

PAC-Bayesian bounds for sparse regression estimation with exponential weights

Pierre Alquier*

CREST - ENSAE
3, avenue Pierre Larousse
92240 Malakoff France
and

LPMA - Université Paris 7
175, rue du Chevaleret
75205 Paris Cedex 13 France
e-mail: alquier@ensae.fr
url: <http://alquier.ensae.net/>

and

Karim Lounici

School of Mathematics
Georgia Institute of Technology
Atlanta, GA 30332-0160 USA
e-mail: klounici@math.gatech.edu
url: <http://people.math.gatech.edu/~klounici6/>

Abstract: We consider the sparse regression model where the number of parameters p is larger than the sample size n . The difficulty when considering high-dimensional problems is to propose estimators achieving a good compromise between statistical and computational performances. The Lasso is solution of a convex minimization problem, hence computable for large value of p . However stringent conditions on the design are required to establish fast rates of convergence for this estimator. Dalalyan and Tsybakov [17–19] proposed an exponential weights procedure achieving a good compromise between the statistical and computational aspects. This estimator can be computed for reasonably large p and satisfies a sparsity oracle inequality in expectation for the empirical excess risk only under mild assumptions on the design. In this paper, we propose an exponential weights estimator similar to that of [17] but with improved statistical performances. Our main result is a sparsity oracle inequality in probability for the true excess risk.

AMS 2000 subject classifications: Primary 62J07; secondary 62J05, 62G08, 62F15, 62B10, 68T05.

Keywords and phrases: Sparsity oracle inequality, high-dimensional regression, exponential weights, PAC-Bayesian inequalities.

Received September 2010.

*Research partially supported by the French “Agence Nationale pour la Recherche” under grant ANR-09-BLAN-0128 “PARCIMONIE”.

Contents

1	Introduction	128
2	Sparsity oracle inequality in expectation	131
3	Sparsity oracle inequality in probability	132
4	Proofs	135
	4.1 Proofs of Section 2	135
	4.2 Proof of Theorem 3.1	139
	References	143

1. Introduction

We observe n independent pairs $(X_1, Y_1), \dots, (X_n, Y_n) \in \mathcal{X} \times \mathbb{R}$ (where \mathcal{X} is any measurable set) such that

$$Y_i = f(X_i) + W_i, \quad 1 \leq i \leq n, \quad (1.1)$$

where $f : \mathcal{X} \rightarrow \mathbb{R}$ is the unknown regression function and the noise variables W_1, \dots, W_n are independent of the design (X_1, \dots, X_n) , satisfy also $\mathbb{E}W_i = 0$ and $\mathbb{E}W_i^2 \leq \sigma^2$ for any $1 \leq i \leq n$ and for some known $\sigma^2 > 0$. The distribution of the sample is denoted by \mathbb{P} , the corresponding expectation is denoted by \mathbb{E} . For any function $g : \mathcal{X} \rightarrow \mathbb{R}$ define $\|g\|_n = (\sum_{i=1}^n g(X_i)^2/n)^{1/2}$ and $\|g\| = (\mathbb{E}\|g\|_n^2)^{1/2}$. Let $\mathcal{F} = \{\phi_1, \dots, \phi_p\}$ be a set—called dictionary—of functions $\phi_j : \mathcal{X} \rightarrow \mathbb{R}$ such that $\|\phi_j\| = 1$ for any j (this assumption can be relaxed). For any $\theta \in \mathbb{R}^p$ define $f_\theta = \sum_{j=1}^p \theta_j \phi_j$, the empirical risk

$$r(\theta) = \frac{1}{n} \sum_{i=1}^n (Y_i - f_\theta(X_i))^2,$$

and the integrated risk

$$R(\theta) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (Y'_i - f_\theta(X'_i))^2 \right],$$

where $\{(X'_1, Y'_1), \dots, (X'_n, Y'_n)\}$ is an independent replication of $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$. Let us choose $\hat{\theta} \in \arg \min_{\theta \in \mathbb{R}^p} R(\theta)$. Note that the minimum may not be unique, however we do not need to treat the identifiability question since we consider in this paper the prediction problem, i.e., find an estimator $\hat{\theta}_n$ such that $R(\hat{\theta}_n)$ is close to $\min_{\theta \in \mathbb{R}^p} R(\theta)$ up to a positive remainder term as small as possible.

It is a known fact that the least-square estimator $\hat{\theta}_n^{LSE} \in \arg \min_{\theta \in \mathbb{R}^p} r(\theta)$ performs poorly in high-dimension $p > n$. Indeed, consider for instance the deterministic design case with i.i.d. noise variables $N(0, \sigma^2)$ and a full-rank design matrix, then $\hat{\theta}_n^{LSE}$ satisfies

$$\mathbb{E} \left[\|f_{\hat{\theta}_n^{LSE}} - f\|_n^2 \right] - \|f_{\hat{\theta}} - f\|_n^2 = \sigma^2.$$

In the same context, assume now there exists a vector $\bar{\theta} \in \arg \min_{\theta \in \mathbb{R}^p} R(\theta)$ with a number of nonzero coordinates $p_0 \leq n$. If the indices of these coordinates are known, then we can construct an estimator $\hat{\theta}_n^0$ such that

$$\mathbb{E} \left[\|f_{\hat{\theta}_n^0} - f\|_n^2 \right] - \|f_{\bar{\theta}} - f\|_n^2 = \sigma^2 \frac{p_0}{n}.$$

The estimator $\hat{\theta}_n^0$ is called oracle estimator since the set of indices of the nonzero coordinates of $\bar{\theta}$ is unknown in practice. The issue is now to build an estimator, when the set of nonzero coordinates of $\bar{\theta}$ is unknown, with statistical performances close to that of the oracle estimator $\hat{\theta}_n^0$.

A possible approach is to consider solutions of penalized empirical risk minimization problems:

$$\hat{\theta}_{\text{pen}} \in \arg \min_{\theta \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n (Y_i - f_{\theta}(X_i))^2 + \text{pen}(\theta) \right\},$$

where the penalization $\text{pen}(\theta)$ is proportional to the number of nonzero components of θ such as for instance AIC, C_p and BIC criteria [1, 31, 37]. Bunea, Tsybakov and Wegkamp [9] established for the BIC estimator $\hat{\theta}_n^{\text{BIC}}$ the following non-asymptotic sparsity oracle inequality. For any $\epsilon > 0$ there exists a constant $C(\epsilon) > 0$ such that for any $p \geq 2, n \geq 1$ we have

$$\mathbb{E} \left[\|f_{\hat{\theta}_n^{\text{BIC}}} - f\|_n^2 \right] \leq (1 + \epsilon) \|f_{\bar{\theta}} - f\|_n^2 + C(\epsilon) \sigma^2 \frac{p_0}{n} \log \left(\frac{ep}{p_0 \vee 1} \right).$$

Despite good statistical properties, these estimators can only be computed in practice for p of the order at most a few tens since they are solutions of non-convex combinatorial optimization problems.

Considering convex penalty function leads to computationally feasible optimization problems. A popular example of convex optimization problem is the Lasso estimator (cf. Frank and Friedman [22], Tibshirani [39], and the parallel work of Chen, Donoho and Saunders [15] on basis pursuit) with the penalty term $\text{pen}(\theta) = \lambda |\theta|_1$, where $\lambda > 0$ is some regularization parameter and, for any integer $d \geq 2$, real $q > 0$ and vector $z \in \mathbb{R}^d$ we define $|z|_q = (\sum_{j=1}^d |z_j^q|)^{1/q}$ and $|z|_{\infty} = \max_{1 \leq j \leq d} |z_j|$. Several algorithms allow to compute the Lasso for very large p , one of the most popular is known as LARS, introduced by Efron, Hastie, Johnstone and Tibshirani [21]. However, the Lasso estimator requires strong assumptions on the matrix $A = (\phi_j(X_i))_{1 \leq i \leq n, 1 \leq j \leq p}$ to establish fast rates of convergence results. Bunea, Tsybakov and Wegkamp [8] and Lounici [30] assume a mutual coherence condition on the dictionary. Bickel, Ritov and Tsybakov [7] and Koltchinskii [25] established sparsity oracle inequalities for the Lasso under a restricted eigenvalue condition. Candès and Tao [11] and Koltchinskii [26] studied the Dantzig Selector which is related to the Lasso estimator and suffers from the same restrictions. See, e.g., Bickel, Ritov and Tsybakov [7] for more details. Several alternative penalties were recently considered. Zou [45] proposed the adaptive Lasso which is the solution of a penalized empirical risk minimization problem with the penalty $\text{pen}(\theta) = \lambda \sum_{j=1}^p \frac{1}{|\hat{w}_j|} |\theta_j|$ where \hat{w} is

an initial estimator of θ . Zou and Hastie [46] proposed the elastic net with the penalty $\text{pen}(\theta) = \lambda_1|\theta|_1 + \lambda_2|\theta|_2^2$, $\lambda_1, \lambda_2 > 0$. Meinshausen and Bühlmann [35] and Bach [6] considered bootstrapped Lasso. See also Ghosh [23] or Cai, Xu and Zhang [10] for more alternatives to the Lasso. All these methods were motivated by their superior performances over the Lasso either from the theoretical or the practical point of view. However, strong assumptions on the design are still required to establish the statistical properties of these methods (when such results exist). A recent paper by van de Geer and Bühlmann [42] provides a complete survey and comparison of all these assumptions.

Simultaneously, the PAC-Bayesian approach for regression estimation was developed by Audibert [4, 5] and Alquier [2, 3], based on previous works in the classification context by Catoni [12–14], Mc Allester [34], Shawe-Taylor and Williamson [38], see also Zhang [44] in the context of density estimation. This framework is very well adapted for studying the excess risk $R(\cdot) - R(\bar{\theta})$ in the regression context since it requires very weak conditions on the dictionary. However, the methods of these papers are not designed to cover the high-dimensional setting under the sparsity assumption. Dalalyan and Tsybakov [16–19] propose an exponential weights procedure related to the PAC-Bayesian approach with good statistical and computational performances. However they consider deterministic design, establishing their statistical result only for the empirical excess risk instead of the true excess risk $R(\cdot) - R(\bar{\theta})$.

In this paper, we propose to study two exponential weights estimation procedures. The first one is an exponential weights combination of the least squares estimators in all the possible sub-models. This estimator was initially proposed by Leung and Barron [28] in the deterministic design setting. Note that in the literature on aggregation and exponential weights, the elements of the dictionary are often arbitrary preliminary estimators computed from a frozen fraction of the initial sample so that these estimators are considered as deterministic functions, the aggregate is then computed using this dictionary and the remaining data. This scheme is referred to as 'data splitting'. See for instance Dalalyan and Tsybakov [18] and Yang [43]. Leung and Barron [28] proved that data splitting is not necessary in order to aggregate least squares estimators and raised the question of computation of this estimator in high dimension. In this paper we explicit the oracle inequality satisfied by this estimator in the high-dimensional case. For the second procedure, the design may be either random or deterministic. We adapt to the regression framework PAC-Bayesian techniques developed by Catoni [14] in the classification framework to build an estimator satisfying a sparsity oracle inequality for the true excess risk. Even though we do not study the computational aspect in this paper, it should be noted that efficient Monte Carlo algorithms are available to compute these exponential weights estimators for reasonably large dimension p ($p \simeq 5000$), see in particular the monograph of Marin and Robert [32] for an introduction to MCMC methods. Note also that in a work parallel to ours, Rigollet and Tsybakov [36] consider also an exponential weights procedure with discrete priors and suggest a version of the Metropolis-Hastings algorithm to compute it.

The paper is organized as follows. In Section 2 we define an exponential weights procedure and derive a sparsity oracle inequality in expectation when the design is deterministic. In Section 3, the design can be either deterministic or random. We propose a modification of the first exponential weights procedure for which we can establish a sparsity oracle inequality in probability for the true excess risk. Finally Section 4 contains all the proofs of our results.

2. Sparsity oracle inequality in expectation

Throughout this section, we assume that the design is deterministic and the noise variables W_1, \dots, W_n are i.i.d. gaussian $N(0, \sigma^2)$.

For any $J \subset \{1, \dots, p\}$ and $K > 0$ define

$$\Theta(J) = \left\{ \theta \in \mathbb{R}^p : \quad \forall j \notin J, \quad \theta_j = 0 \right\}, \quad (2.1)$$

and

$$\Theta_K(J) = \left\{ \theta \in \mathbb{R}^p : \quad |\theta|_1 \leq K \quad \text{and} \quad \forall j \notin J, \quad \theta_j = 0 \right\}. \quad (2.2)$$

For the sake of simplicity we will write $\Theta_K = \Theta_K(\{1, \dots, p\})$.

For any subset $J \subset \{1, \dots, p\}$ define

$$\hat{\theta}_J \in \arg \min_{\theta \in \Theta(J)} r(\theta),$$

where $r(\theta) = \frac{1}{n} \sum_{i=1}^n (Y_i - f_\theta(X_i))^2 = \|Y - f_\theta\|_n^2$ with $Y = (Y_1, \dots, Y_n)^T$. Denote by $\mathcal{P}_n(\{1, \dots, p\})$ the set of all subsets of $\{1, \dots, p\}$ containing at most n elements. The aggregate \hat{f}_n is defined as follows

$$\hat{f}_n = f_{\hat{\theta}_n}, \quad \hat{\theta}_n = \hat{\theta}_n(\lambda, \pi) \triangleq \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J e^{-\lambda(r(\hat{\theta}_J) + \frac{2\sigma^2|J|}{n})} \hat{\theta}_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J e^{-\lambda(r(\hat{\theta}_J) + \frac{2\sigma^2|J|}{n})}} \quad (2.3)$$

where $\lambda > 0$ is the temperature parameter, π is the prior probability distribution on $\mathcal{P}(\{1, \dots, p\})$, the set of all subsets of $\{1, \dots, p\}$, that is, for any $J \in \{1, \dots, p\}$, $\pi_J \geq 0$ and $\sum_{J \in \mathcal{P}(\{1, \dots, p\})} \pi_J = 1$.

The next result is a reformulation in our context of Theorem 8 of [28].

Proposition 2.1. *Assume that the noise variables W_1, \dots, W_n are i.i.d. $N(0, \sigma^2)$. Then the aggregate $\hat{\theta}_n$ defined by (2.3) with $0 < \lambda \leq \frac{n}{4\sigma^2}$ satisfies*

$$\mathbb{E} \left[r(\hat{\theta}_n) \right] \leq \min_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left\{ \mathbb{E}[r(\hat{\theta}_J)] + \frac{1}{\lambda} \log \left(\frac{1}{\pi_J} \right) \right\}. \quad (2.4)$$

Proposition 2.1 holds true for any prior π . We suggest using the following prior. Fix $\alpha \in (0, 1)$ and define

$$\pi_J = \frac{\alpha^{|J|}}{\sum_{j=0}^n \alpha^j} \binom{p}{|J|}^{-1}, \quad \forall J \in \mathcal{P}_n(\{1, \dots, p\}), \quad \text{and} \quad \pi_J = 0 \quad \text{if} \quad |J| > n. \quad (2.5)$$

As a consequence, we obtain the following immediate corollary of Proposition 2.1.

Theorem 2.1. *Assume that the noise variables W_1, \dots, W_n are i.i.d. $N(0, \sigma^2)$. Then the aggregate $\hat{f}_n = f_{\hat{\theta}_n}$, with $\lambda = \frac{n}{4\sigma^2}$ and π taken as in (2.5), satisfies*

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}_n - f\|_n^2 \right] \\ & \leq \min_{\theta \in \mathbb{R}^p} \left\{ \|f_\theta - f\|_n^2 + \frac{\sigma^2 |J(\theta)|}{n} \left(4 \log \left(\frac{pe}{|J(\theta)|\alpha} \right) + 1 \right) + \frac{4\sigma^2 \log\left(\frac{1}{1-\alpha}\right)}{n} \right\}, \end{aligned} \quad (2.6)$$

where for any $\theta \in \mathbb{R}^p$ $J(\theta) = \{j : \theta_j \neq 0\}$.

This result improves upon [17] which established in Theorem 6 for gaussian noise and deterministic design

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}_n - f\|_n^2 \right] \\ & \leq \min_{\theta \in \mathbb{R}^p} \left\{ \|f_\theta - f\|_n^2 + \frac{16\sigma^2 |J(\theta)|}{n} \left(1 + \log_+ \left(\frac{\sqrt{\text{Tr}(A^\top A)}}{M(\theta)\sigma} |\theta|_1 \right) \right) + \frac{\sigma^2}{n} \right\}, \end{aligned}$$

where $\log_+ x = \max\{\log x, 0\}$ and we recall that $A = (\phi_j(X_i))_{1 \leq i \leq n, 1 \leq j \leq p}$. Note that our bound is faster by a factor $\log_+ |\theta|_1$. Note also that the bound in the above display grows worse for large values of $|\theta|_1$.

In order to evaluate the performance of these exponential weights procedures, [40] developed a notion of optimal rate of sparse prediction. In particular, [36] established that there exists a numerical constant $c^* > 0$ such that for all estimator T_n

$$\sup_{\substack{\theta \in \mathbb{R}^p \setminus \{0\}: \\ |J(\theta)| \leq s}} \sup_f \left\{ \mathbb{E} [\|T_n - f\|_n^2] - \|f_\theta - f\|_n^2 \right\} \geq c^* \frac{\sigma^2}{n} \left[\text{rank}(A) \wedge s \log \left(1 + \frac{ep}{s} \right) \right].$$

The above display combined with Theorem 2.1 shows that the exponential weights procedure (2.3) with the prior (2.5) achieves the optimal rate of sparse prediction for any vector θ satisfying $|J(\theta)| \leq \frac{\text{rank}(A)}{\log(1+ep)}$.

3. Sparsity oracle inequality in probability

In Section 2 we assumed the design is deterministic and we established an oracle inequality in expectation with the optimal rate of sparse prediction. We want now to establish an oracle inequality in probability that holds true for deterministic and random design.

From now on, the design can be either deterministic or random. We make the following mild assumption:

$$L = \max_{1 \leq j \leq M} \|\phi_j\|_\infty < \infty.$$

We assume in this section that the noise variables are subgaussian. More precisely we have the following condition.

Assumption 3.1. *The noise variables W_1, \dots, W_n are independent and independent of X_1, \dots, X_n . We assume also that there exist two known constants $\sigma > 0$ and $\xi > 0$ such that*

$$\begin{aligned} \mathbb{E}(W_i^2) &\leq \sigma^2 \\ \forall k \geq 3, \quad \mathbb{E}(|W_i|^k) &\leq \sigma^2 k! \xi^{k-2}. \end{aligned}$$

The estimation method is a version of the Gibbs estimator introduced in [13, 14]. Fix $K \geq 1$. First we define the prior probability distribution as follows. For any $J \subset \{1, \dots, p\}$ let u_J denote the uniform measure on $\Theta_{K+1}(J)$. We define

$$m(d\theta) = \sum_{J \subset \{1, \dots, p\}} \pi_J u_J(d\theta)$$

with π taken as in (2.5).

We are now ready to define our estimator. For any $\lambda > 0$ we consider the probability measure $\tilde{\rho}_\lambda$ admitting the following density w.r.t. the probability measure m

$$\frac{d\tilde{\rho}_\lambda}{dm}(\theta) = \frac{e^{-\lambda r(\theta)}}{\int_{\Theta_K} e^{-\lambda r} dm}. \quad (3.1)$$

The aggregate \tilde{f}_n is defined as follows

$$\tilde{f}_n = f_{\tilde{\theta}_n}, \quad \tilde{\theta}_n = \tilde{\theta}_n(\lambda, m) = \int_{\Theta_K} \theta \tilde{\rho}_\lambda(d\theta). \quad (3.2)$$

Define

$$\mathcal{C}_1 = [8\sigma^2 + (2\|f\|_\infty + L(2K+1))^2] \vee [8[\xi + (2\|f\|_\infty + L(2K+1))]L(2K+1)].$$

We can now state the main result of this section.

Theorem 3.1. *Let Assumption 3.1 be satisfied. Take $K > 1$ and $\lambda = \lambda^* = \frac{n}{2\mathcal{C}_1}$. Assume that $\arg \min_{\theta \in \mathbb{R}^p} R(\theta) \cap \Theta_K \neq \emptyset$. Then we have, for any $\varepsilon \in (0, 1)$ and any $\bar{\theta} \in \arg \min_{\theta \in \mathbb{R}^p} R(\theta) \cap \Theta_K$, with probability at least $1 - \varepsilon$,*

$$\begin{aligned} R(\tilde{\theta}_n) &\leq R(\bar{\theta}) + \frac{3L^2}{n^2} + \frac{8\mathcal{C}_1}{n} \left[|J(\bar{\theta})| \log(K+1) \right. \\ &\quad \left. + \left(|J(\bar{\theta})| \log \left(\frac{enp}{\alpha |J(\bar{\theta})|} \right) + \log \left(\frac{2}{\varepsilon(1-\alpha)} \right) \right) \right]. \end{aligned}$$

The choice $\lambda = \lambda^*$ comes from the optimization of a (rather pessimistic) upper bound on the risk R (see Inequality (4.9) in the proof of this theorem, page 142). However this choice is not necessarily the best choice in practice even though it gives the good order of magnitude for λ . The practitioner may use cross-validation to properly tune the temperature parameter.

Theorem 3.1 improves upon previous results on the following points:

1) Our oracle inequality is sharp in the sense that the leading factor in front of $R(\bar{\theta})$ is equal to 1 and we require only $\max_j \|\phi_j\| < \infty$ whereas l_1 -penalized empirical risk minimization procedures such as Lasso or Dantzig selector have a leading factor strictly larger than 1 and impose in addition stringent conditions on the dictionary (cf. [7, 11]). For instance, [7] imposes the design matrix $A = (\phi_j(X_i))_{1 \leq i \leq n, 1 \leq j \leq p}$ to be deterministic and to satisfy the following restricted eigenvalue condition for the Lasso

$$\kappa(s) \triangleq \min_{\substack{J_0 \subset \{1, \dots, p\}: \\ |J_0| \leq s}} \min_{\substack{\Delta \in \mathbb{R}^p \setminus \{0\}: \\ |\Delta_{J_0^c}|_1 \leq 3|\Delta_{J_0}|_1}} \frac{|A\Delta|_2}{\sqrt{n}|\Delta_{J_0}|_2} > 0,$$

where for any $\Delta \in \mathbb{R}^p$ and $J \subset \{1, \dots, p\}$, we denote by Δ_J the vector in \mathbb{R}^p which has the same components as Δ on J and zero coordinates on the complement J^c . Assuming in addition that the noise is gaussian $N(0, \sigma^2)$ and taking $\lambda = A\sigma\sqrt{\frac{\log p}{n}}$, $A > 8$, [7] proved that the Lasso $\hat{\theta}^L$ satisfies with probability at least $1 - M^{1-A^2/8}$

$$\frac{1}{n} \|f_{\hat{\theta}^L} - f\|_n^2 \leq (1 + \eta) \frac{1}{n} \|f_{\theta} - f\|_n^2 + C(\eta) A^2 \sigma^2 |J(\theta)| \frac{\log p}{n}, \quad \forall \theta \in \mathbb{R}^p,$$

where $\eta, C(\eta) > 0$ and $C(\eta)$ increases to $+\infty$ as η tends to 0.

On the downside, our estimator requires the additional condition $|\bar{\theta}|_1 \leq K$. This condition is common in the PAC-bayesian literature. Removing this condition is a difficult problem and does not seem possible with the actual techniques of proof where this condition is needed in order to apply Bernstein's inequality.

2) We establish a sparsity oracle inequality in probability for the integrated risk $R(\cdot)$ whereas previous results on the exponential weights are given in expectation [16, 17, 19, 24, 29, 36].

3) Unlike mirror averaging or progressive mixture rules, satisfying similar inequalities in expectation, our estimator does not involve an averaging step [18, 24, 29]. As a consequence, its computational complexity is significantly reduced as compared to those procedures with averaging step. For instance [29] considered the model (1.1) with random design and i.i.d. observations $(X_1, Y_1), \dots, (X_n, Y_n)$, $n \geq 2$. The studied estimator is the following mirror averaging scheme

$$f_{\hat{\theta}^{MA}}, \quad \theta^{MA} = \frac{1}{n} \sum_{k=0}^{n-1} \tilde{\theta}_k,$$

where $\tilde{\theta}_0 = \int_{\Theta(K)} \theta d\Pi$ and for any $1 \leq k \leq n-1$, $\tilde{\theta}_k$ is defined similarly as $\tilde{\theta}_n$ in (3.1)–(3.2) with $r(\theta)$ replaced by $r_k(\theta) = \frac{1}{k} \sum_{i=1}^k (Y_i - f_{\theta}(X_i))^2$. These estimators can be implemented for example by MCMC. In this case, computing the integral $\int_{\Theta(K)} \theta \tilde{\rho}_{\lambda}(d\theta)$ is the most time-consuming part of the procedure.

The procedure (3.1)–(3.2) requires computing this integral only once whereas the mirror averaging procedure of [29] requires computing integrals of this form n times.

4) Under the assumption $|\bar{\theta}|_1 \leq K$ for some absolute constant K and taking $\epsilon = n^{-1}$ we have with probability at least $1 - n^{-1}$

$$R(\tilde{\theta}_n) \leq R(\bar{\theta}) + C \frac{\mathcal{C}_1}{n} |J(\bar{\theta})| \log \left(\frac{enp}{\alpha |J(\bar{\theta})|} \right) + \frac{3L^2}{n^2}, \quad (3.3)$$

for some absolute constant $C > 0$. In [36] a minimax lower bound in expectation is established for deterministic design and Gaussian noise. A similar result holds in probability with the same proof combined with Theorem 2.7 of [41]. There exists absolute constants $c_1, c_2 > 0$ such that

$$\inf_{T_n} \sup_{\substack{\theta \in \mathbb{R}^p \\ |J(\theta)| \leq s}} \sup_f \mathbb{P} \left[\|T_n - f\|_n^2 \geq \|f_\theta - f\|_n^2 + c_1 \frac{\sigma^2}{n} \text{rank}(A) \wedge s \log \left(1 + \frac{ep}{s} \right) \right] \geq c_2.$$

If $s \leq \frac{\text{rank}(A)}{\log(1+ep)}$ then we observe that the upper bound in (3.3) is optimal up to the additional logarithmic factor $\log n$. Note however that if $p \geq n^{1+\delta}$ for some $\delta > 0$, which is relevant with the high-dimensional setting we consider in this paper, then our bound is rate optimal.

4. Proofs

4.1. Proofs of Section 2

This proof uses an argument from Leung and Barron [28].

Proof of Proposition 2.1. The mapping $Y \rightarrow \hat{f}_n(Y) \triangleq (\hat{f}_n(X_1, Y), \dots, \hat{f}_n(X_n, Y))^T$ is clearly continuously differentiable by composition of elementary differentiable functions. For any subset $J \subset \{1, \dots, p\}$ define $A_J = (\phi_j(X_i))_{1 \leq i \leq n, j \in J}$, $\Sigma_J = \frac{1}{n} A_J^T A_J$, $\Phi_J(\cdot) = (\phi_j(\cdot))_{j \in J}$ and

$$g_J = e^{-\lambda \left(\|Y - f_J\|_n^2 + \frac{2\sigma^2 |J|}{n} \right)}$$

where

$$f_J(x, Y) = \frac{1}{n} Y^T A_J \Sigma_J^+ \Phi_J(x)^T,$$

and Σ_J^+ denotes the pseudo-inverse of Σ_J . Denote by ∂_i the derivative w.r.t. Y_i . Simple computations give

$$\begin{aligned} \partial_i f_J(x, Y) &= \frac{1}{n} \Phi_J(X_i) \Sigma_J^+ \Phi_J(x)^T, \\ (\partial_i f_J(X_1, Y), \dots, \partial_i f_J(X_n, Y)) Y &= f_J(X_i, Y), \end{aligned}$$

and

$$\sum_{l=1}^n f_J(X_l, Y) \partial_i f_J(X_l, Y) = f_J(X_i, Y).$$

Thus we have

$$\begin{aligned} \partial_i(g_J) &= -\lambda \partial_i(\|Y - f_J\|_n^2) g_J \\ &= -\frac{2\lambda}{n} \left((Y_i - f_J(X_i, Y)) - \sum_{l=1}^n \partial_i f_J(X_l, Y) (Y_l - f_J(X_l, Y)) \right) g_J \\ &= -\frac{2\lambda}{n} (Y_i - f_J(X_i, Y)) g_J, \end{aligned}$$

Recall that

$$\hat{f}_n(\cdot, Y) = \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J f_J(\cdot, Y)}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J}.$$

We have

$$\begin{aligned} \partial_i \hat{f}_n(X_i, Y) &= \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J (\partial_i(g_J) f_J(X_i, Y) + g_J \partial_i(f_J(X_i, Y)))}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} \\ &\quad - \frac{(\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J f_J(X_i, Y)) (\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J \partial_i(g_J))}{(\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J)^2} \\ &= -\frac{2\lambda}{n} Y_i \hat{f}_n + \frac{2\lambda}{n} \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} f_J(X_i, Y)^2 \pi_J g_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} \\ &\quad + \frac{1}{n} \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \Phi_J(X_i) \Sigma_J^+ \Phi_J(X_i)^T \pi_J g_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} \\ &\quad - \frac{2\lambda}{n} Y_i \hat{f}_n(X_i, Y) - \frac{2\lambda}{n} \hat{f}_n^2(X_i, Y) \\ &= \frac{2\lambda}{n} \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} (f_J(X_i, Y) - \hat{f}_n(X_i, Y))^2 \pi_J g_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} \\ &\quad + \frac{1}{n} \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \Phi_J(X_i) \Sigma_J^+ \Phi_J(X_i)^T \pi_J g_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} \geq 0. \end{aligned} \quad (4.1)$$

Consider the following estimator of the risk

$$\hat{r}_n(Y) = \|\hat{f}_n(Y) - Y\|_n^2 + \frac{2\sigma^2}{n} \sum_{i=1}^n \partial_i \hat{f}_n(X_i, Y) - \sigma^2. \quad (4.2)$$

Using an argument based on Stein's identity as in [27] we now prove that

$$\mathbb{E}[\hat{r}_n(Y)] = \mathbb{E}[\|\hat{f}_n(Y) - f\|_n^2].$$

We have

$$\begin{aligned}\mathbb{E} \left[\|\hat{f}_n(Y) - f\|_n^2 \right] &= \mathbb{E} \left[\|\hat{f}_n(Y) - Y\|_n^2 + \frac{2}{n} \sum_{i=1}^n W_i (\hat{f}_n(X_i, Y) - f(X_i)) \right] - \sigma^2 \\ &= \mathbb{E} \left[\|\hat{f}_n(Y) - Y\|_n^2 + \frac{2}{n} \sum_{i=1}^n W_i \hat{f}_n(X_i, Y) \right] - \sigma^2. \quad (4.3)\end{aligned}$$

For $\mathbf{z} = (z_1, \dots, z_n)^T \in \mathbb{R}^n$ write $F_{W,i}(\mathbf{z}) = \prod_{j \neq i} F_{W,i}(z_j)$, where F_W denotes the c.d.f. of the random variable W_1 . Since $\mathbb{E}(W_i) = 0$ we have

$$\begin{aligned}\mathbb{E} \left[W_i \hat{f}_n(X_i, Y) \right] &= \mathbb{E} \left[W_i \int_0^{W_i} \partial_i \hat{f}_n(X_i, Y_1, \dots, Y_{i-1}, f(X_i) + z, Y_{i+1}, \dots, Y_n) dz \right] \\ &= \int_{\mathbb{R}^{n-1}} \left(\int_{\mathbb{R}} y \int_0^y \partial_i \hat{f}_n(X_i, f + \mathbf{z}) dz_i dF_W(y) \right) dF_{W,i}(\mathbf{z}). \quad (4.4)\end{aligned}$$

In view of (4.1) we can apply Fubini's Theorem to the right-hand-side of (4.4). We obtain under the assumption $W \sim \mathbb{N}(0, \sigma^2)$ that

$$\begin{aligned}\int_{\mathbb{R}^+} \int_0^y \partial_i \hat{f}_n(X_i, f + \mathbf{z}) dz_i dF_W(y) &= \int_{\mathbb{R}^+} \int_{z_i}^{\infty} y dF_W(y) \partial_i \hat{f}_n(X_i, f + \mathbf{z}) dz_i \\ &= \int_{\mathbb{R}^+} \sigma^2 \partial_i \hat{f}_n(X_i, f + \mathbf{z}) dF_W(z_i),\end{aligned}$$

A Similar equality holds for the integral over \mathbb{R}^- . Thus we obtain

$$\mathbb{E} \left[W_i \hat{f}_n(X_i, Y) \right] = \sigma^2 \mathbb{E} \left[\partial_i \hat{f}_n(X_i, Y) \right].$$

Combining (4.2), (4.3) and the above display gives

$$\mathbb{E} [\hat{r}_n(Y)] = \mathbb{E} \left[\|\hat{f}_n(Y) - f\|_n^2 \right].$$

Since $\hat{f}_n(\cdot, Y)$ is the expectation of $f_J(\cdot, Y)$ w.r.t. the probability distribution $\propto g \cdot \pi$, we have

$$\|\hat{f}_n(\cdot, Y) - Y\|_n^2 = \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} (\|f_J(\cdot, Y) - Y\|_n^2 - \|f_J(\cdot, Y) - \hat{f}_n(Y)\|_n^2) g_J \pi_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} g_J \pi_J}.$$

For the sake of simplicity set $f_J = f_J(\cdot, Y)$ and $\hat{f}_n = \hat{f}_n(\cdot, Y)$. Combining (4.2), the above display and $\lambda \leq \frac{n}{4\sigma^2}$ yields

$$\begin{aligned}\hat{r}_n(Y) &= \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left(\|f_J - Y\|_n^2 + \sum_{i=1}^n \left(\frac{4\lambda\sigma^2}{n} - 1 \right) \|f_J - \hat{f}_n\|_n^2 \right) g_J \pi_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} \\ &\quad + \frac{2\sigma^2}{n^2} \sum_{i=1}^n \frac{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \Phi_J(X_i) \Sigma_J^+ \Phi_J(X_i)^T \pi_J g_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \pi_J g_J} - \sigma^2 \\ &\leq \sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left(\|f_J - Y\|_n^2 + \frac{2\sigma^2}{n} |J| \right) g_J \pi_J - \sigma^2.\end{aligned}$$

By definition of g_J we have

$$\begin{aligned} \|f_J - Y\|_n^2 + \frac{2\sigma^2|J|}{n} &= -\frac{1}{\lambda} \log \left(\frac{g_J}{\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} g_J \pi_J} \right) \\ &\quad - \frac{1}{\lambda} \log \left(\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} g_J \pi_J \right). \end{aligned}$$

Integrating the above inequality w.r.t. the probability distribution $\frac{1}{C} g \cdot \pi$ (where $C = \sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} g_J \pi_J$ is the normalization factor) and using the fact that

$$\sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \frac{1}{C} g_J \pi_J \log \left(\frac{1}{C} g_J \right) = K \left(\frac{g \cdot \pi}{C}, \pi \right) \geq 0$$

as well as a convex duality argument (cf., e.g., [20], p. 264) we get

$$\hat{r}_n(Y) \leq \sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left(\|Y - f_J\|_n^2 + \frac{2\sigma^2}{n}|J| \right) \pi'_J + \frac{1}{\lambda} K(\pi', \pi) - \sigma^2,$$

for all probability measure π' on $\mathcal{P}(\{1, \dots, p\})$. Taking the expectation in the last inequality we get for any π'

$$\begin{aligned} \mathbb{E} \left[\|\hat{f}_n - f\|_n^2 \right] &= \mathbb{E}[\hat{r}_n(Y)] \\ &\leq \sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left(\mathbb{E}[\|f_J - Y\|_n^2] + \frac{2\sigma^2}{n}|J| \right) \pi'_J + \frac{1}{\lambda} K(\pi', \pi) - \sigma^2 \\ &\leq \sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left(\mathbb{E}[\|f_J - f\|_n^2] + \frac{2}{n} \sum_{i=1}^n \mathbb{E}[W_i f_J(X_i, Y)] + \frac{2\sigma^2}{n}|J| \right) \pi'_J \\ &\quad + \frac{1}{\lambda} K(\pi', \pi) \\ &\leq \sum_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left(\mathbb{E}[\|f_J - f\|_n^2] + \frac{4\sigma^2}{n}|J| \right) \pi'_J + \frac{1}{\lambda} K(\pi', \pi), \end{aligned}$$

where we have used Stein's argument $\mathbb{E}[W_i f_J(X_i, Y)] = \sigma^2 \mathbb{E}[\partial_i f_J(X_i, Y)]$ and the fact that $\sum_{i=1}^n \partial_i f_J(X_i, Y) = 1$ in the last line. Finally taking π' in the set of Dirac distributions on the subset J of $\{1, \dots, p\}$ yields the theorem. \square

Proof of Theorem 2.1. First note that any minimizer $\theta \in \mathbb{R}^p$ of the right-hand-side in (2.6) is such that $|J(\theta)| \leq \text{rank}(A) \leq n$ where we recall that $A = (\phi_j(X_i))_{1 \leq i \leq n, 1 \leq j \leq p}$. Indeed, for any $\theta \in \mathbb{R}^p$ such that $|J(\theta)| > \text{rank}(A)$ we can construct a vector $\theta' \in \mathbb{R}^p$ such that $f_\theta = f_{\theta'}$ and $|J(\theta')| \leq \text{rank}(A)$ and the mapping $x \rightarrow x \log \left(\frac{e p \alpha}{x} \right)$ is nondecreasing on $(0, p]$.

Next for any $J \in \mathcal{P}_n(\{1, \dots, p\})$ we have

$$\mathbb{E}[\|f_J - f\|_n^2] = \min_{\theta \in \Theta(J)} \left\{ \|f_\theta - f\|_n^2 \right\} + \frac{\sigma^2|J|}{n} = \min_{\theta \in \Theta(J)} \left\{ \|f_\theta - f\|_n^2 + \frac{\sigma^2|J(\theta)|}{n} \right\}.$$

Thus

$$\begin{aligned}
& \min_{J \in \mathcal{P}_n(\{1, \dots, p\})} \left\{ \mathbb{E}[\|f_J - f\|_n^2] + \frac{1}{\lambda} \log \left(\frac{1}{\pi_J} \right) + \frac{\sigma^2 J}{n} \right\} \\
&= \min_{J \in \mathcal{P}_n(\{1, \dots, p\})} \min_{\theta \in \Theta(J)} \left\{ \|f_\theta - f\|_n^2 + \frac{1}{\lambda} \log \left(\frac{1}{\pi_{J(\theta)}} \right) + \frac{\sigma^2 |J(\theta)|}{n} \right\} \\
&= \min_{\theta \in \mathbb{R}^p} \left\{ \|f_\theta - f\|_n^2 + \frac{1}{\lambda} \log \left(\frac{1}{\pi_{J(\theta)}} \right) + \frac{\sigma^2 |J(\theta)|}{n} \right\}.
\end{aligned}$$

Combining the above display with Proposition 2.1 and our definition of the prior π gives the result. \square

4.2. Proof of Theorem 3.1

We state below a version of Bernstein's inequality useful in the proof of Theorem 3.1. See Proposition 2.9 page 24 in [33], more precisely Inequality (2.21).

Lemma 4.1. *Let T_1, \dots, T_n be independent real valued random variables. Let us assume that there is two constants v and w such that*

$$\sum_{i=1}^n \mathbb{E}[T_i^2] \leq v$$

and for all integers $k \geq 3$,

$$\sum_{i=1}^n \mathbb{E}[(T_i)_+^k] \leq v \frac{k! w^{k-2}}{2}.$$

Then, for any $\zeta \in (0, 1/w)$,

$$\mathbb{E} \exp \left[\zeta \sum_{i=1}^n [T_i - \mathbb{E}(T_i)] \right] \leq \exp \left(\frac{v \zeta^2}{2(1 - w \zeta)} \right).$$

Proof of Theorem 3.1. For any $\theta \in \Theta_{K+1}$ define the random variables

$$T_i = T_i(\theta) = -(Y_i - f_\theta(X_i))^2 + (Y_i - f_{\bar{\theta}}(X_i))^2.$$

Note that these variables are independent. We have

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E}[T_i^2] &= \sum_{i=1}^n \mathbb{E} \left[[2Y_i - f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 [f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 \right] \\
&= \sum_{i=1}^n \mathbb{E} \left[[2W_i + 2f(X_i) - f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 [f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 \right]
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i=1}^n \mathbb{E} \left[[8W_i^2 + 2(2\|f\|_\infty + L(2K+1))^2] [f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 \right] \\
&= \sum_{i=1}^n \mathbb{E} [8W_i^2 + 2(2\|f\|_\infty + L(2K+1))^2] \mathbb{E} [f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 \\
&\leq n [8\sigma^2 + 2(2\|f\|_\infty + L(2K+1))^2] [R(\theta) - R(\bar{\theta})] =: v(\theta, \bar{\theta}) = v,
\end{aligned}$$

where we have used in the last line Pythagore's Theorem to prove $\|f_\theta - f_{\bar{\theta}}\|^2 = R(\theta) - R(\bar{\theta})$. Next we have, for any integer $k \geq 3$, that

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E} [(T_i)_+^k] &\leq \sum_{i=1}^n \mathbb{E} [|2Y_i - f_{\bar{\theta}}(X_i) - f_\theta(X_i)|^k |f_{\bar{\theta}}(X_i) - f_\theta(X_i)|^k] \\
&\leq \sum_{i=1}^n \mathbb{E} [2^{2k-1} [|W_i|^k + (\|f\|_\infty + L(K+1/2))^k] |f_{\bar{\theta}}(X_i) - f_\theta(X_i)|^k] \\
&\leq \sum_{i=1}^n \mathbb{E} [2^{2k-1} [|W_i|^k + (\|f\|_\infty + L(K+1/2))^k] [L(2K+1)]^{k-2} [f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2] \\
&\leq 2^{2k-1} [\sigma^2 k! \xi^{k-2} + (\|f\|_\infty + L(K+1/2))^k] [L(2K+1)]^{k-2} \\
&\quad \times \sum_{i=1}^n \mathbb{E} [f_{\bar{\theta}}(X_i) - f_\theta(X_i)]^2 \\
&\leq \frac{(\sigma^2 k! \xi^{k-2} + (\|f\|_\infty + L(K+1/2))^k) (4L(2K+1))^{k-2} v}{4(\sigma^2 + (\|f\|_\infty + L(K+1/2))^2)} \\
&\leq \frac{1}{4} (k! \xi^{k-2} + [\|f\|_\infty + L(K+1/2)]^{k-2}) [4L(2K+1)]^{k-2} v \\
&\leq \frac{2}{4} k! (\xi + [\|f\|_\infty + L(K+1/2)])^{k-2} [4L(2K+1)]^{k-2} v \leq v \frac{k! w^{k-2}}{2},
\end{aligned}$$

with $w := 8(\xi + [\|f\|_\infty + L(K+1/2)])L(K+1/2)$.

Next, for any $\lambda \in (0, n/w)$ and $\theta \in \Theta_{K+1}$, applying Lemma 4.1 with $\zeta = \lambda/n$ gives

$$\mathbb{E} \exp \left[\lambda \left(R(\theta) - R(\bar{\theta}) - r(\theta) + r(\bar{\theta}) \right) \right] \leq \exp \left[\frac{v\lambda^2}{2n^2 \left(1 - \frac{w\lambda}{n}\right)} \right].$$

Set $C = 8(\sigma^2 + [\|f\|_\infty + L(K+1/2)]^2)$. For the sake of simplicity let us put

$$\beta = \left(\lambda - \frac{\lambda^2 C}{2n \left(1 - \frac{w\lambda}{n}\right)} \right). \tag{4.5}$$

For any $\varepsilon > 0$ the last display yields

$$\mathbb{E} \exp \left[\beta \left(R(\theta) - R(\bar{\theta}) \right) + \lambda \left(-r(\theta) + r(\bar{\theta}) \right) - \log \frac{2}{\varepsilon} \right] \leq \frac{\varepsilon}{2}.$$

Integrating w.r.t. the probability distribution $m(\cdot)$ we get

$$\int \mathbb{E} \exp \left[\beta \left(R(\theta) - R(\bar{\theta}) \right) + \lambda \left(-r(\theta) + r(\bar{\theta}) \right) - \log \frac{2}{\varepsilon} \right] m(d\theta) \leq \frac{\varepsilon}{2}.$$

Next, Fubini's theorem gives

$$\begin{aligned} & \mathbb{E} \int \exp \left[\beta \left(R(\theta) - R(\bar{\theta}) \right) + \lambda \left(-r(\theta) + r(\bar{\theta}) \right) - \log \frac{2}{\varepsilon} \right] m(d\theta) \\ &= \mathbb{E} \int \exp \left[\beta \left(R(\theta) - R(\bar{\theta}) \right) + \lambda \left(-r(\theta) + r(\bar{\theta}) \right) - \log \left[\frac{d\tilde{\rho}_\lambda}{dm}(\theta) \right] - \log \frac{2}{\varepsilon} \right] \tilde{\rho}_\lambda(d\theta) \\ & \leq \frac{\varepsilon}{2}. \end{aligned}$$

Jensen's inequality yields

$$\mathbb{E} \exp \left[\beta \left(\int R d\tilde{\rho}_\lambda - R(\bar{\theta}) \right) + \lambda \left(- \int r d\tilde{\rho}_\lambda + r(\bar{\theta}) \right) - \mathcal{K}(\tilde{\rho}_\lambda, m) - \log \frac{2}{\varepsilon} \right] \leq \frac{\varepsilon}{2}.$$

Now, using the basic inequality $\exp(x) \geq \mathbf{1}_{\mathbb{R}_+}(x)$ we get

$$\mathbb{P} \left\{ \beta \left(\int R d\tilde{\rho}_\lambda - R(\bar{\theta}) \right) + \lambda \left(- \int r d\tilde{\rho}_\lambda + r(\bar{\theta}) \right) - \mathcal{K}(\tilde{\rho}_\lambda, m) - \log \frac{2}{\varepsilon} \geq 0 \right\} \leq \frac{\varepsilon}{2}.$$

Using Jensen's inequality again gives

$$\int R d\tilde{\rho}_\lambda \geq R \left(\int \theta \tilde{\rho}_\lambda(d\theta) \right) = R(\tilde{\theta}_\lambda).$$

Combining the last two displays we obtain

$$\mathbb{P} \left\{ R(\tilde{\theta}_\lambda) - R(\bar{\theta}) \leq \frac{\int r d\tilde{\rho}_\lambda - r(\bar{\theta}) + \frac{1}{\lambda} [\mathcal{K}(\tilde{\rho}_\lambda, m) + \log \frac{2}{\varepsilon}]}{\frac{\beta}{\lambda}} \right\} \geq 1 - \frac{\varepsilon}{2}.$$

Now, using Lemma 1.1.3 in Catoni [14] we obtain that

$$\mathbb{P} \left\{ R(\tilde{\theta}_\lambda) - R(\bar{\theta}) \leq \inf_{\rho \in \mathcal{M}_+^1(\Theta_{K+1})} \frac{\int r d\rho - r(\bar{\theta}) + \frac{1}{\lambda} [\mathcal{K}(\rho, m) + \log \frac{2}{\varepsilon}]}{\frac{\beta}{\lambda}} \right\} \geq 1 - \frac{\varepsilon}{2}. \quad (4.6)$$

We now want to bound from above $r(\theta) - r(\bar{\theta})$ by $R(\theta) - R(\bar{\theta})$. Applying Lemma 4.1 to $\tilde{T}_i(\theta) = -T_i(\theta)$ and similar computations as above yield successively

$$\mathbb{E} \exp \left[\lambda \left(R(\bar{\theta}) - R(\theta) + r(\theta) - r(\bar{\theta}) \right) \right] \leq \exp \left[\frac{v\lambda^2}{2n^2 \left(1 - \frac{w\lambda}{n} \right)} \right],$$

and so for any (data-dependent) ρ ,

$$\mathbb{E} \exp \left[\gamma \left(- \int R d\rho + R(\bar{\theta}) \right) + \lambda \left(\int r d\rho - r(\bar{\theta}) \right) - \mathcal{K}(\rho, \mathfrak{m}) - \log \frac{2}{\varepsilon} \right] \leq \frac{\varepsilon}{2},$$

where

$$\gamma = \left(\lambda + \frac{\lambda^2 C}{2n(1 - \frac{w\lambda}{n})} \right). \quad (4.7)$$

Now,

$$\mathbb{P} \left\{ \int r d\rho - r(\bar{\theta}) \leq \frac{\gamma}{\lambda} \left[\int R d\rho - R(\bar{\theta}) \right] + \frac{1}{\lambda} \left[\mathcal{K}(\rho, \mathfrak{m}) + \log \frac{2}{\varepsilon} \right] \right\} \geq 1 - \frac{\varepsilon}{2}. \quad (4.8)$$

Combining (4.8) and (4.6) with a union bound argument gives

$$\begin{aligned} & \mathbb{P} \left\{ R(\tilde{\theta}_\lambda) - R(\bar{\theta}) \right. \\ & \quad \left. \leq \inf_{\rho \in \mathcal{M}_+^1(\Theta_{K+1})} \frac{\gamma \left[\int R d\rho - R(\bar{\theta}) \right] + 2 \left[\mathcal{K}(\rho, \mathfrak{m}) + \log \frac{2}{\varepsilon} \right]}{\beta} \right\} \\ & \geq 1 - \varepsilon, \end{aligned}$$

where $\mathcal{M}_+^1(\Theta_{K+1})$ is the set of all probability measures over Θ_{K+1} .

Now for any $\delta \in (0, 1]$ taking ρ as the uniform probability measure on the set $\{t \in \Theta(J(\bar{\theta})) : |t - \bar{\theta}|_1 \leq \delta\} \subset \Theta_{K+1}(J(\bar{\theta}))$ gives

$$\begin{aligned} & \mathbb{P} \left\{ R(\tilde{\theta}_\lambda) \leq R(\bar{\theta}) + \frac{1}{\beta} \left[\gamma L^2 \delta^2 + 2 \left(|J(\bar{\theta})| \log \frac{K+1}{\delta} \right. \right. \right. \\ & \quad \left. \left. \left. + |J(\bar{\theta})| \log \frac{1}{\alpha} + \log \left(\frac{1}{1-\alpha} \right) + \log \left(\frac{p}{|J(\bar{\theta})|} \right) + \log \frac{2}{\varepsilon} \right) \right] \right\} \geq 1 - \varepsilon. \end{aligned}$$

Taking $\delta = n^{-1}$ and the inequality $\log \binom{p}{|J(\bar{\theta})|} \leq |J(\bar{\theta})| \log \frac{pe}{|J(\bar{\theta})|}$ gives

$$\begin{aligned} & \mathbb{P} \left\{ R(\tilde{\theta}_\lambda) \leq R(\bar{\theta}) + \frac{1}{1 - \frac{\lambda C}{2(n-w\lambda)}} \left[\left(1 + \frac{\lambda C}{2(n-w\lambda)} \right) \frac{L^2}{n^2} \right. \right. \\ & \quad \left. \left. + \frac{2}{\lambda} \left(|J(\bar{\theta})| \log(K+1) + |J(\bar{\theta})| \log \left(\frac{epn}{\alpha |J(\bar{\theta})|} \right) + \log \left(\frac{2}{\varepsilon(1-\alpha)} \right) \right) \right] \right\} \geq 1 - \varepsilon \end{aligned} \quad (4.9)$$

where we replaced γ and β by their definitions, see (4.5) and (4.7). Taking now $\lambda = n/(2\mathcal{C}_1)$ (where we recall that $\mathcal{C}_1 = C \vee w$) in (4.9) gives

$$\mathbb{P}\left\{R(\tilde{\theta}_\lambda) \leq R(\bar{\theta}) + \frac{3L^2}{n^2} + \frac{8C_1}{n} \left[|J(\bar{\theta})| \log(K+1) + \left(|J(\bar{\theta})| \log\left(\frac{enp}{\alpha|J(\bar{\theta})|}\right) + \log\left(\frac{2}{\varepsilon(1-\alpha)}\right) \right) \right] \right\} \geq 1 - \varepsilon,$$

where we have used that $1 - \frac{\lambda C}{2(n-w\lambda)} \geq 1/2$ and $1 + \frac{\lambda C}{2(n-w\lambda)} \leq 3/2$. \square

References

- [1] AKAIKE, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. *2nd International Symposium on Information Theory*. Budapest: Akademia Kiado. B. N. Petrov and F. Csaki, editors. 267–281.
- [2] ALQUIER, P. (2006). *Transductive and Inductive Adaptive Inference for Regression and Density Estimation*. PhD Thesis, Université Paris 6.
- [3] ALQUIER, P. (2008). PAC-Bayesian bounds for randomized empirical risk minimizers. *Mathematical Methods of Statistics* **17** (4) 279–304. [MR2483458](#)
- [4] AUDIBERT, J.-Y. (2004). Aggregated Estimators and Empirical Complexity for Least Square Regression. *Annales de l'Institut Henri Poincaré B: Probability and Statistics* **40** (6) 685–736. [MR2096215](#)
- [5] AUDIBERT, J.-Y. (2004). *PAC-Bayesian Statistical Learning Theory*. PhD Thesis, Université Paris 6.
- [6] BACH, F. (2008). Bolasso: model consistent Lasso estimation through the bootstrap. *Proceedings of the 25th international conference on Machine learning*, ACM, New-York.
- [7] BICKEL, P. J., RITOV, Y. AND TSYBAKOV, A. (2009). Simultaneous Analysis of Lasso and Dantzig Selector. *The Annals of Statistics* **37** (4) 1705–1732. [MR2533469](#)
- [8] BUNEA, F., TSYBAKOV, A. AND WEGKAMP, M. (2007). Sparsity oracle inequalities for the Lasso. *Electronic Journal of Statistics* **1** 169–194. [MR2397610](#)
- [9] BUNEA, F., TSYBAKOV, A. AND WEGKAMP, M. (2007). Aggregation for Gaussian regression. *The Annals of Statistics* **35** 1674–1697. [MR2351101](#)
- [10] CAI, T., XU, G. AND ZHANG, J. (2009). On Recovery of Sparse Signals via ℓ_1 Minimization. *IEEE Transactions on Information Theory* **55** 3388–3397. [MR2598028](#)
- [11] CANDÈS, E. AND TAO, T. (2007) The Dantzig selector: statistical estimation when p is much larger than n . *The Annals of Statistics* **35**. [MR2382644](#)
- [12] CATONI, O. (2003) *A PAC-Bayesian approach to adaptive classification*. Preprint Laboratoire de Probabilités et Modèles Aléatoires PMA-840.
- [13] CATONI, O. (2004) *Statistical Learning Theory and Stochastic Optimization*. Saint-Flour Summer School on Probability Theory 2001 (Jean Picard ed.), Lecture Notes in Mathematics, Springer. [MR2163920](#)

- [14] CATONI, O. (2007) *PAC-Bayesian Supervised Classification (The Thermodynamics of Statistical Learning)*. IMS Lecture Notes - Monograph Series **56**. [MR2483528](#)
- [15] CHEN, S. S., DONOHO, D. L. AND SAUNDERS, M. A. (2001). Atomic Decomposition by Basis Pursuit. *SIAM Review* **43** (1) 129–159. [MR1639094](#)
- [16] DALALYAN, A. AND TSYBAKOV, A. (2007). Aggregation by exponential weighting and sharp oracle inequalities. *Proceedings of the 20th annual conference on Learning theory*, Springer-Verlag Berlin, Heidelberg. [MR2397581](#)
- [17] DALALYAN, A. AND TSYBAKOV, A. (2008). Aggregation by exponential weighting, sharp oracle inequalities and sparsity. *Machine Learning* **72** (1-2) 39–61.
- [18] DALALYAN, A. AND TSYBAKOV, A. (2010). *Mirror averaging with sparsity priors*. In minor revision for Bernoulli, [arXiv:1003.1189](#).
- [19] DALALYAN, A. AND TSYBAKOV, A. (2010). *Sparse Regression Learning by Aggregation and Langevin Monte-Carlo*. Preprint [arXiv:0903.1223](#).
- [20] DEMBO, A. AND ZEITOUNI, O. (1998). *Large Deviation Techniques and Applications*. Springer. [MR1619036](#)
- [21] EFRON, B., HASTIE, T., JOHNSTONE, I. AND TIBSHIRANI, R. (2004). Least Angle Regression. *The Annals of Statistics* **32** (2) 407–499. [MR2060166](#)
- [22] FRANK, L. AND FRIEDMAN, J. (1993). A Statistical View on Some Chemometrics Regression Tools. *Technometrics* **16** 199–511.
- [23] GHOSH, S. (2007). *Adaptive Elastic Net: an Improvement of Elastic Net to achieve Oracle Properties*. Preprint.
- [24] JUDITSKY, A., RIGOLLET, P. AND TSYBAKOV, A. (2008). Learning by mirror averaging. *The Annals of Statistics* **36** (5) 2183–2206. [MR2458184](#)
- [25] KOLTCHINSKII, V. (2010). Sparsity in Empirical Risk Minimization. *Annales de l'Institut Henri Poincaré B: Probability and Statistics*, to appear.
- [26] KOLTCHINSKII, V. (2009). The Dantzig selector and sparsity oracle inequalities. *Bernoulli* **15** (3) 799–828. [MR2555200](#)
- [27] LEHMANN, E. L. AND CASELLA, G. (1998). *Theory of Point Estimation*. Springer. [MR1639875](#)
- [28] LEUNG, G. AND BARRON, A. R. (2006). Information theory and mixing least-squares regressions. *IEEE Transactions on Information Theory* **52** (8) 3396–3410. [MR2242356](#)
- [29] LOUNICI, K. (2007). Generalized mirror averaging and d -convex aggregation. *Mathematical Methods of Statistics* **16** (3). [MR2356820](#)
- [30] LOUNICI, K. (2008). Sup-norm convergence rate and sign concentration property of Lasso and Dantzig estimators. *Electronic Journal of Statistics* **2** 90–102. [MR2386087](#)
- [31] MALLOWS, C. L. (1973). Some comments on c_p . *Technometrics* **15** 661–676. [MR1365719](#)
- [32] MARIN, J.-M. AND ROBERT, C. P. (2007). *Bayesian Core: A practical approach to computational Bayesian analysis*. Springer. [MR2289769](#)

- [33] MASSART, P. (2007). *Concentration Inequalities and Model Selection*. Saint-Flour Summer School on Probability Theory 2003 (Jean Picard ed.), Lecture Notes in Mathematics, Springer. [MR2319879](#)
- [34] MCALLESTER D. A. (1998). Some pac-bayesian theorems. *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*. ACM, 230–234. [MR1811587](#)
- [35] MEINSHAUSEN, N. AND BÜHLMANN, P. (2010). Stability selection. *Journal of the Royal Statistical Society B*, to appear.
- [36] RIGOLLET, P. AND TSYBAKOV, A. (2010). Exponential Screening and optimal rates of sparse estimation. *The Annals of Statistics*, to appear.
- [37] SCHWARZ, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, **6** 461–464. [MR0468014](#)
- [38] SHAWE-TAYLOR, J. AND WILLIAMSON, R. (1997). A PAC analysis of a bayes estimator. *Proceedings of the Tenth Annual Conference on Computational Learning Theory*. ACM, 2–9.
- [39] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society B* **58** (1) 267–288. [MR1379242](#)
- [40] TSYBAKOV, A. (2003). Optimal rates of aggregation. *Computational Learning theory and Kernel Machines*, Lecture Notes in Artificial Intelligence n.2777, Springer, Heidelberg. 303–313.
- [41] TSYBAKOV, A. (2009). *Introduction to nonparametric estimation*. Springer, New York. [MR2724359](#)
- [42] VAN DE GEER, S. A. AND BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the LASSO. *Electronic Journal of Statistics* **3** 1360–1392. [MR2576316](#)
- [43] YANG, Y. (2004). Aggregating regression procedures to improve performance. *Bernoulli*, **10** (1) 25–47. [MR2044592](#)
- [44] ZHANG, T. (2008). From epsilon-entropy to KL-entropy: analysis of minimum information complexity density estimation. *The Annals of Statistics* **34** 2180–2210. [MR2291500](#)
- [45] ZOU, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, **101** 1418–1429. [MR2279469](#)
- [46] ZOU, H. AND HASTIE, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society B* **67** (2) 301–320. [MR2137327](#)