k+1}}{a_k} &= \exp\left(-\nu C_M^{(1)} 2^{k+1} \rho^* \sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{2^{k+1} \rho^* \sqrt{N}}\right)} + \nu C_M^{(1)} 2^k \rho^* \sigma \sqrt{\frac{\log(ed)}{N}}\right) \\ &\leq \exp\left(-\nu C_M^{(1)} 2^k \rho^* \frac{\sigma}{\sqrt{N}} \left(2\sqrt{\log\left(\frac{ed}{2\sqrt{\log d}}\right)} - \sqrt{\log ed}\right)\right) \\ &\leq \exp\left(-\nu C_M^{(1)} 2^k \rho^* \frac{\sigma}{2\sqrt{N}}\right) \leq \exp\left(-\nu C_M^{(1)} \rho^* \frac{\sigma}{2\sqrt{N}}\right) \end{aligned}$$

The second inequality relies on a straightforward analysis fact:

$$\forall d \geq 2, \quad \left(2\sqrt{\log\left(\frac{ed}{2\sqrt{\log d}}\right)} - \sqrt{\log ed}\right) > \frac{1}{2}$$

Let us tackle now the case $(2^k \rho^*)^2 > \sigma^2 \log(d)/N$. We have $a_{k+1}/a_k = \exp(b_k \nu C_M^{(1)} 2^k \rho^* \sigma / \sqrt{N})$ where

$$b_k := -2\sqrt{\log\left(\frac{e\sigma d}{2^{k+1} \rho^* \sqrt{N}}\right)} + \sqrt{\log\left(\frac{e\sigma d}{2^{k+1} \rho^* \sqrt{N}}\right)} + \log(2).$$

Since $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ for all $x, y \geq 0$, one has (still for $k \leq K_0 - 9$):

$$b_k \leq \sqrt{\log(2)} - \sqrt{\log\left(\frac{e\sigma d}{2^{k+1} \rho^* \sqrt{N}}\right)} \leq \sqrt{\log(2)} - \sqrt{\log\left(\frac{e\sigma d}{2^{K_0-8} \rho^* \sqrt{N}}\right)} \leq (1 - \sqrt{2})\sqrt{\log(2)}$$

because $2^{K_0-8} \rho^* \leq 2^{-6} \rho_0 \leq 2^{-6} \sigma d / (\eta \sqrt{N}) \leq \sigma d / (2\sqrt{N})$ (we recall that $\rho_0 = \sigma d / (\eta \sqrt{N})$ and $\eta = 1/(16\sqrt{2})$). It follows that

$$\frac{a_{k+1}}{a_k} = \exp\left(\nu C_M^{(1)} 2^k \rho^* \frac{\sigma}{\sqrt{N}} b_k\right) \leq \exp\left(-\nu C_M^{(1)} \rho^* \frac{\sigma}{\sqrt{N}} (-b_k)\right) \leq \exp(-\nu C_M^{(1)} \rho^* \beta_2)$$

with $\beta_2 = (\sqrt{2} - 1)\sqrt{\log(2)}\sigma/\sqrt{N} \geq \sigma/(4\sqrt{N})$. Then, we conclude that for all $k \leq K_0 - 9$, $a_{k+1}/a_k \leq \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)$ and so

$$\sum_{k=0}^{K_0-9} a_k \leq a_0 \sum_{k=0}^{K_0-9} \exp\left(-k \nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right) = a_0 \frac{1 - \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)}{1 - \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)}.$$

Finally, the result follows, since $a_0 \leq \exp\left(-(C_M^{(1)}/C_M^{(2)})\nu r_M(\rho^*)^2\right)$ and

$$\frac{1 - \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)}{1 - \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)} \leq \frac{1}{1 - \exp\left(-\nu C_M^{(1)} \frac{\sigma}{4\sqrt{N}} \rho^*\right)}.$$

■

3.3 Proof of Proposition 1.6

The proof of Proposition 1.6 relies on key ideas developed in minimax theory. We refer the reader to [39] for a state of the art in minimax theory.

Let \hat{t} satisfy the properties of Proposition 1.6, i.e. a procedure adaptive on the finite set $\{\pm\rho e_1, \dots, \pm\rho e_d\}$ and let $t^* \in (\rho/2)B_1^d$. Denote $\Lambda = \{t^*, \pm\rho e_1, \dots, \pm\rho e_d\} = \{t_0^*, t_1^*, \dots, t_{2d}^*\}$ so that $t_0^* = t^*$. It is straightforward to check that Λ is a $\rho/2$ -separated set in ρB_1^d w.r.t. ℓ_2^d .

Now, let us define the test statistics

$$\hat{\phi} \in \operatorname{argmin}_{j \in \{0, \dots, 2d\}} \|\hat{t}_j^* - \hat{t}\|_2.$$

One has for all $j \in \{0, 1, \dots, 2d\}$ that, if $\hat{\phi} \neq j$ then $\|\hat{t} - t_j^*\|_2 \geq \rho/4$. Indeed, if $\hat{\phi} \neq j$ then there exists $k \in \{0, 1, \dots, 2d\} \setminus \{j\}$ such that $\|t_k^* - \hat{t}\|_2 \leq \|t_j^* - \hat{t}\|_2$. If $\|t_k^* - \hat{t}\|_2 \geq \rho/4$ then the result holds otherwise $\|t_k^* - \hat{t}\|_2 < \rho/4$ and so $\|\hat{t}_j^* - \hat{t}\|_2 \geq \|t_j^* - t_k^*\|_2 - \|t_k^* - \hat{t}\|_2 \geq \rho/4$. Therefore, we have, for all $\tau > 0$

$$\begin{aligned} \mathbb{P}_{t_0^*} [\|\hat{t} - t_0^*\|_2 \geq \rho/4] &\geq \mathbb{P}_{t_0^*} [\hat{\phi} \neq 0] = \sum_{j=1}^{2d} \mathbb{P}_{t_0^*} [\hat{\phi} = j] \geq \sum_{j=1}^{2d} \tau \mathbb{P}_{t_j^*} \left[\hat{\phi} = j \text{ and } \frac{d\mathbb{P}_{t_0^*}}{d\mathbb{P}_{t_j^*}} \geq \tau \right] \\ &\geq \tau \sum_{j=1}^{2d} \mathbb{P}_{t_j^*} [\hat{\phi} = j] - \mathbb{P}_{t_j^*} \left[\frac{d\mathbb{P}_{t_0^*}}{d\mathbb{P}_{t_j^*}} < \tau \right] \geq \tau \sum_{j=1}^{2d} \mathbb{P}_{t_j^*} [\|\hat{t} - t_j^*\|_2 < \rho/4] - \mathbb{P}_{t_j^*} \left[\frac{d\mathbb{P}_{t_0^*}}{d\mathbb{P}_{t_j^*}} < \tau \right]. \end{aligned} \quad (3.1)$$

It follows from the adaptation property of \hat{t} over $\{\pm\rho e_1, \dots, \pm\rho e_d\}$ that for every $j \in \{1, \dots, 2d\}$,

$$\mathbb{P}_{t_j^*} [\|\hat{t} - t_j^*\|_2 < \rho/4] \geq \mathbb{P}_{t_j^*} [\|\hat{t} - t_j^*\|_2^2 < \chi_1 r^2(\rho)] \geq 3/4 \quad (3.2)$$

when $\chi_1 r^2(\rho) \leq \rho^2/16$. Let $j \in \{1, \dots, 2d\}$. Following the same argument as in Proposition 2.3 in [39] (based on second Pinsker inequality), we obtain

$$\begin{aligned} \mathbb{P}_{t_j^*} \left[\frac{d\mathbb{P}_{t_0^*}}{d\mathbb{P}_{t_j^*}} \geq \tau \right] &= \mathbb{P}_{t_j^*} \left[\frac{d\mathbb{P}_{t_j^*}}{d\mathbb{P}_{t_0^*}} \leq \frac{1}{\tau} \right] = 1 - \mathbb{P}_{t_j^*} \left[\log \frac{d\mathbb{P}_{t_j^*}}{d\mathbb{P}_{t_0^*}} > \log(1/\tau) \right] \\ &\geq 1 - \frac{1}{\log(1/\tau)} \int \left(\log \frac{d\mathbb{P}_{t_j^*}}{d\mathbb{P}_{t_0^*}} \right)_+ d\mathbb{P}_{t_j^*} \geq 1 - \frac{1}{\log(1/\tau)} \left[K(\mathbb{P}_{t_j^*}, \mathbb{P}_{t_0^*}) + \sqrt{2K(\mathbb{P}_{t_j^*}, \mathbb{P}_{t_0^*})} \right] \end{aligned}$$

where $K(\mathbb{P}_{t_j^*}, \mathbb{P}_{t_0^*})$ denotes the Kullback-Leiber divergence between $\mathbb{P}_{t_j^*}$ and $\mathbb{P}_{t_0^*}$. Since the noise is Gaussian and independent of X in (1.7) and X is isotropic, we have $K(\mathbb{P}_{t_j^*}, \mathbb{P}_{t_0^*}) = N\|t_j^* - t_0^*\|_2^2/(2\sigma^2) \leq N9\rho^2/(8\sigma^2)$. Hence, if $\rho \leq 2\sigma\sqrt{\log(1/\tau)/(96N)}$ and $\log(1/\tau) \geq 544/225$, one has

$$\mathbb{P}_{t_j^*} \left[\frac{d\mathbb{P}_{t_0^*}}{d\mathbb{P}_{t_j^*}} \geq \tau \right] \geq \frac{3}{4}.$$

Using the latter result together with (3.2) in (3.1) for the values $\tau = 1/(2d)$, we obtain that

$$\mathbb{P}_{t_0^*} [\|\hat{t} - t_0^*\|_2 \geq \rho/4] \geq \frac{1}{2}. \quad \blacksquare$$

4 Minimax regularization function in the fixed design setup

In this section, we consider the Gaussian regression model with fixed design. Our aim is to design a minimax regularization function in this setup. We are therefore given a deterministic $N \times d$ design matrix \mathbb{X} , whose i -th row vector is denoted by $X_i \in \mathbb{R}^d$, and we observe N noisy projections of a vector $t^* \in \mathbb{R}^d$, $Y_i = \langle X_i, t^* \rangle + \xi_i$, $i = 1, \dots, N$, where the noises ξ_1, \dots, ξ_N are independent centered Gaussian variables with variances σ^2 . The data $(Y_i, X_i)_{i=1}^N$ are used to construct estimators of t^* and the only difference with the previous random design setup is that the X_i 's are deterministic vectors whereas so far they were random vectors. We will use most of the time the matrix form of the data: $Y = \mathbb{X}t + \xi$ where $Y = (Y_i)_{i=1}^N$ and $\xi \sim \mathcal{N}(0, \sigma^2 I_N)$ where I_N is the $N \times N$ identity matrix. Note that the fixed design setup is usually considered in signal processing over a finite grid or in experiences where the statistician can chose in advance the design of an experience and then observed an output.

In order to design a minimax regularization function in this setup, we need to adapt the definitions introduced in Section 1 to the fixed design case. We first start with the definition of a minimax rate of convergence over a subset of \mathbb{R}^d . We use the empirical (or normalized) Euclidean inner product and norm: for all $u, v \in \mathbb{R}^N$,

$$\langle u, v \rangle_{L_N^2} = \frac{1}{N} \sum_{i=1}^N u_i v_i \text{ and } \|u\|_{L_N^2} = \sqrt{\frac{1}{N} \sum_{i=1}^N u_i^2}$$

and the associated unit ball $B_{L_N^2} = \{u \in \mathbb{R}^N : \|u\|_{L_N^2} \leq 1\}$.

Definition 4.1 *Let $T \subset \mathbb{R}^d$, \mathbb{X} denote a (deterministic) $N \times d$ design matrix and ξ be a centered Gaussian random vector in \mathbb{R}^N with covariance matrix $\sigma^2 I_N$. For all $t^* \in T$, define the random vector $Y^{t^*} = \mathbb{X}t^* + \xi$ and denote by $\mathcal{Y}^T := \{Y^{t^*} : t^* \in T\}$ the set of all such random vectors.*

*Let \hat{t}_N be a statistics from \mathbb{R}^N to \mathbb{R}^d . Let $0 < \delta_N < 1$ and $\zeta_N > 0$. We say that \hat{t}_N **performs with accuracy ζ_N and confidence $1 - \delta_N$ relative to the set of targets \mathcal{Y}^T** , if for all $t^* \in T$, with probability, w.r.t. to a vector Y distributed as Y^{t^*} , at least $1 - \delta_N$, $\|\mathbb{X}(\hat{t}_N(Y) - t^*)\|_{L_N^2}^2 \leq \zeta_N$.*

*We say that \mathcal{R}_N is a **minimax rate of convergence over T for the confidence $1 - \delta_N$** if the two following hold:*

1. *there exists a statistics \hat{t}_N which performs with accuracy \mathcal{R}_N and confidence $1 - \delta_N$ relative to the set of targets \mathcal{Y}^T*
2. *there exists an absolute constant $g'_0 > 0$ such that if \tilde{t}_N is a statistics which performs with accuracy ζ_N and confidence $1 - \delta_N$ relative to the set of targets \mathcal{Y}^T then necessarily $\zeta_N \geq g'_0 \mathcal{R}_N$.*

Note that we use the empirical L_N^2 -metric $\|\mathbb{X} \cdot\|_{L_N^2}$ (to the square) with respect to the design \mathbb{X} as a measure of performances of estimators in Definition 4.1. The reason we do so is that it is the natural counterpart to the random design case – that is when \mathbb{X} is a standard

Gaussian matrix then $\mathbb{E}\|\mathbb{X}t\|_{L_N^2}^2 = \|t\|_2^2$ and the ℓ_2^d -norm is the metric used to measure the performance of estimators in the random design setup – and that it is the natural metric associated to the prediction of Y problem given that if $R(t) = \mathbb{E}\|Y - \mathbb{X}t\|_{L_N^2}^2$ is the risk of t for all $t \in \mathbb{R}^d$ then we have for any estimator \hat{t}_N , $R(\hat{t}_N) = R(t^*) + \|\mathbb{X}(\hat{t}_N - t^*)\|_{L_N^2}^2$. So that predicting Y is the same problem as estimating t^* with respect to the empirical $\|\mathbb{X} \cdot\|_{L_N^2}$ metric.

Now, we adapt the definition of a minimax regularization function to the fixed design setup in the next definition.

Definition 4.2 *Let $\|\cdot\|$ be a norm on \mathbb{R}^d , $T \subset \mathbb{R}^d$ and $0 < \delta_N < 1$. Let us consider the following RERM for some given function $\Psi : \mathbb{R}_+ \rightarrow \mathbb{R}$:*

$$\hat{t} \in \operatorname{argmin}_{t \in \mathbb{R}^d} \left(\|Y - \mathbb{X}t\|_{L_N^2}^2 + \Psi(\|t\|) \right)$$

constructed from a $N \times d$ deterministic matrix \mathbb{X} and a random vector $Y = \mathbb{X}t^ + \xi$, with $\xi \sim \sigma\mathcal{N}(0, I_N)$. We say that Ψ is a **minimax regularization function for the norm $\|\cdot\|$ over T for the confidence $1 - \delta_N$** , if there exists an absolute constant $g'_1 > 0$ such that for all $t^* \in T$, the RERM \hat{t} is such that with probability at least $1 - \delta_N$, $\|\hat{t} - t^*\|_{L_N^2}^2 \leq g'_1 \mathcal{R}_N$, where \mathcal{R}_N is the minimax rate of convergence over $\{t \in \mathbb{R}^d : \|t\| \leq \|t^*\|\}$.*

The statistical bounds one can prove in the fixed design setup depend generally on the property of the design matrix \mathbb{X} . Many different assumptions have been introduced during the last two decades in high-dimensional statistics and we refer to [43] for some of them. In particular, norm preserving properties like the RIP or weaker assumption on the restricted eigenvalues like the REC or CC have played an important role in statistics (cf. [12, 41, 6, 18, 42, 8]). In this paper, we assume that \mathbb{X} satisfies the ‘‘Restricted Isometry Property’’. It appears that this condition is equivalent (up to constants) to the property satisfied by a standard Gaussian matrix as in Lemma 2.1.

Assumption 4.1 (RIP(s)) *If $N < d$ and $N/\log(ed/N) > 1$ then we set $s = N/\log(ed/N)$. We assume that all t in $\Sigma_s := \{x \in \mathbb{R}^d : \|x\|_0 \leq s\}$ is such that*

$$\frac{1}{2}\|t\|_2 \leq \|\mathbb{X}t\|_{L_N^2} \leq \frac{3}{2}\|t\|_2 \quad (4.1)$$

where $\|x\|_0$ is the size of the support of x . If $N \geq d$, we assume that (4.1) is satisfied for all $t \in \mathbb{R}^d$.

Note that in the high-dimensional case, i.e. $d > N$, only the situation $N/\log(ed/N) > 1$ is considered in Assumption 4.1 to avoid the ultra-high dimensional phenomena discovered in [46]. RIP was introduced in [11] and it has been widely used and discussed (cf. for example [15], [1] or [16]), in particular in the field of Compressed Sensing. From our perspective, we use this result for two reasons: 1) the minimax results over ℓ_1^d balls we need to develop for our proof of minimax regularization function has been obtained in the fixed design under this condition (or an equivalent one) in [32]; 2) the complexity parameter that we will be

using in the fixed design setup have been computed under very general design matrix \mathbb{X} in [4] but they were only proved to be optimal under the RIP assumption. Our main result covers more general design matrix than the one satisfying RIP but it turns out that those results do not allow to conclude on the minimax optimality of the associated regularization function beyond the RIP case; moreover and to our knowledge no sharp closed form are available for the computation of this regularization function beyond the RIP case.

4.0.1 The multiplier process and its associated fixed point

Our analysis is based upon the study of the same regularized excess empirical risk quantity as in the random design section: for all $t \in \mathbb{R}^d$,

$$P_N \mathcal{L}_t^\Psi := \left(\|Y - \mathbb{X}t\|_{L_N^2}^2 + \Psi(t) \right) - \left(\|Y - \mathbb{X}t^*\|_{L_N^2}^2 + \Psi(t^*) \right).$$

We use the same quadratic / multiplier decomposition as in the random design case: for all $t \in \mathbb{R}^d$, $P_N \mathcal{L}_t^\Psi = P_N \mathcal{Q}_{t-t^*} + P_N \mathcal{M}_{t-t^*} + \mathcal{R}_{t-t^*}$ where

- $P_N \mathcal{Q}_{t-t^*} := \|\mathbb{X}(t-t^*)\|_{L_N^2}^2$ is the quadratic part
- $P_N \mathcal{M}_{t-t^*} := 2\langle \xi, \mathbb{X}(t^* - t) \rangle$ is the multiplier part
- $\mathcal{R}_{t,t^*} := \Psi(t) - \Psi(t^*)$ is the regularization part.

Contrary to the random design case, in the fixed design setup the only source of randomness is the Gaussian noise ξ , in particular the quadratic process is fully deterministic. Therefore, there is no need to define a fixed point similar to r_Q for a control on the quadratic process. The only fixed point we introduce is a version of the previous multiplier fixed point $r_M(\cdot)$ adapted to the fixed design setup: for $\eta' = 1/8$, let

$$r_{\mathbb{X}}(\rho) := \inf \left(r > 0 : \sigma \ell^* \left(\frac{\rho}{\sqrt{N}} \mathbb{X}B_1^d \cap rB_2^N \right) \leq \eta' r^2 \sqrt{N} \right) \quad (4.2)$$

where $\mathbb{X}B_1^d := \{\mathbb{X}t : t \in B_1^d\}$ and B_2^N is the unit ball in ℓ_2^N .

Define r'_0 as the non-zero solution to the equation $\sigma \ell^*(rB_2^N \cap \text{Im}(\mathbb{X})) = \eta' r^2 \sqrt{N}$, where $\text{Im}(\mathbb{X})$ is the image of \mathbb{X} in \mathbb{R}^d , i.e.

$$r'_0 = \sigma \ell^*(B_2^N \cap \text{Im}(\mathbb{X})) / (\eta' \sqrt{N}) = \sigma \ell^*(B_2^{\text{Rank}(\mathbb{X})}) / (\eta' \sqrt{N}) = (\sigma / \eta') \sqrt{\text{Rank}(\mathbb{X}) / N}$$

Let ρ'_0 be the smallest ρ such that $(\rho / \sqrt{N}) \mathbb{X}B_1^d$ contains $r'_0 B_2^N \cap \text{Im}(\mathbb{X})$. An argument similar to the one used in the random design case shows that ρ'_0 is the smallest radius such that for all $\rho \geq \rho'_0$, $r_{\mathbb{X}}(\rho) = r_{\mathbb{X}}(\rho'_0) = r'_0$. Finally, we consider $K'_0 = \min\{k \in \mathbb{N} : 2^k \rho^* \geq 2\rho'_0\}$.

The fixed point function $r_{\mathbb{X}}(\cdot)$ depends on the Gaussian mean width of $\mathbb{X}B_1^d$ intersected with rB_2^N for various radii r . This quantity has been recently controlled in Proposition 2 in [4].

Proposition 4.3 (Proposition 2 in [4]) *Let $\mathbb{X} \in \mathbb{R}^{N \times d}$. Assume that the column vectors of \mathbb{X} are in B_2^N . Then, for all $r \geq 0$,*

$$\ell^*(\mathbb{X}B_1^d \cap rB_2^N) \leq \min \left(4\sqrt{\log(8ed)}, 4\sqrt{\log(8edr^2)}, r\sqrt{\text{Rank}(\mathbb{X})} \right).$$

It follows from some calculations (similar to the one used to obtain the closed form of $r_M(\cdot)$ in (1.9)) that there exists an absolute constant $C_{\mathbb{X}}$ such that for all ρ , $r_{\mathbb{X}}(\rho) \leq \bar{r}_{\mathbb{X}}(\rho)$ with

$$\bar{r}_{\mathbb{X}}^2(\rho) = C \begin{cases} \frac{\sigma^2 \text{Rank}(\mathbb{X})}{N} & \text{if } \rho \geq \rho'_0 \\ \rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)} & \text{if } \sigma^2 \log d \leq \rho^2 N \leq \rho_0'^2 N \\ \rho\sigma \sqrt{\frac{\log(ed)}{N}} & \text{if } \rho^2 N \leq \sigma^2 \log d. \end{cases} \quad (4.3)$$

We see that $\bar{r}_{\mathbb{X}}(\rho)$ in (4.3) and $r_M(\rho)$ in (1.9) are very close. The only difference comes from the rank of \mathbb{X} and when $N \geq \zeta' d$ and $\text{Rank}(\mathbb{X}) \sim d$, the two fixed points r_M and $\bar{r}_{\mathbb{X}}$ are equal up to absolute constants. Furthermore, one can check that there are two absolute constants $0 < C_1 < C_2$ and $C_{\mathbb{X}} = C_{\mathbb{X}}(\rho, d, \sigma, N)$ such that $C_1 \leq C_{\mathbb{X}} \leq C_2$ and

$$\bar{r}_{\mathbb{X}}^2(\rho) = C_{\mathbb{X}} \min\left(\frac{\sigma^2 \text{Rank}(\mathbb{X})}{N}, \rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)}, \rho\sigma \sqrt{\frac{\log(ed)}{N}}\right). \quad (4.4)$$

We will use $\rho \rightarrow \bar{r}_{\mathbb{X}}^2(\rho)$ as a regularization function (up to an absolute constant). It will appear that even though $\rho \rightarrow \bar{r}_{\mathbb{X}}^2(\rho)$ is only an upper bound on $\rho \rightarrow r_{\mathbb{X}}(\rho)$, it will be an optimal minimax choice when the design matrix satisfies RIP.

4.0.2 Main result: a minimax regularization function

In this section, we prove that the regularization function $\Psi(\rho) = c'_0 \bar{r}_{\mathbb{X}}^2(\rho)$ is a minimax regularization function for some well-chosen absolute constant c'_0 when the design matrix \mathbb{X} satisfies RIP. To that end, we first need to know the minimax rate over ℓ_1^d -balls in the fixed design setup. Such a result was obtained in [32]. Let us now recall this result in our context (see (5.25) in Section 5.2.2 in [32]).

Proposition 4.4 ([32]) *Let $\mathbb{X} \in \mathbb{R}^{N \times d}$ be a matrix satisfying RIP(2s). For all $\rho \geq 0$, the minimax rate of convergence over ρB_1^d in the Gaussian linear model with fixed design \mathbb{X} in expectation is given by*

$$\min\left(\frac{\sigma^2 \text{Rank}(\mathbb{X})}{N}, \rho\sigma \sqrt{\frac{1}{N} \log\left(\frac{e\sigma d}{\rho\sqrt{N}}\right)}, \rho^2\right). \quad (4.5)$$

The latter result holds in expectation whereas we are interested in deviation results. Even though the minimax rate of convergence in deviation over ρB_1^d in the Gaussian linear model under RIP has not been established, we believe that this rate of convergence in deviation is identical to the one given in Proposition 4.4. Note that a proof of this fact follows from Section 3 in [22] for the minimax lower bound in deviation and from the quadratic / multiplier decomposition of the excess loss together with Lemma 4.6 below to show that the ERM over ρB_1^N achieves the minimax bound in deviation. We do not provide the proof here but we will use (4.5) as a benchmark result to construct a minimax regularization function.

Theorem 4.5 Let $\mathbb{X} \in \mathbb{R}^{N \times d}$ be such that the column vectors of \mathbb{X} are in B_2^N . Consider the following regularization function:

$$\rho \geq 0 \rightarrow \Psi(\rho) = c'_0 \overline{r_{\mathbb{X}}}^2(\rho)$$

where $\overline{r_{\mathbb{X}}}(\rho)$ is defined in (4.4) and $c'_0 \geq 2$ is an absolute constant. Then there exist absolute constants κ'_1, κ'_2 and κ'_3 such that for any $t^* \in \mathbb{R}^d$ the RERM \hat{t} constructed from the data $Y = \mathbb{X}t^* + \xi$:

$$\hat{t} \in \operatorname{argmin}_{t \in \mathbb{R}^d} \left(\|Y - \mathbb{X}t\|_{L_N^2}^2 + \Psi(\|t\|_1) \right)$$

satisfies with probability greater than $1 - \kappa'_1 \exp(-\kappa'_2 N \overline{r_{\mathbb{X}}}(\|t^*\|_1)^2 / \sigma^2)$,

$$\|\mathbb{X}(\hat{t} - t^*)\|_{L_N^2}^2 \leq \kappa'_3 \min \left(\frac{\sigma^2 \operatorname{Rank}(\mathbb{X})}{N}, \frac{\|t^*\|_1 \sigma}{\sqrt{N}} \sqrt{\log \left(\frac{ed\sigma}{\|t^*\|_1 \sqrt{N}} \right)}, \|t^*\|_1 \sigma \sqrt{\frac{\log(ed)}{N}} \right).$$

Note that the probability estimate $1 - \kappa'_1 \exp(-\kappa'_2 N \overline{r_{\mathbb{X}}}(\|t^*\|_1)^2 / \sigma^2) \geq 3/4$ only when $\|t^*\|_1 \geq \Delta'_0 \sigma \sqrt{\log(ed)/N}$ for some absolute constant Δ'_0 large enough. Therefore, if (4.5) is indeed the minimax rate of convergence over ρB_1^d for the deviation $1 - \delta_N = 3/4$ under RIP then Theorem 4.5 proves that $\rho \geq 0 \rightarrow \Psi(\rho) = c'_0 \overline{r_{\mathbb{X}}}^2(\rho)$ is a minimax regularization function for the ℓ_1^d -norm over $\mathbb{R}^d \setminus \Delta'_0 \sigma \sqrt{\log(ed)/N} B_1^d$ for the constant confidence regime.

4.1 Proof of Theorem 4.5

The proof is split into a probabilist part used to identify a high probability event Ω'_0 on which the multiplier process is well controlled on the entire space \mathbb{R}^d and a deterministic part where it is proved that, on the event Ω'_0 , $P_N \mathcal{L}_t^\Psi > 0$ if $\|\mathbb{X}(t - t^*)\|_{L_N^2}^2 \gtrsim \overline{r_{\mathbb{X}}}(\|t^*\|_1)^2$.

4.1.1 Probabilistic control of the multiplier process

The following lemma shows how the fixed point $r_{\mathbb{X}}$ allows to control the multiplier process.

Lemma 4.6 Let $\rho > 0$ and take $\eta' = 1/8$. Then, for all $r \geq r_{\mathbb{X}}(\rho)$, with probability greater than $1 - \exp(-Nr^2/(128\sigma^2))$, for all $t \in t^* + \rho B_1^d$,

$$|P_N \mathcal{M}_{t-t^*}| \leq \frac{1}{2} \max \left(r^2, \|\mathbb{X}(t - t^*)\|_{L_N^2}^2 \right).$$

Proof. Let $r \geq r_{\mathbb{X}}(\rho)$. We denote by $B_{\mathbb{X}}^2$ the unit ball associated with the pseudo-metric $\|\mathbb{X} \cdot\|_{L_N^2}$ and $rB_{\mathbb{X}}^2 = \{t \in \mathbb{R}^d : \|\mathbb{X}t\|_{L_N^2} \leq r\}$ its unit ball of radius r . We have

$$\sup_{t \in t^* + \rho B_1^d \cap rB_{\mathbb{X}}^2} |P_N \mathcal{M}_{t-t^*}| = 2 \sup_{t \in t^* + \rho B_1^d \cap rB_{\mathbb{X}}^2} \left| \langle \mathbb{X}(t - t^*), \xi \rangle_{L_N^2} \right| = \frac{2}{\sqrt{N}} \sup_{v \in V} \langle v, \xi \rangle$$

where $V = [(\rho/\sqrt{N})\mathbb{X}B_1^d] \cap rB_2^N$. Then, it follows from Borell's concentration inequality (cf. [25]) that for all $x > 0$, with probability at least $1 - \exp(-x^2/2)$,

$$\sup_{v \in V} \langle v, \xi \rangle \leq \mathbb{E} \sup_{v \in V} \langle v, \xi \rangle + x\sigma(V)$$

where $\sigma(V) := \sup_{v \in V} \sqrt{\mathbb{E}\langle v, \xi \rangle^2} = \sup_{v \in V} \sigma \|v\|_2 \leq \sigma r$. Moreover, given that $\xi \sim \mathcal{N}(0, \sigma^2 I_N)$, we have

$$\mathbb{E} \sup_{v \in V} \langle v, \xi \rangle = \sigma \ell^*(V) = \sigma \ell^* \left(\frac{\rho}{\sqrt{N}} \mathbb{X} B_1^d \cap r B_2^N \right) \leq \eta' r^2 \sqrt{N}$$

where the last inequality follows from the definition of $r_{\mathbb{X}}(\rho)$ and because $r \geq r_{\mathbb{X}}(\rho)$. Gathering all the pieces together, it follows for $x = \eta' r \sqrt{N} / \sigma$, that, with probability at least $1 - \exp(-r^2 N / (128 \sigma^2))$,

$$\sup_{t \in t^* + \rho B_1^d \cap r B_{\mathbb{X}}^2} |P_N \mathcal{M}_{t-t^*}| \leq 4\eta' r^2 = \frac{1}{2} r^2. \quad (4.6)$$

Now, the proof follows from an homogeneity argument. Indeed, let us assume that (4.6) holds. Let $t \in t^* + \rho B_1^d$ be such that $\|\mathbb{X}(t-t^*)\|_{L_N^2} > r$ and define $t' := t^* + \alpha_t(t-t^*)$ where $\alpha_t = r / \|\mathbb{X}(t-t^*)\|_{L_N^2}$. Note that $\alpha_t < 1$ and $t' \in t^* + \rho B_1^d \cap r B_{\mathbb{X}}^2$. Hence, it follows from (4.6) that $|P_N \mathcal{M}_{t'-t^*}| \leq r^2/2$ and so $|P_N \mathcal{M}_{t-t^*}| = |P_N \mathcal{M}_{t'-t^*}| / \alpha_t \leq r^2 / (2\alpha_t) \leq \|\mathbb{X}(\hat{t}-t^*)\|_{L_N^2}^2 / 2$. \blacksquare

4.1.2 Deterministic part of the proof

We start with two lemmas on the growth behavior of $\overline{r_{\mathbb{X}}}(\cdot)$. Their proofs are almost identical to the one of Lemmas 2.3 and 2.4 and are therefore omitted.

Lemma 4.7 *Let $\phi' = 4$. If $\phi' \rho \leq \rho'_0 \min(1, \eta')$, then for any $\rho' \geq \phi' \rho$, $\overline{r_{\mathbb{X}}}^2(\rho') > 2\overline{r_{\mathbb{X}}}^2(\rho)$.*

Lemma 4.8 *Let $\nu > 0$. If $\nu \geq 1$ then $\overline{r_{\mathbb{X}}}(\nu\rho) \leq \sqrt{\nu} \overline{r_{\mathbb{X}}}(\rho)$. If $\nu \leq 1$ then $\overline{r_{\mathbb{X}}}(\nu\rho) \geq \sqrt{\nu} \overline{r_{\mathbb{X}}}(\rho)$.*

To prove Theorem 4.5, we use the same argument as in the proof of Theorem 1.4 for the random design. Let $\rho^* = 10\|t^*\|_1/\eta'$ and split \mathbb{R}^d into three zones:

- the “central zone” $t^* + \rho^* B_1^d$,
- the intermediate “peeling zone”: $\{t \in \mathbb{R}^d : \rho^* < \|t-t^*\|_1 \leq 2^{K'_0} \rho^*\}$ – to be considered only when $K'_0 \geq 1$. This part of \mathbb{R}^d is itself partitioned into K'_0 shelves: for $k = 1, \dots, K'_0$, $\{t \in \mathbb{R}^d : 2^{k-1} \rho^* < \|t-t^*\|_1 \leq 2^k \rho^*\}$,
- the “exterior zone”: $\{t \in \mathbb{R}^d : \|t-t^*\|_1 > 2^{K'_0} \rho^*\}$ on which $r_{\mathbb{X}}$ is constant equal to r'_0 .

For all $k = 0, \dots, K'_0$, we denote by A'_k the event on which for all $t \in t^* + 2^k \rho^* B_1^d$,

$$|P_N \mathcal{M}_{t-t^*}| \leq \frac{1}{2} \max \left(\overline{r_{\mathbb{X}}}(2^k \rho^*)^2, \|\mathbb{X}(t-t^*)\|_{L_N^2}^2 \right).$$

We consider the event $\Omega'_0 = A'_0 \cap \dots \cap A'_{K'_0}$. It follows from Lemma 4.6 and an argument similar to the one in Lemma 3.2 that for some absolute constants κ'_1, κ'_2 and κ'_4 ,

$$\mathbb{P}[\Omega'_0] \geq 1 - \kappa'_1 \exp \left(-\kappa'_2 N \overline{r_{\mathbb{X}}}(\|t^*\|_1)^2 / \sigma^2 \right)$$

as long as $\|t^*\|_1 \geq \kappa'_4 \sigma / \sqrt{N}$ (which is the case when $\|t^*\|_1 \geq \Delta'_0 \sigma \sqrt{\log(ed)/N}$ for $\Delta'_0 \geq \kappa'_4$).

Let us now assume for the remaining of the proof that Ω'_0 holds. Note that unlike in the random design case, there is no event such as Ω^* in Ω'_0 on which the quadratic process is controlled, because, in the deterministic design case this process is deterministic.

Our strategy is to show that \hat{t} belongs to the ‘‘central zone’’. To that end it is enough to prove that $P_N \mathcal{L}_t^\Psi > 0$ for every $t \in \mathbb{R}^d$ such that $\|t - t^*\|_1 > \rho^*$ because by definition $P_N \mathcal{L}_{\hat{t}}^\Psi \leq 0$.

Let t be in the intermediate peeling zone (which can happen only if $\rho^* \leq 2\rho'_0$), say in the k -th shell for some $k \in \{0, \dots, K'_0\}$: $2^{k-1}\rho^* < \|t - t^*\|_1 \leq 2^k \rho^*$. In particular $\|t\|_1 > \|t^*\|_1$ and $\mathcal{R}_{t,t^*} \geq 0$. Therefore, if $\|\mathbb{X}(t - t^*)\|_{L_N^2} \geq \overline{r_{\mathbb{X}}}(2^k \rho^*)$ then by Lemma 4.6, $|P_N \mathcal{M}_{t-t^*}| \leq \|\mathbb{X}(t - t^*)\|_{L_N^2}^2 = P_N \mathcal{Q}_{t-t^*}$ and so $P_N \mathcal{L}_t^\Psi > 0$. Now, if $\|\mathbb{X}(t - t^*)\|_{L_N^2} \leq \overline{r_{\mathbb{X}}}(2^k \rho^*)$ then by Lemmas 4.6 and 4.8, for $c'_0 \geq 2$,

$$|P_N \mathcal{M}_{t-t^*}| \leq \frac{1}{2} \overline{r_{\mathbb{X}}}(2^k \rho^*)^2 < \frac{c'_0}{2} \overline{r_{\mathbb{X}}}^2(\|t - t^*\|_1)$$

and since $\|t\|_1 \geq \|t - t^*\|_1 - \|t^*\|_1 \geq 4\|t^*\|_1$, and $4\|t^*\|_1 \leq \eta' \rho^* / 2 \leq \eta' \rho'_0$, by Lemma 4.7, one has $c'_0 \overline{r_{\mathbb{X}}}^2(\|t\|_1) \geq 2c'_0 \overline{r_{\mathbb{X}}}^2(\|t^*\|_1)$. As a consequence, $\mathcal{R}_{t,t^*} \geq c'_0 \overline{r_{\mathbb{X}}}^2(\|t\|_1) / 2 > |P_N \mathcal{M}_{t-t^*}|$ and so $P_N \mathcal{L}_t^\Psi > 0$.

Let us now tackle the exterior zone in both cases $\rho^* \leq 2\rho'_0$ and $\rho^* > 2\rho'_0$. Let $t \in \mathbb{R}^d$ be such that $\|t - t^*\|_1 > 2^{K'_0} \rho^*$. We have $\mathcal{R}_{t,t^*} \geq 0$ because $\Psi(\|t\|_1) = c'_0 \overline{r_{\mathbb{X}}}^2(\|t\|_1) = c'_0 r_0'^2 \geq c'_0 \overline{r_{\mathbb{X}}}^2(\|t^*\|_1) = \Psi(\|t^*\|_1)$. Let $t' = t^* + \alpha_t(t - t^*)$ for some $0 < \alpha_t < 1$ be such that $\|t' - t^*\|_1 = 2^{K'_0} \rho^*$. By definition of K'_0 and ρ'_0 , we have $\|\mathbb{X}(t' - t^*)\|_{L_N^2} \geq r'_0$. Therefore, since $A_{K'_0} \subset \Omega'_0$, we have $|P_N \mathcal{M}_{t'-t^*}| \leq (1/2) \|\mathbb{X}(t' - t^*)\|_{L_N^2}^2$ which implies that $P_N \mathcal{Q}_{t'-t^*} + P_N \mathcal{M}_{t'-t^*} > 0$ and therefore by an homogeneity argument that $P_N \mathcal{Q}_{t-t^*} + P_N \mathcal{M}_{t-t^*} > 0$. Finally, given that $\mathcal{R}_{\hat{t},t^*} \geq 0$ we conclude that $P_N \mathcal{L}_{\hat{t}}^\Psi > 0$.

This proves that \hat{t} lies in the central zone in both cases $\rho^* \leq 2\rho'_0$ and $\rho^* > 2\rho'_0$. But, now given that $A'_0 \subset \Omega'_0$, we have

$$|P_N \mathcal{M}_{\hat{t}-t^*}| \leq \frac{1}{2} \max\left(\overline{r_{\mathbb{X}}}(\rho^*)^2, \|\mathbb{X}(\hat{t} - t^*)\|_{L_N^2}^2\right).$$

If $\|\mathbb{X}(\hat{t} - t^*)\|_{L_N^2}^2 \leq \overline{r_{\mathbb{X}}}(\rho^*)^2$ the proof is over and otherwise $|P_N \mathcal{M}_{\hat{t}-t^*}| \leq (1/2) \|\mathbb{X}(\hat{t} - t^*)\|_{L_N^2}^2$ which implies that $\|\hat{t} - t^*\|_{L_N^2}^2 \leq 2\Psi(\|t^*\|_1) \leq 2c'_0 \overline{r_{\mathbb{X}}}(\rho^*)^2$ because

$$0 \geq P_N \mathcal{L}_{\hat{t}}^\Psi \geq \frac{1}{2} \|\hat{t} - t^*\|_{L_N^2}^2 + \Psi(\|\hat{t}\|_1) - \Psi(\|t^*\|_1).$$

This proves, on the event Ω'_0 , that $\|\hat{t} - t^*\|_{L_N^2}^2 \leq 2c'_0 \overline{r_{\mathbb{X}}}(\rho^*)^2$. ■

References

- [1] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [2] Andrew Barron, Lucien Birgé, and Pascal Massart. Risk bounds for model selection via penalization. *Probab. Theory Related Fields*, 113(3):301–413, 1999.

- [3] Peter L. Bartlett and Shahar Mendelson. Empirical minimization. Probab. Theory Related Fields, 135(3):311–334, 2006.
- [4] Pierre Bellec. Localized gaussian width of m -convex hulls with applications to lasso and convex aggregation. Technical report, Rutgers University, 2017.
- [5] Pierre Bellec, Guillaume Lecué, and Alexandre Tsybakov. Slope meets lasso: improved oracle bounds and optimality. Technical report, CNRS, Université’ Paris-Saclay, CREST, 2016.
- [6] Peter J. Bickel, Ya’acov Ritov, and Alexandre B. Tsybakov. Simultaneous analysis of lasso and Dantzig selector. Ann. Statist., 37(4):1705–1732, 2009.
- [7] Lucien Birgé and Pascal Massart. Gaussian model selection. Journal of the European Mathematical Society, 3(3):203–268, 2001.
- [8] Peter Bühlmann and Sara van de Geer. Statistics for high-dimensional data. Springer Series in Statistics. Springer, Heidelberg, 2011. Methods, theory and applications.
- [9] Emmanuel J Candès and Yaniv Plan. Near-ideal model selection by l_1 minimization. The Annals of Statistics, 37(5A):2145–2177, 2009.
- [10] Emmanuel J. Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inform. Theory, 52(2):489–509, 2006.
- [11] Emmanuel J. Candès and Terence Tao. Decoding by linear programming. IEEE Trans. Inform. Theory, 51(12):4203–4215, 2005.
- [12] Emmanuel J. Candès and Terence Tao. The power of convex relaxation: near-optimal matrix completion. IEEE Trans. Inform. Theory, 56(5):2053–2080, 2010.
- [13] Djilil Chafaï, Olivier Guédon, Guillaume Lecué, and Alain Pajor. Interactions between compressed sensing random matrices and high dimensional geometry, volume 37 of Panoramas et Synthèses [Panoramas and Syntheses]. Société Mathématique de France, Paris, 2012.
- [14] Sourav Chatterjee. A new perspective on least squares under convex constraint. The Annals of Statistics, 42(6):2340–2381, 2014.
- [15] Mark A Davenport and Michael B Wakin. Analysis of orthogonal matching pursuit using the restricted isometry property. IEEE Transactions on Information Theory, 56(9):4395–4401, 2010.
- [16] Rahul Garg and Rohit Khandekar. Gradient descent with sparsification: an iterative algorithm for sparse recovery with restricted isometry property. In Proceedings of the 26th Annual International Conference on Machine Learning, pages 337–344. ACM, 2009.
- [17] A. Yu. Garnaev and E. D. Gluskin. The widths of a Euclidean ball. Dokl. Akad. Nauk SSSR, 277(5):1048–1052, 1984.
- [18] Christophe Giraud. Introduction to high-dimensional statistics, volume 138. CRC Press, 2014.
- [19] Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. Ann. Statist., 34(6):2593–2656, 2006.
- [20] 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].
- [21] Guillaume Lecué. Interplay between concentration, complexity and geometry in learning theory with applications to high dimensional data analysis. Habilitation à diriger des recherches. 2011.
- [22] Guillaume Lecué and Shahar Mendelson. Learning subgaussian classes: Upper and minimax bounds. Technical report, CNRS, Ecole polytechnique and Technion, 2013.
- [23] Guillaume Lecué and Shahar Mendelson. Regularization and the small-ball method i: sparse recovery. arXiv preprint arXiv:1601.05584, 2016.
- [24] Guillaume Lecué and Shahar Mendelson. Sparse recovery under weak moment assumptions. J. Eur. Math. Soc. (JEMS), 19(3):881–904, 2017.

- [25] Michel Ledoux. The concentration of measure phenomenon, volume 89 of Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2001.
- [26] Stéphane Mallat. A wavelet tour of signal processing. Elsevier/Academic Press, Amsterdam, third edition, 2009. The sparse way, With contributions from Gabriel Peyré.
- [27] Pascal Massart. Concentration inequalities and model selection, volume 1896 of Lecture Notes in Mathematics. Springer, Berlin, 2007. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003, With a foreword by Jean Picard.
- [28] Shahar Mendelson. Upper bounds on product and multiplier empirical processes. Technical report, Technion, I.I.T., 2013. To appear in Stochastic Processes and their Applications.
- [29] Balas Kausik Natarajan. Sparse approximate solutions to linear systems. SIAM journal on computing, 24(2):227–234, 1995.
- [30] Arkadii Nemirovski. Lectures on probability theory and statistics, volume 1738 of Lecture Notes in Mathematics. Springer-Verlag, Berlin, 2000. Lectures from the 28th Summer School on Probability Theory held in Saint-Flour, August 17–September 3, 1998, Edited by Pierre Bernard.
- [31] Allan Pinkus. n -widths in approximation theory, volume 7 of Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1985.
- [32] Philippe Rigollet and Alexandre Tsybakov. Exponential screening and optimal rates of sparse estimation. Ann. Statist., 39(2):731–771, 2011.
- [33] Angelika Rohde and Alexandre B. Tsybakov. Estimation of high-dimensional low-rank matrices. Ann. Statist., 39(2):887–930, 2011.
- [34] A. Saumard. Optimal upper and lower bounds for the true and empirical excess risks in heteroscedastic least-squares regression. Electron. J. Statist., 6(1-2):579–655, 2012.
- [35] Adrien Saumard. A concentration inequality for the excess risk in least-squares regression with random design and heteroscedastic noise. arXiv preprint arXiv:1702.05063, 2017.
- [36] Charles Stein. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In Proceedings of the Third Berkeley symposium on mathematical statistics and probability, volume 1, pages 197–206, 1956.
- [37] Michel Talagrand. Upper and lower bounds for stochastic processes, volume 60 of Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]. Springer, Heidelberg, 2014. Modern methods and classical problems.
- [38] Robert Tibshirani. Regression shrinkage and selection via the lasso. J. Roy. Statist. Soc. Ser. B, 58(1):267–288, 1996.
- [39] Alexandre B. Tsybakov. Introduction to nonparametric estimation. Springer Series in Statistics. Springer, New York, 2009. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.
- [40] Sara A. van de Geer. Applications of empirical process theory, volume 6 of Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2000.
- [41] Sara A. van de Geer. The deterministic lasso. Technical report, ETH Zürich, 2007. <http://www.stat.math.ethz.ch/geer/lasso.pdf>.
- [42] Sara A. van de Geer and Peter Bühlmann. On the conditions used to prove oracle results for the Lasso. Electron. J. Stat., 3:1360–1392, 2009.
- [43] Sara A Van De Geer, Peter Bühlmann, et al. On the conditions used to prove oracle results for the lasso. Electronic Journal of Statistics, 3:1360–1392, 2009.
- [44] Aad W. van der Vaart and Jon A. Wellner. Weak convergence and empirical processes. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics.
- [45] 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.

- [46] Nicolas Verzelen. Minimax risks for sparse regressions: ultra-high dimensional phenomenons. Electron. J. Stat., 6:38–90, 2012.
- [47] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. J. R. Stat. Soc. Ser. B Stat. Methodol., 67(2):301–320, 2005.