Learning Out of Leaders

Gerard Kerkyacharian, Mathilde Mougeot, Dominique Picard, Karine
Tribouley

HAL Id: hal-00445690

https://hal.science/hal-00445690

Preprint submitted on 11 Jan 2010

Learning Out of Leaders

Gérard KERKYACHARIAN[a], Mathilde MOUGEOT[a],

Dominique PICARD[a] and Karine TRIBOULEY[b]

[a] Université Paris-Diderot, CNRS LPMA, 175 rue du Chevaleret
75013 Paris, France
kerkmath.jussieu.fr, mougeotmath.jussieu.fr, picardmath.jussieu.fr

[b] Université Nanterre Paris Ouest, MODALX & LPMA
Rue de la république
92001 Nanterre, France
karine.tribouleyu-paris10.fr

Résumé

This paper investigates the problem of selection and estimation in a high dimensional regression-type model. We propose a procedure with no optimization called LOL, for Learning Out of Leaders. LOL is an auto-driven algorithm with two thresholding steps. A first adaptive thresholding helps to select leaders among the initial regressors in such a way to reduce the dimensionality. Then a second thresholding follows the estimations and predictions performed by linear regression on the leaders. Theoretical results are proved. As an estimation procedure, LOL is optimal since the upper exponential bounds are achieved. Rates of convergence are provided and show that LOL is also consistent as a selection procedure. An extensive computational experiment is conducted to emphasize the practical good performances of LOL.

1   Introduction

The general linear model is considered in this paper, with a focus on cases where the number p of regressors is large compared to the number n of the observations (although there is no such restrictions). These type of models have lots of practical applications in many areas of science and engineering including collaborative filtering, machine learning, control, remote sensing, and computer vision just to name a few. Examples in statistical signal processing and nonparametric estimation include the recovery of a continuous-time curve or surface from a finite number of noisy samples. Other interesting fields of application are radiology and biomedical imaging when fewer measurements about an image are available compared to the unknown number of pixels collected. In biostatistics, high dimensional data frequently arise in genomics to study gene expression given a huge number of initial genes and a relatively low number of observations.

A considerable amount of work have been produced in this domain in the last years, which has been a large source of inspiration for this paper. We have especially considered the algorithms coming from the learning framework ([Barron et al., 2008], [Binev et al., 2005], [Binev et al., 2007a], [Binev et al., 2007b]), as well as the extraordinary explosive domain of ℓ_1 penalties (among many others [Tibshirani, 1996], [Candes and Tao, 2007], [Bickel et al., 2007], [Bunea et al., 2007a], [Bunea et al., 2007b], [Fan and Lv, 2008] and [Candès and Plan, 2008]. See also [Lounici, 2008] and [Alquier and Hebiri, 2009]).

The essential motivation of this work is to provide one of the simplest procedures which achieves, in the same time, good performances. LOL algorithm (for Learning Out of Leaders) consists in a two steps thresholding procedure. As we do not perform any optimization step, it is important to address in which domains the procedure is competitive to more sophisticated algorithms and especially to algorithms performing a one or two steps ℓ_1 minimization. One of our aim is to delimit where LOL is performant and where its simplicity induces a slight lack of efficiency from a theoretical point of view as from a practical aspect.

Let us start by introducing the ideas of the emergence of LOL algorithm. This simple procedure can be viewed as an 'explanation' or as a 'cartoon' of ℓ_1 minimizations. It is well known that when the regressors are exactly normalized and orthogonal, ℓ_1 minimization corresponds to soft thresholding which itself is close to hard thresholding. Hence, it is quite natural to expect that thresholding should perform well, at least in cases not so far from these orthonormal conditions which correspond, as noted below, to small coherence conditions. A tricky problem occurs when the regressors are not orthonormal or when the number of regressors is large. Then, the minimum least squares estimator has a non unique solution and is very unstable. This stays the main difficulty for the ℓ_1 minimizers or more generally for all methods based on sparsity assumptions. Moreover, this is the part of the algorithms where the computation cost shows up. Obviously a simple thresholding would not fit, but assuming some sparsity conditions, in this case, ensures that it is possible to choose some regressors and exclude some others. LOL algorithm solves the difficult problem of the choice of the regressors in a quite crude way by adaptively selecting N regressors which are the most correlated to the target : this defines the first step thresholding of LOL, determining the N leaders. The number N is chosen using a fine tuning depending on the coherence, and it has to be emphasized that the choice is auto driven. In a second thresholding step, LOL regresses on the previous leaders and thresholds the result to take into account the noise of the model.

Properties of LOL procedure are investigated through two different points of view : the prediction problem and the estimation problem. More precisely, it is established that LOL procedure has a prediction error which is going to zero in probability with exponential rates. These types of

2

results are often called Bahadur type efficiency. Although Bahadur efficiency of test and estimation procedures goes back to the sixties (see [Bahadur, 1960]), it has seen recently a revival in learning theory, where the rates of convergence (preferably exponential) of being at some fixed distance of the target are investigated and compared to optimality. This is also the connection to learning theory which guides here the choice to measure LOL performances as the mean of the empirical quadratic distance between the observations and the predicted values. We also establish that LOL procedure works quite well regarding the detection since the number of false negative as well as false positive are going to zero in probability with pretty fast rates.

Of course, because of the simplicity of the method, some loss of efficiency can be expected compare to more elaborate and costly procedures. But even when there is a loss, the limitations of the procedure could be an interesting information on the ℓ_1 minimizers themselves. From both a theoretical and a practical point of view, when the coherence is small, LOL procedure is as powerful as the best procedures. Also when there is a loss in the rate, a positive aspect of the method is that the practitioner is informed of the possible instability since the coherence is provided by the observations. An intensive calculation program is performed to show the advantages and the limitations of LOL procedure in several practical aspects. In Section 6, the case where the regressors are forming a random design matrix with i.i.d. entries is investigated. Different laws of the entries are considered (Gaussian, Uniform, Bernoulli or Student laws) inducing specific coherence for the design matrix. Several interesting features are discussed in this section. The impact of the sparsity and the undetermination of the regression on LOL performances are studied. A comparison with two others two-step procedures namely [Fan and Lv, 2008] and [Candès and Plan, 2008] is also provided and shows the additional benefits brought by LOL. The most interesting conclusion being that the practical results are even better and more comforting than the theoretical ones in the sense that even when the coherence is pretty high, LOL procedure shows good performances.

The paper is organized as follows. In Section 2, the general model and the notations are presented. In Section 3, LOL procedure is detailed as other procedures with a ℓ_1 optimization step ; practical comparisons with other procedures are later discussed in Section 5. In Section 4, after stating the hypotheses needed in the model, theoretical results are established. Practical performances of the LOL procedure are investigated in Section 6 and the proofs are detailed in Section ??.

2   Description of the models

In this part, the model of interest is presented with a focus on two specific cases : the random matrices design and the functional regression.

2.1 General model

A Gaussian (or sub-gaussian) high dimensional linear model is here considered and more precisely data $Y = (Y_1, \ldots, Y_n)^t$ are observed coming from the following regression model

$$Y = \Phi\alpha + u + \varepsilon \tag{1}$$

where the parameter $\alpha \in \mathbb{R}^p$ is the unknown vector to be estimated and

- the vector $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n)^t$ is a (non observed) vector of random errors. It is assumed to be independent Gaussian variables $N(0, \sigma^2)$ but essentially comparable results can be obtained in the case of zero mean subgaussian errors (see the remark before Lemma **??**).
- the vector $u = (u_1, \ldots, u_n)^t$ is a non observed vector of (possibly) random errors. Its amplitude is assumed to be small. The differences between the two previously described "errors" lies in the fact that the ε_i's are centered but unbounded and independent, while the u_i's are only bounded. The importance of introducing these two types of errors becomes clear in the functional regression example (see section 2.3).
- Φ is a $n \times p$ known matrix. This paper focuses on the interesting case where $p >> n$ but it is not necessary. We assume that Φ has normalized columns (or normalize them) in the following sense :

$$\frac{1}{n}\sum_{i=1}^n \Phi_{i\ell}^2 = 1, \quad \forall \, \ell = 1 \ldots, p. \tag{2}$$

2.2 Coherence

The following Gram $p \times p$ matrix is

$$M := \frac{1}{n}\Phi^t\Phi.$$

The quantity

$$\tau_n = \sup_{\ell \neq m} |M_{\ell m}| = \sup_{\ell \neq m} |\frac{1}{n}\sum_{i=1}^n \Phi_{i\ell}\Phi_{im}|$$

is called the coherence of the matrix M. This quantity is important because it induces a bound on the size of the invertible matrices built with the columns of M. More precisely, fix $0 < \nu < 1$ and let $\mathcal{C}$ be a subset of indices of $\{1, \ldots, p\}$ with cardinality m. Denote $\Phi_\mathcal{C}$ the matrix restricted to the columns of Φ whose indices are in $\mathcal{C}$. If $2\tau_n \leq \nu$, the associated Gram matrix

$$M(\mathcal{C}) := \frac{1}{n}\Phi_\mathcal{C}^t\Phi_\mathcal{C}$$

4

is almost diagonal as soon as m is smaller than $\lfloor \nu/\tau_n \rfloor$ in the sense that it satisfies the following so called Restricted Isometry Property (RIP).

$$\forall x \in \mathbb{R}^m, \ \|x\|_{l_2(m)}^2(1 - \nu) \leq x^t \mathcal{M}(\mathcal{C})x \leq \|x\|_{l_2(m)}^2(1 + \nu), \tag{3}$$

This proves in particular that the matrix $M(\mathcal{C})$ is invertible.

2.3 Models of interest

Although these results apply in the general case, two typical cases of applications are especially considered.

The first application concerns a random matrix Φ composed of n independent random vectors of size p. The important role played by the distribution of these random vectors is detailed in the simulation study, Section 6.

The second application is the learning (also called functional regression) framework

$$Y_i = f(X_i) + \varepsilon_i, \ i = 1 \ldots n \tag{4}$$

where f is the functional parameter of interest to be estimated. The $X_i's$ are i.i.d. random variables living in a compact domain of $\mathbb{R}^d$. The errors $\varepsilon_i's$, are i.i.d. standard Gaussian random variables and independent of the $X_i's$ (or centered sub-gaussian variables). ρ denotes the common (unknown) law of the $(1+d)-$vectors $Z_i = (X_i, Y_i)'s$.

To relate this framework to our model, let us consider a dictionary $\mathcal{D}$ of size p, of real functions defined on $\mathbb{R}^d$. Assume that f can be reasonably well approximated using the elements of the dictionary which means that there exists a sequence $\{\alpha_g, g \in \mathcal{D}\}$ such that

$$f = \sum_{g \in \mathcal{D}} \alpha_g g + h$$

where h is hopefully small. Then the regression model becomes

$$Y_i = \sum_{g \in \mathcal{D}} \alpha_g g(X_i) + h(X_i) + \varepsilon_i, \ i = 1, \ldots, n$$

which coincides with the general model

$$Y = \Phi\alpha + u + \varepsilon,$$

setting $u_i = h(X_i)$ for any $i = 1, \ldots, n$ and Φ being the matrix with general terms $\Phi_{i\ell} = g_\ell(X_i)$ (after choosing an enumeration of $\mathcal{D}$). Again, the dictionary has to be normalized and (2) translates here as

$$\frac{1}{n}\sum_{i=1}^n g^2(X_i) = 1, \ \forall \, g \in \mathcal{D}.$$

3 Estimation procedures

As explained in the introduction, the essential motivation of this work is to provide one of the simplest procedures, finding its inspiration among a lot of works around the same theme. First, the estimation of the unknown parameter α using LOL is described. The procedure has the particularity to perform a selection method of the regressors in the same time. Next, a short review on the procedures directly connected to LOL is proposed.

Once for all, the constant ν is fixed. This constant will obviously be related to the precision of LOL main procedure (for instance $\nu = 1/2$ can be taken as default value).

3.1 LOL Procedure

Once τ_n (or a bound for τ_n) is evaluated and $N = \lfloor \nu/\tau_n \rfloor$ is computed, LOL procedure has three successive steps : Find N leaders, Regress on the leaders, Threshold.

1. Find the leaders :
   - For some constant $T_1 > 0$, fix a threshold
   $$\lambda_n(1) = T_1 \left(\left(\frac{\log p}{n} \right)^{1/2} \vee \tau_n \right). \tag{5}$$
   - Compute the 'correlations'
   $$K_\ell = |\frac{1}{n} \sum_{i=1}^{n} \Phi_{i\ell} Y_i|$$
   and consider the ordered sequence $K_{(1)} \geq K_{(2)} \geq \ldots \geq K_{(N)}$ of the N largest, and the associated set of indices $\mathcal{K} = \{\kappa_{(1)}, \kappa_{(2)}, \ldots, \kappa_{(N)}\}$.
   - The final set of the leaders is defined by the following set of columns Φ_ℓ of the matrix Φ :
   $$B = \{\Phi_\ell, \ \ell \in \mathcal{K} \text{ and } K_\ell \geq \lambda_n(1)\}$$
   and $\mathcal{B}$ denotes the set of their indices (which might of course be different from $\mathcal{K}$). It is clear from this construction that N appears as a bound for the number of leaders (equal to the cardinal of $\mathcal{B}$).

2. Regress on the leaders :
   - Consider the pseudo-regression model :
   $$Y_i = \sum_{\ell \in \mathcal{B}} \Phi_{i\ell} \alpha_\ell + e_i$$
   and define the extracted matrix $\Phi_\mathcal{B}$ by
   $$(\Phi_\mathcal{B})_{\ell, i} = \Phi_{i\ell} \quad \text{for any } \ell \in \mathcal{B} \text{ and } i \in \{1, \ldots, n\}. \tag{6}$$

- Let $\widehat{\alpha}(\mathcal{B}) = (\widehat{\alpha_\ell}(\mathcal{B}), \ \ell \in \mathcal{B})$ be the minimum least square error in this model :

$$\hat{\alpha}(\mathcal{B}) = \mathrm{Arg} \min_{\alpha = (\alpha_\ell)_{\ell \in \mathcal{B}}} \left(\sum_{i=1}^{n} (Y_i - \sum_{\ell \in \mathcal{B}} \Phi_{i\ell} \alpha_\ell)^2 \right) = (\Phi_{\mathcal{B}}^t \Phi_{\mathcal{B}})^{-1} \Phi_{\mathcal{B}}^t Y .$$

- Define the vector $\hat{\alpha}$ of $\mathbb{R}^p$ by

$$\widehat{\alpha_\ell} := \left\{ \begin{array}{ll} \widehat{\alpha_\ell}(\mathcal{B}) & \text{if} \quad \ell \in \mathcal{B} \\ 0 & \text{if} \quad \ell \notin \mathcal{B} \end{array} \right.$$

3. Threshold :

   For some constant $T_2 > 0$, fix a threshold

$$\lambda_n(2) = T_2 \left(\frac{\log n}{n} \right)^{1/2} \tag{7}$$

   and threshold again the estimated coefficients to obtain the final predictor $\hat{\alpha}^*$ whose coordinates are

$$\hat{\alpha}_\ell^* = \widehat{\alpha_\ell} \, I\{|\widehat{\alpha_\ell}| \geq \lambda_n(2)\}.$$

   The selected regressors are then the columns of Φ whose indices belong to

$$\mathcal{L} = \{\ell = 1, \ldots, p, \ \widehat{\alpha_\ell}^* \neq 0\}$$

Notice that the formula (5) and (7) are the 'default' values for the tuning sequences $\lambda_n(1)$ and $\lambda_n(2)$ given for the procedure. However, the presentation as well as the theoretical results in sequel are given for arbitrary sequences $\lambda_n(1)$ and $\lambda_n(2)$.

3.2 Several inspirations

Although it is impossible to be exhaustive in such a productive domain, some of the works directly in relation to our construction are hereafter mentioned. We apologize in advance for all the works that are not mentioned but still in connection. For a comprehensive overview, we refer to [Fan and Lv, 2009].

In the context of the learning theory (second application), various methods are already been proposed, including kernel methods and search within dictionaries. Let us especially mention following works providing greedy algorithms [Barron et al., 2008], or adjusting tree algorithms [Binev et al., 2005], [Binev et al., 2007a]. A one step algorithm rough version of LOL is given in [Kerkyacharian and Picard, 2007] as well as in [Kerkyacharian et al., 2009] for the case $p \leq n$. [Bunea, 2009] also proposes an estimation procedure based on the lasso and derives a selection procedure by keeping the non zero estimated coefficients.

In the context of the linear regression (first application), several authors propose procedures to solve the selection problem and the estimation problem in the case where the vector α has only a small number of non zero components, and (often) when the design matrix Φ is composed of i.i.d. random vectors : see among many others [Tibshirani, 1996], [Candes and Tao, 2007], [Bickel et al., 2007], [Bunea et al., 2007a] and [Bunea et al., 2007b].

We especially refer to the 2-steps procedures which are also commonly used. Apparently, as soon as in 1959 such a procedure is already discussed (see [Satterthwaite, 1959]). In [Candes and Tao, 2007] and [Candès and Plan, 2008], the leaders are selected with (respectively) the Danzig procedure and the lasso procedure. Then, the estimated coefficients are obtained via a linear regression on the leaders. Using an intensive simulation program, [Fan and Lv, 2008] show that it could be unfavorable to use the procedures lasso or Danzig *before* the reduction of the dimension. They also provide a search among leaders called Sure Independence Screening (SIS) procedure. This procedure is very close to the one discussed in this paper : the leaders are the $N = \lfloor \gamma_n n \rfloor$ columns of Φ with largest correlations to the target variable Y (γ_n is a tuning sequence tending to zero). This step is followed with a subsequent estimation procedure using Danzig or lasso. All these methods focus on the complexity of the algorithms.

4   Main theoretical results

This section states the theoretical results of the procedure LOL. First, the assumptions on the model are described. Next, the quantities allowing to measure the performances of the procedures are defined. The consistency of LOL is shown using two different points of view : the prediction problem and the estimation problem.

4.1   Sparsity conditions on the model

Recall that the model specifies a gaussian (or sub-gaussian) observation of the following form : $Y = \Phi \alpha + u + \varepsilon$. The following sparsity conditions are assumed. There exist $S \leq N$ and constants

M, c_1, c_t, c'_t, c_0, such that the sequences $\{\alpha_\ell\}_{\ell \leq p}$ and $\{u_i\}_{i \leq n}$ satisfy the following conditions

$$\sup_{i=1,\ldots,n} |u_i| \leq c_1 \left(\frac{S}{n}\right)^{1/2} \tag{8}$$

$$\sum_{\ell=1}^{p} |\alpha_\ell| \leq M, \tag{9}$$

$$\# \left\{\ell \in \{1,\ldots,p\}, \ |\alpha_\ell| \geq \lambda_n(2)/2\right\} \leq S \tag{10}$$

$$\sum_{(\ell)>N} |\alpha_{(\ell)}| \leq c_t \left(\frac{S}{n\tau_n}\right)^{1/2} \tag{11}$$

$$\sum_{(\ell)>N} |\alpha_{(\ell)}|^2 \leq c_t'^2 \frac{S}{n} \tag{12}$$

$$\sum_{\ell=1}^{p} |\alpha_\ell|^2 \, I\{|\alpha_\ell| \leq 2\lambda_n(1)\} \leq c_0^2 \frac{S}{n} \tag{13}$$

Recall that $(\alpha_{(\ell)})$ is the ordered sequence (for the modulus) $|\alpha_{(1)}| \geq |\alpha_{(2)}| \geq \ldots |\alpha_{(p)}|$. For $S, M > 0$, $V(S, M)$ denotes the class of models of type (1) satisfying the sparsity conditions (9), (8), (10), (11), (13).

A very important example of such a class occurs when all the coefficients of α are 0 except S coefficients (with $S \leq N$) with a modulus greater that $\lambda_n(2)/2$ but bounded : $\mathcal{S}pars(S, M)$ denotes such a class.

The conditions (9)–(13) are also satisfied if the l_q conditions are assumed, as in [Raskutti et al., 2009] which provide upper and lower bounds. More precisely, for $q \in (0, 1]$, define the l_q-balls as the sets

$$B_q(M) := \{\alpha \in \mathbb{R}^p, \ \sum_{j=1}^{p} |\alpha_j|^q \leq M^q\}. \tag{14}$$

It is not difficult to prove that if α belongs to $B_q(M)$ then (9)–(13) are verified for

$$S \geq \lambda_n(2)^{-q} \vee n\lambda_n(1)^{2-q} \vee n\tau_n^{(2-q)/q}.$$

In particular, in order to compare our results to the lower bounds in [Raskutti et al., 2009], it is important to verify that the conditions are verified for $\tau_n = O\left(\sqrt{\frac{\log p}{n}}\right)$ and for the defaults values for $\lambda_n(1)$ and $\lambda_n(2)$. In this precise case, this means that S/n has to be of order τ_n^{2-q}.

In the context of the learning theory (second application), the sparsity conditions are required on the target function f. The above assumptions are easily translated by replacing the condition (8) by the following one :

$$\|h\|_\infty^2 \leq c_1 \left(\frac{S}{n}\right)^{1/2}.$$

The other conditions are quite usual in functional analysis and relate to Lorentz spaces.

4.2 Measures of performances

First, let us define loss functions to measure the difference between the true value $\alpha \in \mathbb{R}^p$ and the result $\hat{\alpha}^*$ of LOL procedure. Denote Φ_i the $i-$th line of the matrix Φ and recall that the $i-$th observation is given by the model

$$Y_i = \Phi_i \alpha + u_i + \varepsilon_i.$$

The predicted $i-$th observation is $\widehat{Y}_i = \Phi_i \hat{\alpha}^*$. The empirical quadratic distance between the predicted observations and the expected value is here considered

$$d(\hat{\alpha}^*, \alpha)^2 = \frac{1}{n} \sum_{i=1}^{n} \left(\widehat{Y}_i - \mathbb{E}Y_i \right)^2 = \frac{1}{n} \sum_{i=1}^{n} \left(\sum_{\ell=1}^{p} (\hat{\alpha}_\ell^* - \alpha_\ell) \Phi_{i\ell} + u_i \right)^2. \tag{15}$$

Notice that in the functional regression case, this error coincides with the L_2 error with respect to the empirical measure

$$\hat{\rho} = \frac{1}{n} \sum \delta_{X_i}$$

where δ_x denotes the Dirac measure at point x. Indeed, we get

$$d(\hat{\alpha}^*, \alpha)^2 = \frac{1}{n} \sum_{i=1}^{n} \left(\hat{f}(X_i) - f(X_i) \right)^2 = \|\hat{f} - f\|_{\hat{\rho}}^2.$$

With a slight abuse of notations, we also write the distance defined in (15) in the general model

$$d(\hat{\alpha}^*, \alpha) := \| \sum_{l=1}^{p} (\hat{\alpha}_\ell^* - \alpha_\ell) \Phi_{\bullet\ell} + u_\bullet \|_{\hat{\rho}}$$

where $\Phi_{\bullet\ell}$ is the $\ell-$th row of Φ.

The first measure of performance under consideration is issued from the Bahadur efficiency of test and estimation procedures and is defined for any tolerance $\eta > 0$ as

$$AC_n(LOL, \eta) = P\left(d(\hat{\alpha}^*, \alpha) > \eta \right). \tag{16}$$

Obviously, if the tolerance is low (smaller than a critical value η_n), this quantity is large. In the opposite, for $\eta \geq \eta_n$, the quality of the procedure is given by the rate of convergence of $AC_n(LOL, \eta)$ towards zero. Observe that the value of the critical value η_n is essential since it yields, as a consequence, bounds for $Ed(\hat{\alpha}^*, \alpha)$ which is another (more standard) measure of performance of the procedure.

More generally, in the learning framework, given priors Θ on the class of probability distributions generating the observations, it has been defined in [DeVore et al., 2006] the accuracy confidence function of the procedure $\hat{f}$:

$$AC_n(\Theta, \hat{f}, \eta) := \sup_{\rho \in \Theta} \rho^{\otimes n} \{ \|f - \hat{f}\|_{\rho_X} > \eta \}. \tag{17}$$

10

This quantity measures a uniform confidence (over the class Θ) that the estimator $\hat{f}$ is accurate to the tolerance η. In most examples, there exist a phase transition and a critical value η_n depending on n and Θ such that $AC_n(\Theta, \hat{f}, \eta)$ decreases exponentially for any $\eta > \eta_n$. More precisely, in terms of lower bound, it is proved in [DeVore et al., 2006]

$$\inf_{\hat{f}} AC_n(\Theta, \hat{f}, \eta) \geq C \sqrt{\bar{N}(\Theta, \eta)} e^{-cn\eta^2}, \tag{18}$$

where $\bar{N}(\Theta, \eta)$ is the tight entropy analogue of the Sobolev covering numbers. The results in [DeVore et al., 2006] are obtained in the learning framework; however identical bounds can easily be obtained in the setting (1) of this paper, leading to $\eta_n = O(\sqrt{S/n})$.

If the focus is made on the case where $\alpha \in \mathcal{S}pars(S, M)$, it could be interesting to adopt the point of view of the "detection" instead of the "prediction". Two quantities become then crucial in view to measure the "similarity" between the true value and its estimator. The number of False Positive decisions (FP) and the number of False Negative decisions (FN) are given by

$$FP := \sum_{\ell=1}^{p} I\{\alpha_\ell = 0\} I\{\hat{\alpha}_\ell^* \neq 0\} \quad \text{and} \quad FN := \sum_{\ell=1}^{p} I\{\alpha_\ell \neq 0\} I\{\hat{\alpha}_\ell^* = 0\}.$$

In order to evaluate the performances of LOL selection procedure using these distances between α and $\hat{\alpha}$, the quantity $P(FP > p\eta) + P(FN > p\eta)$ for $\eta \geq 0$ is studied. A selection procedure is said consistent if $P(\ \{\ell, \ \alpha_\ell \neq 0\} = \{\ell, \ \hat{\alpha}_\ell^* \neq 0\}\ )$ is tending towards 1.

4.3 Performances of the procedure LOL

The performances of the LOL procedure are summarized in the following theorems. In Theorem 1, we establish that LOL procedure is a good procedure for estimation since the prediction error is going to zero in probability with exponential rates. Indeed, the LOL estimator is optimal (up to a logarithmic factor) in terms of the critical value $\eta_n \sim \sqrt{S/n}$, as well as in terms of exponential rates if the coherence is small enough (see the discussion below). In Theorem 2, we establish that LOL procedure works also quite well for detection since quantities FN and FP are going to zero in probability with pretty fast rates.

Theorem 1. *Let $S, M > 0$ and fix ν in $]0, 1[$. Suppose $p \leq n^a$, for some constant $a > 0$ and choose the thresholds $\lambda_n(1)$ and $\lambda_n(2)$ such that*

$$\lambda_n(1) \geq \left(T_{11} \left(\frac{\log p}{n} \right)^{1/2} \vee T_{12}\, \tau_n \right) \quad and \quad \lambda_n(2) \leq \lambda_n(1)$$

for $T_{11} = 16\sqrt{2\sigma^2}/(1+\nu)$ and $T_{12} = M\left(\frac{(1-\nu)^{1/2}}{4} \vee 4\sqrt{2}\right)$. Then, the model is of class $V(S,M)$ defined above, there exist positive constants D and γ, such that

$$
\sup_{V(S,M)} \mathbb{P}\left(d(\hat{\alpha}^*, \alpha) > \eta\right) \leq
\begin{cases}
4e^{-\gamma n \eta^2} & for \quad \eta^2 \geq D\left(\frac{S}{n} \vee \frac{S|\log \tau_n|}{n} \vee S\tau_n^2\right), \\
\\
1 & for \quad \eta^2 \leq D\left(\frac{S}{n} \vee \frac{S|\log \tau_n|}{n} \vee S\tau_n^2\right)
\end{cases}
\tag{19}
$$

Observe that the result given in Theorem 1 is concerning LOL procedures associated with more general thresholds than $\lambda_n(1), \lambda_n(2)$ than those prescribed in (5) and (7). It is interesting to notice the very few conditions on the threshold $\lambda_n(2)$ ($\lambda_n(2) \leq \lambda_n(1)$ and Condition (10) relating to the considered set of α's).

The constants D and γ are precisely given at the end of the proof of Theorem 1. For a sake of completeness, precision on the constants is given. However, it is obvious that the constants provided here are not optimal : for instance in the proof, in order to avoid unnecessary technicalities, most of the events are divided as if they had equal importance, leading to constants which are each time divided by 2. Obviously there is room for improvement at any of these stages.

An elementary consequence of Theorem 1 is the following corollary which details the behavior of the expectation of $d(\hat{\alpha}^*, \alpha)$. Notice also that we did not give here explicit oracle inequalities, which however could be derived from the proof of Theorem 1.

Corollary 1. *For $r \geq 1$ arbitrary, under the same assumptions as in Theorem 1, we get*

$$
\sup_{V(S,M)} \mathbb{E}d(\hat{\alpha}^*, \alpha)^r \leq D'\left(\frac{S}{n} \vee \frac{S|\log \tau_n|}{n} \vee S\tau_n^2\right)^{r/2}
$$

for some positive constant D'.

Notice that in the case of the l_q balls $B_q(M)$ for $q \in (0, 1]$) (see (14)) and taking the defaults values for $\lambda_n(1)$ and $\lambda_n(2)$, LOL procedure has optimal rates in the minimax sense (compare the upper bound to the lower bounds in [Raskutti et al., 2009]) as soon as $\tau_n = O\left(\sqrt{\frac{\log p}{n}}\right)$.

Let us now focus on the selection point of view. As usual, an additional assumption is needed on the non zero coefficients : they have to be large enough to be detected. Theorem 2 establishes that LOL procedure is consistent as a selection procedure.

Theorem 2. *Let k be a given positive number. Let $S, M > 0$ and fix ν in $]0, 1[$. Suppose $p \leq n^a$, for some constant $a > 0$, choose $\lambda_n(1) \geq \lambda_n(2)$ and assume that the model is of class $Spars(S, M)$ described above, then*

– **False Positive :** *Assume that*

$$
\min_{\ell=1,\ldots,p} |\alpha_\ell| \mathbb{I}\{\alpha_\ell \neq 0\} \geq \mu_n
$$

12

where μ_n satisfies

$$\mu_n = T_3 \left(\lambda_n(2) \vee \tau_n \sqrt{\frac{S}{k}} \vee \sqrt{\frac{S}{nk}} \vee \sqrt{\frac{S|\log \tau_n|}{nk}} \right)$$

where T_3 is a constant large enough. Then there exists a constant $c > 0$ such that

$$P(FP > k) \leq c \exp\{-c\, kn\lambda_n^2(2)\}.$$

– **False Negative :** *Choose the thresholds such that*

$$\lambda_n(1) \geq \left(T_{11} \left(\frac{\log p}{n} \right)^{1/2} \vee T_{12}\, \tau_n \right)$$

where the constants T_{11}, T_{12} defined as in Theorem 1 and

$$\lambda_n(2) \geq \sigma \left(32\, c_1' \sqrt{c'} \vee 256 c_1 \frac{(1+\nu)^{1/2}}{(1-\nu)^{1/2}} \right) \sqrt{\frac{S}{nk}}.$$

There exists some constant $c > 0$ such that

$$P(FN > k) \leq c \exp\{-c\, kn\lambda_n^2(2)\}$$

As for Theorem 1, Theorem 2 states for general thresholds $\lambda_n(1), \lambda_n(2)$ (which are valid for (5) and (7) but also more widely). Observe that the choice of $\lambda_n(2)$ is crucial from a detection point of view. For the specific choices (5) and (7), we get

Corollary 2. *Assume that*

$$\min_{\ell=1,\dots,p} |\alpha_\ell|\, \mathbb{I}\{\alpha_\ell \neq 0\} \geq O\left(\frac{\log n}{\sqrt{n}} \right).$$

Let $S, M > 0$ and fix ν in $]0,1[$. Suppose $p \leq n^a$, for some constant $a > 0$ and assume that the model is of class $Spars(S, M)$ described above. The LOL procedure the specific choices (5) and (7) satisfies

$$P(FN + FP > k) \leq c' n^{-c'k}.$$

for k larger than $O(S/\log n)$.

Note that LOL procedure works better and better as S gets smaller, as it is confirmed by the practical simulations.

5 Discussion and Comparisons

Comparison with other theoretical results in the literature are hereafter presented with a specific focus on domains where LOL is competitive to more sophisticated algorithms and where its simplicity induces a slight lack of efficiency. To summarize, the great benefits of LOL is to produce a very simple and auto driven algorithm with no optimization step, and with quite elementary assumptions leading to optimal exponential rates.

5.1 Estimation bounds in learning theory

As mentioned in the previous section, LOL finds its inspiration in the learning framework, especially in [Barron et al., 2008], [Binev et al., 2005], [Binev et al., 2007a],[Binev et al., 2007b]. In all these papers, consistency results are obtained under fewer assumptions but with no exponential bounds and a higher cost in implementation.

In the learning context, [Temlyakov, 2008] provides optimal critical value η_n as well as exponential bounds with fewer assumptions since there is no coherence restriction. However, the procedure is very difficult to implement for large values of p and n (N-P hard).

5.2 Comparison with other penalization procedure and coherence conditions

Comparisons has to be conducted with various procedures affiliated to the Lasso or Danzig procedures for instance [Tibshirani, 1996], [Candes and Tao, 2007], [Bickel et al., 2007], [Bunea et al., 2007a], [Bunea et al., 2007b]. First, the normalization needs to be stressed since it plays a crucial role. In many papers, the model is $Y = X\beta + \varepsilon$ and the columns of X are normalized. For comparison, our model needs to be identified in the following way

$$X := \frac{\Phi}{\sqrt{n}}, \quad \beta := \sqrt{n}\alpha.$$

Of course, each normalization brings its own benefit. Our choice has a natural interpretation in terms of prediction in the functional learning model. However, it is interesting to notice that precisely because of this normalization, the sparsity conditions on the function (model $V(S, M)$) are lighter for LOL.

LOL estimation bounds are compared with the lower bounds produced in [DeVore et al., 2006], LOL procedure gives optimal results when the coherence satisfies $\tau_n \leq O(\sqrt{\log n/n})$. This is to be compared with conditions of type $\tau_n \leq O(S^{-1})$ (see for instance [Bickel et al., 2007], [Bunea et al., 2007a], [Bunea et al., 2007b]) which are lighter except for large S, or $\tau_n \leq O(1/\log p)$ in [Candès and Plan, 2008] which is better. However, in these papers, there is generally additional assumptions

– either on the matrix X itself which generally are not possible to verify in practice. In the opposite, notice that the coherence can always be calculated.

– or on the way X as well as the β coefficients are produced, namely all these values are in fact random and independent. In our case, it can allow to less drastic coherence conditions. We infer that conditions of type $\tau_n \leq O(\sqrt{S \log n/n})$ could suffice in this case, but these precise types of models are not the scope of this paper.

5.3 Selection properties

[Meinshausen and Buhlmann, 2006] show that the selection by lasso type algorithm is consistent in graphical models, under assumptions that are tailored to models for which the vector $(Y, \Phi_1, \ldots, \Phi_p)$ is gaussian. Basically, to establish the consistency property of the selection procedure, a minimal size of the (non zero) coordinates of α is required : it is generally assumed that there exists some sequence $v_n > 0$ such that

$$\min_{\ell \in \mathcal{I}^*} |\alpha_\ell| \geq O(v_n). \tag{20}$$

[Zhao and Yu, 2006] establish the consistency of selection for fixed design linear regression models, assuming that Hypothesis (20) holds for $v_n = n^{-\kappa}$ for some $\kappa \in (0, 1/2)$. Under the same hypothesis, [Fan and Lv, 2008] prove that Sure Independence Screening (SIS) is accurate in the sense that SIS selects (with large probability) at least the regressors which have to be selected. They need to assume that there exists some $\tau > 0$ which is the indicator of the growth of the largest eigenvalue of the variance matrix Σ of Φ defined by $\lambda_{\max}(\Sigma) \leq O(n^\tau)$. The main advantage of [Fan and Lv, 2008] is that their results are basically concerning a linear model in ultra high dimension $p = \exp(cn^\xi)$ for constants $c, \xi > 0$ with the restriction $\xi \in (0, 1 - 2\kappa)$. Practical inconvenient is that the tuning sequence γ_n is not auto driven since it has to verify $n^{1-2\kappa-\tau} \longrightarrow \infty$. The selection procedure of [Bunea, 2009] proposed in the learning framework is also shown to be consistent. Hypothesis (20) is required for $v_n = S\sqrt{\log n/n}$ imposing some restriction on S because v_n is supposed to tend to zero. Finally, [Candès and Plan, 2008] prove consistency results as soon as Hypothesis (20) is satisfied for $v_n = 8\sigma\sqrt{2 \log p}$ and if $S \leq O(p/[\|\Phi\|^2 \log p])$ (in a non asymptotical framework, for any dimension p). When the selection procedure is derived from an estimation procedure, a coherence restriction could be asked. In [Bunea, 2009] and [Bunea et al., 2007b], it is assumed in addition that

$$\sup_{\ell \in \mathcal{I}^*, \, m \notin \mathcal{I}^*} \frac{1}{n} \sum_{i=1}^n |\Phi_{i\ell}\Phi_{im}| \leq O(S^{-1}).$$

and an exponential bound (tending to zero) is established for $P\left(\sum_{\ell=1}^p |\hat{\alpha}_\ell - \alpha_\ell| > \sqrt{S}\,\eta\right)$ when $\eta \geq \sqrt{S \log p/n}$. In [Candès and Plan, 2008], if $\tau_n \leq O(c/\log p)$ and again $\eta \geq \sqrt{S \log p/n}$, it is

proved that

$$P\left(d(\hat{\alpha}_\ell, \alpha_\ell) > \eta\right) \le 6p^{-2\log 2} - p^{-1}(2\pi \log p)^{-1/2}.$$

[Temlyakov, 2008] provides optimal critical value η_n as well as exponential bounds with fewer assumptions : there is no coherence restriction and the setting is the learning framework. In [Fan and Lv, 2008], under some hypothesis of RIP type, the procedure SIS-D (SIS followed by danzig) is asymptotically consistent

$$P\left(\sum_{\ell=1}^{p}(\hat{\alpha}_\ell^{SIS-D} - \alpha_\ell)^2 > S\sqrt{\log N}\right) \longrightarrow 0$$

6 Practical results

In this section, an extensive computational experiment is conducted using LOL. The procedure is dedicated to find sparse solutions of linear models assuming that the target variable Y is a linear combination of only S predictors among p. The performances of LOL procedure are studied over various ranges of level of indeterminacy $\delta = 1 - n/p$ and over various ranges of sparsity rates $\rho = S/n$ (see [Maleki and Donoho, 2009]). The influence of the choice of the distribution family for the design matrix is analyzed through the performances. LOL procedure is finally compared to some others two-steps procedures described in Section 3.2.

6.1 Experimental design

The design matrix Φ has p i.i.d. columns of size n. Different distributions are studied : Gaussian, Uniform, Bernoulli, or Student laws. It is important to notice that this choice of laws yields different values of the coherence τ_n and then different behaviors of the procedure. Each column vector of Φ is normalized to have unit norms. Given Φ, the target observations are $Y = \Phi\alpha + \varepsilon$ for ε i.i.d. variables with a normal distribution $N(0, \sigma_\varepsilon)$, σ_ε chosen such that the signal over noise ratio is close to 2. The vector α is built as follows : all coordinates are zero except S non zero coordinates with $\alpha_\ell = (-1)^b |z|$ where b is drawn from a Bernoulli distribution with parameter 0.5 and z from a $N(2, 1)$ (see [Fan and Lv, 2008]).

To evaluate the quality of the prediction, the relative l_2 error E_Y is computed on the target Y and the relative quadratic error E_α is computed on the α coefficients

$$E_Y = \|Y - \widetilde{Y}\|_2^2 / \|Y\|_2^2 \quad \text{and} \quad E_\alpha = \|\alpha - \hat{\alpha}\|_2^2 / \|\alpha\|_2^2.$$

The sparsity S is estimated by the cardinal of $\mathcal{L} = \{\ell = 1, \ldots, p, \widehat{\alpha}^* \neq 0\}$ where $\widehat{\alpha}^*$ is the LOL estimator. The number of False Positive and of False Negative as defined in Section 4.2 are also

computed. All these quantities are estimated by averaging the results obtained over $K = 200$ replications of the experiment.

6.2 Algorithm

Let us explain how to determine in a really adaptive way the thresholds $\lambda_n(1)$ and $\lambda_n(2)$. These are critical values quite hard to tune practically since they depend on inaccessible constants (see the theoretical results). Since the first threshold $\lambda_n(1)$ is used to select the candidates to the regression, the aim is to split the set of 'correlations' $\{K_\ell, \ell = 1, \ldots, p\}$, in two clusters in such a way to pick up the regression candidates in one group. Here, the sparsity assumption is used : some predictors are more correlated to the target Y than some others associated to a weak correlation value, close to zero. This remark implies that the distribution of correlations (in absolute value) should be distributed in two clusters : one for the leaders (high correlations) and one for the others (very small correlations). The frontier between the clusters is adaptively computed by minimizing the deviance of the absolute value correlations for two classes as described in [Kerkyacharian et al., 2009].

The same procedure is used to threshold adaptively the estimated coefficients $\widehat{\alpha}_\ell$ obtained by linear regression on the leaders. Indeed, notice that the distribution of the $\widehat{\alpha}_\ell$ provides two clusters : one cluster associated to the largest coefficients (in absolute value) corresponding to the non zero coefficients and one cluster composed of coefficients closed to zero, which should not be involved in the model. The frontier between the two clusters, which defines $\lambda_n(2)$, is again computed by minimizing the deviance between the two classes of regression coefficients.

Finally, an improvement for LOL is proposed. It seems more appropriate to perform a second regression using the final set $\mathcal{L}$ of selected predictors involved in the model : the estimators of the (non zero) coefficients should be more accurate. This updating procedure is denoted LOL^+ in the sequel.

6.3 Results with random gaussian design matrices

First, the design matrix Φ is defined with i.i.d. gaussian variables. The computed coherence is also $\tau_n = 0.33$ (see Figure 5). As we are interested in quantifying LOL performance in an overwhelming majority of cases, we study the impact of the level of indeterminacy δ from 0 to 0.9 by 0.05 step and the impact of the the sparsity rate ρ from 0.01 to 0.16 by 20 steps. $p = 1000$ is chosen and for specific studies $n = 250$.

Influence of the indeterminacy level : Figure 1 studies LOL prediction and estimation performances when the indeterminacy level is varying ($p = 1000$, n varying). Both errors E_Y and E_α continuously increase with the indeterminacy δ, as the number of available observations

decreases compared to the number of variables. For a given value of δ, E_Y decreases as the sparsity does. For $\delta \leq 0.75$ ($n \geq 0.25p$), the prediction error is weak, below 5%. In this case, the estimation error on the coefficients is less than 10%. When the number of available observations is at least higher than half of the number of potential predictors ($\delta < 0.5$), the prediction and the estimation errors are negligible : LOL performances are in this case exceptionally good. For a given number of observations and potential predictors, the prediction is more accurate as the sparsity rate decreases. For a fixed number observations, regarding the joint values of both indeterminacy and sparsity parameters, the errors tends to be null as δ and/or ρ decrease.

Influence of the sparsity rate : Figure 2 illustrates LOL prediction and estimation performances when the sparsity rate is varying. For small values of sparsity rate ($\rho \leq 5\%$), both prediction and estimation errors are very good (less than 5%). For an extreme level of sparsity ($\rho \leq 2\%$), the performances are, as expected, excellent. As observed before, for a given sparsity rate value, the performances are improved as the indeterminacy decreases.

Estimator of the Sparsity S : Figure 3 shows the estimated sparsity as a function of the effective sparsity S. For weak sparsity values ($\rho \leq 5\%$), LOL procedure is excellent because it estimates exactly (with no error) the sparsity S and that for all studied indeterminacy levels. As the sparsity increases, LOL procedure tends to underestimate the parameter S. For a given sparsity value, the underestimation becomes weaker as the indeterminacy level δ decreases. This observation is detailed in Table 1 where the False Positive and False Negative numbers are computed for different values of sparsity. Two different cases of indeterminacy are presented ($\delta = 0.75, 0.5$). For each indeterminacy level, we observe that False Negative and False Positive numbers increase with S both in mean and variability. As the indeterminacy level decreases from $\delta = 0.75$ to $\delta = 0.5$, meaning that more observations are available relatively to the number of potential predictors, the detection of True Positive is improved.

Estimator of the coefficients : Figure 4 presents the improvements provided by LOL^+ compared to LOL as a function of sparsity rate for the prediction error. For all indeterminacy and sparsity values, the prediction error decreases using LOL^+ procedure instead of LOL. Improvements are stronger as both sparsity rate and indeterminacy level increase. The prediction improvements are observed as ρ increases given all studied indeterminacy levels δ. Obviously, the estimated sparsity in the same for both procedures LOL and LOL+ (see Table 1).

6.4 Impact of the variable distribution in the design matrix

This section investigates the impact of the law of the regressor variables. Eight different distributions are studied : Gaussian ($N(0,1)$), Uniform ($U[-1,1]$), Bernoulli ($B\{-1,+1\}$) and Student

($T(m)$ with $m \in \{5, 4, 3, 2, 1\}$). The column of the design matrix Φ are empirically normalized.

Figure 5 shows the empirical density of the coherence τ_n computed for each law. Similar distributions are observed for Gaussian, Uniform or Bernoulli laws with a mode of the coherence equal to $\tau_n = 0.30$. For Student's families, a shift of the mode of the empirical distributions can be observed from left to right equaled to 0.36 for $T(5)$, 0.47 for T(4), 0.68 for T(3), 0.92 for T(2) to 0.99 for T(1). Figure 6 studies the estimation of S as a function of the sparsity rate ρ for those distributions. All the curves, except the one for the Student law T(1), are confounded and show similar evolution as the one observed for gaussian predictors (see Figure 3 for $\delta = 0.25$). LOL provides similar results for Gaussian, Uniform, Bernoulli, or Student laws, T(m) with m large enough. It is amazing to observe that the procedure works fine even when the empirical coherence of the distribution τ_n reaches large values closed to 0.99. But LOL procedure does not work fine for heavy tailed variables as for $T(1)$.

Figure 7 shows the coherence of the matrix restricted to the N leaders. This "restricted" coherence is much lower than the coherence computed on all the predictors. For the Student $T(1)$ law, $\tau_n = 0.99$ (see Figure 5) while the coherence computed just on the leaders is 0.3 (see Figure 7 by instance for $S = 10$). LOL procedure provides also good results even when the global coherence approaches 1 : it seems that the practical results are much more optimistic (although they do show some deterioration under high coherence). Conclusions would be that it could be interesting to find new measures of collinearity to best reflect the performances of the method. This is true in general, for all the methods concerned with high dimension.

Table 2 shows the false detections FP and FN estimated for different distributions and values of sparsity S. For a given distribution, they increase with sparsity. This increment is stronger for distributions with high coherence. For a given sparsity number S, False Positive and False Negative increase as the coherence τ_n does. LOL tends to underestimate the number of non-zero coefficients. The underestimation is stronger as the coherence of the predictors increases.

6.5   Comparison with other two-steps procedures

In this part, the performances of LOL and LOL+ are compared with the performances of two two-step procedures. The first one referred as SIS-Lasso is coming from [Fan and Lv, 2008] : the selection step called SIS is followed by the Lasso procedure. The second one, called Lasso-Reg, is proposed in [Candès and Plan, 2008]. First, the Lasso algorithm performs the selection of the leaders and then the coefficients are estimated by regression.

The performances of the four procedures (LOL, LOL+, SIS-Lasso, Lasso-Reg) are studied over a large range of sparsity rates in order to merge previous results already presented in [Fan and Lv, 2008]

and [Candès and Plan, 2008]. In this section, the sparsity S varies from 5 to 50 in 10 steps and the number of initial predictors is $p = 1000$. This experimental design let us analyze extreme sparsity values ($0.02 \le \rho \le 0.05$) (as in [Fan and Lv, 2008]) as values as large as $1/log(p)$ ($\rho = 0.20$) (as in [Candès and Plan, 2008]). For the Lasso procedures, the regularization parameter is chosen by crossvalidation.

Figure 8 presents the prediction error for the different design matrices distributions presented in the previous section. For extreme sparsity levels, $\rho < 5\%$, all the procedures performs extremely well. For middle sparsity levels ($5\% \le \rho \le 15\%$), the Lasso-Reg performs better than the others ones, as the design matrix is defined with Gaussian, Uniform, Bernoulli or Student distributions ($m = 4, 5$). For this range of sparsity levels, the Lasso-Reg procedure seems to be more efficient to select the leaders than the SIS-Lasso and the LOL procedures. For largest values of the sparsity level $\rho \ge 0.15$, it appears that SIS-Lasso and LOL are better than Lasso-Reg. A phase transition can be observed for the Lasso-Reg procedure as described in [Maleki and Donoho, 2009]. As the coherence of the design matrix increases, the phase transition appears sooner for smallest ρ values. The performances of the SIS-Lasso and LOL are globally similar. Note that LOL+ procedure improves continuously the performances compared to LOL and SIS-Lasso.

Références

[Alquier and Hebiri, 2009] Alquier, P. and Hebiri, M. (2009). Transductive versions of the lasso and the dantzig selector.

[Bahadur, 1960] Bahadur, R. R. (1960). On the asymptotic efficiency of tests and estimates. *Sankhyā*, 22 :229–252.

[Barron et al., 2008] Barron, A. R., Cohen, A., Dahmen, W., and DeVore, R. A. (2008). Approximation and learning by greedy algorithms. *Ann. Statist.*, 36(1) :64–94.

[Bickel et al., 2007] Bickel, P. J., Ritov, Y., and Tsybakov, A. (2007). Simultaneous analysis of lasso and dantzig selector. Preprint Submitted to the Annals of Statistics.

[Binev et al., 2007a] Binev, P., Cohen, A., Dahmen, W., and DeVore, R. (2007a). Universal algorithms for learning theory. II. Piecewise polynomial functions. *Constr. Approx.*, 26(2) :127–152.

[Binev et al., 2007b] Binev, P., Cohen, A., Dahmen, W., and DeVore, R. (2007b). Universal piecewise polynomial estimators for machine learning. In *Curve and surface design : Avignon 2006*, Mod. Methods Math., pages 48–77. Nashboro Press, Brentwood, TN.

[Binev et al., 2005] Binev, P., Cohen, A., Dahmen, W., DeVore, R., and Temlyakov, V. (2005). Universal algorithms for learning theory. I. Piecewise constant functions. *J. Mach. Learn. Res.*, 6 :1297–1321 (electronic).

[Bunea, 2009] Bunea, F. (2009). Consistent selection via the lasso for high dimensional approximation regression models.

[Bunea et al., 2007a] Bunea, F., Tsybakov, A., and Wegkamp, M. (2007a). Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.*, 1 :169–194 (electronic).

[Bunea et al., 2007b] Bunea, F., Tsybakov, A. B., and Wegkamp, M. H. (2007b). Sparse density estimation with ℓ_1 penalties. In *Learning theory*, volume 4539 of *Lecture Notes in Comput. Sci.*, pages 530–543. Springer, Berlin.

[Candès and Plan, 2008] Candès, E. and Plan, Y. (2008). Near-ideal model selection by ℓ_1 minimization.

[Candes and Tao, 2007] Candes, E. and Tao, T. (2007). The Dantzig selector : statistical estimation when p is much larger than n. *Ann. Statist.*, 35(6) :2313–2351.

[DeVore et al., 2006] DeVore, R., Kerkyacharian, G., Picard, D., and Temlyakov, V. (2006). Approximation methods for supervised learning. *Found. Comput. Math.*, 6(1) :3–58.

[Fan and Lv, 2008] Fan, J. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Statist. Soc. B*, 70 :849–911.

[Fan and Lv, 2009] Fan, J. and Lv, J. (2009). A selective overview of variable selection in high dimensional feature space.

[Kerkyacharian et al., 2009] Kerkyacharian, G., Mougeot, M., Picard, D., and Tribouley, K. (2009). Learning out leaders : exponential rates of convergence in high dimensional regression.

[Kerkyacharian and Picard, 2007] Kerkyacharian, G. and Picard, D. (2007). Thresholding in learning theory. *Constr. Approx.*, 26(2) :173–203.

[Lounici, 2008] Lounici, K. (2008). High-dimensional stochastic optimization with the generalized dantzig estimator.

[Maleki and Donoho, 2009] Maleki, A. and Donoho, D. L. (2009). Optimally tuned iterative thresholding algorithm for compressed sensing. *IEEE journal in signal processing*, in press.

[Meinshausen and Buhlmann, 2006] Meinshausen, N. and Buhlmann, P. (2006). High dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34 :1436–1462.

[Raskutti et al., 2009] Raskutti, G., Wainwright, M. J., and Yu, B. (2009). Minimax rates of estimation for high-dimensional linear regression over $
ell_q$-balls.

[Satterthwaite, 1959] Satterthwaite, F. E. (1959). Random balance experimentation. *Technometrics*, 1 :111–137.

[Temlyakov, 2008] Temlyakov, V. N. (2008). Approximation in learning theory. *Constr. Approx.*, 27(1) :33–74.

[Tibshirani, 1996] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society B*, 58(1) :267–288.

[Zhao and Yu, 2006] Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *Journal of Machine Learning Journal*, 7 :2541–2563.

FIG. 1 – X−axis : indeterminacy level δ, Y−axis : Prediction error (left) and estimation error (right). $S = 10$ (solid line-red) ; $S = 12$ (dot dash line-blue) ; $S = 15$ (dashed line -green) ; $S = 20$ (dot line-black).

Fig. 2 – X−axis : sparsity rate ρ, Y−axis : Prediction error (left) and estimation error (right). $\delta = 0.4$ (dot line-black) ; $\delta = 0.7$ (dot dash line-blue) ; $\delta = 0.75$ (solid line-red) ; $\delta = 0.875$, (dashed line-green).

FIG. 3 – LOL Sparsity Estimation (ρ : bottom, left ; S : right, top). $\delta = 0.875$ (dashed line-green) ; $\delta = 0.75$ (solid line-red) ; $\delta = 0.7$ (dot dash line-blue) ; $\delta = 0.4$ (dot line-black). The columns of Φ are Gaussian of size $n = 250$.

δ	$S\ (\rho)$	TP	FN	FP
0.5	5	4.98 (0.14)	0.00 (0.00)	0.00 (0.00)
0.5	10	9.88 (0.35)	0.03 (0.18)	0.00 (0.00)
0.5	15	14.54 (0.66)	0.24 (0.47)	0.01 (0.09)
0.5	20	18.78 (1.04)	0.76 (0.88)	0.03 (0.17)
0.5	25	22.74 (1.42)	1.67 (1.26)	0.07 (0.25)
0.75	5	4.98 (0.12)	0.00 (0.00)	0.00 (0.00)
0.75	10	9.90 (0.32)	0.04 (0.19)	0.00 (0.00)
0.75	15	14.57 (0.56)	0.30 (0.51)	0.01 (0.12)
0.75	20	18.77 (0.95)	1.03 (0.89)	0.05 (0.24)
0.75	25	21.94 (1.91)	2.81 (1.90)	0.19 (0.48)

TAB. 1 – Detection, $n = 250$. The columns of Φ are i.i.d. gaussian. True Positive, False positive and False negative. Means over $K = 200$ replications, variances into the brackets.

Fig. 4 – Error. $X-$axis : sparsity rate ρ. $Y-$axis : Prediction errors for LOL (dot lines) and LOL+ (solid lines). $\delta = 0.4$ (blue color) ; $\delta = 0.75$ (red color) ; $\delta = 0.875$ (green color). The columns of Φ are Gaussian of size $n = 250$.

S		G	U	B	T(5)	T(4)	T(3)	T(2)	T(1)
5	FP	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.01 (0.2)
	FN	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.01 (0.1)
10	FP	0.01 (0.1)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.00 (0.0)	0.01 (0.1)	0.01 (0.1)	0.19 (0.8)
	FN	0.04 (0.2)	0.05 (0.2)	0.05 (0.2)	0.06 (0.3)	0.04 (0.2)	0.06 (0.2)	0.04 (0.2)	0.16 (0.8)
15	FP	0.03 (0.2)	0.02 (0.1)	0.02 (0.1)	0.03 (0.2)	0.01 (0.1)	0.03 (0.2)	0.11 (0.4)	1.08 (1.9)
	FN	0.41 (0.6)	0.35 (0.6)	0.36 (0.6)	0.31 (0.6)	0.30 (0.5)	0.34 (0.6)	0.35 (0.6)	0.56 (1.3)
20	FP	0.09 (0.3)	0.07 (0.3)	0.08 (0.3)	0.07 (0.3)	0.14 (0.4)	0.08 (0.3)	0.39 (0.6)	1.91 (2.2)
	FN	1.26 (1.0)	1.26 (1.1)	1.25 (1.0)	1.24 (1.0)	1.37 (1.1)	1.25 (1.0)	1.26 (1.2)	1.59 (2.8)
25	FP	0.19 (0.4)	0.10 (0.3)	0.17 (0.4)	0.17 (0.4)	0.21 (0.5)	0.23 (0.6)	0.53 (0.7)	3.92 (2.7)
	FN	2.78 (1.5)	2.93 (1.8)	2.61 (1.7)	2.69 (1.8)	2.84 (1.7)	2.75 (1.8)	2.92 (1.9)	4.12 (3.9)
30	FP	0.39 (0.8)	0.42 (0.9)	0.39 (0.6)	0.34 (0.6)	0.36 (0.7)	0.41 (0.7)	0.83 (1.0)	4.69 (2.7)
	FN	5.90 (2.9)	6.05 (3.0)	5.45 (2.5)	5.93 (2.8)	5.29 (2.8)	5.42 (2.7)	5.47 (3.0)	8.76 (7.4)
35	FP	0.70 (1.5)	0.61 (1.0)	0.78 (1.3)	0.68 (1.1)	0.63 (1.0)	0.84 (1.7)	1.02 (1.3)	5.71 (3.0)
	FN	9.44 (3.7)	9.19 (3.9)	9.54 (3.6)	9.63 (4.3)	10.02 (4.0)	9.73 (4.1)	10.01 (3.9)	14.77 (8.6)
40	FP	1.24 (1.5)	1.21 (1.5)	1.18 (1.4)	1.06 (1.5)	1.15 (1.5)	1.31 (1.7)	1.60 (2.1)	6.24 (3.0)
	FN	14.73 (4.5)	14.98 (4.8)	14.72 (5.0)	15.15 (4.6)	14.56 (4.9)	15.53 (4.3)	15.34 (5.1)	21.70 (9.0)

Tab. 2 – False Detection, $n = 250, p = 1000$. First line : Common law of the columns of Φ. First column : Sparsity S. First lines : False positive, Second lines : False negative. Means over $K = 200$ replications, variances into the brackets.

FIG. 5 – $n = 250, p = 1000$. Empirical densities of the coherence. The columns of Φ are Gaussian (solid line-red) ; uniform (solid line-blue) ; Bernoulli (solid line-green) ; Student $5, 4, 3, 2, 1$ black lines from left to right.

FIG. 6 – LOL Sparsity estimation for different families of laws for the predictors. *Gauss* (solid line-red) ; *Uniform* (solid line-blue) ; *Bernoulli* : (solid line-green) ; T(1-5) (black-lines). $n = 250$, $p = 1000$. ($K = 200$)

FIG. 7 – Coherence computed for the N selected Leaders. $Gauss$ (solid line-red) ; $Uniform$ (solid line-blue) ; $Bernoulli$: (solid line-green) ; T(1) (dot line-black). $n = 250$, $p = 1000$. ($K = 200$)

FIG. 8 – X-axis : Sparsity rate. Y-Axis : Prediction error for different design matrices. LOL (red solid line), LOL^+ (red dotted lines), $SIS - Lasso$ (green solid lines), and $Lasso - Reg$ (blue solid line). $n = 250$, $p = 1000$. ($K = 200$)