

EXPONENTIAL SCREENING AND OPTIMAL RATES OF SPARSE ESTIMATION

BY PHILIPPE RIGOLLET¹ AND ALEXANDRE TSYBAKOV

Princeton University and CREST

In high-dimensional linear regression, the goal pursued here is to estimate an unknown regression function using linear combinations of a suitable set of covariates. One of the key assumptions for the success of any statistical procedure in this setup is to assume that the linear combination is sparse in some sense, for example, that it involves only few covariates. We consider a general, nonnecessarily linear, regression with Gaussian noise and study a related question, that is, to find a linear combination of approximating functions, which is at the same time sparse and has small mean squared error (MSE). We introduce a new estimation procedure, called *Exponential Screening*, that shows remarkable adaptation properties. It adapts to the linear combination that optimally balances MSE and sparsity, whether the latter is measured in terms of the number of nonzero entries in the combination (ℓ_0 norm) or in terms of the global weight of the combination (ℓ_1 norm). The power of this adaptation result is illustrated by showing that Exponential Screening solves optimally and simultaneously all the problems of aggregation in Gaussian regression that have been discussed in the literature. Moreover, we show that the performance of the Exponential Screening estimator cannot be improved in a minimax sense, even if the optimal sparsity is known in advance. The theoretical and numerical superiority of Exponential Screening compared to state-of-the-art sparse procedures is also discussed.

1. Introduction. The theory of estimation in high-dimensional statistical models under the sparsity scenario has been considerably developed during the recent years. One of the main achievements was to derive *sparsity oracle inequalities* (SOI), that is, bounds on the risk of various sparse estimation procedures in terms of the ℓ_0 norm (number of nonzero components) of the estimated vectors or their approximations [see Bickel, Ritov and Tsybakov (2009), Bunea, Tsybakov and Wegkamp (2007a, 2007b), Candès and Tao (2007), Koltchinskii (2010, 2009a, 2009b), van de Geer (2008), Zhang and Huang (2008), Zhang (2009) and references therein]. The main message of these results was to demonstrate that if the number of nonzero components of a high-dimensional target vector is small, then it can be reasonably well estimated even when the ambient dimension is larger

Received March 2010; revised July 2010.

¹Supported in part by the NSF Grant DMS-09-06424.

MSC2010 subject classifications. Primary 62G08; secondary 62G05, 62J05, 62C20, 62G20.

Key words and phrases. High-dimensional regression, aggregation, adaptation, sparsity, sparsity oracle inequalities, minimax rates, Lasso, BIC.

than the sample size. However, there were relatively few discussions of the optimality of these bounds; they were mainly based on specific counter-examples or referred to the paper by [Donoho and Johnstone \(1994a\)](#), which treats the Gaussian sequence model. The latter approach is, in general, insufficient as we will show below. An interesting point related to the optimality issue is that some of the bounds in the papers mentioned above involve not only the ℓ_0 norm but also the ℓ_1 norm of the target vector, which is yet another characteristic of sparsity. Thus, a natural question is whether the ℓ_1 norm plays an intrinsic role in the SOI or it appears there due to the techniques employed in the proof.

In this paper, considering the regression model with fixed design, we will show that the role of ℓ_1 norm is indeed intrinsic. Once we have a “rather general SOI” in terms the ℓ_0 norm, a SOI in terms of the ℓ_1 norm follows as a consequence. This means that we can write the resulting bound with the rate which is equal to the minimum of the ℓ_0 and ℓ_1 rates (see [Theorem 3.2](#)). Unfortunately, the above mentioned “rather general SOI” is not available in the literature for the previously known sparse estimation procedures. We therefore suggest a new procedure called the *Exponential Screening* (ES), which satisfies the desired bound. It is based on exponentially weighted aggregation of least squares estimators with suitably chosen prior. The idea of using exponentially weighted aggregation for sparse estimation is due to [Dalalyan and Tsybakov \(2007\)](#). [Dalalyan and Tsybakov \(2007, 2008, 2009, 2010\)](#) suggested several procedures of this kind based on continuous sparsity priors. Our approach is different because we use a discrete prior in the spirit of earlier work by [George \(1986a, 1986b\)](#), [Leung and Barron \(2006\)](#), [Giraud \(2008\)](#). Unlike [George \(1986a, 1986b\)](#), [Leung and Barron \(2006\)](#), [Giraud \(2008\)](#), we focus on high-dimensional models and treat explicitly the sparsity issue. Because of the high dimensionality of the problem, we need efficient computational algorithms, and therefore we suggest a version of the Metropolis–Hastings algorithm to approximate our estimators ([Section 7.1](#)). Regarding the sparsity issue, we prove that our method benefits simultaneously from three types of sparsity. The first one is expressed by the small rank of the design matrix \mathbf{X} , the second by the small number of nonzero components of the target vector and the third by its small ℓ_1 norm. Finally, we mention that in a work parallel to ours, [Alquier and Lounici \(2010\)](#) consider exponentially weighted aggregates with priors involving both discrete and continuous components and suggest another version of the Metropolis–Hastings algorithm to compute them.

The contributions of this paper are the following:

(i) We propose the ES estimator which benefits simultaneously from the above-mentioned three types of sparsity. This follows from the oracle inequalities that we prove in [Section 3](#). We also provide an efficient and fast algorithm to approximately compute the ES estimator and show that it outperforms several other competitive estimators in a simulation study.

(ii) We show that the ES estimator attains the *optimal rate of sparse estimation*. To this end, we establish a minimax lower bound which coincides with the upper bound on the risk of the ES estimator on the intersection of the ℓ_0 and ℓ_1 balls (Theorem 5.3).

(iii) As a consequence, we find *optimal rates of aggregation* for the regression model with fixed design. We consider the five main types of aggregation, which are the linear, convex, model selection, subset selection and D -convex aggregation [cf. Nemirovski (2000), Tsybakov (2003), Bunea, Tsybakov and Wegkamp (2007b), Lounici (2007)]. We show that the optimal rates are different from those for the regression model with random design established in Tsybakov (2003). Indeed, they turn out to be moderated by the rank of the regression matrix \mathbf{X} . The rates are faster for the smaller ranks. See Section 6.

This paper is organized as follows. After setting the problem and the notation in Section 2, we introduce the ES estimator in Section 3 and prove that it satisfies a SOI with a remainder term obtained as the minimum of the ℓ_0 and the ℓ_1 rate. This result holds with no assumption on the design matrix \mathbf{X} , except for simple normalization. We put it into perspective in Section 4 where we compare it with weaker SOI for the BIC and the Lasso estimators. In Sections 5.1 and 5.2 we discuss the optimality of SOI. In particular, Section 5.1 comments on why a minimax result in Donoho and Johnstone (1994a) with normalization depending on the unknown parameter is not suitable to treat optimality. Instead, we propose to consider minimax optimality on the intersection of ℓ_0 and ℓ_1 balls. In Section 5.2 we prove the corresponding minimax lower bound for all estimators and show rate optimality of the ES estimator in this sense. Section 6 discusses corollaries of our main results for the problem of aggregation; we show that the ES estimator solves simultaneously and optimally the five problems of aggregation mentioned in (iii) above. Finally, Section 7 presents a simulation study demonstrating a good performance of the ES estimator in numerical experiments.

2. Model and notation. Let $\mathcal{Z} := \{(x_1, Y_1), \dots, (x_n, Y_n)\}$ be a collection of independent random couples such that $(x_i, Y_i) \in \mathcal{X} \times \mathbb{R}$, where \mathcal{X} is an arbitrary set. Assume the regression model

$$Y_i = \eta(x_i) + \xi_i, \quad i = 1, \dots, n,$$

where $\eta: \mathcal{X} \rightarrow \mathbb{R}$ is the unknown regression function and the errors ξ_i are independent Gaussian $\mathcal{N}(0, \sigma^2)$. The covariates are deterministic elements x_1, \dots, x_n of \mathcal{X} . Consider the equivalence relation \sim on the space of functions $f: \mathcal{X} \rightarrow \mathbb{R}$ such that $f \sim g$ if and only if $f(x_i) = g(x_i)$ for all $i = 1, \dots, n$. Denote by $\mathcal{Q}_{1:n}$ the quotient space associated with this equivalence relation and define the norm $\|\cdot\|$ by

$$\|f\|^2 := \frac{1}{n} \sum_{i=1}^n f^2(x_i), \quad f \in \mathcal{Q}_{1:n}.$$

Notice that $\|\cdot\|$ is a norm on the quotient space but only a seminorm on the whole space of functions $f: \mathcal{X} \rightarrow \mathbb{R}$. Hereafter, we refer to it as a norm. We also define the associated inner product

$$\langle f, g \rangle := \frac{1}{n} \sum_{i=1}^n f(x_i)g(x_i).$$

Let $\mathcal{H} := \{f_1, \dots, f_M\}$, be a dictionary of $M \geq 1$ given functions $f_j: \mathcal{X} \rightarrow \mathbb{R}$. We approximate the regression function η by a linear combination $f_\theta(x) = \sum_{j=1}^M \theta_j f_j(x)$ with weights $\theta = (\theta_1, \dots, \theta_M)$, where possibly $M \gg n$.

We denote by \mathbf{X} , the $n \times M$ design matrix with elements $\mathbf{X}_{i,j} = f_j(x_i)$, $i = 1, \dots, n, j = 1, \dots, M$. We also introduce the column vectors $\mathbf{f} = (\eta(x_1), \dots, \eta(x_n))^\top$, $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ and $\xi = (\xi_1, \dots, \xi_n)^\top$. Let $|\cdot|_p$ denote the ℓ_p norm in \mathbb{R}^d for $p, d \geq 1$ and $M(\theta)$ denote the ℓ_0 norm of $\theta \in \mathbb{R}^M$, that is, the number of nonzero elements of $\theta \in \mathbb{R}^M$. For two real numbers a and b we use the notation $a \wedge b := \min(a, b)$, $a \vee b := \max(a, b)$; we denote by $[a]$ the integer part of a and by $\lceil a \rceil$ the smallest integer greater than or equal to a .

3. Sparsity pattern aggregation and Exponential Screening. A *sparsity pattern* is a binary vector $\mathbf{p} \in \mathcal{P} := \{0, 1\}^M$. The terminology comes from the fact that the coordinates of any such vectors can be interpreted as indicators of presence ($\mathbf{p}_j = 1$) or absence ($\mathbf{p}_j = 0$) of a given feature indexed by $j \in \{1, \dots, M\}$. We denote by $|\mathbf{p}|$ the number of ones in the sparsity pattern \mathbf{p} and by $\mathbb{R}^{\mathbf{p}}$ the space defined by

$$\mathbb{R}^{\mathbf{p}} = \{\theta \cdot \mathbf{p} : \theta \in \mathbb{R}^M\} \subset \mathbb{R}^M,$$

where $\theta \cdot \mathbf{p} \in \mathbb{R}^M$ denotes the Hadamard product between θ and \mathbf{p} and is defined as the vector $(\theta \cdot \mathbf{p})_j = \theta_j \mathbf{p}_j, j = 1, \dots, M$.

For any $\mathbf{p} \in \mathcal{P}$, let $\hat{\theta}_{\mathbf{p}}$ be any least squares estimator defined by

$$(3.1) \quad \hat{\theta}_{\mathbf{p}} \in \arg \min_{\theta \in \mathbb{R}^{\mathbf{p}}} \|\mathbf{Y} - \mathbf{X}\theta\|_2^2.$$

The following simple lemma gives an oracle inequality for the least squares estimator. Let $\text{rk}(\mathbf{X}) \leq M \wedge n$ denote the rank of the design matrix \mathbf{X} .

LEMMA 3.1. *Fix $\mathbf{p} \in \mathcal{P}$. Then any least squares estimator $\hat{\theta}_{\mathbf{p}}$ defined in (3.1) satisfies*

$$(3.2) \quad \mathbb{E} \|\mathbf{f}_{\hat{\theta}_{\mathbf{p}}} - \eta\|^2 = \min_{\theta \in \mathbb{R}^{\mathbf{p}}} \|\mathbf{f}_\theta - \eta\|^2 + \sigma^2 \frac{R_{\mathbf{p}}}{n} \leq \min_{\theta \in \mathbb{R}^{\mathbf{p}}} \|\mathbf{f}_\theta - \eta\|^2 + \sigma^2 \frac{|\mathbf{p}| \wedge R}{n},$$

where $R_{\mathbf{p}}$ is the dimension of the linear subspace $\{\mathbf{X}\theta : \theta \in \mathbb{R}^{\mathbf{p}}\}$ and $R = \text{rk}(\mathbf{X})$. Moreover, the random variables ξ_1, \dots, ξ_n need not be Gaussian for (3.2) to hold.

Proof of the lemma is straightforward in view of the Pythagorean theorem.

Let $\pi = (\pi_p)_p$ be a probability measure on \mathcal{P} , which we will further call a prior. The *sparsity pattern aggregate* (SPA) estimator is defined as $\hat{f}_{\tilde{\theta}^{\text{SPA}}}$, where

$$\tilde{\theta}^{\text{SPA}} := \frac{\sum_{p \in \mathcal{P}} \hat{\theta}_p \exp((-1/(4\sigma^2)) \sum_{i=1}^n (Y_i - f_{\hat{\theta}_p}(x_i))^2 - R_p/2) \pi_p}{\sum_{p \in \mathcal{P}} \exp((-1/(4\sigma^2)) \sum_{i=1}^n (Y_i - f_{\hat{\theta}_p}(x_i))^2 - R_p/2) \pi_p}.$$

As shown in [Leung and Barron \(2006\)](#), the following oracle inequality holds:

$$(3.3) \quad \mathbb{E} \|\hat{f}_{\tilde{\theta}^{\text{SPA}}} - \eta\|^2 \leq \min_{p \in \mathcal{P}: \pi_p \neq 0} \left\{ \mathbb{E} \|\hat{f}_{\hat{\theta}_p} - \eta\|^2 + \frac{4\sigma^2 \log(\pi_p^{-1})}{n} \right\}.$$

Now, we consider a specific choice of the prior π

$$(3.4) \quad \pi_p := \begin{cases} \frac{1}{H} \left(\frac{|\mathbf{p}|}{2eM} \right)^{|\mathbf{p}|}, & \text{if } |\mathbf{p}| < R, \\ \frac{1}{2}, & \text{if } |\mathbf{p}| = M, \\ 0, & \text{otherwise,} \end{cases}$$

where $R = \text{rk}(\mathbf{X})$, we use the convention $0^0 = 1$ and $H = 2 \sum_{k=0}^R \binom{M}{k} \left(\frac{k}{2eM}\right)^k$ is a normalization factor. In this paper we study the SPA estimator with the prior defined in (3.4). We call it the *Exponential Screening* (ES) estimator, and denote by $\tilde{\theta}^{\text{ES}}$ the estimator $\tilde{\theta}^{\text{SPA}}$ with the prior (3.4). The ES estimator is a mixture of least squares estimators corresponding essentially to sparsity patterns \mathbf{p} with small size and small residual sum of squares. Note that the weight 1/2 is assigned to the least squares estimator on the whole space (case where $|\mathbf{p}| = M$) and can be changed to any other constant in $(0, 1)$ without modifying the rates presented below, as long as H is modified accordingly.

Since $\binom{M}{k} \leq \left(\frac{eM}{k}\right)^k$, we obtain that $H \leq 4$. Using this and considering separately the cases $|\mathbf{p}| \leq 1$ and $|\mathbf{p}| \geq 2$, we obtain that the remainder term in (3.3) satisfies

$$(3.5) \quad \begin{aligned} \frac{4\sigma^2 \log(\pi_p^{-1})}{n} &\leq \frac{4\sigma^2}{n} \left[|\mathbf{p}| \log\left(\frac{2eM}{|\mathbf{p}| \vee 1}\right) + \log 4 \right] \\ &\leq \frac{8\sigma^2 |\mathbf{p}|}{n} \log\left(1 + \frac{eM}{|\mathbf{p}| \vee 1}\right) \\ &\quad + \frac{8\sigma^2}{n} \log 2 \end{aligned}$$

for sparsity patterns \mathbf{p} such that $|\mathbf{p}| < R$. Together with (3.3), this inequality yields the following theorem.

THEOREM 3.1. *For any $M \geq 1, n \geq 1$, the Exponential Screening estimator satisfies the following sparsity oracle inequality:*

$$(3.6) \quad \mathbb{E} \|f_{\hat{\theta}^{ES}} - \eta\|^2 \leq \min_{\theta \in \mathbb{R}^M} \left\{ \|f_{\theta} - \eta\|^2 + \frac{\sigma^2 R}{n} \right. \\ \left. \wedge \frac{9\sigma^2 M(\theta)}{n} \log \left(1 + \frac{eM}{M(\theta) \vee 1} \right) \right\} \\ + \frac{8\sigma^2}{n} \log 2,$$

where $R \leq M \wedge n$ denotes the rank of the design matrix \mathbf{X} .

PROOF. Combining the result of Lemma 3.1 and (3.3) with the sparsity prior defined in (3.4), we obtain that $\mathbb{E} \|f_{\hat{\theta}^{ES}} - \eta\|^2$ is bounded from above by

$$(3.7) \quad \min_{\substack{\theta \in \mathbb{R}^M \\ M(\theta) < R}} \left\{ \|f_{\theta} - \eta\|^2 + 9\sigma^2 \frac{M(\theta)}{n} \log \left(1 + \frac{eM}{M(\theta) \vee 1} \right) \right\} + \frac{8\sigma^2}{n} \log 2$$

and by

$$(3.8) \quad \min_{\theta \in \mathbb{R}^M} \left\{ \|f_{\theta} - \eta\|^2 + \sigma^2 \frac{R}{n} \right\} + \frac{4\sigma^2}{n} \log 2.$$

Combining (3.7) and (3.8) concludes the proof. \square

An interesting corollary of Theorem 3.1 is obtained for the linear regression model where it is assumed that $\eta = f_{\theta^*}$ for some $\theta^* \in \mathbb{R}^M$. In this case (3.6) yields

$$\mathbb{E} \|f_{\hat{\theta}^{ES}} - f_{\theta^*}\|^2 \leq \frac{\sigma^2 R}{n} \wedge \frac{9\sigma^2 M(\theta^*)}{n} \log \left(1 + \frac{eM}{M(\theta^*) \vee 1} \right) + \frac{8\sigma^2}{n} \log 2.$$

However, even in this parametric case, Theorem 3.1 provides a stronger result. Indeed, if there exists $\theta' \in \mathbb{R}^M$, such that

$$(3.9) \quad \|f_{\theta'} - f_{\theta^*}\|^2 + \frac{\sigma^2 R}{n} \wedge \frac{9\sigma^2 M(\theta')}{n} \log \left(1 + \frac{eM}{M(\theta') \vee 1} \right) \\ < \frac{\sigma^2 R}{n} \wedge \frac{9\sigma^2 M(\theta^*)}{n} \log \left(1 + \frac{eM}{M(\theta^*) \vee 1} \right),$$

then Theorem 3.1 gives a tighter bound on $\mathbb{E} \|f_{\hat{\theta}^{ES}} - f_{\theta^*}\|^2$. A vector $\theta' \in \mathbb{R}^M$ that satisfies (3.9) exists when f_{θ^*} can be well approximated by $f_{\theta'}$, and θ' is much sparser than θ^* .

While the sparsity oracle inequality (3.6) indicates that the ES estimator adapts to the underlying sparsity when measured in terms of the number of nonzero coefficients $M(\theta)$, it is also adaptive to the sparsity when measured in terms of the

ℓ_1 norm $|\theta|_1 = \sum_j |\theta|_j$. This can become an advantage when θ has many small coefficients so that $|\theta|_1 \ll M(\theta)$. Indeed, the following theorem shows that the ES estimator also enjoys adaptation in terms of its ℓ_1 norm.

THEOREM 3.2. *Assume that $\max_{1 \leq j \leq M} \|f_j\| \leq 1$. Then for any $M \geq 1, n \geq 1$ the Exponential Screening estimator satisfies*

$$(3.10) \quad \mathbb{E} \|f_{\hat{\theta}^{ES}} - \eta\|^2 \leq \min_{\theta \in \mathbb{R}^M} \{ \|f_\theta - \eta\|^2 + \bar{\varphi}_{n,M}(\theta) \} + \frac{\sigma^2}{n} (9 \log(1 + eM) + 8 \log 2),$$

where $\varphi_{n,M}(0) := 0$, and, for $\theta \neq 0$,

$$(3.11) \quad \varphi_{n,M}(\theta) := \frac{\sigma^2 R}{n} \wedge \frac{9\sigma^2 M(\theta)}{n} \log \left(1 + \frac{eM}{M(\theta) \vee 1} \right) \wedge \frac{11\sigma |\theta|_1}{\sqrt{n}} \sqrt{\log \left(1 + \frac{3eM\sigma}{|\theta|_1 \sqrt{n}} \right)}.$$

Furthermore, for any $\theta \in \mathbb{R}^M$, such that $\langle f_\theta, \eta \rangle \leq \|f_\theta\|^2$, we have

$$(3.12) \quad \mathbb{E} \|f_{\hat{\theta}^{ES}} - \eta\|^2 \leq \|f_\theta - \eta\|^2 + \psi_{n,M}(\theta) + \frac{8\sigma^2}{n} \log 2,$$

where $\psi_{n,M}(0) := 0$, and, for $\theta \neq 0$,

$$(3.13) \quad \psi_{n,M}(\theta) := \frac{\sigma^2 R}{n} \wedge \frac{9\sigma^2 M(\theta)}{n} \log \left(1 + \frac{eM}{M(\theta) \vee 1} \right) \wedge \frac{11\sigma |\theta|_1}{\sqrt{n}} \sqrt{\log \left(1 + \frac{3eM\sigma}{|\theta|_1 \sqrt{n}} \right)} \wedge 4|\theta|_1^2.$$

In particular, if there exists $\theta^* \in \mathbb{R}^M$ such that $\eta = f_{\theta^*}$, we have

$$(3.14) \quad \mathbb{E} \|f_{\hat{\theta}^{ES}} - f_{\theta^*}\|^2 \leq \psi_{n,M}(\theta^*) + \frac{8\sigma^2}{n} \log 2.$$

The proof of Theorem 3.2 is obtained by combining Theorem 3.1 and Lemma A.2 in the Appendix. For brevity, the constants derived from Lemma A.2 are rounded up to the closest integer.

It is easy to see that in fact Lemma A.2 implies a more general result. Not necessarily $f_{\hat{\theta}^{ES}}$, but, in general, any estimator satisfying a SOI of the type (3.6) also obeys the oracle inequality of the form (3.10), that is, enjoys adaptation simultaneously in terms of the ℓ_0 and ℓ_1 norms. This remains still a theoretical proposal, since we are not aware of estimators satisfying (3.6) apart from $f_{\hat{\theta}^{ES}}$. However, there are estimators for which coarser versions of (3.6) are available as discussed in the next section.

4. Sparsity oracle inequalities for the BIC and Lasso estimators. The aim of this section is to put Theorem 3.2 in perspective by discussing weaker results in the same spirit for two popular estimators, namely, the BIC and the Lasso estimators.

We consider the following version of the BIC estimator [cf. Bunea, Tsybakov and Wegkamp (2007b)]:

$$(4.1) \quad \hat{\theta}^{\text{BIC}} \in \arg \min_{\theta \in \mathbb{R}^M} \left\{ \frac{1}{n} \|\mathbf{Y} - \mathbf{X}\theta\|_2^2 + \text{pen}(\theta) \right\},$$

where

$$\text{pen}(\theta) := \frac{2\sigma^2}{n} \left\{ 1 + \frac{2+a}{1+a} \sqrt{L(\theta)} + \frac{1+a}{a} L(\theta) \right\} M(\theta)$$

with for some $a > 0$ and $L(\theta) = 2 \log(\frac{eM}{M(\theta)\sqrt{1}})$. Combining Theorem 3.1 in Bunea, Tsybakov and Wegkamp (2007b) and Lemma A.2 in the Appendix we get the following corollary.

COROLLARY 4.1. *Assume that $\max_{1 \leq j \leq M} \|f_j\| \leq 1$. Then there exists a positive numerical constant C such that for any $M \geq 2, n \geq 1$ and any $a > 0$ the BIC estimator satisfies*

$$(4.2) \quad \mathbb{E} \|\mathbf{f}_{\hat{\theta}^{\text{BIC}}} - \eta\|^2 \leq (1+a) \min_{\theta \in \mathbb{R}^M} \left\{ \|\mathbf{f}_\theta - \eta\|^2 + C \frac{1+a}{a} \varphi_{n,M}(\theta) \right\} + \frac{C\sigma^2}{n},$$

where $\varphi_{n,M}$ is defined in (3.11).

We note that Theorem 3.1 in Bunea, Tsybakov and Wegkamp (2007b) is stated with $R = M$ and with the additional assumption that all the functions f_j are uniformly bounded. Nevertheless, this last condition is not used in the proof in Bunea, Tsybakov and Wegkamp (2007b), and the result trivially extends to the framework that we consider here. The SOI (4.2) ensures adaptation to sparsity simultaneously in terms of the ℓ_0 and ℓ_1 norms. However, it is less precise than the SOI in Theorem 3.2 because the leading constant $(1+a)$ is strictly greater than 1 and the rate deteriorates as the leading constant approaches 1, that is, as $a \rightarrow 0$. Also the computation of the BIC estimator is a hard combinatorial problem, exponential in M , and it can be efficiently solved only when the dimension M is small.

Consider now the Lasso estimator $\hat{\theta}^{\text{L}}$, that is, a solution of the minimization problem

$$(4.3) \quad \hat{\theta}^{\text{L}} \in \arg \min_{\theta \in \mathbb{R}^M} \left\{ \frac{1}{n} \|\mathbf{Y} - \mathbf{X}\theta\|_2^2 + \lambda |\theta|_1 \right\},$$

where $\lambda > 0$ is a tuning parameter. This problem is convex, and there exist several efficient algorithms of computing $\hat{\theta}^{\text{L}}$ in polynomial time.

Our aim here is to present results in the spirit of Theorem 3.2 for the Lasso. They have a weaker form than for the ES estimator and for the BIC. In the next theorem, we give a SOI in terms of the ℓ_1 norm that is similar to those that we have presented for the ES and BIC estimators, but it is stated in probability rather than in expectation and the logarithmic factor in the rate is less accurate. Note that it does not require any restrictive condition on the dictionary f_1, \dots, f_M .

THEOREM 4.1. *Assume that $\max_{1 \leq j \leq M} \|f_j\| \leq 1$. Let $M \geq 2, n \geq 1$ and let $\hat{\theta}^L$ be the Lasso estimator defined by (4.3) with $\lambda = A\sigma\sqrt{\frac{\log M}{n}}$, where $A > 2\sqrt{2}$. Then with probability at least $1 - M^{1-A^2/8}$ we have*

$$(4.4) \quad \|\mathbf{f}_{\hat{\theta}^L} - \eta\|^2 \leq \min_{\theta \in \mathbb{R}^M} \left\{ \|\mathbf{f}_\theta - \eta\|^2 + 2A\sigma \frac{|\theta|_1}{\sqrt{n}} \sqrt{\log M} \right\}.$$

PROOF. From the definition of $\hat{\theta}^L$ by a simple algebra we get

$$\|\mathbf{f}_{\hat{\theta}^L} - \eta\|^2 \leq \|\mathbf{f}_\theta - \eta\|^2 + \frac{2}{n} |(\hat{\theta}^L - \theta)^\top \mathbf{X}^\top \xi| + \lambda(|\theta|_1 - |\hat{\theta}^L|_1) \quad \forall \theta \in \mathbb{R}^M.$$

Next, note that $P(\mathcal{A}) \geq 1 - M^{1-A^2/8}$ for the random event $\mathcal{A} = \{|\frac{2}{n} \mathbf{X}^\top \xi|_\infty \leq \lambda\}$ [cf. [Bickel, Ritov and Tsybakov \(2009\)](#), equation (B.4)]. Therefore,

$$\|\mathbf{f}_{\hat{\theta}^L} - \eta\|^2 \leq \|\mathbf{f}_\theta - \eta\|^2 + \lambda|\hat{\theta}^L - \theta|_1 + \lambda(|\theta|_1 - |\hat{\theta}^L|_1) \quad \forall \theta \in \mathbb{R}^M$$

with probability at least $1 - M^{1-A^2/8}$. Thus, (4.4) follows by the triangle inequality and the definition of λ . \square

The rate $\frac{|\theta|_1}{\sqrt{n}} \sqrt{\log M}$ in (4.4) is slightly worse than the corresponding ℓ_1 term of the rate of ES estimator [cf. (3.11) and (3.13)].

In contrast to Theorem 4.1, a SOI in terms of the ℓ_0 norm for the Lasso is available only under strong conditions on the dictionary f_1, \dots, f_M . Following [Bickel, Ritov and Tsybakov \(2009\)](#), we say that the restricted eigenvalue condition $\text{RE}(s, c_0)$ is satisfied for some integer s such that $1 \leq s \leq M$, and a positive number c_0 if we have

$$\kappa(s, c_0) := \min_{\substack{J_0 \subseteq \{1, \dots, M\}, \\ |J_0| \leq s}} \min_{\substack{\Delta \neq 0, \\ |\Delta_{J_0^c}|_1 \leq c_0 |\Delta_{J_0}|_1}} \frac{|\mathbf{X}\Delta|_2}{\sqrt{n} |\Delta_{J_0}|_2} > 0.$$

Here $|J|$ is the cardinality of the index set J and we denote by Δ_J the vector in \mathbb{R}^M that has the same coordinates as Δ on J and zero coordinates on the complement J^c of J . A typical SOI in terms of the ℓ_0 norm for the Lasso is given in Theorem 6.1 of [Bickel, Ritov and Tsybakov \(2009\)](#). It guarantees that, under the

condition $\text{RE}(s, 3 + 4/a)$ and the assumptions of Theorem 4.1, with probability at least $1 - M^{1-A^2/8}$, we have

$$(4.5) \quad \begin{aligned} & \|f_{\hat{\theta}_L} - \eta\|^2 \leq (1 + a) \\ & \times \min_{\theta \in \mathbb{R}^M : M(\theta) \leq s} \left\{ \|f_\theta - \eta\|^2 \right. \\ & \left. + \frac{C(1 + a)}{a\kappa^2(s, 3 + 4/a)} \frac{M(\theta) \log M}{n} \right\} \end{aligned}$$

for all $a > 0$ and some constant $C > 0$ depending only on A and σ . This oracle inequality is substantially weaker than (3.10) and (4.2). Indeed, it is valid under assumption $\text{RE}(s, 3 + 4/a)$, which is a strong condition. Furthermore, the rank of the matrix \mathbf{X} does not appear, the minimum in (4.5) is taken over the set of sparsity s linked to the properties of the matrix \mathbf{X} , and the minimal restricted eigenvalue $\kappa(s, 3 + 4/a)$ appears in the denominator. This contrasts with inequalities (3.10), (4.2) and (4.4) which hold under no assumption on \mathbf{X} , except for simple normalization, $\max_{1 \leq j \leq M} \|f_j\| \leq 1$. Finally, the leading constant in (4.5) is strictly larger than 1, and the same comments as for the BIC apply in this respect.

5. Discussion of the optimality.

5.1. *Deficiency of the approach based on function normalization.* Section 3 provides upper bounds on the risk of ES estimator. A natural question is whether these bounds are optimal. At first sight, to show the optimality it seems sufficient to prove that there exists $\theta \in \mathbb{R}^M$ and η such that, for any estimator T ,

$$\mathbb{E} \|T - \eta\|^2 \geq \|f_\theta - \eta\|^2 + c\psi_{n,M}(\theta),$$

where $c > 0$ is some constant independent of n and M . This can be also written in the form

$$(5.1) \quad \inf_T \sup_{\eta} \sup_{\theta \in \mathbb{R}^M} \frac{\mathbb{E} \|T - \eta\|^2 - \|f_\theta - \eta\|^2}{\psi_{n,M}(\theta)} \geq c,$$

where \inf_T denotes the infimum over all estimators. We note that it is possible to prove (5.1) under some assumptions on the dictionary f_1, \dots, f_M . However, we do not consider these type of results because they do not lead to a valid notion of optimality. Indeed, since the rate $\psi_{n,M}(\theta)$ is a function of parameter θ , there exist infinitely many different rate functions $\psi_{n,M}(\cdot)$ for which (5.1) can be proved and complemented by the corresponding upper bounds. To illustrate this point, consider a basic example defined by the following conditions:

- (i) $M = n$;
- (ii) $\eta = f_{\theta^*}$ for some $\theta^* \in \mathbb{R}^n$;
- (iii) the Gram matrix $\Psi = \mathbf{X}^\top \mathbf{X}/n$ is equal to the $n \times n$ identity matrix;

(iv) $\sigma^2 = 1$.

This will be further referred to as the *diagonal model*. It can be equivalently written as a Gaussian sequence model,

$$(5.2) \quad y_i = \theta_i + \frac{1}{\sqrt{n}}\varepsilon_i, \quad i = 1, \dots, n,$$

where $(y_1, \dots, y_n)^\top = \mathbf{X}^\top \mathbf{Y}/n$ and $\varepsilon_1, \dots, \varepsilon_n$ are i.i.d. standard Gaussian random variables.

Clearly, estimation of η in the diagonal model is equivalent to estimation of θ^* in model (5.2), and we have the isometry $\|f_\theta - \eta\| = |\theta - \theta^*|_2$. Moreover, it is easy to see that we can consider w.l.o.g. only estimators T of the form $T = f_{\hat{\theta}}$ for some statistic $\hat{\theta}$, and that (5.1) for the diagonal model follows from a simplified bound

$$(5.3) \quad \inf_{\hat{\theta}} \sup_{\theta \in \mathbb{R}^n} \frac{E_\theta |\hat{\theta} - \theta|_2^2}{\psi_n(\theta)} \geq c,$$

where we write E_θ to specify the dependence of the expectation upon θ , $\inf_{\hat{\theta}}$ denotes the infimum over all estimators, and for brevity $\psi_n(\theta) = \psi_{n,n}(\theta)$.

Results such as (5.3) are available in Donoho and Johnstone (1994a) where it is proved that, for the diagonal model,

$$(5.4) \quad \inf_{\hat{\theta}} \sup_{\theta \in \mathbb{R}^n} \frac{E_\theta |\hat{\theta} - \theta|_2^2}{\psi_n^{01}(\theta)} = 1 + o(1)$$

as $n \rightarrow \infty$, where

$$(5.5) \quad \psi_n^{01}(\theta) = 2 \log n \left\{ \frac{1}{n} + \sum_{i=1}^n \min\left(\theta_i^2, \frac{1}{n}\right) \right\}.$$

The expression in curly brackets in (5.5) is the risk of 0–1 (or “keep-or-kill”) oracle, that is, the minimal risk of the estimators $\hat{\theta}$ whose components $\hat{\theta}_j$ are either equal to y_j or to 0. A relation similar to (5.4), with the infimum taken over a class of thresholding rules, is proved in Foster and George (1994).

Result (5.4) is often wrongly interpreted as the fact that the factor $2 \log n$ is the “unavoidable” price to pay for sparse estimation. In reality this is not true, and (5.4) cannot be considered as a basis of valid notion of optimality. Indeed, using the results of Section 3, we are going to construct an estimator whose risk is $O(\psi_n^{01}(\theta))$ for all θ , and is of order $o(\psi_n^{01}(\theta))$ for some θ (cf. Theorem 5.2 below). So, this estimator improves upon (5.4) not only in constants but in the rate; in particular, the exact asymptotic constant appearing in (5.4) is of no importance. The reason is that the lower bound for (5.4) in Donoho and Johnstone (1994a) is proved by restricting θ to a small subset of \mathbb{R}^n , and the behavior of the risk on other subsets of \mathbb{R}^n can be much better.

Define the rate

$$\psi_n^*(\theta) = \min \left[\frac{M(\theta) \log n}{n}, |\theta|_1 \sqrt{\frac{\log n}{n}}, |\theta|_1^2 \right] + \frac{1}{n},$$

which is an asymptotic upper bound on the rate in (3.13) for $M = n, n \rightarrow \infty$.

THEOREM 5.1. *Consider the diagonal model. Then the Exponential Screening estimator satisfies*

$$(5.6) \quad \limsup_{n \rightarrow \infty} \sup_{\theta \in \mathbb{R}^n} \frac{E_\theta |\tilde{\theta}^{\text{ES}} - \theta|_2^2}{\psi_n^{01}(\theta)} \leq 2$$

and

$$(5.7) \quad \liminf_{n \rightarrow \infty} \inf_{\theta \in \mathbb{R}^n} \frac{E_\theta |\tilde{\theta}^{\text{ES}} - \theta|_2^2}{\psi_n^{01}(\theta)} = 0.$$

Furthermore,

$$(5.8) \quad \lim_{n \rightarrow \infty} \inf_{\theta \in \mathbb{R}^n} \frac{\psi_n^*(\theta)}{\psi_n^{01}(\theta)} = 0.$$

PROOF. We first prove (5.6). From (3.3), Lemma 3.1 and (3.5) we obtain

$$E_\theta |\tilde{\theta}^{\text{ES}} - \theta^*|_2^2 \leq \min_{\theta \in \mathbb{R}^n} \left\{ |\theta - \theta^*|_2^2 + \frac{M(\theta)}{n} (1 + 4 \log(2en)) \right\} + \frac{4 \log 2}{n}$$

for any $\theta^* \in \mathbb{R}^n$. Let $\bar{\theta} \in \mathbb{R}^n$ be the vector with components $\bar{\theta}_j = \theta_j^* I(|\theta_j^*| > 1/\sqrt{n})$ where $I(\cdot)$ denotes the indicator function. Then

$$|\bar{\theta} - \theta^*|_2^2 = \sum_{j=1}^n |\theta_j^*|^2 I(|\theta_j^*| \leq 1/\sqrt{n})$$

and

$$\frac{M(\bar{\theta})}{n} = \sum_{j=1}^n \frac{1}{n} I(|\theta_j^*| > 1/\sqrt{n}).$$

Therefore,

$$E_\theta |\tilde{\theta}^{\text{ES}} - \theta^*|_2^2 \leq (1 + 4 \log(2en)) \sum_{j=1}^n \min \left(|\theta_j^*|^2, \frac{1}{n} \right) + \frac{4 \log 2}{n},$$

which implies (5.6). Next, (5.7) is an immediate consequence of (5.8). To prove (5.8) we consider, for example, the set $\Theta_n = \{\theta \in \mathbb{R}^n : a/\sqrt{n} \leq |\theta_j| \leq b/\sqrt{n} \text{ for all } \theta_j \neq 0\}$ where $0 < a < b < \infty$ are constants. For all $\theta \in \Theta_n$ we have

$$\psi_n^*(\theta) \leq |\theta|_1 \sqrt{\frac{\log n}{n}} + \frac{1}{n} \leq bM(\theta) \frac{\sqrt{\log n}}{n} + \frac{1}{n}$$

and

$$\psi_n^{01}(\theta) \geq 2(\min(a^2, 1)M(\theta) + 1)\frac{\log n}{n},$$

so that

$$(5.9) \quad \lim_{n \rightarrow \infty} \sup_{\theta \in \Theta_n} \frac{\psi_n^*(\theta)}{\psi_n^{01}(\theta)} = 0.$$

Hence, (5.8) follows. \square

Theorem 5.1 shows that the normalizing function (rate) $\psi_n^{01}(\theta)$ and the result (5.4) cannot be considered as a benchmark. Indeed, the risk of the ES estimator is strictly below this bound. It attains the rate $\psi_n^{01}(\theta)$ everywhere on \mathbb{R}^n [cf. (5.6)] and has strictly better rate on some subsets of \mathbb{R}^n [cf. (5.8), (5.9)]. In particular, the ES estimator improves upon the soft thresholding estimator, which is known to asymptotically attain the bound (5.4) [cf. Donoho and Johnstone (1994a)]. This is a kind of inadmissibility statement for the rate $\psi_n^{01}(\theta)$.

Observe also that the improvement that we obtain is not a “marginal” effect regarding signals θ with small intensity. Indeed, (5.9) is stronger than (5.8), and the set Θ_n is rather massive. In particular, the ℓ_0 norm $M(\theta)$ in the definition of Θ_n can be arbitrary, so that Θ_n contains elements θ with the whole spectrum of ℓ_1 norms, from small $|\theta|_1 = a/\sqrt{n}$ to very large $|\theta|_1 = bM/\sqrt{n} = b\sqrt{n}$. Various other examples of Θ_n satisfying (5.9) can be readily constructed.

So far, we were interested only in the rates. The fact that the constant in (5.6) is equal to 2 was of no importance in this argument since on some subsets of \mathbb{R}^n we can improve the rate. Notice that one can construct estimators having the same properties as those proved for $\tilde{\theta}^{\text{ES}}$ in Theorem 5.1 with constant 1 instead of 2 in (5.6). In other words, one can construct an estimator $\tilde{\theta}^*$ whose risk is at least as small as $\psi_n^{01}(\theta)(1 + o(1))$ everywhere on \mathbb{R}^n and attains strictly faster rate $o(\psi_n^{01}(\theta))$ on some subsets of \mathbb{R}^n . Such an estimator $\tilde{\theta}^*$ can be obtained by aggregating $\tilde{\theta}^{\text{ES}}$ with the soft thresholding estimator, as shown in the next theorem.

THEOREM 5.2. *Consider the diagonal model. Then there exists a randomized estimator $\tilde{\theta}^*$ such that*

$$(5.10) \quad \limsup_{n \rightarrow \infty} \sup_{\theta \in \mathbb{R}^n} \frac{E_\theta |\tilde{\theta}^* - \theta|_2^2}{\tilde{\psi}_n(\theta)} \leq 1,$$

where the expectation includes that over the randomizing distribution, and where the normalizing functions $\tilde{\psi}_n$ satisfy

$$(5.11) \quad \liminf_{n \rightarrow \infty} \inf_{\theta \in \mathbb{R}^n} \frac{\psi_n^{01}(\theta)}{\tilde{\psi}_n(\theta)} \geq 1$$

and

$$(5.12) \quad \liminf_{n \rightarrow \infty} \sup_{\theta \in \mathbb{R}^n} \frac{\psi_n^{01}(\theta)}{\tilde{\psi}_n(\theta)} = \infty.$$

The proof of this theorem is given in the [Appendix](#).

5.2. *Minimax optimality on the intersection of ℓ_0 and ℓ_1 balls.* The rate in the upper bound of Theorem 3.2 is the minimum of terms depending on the ℓ_0 norm $M(\theta)$ and on the ℓ_1 norm $|\theta|_1$ [cf. (3.13)]. We would like to derive a corresponding lower bound, that is, to show that this rate of convergence cannot be improved in a minimax sense. Since both ℓ_0 and ℓ_1 norms are present in the upper bound, a natural approach is to consider minimax lower bounds on the intersection of ℓ_0 and ℓ_1 balls. Here we prove such a lower bound under some assumptions on the dictionary $\mathcal{H} = \{f_1, \dots, f_M\}$ or, equivalently, on the matrix \mathbf{X} . Along with the lower bound for one “worst case” dictionary \mathcal{H} , we also state it uniformly for all dictionaries in a certain class.

5.2.1. *Assumptions on the dictionary.* Recall first that all the results from Section 3 hold under the only condition that the dictionary \mathcal{H} is composed of functions f_j such that $\|f_j\| \leq 1$. This condition is very mild compared to the assumptions that typically appear in the literature on sparse recovery using ℓ_1 penalization such as the Lasso or the Dantzig selector. Bühlmann and van de Geer (2009) review a long list of such assumptions, including the restricted isometry (RI) property given, for example, in Candes (2008) and the restricted eigenvalue (RE) condition of Bickel, Ritov and Tsybakov (2009) described in Section 4. We call them for brevity the L -conditions. Loosely speaking, they ensure that for some integer $S \leq M$, the design matrix \mathbf{X} forms a quasi-isometry from a suitable subset \mathcal{A}_p of \mathbb{R}^p into \mathbb{R}^n for any p such that $|p| \leq S$. Here “quasi-isometry” means that there exist two positive constants $\underline{\kappa}$ and $\bar{\kappa}$ such that

$$(5.13) \quad \underline{\kappa}|\theta|_2^2 \leq \frac{|\mathbf{X}\theta|_2^2}{n} \leq \bar{\kappa}|\theta|_2^2 \quad \forall \theta \in \mathcal{A}_p.$$

While the general thinking is that a design matrix \mathbf{X} satisfying an L -condition is favorable, we establish below that, somewhat surprisingly, such matrices correspond to the least favorable case.

We now formulate a weak version of the RI condition. For any integer $M \geq 2$ and any $0 < u \leq M$ let \mathcal{P}_u denote the set of vectors $\theta \in \{-1, 0, 1\}^M$ such that $M(\theta) \leq u$. For any constants $0 < \kappa \leq 1$ and $0 < t \leq (M \wedge n)/2$ let $\mathcal{D}(t, \kappa)$ be the class of design matrices \mathbf{X} defined by the following conditions:

- (i) $\max_{1 \leq j \leq M} \|f_j\| \leq 1$;

(ii) there exist $\underline{\kappa}, \bar{\kappa} > 0$, such that $\underline{\kappa}/\bar{\kappa} \geq \kappa$ and

$$(5.14) \quad \underline{\kappa}|\theta|_2^2 \leq \frac{|\mathbf{X}\theta|_2^2}{n} \leq \bar{\kappa}|\theta|_2^2 \quad \forall \theta \in \mathcal{P}_{2t}.$$

Note that $t \leq t'$ implies $\mathcal{D}(t', \kappa) \leq \mathcal{D}(t, \kappa)$. Examples of matrices \mathbf{X} that satisfy (5.14) are given in the next subsection.

In the next subsection we show that the upper bound of Theorem 3.2 matches a minimax lower bound which holds uniformly over the class of design matrices $\mathcal{D}(S, \kappa)$.

5.2.2. Minimax lower bound. Denote by P_η the distribution of (Y_1, \dots, Y_n) where $Y_i = \eta(x_i) + \xi_i, i = 1, \dots, n$, and by E_η the corresponding expectation. For any $\delta > 0$ and any integers $S \geq 1, n \geq 1, M \geq 1, R \geq 1$ such that $R \leq M \wedge n$, define the quantity

$$(5.15) \quad \begin{aligned} \zeta_{n,M,R}(S, \delta) := & \frac{\sigma^2 R}{n} \wedge \frac{\sigma^2 S}{n} \log\left(1 + \frac{eM}{S}\right) \\ & \wedge \frac{\sigma \delta}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eM\sigma}{\delta\sqrt{n}}\right)} \wedge \delta^2. \end{aligned}$$

Note that $\zeta_{n,M,R}(S, \delta) = \psi_{n,M}(\theta)$ where $\psi_{n,M}$ is the function (3.13) with $M(\theta) = S$ and $|\theta|_1 = \delta$. Let $m \geq 1$ be the largest integer satisfying

$$(5.16) \quad m \leq \frac{\delta\sqrt{n}}{\sigma\sqrt{\log(1 + eM/m)}},$$

if such an integer exists. If there is no $m \geq 1$ such that (5.16) holds, we set $m = 0$. Note that $m \leq \delta\sqrt{n}/\sigma$.

THEOREM 5.3. *Fix $\delta > 0$ and integers $n \geq 1, M \geq 2, 1 \leq S \leq M$. Fix $\kappa > 0$ and let \mathcal{H} be any dictionary with design matrix $\mathbf{X} \in \mathcal{D}(S \wedge \bar{m}, \kappa)$, where $\bar{m} = m \vee 1$ and m is defined in (5.16). Then, for any estimator T_n , possibly depending on δ, S, n, M and \mathcal{H} , there exists a numerical constant $c^* > 0$, such that*

$$(5.17) \quad \sup_{\substack{\theta \in \mathbb{R}_+^M \setminus \{0\} \\ M(\theta) \leq S \\ |\theta|_1 \leq \delta}} \sup_{\eta} \{E_\eta \|T_n - \eta\|^2 - \|\mathbf{f}_\theta - \eta\|^2\} \geq c^* \kappa \zeta_{n,M,\text{rk}(\mathbf{X})}(S, \delta),$$

where $\text{rk}(\mathbf{X})$ denotes the rank of \mathbf{X} and \mathbb{R}_+^M is the positive cone of \mathbb{R}^M . Moreover,

$$(5.18) \quad \sup_{\substack{\theta \in \mathbb{R}_+^M \setminus \{0\} \\ M(\theta) \leq S \\ |\theta|_1 \leq \delta}} E_{\mathbf{f}_\theta} \|T_n - \mathbf{f}_\theta\|^2 \geq c^* \kappa \zeta_{n,M,\text{rk}(\mathbf{X})}(S, \delta).$$

The proof of this theorem is given in Section A.3 of the Appendix. It is worth mentioning that the result of Theorem 5.3 is stronger than the minimax lower bounds discussed in Section 5.1 [cf. (5.3)] in the sense that even if $\eta = f_{\theta^*}$, $\theta^* \in \mathbb{R}^M$, where $M(\theta^*)$ and $|\theta^*|_1$ are known a priori, the rate cannot be improved.

Let us fix some positive constant $\kappa \leq 1$ and define $\tilde{R} = 1 + \lceil \frac{R}{C_0} \log(1 + \frac{eM}{R}) \rceil$ for some constant $C_0 > 0$ to be chosen small enough. We now show that for each choice of $R \geq 1$ such that $\tilde{R} \leq M \wedge n$, there exists at least one matrix $\mathbf{X} \in \mathcal{D}(R/2, \kappa)$ such that $R \leq \text{rk}(\mathbf{X}) \leq \tilde{R}$. A basic example is the following. Take the elements $\mathbf{X}_{i,j} = f_j(x_i)$, $i = 1, \dots, n$, $j = 1, \dots, M$, of matrix \mathbf{X} as

$$(5.19) \quad \mathbf{X}_{i,j} = \begin{cases} \epsilon_{i,j} \sqrt{\frac{n}{\tilde{R}}}, & \text{if } i \leq \tilde{R}, \\ 0, & \text{otherwise,} \end{cases}$$

where $\epsilon_{i,j}$, $1 \leq i \leq n$, $1 \leq j \leq M$ are i.i.d. Rademacher random variables, that is, random variables taking values 1 and -1 with probability $1/2$. First, it is clear that then $\|f_j\| \leq 1$, $j = 1, \dots, M$. Next, condition (ii) in the definition of $\mathcal{D}(R/2, \kappa)$ follows from the results on RI properties of Rademacher matrices. Many such results have been derived and we focus only on that of Baraniuk et al. (2008) because of its simplicity. Indeed, Theorem 5.2 in Baraniuk et al. (2008) ensures not only that for an integer $S' \leq M \wedge n$ there exist design matrices in $\mathcal{D}(S'/2, \kappa)$ but also that most of the design matrices \mathbf{X} with i.i.d. Rademacher entries $\epsilon_{i,j}$ are in $\mathcal{D}(S'/2, \kappa)$ for some $\kappa > 0$ as long as there exists a constant C_0 small enough such that the condition

$$(5.20) \quad \frac{S'}{M \wedge n} \log\left(1 + \frac{eM}{S'}\right) < C_0$$

is satisfied. Specifically, Theorem 5.2 in Baraniuk et al. (2008) ensures that if \mathbf{X}' is the $\tilde{R} \times M$ matrix composed of the first \tilde{R} rows of \mathbf{X} with elements as defined in (5.19), and

$$(5.21) \quad \frac{R}{\tilde{R}} \log\left(1 + \frac{eM}{R}\right) \leq C_0$$

holds for small enough C_0 , then

$$\kappa \frac{n}{\tilde{R}} |\theta|_2^2 \leq \frac{|\mathbf{X}'\theta|_2^2}{\tilde{R}} \leq \bar{\kappa} \frac{n}{\tilde{R}} |\theta|_2^2 \quad \forall \theta : M(\theta) \leq R$$

with probability close to 1 which in turn implies (ii) with $t = R/2$. As a result, the above construction yields $\mathbf{X} \in \mathcal{D}(R/2, \kappa)$ that has rank bracketed by R and \tilde{R} since (5.21) holds by our definition of \tilde{R} .

In what follows, C_0 is the constant in (5.20) small enough to ensure that Theorem 5.2 in Baraniuk et al. (2008) holds, and we assume w.l.o.g. that $C_0 < 1$.

Using the above remarks and Theorem 5.3 we obtain the following result.

THEOREM 5.4. Fix $\delta > 0$ and integers $n \geq 1, M \geq 2, 1 \leq S \leq M, R \geq 1$. Moreover, assume that $1 + \frac{R}{C_0} \log(1 + eM/R) \leq M \wedge n$. Then there exists a dictionary \mathcal{H} composed of functions f_j with $\max_{1 \leq j \leq M} \|f_j\| \leq 1, R \leq \text{rk}(\mathbf{X}) \leq 1 + \frac{R}{C_0} \log(1 + eM/R)$, and a constant $c_* > 0$ such that

$$(5.22) \quad \inf_{T_n} \sup_{\substack{\theta \in \mathbb{R}_+^M \setminus \{0\} \\ M(\theta) \leq S \\ |\theta|_1 \leq \delta}} \sup_{\eta} \{E_{\eta} \|T_n - \eta\|^2 - \|f_{\theta} - \eta\|^2\} \geq c_* \zeta_{n, M, \text{rk}(\mathbf{X})}(S, \delta),$$

where the infimum is taken over all estimators. Moreover,

$$(5.23) \quad \inf_{T_n} \sup_{\substack{\theta \in \mathbb{R}_+^M \setminus \{0\} \\ M(\theta) \leq S \\ |\theta|_1 \leq \delta}} E_{f_{\theta}} \|T_n - f_{\theta}\|^2 \geq c_* \zeta_{n, M, \text{rk}(\mathbf{X})}(S, \delta).$$

PROOF. Fix a positive constant $\kappa \leq 1$ and let \mathbf{X} be a random matrix constructed as in (5.19) so that the rank of \mathbf{X} is bracketed by R and \tilde{R} and $\mathbf{X} \in \mathcal{D}(R/2, \kappa)$. We consider two cases. Assume first that $S \leq R/2$ so that $\mathbf{X} \in \mathcal{D}(R/2, \kappa) \subset \mathcal{D}(S, \kappa) \subseteq \mathcal{D}(S \wedge \bar{m}, \kappa)$ and the result follows trivially from Theorem 5.3. Next, if $S \geq R/2$, observe that

$$\text{rk}(\mathbf{X}) \leq \tilde{R} \leq 1 + \frac{R}{C_0} \log\left(1 + \frac{eM}{R}\right) \leq \frac{2}{C_0} R \log\left(1 + \frac{2eM}{R}\right)$$

(we used here that $C_0 < 1$), so that

$$\text{rk}(\mathbf{X}) \wedge S \log\left(1 + \frac{eM}{S}\right) \leq \text{rk}(\mathbf{X}) \leq \frac{2}{C_0} \left(\text{rk}(\mathbf{X}) \wedge R \log\left(1 + \frac{2eM}{R}\right)\right).$$

It yields $\zeta_{n, M, \text{rk}(\mathbf{X})}(S, \delta) \leq C \zeta_{n, M, \text{rk}(\mathbf{X})}(R/2, \delta)$ and the result follows from Theorem 5.3, which ensures that

$$\begin{aligned} & \inf_{T_n} \sup_{\substack{\theta \in \mathbb{R}_+^M \setminus \{0\} \\ M(\theta) \leq S \\ |\theta|_1 \leq \delta}} \sup_{\eta} \{E_{\eta} \|T_n - \eta\|^2 - \|f_{\theta} - \eta\|^2\} \\ & \geq c_* \kappa \zeta_{n, M, \text{rk}(\mathbf{X})}(R/2, \delta) \geq c_* \zeta_{n, M, \text{rk}(\mathbf{X})}(S, \delta). \end{aligned} \quad \square$$

As a consequence of Theorem 5.4 we get a lower bound on the ℓ_0 ball $B_0(S) = \{\theta : M(\theta) \leq S\}$ by formally setting $\delta = \infty$ in (5.22)

$$(5.24) \quad \begin{aligned} & \sup_{\eta} \sup_{\substack{\theta \in \mathbb{R}^M \\ M(\theta) \leq S}} \{E_{\eta} \|T_n - \eta\|^2 - \|f_{\theta} - \eta\|^2\} \\ & \geq c_* \frac{\sigma^2}{n} \left[\text{rk}(\mathbf{X}) \wedge S \log\left(1 + \frac{eM}{S}\right) \right] \end{aligned}$$

and the same type of bound can be derived from (5.23). Analogous considerations lead to the following lower bound on the ℓ_1 ball $B_1(\delta) = \{\theta : |\theta|_1 \leq \delta\}$ when setting $S = M$:

$$(5.25) \quad \begin{aligned} & \sup_{\eta} \sup_{\substack{\theta \in \mathbb{R}^M \\ |\theta|_1 \leq \delta}} \{E_{\eta} \|T_n - \eta\|^2 - \|\mathbf{f}_{\theta} - \eta\|^2\} \\ & \geq c_* \left(\frac{\sigma^2 \text{rk}(\mathbf{X})}{n} \wedge \frac{\sigma \delta}{\sqrt{n}} \sqrt{\log \left(1 + \frac{eM\sigma}{\delta \sqrt{n}} \right)} \wedge \delta^2 \right), \end{aligned}$$

and the same type of bound can be derived from (5.23).

Consider now the linear regression, that is, assume that there exists θ^* such that $\eta = \mathbf{f}_{\theta^*}$. Comparing (3.14) with (5.18) we find that for $\delta \geq 1/\sqrt{n}$ the rate $\zeta_{n,M,\text{rk}(\mathbf{X})}(S, \delta)$ is the minimax rate of convergence on $B_0(S) \cap B_1(\delta)$ and that the ES estimator is rate optimal. Moreover, it is rate optimal separately on $B_0(S)$ and $B_1(\delta)$, and the minimax rates on these sets are given by the right-hand sides of (5.24) and (5.25), respectively.

For the *diagonal model* (cf. Section 5.1), asymptotic lower bounds and exact asymptotics of the minimax risk on ℓ_q balls were studied by Donoho et al. (1992) for $q = 0$ and by Donoho and Johnstone (1994b) for $0 < q < \infty$. These results were further refined by Abramovich et al. (2006). In the ℓ_0 case, Donoho et al. (1992) exhibit a minimax rate over $B_0(S)$ that is asymptotically equivalent to

$$2\sigma^2 \frac{S}{n} \log \left(\frac{n}{S} \right) \quad \text{as } M = n \rightarrow \infty.$$

In the ℓ_1 case, Donoho and Johnstone (1994b) prove that the minimax rate over an ℓ_1 ball with radius δ is asymptotically equivalent to

$$\frac{\delta \sigma}{\sqrt{n}} \sqrt{2 \log \left(\frac{\sigma \sqrt{n}}{\delta} \right)} \quad \text{as } M = n \rightarrow \infty.$$

In both cases, the above rates are equivalent, up to a numerical constant, to the asymptotics of the right-hand sides of (5.24) and (5.25) under the diagonal model. We note that the results of those papers are valid under some restrictions on asymptotical behavior of S (resp., δ) as a function of n .

Recently Raskutti, Wainwright and Yu (2009) extended the study of asymptotic lower bounds on ℓ_q balls ($0 \leq q \leq 1$) to the nondiagonal case with $M \neq n$. Their results hold under some restrictions on the joint asymptotic behavior of n , M and S (resp., δ). The minimax rates on the ℓ_0 and ℓ_1 balls obtained in Raskutti, Wainwright and Yu [(2009), Theorem 3] are similar to (5.24) and (5.25) but, because of the specific asymptotics, some effects are wiped out there. For example, the ℓ_1 rate in Raskutti, Wainwright and Yu (2009) is $\delta \sqrt{(\log M)/n}$, whereas (5.25) reveals an elbow effect that translates into different rates for $\sigma \text{rk}(\mathbf{X}) \leq \delta \sqrt{n}$. Furthermore, the dependence on the rank of \mathbf{X} does not appear in Raskutti, Wainwright and

Yu (2009), since under their assumptions $\text{rk}(\mathbf{X}) = n$. Theorem 5.3 above gives a stronger result since it is (i) nonasymptotic, (ii) it explicitly depends on the rank $\text{rk}(\mathbf{X})$ of the design matrix and (iii) it holds on the intersection of the ℓ_0 and ℓ_1 balls. Moreover, Theorem 3.2 shows that the $\ell_0 - \ell_1$ lower bound is attained by one single estimator: the Exponential Screening estimator. Alternatively, Raskutti, Wainwright and Yu (2009) treat the two cases separately, providing two lower bounds and two different estimators that attain them in some specific asymptotics.

6. Universal aggregation. Combining the elements of a dictionary $\mathcal{H} = \{f_1, \dots, f_M\}$ to estimate a regression function η originates from the problem of aggregation introduced by Nemirovski (2000). It can be generally described as follows. Given $\Theta \subset \mathbb{R}^M$, the goal of aggregation is to construct an estimator \hat{f}_n that satisfies an oracle inequality of the form

$$(6.1) \quad \mathbb{E} \|\hat{f}_n - \eta\|^2 \leq \min_{\theta \in \Theta} \|\mathbf{f}_\theta - \eta\|^2 + C \Delta_{n,M}(\Theta), \quad C > 0,$$

with the smallest possible (in a minimax sense) remainder term $\Delta_{n,M}(\Theta)$, in which case $\Delta_{n,M}(\Theta)$ is called optimal rate of aggregation [cf. Tsybakov (2003)]. Nemirovski (2000) identified three types of aggregation: (MS) for *model selection*, (C) for *convex* and (L) for *linear*. Bunea, Tsybakov and Wegkamp (2007b) also considered another collection of aggregation problems, denoted by (L_D) for *subset selection* and indexed by $D \in \{1, \dots, M\}$. To each of these problems corresponds a given set $\Theta \subset \mathbb{R}^M$ and an optimal remainder term $\Delta_{n,M}(\Theta)$. For (MS) aggregation, $\Theta = \Theta_{(\text{MS})} = B_0(1) \cap B_1(1) = \{e_1, \dots, e_M\}$, where e_j is the j th vector of the canonical basis of \mathbb{R}^M . For (C) aggregation, $\Theta = \Theta_{(\text{C})}$ is a convex compact subset of the simplex $B_1(1) = \{\theta \in \mathbb{R}^M : |\theta|_1 \leq 1\}$. The main example of $\{\mathbf{f}_\theta, \theta \in \Theta_{(\text{C})}\}$ is the set of all convex combinations of the f_j 's. For (L) aggregation, $\Theta = \Theta_{(\text{L})} = \mathbb{R}^M = B_0(M)$, so that $\{\mathbf{f}_\theta, \theta \in \Theta_{(\text{L})}\}$ is the set of all linear combinations of the f_j 's. Given an integer $D \in \{1, \dots, M\}$, for (L_D) aggregation, $\Theta = \Theta_{(\text{L}_D)} = B_0(D) = \{\theta \in \mathbb{R}^M : M(\theta) \leq D\}$. For this problem, $\{\mathbf{f}_\theta, \theta \in \Theta_{(\text{L}_D)}\}$ is the set of all linear combinations of at most D of the f_j 's.

Note that all these sets Θ are of the form $B_0(S) \cap B_1(\delta)$ for specific values of S and δ . This allows us to apply the previous theory.

Table 1 presents the four different choices for Θ together with the optimal remainder terms given by Bunea, Tsybakov and Wegkamp (2007b). For (MS), (C) and (L) aggregation they coincide with optimal rates of aggregation originally proved in Tsybakov (2003) for the regression model with i.i.d. random design and integral L_2 norm in the risk. A fifth type of aggregation called the *D-convex* aggregation, which we denote by (C_D) was studied by Lounici (2007). In this case, $\Theta = \Theta_{(\text{L}_D)}$ is a convex compact subset of $B_1(1) \cap B_0(D)$, so that $\{\mathbf{f}_\theta, \theta \in \Theta_{(\text{L}_D)}\}$ can be, as a typical example, the set of convex combinations of at

TABLE 1
Sets of parameters $\Theta_{(MS)}$, $\Theta_{(C)}$, $\Theta_{(L)}$ and $\Theta_{(L_D)}$ and corresponding optimal rates of aggregation presented in Bunea, Tsybakov and Wegkamp (2007b). Note that Bunea, Tsybakov and Wegkamp (2007b) considered a slightly different definition in the (C) case: $\Theta_{(C)} = B_1(1) \cap \mathbb{R}_+^M$ leading to the same rate

Problem	Θ	$\Delta_{n,M}(\Theta)$
(MS)	$\Theta_{(MS)} = B_0(1) \cap B_1(1)$	$\frac{\log M}{n}$
(C)	$\Theta_{(C)} = B_1(1)$	$\frac{M}{n} \wedge \sqrt{\frac{1}{n} \log(1 + \frac{eM}{\sqrt{n}})}$
(L)	$\Theta_{(L)} = B_0(M)$	$\frac{M}{n}$
(L _D)	$\Theta_{(L_D)} = B_0(D)$	$\frac{D}{n} \log(1 + \frac{eM}{D})$

most D of the f_j 's. Lounici (2007) proves minimax lower bounds together with an upper bound that departs from the lower bound by logarithmic terms. However, the results hold in the i.i.d. random design setting and do not extend to our setup. While several papers use different estimators for different aggregation problems [see Tsybakov (2003), Rigollet (2009)], one contribution of Bunea, Tsybakov and Wegkamp (2007b) was to show that the BIC estimator defined in Section 4 satisfies oracle inequalities of the form

$$(6.2) \quad \mathbb{E} \|f_{\hat{\theta}_{BIC}} - \eta\|^2 \leq (1 + a) \min_{\theta \in \Theta} \|f_{\theta} - \eta\|^2 + C \frac{1 + a}{a^2} \Delta_{n,M}(\Theta),$$

simultaneously for all the sets Θ presented in Table 1. Here a and C are positive constants. Moreover, for the Lasso estimator defined in (4.3), Bunea, Tsybakov and Wegkamp (2007b) show less precise inequalities under the assumption the matrix $\mathbf{X}^T \mathbf{X}$ is positive definite, where \mathbf{X} is the design matrix defined in Section 2. Note that these oracle inequalities are not sharp since the leading constant is $1 + a$ and not 1, whereas letting $a \rightarrow 0$ results in blowing up the remainder term. The following theorem shows that the Exponential Screening estimator satisfies sharp oracle inequalities (i.e., with leading constant 1) that hold simultaneously for the five problems of aggregation.

THEOREM 6.1. *Assume that $\max_{1 \leq j \leq M} \|f_j\| \leq 1$. Then for any $M \geq 2$, $n \geq 1$, $D \leq M$ and $\Theta \in \{\Theta_{(MS)}, \Theta_{(C)}, \Theta_{(L)}, \Theta_{(L_D)}, \Theta_{(L_D)}\}$ the Exponential Screening estimator satisfies the following oracle inequality:*

$$\mathbb{E} \|f_{\hat{\theta}_{ES}} - \eta\|^2 \leq \min_{\theta \in \Theta} \|f_{\theta} - \eta\|^2 + C \Delta_{n,M}^*(\Theta),$$

where $C > 0$ is a numerical constant and

$$\Delta_{n,M}^*(\Theta) = \begin{cases} \frac{\sigma^2 R}{n} \wedge \frac{\sigma^2 \log M}{n}, & \text{if } \Theta = \Theta(\text{MS}), \\ \frac{\sigma^2 R}{n} \wedge \sqrt{\frac{\sigma^2}{n} \log\left(1 + \frac{eM\sigma}{\sqrt{n}}\right)}, & \text{if } \Theta = \Theta(\text{C}), \\ \frac{\sigma^2 R}{n}, & \text{if } \Theta = \Theta(\text{L}), \\ \frac{\sigma^2 R}{n} \wedge \frac{\sigma^2 D}{n} \log\left(1 + \frac{eM}{D}\right), & \text{if } \Theta = \Theta(\text{L}_D), \\ \frac{\sigma^2 R}{n} \wedge \sqrt{\frac{\sigma^2}{n} \log\left(1 + \frac{eM\sigma}{\sqrt{n}}\right)} \\ \wedge \frac{\sigma^2 D}{n} \log\left(1 + \frac{eM}{D}\right), & \text{if } \Theta = \Theta(\text{L}_D). \end{cases}$$

The proof of Theorem 6.1 follows directly from (3.10) and (3.11).

We also observe that $\Delta_{n,M}^*(\Theta)$ is to within a constant factor of $\Delta_{n,M}^*(\Theta) \wedge 1$ since $\frac{\sigma^2 R}{n} \leq \frac{\sigma^2(M \wedge n)}{n} \leq \sigma^2$.

Using Theorems 5.3 and 5.4 it is not hard to show that the rates $\Delta_{n,M}^*(\Theta) \wedge 1$ for $\Delta_{n,M}^*(\Theta)$ listed in Theorem 6.1 are *optimal rates of aggregation* in the sense of Tsybakov (2003). Indeed, it means to prove that there exists a dictionary \mathcal{H} satisfying the assumptions of Theorem 5.4, and a constant $c > 0$ such that the following lower bound holds:

$$(6.3) \quad \inf_{T_n} \sup_{\eta} \left\{ E_{\eta} \|T_n - \eta\|^2 - \min_{\theta \in \Theta} \|f_{\theta} - \eta\|^2 \right\} \geq c(\Delta_{n,M}^*(\Theta) \wedge 1),$$

where the infimum is taken over all estimators. An important observation here is that the left-hand side of (6.3) is greater than or equal to

$$(6.4) \quad \inf_{T_n} \sup_{\theta \in \Theta} E_{f_{\theta}} \|T_n - f_{\theta}\|^2.$$

It remains to note that a lower bound for (6.4) with the rate $\Delta_{n,M}^*(\Theta) \wedge 1$ follows directly from Theorem 5.4 [cf. also (5.24) and (5.25)] applied with the values S and δ corresponding to the definition of Θ .

Interestingly, the rates given in Theorem 6.1 are different from those in Table 1, and also from those for the regression model with i.i.d. random design established in Tsybakov (2003) and Lounici (2007). Indeed, they depend on the rank R of the regression matrix \mathbf{X} , and the bounds are better when the rank is smaller. This is quite natural since the distance $\|f_{\tilde{\theta}^{\text{ES}}} - \eta\|$ is the “empirical distance” depending on \mathbf{X} . One can easily understand it from the analogy with the behavior of the ordinary least squares estimator [cf. Lemma 3.1]. Alternatively, the distance used in Tsybakov (2003) and Lounici (2007) for the i.i.d. random design setting is the

$L_2(P_X)$ -distance where P_X is the marginal distribution of X_i 's, and no effects related to the rank can occur. As concerns Table 1, the optimality of the rates given there is proved in Bunea, Tsybakov and Wegkamp (2007b) only for $M \leq n$ and $\mathbf{X}^\top \mathbf{X}/n$ equal to the identity matrix, in which case $R = M$ and thus the effect of R is not visible.

7. Implementation and numerical illustration. In this section, we propose an implementation of the ES estimator together with a numerical experiment both on artificial and real data. We suppose throughout that the sample is fixed, so that the least squares estimators $\theta_p, p \in \mathcal{P}$, are fixed vectors.

7.1. *Implementation via Metropolis approximation.* Recall that the ES estimator $\tilde{\theta}^{\text{ES}}$ is the following mixture of least squares estimators:

$$(7.1) \quad \tilde{\theta}^{\text{ES}} := \frac{\sum_{p \in \mathcal{P}} \hat{\theta}_p \exp((-1/(4\sigma^2)) \sum_{i=1}^n (Y_i - f_{\hat{\theta}_p}(x_i))^2 - R_p/2) \pi_p}{\sum_{p \in \mathcal{P}} \exp((-1/(4\sigma^2)) \sum_{i=1}^n (Y_i - f_{\hat{\theta}_p}(x_i))^2 - R_p/2) \pi_p},$$

where $\mathcal{P} := \{0, 1\}^M$, π is the prior (3.4), and $\hat{\theta}_p$ is the least squares estimator on \mathbb{R}^p .

Recall also that the prior π defined in (3.4) assigns weight 1/2 to the ordinary least squares estimator $\hat{\theta}_{\mathbf{1}}$, where $\mathbf{1} = (1, \dots, 1) \in \mathcal{P}$. It is not hard to check from the proof of Theorem 3.1 that it allows us to cap the rates by $\sigma^2 R/n$. While this upper bound has important theoretical consequences, in the examples that we consider in this section, we typically have $R = n$ so that the dependence of the rates in R is inconsequential. As a result, in the rest of the section, we consider the following simpler prior:

$$(7.2) \quad \tilde{\pi}_p := \begin{cases} \frac{2}{H} \left(\frac{|p|}{2eM} \right)^{|p|}, & \text{if } |p| < R, \\ 0, & \text{otherwise.} \end{cases}$$

Exact computation of $f_{\tilde{\theta}^{\text{ES}}}$ requires the computation of 2^{R-1} least squares estimators. In many applications this number is prohibitively large and we need to resort to a numerical approximation. Notice that $\tilde{\theta}^{\text{ES}}$ is obtained as the expectation of the random variable $\hat{\theta}_{\mathbf{P}}$ where \mathbf{P} is a random variable taking values in \mathcal{P} with probability mass function ν given by

$$\nu_p \propto \exp\left(-\frac{1}{4\sigma^2} \sum_{i=1}^n (Y_i - f_{\hat{\theta}_p}(x_i))^2 - \frac{R_p}{2}\right) \tilde{\pi}_p, \quad p \in \mathcal{P}.$$

This Gibbs-type distribution can be expressed as the stationary distribution of the Markov chain generated by the Metropolis–Hastings (MH) algorithm [see, e.g., Robert and Casella (2004), Section 7.3]. We now describe the MH algorithm employed here. Consider the M -hypercube graph \mathcal{G} with vertices given by \mathcal{P} . For

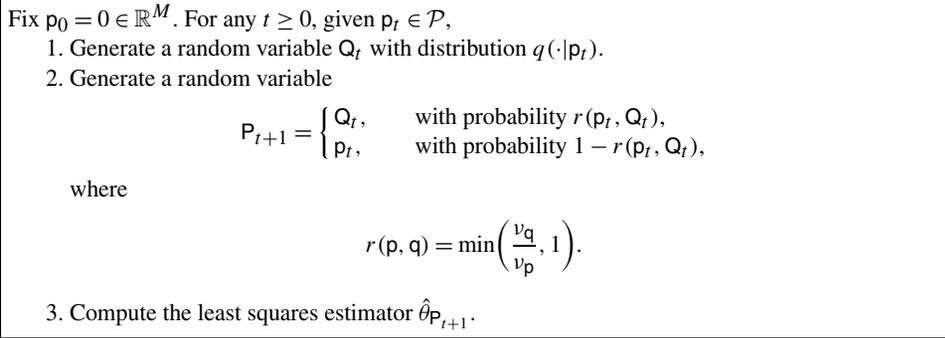


FIG. 1. *The Metropolis–Hastings algorithm on the M-hypercube.*

any $p \in \mathcal{P}$, define the instrumental distribution $q(\cdot|p)$ as the uniform distribution on the neighbors of p in \mathcal{G} , and notice that since each vertex has the same number of neighbors, we have $q(p|q) = q(q|p)$ for any $p, q \in \mathcal{P}$. The MH algorithm is defined in Figure 1. We use here the uniform instrumental distribution for the sake of simplicity. Our simulations show that it yields satisfactory results both in performance and in the speed. Another choice of $q(\cdot|\cdot)$ can potentially further accelerate the convergence of the MH algorithm.

The following theorem ensures the ergodicity of the Markov chain generated by the MH algorithm.

THEOREM 7.1. *For any function $p \mapsto \theta_p \in \mathbb{R}^M$, the Markov chain $(P_t)_{t \geq 0}$ defined by the MH algorithm satisfies*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=T_0+1}^{T_0+T} \theta_{P_t} = \sum_{p \in \mathcal{P}} \theta_p \nu_p, \quad \nu\text{-almost surely,}$$

where $T_0 \geq 0$ is an arbitrary integer.

PROOF. The chain is clearly ν -irreducible, so the result follows from Robert and Casella (2004), Theorem 7.4, page 274. \square

In view of this result, we approximate $\tilde{\theta}^{ES} = \sum_{p \in \mathcal{P}} \theta_p \nu_p$ by

$$\tilde{\theta}_T^{ES} = \frac{1}{T} \sum_{t=T_0+1}^{T_0+T} \hat{\theta}_{P_t},$$

which is close to $\tilde{\theta}^{ES}$ for sufficiently large T . One salient feature of the MH algorithm is that it involves only the ratios ν_q/ν_p where p and q are two neighbors in \mathcal{G} . Since

$$\frac{\nu_q}{\nu_p} = \exp\left(\frac{1}{4\sigma^2} \sum_{i=1}^n [(Y_i - f_{\hat{\theta}_p}(x_i))^2 - (Y_i - f_{\hat{\theta}_q}(x_i))^2] + \frac{R_p - R_q}{2}\right) \frac{\tilde{\pi}_q}{\tilde{\pi}_p},$$

the MH algorithm benefits from the choice (7.2) of the prior $\tilde{\pi}$ in terms of speed. Indeed, for this prior, we have

$$\frac{\tilde{\pi}_q}{\tilde{\pi}_p} = \left(1 + \frac{\omega}{|p|}\right)^{|q|} \left(\frac{|p|}{2eM}\right)^\omega,$$

and $\omega = |q| - |p| \in \{-1, 1\}$ when p and q are two neighbors in \mathcal{G} . In this respect, the choice of the prior $\tilde{\pi}$ as in (7.2) is better than the suggestions in [Leung and Barron \(2006\)](#) and [Giraud \(2008\)](#) who consider priors that require the computation of the combinatoric quantity $\binom{M}{|p|}$. Moreover, the choice (7.2) yields slightly better constants and improves the remainder terms in the oracle inequalities of Section 3, as compared to what would be obtained with those priors.

As a result, the MH algorithm in this case takes the form of a stochastic greedy algorithm with averaging, which measures a tradeoff between sparsity and prediction to decide whether to add or remove a variable. In all subsequent examples, we use a pure MATLAB implementation of the ES estimator. While the benchmark estimators considered below employ a C-based code optimized for speed, we observed that a safe implementation of the MH algorithm (three time more iterations than needed) exhibited an increase of computation time of at most a factor two.

7.2. Numerical experiments.

7.2.1. Sparse recovery. While our results for the ES estimator hold under no assumption on the dictionary, we first compare the behavior of our algorithm in a well-known example where the L -conditions on the dictionary are satisfied and therefore sparse recovery by ℓ_1 -penalized techniques is theoretically achievable.

Consider the model $\mathbf{Y} = \mathbf{X}\theta^* + \sigma\xi$, where \mathbf{X} is an $n \times M$ matrix with independent Rademacher or standard Gaussian entries and $\xi \in \mathbb{R}^n$ is a vector of independent standard Gaussian random variables and is independent of \mathbf{X} . The vector θ^* is given by $\theta_j^* = \mathbb{1}(j \leq S)$ for some fixed S so that $M(\theta^*) = S$. The variance is chosen as $\sigma^2 = S/9$ following the numerical experiments of [Candes and Tao \(2007\)](#), Section 4. For different values of (n, M, S) , we run the ES algorithm on 500 replications of the problem and compare our results with several other popular estimators in the sparse recovery literature. We limit our choice to estimators that are readily implemented in R or MATLAB. The considered estimators are:

1. the Lasso estimator with regularization parameter $\sigma\sqrt{8(\log M)/n}$ as indicated in [Bickel, Ritov and Tsybakov \(2009\)](#);
2. the cross-validated Lasso estimator (LassoCV) with regularization parameter obtained by ten-fold cross-validation;
3. the Lasso–Gauss estimator (Lasso–G) corresponding to the Lasso estimator computed in 1 and threshold value given by $\sigma\sqrt{(2\log M)/n}$;
4. the cross-validated Lasso–Gauss estimator (LassoCV–G) corresponding to the Lasso estimator computed in 2 and threshold value given by $\sigma\sqrt{(2\log M)/n}$;

5. the MC+ estimator of Zhang (2010) with regularization parameter $\sigma \times \sqrt{(2 \log M)/n}$;
6. the SCAD estimator of Fan and Li (2001) with regularization parameter $\sigma \sqrt{(2 \log M)/n}$.

The Lasso–Gauss estimators in 3 and 4 are obtained using the following two-step procedure. In the first step, a Lasso estimator (Lasso or LassoCV) is computed and only coordinates larger than the threshold $\sigma \sqrt{2(\log M)/n}$ are retained in a set \mathcal{J} . In the second step, the Lasso–Gauss estimators are obtained by constrained least squares under the constraint that coordinates $\hat{\beta}_j \notin \mathcal{J}$ are equal to 0. Indeed, it is usually observed that the Lasso estimator induces a strong bias by over-shrinking large coefficients and the Lasso–Gauss procedure is a practically efficient remedy to this issue. By construction, the SCAD and MC+ estimators should not suffer from such a shrinkage. The Lasso estimators are based on the `l1-ls` package in MATLAB [Koh, Kim and Boyd (2008)]. The MC+ and SCAD estimators are implemented in the `plus` package in R [Zhang and Melnik (2009)].

The performance of each of the seven estimators generically denoted by $\hat{\theta}$ is measured by its prediction error $|\mathbf{X}(\hat{\theta} - \theta^*)|_2^2/n = \|\mathbf{f}_{\hat{\theta}} - \mathbf{f}_{\theta^*}\|^2$. Moreover, even though the estimation error $|\hat{\theta} - \theta^*|_2^2$ is not studied above, we also report its values in Table 3, for a better comparison with other simulation studies. We considered the cases $(n, M, S) \in \{(100, 200, 10), (200, 500, 20)\}$. The Metropolis approximation $\tilde{\theta}_T^{\text{ES}}$ was computed with $T_0 = 3000, T = 7000$, which should be in the asymptotic regime of the Markov chain since Figure 3 shows that on a typical example, the right sparsity pattern is recovered after about 2000 iterations.

Figure 2 displays comparative boxplots for both Gaussian and Rademacher design matrix. In particular, it shows that ES outperforms all six other estimators and has less variability across repetitions.

Figure 3 illustrates a typical behavior of the ES estimator for one particular realization of \mathbf{X} and ξ . For better visibility, both displays represent only the 50 first coordinates of $\tilde{\theta}_T^{\text{ES}}$, with $T = 7000$ and $T_0 = 3000$. The left-hand side display shows that the sparsity pattern is well recovered and the estimated values are close to one. The right-hand side display illustrates the evolution of the intermediate parameter $\hat{\theta}_{P_t}$ for $t = 1, \dots, 5000$. It is clear that the Markov chain that runs on the M -hypercube graph gets *trapped* in the vertex that corresponds to the sparsity pattern of θ^* after only 2000 iterations. As a result, while the ES estimator is not sparse itself, the MH approximation to the ES estimator may output a sparse solution. A covariate X_j is considered to be *selected* by an estimator $\hat{\theta}$, if $|\hat{\theta}_j| > 1/n$. Hence, for any two vectors $\theta^{(1)}, \theta^{(2)} \in \mathbb{R}^M$ define $\theta^{(1)} \Delta \theta^{(2)} \in \{0, 1\}^M$ as the binary vector with j th coordinate given by

$$(\theta^{(1)} \Delta \theta^{(2)})_j = \mathbb{1}(|\theta_j^{(1)}| > 1/n, \theta_j^{(2)} = 0) + \mathbb{1}(\theta_j^{(1)} = 0, |\theta_j^{(2)}| > 1/n).$$

The performance of an estimator $\hat{\theta}$ in terms of model selection is measured by the number $M(\hat{\theta} \Delta \theta^*)$ of variables that are incorrectly selected or incorrectly left out of

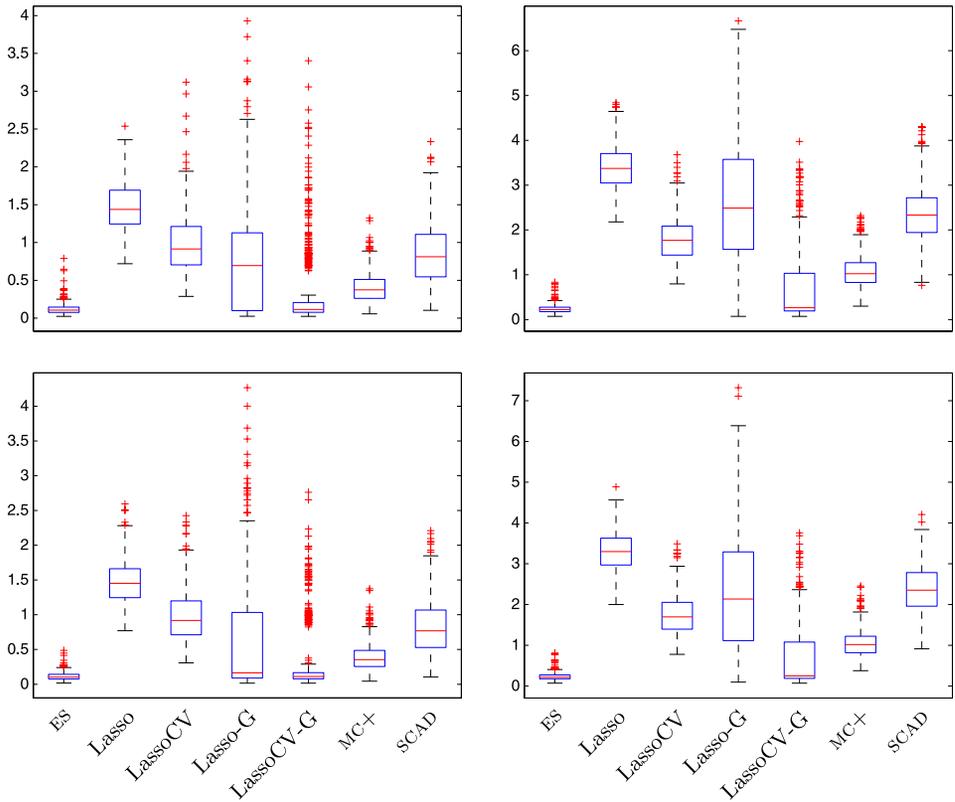


FIG. 2. Boxplots of $|\mathbf{X}(\hat{\theta} - \theta^*)|_2^2/n$ over 500 realizations for the ES, Lasso, cross-validated Lasso (LassoCV), Lasso–Gauss (Lasso–G), cross-validated Lasso–Gauss (LassoCV–G), MC+ and SCAD estimators. Left: $(n, M, S) = (100, 200, 10)$, right: $(n, M, S) = (200, 500, 20)$, top: Gaussian design, bottom: Rademacher design.

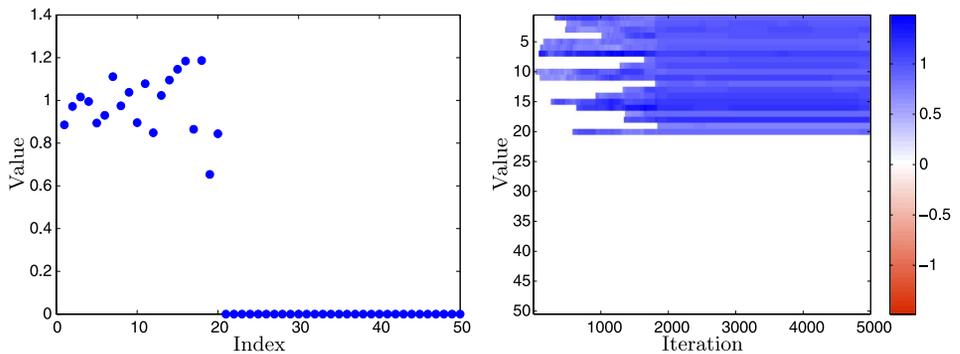


FIG. 3. Typical realization for $(M, n, S) = (500, 200, 20)$ and Gaussian design. Left: value of the $\tilde{\theta}_T^{\text{ES}}$, $T = 7000$, $T_0 = 3000$. Right: value of $\hat{\theta}_{P_t}$ for $t = 1, \dots, 5000$. Only the first 50 coordinates are shown for each vector.

the model. Among the four procedures considered here, MC+ uniformly dominates the other three in terms of model selection. Table 4 displays the relative average model selection error (RAMS) over 500 repetitions of each of the experiments described above

$$(7.3) \quad \text{RAMS}(\hat{\theta}) = \frac{\sum_{i=1}^{500} M(\hat{\theta}^{(i)} \Delta \theta^*)}{\sum_{i=1}^{500} M(\hat{\theta}^{(i), \text{MC+}} \Delta \theta^*)},$$

where for each repetition i of the experiment, $\hat{\theta}^{(i), \text{MC+}}$ denotes the MC+ estimator and $\hat{\theta}^{(i)}$ is one of the four estimators: ES, Lasso, MC+ or SCAD.

While MC+ uniformly dominates the three other procedures, the model selection properties of ES are better than Lasso but not as good as SCAD and the relative performance of ES improves when the problem size increases. The superiority of MC+ and SCAD does not come as a surprise as these procedures are designed for variable selection. However, ES makes up for this deficiency by having much better estimation and prediction properties.

To conclude this numerical experiment in the linear regression model, notice that we used the knowledge of the variance parameter σ^2 to construct the estimators, except for those based on cross-validation. In particular, ES depends on σ^2 and it is necessary to be able to implement it without such a knowledge. While an obvious solution consists in resorting to cross-validation or bootstrap, such procedures tend to become computationally burdensome. We propose the following estimator for σ^2 . Let $\bar{\theta}^{\text{ES}}(\bar{\sigma}^2)$ denote the estimator obtained by replacing σ^2 with any upper bound $\bar{\sigma}^2 \geq \sigma^2$ in the definition (7.1) of the ES estimator. Define

$$\hat{\sigma}^2 = \inf \left\{ s^2 : \left| \frac{\|\mathbf{Y} - \mathbf{X}\bar{\theta}^{\text{ES}}(s^2)\|_2^2}{n - M_n(\bar{\theta}^{\text{ES}}(s^2))} - s^2 \right| > \alpha \right\},$$

where $\alpha > 0$ is a tolerance parameter and for any $\theta \in \mathbb{R}^M$, $M_n(\theta) = \sum_{j=1}^M \mathbb{1}(|\theta_j| > 1/n)$. As a result, the proposed estimator $\hat{\sigma}^2$ is the smallest positive value that departs from the usual estimator for the variance by more than α . The motivation for this estimator comes from the following heuristics, which is loosely inspired by Zhang (2010), Section 5.2. It follows from the results of Leung and Barron (2006) that $\bar{\theta}^{\text{ES}}(\bar{\sigma}^2)$ satisfies the oracle inequalities of Section 3 and thus of Section 6 with σ^2 replaced by $\bar{\sigma}^2$. As a consequence, we can use any upper bound $\bar{\sigma}^2 \geq \sigma^2$ to compute an estimator $\bar{\theta}^{\text{ES}}(\bar{\sigma}^2)$ and thus, an estimator of the variance based on the residuals. Our heuristics consists in choosing the smallest upper bound that is inconsistent with the estimator based on the residuals. Figure 4 and Table 5 summarize the performance of the variance estimator $\hat{\sigma}^2$ and the corresponding ES estimator $\bar{\theta}^{\text{ES}}(\hat{\sigma}^2)$ for $\alpha = 1$.

Note that in Table 5, the obtained values are comparable to those in Tables 2 and 3. In particular, it is remarkable that the ES estimator with estimated variance still has smaller prediction and estimation errors in these experiments than the other six estimators considered methods.

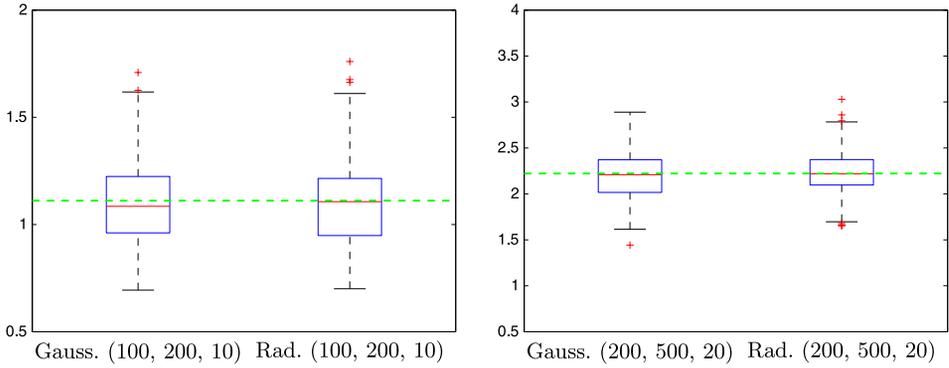


FIG. 4. Boxplots of the estimated variance $\hat{\sigma}^2$ based on 500 replications of each of the four experiments described above. The horizontal dashed lines indicate the value of the true parameter $\sigma^2 = S/9$. Left: $\sigma^2 = 1.11$. Right: $\sigma^2 = 2.22$.

7.2.2. *Handwritten digits dataset.* The aim of this subsection is to illustrate the performance of the ES algorithm on a real dataset and to compare it with the state-of-the-art procedure in sparse estimation, namely the Lasso. While sparse estimation is the object of many recent statistical studies, it is still hard to find a freely available benchmark dataset where $M \gg n$. We propose the following real dataset originally introduced in LeCun et al. (1990) and, in the particular instance of this paper, obtained from the webpage of the book by Hastie, Tibshirani and Friedman (2001). We observe a grayscale image of size 16×16 pixels of the handwritten digit “6” (see Figure 6) which is artificially corrupted by a Gaussian noise. Formally, we can write

$$(7.4) \quad \mathbf{Y} = \mu + \sigma \xi,$$

TABLE 2
Means and standard deviations of $|\mathbf{X}(\hat{\theta} - \theta^*)|_2^2/n$ over 500 realizations for the ES, Lasso, cross-validated Lasso (LassoCV), Lasso–Gauss (Lasso–G), cross-validated Lasso–Gauss (LassoCV–G), MC+ and SCAD estimators. Top: Gaussian design, bottom: Rademacher design

(M, n, S)	ES	Lasso	LassoCV	Lasso–G	LassoCV–G	MC+	SCAD
(100, 200, 10)	0.12 (0.07)	1.47 (0.31)	0.99 (0.40)	0.75 (0.77)	0.35 (0.53)	0.41 (0.20)	0.86 (0.40)
(200, 500, 20)	0.24 (0.10)	3.39 (0.50)	1.81 (0.50)	2.55 (1.45)	0.70 (0.76)	1.07 (0.35)	2.37 (0.64)
(100, 200, 10)	0.12 (0.06)	1.48 (0.31)	0.99 (0.38)	0.70 (0.79)	0.30 (0.47)	0.39 (0.19)	0.83 (0.39)
(200, 500, 20)	0.24 (0.09)	3.32 (0.49)	1.76 (0.49)	2.34 (1.44)	0.66 (0.74)	1.05 (0.33)	2.37 (0.61)

TABLE 3
Means and standard deviations of $|\hat{\theta} - \theta^|_2^2$ over 500 realizations for the ES, Lasso, cross-validated Lasso (LassoCV), Lasso-Gauss (Lasso-G), cross-validated Lasso-Gauss (LassoCV-G), MC+ and SCAD estimators. Top: Gaussian design, bottom: Rademacher design*

(M, n, S)	ES	Lasso	LassoCV	Lasso-G	LassoCV-G	MC+	SCAD
(100, 200, 10)	0.14 (0.12)	2.06 (0.72)	1.42 (0.66)	1.08 (1.22)	0.48 (0.84)	0.56 (0.34)	1.30 (0.81)
(200, 500, 20)	0.27 (0.13)	4.72 (1.24)	2.73 (0.88)	3.62 (2.29)	0.93 (1.13)	1.45 (0.63)	3.51 (1.33)
(100, 200, 10)	0.13 (0.07)	1.99 (0.71)	1.37 (0.60)	0.94 (1.19)	0.38 (0.68)	0.51 (0.35)	1.21 (0.81)
(200, 500, 20)	0.26 (0.11)	4.50 (1.14)	2.60 (0.80)	3.20 (2.20)	0.82 (1.00)	1.38 (0.56)	3.44 (1.22)

TABLE 4
Relative average model selection error (RAMS) defined in (7.3) over 500 realizations for the ES, Lasso, MC+ and SCAD estimators. Top: Gaussian design, bottom: Rademacher design

Design (M, n, S)	ES	Lasso	MC+	SCAD
Gauss. (100, 200, 10)	10.54	12.43	1.00	3.56
Gauss. (200, 500, 20)	9.26	15.81	1.00	6.04
Rad. (100, 200, 10)	13.18	15.80	1.00	3.59
Rad. (200, 500, 20)	10.07	16.18	1.00	6.18

TABLE 5
Means and standard deviations of prediction error $|\mathbf{X}(\bar{\theta}^{\text{ES}}(\hat{\sigma}^2) - \theta^)|_2^2/n$ and estimation error $|\bar{\theta}^{\text{ES}}(\hat{\sigma}^2) - \theta^*|_2^2$ over 500 realizations for the ES estimator $\bar{\theta}^{\text{ES}}(\hat{\sigma}^2)$ with estimated variance*

Design (M, n, S)	$ \mathbf{X}(\bar{\theta}^{\text{ES}}(\hat{\sigma}^2) - \theta^*) _2^2/n$	$ \bar{\theta}^{\text{ES}}(\hat{\sigma}^2) - \theta^* _2^2$
Gauss. (100, 200, 10)	0.12 (0.09)	0.14 (0.14)
Gauss. (200, 500, 20)	0.26 (0.19)	0.31 (0.32)
Rad. (100, 200, 10)	0.12 (0.07)	0.13 (0.08)
Rad. (200, 500, 20)	0.25 (0.11)	0.28 (0.14)

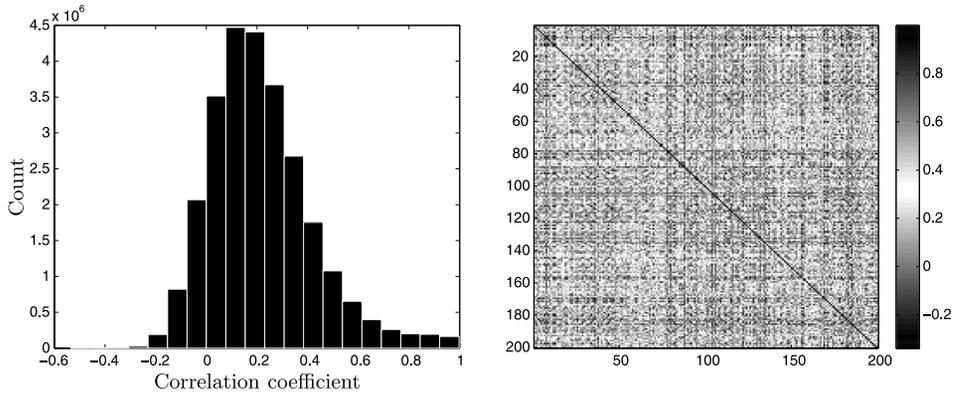


FIG. 5. Left: histogram of the $M(M - 1)/2$ correlation coefficients between different images in the database. Right: the upper left corner of size 200×200 of the full correlation matrix. Notice that only the absolute value of the correlation coefficients is discriminative in terms of color. The dark, off-diagonal regions are characteristic of correlated features.

where $\mathbf{Y} \in \mathbb{R}^{256}$ is the observed image, $\mu \in [0, 1]^{256}$ is the true image, $\sigma > 0$ and $\xi \in \mathbb{R}^{256}$ is a standard Gaussian vector. Therefore the number of observations is equal to the number of pixels: $n = 256$. The goal is to reconstruct μ using linear combinations of vectors $x_1, \dots, x_M \in [0, 1]^{256}$ that form a dictionary of size $M = 7290$. Each vector x_j is a 16×16 grayscale image of a handwritten digit from 0 to 9. As a result, x_j 's are strongly correlated as illustrated by the correlation matrix displayed in Figure 5. The digit “6” is a notably hard instance due to its similarity with the digits “0” and with some instances of the digit “5” (see Figure 9). Given an estimator $\hat{\theta}$, the performance is measured by the prediction error $\|\mu - \mathbf{X}\hat{\theta}\|_2^2$, where \mathbf{X} is the $n \times M$ design matrix formed by horizontal concatenation of the column vectors $x_1, \dots, x_M \in \mathbb{R}^n$.

Figures 6 and 7 illustrate the reconstruction of this digit by the ES, Lasso and Lasso–Gauss estimators for $\sigma = 0.5$ and $\sigma = 1$, respectively. The latter two estimators were computed with fixed regularization parameter equal to $\sigma\sqrt{8(\log M)/n}$ and the threshold for the Lasso–Gauss estimator was taken equal to $\sigma\sqrt{2(\log M)/n}$. It is clear from those figures that the Lasso estimator reconstructs the noisy image and not the true one indicating that the regularization parameter $\sigma\sqrt{8(\log M)/n}$ may be too small for this problem.

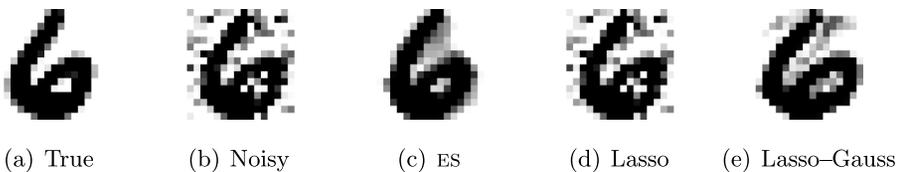


FIG. 6. Reconstruction of the digit “6” with $\sigma = 0.5$.

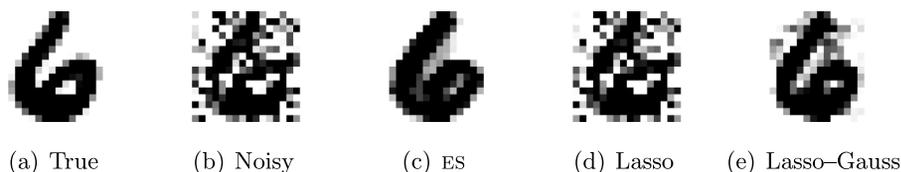


FIG. 7. Reconstruction of the digit “6” with $\sigma = 1.0$.

For both $\sigma = 0.5$ and $\sigma = 1$, the experiment was repeated 250 times and the predictive performance of ES was compared with that of the Lasso and Lasso–Gauss estimators. The results are represented in Figure 8 and Table 6.

To conclude, we mention a byproduct of this simulation study. The coefficients of $\tilde{\theta}^{\text{ES}}$ can be used to perform multi-class classification following the idea of Wright et al. (2009). The procedure consists of performing a majority vote on the features x_j that are positively weighted by $\tilde{\theta}^{\text{ES}}$, that is, such that $\tilde{\theta}_j^{\text{ES}} > 0$. For the particular instance illustrated in Figure 6(c), we see in Figure 9 that only a few features x_j receive a large positive weight and that a majority of those correspond to the digit “6.”

APPENDIX

A.1. Lemmas for the upper bound. The following lemma is obtained by a variant of the “Maurey argument” [cf. also Barron (1993), Bunea, Tsybakov and Wegkamp (2007b), Bickel, Ritov and Tsybakov (2008) for similar but somewhat different results].

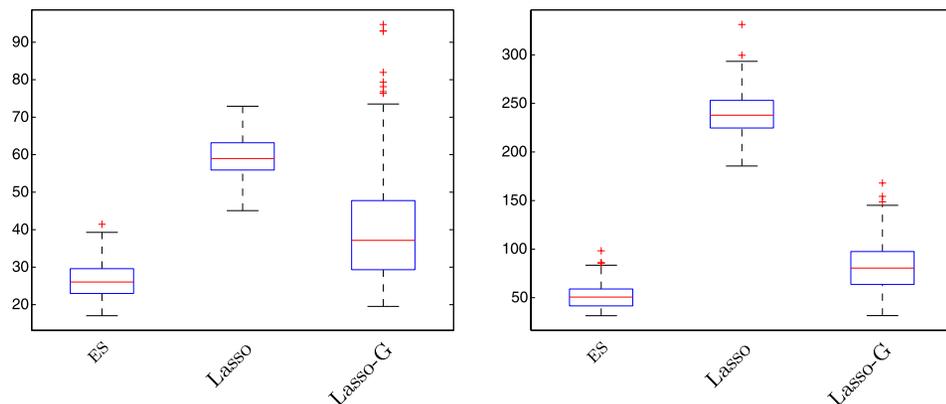


FIG. 8. Boxplots of the predictive performance $|\mu - \mathbf{X}\hat{\theta}|_2^2$ of the ES, Lasso and Lasso–Gauss (Lasso–G) estimators computed from 250 replications of the model (7.4) with μ corresponding to the digit “6.” Left: $\sigma = 0.5$. Right: $\sigma = 1$. Notice that each graph uses a different scale.

TABLE 6
Means and standard deviations for $|\mu - \mathbf{X}\hat{\theta}|_2^2$ over 250 realizations of the ES, Lasso and Lasso–Gauss estimators to reconstruct the digit “6”

	ES	Lasso	Lasso–Gauss
$\sigma = 0.5$	26.57 (4.57)	59.49 (5.28)	40.55 (14.58)
$\sigma = 1.0$	51.70 (12.32)	239.39 (22.12)	82.95 (24.40)

LEMMA A.1. For any $\theta^* \in \mathbb{R}^M \setminus \{0\}$, any integer $k \geq 1$ and any function f we have

$$\min_{\substack{\theta: |\theta|_1 = |\theta^*|_1 \\ M(\theta) \leq k}} \|f - \mathbf{f}_\theta\|^2 \leq \|f - \mathbf{f}_{\theta^*}\|^2 + \frac{|\theta^*|_1^2}{\min(k, M(\theta^*))}.$$

PROOF. Fix $\theta^* \in \mathbb{R}^M \setminus \{0\}$ and an integer $k \geq 1$. Set $K = \min(k, M(\theta^*))$. Consider the multinomial parameter $p = (p_1, \dots, p_M)^\top \in [0, 1]^M$ with $p_j = |\theta_j^*|/|\theta^*|_1, j = 1, \dots, M$, where θ_j^* are the components of θ^* . Let $\kappa = (\kappa_1, \dots, \kappa_M)^\top \in \{0, 1, \dots, M\}^M$ be the random vector with multinomial distribution $\mathcal{M}(K, p)$, that is, let $\kappa_j = \sum_{s=1}^K \mathbb{1}(I_s = j)$ where I_1, \dots, I_K are i.i.d. random variables taking value $j \in \{1, \dots, M\}$ with probability $p_j, j = 1, \dots, M$. In particular, we have $\mathcal{E}(\kappa_j) = Kp_j, j = 1, \dots, M$, where \mathcal{E} denotes the expectation with respect to the multinomial distribution. As a result, for the random vector $\bar{\theta} \in \mathbb{R}^M$ with the components $\bar{\theta}_j = \kappa_j \text{sign}(\theta_j^*)|\theta^*|_1/K$ we have $\mathcal{E}(\bar{\theta}_j) = \theta_j^*$ for $j = 1, \dots, M$ with the convention that $\text{sign}(0) = 0$. Moreover, using the fact that $\text{Var}(\kappa_j) = Kp_j(1 - p_j)$ and $\text{Cov}(\kappa_j, \kappa_l) = -np_j p_l$ for $j \neq l$ [see, e.g., Bickel and

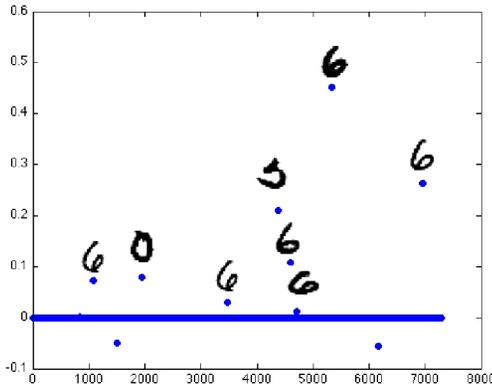


FIG. 9. Coefficients of $\tilde{\theta}_T^{\text{ES}}, T = 10,000$ and the corresponding image.

Doksum (2006), equation (A.13.15), page 462] we find that the covariance matrix of $\bar{\theta}$ is given by

$$\Sigma^* = \mathcal{E}[(\bar{\theta} - \theta^*)(\bar{\theta} - \theta^*)^\top] = \frac{|\theta^*|_1}{K} \text{diag}(|\theta_j^*|) - \frac{1}{K} |\theta^*| |\theta^*|^\top,$$

where $|\theta^*| = (|\theta_j^*|, \dots, |\theta_j^*|)^\top$. Using a bias–variance decomposition together with the assumption $\max_j \|f_j\| \leq 1$, it yields that, for any function f ,

$$\mathcal{E} \|f - f_\theta\|^2 = \|f - f_{\theta^*}\|^2 + \frac{1}{n} \sum_{i=1}^n F(x_i)^\top \Sigma^* F(x_i) \leq \|f - f_{\theta^*}\|^2 + \frac{|\theta^*|_1^2}{K},$$

where $F(x_i) = (f_1(x_i), \dots, f_M(x_i))^\top, i = 1, \dots, n$. Moreover, since $\bar{\theta}$ is such that $|\bar{\theta}|_1 = |\theta^*|_1$ and $M(\bar{\theta}) \leq K$, the lemma follows. \square

LEMMA A.2. Fix $M, n \geq 1$ and assume that $\max_j \|f_j\| \leq 1$. For any function η and any constant $v > 0$ we have

$$\begin{aligned} \min_{\theta \in \mathbb{R}^M} \left\{ \|f_\theta - \eta\|^2 + v^2 \frac{M(\theta)}{n} \log \left(1 + \frac{eM}{M(\theta) \vee 1} \right) \right\} \\ \leq \min_{\theta \in \mathbb{R}^M} \{ \|f_\theta - \eta\|^2 + \tilde{c} \bar{\varphi}_{n,M}(\theta) \}, \end{aligned} \tag{A.1}$$

where $\tilde{c} = (3 + \frac{1}{e})$, $\bar{\varphi}_{n,M}(0) = 0$ and for $\theta \neq 0$,

$$\bar{\varphi}_{n,M}(\theta) = \begin{cases} \min \left[\frac{v|\theta|_1}{\sqrt{n}} \sqrt{\log \left(1 + \frac{eMv}{|\theta|_1 \sqrt{n}} \right)}, |\theta|_1^2 \right], & \text{if } \langle f_\theta, \eta \rangle \leq \|f_\theta\|^2, \\ \frac{v|\theta|_1}{\sqrt{n}} \sqrt{\log \left(1 + \frac{eMv}{|\theta|_1 \sqrt{n}} \right)} + \frac{v^2 \log(1 + eM)}{\tilde{c}n}, & \text{otherwise.} \end{cases} \tag{A.2}$$

PROOF. Set

$$A = \min_{\theta \in \mathbb{R}^M} \left\{ \|f_\theta - \eta\|^2 + v^2 \frac{M(\theta)}{n} \log \left(1 + \frac{eM}{M(\theta) \vee 1} \right) \right\}.$$

It suffices to consider $\mathbb{R}^M \setminus \{0\}$ instead of \mathbb{R}^M since $A \leq \|f_0 - \eta\|^2 + \tilde{c} \bar{\varphi}_{n,M}(0) = \|\eta\|^2$. Fix $\theta^* \in \mathbb{R}^M \setminus \{0\}$ and define

$$x^* = |\theta^*|_1 \sqrt{n} / \ell \quad \text{where } \ell = v \sqrt{\log \left(1 + \frac{eMv}{|\theta^*|_1 \sqrt{n}} \right)}.$$

Assume first $x^* \leq 1$. In this case we have

$$|\theta^*|_1^2 \leq v \frac{|\theta^*|_1}{\sqrt{n}} \sqrt{\log \left(1 + \frac{eMv}{|\theta^*|_1 \sqrt{n}} \right)}. \tag{A.3}$$

The previous display yields that $\bar{\varphi}_{n,M}(\theta^*) = |\theta^*|_1^2$. Moreover, if $\langle f_{\theta^*}, \eta \rangle \leq \|f_{\theta^*}\|^2$, it holds

$$\|\eta\|^2 \leq \|f_{\theta^*} - \eta\|^2 + \|f_{\theta^*}\|^2 \leq \|f_{\theta^*} - \eta\|^2 + |\theta^*|_1^2.$$

As a result,

$$(A.4) \quad A \leq \|\eta\|^2 \leq \bar{\varphi}_{n,M}(\theta^*) \quad \text{if } \langle f_{\theta^*}, \eta \rangle \leq \|f_{\theta^*}\|^2 \text{ and } x^* \leq 1.$$

Set $k^* = \lceil x^* \rceil$, that is, k^* is the minimal integer greater than or equal to x^* . Using the monotonicity of the mapping $t \mapsto \frac{1}{n} \log(1 + \frac{eM}{t})$ for $t > 0$, and Lemma A.1 we get, for any $\theta^* \in \mathbb{R}^M \setminus \{0\}$ such that $k^* \leq M(\theta^*)$,

$$\begin{aligned} A &\leq \min_{\theta \in \mathbb{R}^M} \left\{ \|f_{\theta} - \eta\|^2 + v^2 \frac{M(\theta)}{n} \log\left(1 + \frac{eM}{M(\theta) \vee 1}\right) \right\} \\ &\leq \min_{1 \leq k \leq M(\theta^*)} \min_{\theta: M(\theta) \leq k} \left\{ \|f_{\theta} - \eta\|^2 + v^2 \frac{k}{n} \log\left(1 + \frac{eM}{k}\right) \right\} \\ &\leq \|f_{\theta^*} - \eta\|^2 + \min_{1 \leq k \leq M(\theta^*)} \left\{ v^2 \frac{k}{n} \log\left(1 + \frac{eM}{k}\right) + \frac{|\theta^*|_1^2}{k} \right\} \\ &\leq \|f_{\theta^*} - \eta\|^2 + v^2 \frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right) + \frac{|\theta^*|_1^2}{k^*}. \end{aligned}$$

On the other hand, if $\theta^* \in \mathbb{R}^M \setminus \{0\}$ and $k^* > M(\theta^*)$, we use the simple bound

$$\begin{aligned} A &\leq \|f_{\theta^*} - \eta\|^2 + v^2 \frac{M(\theta^*)}{n} \log\left(1 + \frac{eM}{M(\theta^*) \vee 1}\right) \\ &\leq \|f_{\theta^*} - \eta\|^2 + v^2 \frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right). \end{aligned}$$

In view of the last two displays, to conclude the proof it suffices to show that

$$(A.5) \quad v^2 \frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right) + \frac{|\theta^*|_1^2}{k^*} \leq \tilde{c} \bar{\varphi}_{n,M}(\theta^*)$$

for all $\theta^* \in \mathbb{R}^M \setminus \{0\}$. Note first that if $x^* \leq 1$, then $k^* = 1$ and

$$\begin{aligned} v^2 \frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right) + \frac{|\theta^*|_1^2}{k^*} &\leq \frac{v^2 \log(1 + eM)}{n} + v \frac{|\theta^*|_1}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eMv}{|\theta^*|_1 \sqrt{n}}\right)} \\ &\leq \tilde{c} \bar{\varphi}_{n,M}(\theta^*), \end{aligned}$$

where we used (A.3) in the first inequality. Together with (A.4), this proves that $A \leq \|f_{\theta^*} - \eta\|^2 + \tilde{c} \bar{\varphi}_{n,M}(\theta^*)$ for all $\theta^* \in \mathbb{R}^M \setminus \{0\}$ such that $x^* \leq 1$. Thus, to complete the proof of the lemma we only need to consider the case $x^* > 1$. For $x^* > 1$ we have

$$\bar{\varphi}_{n,M}(\theta^*) \geq v \frac{|\theta^*|_1}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eMv}{|\theta^*|_1 \sqrt{n}}\right)}.$$

As a result, we have

$$(A.6) \quad \frac{|\theta^*|_1^2}{k^*} \leq \frac{|\theta^*|_1 \ell}{\sqrt{n}} \leq \bar{\varphi}_{n,M}(\theta^*).$$

Moreover, it holds $k^* \leq 2x^* = 2|\theta^*|_1 \sqrt{n}/\ell$ and since the function $t \mapsto \frac{t}{n} \log(1 + \frac{eM}{t})$ is increasing, we obtain

$$\frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right) \leq \frac{2|\theta^*|_1}{\ell\sqrt{n}} \log\left(1 + \frac{eM\ell}{2|\theta^*|_1\sqrt{n}}\right).$$

Thus, for $\ell \leq \nu$ we have

$$\frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right) \leq \frac{2|\theta^*|_1}{\ell\sqrt{n}} \log\left(1 + \frac{eM\nu}{2|\theta^*|_1\sqrt{n}}\right) \leq \frac{2|\theta^*|_1\ell}{\nu^2\sqrt{n}} \leq \frac{2}{\nu^2} \bar{\varphi}_{n,M}(\theta^*).$$

For $\ell > \nu$ we use the inequality $\log(1 + ab) \leq \log(1 + a) + \log b, \forall a \geq 0, b \geq 1$, to obtain

$$\begin{aligned} \frac{k^*}{n} \log\left(1 + \frac{eM}{k^*}\right) &\leq \frac{2|\theta^*|_1}{\ell\sqrt{n}} \left[\log\left(1 + \frac{eM\nu}{2|\theta^*|_1\sqrt{n}}\right) + \log\left(\frac{\ell}{\nu}\right) \right] \\ &\leq \frac{2|\theta^*|_1}{\sqrt{n}} \left(\frac{\ell}{\nu^2} + \frac{\log(\ell/\nu)}{\ell} \right) \leq \left(2 + \frac{1}{e}\right) \frac{|\theta^*|_1\ell}{\nu^2\sqrt{n}} \\ &\leq \left(2 + \frac{1}{e}\right) \frac{1}{\nu^2} \bar{\varphi}_{n,M}(\theta^*). \end{aligned}$$

Thus, in both cases $\frac{k^*}{n} \log(1 + \frac{eM}{k^*}) \leq (2 + 1/e)\nu^{-2} \bar{\varphi}_{n,M}(\theta^*)$. Combining this with (A.6) we get (A.5). \square

A.2. Proof of Theorem 5.2. Applying the randomization scheme described in Nemirovski [(2000), page 211] we create from the sample y_1, \dots, y_n satisfying (5.2) two independent subsamples with “equivalent” sizes $\lceil n(1 - 1/\log \log n) \rceil$ and $n - \lceil n(1 - 1/\log \log n) \rceil$. We use the first subsample to construct the ES estimator and the soft thresholding estimator $\hat{\theta}^{\text{SOFT}}$, the latter attaining asymptotically the rate $\psi_n^{01}(\theta)$ for all $\theta \in \mathbb{R}^n$. We then use the second subsample to aggregate them, for example, as described in Nemirovski (2000). Then the aggregated estimator denoted by $\tilde{\theta}^*$ satisfies, for all $\theta \in \mathbb{R}^n$,

$$\begin{aligned} E_\theta |\tilde{\theta}^* - \theta|_2^2 &\leq \min\{E_\theta |\hat{\theta}^{\text{SOFT}} - \theta|_2^2, E_\theta |\tilde{\theta}^{\text{ES}} - \theta|_2^2\} + \frac{C \log \log n}{n} \\ &\leq \min(\psi_n^{01}(\theta), \psi_n^*(\theta))(1 + o(1)) + \frac{C \log \log n}{n}, \end{aligned}$$

where $C > 0$ is an absolute constant and $o(1) \rightarrow 0$ as $n \rightarrow \infty$ uniformly in $\theta \in \mathbb{R}^n$. Set

$$\tilde{\psi}_n(\theta) = \min(\psi_n^{01}(\theta), \psi_n^*(\theta)) + \frac{C \log \log n}{n}.$$

Then (5.10) follows immediately. Next, $\psi_n^{01}(\theta) \geq 2(\log n)/n$, so that for all $\theta \in \mathbb{R}^n$,

$$\frac{\psi_n^{01}(\theta)}{\tilde{\psi}_n(\theta)} \geq \frac{\psi_n^{01}(\theta)}{\psi_n^{01}(\theta) + C(\log \log n)/n} \geq \frac{2(\log n)/n}{2(\log n)/n + C(\log \log n)/n},$$

which implies (5.11). Finally, to prove (5.12) it is enough to notice that since $\psi_n^{01}(\theta) \geq 2(\log n)/n$,

$$\frac{\tilde{\psi}_n(\theta)}{\psi_n^{01}(\theta)} \leq \frac{\psi_n^*(\theta) + C(\log \log n)/n}{\psi_n^{01}(\theta)} \leq \frac{\psi_n^*(\theta)}{\psi_n^{01}(\theta)} + \frac{C \log \log n}{2 \log n}$$

and to use (5.8).

A.3. Proof of Theorem 5.3. Clearly (5.17) follows from (5.18) since in the latter η is fixed and equal to one particular function $\eta = f_\theta$.

We now prove (5.18). Let $\mathcal{H} = \{f_1, \dots, f_M\}$ be any dictionary in $\mathcal{D}(S \wedge \bar{m}, \kappa)$ with the corresponding $\underline{\kappa}$ and $\bar{\kappa}$ such that $\underline{\kappa}/\bar{\kappa} = \kappa$. For any $k \in \{1, \dots, M\}$, let Ω_k be the subset of $\mathcal{P} = \{0, 1\}^M$ defined by

$$(A.7) \quad \Omega_k := \{\mathbf{p} \in \mathcal{P} : |\mathbf{p}| = k\}.$$

We consider the class of functions

$$\mathcal{F}_k(\delta) := \left\{ f = \frac{\delta}{k} \tau \mathbf{f}_{\mathbf{p}} : \mathbf{p} \in \Omega_k \right\},$$

where $0 < \tau \leq 1$ will be chosen later. Note that functions in $\mathcal{F}_k(\delta)$ are of the form f_θ with $\theta \in \mathbb{R}_+^M \setminus \{0\}$, $M(\theta) = k$ and $|\theta|_1 = \tau \delta \leq \delta$. Thus, to prove (5.18), it is sufficient to show that, for any estimator T_n ,

$$(A.8) \quad \sup_{\eta \in \mathcal{G}} E_\eta \|T_n - \eta\|^2 > c_* \kappa \zeta_{n, M, \text{rk}(\mathbf{X})}(S, \delta)$$

for some subset $\mathcal{G} \subset \mathcal{F}_{\bar{S}}(\delta)$ where $\bar{S} = [S \wedge \bar{m} \wedge (M/2)]$ and $[\cdot]$ denotes the integer part. Note that $\bar{S} \geq 1$ since $M \geq 2$ and $S \wedge \bar{m} \geq 1$.

In what follows we will use the fact that for $f, g \in \mathcal{F}_{\bar{S}}(\delta)$ the difference $f - g$ is of the form f_θ with some $\theta \in \mathcal{P}_{2\bar{S}}$, so that in view of (5.14), $\|f - g\|^2$ is bracketed by the multiples of $|\theta|_2^2$ with this value of θ .

We now consider three cases, depending on the value of the integer m defined in (5.16).

Case (1): $m = 0$. Use Lemma A.3 to construct a subset $\mathcal{G}_{(1)} \subseteq \mathcal{F}_1(\delta) \subseteq \mathcal{F}_{\bar{S}}(\delta)$ with cardinality $s_{(1)} \geq (1 + eM)^{C_1}$ and such that

$$(A.9) \quad \|f - g\|^2 \geq \frac{\tau^2 \delta^2 \underline{\kappa}}{2} \quad \forall f, g \in \mathcal{G}_{(1)}, f \neq g.$$

Since $m = 0$, inequality (5.16) is violated for $m = 1$, so that

$$(A.10) \quad \delta^2 < \frac{\sigma^2}{n} \log(1 + eM) \leq \frac{\sigma^2}{nC_1} \log(s_{(1)}).$$

Case (2): $m \geq 1$, $S \wedge (M/2) \geq m$. Then $\bar{m} = m = \bar{S}$ and $m \leq M/2$, so that we have $\min(m, M - m) = m$, and Lemma A.3 guarantees that there exists $\mathcal{G}_{(2)} \subseteq \mathcal{F}_m(\delta) = \mathcal{F}_{\bar{S}}(\delta)$ with cardinality $s_{(2)} \geq (1 + eM/m)^{C_1 m}$ and such that

$$\|f - g\|^2 \geq \frac{\tau^2 \delta^2 \kappa}{4m} \quad \forall f, g \in \mathcal{G}_{(2)}, f \neq g.$$

To bound from below the quantity δ^2/m , observe that from the definition of m we have

$$(A.11) \quad \frac{\delta^2}{m} \geq \frac{\delta\sigma}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eM}{m}\right)} \geq \frac{\delta\sigma}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eM\sigma}{\delta\sqrt{n}}\right)}.$$

The previous two displays yield

$$(A.12) \quad \|f - g\|^2 \geq \frac{\tau^2 \kappa \delta \sigma}{4\sqrt{n}} \sqrt{\log\left(1 + \frac{eM\sigma}{\delta\sqrt{n}}\right)}.$$

Note that in this case

$$m + 1 > \frac{\delta\sqrt{n}}{\sigma\sqrt{\log(1 + eM/(m + 1))}},$$

so that

$$(A.13) \quad \begin{aligned} \frac{\delta^2}{m} &\leq \frac{2\delta^2}{m + 1} < 2(m + 1) \frac{\sigma^2}{n} \log\left(1 + \frac{eM}{m + 1}\right) \\ &\leq 4m \frac{\sigma^2}{n} \log\left(1 + \frac{eM}{m}\right) \leq \frac{4\sigma^2}{nC_1} \log(s_{(2)}). \end{aligned}$$

Case (3): $m \geq 1$, $S \wedge (M/2) < m$. Then $\bar{S} = [S \wedge (M/2)] < m$. Moreover, we have $\min(\bar{S}, M - \bar{S}) = \bar{S}$ and using Lemma A.3, for any positive $\bar{\delta} \leq \delta$ we can construct $\mathcal{G}_{(3)} \subseteq \mathcal{F}_{\bar{S}}(\bar{\delta})$ with cardinality $s_{(3)} \geq (1 + eM/\bar{S})^{C_1 \bar{S}}$ and such that

$$\|f - g\|^2 \geq \frac{\tau^2 \bar{\delta}^2 \kappa}{4\bar{S}} \quad \forall f, g \in \mathcal{G}_{(3)}, f \neq g.$$

Take

$$\bar{\delta} := \sigma \frac{\bar{S}}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eM}{\bar{S}}\right)} \leq \sigma \frac{m}{\sqrt{n}} \sqrt{\log\left(1 + \frac{eM}{m}\right)} \leq \delta,$$

where, in the last inequality, we used the definition of m . Next, note that $\bar{S} = [S \wedge (M/2)] \geq S/4$ since $M \geq 2$. Then

$$(A.14) \quad \|f - g\|^2 \geq \frac{\tau^2 \kappa \sigma^2 \bar{S}}{4n} \log\left(1 + \frac{eM}{\bar{S}}\right) \geq \frac{\tau^2 \kappa \sigma^2 S}{16n} \log\left(1 + \frac{4eM}{S}\right).$$

In addition, we have

$$(A.15) \quad \frac{\bar{\delta}^2}{\bar{S}} = \bar{S} \frac{\sigma^2}{n} \log\left(1 + \frac{eM}{\bar{S}}\right) \leq \frac{\sigma^2}{nC_1} \log(s_{(3)}).$$

Since the random variables $\xi_i, i = 1, \dots, n$, are i.i.d. Gaussian $\mathcal{N}(0, \sigma^2)$, for any $f, g \in \mathcal{G}_{(j)}, j \in \{1, 2, 3\}$, the Kullback–Leibler divergence $\mathcal{K}(P_f, P_g)$ between P_f and P_g is given by

$$\mathcal{K}(P_f, P_g) = \frac{n}{2\sigma^2} \|f - g\|^2 \leq \frac{n\tau^2 \delta_{(j)}^2 \bar{\kappa}}{\sigma^2 k_{(j)}},$$

where $\delta_{(1)} = \delta_{(2)} = \delta, \delta_{(3)} = \bar{\delta}, k_{(1)} = 1, k_{(2)} = m, k_{(3)} = \bar{S}$. Using, respectively, (A.10) in case (1), (A.13) in case (2) and (A.15) in case (3), and choosing $\tau^2 = \min(C_1/(32\bar{\kappa}), 1)$ (note that we need $\tau \leq 1$ by construction) we obtain

$$(A.16) \quad \mathcal{K}(P_f, P_g) \leq \frac{4\tau^2 \bar{\kappa}}{C_1} \log s_{(j)} \leq \frac{\log s_{(j)}}{8} \quad \forall f, g \in \mathcal{G}_{(j)}, j = 1, 2, 3.$$

Combining (A.9), (A.12) and (A.14) together with (A.16), we find that the conditions of Theorem 2.7 in Tsybakov (2009) are satisfied and use it to obtain (A.8).

A.4. A lemma for minimax lower bound. Here we give a result related to subset extraction, which is a generalization of the Varshamov–Gilbert lemma used to prove minimax lower bounds [see, e.g., a recent survey in Tsybakov (2009), Chapter 2]. For any $M \geq 1, k \in \{1, \dots, M - 1\}$, let Ω_k^M be the subset of $\{0, 1\}^M$ defined by

$$\Omega_k^M := \left\{ \omega \in \{0, 1\}^M : \sum_{j=1}^M \omega_j = k \right\}.$$

The next lemma belongs to Reynaud-Bouret (2003). The version presented here can be found in Rigollet and Tsybakov (2010), Lemma 8.3.

LEMMA A.3. *Let $M \geq 2$ and $1 \leq k \leq M$ be two integers and define $\bar{k} = \min(k, M - k)$. Then there exists a subset Ω of Ω_k^M such that the Hamming distance $\rho(\omega, \omega') = \sum_{j=1}^M \mathbb{1}(\omega_j \neq \omega'_j)$ satisfies*

$$\rho(\omega, \omega') \geq \frac{\bar{k} + 1}{4} \quad \forall \omega, \omega' \in \Omega : \omega \neq \omega',$$

and $s = \text{card}(\Omega)$ satisfies

$$\log(s) \geq C_1 \bar{k} \log\left(1 + \frac{eM}{\bar{k}}\right)$$

for some numerical constant $C_1 > 0$.

Acknowledgment. We would like to thank Victor Chernozhukov and Joseph Salmon for helpful discussions of the paper.

REFERENCES

- ABRAMOVICH, F., BENJAMINI, Y., DONOHO, D. L. and JOHNSTONE, I. M. (2006). Adapting to unknown sparsity by controlling the false discovery rate. *Ann. Statist.* **34** 584–653. [MR2281879](#)
- ALQUIER, P. and LOUNICI, K. (2010). PAC-Bayesian bounds for sparse regression estimation with exponential weights. HAL.
- BARANIUK, R., DAVENPORT, M., DEVORE, R. and WAKIN, M. (2008). A simple proof of the restricted isometry property for random matrices. *Constr. Approx.* **28** 253–263. [MR2453366](#)
- BARRON, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Trans. Inform. Theory* **39** 930–945. [MR1237720](#)
- BICKEL, P., RITOV, Y. and TSYBAKOV, A. (2008). Hierarchical selection of variables in sparse high-dimensional regression. Available at [ArXiv:0801.1158](#).
- BICKEL, P. J. and DOKSUM, K. A. (2006). *Mathematical Statistics: Basic Ideas and Selected Topics, Vol. 1*, 2nd ed. Prentice-Hall, Upper Saddle River, NJ.
- BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. [MR2533469](#)
- BÜHLMANN, P. and VAN DE GEER, S. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Statist.* **3** 1360–1392. [MR2576316](#)
- BUNEA, F., TSYBAKOV, A. and WEGKAMP, M. (2007a). Sparsity oracle inequalities for the Lasso. *Electron. J. Statist.* **1** 169–194 (electronic). [MR2312149](#)
- BUNEA, F., TSYBAKOV, A. B. and WEGKAMP, M. H. (2007b). Aggregation for Gaussian regression. *Ann. Statist.* **35** 1674–1697. [MR2351101](#)
- CANDES, E. (2008). The restricted isometry property and its implications for compressed sensing. *C. R. Math. Acad. Sci. Paris* **346** 589–592. [MR2412803](#)
- CANDES, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35** 2313–2351. [MR2382644](#)
- DALALYAN, A. and TSYBAKOV, A. (2008). Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity. *Machine Learning* **72** 39–61.
- DALALYAN, A. and TSYBAKOV, A. B. (2010). Mirror averaging with sparsity priors. Available at [ArXiv.org:1003.1189](#).
- DALALYAN, A. S. and TSYBAKOV, A. B. (2007). Aggregation by exponential weighting and sharp oracle inequalities. In *Learning Theory. Lecture Notes in Computer Science* **4539** 97–111. Springer, Berlin. [MR2397581](#)
- DALALYAN, A. S. and TSYBAKOV, A. B. (2009). Sparse regression learning by aggregation and Langevin Monte Carlo. Available at [ArXiv:0903.1223](#).
- DONOHO, D. L. and JOHNSTONE, I. M. (1994a). Ideal spatial adaptation by wavelet shrinkage. *Biometrika* **81** 425–455. [MR1311089](#)
- DONOHO, D. L. and JOHNSTONE, I. M. (1994b). Minimax risk over l_p -balls for l_q -error. *Probab. Theory Related Fields* **99** 277–303. [MR1278886](#)

- DONOHO, D. L., JOHNSTONE, I. M., HOCH, J. C. and STERN, A. S. (1992). Maximum entropy and the nearly black object. *J. Roy. Statist. Soc. Ser. B* **54** 41–81. With discussion and a reply by the authors. [MR1157714](#)
- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- FOSTER, D. P. and GEORGE, E. I. (1994). The risk inflation criterion for multiple regression. *Ann. Statist.* **22** 1947–1975. [MR1329177](#)
- GEORGE, E. I. (1986a). Combining minimax shrinkage estimators. *J. Amer. Statist. Assoc.* **81** 437–445. [MR0845881](#)
- GEORGE, E. I. (1986b). Minimax multiple shrinkage estimation. *Ann. Statist.* **14** 188–205. [MR0829562](#)
- GIRAUD, C. (2008). Mixing least-squares estimators when the variance is unknown. *Bernoulli* **14** 1089–1107. [MR2543587](#)
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York. [MR1851606](#)
- KOH, K., KIM, S.-J. and BOYD, S. (2008). l1_ls: A Matlab solver for l1-regularized least squares problems. BETA version. Stanford Univ. Available at http://www.stanford.edu/~boyd/l1_ls.
- KOLTCHINSKII, V. (2010). Oracle inequalities in empirical risk minimization and sparse recovery problems. *St. Flour Lecture Notes*. To appear.
- KOLTCHINSKII, V. (2009a). The Dantzig selector and sparsity oracle inequalities. *Bernoulli* **15** 799–828. [MR2555200](#)
- KOLTCHINSKII, V. (2009b). Sparsity in penalized empirical risk minimization. *Ann. Inst. H. Poincaré Probab. Statist.* **45** 7–57. [MR2500227](#)
- LECUN, Y., BOSER, B., DENKER, J. S., HENDERSON, D., HOWARD, R. E., HUBBARD, W. and JACKEL, L. D. (1990). Handwritten digit recognition with a back-propagation network. In *Advances in Neural Information Processing Systems* 396–404. Morgan Kaufmann, San Francisco.
- LEUNG, G. and BARRON, A. R. (2006). Information theory and mixing least-squares regressions. *IEEE Trans. Inform. Theory* **52** 3396–3410. [MR2242356](#)
- LOUNICI, K. (2007). Generalized mirror averaging and D -convex aggregation. *Math. Methods Statist.* **16** 246–259. [MR2356820](#)
- NEMIROVSKI, A. (2000). Topics in non-parametric statistics. In *Lectures on Probability Theory and Statistics (Saint-Flour, 1998)*. *Lecture Notes in Math.* **1738** 85–277. Springer, Berlin. [MR1775640](#)
- RASKUTTI, G., WAINWRIGHT, M. J. and YU, B. (2009). Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. Available at [ArXiv:0910.2042](#).
- REYNAUD-BOURET, P. (2003). Adaptive estimation of the intensity of inhomogeneous Poisson processes via concentration inequalities. *Probab. Theory Related Fields* **126** 103–153.
- RIGOLLET, P. (2009). Maximum likelihood aggregation and misspecified generalized linear models. Technical report. Available at [ArXiv:0911.2919](#).
- RIGOLLET, P. and TSYBAKOV, A. (2010). Exponential screening and optimal rates of sparse estimation. Available at [ArXiv:1003.2654v3](#).
- ROBERT, C. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*. Springer, New York.
- TSYBAKOV, A. B. (2003). Optimal rates of aggregation. In *COLT* (B. Schölkopf and M. K. Warmuth, eds.). *Lecture Notes in Computer Science* **2777** 303–313. Springer, Berlin.
- TSYBAKOV, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer, New York. [MR2724359](#)
- VAN DE GEER, S. A. (2008). High-dimensional generalized linear models and the Lasso. *Ann. Statist.* **36** 614–645. [MR2396809](#)
- WRIGHT, J., YANG, A. Y., GANESH, A., SASTRY, S. S. and MA, Y. (2009). Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **31** 210–227.
- ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. [MR2604701](#)

- ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the LASSO selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. MR2435448
- ZHANG, C.-H. and MELNIK, O. (2009). Plus: Penalized linear unbiased selection. R package version 0.8. CRAN. Available at <http://CRAN.R-project.org/package=plus>.
- ZHANG, T. (2009). Some sharp performance bounds for least squares regression with L_1 regularization. *Ann. Statist.* **37** 2109–2144. MR2543687

DEPARTMENT OF OPERATIONS RESEARCH
AND FINANCIAL ENGINEERING
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: rigollet@princeton.edu

LABORATOIRE DE STATISTIQUE
CREST
3, AV. PIERRE LAROUSSE
F-92240 MALAKOFF CEDEX
FRANCE
E-MAIL: alexandre.tsybakov@upmc.fr