

HIGH-DIMENSIONAL ROBUST REGRESSION AND OUTLIERS DETECTION WITH SLOPE

BY ALAIN VIROULEAU^{*,‡}, AGATHE GUILLOUX^{‡,§}, STÉPHANE
GAÏFFAS^{*,‡,‡} AND MALGORZATA BOGDAN[¶]

*Université Paris Diderot[†], Ecole polytechnique[‡], Université
d'Evry[§] and University of Wrocław[¶]*

The problems of outliers detection and robust regression in a high-dimensional setting are fundamental in statistics, and have numerous applications. Following a recent set of works providing methods for simultaneous robust regression and outliers detection, we consider in this paper a model of linear regression with individual intercepts, in a high-dimensional setting. We introduce a new procedure for simultaneous estimation of the linear regression coefficients and intercepts, using two dedicated sorted- ℓ_1 penalizations, also called SLOPE [5]. We develop a complete theory for this problem: first, we provide sharp upper bounds on the statistical estimation error of both the vector of individual intercepts and regression coefficients. Second, we give an asymptotic control on the False Discovery Rate (FDR) and statistical power for support selection of the individual intercepts. As a consequence, this paper is the first to introduce a procedure with guaranteed FDR and statistical power control for outliers detection under the mean-shift model. Numerical illustrations, with a comparison to recent alternative approaches, are provided on both simulated and several real-world datasets. Experiments are conducted using an open-source software written in Python and C++.

1. Introduction. Outliers are a fundamental problem in statistical data analysis. Roughly speaking, an outlier is an observation point that differs from the data's "global picture" [27]. A rule of thumb is that a typical dataset may contain between 1% and 10% of outliers [26], or much more than that in specific applications such as web data, because of the inherent complex nature and highly uncertain pattern of users' web browsing [22]. This outliers problem was already considered in the early 50's [15, 21] and it motivated in the 70's the development of a new field called robust statistics [28, 29].

In this paper, we consider the problem of linear regression in the presence of outliers. In this setting, classical estimators, such as the least-squares, are

*Data Science Initiative of Ecole polytechnique

MSC 2010 subject classifications: Primary 62J05; secondary 62F35; 62J07; 62H15

Keywords and phrases: Linear regression, Outliers detection, Robust statistics, High-dimensional statistics, Sparsity, False Discovery Rate, Sorted-L1 norm

known to fail [28]. In order to conduct regression analysis in the presence of outliers, roughly two approaches are well-known. The first is based on detection and removal of the outliers to fit least-squares on “clean” data [49]. Popular methods rely on leave-one-out methods (sometimes called case-deletion), first described in [14] with the use of residuals in linear regression. The main issue about these methods is that they are theoretically well-designed for the situations where only one given observation is an outlier. Repeating the process across all locations can lead to well-known masking and swamping effects [25]. An interesting recent method that does not rely on a leave-one-out technique is the so-called IPOD [39], a penalized least squares method with the choice of tuning parameter relying on a BIC criterion. A second approach is based on robust regression, that considers loss functions that are less sensitive to outliers [29]. This relies on the M -estimation framework, that leads to good estimators of regression coefficients in the presence of outliers, thanks to the introduction of robust losses replacing the least-squares. However, the computation of M -estimates is substantially more involved than that of the least-squares estimates, which to some extent counter-balance the apparent computational gain over previous methods. Moreover, robust regression focuses only on the estimation of the regression coefficients, and does not allow directly to localize the outliers, see also for instance [50] for a recent review.

Alternative approaches have been proposed to perform simultaneously outliers detection and robust regression. Such methods involve median of squares [40], S-estimation [38] and more recently robust weighted least-squares [19], among many others, see also [24] for a recent review on such methods. The development of robust methods intersected with the development of sparse inference techniques recently. Such inference techniques, in particular applied to high-dimensional linear regression, are of importance in statistics, and have been an area of major developments over the past two decades, with deep results in the field of compressed sensing, and more generally convex relaxation techniques [45, 9, 10, 12, 11]. These led to powerful inference algorithms working under a sparsity assumption, thanks to fast and scalable convex optimization algorithms [2]. The most popular method allowing to deal with sparsity and variable selection is the LASSO [46], which is ℓ_1 -penalized least-squares, with improvements such as the Adaptive LASSO [52], among a large set of other sparsity-inducing penalizations [7, 3].

Within the past few years, a large amount of theoretical results have been established to understand regularization methods for the sparse linear regression model, using so-called oracle inequalities for the prediction and estimation errors [30, 31, 33], see also [7, 20] for nice surveys on this topic.

Another line of works focuses on variable selection, trying to recover the support of the regression coefficients with a high probability [32, 31, 13]. Other types of loss functions [47] or penalizations [17, 5] have also been considered. Very recently, the sorted- ℓ_1 norm penalization has been introduced [5, 6, 42] and very strong statistical properties have been shown. In particular, when covariates are orthogonal, SLOPE allows to recover the support of the regression coefficients with a control on the False Discovery Rate [5]. For i.i.d covariates with a multivariate Gaussian distribution, oracle inequalities with optimal minimax rates have been shown, together with a control on a quantity which is very close to the FDR [42]. For more general covariate distributions, oracle inequalities with an optimal convergence rate are obtained in [8].

However, many high-dimensional datasets, with hundreds or thousands of covariates, do suffer from the presence of outliers. Robust regression and detection of outliers in a high-dimensional setting is therefore important. Increased dimensionality and complexity of the data may amplify the chances of an observation being an outlier, and this can have a strong negative impact on the statistical analysis. In such settings, many of the aforementioned outlier detection methods do not work well. A new technique for outliers detection in a high-dimensional setting is proposed in [1], which tries to find the outliers by studying the behavior of projections from the data set. A small set of other attempts to deal with this problem have been proposed in literature [48, 37, 23, 39, 18], and are described below with more details.

2. Contributions of the paper. Our focus is on possibly high dimensional linear regression where observations can be contaminated by gross errors. This so-called mean-shifted outliers model can be described as follows:

$$(2.1) \quad y_i = x_i^\top \beta^\star + \mu_i^\star + \varepsilon_i$$

for $i = 1, \dots, n$, where n is the sample size. A non-zero μ_i^\star means that observation i is an outlier, and $\beta^\star \in \mathbb{R}^p$, $x_i \in \mathbb{R}^p$, $y_i \in \mathbb{R}$ and $\varepsilon_i \in \mathbb{R}$ respectively stand for the linear regression coefficients, vector of covariates, label and noise of sample i . For the sake of simplicity we assume throughout the paper that the noise is i.i.d centered Gaussian with known variance σ^2 .

2.1. Related works. We already said much about the low-dimensional problem so we focus in this part on the high-dimensional one. The leave-one-out technique has been extended in [48] to high-dimension and general regression cases, but the masking and swamping problems remains. In other

models, outliers detection in high-dimension also includes distance-based approaches [37] where the idea is to find the center of the data and then apply some thresholding rule. The model (2.1) considered here has been recently studied with LASSO penalizations [23] and hard-thresholding [39]. LASSO was used also in [18], but here outliers are modeled in the variance of the noise. In [23, 39], that are closer to our approach, the penalization is applied differently: in [23], the procedure named Extended-LASSO uses two different ℓ_1 penalties for β and μ , with tuning parameters that are fixed according to theoretical results, while the IPOD procedure from [39] applies the same penalization to both vectors, with a regularization parameter tuned with a modified BIC criterion. In [23], error bounds and a signed support recovery result are obtained for both the regression and intercepts coefficients. However, these results require that the magnitude of the coefficients is very large, which is one of the issues that we want to overcome with this paper.

It is worth mentioning that model (2.1) can be written in a concatenated form $y = Z\gamma^* + \varepsilon$, with Z being the concatenation of the covariates matrix X (with lines given by the x_i 's) and the identity matrix I_n in \mathbb{R}^n , and γ^* being the concatenation of β^* and μ^* . This leads to a regression problem with a very high dimension $n + p$ for the vector γ^* . Working with this formulation, and trying to estimate γ^* directly is actually a bad idea. This point is illustrated experimentally in [23], where it is shown that applying two different LASSO penalizations on β and μ leads to a procedure that outperforms the LASSO on the concatenated vector. The separate penalization is even more important in case of SLOPE, whose aim is FDR control for the support recovery of μ^* . Using SLOPE directly on γ^* would mix the entries of μ and β together, which would make FDR control practically impossible due to the correlations between covariates in the X matrix.

2.2. Main contributions. Given a vector $\lambda = [\lambda_1 \cdots \lambda_m] \in \mathbb{R}_+^m$ with non-negative and non-increasing entries, we define the sorted- ℓ_1 norm of a vector $x \in \mathbb{R}^m$ as

$$(2.2) \quad \forall x \in \mathbb{R}^m, J_\lambda(x) = \sum_{j=1}^m \lambda_j |x|_{(j)},$$

where $|x|_{(1)} \geq |x|_{(2)} \geq \cdots \geq |x|_{(m)}$. In [5] and [6] the sorted- ℓ_1 norm was used as a penalty in the Sorted L-One Penalized Estimator (SLOPE) of coefficients in the multiple regression. Degenerate cases of SLOPE are ℓ_1 -penalization whenever λ_j are all equal to a positive constant, and null-penalization if this constant is zero. We apply two different SLOPE penal-

izations on β and μ , by considering the following optimization problem:

$$(2.3) \quad \min_{\beta \in \mathbb{R}^p, \mu \in \mathbb{R}^n} \left\{ \|y - X\beta - \mu\|_2^2 + 2\rho_1 J_{\tilde{\lambda}}(\beta) + 2\rho_2 J_{\lambda}(\mu) \right\}$$

where ρ_1 and ρ_2 are positive parameters, X is the $n \times p$ covariates matrix with rows x_1, \dots, x_n , $y = [y_1 \cdots y_n]^T$, $\mu = [\mu_1 \cdots \mu_n]^T$, $\|u\|_2$ is the Euclidean norm of a vector u and $\lambda = [\lambda_1 \cdots \lambda_p]$ and $\tilde{\lambda} = [\tilde{\lambda}_1 \cdots \tilde{\lambda}_n]$ are two vectors with non-increasing and non-negative entries.

In this article we provide the set of sequences λ and $\tilde{\lambda}$ which allow to obtain better error bounds for estimation of μ^* and β^* than previously known ones [23], see Section 3 below. Moreover, in Section 4 we provide specific sequences which, under some asymptotic regime, lead to a control of the FDR for the support selection of μ^* , and such that the power of the procedure (2.3) converges to one. Procedure (2.3) is therefore, to the best of our knowledge, the first proposed in literature to *robustly estimate β^* , estimate and detect outliers at the same time, with a control on the FDR for the multi-test problem of support selection of μ^* , and power consistency*.

We compare in Section 5 our procedure to the recent alternatives for this problem, that is the IPOD procedure [39] and the Extended-Lasso [23]. The numerical experiments given in Section 5 confirm the theoretical findings from Sections 3 and 4. As shown in our numerical experiments, the other procedures fail to guarantee FDR control or exhibit a lack of power when outliers are difficult to detect, namely when their magnitude is not far enough from the noise-level. It is particularly noticeable that our procedure overcomes this issue.

The theoretical results proposed in this paper are based on two popular assumptions in compressed sensing or other sparsity problems, similar to the ones from [23]: first, a Restricted Eigenvalues (RE) condition [30] on X , then a mutual incoherence assumption [34] between X and I_n , which is natural since it excludes settings where the column spaces of X and I_n are impossible to distinguish. Proofs of results stated in Sections 3 and 4 are given in Section C and D, while preliminary results are given in Sections A and B. Section E provides contains supplementary extra numerical results.

3. Upper bounds for the estimation of β^* and μ^* . Throughout the paper, n is the sample size whereas p is the number of covariables, so that $X \in \mathbb{R}^{n \times p}$. For any vector u , $|u|_0$, $\|u\|_1$ and $\|u\|_2$ denote respectively the number of non-zero coordinates of u (also called sparsity), the ℓ_1 -norm and the Euclidean norm. We denote respectively by $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ the smallest and largest eigenvalue of a symmetric matrix A . We work under the following assumption

ASSUMPTION 3.1. *We assume the following sparsity assumption:*

$$(3.1) \quad |\beta^*|_0 \leq k \quad \text{and} \quad |\mu^*|_0 \leq s$$

for some positive integers k and s , and we assume that the columns of X are normalized, namely $\|Xe_i\|_2 = 1$ for $i = 1, \dots, n$, where e_i stands for the i -th element of the canonical basis.

For the results of this Section, we consider procedure (2.3) with the following choice of λ :

$$(3.2) \quad \lambda_i = \sigma \sqrt{\log\left(\frac{2n}{i}\right)},$$

for $i = 1, \dots, n$, and we consider three possibilities for $\tilde{\lambda}$, corresponding to no penalization, ℓ_1 penalization and SLOPE penalization on β .

Table 1 below gives a quick view of the convergence rates of the squared ℓ_2 estimation errors of β^* and μ^* obtained in Theorems 3.4, 3.5 and 3.6. We give also the convergence rate obtained in [23] for ℓ_1 penalization applied to β and μ . In particular, we see that using two SLOPE penalizations leads to a better convergence rate than the use of ℓ_1 penalizations. Condition 3.2 below

TABLE 1

Convergence rates, up to constants, associated to several penalization techniques. NO means no-penalization, L1 stands for ℓ_1 penalization, while SL1 stands for SLOPE. We observe that SL1 + SL1 leads to a better convergence rate than L1 + L1.

Penalization (β / μ)	Convergence rates	Reference
NO/SL1	$p \vee s \log(n/s)$	Theorem 3.4
L1/L1	$k \log p \vee s \log n$	[23]
L1/SL1	$k \log p \vee s \log(n/s)$	Theorem 3.5
SL1/SL1	$k \log(p/k) \vee s \log(n/s)$	Theorem 3.6

is a Restricted Eigenvalue (RE) type of condition which is adapted to our problem. Such an assumption is known to be mandatory in order to derive fast rates of convergence for penalizations based on the convex-relaxation principle [51].

CONDITION 3.2. *Consider two vectors $\lambda = (\lambda_i)_{i=1, \dots, n}$ and $\tilde{\lambda} = (\tilde{\lambda}_i)_{i=1, \dots, p}$ with non-increasing and positive entries, and consider positive integers k, s and $c_0 > 0$. We define the cone $\mathcal{C}(k, s, c_0)$ of all vectors $[\beta^\top, \mu^\top]^\top \in \mathbb{R}^{p+n}$*

satisfying

$$(3.3) \quad \sum_{j=1}^p \frac{\tilde{\lambda}_j}{\tilde{\lambda}_p} |\beta|_{(j)} + \sum_{j=1}^n \frac{\lambda_j}{\lambda_n} |\mu|_{(j)} \leq (1 + c_0)(\sqrt{k} \|\beta\|_2 + \sqrt{s} \|\mu\|_2).$$

We also define the cone $\mathcal{C}^p(s, c_0)$ of all vectors $[\beta^\top, \mu^\top]^\top \in \mathbb{R}^{p+n}$ satisfying

$$(3.4) \quad \sum_{j=1}^n \frac{\lambda_j}{\lambda_n} |\mu|_{(j)} \leq (1 + c_0)(\sqrt{p} \|\beta\|_2 + \sqrt{s} \|\mu\|_2).$$

We assume that there are constants $\kappa_1, \kappa_2 > 0$ with $\kappa_1 > 2\kappa_2$ such that X satisfies the following, either for all $[\beta^\top, \mu^\top]^\top \in \mathcal{C}(k, s, c_0)$ or for all $[\beta^\top, \mu^\top]^\top \in \mathcal{C}^p(s, c_0)$:

$$(3.5) \quad \|X\beta\|_2^2 + \|\mu\|_2^2 \geq \kappa_1 (\|\beta\|_2^2 + \|\mu\|_2^2)$$

$$(3.6) \quad |\langle X\beta, \mu \rangle| \leq \kappa_2 (\|\beta\|_2^2 + \|\mu\|_2^2).$$

Equation (3.4) corresponds to the particular case where we do not penalize the regression coefficient β , namely $\tilde{\lambda}_i = 0$ for all i . Note also that Condition 3.2 entails

$$\|X\beta + \mu\|_2 \geq \sqrt{\kappa_1 - 2\kappa_2} \sqrt{\|\beta\|_2^2 + \|\mu\|_2^2},$$

which actually corresponds to a RE condition on $[X^\top I_n]^\top$ and that Equation (3.5) is satisfied if X satisfies a RE condition with constant $\kappa < 1$. Finally, note that Equation (3.6), called mutual incoherence in the literature of compressed sensing, requires in this context that for all β and μ from the respective cones the potential regression predictor $X\beta$ is sufficiently not-aligned with potential outliers μ . An extreme case of violation of this assumption occurs when $X = I_n$, where we cannot separate the regression coefficients from the outliers.

The Condition 3.2 is rather mild and e.g. for a wide range of random designs. Specifically, Theorem 3.3 below, shows that it holds with large probability whenever X has i.i.d $\mathcal{N}(0, \Sigma)$ rows, with $\lambda_{\min}(\Sigma) > 0$, and the vectors β and μ are sufficiently sparse.

THEOREM 3.3. *Let $X' \in \mathbb{R}^{n \times p}$ be a random matrix with i.i.d $\mathcal{N}(0, \Sigma)$ rows and $\lambda_{\min}(\Sigma) > 0$. Let X be the corresponding matrix with normalized columns. Given positive integers k, s and $c_0 > 0$, define $r = s \vee k(1 + c_0)^2$. If*

$$\sqrt{n} \geq C\sqrt{r} \quad \text{and} \quad \sqrt{n} \geq C'\sqrt{r \log(p \vee n)}$$

with

$$C \geq 30 \frac{\sqrt{\lambda_{\max}(\Sigma)}}{\min_j \Sigma_{jj}} \left(\frac{256 \times 5 \max_j \Sigma_{jj}}{\lambda_{\min}(\Sigma)} \sqrt{16} \right) \quad \text{and} \quad C' \geq 72 \sqrt{10} \frac{(\max_j \Sigma_{jj})^{3/2}}{\min_j \Sigma_{jj} \sqrt{\lambda_{\min}(\Sigma)}},$$

then there are $c, c' > 0$ such that for any $[\beta^\top, \mu^\top]^\top \in \mathcal{C}(k, s, c_0)$, we have

$$\begin{aligned} \|X\beta\|_2^2 + \|\mu\|_2^2 &\geq \min \left\{ \frac{\lambda_{\min}(\Sigma)}{128(\max_j \Sigma_{jj})^2}, \frac{1}{8} \right\} (\|\beta\|_2^2 + \|\mu\|_2^2) \\ 2|\langle X\beta, \mu \rangle| &\leq \min \left\{ \frac{\lambda_{\min}(\Sigma)}{256 \times 5(\max_j \Sigma_{jj})^2}, \frac{1}{16} \right\} (\|\beta\|_2^2 + \|\mu\|_2^2) \end{aligned}$$

with a probability greater than $1 - cn \exp(-c'n)$. These inequalities also hold for any $[\beta^\top, \mu^\top]^\top \in \mathcal{C}^p(s, c_0)$ when k is replaced by p in the above conditions.

The proof of Theorem 3.3 is given in Appendix C.1. It is based on recent bounds results for Gaussian random matrices [36]. The numerical constants in Theorem 3.3 are far from optimal and chosen for simplicity so that $\kappa_1 > 2\kappa_2$ as required in Assumption 3.2. A typical example for Σ is the Toeplitz matrix $[a^{|i-j|}]_{i,j}$ with $a \in [0, 1)$, for which $\lambda_{\min}(\Sigma)$ is equal to $1 - a$ [36]. The required lower bound on n is non-restrictive, since k and s correspond to the sparsity of β^* and μ^* , that are typically much smaller than n . Note also that $\mathcal{C}^p(s, c_0)$ will only be used in low dimension, and in this case p is again much smaller than n .

Let us define $\kappa = \sqrt{\kappa_1 - 2\kappa_2}$ for the whole Section, with κ_1 and κ_2 defined in Assumption 3.2. The three theorem below and their proof are very similar in nature, but differ in some details, therefore are stated and proved separately. We emphasize that the proof give slightly more general versions of the theorems, allowing the same result with $\hat{\mu}$ having any given support containing $\text{Supp}(\mu^*)$. This is of great theoretical interest and is a key point for the support detection of μ^* investigated in 4. The proof use a very recent bound on the inner product between a white Gaussian noise and any vector, involving the sorted ℓ_1 norm [8]. Our first result deals with linear regression with outliers and no sparsity assumption on β^* . We consider procedure (2.3) with no penalization on β , namely

$$(3.7) \quad (\hat{\beta}, \hat{\mu}) \in \underset{\beta \in \mathbb{R}^p, \mu \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ \|y - X\beta - \mu\|_2^2 + 2\rho J_\lambda(\mu) \right\},$$

with J_λ given by (2.2) and weights λ given by (3.2), and with $\rho \geq 2(4 + \sqrt{2})$. Theorem 3.4, below, shows that a convergence rate for procedure (3.7) is indeed $p \vee s \log(n/s)$, as reported in Table 1 above.

THEOREM 3.4. *Suppose that Assumption 3.1 is met with $k = p$, and that X satisfies Assumption 3.2 on the cone $\mathcal{C}(k_1, s_1, 4)$ with $k_1 = p/\log 2$ and $s_1 = s \log(2en/s)/\log 2$. Then, the estimators $(\hat{\beta}, \hat{\mu})$ given by (3.7) satisfy*

$$\|\hat{\beta} - \beta^*\|_2^2 + \|\hat{\mu} - \mu^*\|_2^2 \leq \frac{4\rho^2}{\kappa^4} \sum_{j=1}^s \lambda_j^2 + \frac{5\sigma^2}{\kappa^4} p \leq \frac{\sigma^2}{\kappa^4} \left(4\rho^2 s \log\left(\frac{2en}{s}\right) + 5p \right),$$

with a probability larger than $1 - (s/2n)^s/2 - e^{-p}$.

The proof of Theorem 3.4 is given in Appendix C.2. The second result involves a sparsity assumption on β^* and considers ℓ_1 penalization for β . We consider this time

$$(3.8) \quad (\hat{\beta}, \hat{\mu}) \in \operatorname{argmin}_{\beta, \mu} \left\{ \|y - X\beta - \mu\|_2^2 + 2\nu \|\beta\|_1 + 2\rho J_\lambda(\mu) \right\},$$

where $\nu = 4\sigma\sqrt{\log p}$ is the regularization level for ℓ_1 penalization, $\rho \geq 2(4 + \sqrt{2})$ and J_λ is given by (2.2). Theorem 3.5, below, shows that a convergence rate for procedure (3.8) is indeed $k \log p \vee s \log(n/s)$, as reported in Table 1 above.

THEOREM 3.5. *Suppose that Assumption 3.1 is met and that X satisfies Assumption 3.2 on the cone $\mathcal{C}(k_1, s_1, 4)$ with $k_1 = 16k \log p/\log 2$ and $s_1 = s \log(2en/s)/\log 2$. Suppose also that $\sqrt{\log p} \geq \rho \log 2/4$. Then, the estimators $(\hat{\beta}, \hat{\mu})$ given by (3.8) satisfy*

$$\|\hat{\beta} - \beta^*\|_2^2 + \|\hat{\mu} - \mu^*\|_2^2 \leq \frac{36}{\kappa^4} \sigma^2 k \log p + \frac{4\rho^2}{\kappa^4} \sum_{j=1}^s \lambda_j^2 \leq \frac{4\sigma^2}{\kappa^4} \left(9k \log p + \rho^2 s \log\left(\frac{2en}{s}\right) \right),$$

with a probability larger than $1 - (s/2n)^s/2 - 1/p$.

The proof of Theorem 3.5 is given in Appendix C.4. The third result is obtained using SLOPE both on β and μ , namely

$$(3.9) \quad (\hat{\beta}, \hat{\mu}) \in \operatorname{argmin}_{\beta, \mu} \left\{ \|y - X\beta - \mu\|_2^2 + 2\rho J_{\tilde{\lambda}}(\beta) + 2\rho J_\lambda(\mu) \right\}$$

where $\rho \geq 2(4 + \sqrt{2})$, J_λ is given by (2.2), and where

$$\tilde{\lambda}_j = \sigma \sqrt{\log\left(\frac{2p}{j}\right)}$$

for $j = 1, \dots, p$. Theorem 3.6, below, shows that the rate of convergence of estimators provided by (3.9) is indeed $k \log(p/k) \vee s \log(n/s)$, as presented in Table 1.

THEOREM 3.6. *Suppose that Assumption 3.1 is met and that X satisfies Assumption 3.2 on the cone $\mathcal{C}(k_1, s_1, 4)$ with $k_1 = k \log(2ep/k)/\log 2$ and $s_1 = s \log(2en/s)/\log 2$. Then, the estimators $(\hat{\beta}, \hat{\mu})$ given by (3.9) satisfy*

$$(3.10) \quad \begin{aligned} \|\hat{\beta} - \beta^*\|_2^2 + \|\hat{\mu} - \mu^*\|_2^2 &\leq \frac{C'}{\kappa^4} \left(\sum_{j=1}^k \tilde{\lambda}_j^2 + \sum_{j=1}^s \lambda_j^2 \right) \\ &\leq \frac{C' \sigma^2}{\kappa^4} \left(k \log \left(\frac{2ep}{k} \right) + s \log \left(\frac{2en}{s} \right) \right), \end{aligned}$$

with a probability greater than $1 - (s/2n)^s/2 - (k/2p)^k/2$, where $C' = 4\rho^2 \vee (3 + C)^2/2$.

The proof of Theorem 3.6 is given in Appendix C.4. Note that according to Theorem 3.3, the assumptions of Theorem 3.6 are satisfied with probability converging to one when the rows of X are i.i.d from the multivariate Gaussian distribution with the positive definite covariance matrix, and when the signal is sparse such that $(k \vee s) \log(n \vee p) = o(n)$.

4. Asymptotic FDR control and power for the selection of the support of μ^* . We consider the multi-test problem with null-hypotheses

$$H_i : \mu_i^* = 0$$

for $i = 1, \dots, n$, and we consider the multi-test that rejects H_i whenever $\hat{\mu}_i \neq 0$, where $\hat{\mu}$ (and $\hat{\beta}$) are given either by (3.7), (3.8) or (3.9). When H_i is rejected, or “discovered”, we consider that sample i is an outlier. Note however that in this case, the value of $\hat{\mu}_i$ gives extra information on how much sample i is outlying.

We use the FDR as a standard Type I error for this multi-test problem [4]. The FDR is the expectation of the proportion of false discoveries among all discoveries. Letting V (resp. R) be the number of false rejections (resp. the number of rejections), the FDR is defined as

$$(4.1) \quad \text{FDR}(\hat{\mu}) = \mathbb{E} \left[\frac{V}{R \vee 1} \right] = \mathbb{E} \left[\frac{\#\{i : \mu_i^* = 0, \hat{\mu}_i \neq 0\}}{\#\{i : \hat{\mu}_i \neq 0\}} \right].$$

We use the Power to measure the Type II error for this multi-test problem. The Power is the expectation of the proportion of true discoveries. It is defined as

$$(4.2) \quad \Pi(\hat{\mu}) = \mathbb{E} \left[\frac{\#\{i : \mu_i^* \neq 0, \hat{\mu}_i \neq 0\}}{\#\{i : \mu_i^* \neq 0\}} \right],$$

the Type II error is then given by $1 - \Pi(\hat{\mu})$.

For the linear regression model without outliers, a multi-test for the support selection of β^* with controlled FDR based on SLOPE is given in [5] and [6]. Specifically, it is shown that SLOPE with weights

$$(4.3) \quad \lambda_i^{\text{BH}} = \sigma \Phi^{-1} \left(1 - \frac{iq}{2n} \right)$$

for $i = 1, \dots, n$, where Φ is the cumulative distribution function of $\mathcal{N}(0, 1)$ and $q \in (0, 1)$, controls FDR at the level q in the multiple regression problem with orthogonal design matrix $X^T X = I$. It is also observed that when the columns of X are not orthogonal but independent the weights have to be substantially increased to guarantee FDR control. This effect results from the random correlations between columns of X and the shrinkage of true nonzero coefficients, and in context of LASSO have been thoroughly discussed in [41].

In this paper we substantially extend current results on FDR controlling properties of SLOPE. Specifically, Theorem 4.1 below gives asymptotic controls of $\text{FDR}(\hat{\mu})$ and $\Pi(\hat{\mu})$ for the procedures (3.7), (3.8) and (3.9), namely different penalizations on β and SLOPE applied on μ , with slightly increased weights

$$(4.4) \quad \lambda = (1 + \epsilon) \lambda^{\text{BH}},$$

where $\epsilon > 0$, see also [42]. This choice of λ also yields optimal convergence rates, however considering it in Section 3 would lead to some extra technical difficulties. Under appropriate assumptions on p, n , the signal sparsity and the magnitude of outliers, Theorem 4.1 not only gives FDR control, but also proves that the Power is actually going to 1.

Note that all asymptotics considered here are with respect to the sample size n , namely the statement $d \rightarrow +\infty$ means that $d = d_n \rightarrow +\infty$ with $n \rightarrow +\infty$.

THEOREM 4.1. *Suppose that there is a constant M such that the entries of X satisfy $|x_{i,j}| \sqrt{n} \leq M$ for all $i, j \in \{1, \dots, n\}$, and suppose that*

$$|\mu_i^*| \geq (1 + \rho(1 + 2\epsilon)) 2\sigma \sqrt{\log n}$$

for any $i = 1, \dots, n$ such that $\mu_i^ \neq 0$. Suppose also that $s \rightarrow +\infty$. Then, consider $(\hat{\beta}, \hat{\mu})$ given either by (3.7), (3.8) and (3.9), with λ given by (4.4). For Procedure (3.7), assume the same as in Theorem 3.4, and that*

$$\frac{p(s \log(n/s) \vee p)}{n} \rightarrow 0.$$

For Procedure (3.8), assume the same as in Theorem 3.5, and that

$$\frac{(s \log(n/s) \vee k \log p)^2}{n} \rightarrow 0.$$

For Procedure (3.9), assume the same as in Theorem 3.6, and that

$$\frac{(s \log(n/s) \vee k \log(p/k))^2}{n} \rightarrow 0.$$

Then, the following properties hold:

$$(4.5) \quad \Pi(\hat{\mu}) \rightarrow 1, \quad \limsup \text{FDR}(\hat{\mu}) \leq q.$$

The proof of Theorem 4.1 is given in Appendix D. It relies on a careful look at the KKT conditions, also known as the dual-certificate method [32] or resolvent solution [42]. The assumptions of Theorem 4.1 are natural. The boundedness assumption on the entries of X are typically satisfied with a large probability when X is Gaussian for instance. When $n \rightarrow +\infty$, it is also natural to assume that $s \rightarrow +\infty$ (let us recall that s stands for the sparsity of the sample outliers $\mu \in \mathbb{R}^n$). The asymptotic assumptions roughly ask for the rates in Table 1 to converge to zero. Finally, the assumption on the magnitude of the non-zero entries of μ^* is somehow unavoidable, since it allows to distinguish outliers from the Gaussian noise. We emphasize that good numerical performances are actually obtained with lower magnitudes, as illustrated in Section 5.2.

5. Numerical experiments. In this section, we illustrate the performance of procedure (3.7) and procedure (3.9) both on simulated and real-world datasets, and compare them to several state-of-the-art baselines described below. Experiments are done using the open-source `tick` library, available at <https://x-datainitiative.github.io/tick/>, notebooks allowing to reproduce our experiments are available on demand to the authors.

5.1. *Considered procedures.* We consider the following baselines, featuring the best methods available in literature for the joint problem of outlier detection and estimation of the regression coefficients, together with the methods introduced in this paper.

E-SLOPE. It is procedure (3.9). The weights used in both SLOPE penalizations are given by (4.3), with $q = 5\%$ (target FDR), except in low-dimensional setting where we do not apply any penalization on β . Similar results for $q = 10\%$ and $q = 20\%$ are provided in Appendix E.

E-LASSO. This is Extended LASSO from [23], that uses two dedicated ℓ_1 -penalizations for β and μ with respective tuning parameters $\lambda_\beta = 2\sigma\sqrt{\log p}$ and $\lambda_\mu = 2\sigma\sqrt{\log n}$.

IPOD. This is (soft-)IPOD from [39]. It relies on a nice trick, based on a QR decomposition of X . Indeed, write $X = QR$ and let P be formed by $n - p$ column vectors completing the column vectors of Q into an orthonormal basis, and introduce $\tilde{y} = P^\top y \in \mathbb{R}^{n-p}$. Model (2.1) can be rewritten as $\tilde{y} = P^\top \mu^* + \varepsilon$, a new high-dimensional linear model, where only μ^* is to be estimated. The IPOD then considers the Lasso procedure applied to this linear model, and a BIC criterion is used to choose the tuning parameter of ℓ_1 -penalization. Note that this procedure, which involves a QR decomposition of X , makes sense only for p significantly smaller than n , so that we do not report the performances of IPOD on simulations with a large p .

LassoCV. Same as IPOD but with tuning parameter for the penalization of individual intercepts chosen by cross-validation. As explained above, cross-validation is doomed to fail in the considered model, but results are shown for the sake of completeness.

SLOPE. It is SLOPE applied to the concatenated problem, namely $y = Z\gamma^* + \varepsilon$, where Z is the concatenation of X and I_n and γ^* is the concatenation of β^* and μ^* . We use a single SLOPE penalization on γ , with weights equal to (4.3). We report the performances of this procedure only in high-dimensional experiments, since it always penalizes β . This is considered mostly to illustrate the fact that working on the concatenated problem is indeed a bad idea, and that two distinct penalizations must be used on β and μ .

Note that the difference between IPOD and E-LASSO is that, as explained in [23], the weights used for E-LASSO to penalize μ (and β in high-dimension) are fixed, while the weights in IPOD are data-dependent. Another difference is that IPOD do not extend well to a high-dimensional setting, since its natural extension (considered in [39]) is a thresholding rule on the the concatenated problem, which is, as explained before, and as illustrated in our numerical experiments, poorly performing. Another problem is that there is no clear extension of the modified BIC criterion proposed in [39] for high-dimensional problems.

The tuning of the SLOPE or ℓ_1 penalizations in the procedure described above require the knowledge of the noise level. We overcome this simply by plugging in (4.3) or wherever it is necessary a robust estimation of the variance: we first fit a Huber regression model, and apply a robust estimation

of the variance of its residuals. All procedures considered in our experiments use this same variance estimate.

REMARK 5.1. *The noise level can be estimated directly by the Huber regression since in our simulations $p < n$. When p is comparable or larger than n and the signal (both β^* and μ^*) is sufficiently sparse one can jointly estimate the noise level and other model parameters in the spirit of scaled LASSO [43]. The corresponding iterative procedure for SLOPE was proposed and investigated in [5] in the context of high-dimensional regression with independent regressors. However, we feel that the problem of selection of the optimal estimator of σ in ultra-high dimensional settings still requires a separate investigation and we postpone it for a further research.*

5.2. *Simulation settings.* The matrix X is simulated as a matrix with i.i.d row distributed as $\mathcal{N}(0, \Sigma)$, with Toeplitz covariance $\Sigma_{i,j} = \rho^{|i-j|}$ for $1 \leq i, j \leq p$, with moderate correlation $\rho = 0.4$. Some results with higher correlation $\rho = 0.8$ are given in Appendix E. The columns of X are normalized to 1. We simulate $n = 5000$ observations according to model (2.1) with $\sigma = 1$ and $\beta_i^* = \sqrt{2 \log p}$. Two levels of magnitude are considered for μ^* : *low-magnitude*, where $\mu_i^* = \sqrt{2 \log n}$ and *large-magnitude*, where $\mu_i^* = 5\sqrt{2 \log n}$. In all reported results based on simulated datasets, the sparsity of μ^* varies between 1% to 50%, and we display the averages of FDR, MSE and power over 100 replications.

Setting 1 (low-dimension). This is the setting described above with $p = 20$. Here $\beta_1^* = \dots = \beta_{20}^* = \sqrt{2 \log 20}$.

Setting 2 (high-dimension). This is the setting described above with $p = 1000$ and a sparse β^* with sparsity $k = 50$, with non-zero entries chosen uniformly at random.

5.3. *Metrics.* In our experiments, we report the ‘‘MSE coefficients’’, namely $\|\hat{\beta} - \beta^*\|_2^2$ and the ‘‘MSE intercepts’’, namely $\|\hat{\mu} - \mu^*\|_2^2$. We report also the FDR (4.1) and the Power (4.2) to assess the procedures for the problem of outliers detection, where the expectations are approximated by averages over 100 simulations.

5.4. *Results and conclusions on simulated datasets.* We comment the displays provided in Figures 1, 2 and 3 below. On Simulation Setting 2 we only display results for the low magnitude case, since it is the most challenging one.

- In Figures 1 and 2, LassoCV is very unstable, which is expected since cross-validation cannot work in the considered setting: data splits are highly non-stationary, because of the significant amount of outliers.
- In the low dimensional setting, our procedure $E - SLOPE$ allows for almost perfect FDR control and its MSE is the smallest among all considered methods. Note that in this setting the MSE is plotted after debiasing the estimators, performing ordinary least squares on the selected support.
- In the sparse (on β) high dimensional setting with correlated regressors, E-SLOPE allows to keep FDR below the nominal level even when the outliers consist 50% of the total data points. It also allows to maintain a small MSE and high power. The only procedure that improves E-SLOPE in terms of MSE for μ is SLOPE in Figure 3, at the cost of a worse FDR control.
- E-SLOPE provides a massive gain of power compared to previous state-of-the-art procedures (power is increased by more than 30%) in settings where outliers are difficult to detect.

5.5. *PGA/LPGA dataset.* This dataset contains Distance and Accuracy of shots, for PGA and LPGA players in 2008. This will allow us to visually compare the performance of IPOD, E-LASSO and E-SLOPE. Our data contain 197 points corresponding to PGA (men) players, to which we add 8 points corresponding to LPGA (women) players, injecting outliers. We apply SLOPE and LASSO on μ with several levels of penalization. This leads to the “regularization paths” given in the top plots of Figure 4, that shows the value of the 205 sample intercepts $\hat{\mu}$ as a function of the penalization level used in SLOPE and LASSO. Vertical lines indicate the choice of the parameter according the corresponding method (E-SLOPE, E-LASSO, IPOD). We observe that E-SLOPE correctly discover the confirmed outliers (women data), together with two men observations that can be considered as outliers in view of the scatter plot. IPOD procedure does quite good, with no false discovery, but misses some real outliers (women data) and the suspicious point detected by E-SLOPE. E-LASSO does not make any false discovery but clearly reveals a lack of power, with only one discovery.

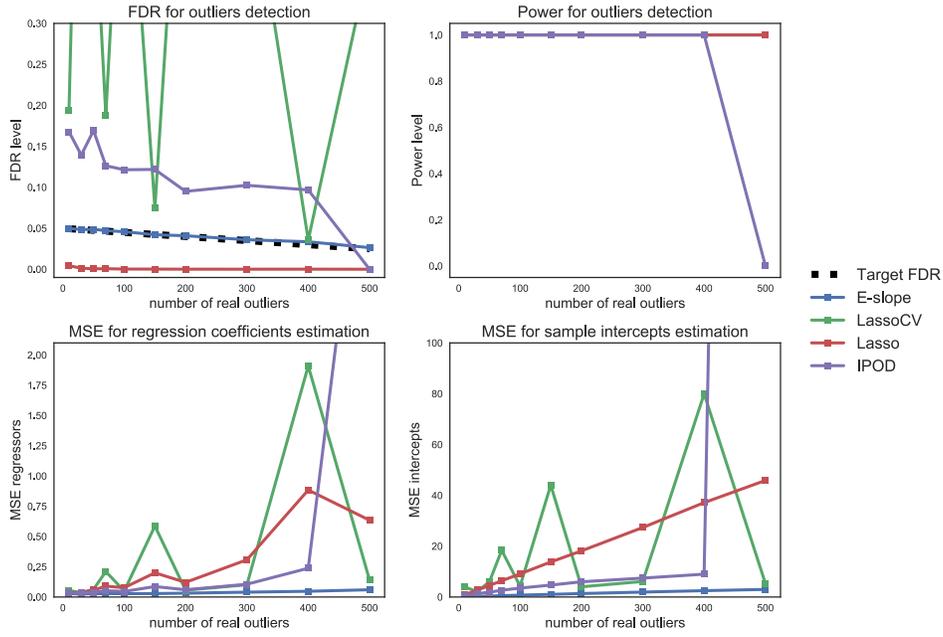


Fig 1: Results for Simulation Setting 1 with high-magnitude outliers. First row gives the FDR (left) and power (right) of each considered procedure for outliers discoveries. Second row gives the MSE for regressors (left) and intercepts (right). E-SLOPE gives perfect power, is the only one to respect the required FDR, and provides the best MSEs.

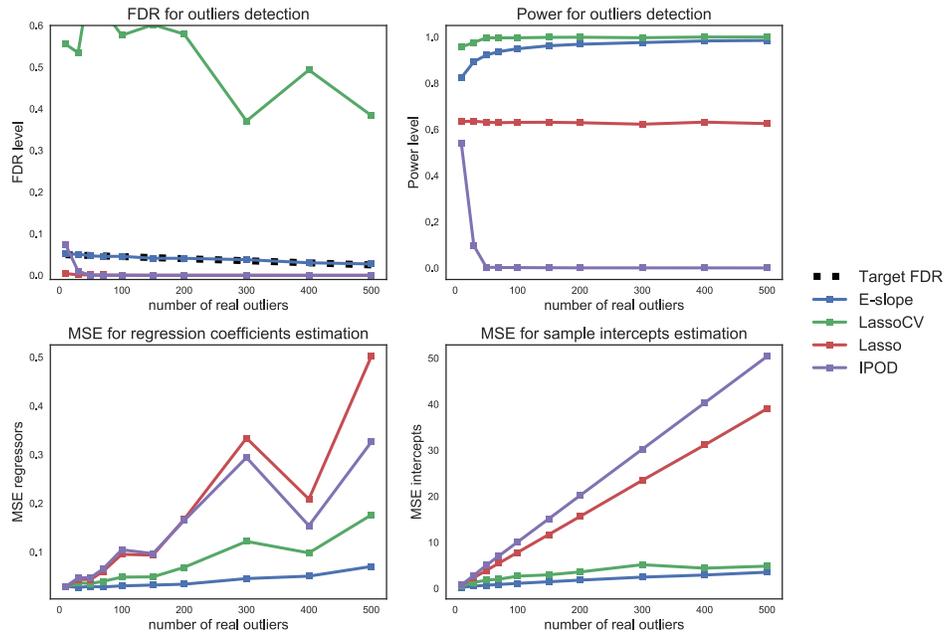


Fig 2: Results for Simulation Setting 1 with low-magnitude outliers. First row gives the FDR (left) and power (right) of each considered procedure for outliers discoveries. Second row gives the MSE for regressors (left) and intercepts (right). Once again E-SLOPE nearly gives the best power, but is the only one to respect the required FDR, and provides the best MSEs.

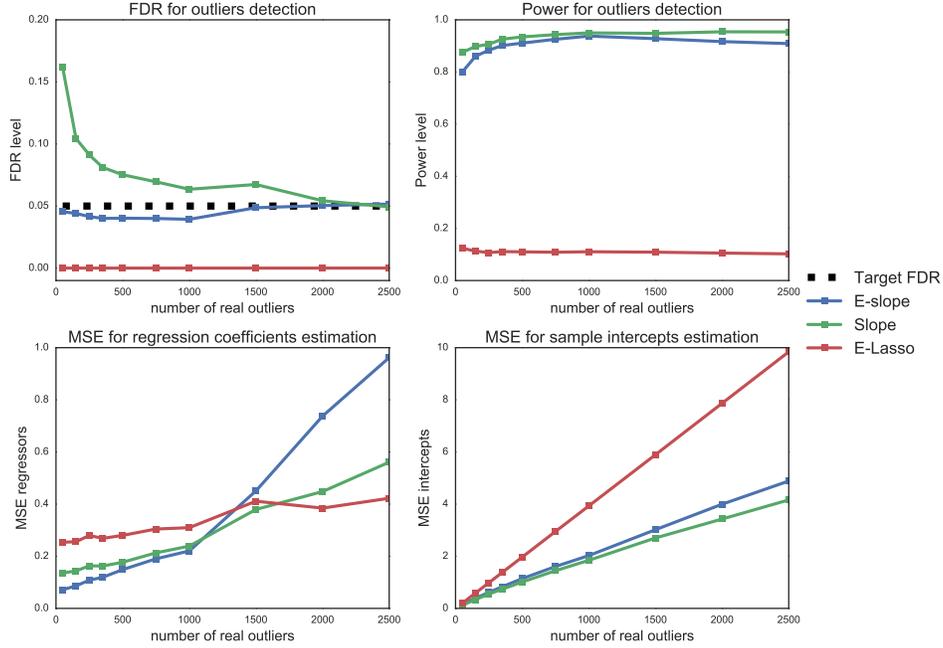


Fig 3: Results for Simulation Setting 2 with low-magnitude outliers. First row gives the FDR (left) and power (right) of each considered procedure for outliers discoveries. Second row gives the MSE for regressors (left) and intercepts (right). Once again E-SLOPE nearly gives the best power, but is the only one to respect the required FDR. It gives the best MSE for outliers estimation, and is competitive for regressors estimation. All procedures have a poor MSE when the number of outliers is large, since the simulation setting considered in this experiment is hard: low-magnitude outliers and high-dimension.

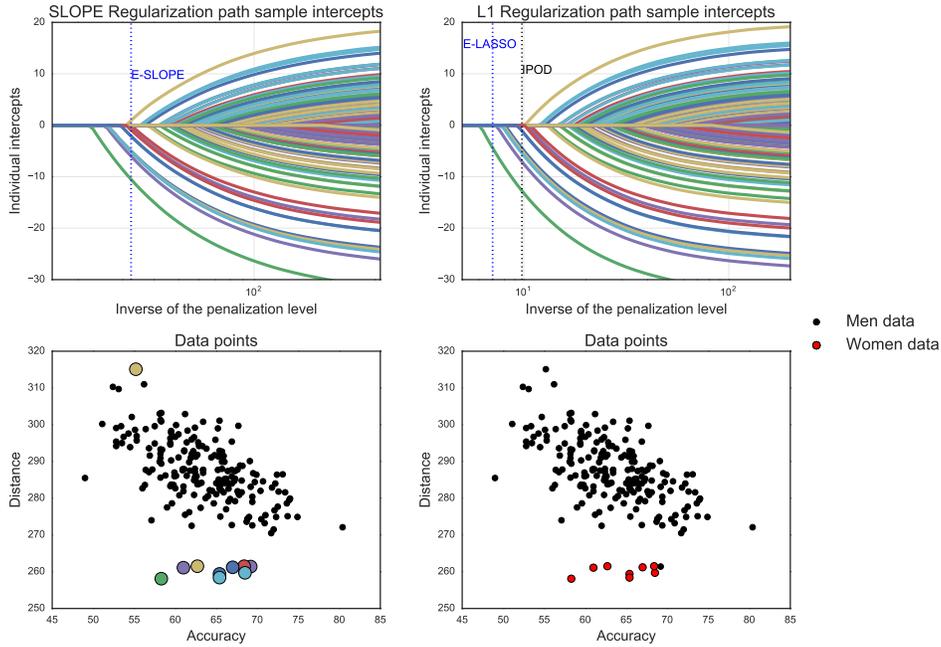


Fig 4: PGA/LPGA dataset: top plots show the regularization paths for both types of penalization, bottom-left plot is a scatter plot of the data, with colored points corresponding to the discoveries made by E-SLOPE, bottom-right plot show the original data and the true outliers.

5.6. *Retail Sales Data.* This dataset is from the U.S. census Bureau, for year 1992. The informations contained in it are the per capita retail sales of 845 US counties (in \$1000s). It also contains five covariates: the per capita retail establishments, the per capita income (in \$1000s), per capita federal expenditures (in \$1000s), and the number of males per 100 females. No outliers are known, so we artificially create outliers by adding a small amount (magnitude 8, random sign) to the retail sales of counties chosen uniformly at random. We consider various scenarii (from 1% to 20% of outliers) and compute the false discovery proportion and the power. Figure 5 below summarizes the results for the three procedures.

The results are in line with the fact that E-SLOPE is able to discover more outliers than its competitors. E-SLOPE has the highest power, and the FDP remains under the target level.

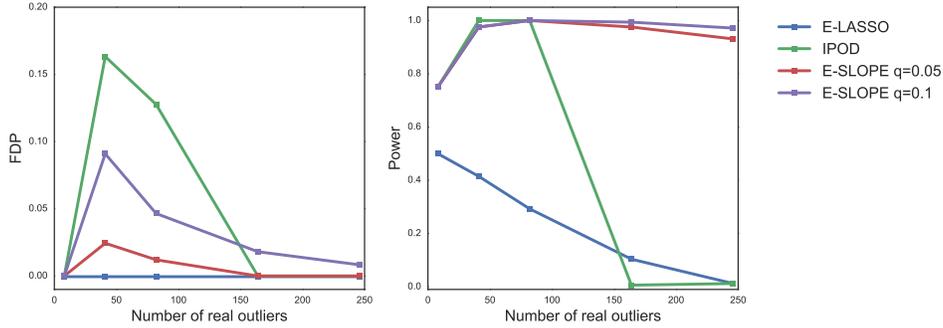


Fig 5: *Left*: False Discovery proportion, E-SLOPE remains under the target level; *Right*: power, E-SLOPE performs better than the competitors.

5.7. *Colorectal Cancer Data.* We consider whole exome sequencing data for 47 primary colorectal cancer tumors, characterized by a global genomic instability affecting repetitive DNA sequences (also known as microsatellite instable tumors, see [16]). In what follows, we restrict ourselves to repetitive sequences whose base motif is the single nucleotide A, and which are in regulatory regions (following the coding regions) that influence gene expression (UTR3). Same analysis could be run with different base motifs and different regions (exonic, intronic). It has been shown in recent publications (see [44]), that the probability of mutation of a sequence is dependent of the length of the repeat. So we fit, after a rescaled probit transformation, our mean-shift model with an intercept and the length of the repeat as covariates. The aim of the analysis is to find two categories of sequences: survivors (multi-satellites that mutated less than expected) and transformers (multi-satellites that mutated more than expected), with the idea that those sequences must play a key role in the cancer development.

We fix the FDR level $\alpha = 5\%$, results are shown in Figure 6: blue dots are the observed mutation frequencies of each gene among the 47 tumors, plotted as a function of the repeat length of the corresponding gene, our discoveries are highlighted in red.

We made 37 discoveries, and it is of particular interest to see that our procedure select both "obvious" outliers and more "challenging" observations that are discutable with the unaided eye. We also emphasize that with the IPOD procedure and the LASSO procedure described in the previous paragraph, respectively 32 and 22 discoveries were made, meaning that even with this stringent FDR level, our procedure allow us to make about 16% more discoveries than IPOD.

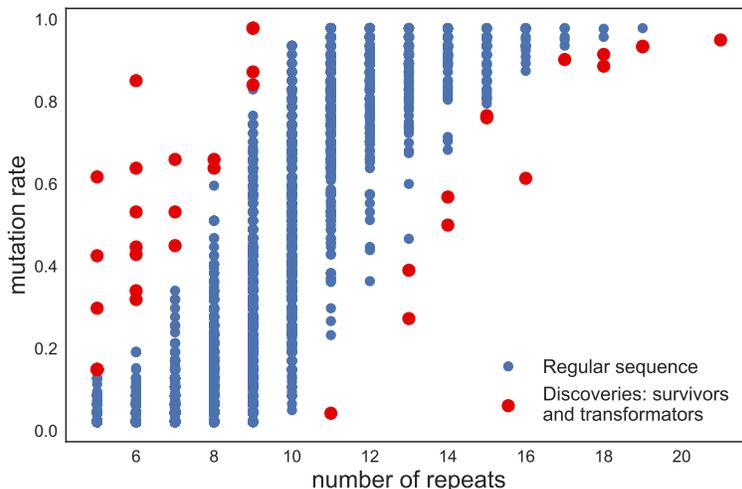


Fig 6: Colorectal Cancer data: 37 discoveries made by E-SLOPE, namely sequence considered by our procedure as transformers or mutators (see main text for details). The procedure selects both “obvious” and more “challenging” observations for the eye. IPOD and Lasso procedures led only to 32 and 22 discoveries, despite the fact that we restricted E-SLOPE to the stringent FDR level of 5%

6. Conclusion. This paper introduces a novel approach for simultaneous robust estimation and outliers detection in the linear regression model. Three main results are provided: optimal bounds for the estimation problem in Section 3, that improve in particular previous results obtained with LASSO penalization [23], and asymptotic FDR control and power consistency for the outlier detection problem in Section 4. To the best of our knowledge, this is the first result involving FDR control in this context.

Our theoretical findings are confirmed on intensive experiments both on real and synthetic datasets, showing that our procedure outperforms existing procedure in terms of power, while maintaining a control on the FDR, even in challenging situations such as low-magnitude outliers, a high-dimensional setting and highly correlated features.

Finally, this work extends the understanding of the deep connection between the SLOPE penalty and FDR control, previously studied in linear regression with orthogonal [5] or i.i.d gaussian [42] features, which distinguishes SLOPE from other popular convex penalization methods.

APPENDIX A: TECHNICAL INEQUALITIES

The following technical inequalities are borrowed from [8], where proofs can be found. Let m, n, p be positive integers. In the following lemma, an inequality for the sorted ℓ_1 -norm J_λ (defined in equation 2.2) is stated.

LEMMA A.1. *For any two $x, y \in \mathbb{R}^m$, and any $s \in 1, \dots, m$ such that $|x|_0 \leq s$ we have*

$$J_\lambda(x) - J_\lambda(y) \leq \Lambda(s) \|x - y\|_2 - \sum_{j=s+1}^m \lambda_j |x - y|_{(j)},$$

where

$$\Lambda(s) = \sqrt{\sum_{j=1}^s \lambda_j^2}.$$

The following lemma gives a preliminary bound for the prediction error in our context, that are the starting point of our proof.

LEMMA A.2. *Let $h : \mathbb{R}^p \rightarrow \mathbb{R}$ be a convex function. Consider a $n \times p$ design matrix X , a vector $\varepsilon \in \mathbb{R}^n$ and define $y = X\beta^* + \varepsilon$ where $\beta^* \in \mathbb{R}^p$. If $\hat{\beta}$ is a solution of the minimization problem $\min_{\beta \in \mathbb{R}^p} (\|y - X\beta\|_2^2 + 2h(\beta))$, then $\hat{\beta}$ satisfies:*

$$\|X\hat{\beta} - X\beta^*\|_2^2 \leq \varepsilon^\top X(\hat{\beta} - \beta^*) + h(\beta^*) - h(\hat{\beta}).$$

PROOF. Because the proof in [8] is more general, we give a proof adapted to our context. Optimality of $\hat{\beta}$ allows to choose v in the subdifferential of h such that

$$0 = X^\top(X\hat{\beta} - y) + v = X^\top(X\hat{\beta} - X\beta^* - \varepsilon) + v.$$

Therefore,

$$\begin{aligned} \|X\hat{\beta} - X\beta^*\|_2^2 &= (\hat{\beta} - \beta^*)^\top X^\top X(\hat{\beta} - \beta^*) \\ &= (\hat{\beta} - \beta^*)^\top (X^\top \varepsilon - v) \\ &= \varepsilon^\top X(\hat{\beta} - \beta^*) + \langle v, \beta^* - \hat{\beta} \rangle. \end{aligned}$$

Now, by definition of subdifferential, $h(\beta^*) \geq h(\hat{\beta}) + \langle v, \beta^* - \hat{\beta} \rangle$. Combining this inequality with the previous equality leads to the conclusion. \square

The following lemma allows to bound the inner product between a white Gaussian noise and any vector. The resulting bound involved the sorted ℓ_1 norm.

LEMMA A.3. *Let $\delta_0 \in (0, 1)$ and let $X \in \mathbb{R}^{n \times p}$ with columns normed to 1. If ε is $\mathcal{N}(0, I_n)$ distributed, then the event*

$$\left\{ \forall u \in \mathbb{R}^p, \varepsilon^\top X u \leq \max(H(u), G(u)) \right\}$$

is of probability at least $1 - \delta_0/2$, where

$$H(u) = (4 + \sqrt{2}) \sum_{j=1}^p |u|_{(j)} \sigma \sqrt{\log(2p/j)}$$

and

$$G(u) = (4 + \sqrt{2}) \sigma \sqrt{\log(1/\delta_0)} \|u\|_2.$$

APPENDIX B: RESULTS RELATED TO GAUSSIAN MATRICES

Inequalities for Gaussian random matrices are needed in what follows. They are stated here for the sake of clarity and we refer the reader to [20] for proof (except bounds B.4 and B.5 that are taken from Lemma 1 in [35]). Again, n and p denote positive integers.

LEMMA B.1. *Let $X \in \mathbb{R}^{n \times p}$ with i.i.d $\mathcal{N}(0, I_p)$ rows. Denote by σ_{max} the largest singular value of X . Then, for all $\tau \geq 0$,*

$$(B.1) \quad \mathbb{P} \left(\frac{\sigma_{max}}{\sqrt{n}} \geq 1 + \sqrt{\frac{p}{n}} + \tau \right) \leq \exp \left(-\frac{n\tau^2}{2} \right).$$

LEMMA B.2. *Concentration inequalities:*

- *Let Z be $\mathcal{N}(0, 1)$ distributed. Then for all $q \geq 0$:*

$$(B.2) \quad \mathbb{P} (|Z| \geq q) \leq \exp \left(-\frac{q^2}{2} \right).$$

- *Let Z_1, Z_2, \dots, Z_p be independent and $\mathcal{N}(0, \sigma^2)$ distributed. Then for all $L > 0$:*

$$(B.3) \quad \mathbb{P} \left(\max_{i=1, \dots, p} |Z_i| > \sigma \sqrt{2 \log p + 2L} \right) \leq e^{-L}.$$

- Let X be $\chi_2(n)$ distributed. Then, for all $x > 0$:

$$(B.4) \quad \mathbb{P}(X - n \geq 2\sqrt{nx} + 2x) \leq \exp(-x).$$

$$(B.5) \quad \mathbb{P}(n - X \geq 2\sqrt{nx}) \leq \exp(-x).$$

The following recent result ([36], Theorem 1) will also be useful.

LEMMA B.3. *Let $X \in \mathbb{R}^{n \times p}$ with i.i.d $\mathcal{N}(0, \Sigma)$ rows. There exists positive constants c and c' such that with probability greater than $1 - c' \exp(-cn)$, we have for all $z \in \mathbb{R}^p$:*

$$(B.6) \quad \frac{\|Xz\|_2}{\sqrt{n}} \geq \frac{1}{4} \sqrt{\lambda_{\min}(\Sigma)} \|z\|_2 - 9 \sqrt{\max_j \Sigma_{jj} \frac{\log p}{n}} \|z\|_1,$$

where $\lambda_{\min}(\Sigma)$ is the lowest eigenvalue of Σ .

APPENDIX C: PROOF OF SECTION 3

This section is devoted to the proof of our main results, stated in Section 3.

C.1. Proof of Theorem 3.3. Define D the diagonal matrix such that $X = X'D$ (D is the diagonal matrix formed by the inverse of the norm of each column of X'). Applying now Lemma B.3 for X' and Dz we obtain for all $z \in \mathbb{R}^p$

$$\begin{aligned} \|Xz\|_2 &\geq \frac{1}{4} \sqrt{\lambda_{\min}(\Sigma)} \|\sqrt{n}Dz\|_2 - 9 \sqrt{\max_j \Sigma_{jj} \frac{\log p}{n}} \|\sqrt{n}Dz\|_1 \\ &\geq \frac{\sqrt{n} \sqrt{\lambda_{\min}(\Sigma)}}{4M} \|z\|_2 - 9 \frac{\sqrt{\max_j \Sigma_{jj} \log p}}{m} \|z\|_1, \end{aligned}$$

with probability greater than $1 - c' \exp(-cn)$, where M and m denote respectively the maximum and minimum of the norms of the columns of X' . Note that for all $1 \leq i \leq p$, the squared norm of the i^{th} column of X' is $\sigma_i^2 \chi_2(n)$ distributed, so using the bounds B.4 and B.5 of Lemma B.2 (respectively with $x = n$ and $x = n/16$), together with a union bound we obtain that with probability greater than $1 - ne^{-n} - ne^{-n/16}$

$$M \leq (\max_j \Sigma_{jj}) \sqrt{5n}, \quad m \geq (\min_j \Sigma_{jj}) \sqrt{\frac{n}{2}},$$

and we eventually obtain

$$(C.1) \quad \|Xz\|_2 \geq \frac{\sqrt{\lambda_{\min}(\Sigma)}}{4\sqrt{5} \max_j \Sigma_{jj}} \|z\|_2 - \frac{9}{\min_j \Sigma_{jj}} \sqrt{\frac{2 \log p \max_j \Sigma_{jj}}{n}} \|z\|_1.$$

Let define $v = [\beta^\top, \mu^\top]^\top \in \mathbb{R}^{p+n} \in \mathcal{C}(k, s, c_0)$ (see Definition 3.3). Then,

$$(C.2) \quad \|\beta\|_1 \leq \sum_{j=1}^p \frac{\tilde{\lambda}_j}{\tilde{\lambda}_p} |\beta|_{(j)} \leq (1 + c_0) \left(\sqrt{k} \|\beta\|_2 + \sqrt{s} \|\mu\|_2 \right).$$

Thus we obtain

$$(C.3) \quad \|\beta\|_1 \leq (1 + c_0) \left(\sqrt{k} \|\beta\|_2 + \sqrt{s} \|\mu\|_2 \right).$$

Injecting (C.3) in (C.1) applied to the vector β now leads to

$$(C.4) \quad \begin{aligned} \|X\beta\|_2 + \|\mu\|_2 &\geq \|\beta\|_2 \left(\frac{\sqrt{\lambda_{\min}(\Sigma)}}{4\sqrt{5} \max_j \Sigma_{jj}} - \frac{9}{\min_j \Sigma_{jj}} (1 + c_0) \sqrt{\frac{2k(\log p) \max_j \Sigma_{jj}}{n}} \right) \\ &\quad + \|\mu\|_2 \left(1 - \frac{9}{\min_j \Sigma_{jj}} (1 + c_0) \sqrt{\frac{2s(\log p) \max_j \Sigma_{jj}}{n}} \right). \end{aligned}$$

For n large enough as explicited in the assumption of Theorem 3.3, Equation (C.4) turns to

$$\|X\beta\|_2 + \|\mu\|_2 \geq \frac{\sqrt{\lambda_{\min}(\Sigma)}}{8\sqrt{5} \max_j \Sigma_{jj}} \|\beta\|_2 + \frac{1}{2} \|\mu\|_2,$$

and thus, using the fact that $2(a^2 + b^2) \geq (a + b)^2$,

$$(C.5) \quad \|X\beta\|_2^2 + \|\mu\|_2^2 \geq \min \left\{ \frac{\lambda_{\min}(\Sigma)}{128 \times 5 (\max_j \Sigma_{jj})^2}, \frac{1}{8} \right\} \|v\|_2^2.$$

Now if $v = [\beta^\top, \mu^\top]^\top \in \mathbb{R}^{p+n} \in \mathcal{C}^p(s, c_0)$, Equation (C.1) together with the inequality $\|\beta\|_1 \leq \sqrt{p} \|\beta\|_2$ lead to

$$\|X\beta\|_2 + \|\mu\|_2 \geq \|\beta\|_2 \left(\frac{\sqrt{\lambda_{\min}(\Sigma)}}{4\sqrt{5} \max_j \Sigma_{jj}} - \frac{9}{\min_j \Sigma_{jj}} \sqrt{\frac{2p \log p \max_j \Sigma_{jj}}{n}} \right) + \|\mu\|_2,$$

and we conclude as above. Thus the first part of the theorem is satisfied.

Now, we must lower bound the scalar product $\langle X\beta, \mu \rangle$.

Divide $\{1, \dots, p\} = T_1 \cup T_2 \cup \dots \cup T_t$ with T_i ($1 \leq i \leq t-1$) of cardinality k'

containing the support of the k' largest absolute values of $b_{(\cup_{j=1}^{i-1} T_j)^c}$ and T_t of cardinality $k'' \leq k'$ the support of the remaining values. Divide in the same way $\{1, \dots, n\} = S_1 \cup S_2 \cup \dots \cup S_q$ (of cardinalities $s', \dots, s', s'' \leq s'$) with respect to the largest absolute values of μ (k' and s' to be chosen later). We use this to lower bound the scalar product:

$$|\langle X\beta, \mu \rangle| = |\langle X'D\beta, \mu \rangle| \leq \sum_{i=1}^q \sum_{j=1}^t |\langle X'_{S_i, T_j}(D\beta)_{T_j}, \mu_{S_i} \rangle|,$$

so

$$(C.6) \quad |\langle X\beta, \mu \rangle| \leq \max_{i,j} \|X'_{S_i, T_j}\|_2 \frac{1}{m} \sum_{j=1}^t \|\beta_{T_j}\|_2 \sum_{i=1}^q \|\mu_{S_i}\|_2,$$

where we recall that m is the minimal value of the column norms of X' . According to Lemma B.1, conditionnally on S_i and T_j , we have with probability greater than $1 - \exp(-n\tau^2/2)$,

$$\|X'_{S_i, T_j}\|_2 \leq \|\Sigma_{T_j, T_j}^{1/2}\|(\sqrt{k'} + \sqrt{s'} + \sqrt{s'}\tau) \leq \sqrt{\lambda_{max}(\Sigma)}(\sqrt{k'} + \sqrt{s'} + \sqrt{s'}\tau).$$

Considering all possibilities for S_i and T_j , we have with probability greater than $1 - \binom{p}{k'} \binom{n}{s'} e^{-n\tau^2/2}$,

$$(C.7) \quad \max_{i,j} \|X'_{S_i, T_j}\|_2 \leq \sqrt{\lambda_{max}(\Sigma)}(\sqrt{k'} + (1 + \tau)\sqrt{s'}).$$

Moreover, thanks to the decreasing value along the subset T_j we can use the trick of [31], writing for all $j \in \{1, \dots, t-1\}$ and all $x \in \{1, \dots, |T_{j+1}|\}$:

$$|(\beta_{T_{j+1}})_x| \leq \frac{\|\beta_{T_j}\|_1}{|T_j|}.$$

Squaring this inequality and summing over x gives:

$$\|\beta_{T_{j+1}}\|_2^2 \leq \frac{\|b_{T_j}\|_1^2 |T_{j+1}|}{|T_j|} \leq \frac{\|b_{T_j}\|_1^2}{|T_j|} = \frac{\|\beta_{T_j}\|_1^2}{k'}.$$

Then,

$$\sum_{j=1}^t \|\beta_{T_j}\|_2 \leq \|\beta\|_2 + \sum_{j=2}^t \|\beta_{T_j}\|_2 \leq \|\beta\|_2 + \frac{1}{\sqrt{k'}} \sum_{j=1}^{t-1} \|\beta_{T_j}\|_1 \leq \|\beta\|_2 + \frac{1}{\sqrt{k'}} \|\beta\|_1,$$

and so

$$(C.8) \quad \sum_{j=1}^t \|\beta_{T_j}\|_2 \leq \|\beta\|_2 + \frac{1}{\sqrt{k'}} \sum_{j=1}^p \frac{\tilde{\lambda}_j}{\tilde{\lambda}_p} |\beta|_{(j)}.$$

In the same way we obtain:

$$(C.9) \quad \sum_{i=1}^q \|\mu_{S_i}\|_2 \leq \|\mu\|_2 + \frac{1}{\sqrt{s'}} \sum_{j=1}^n \frac{\lambda_j}{\lambda_n} |\mu|_{(j)}.$$

Now if $v = [\beta^\top, \mu^\top]^\top \in \mathbb{R}^{p+n} \in \mathcal{C}(k, s, c_0)$,

$$\begin{aligned} \sum_{j=1}^t \|\beta_{T_j}\|_2 \sum_{i=1}^q \|\mu_{S_i}\|_2 &\leq (\|\beta\|_2 + \frac{1}{\sqrt{k'}}(1+c_0)(\sqrt{k}\|\beta\|_2 + \sqrt{s}\|\mu\|_2)) \\ &\quad \times (\|\mu\|_2 + \frac{1}{\sqrt{s'}}(1+c_0)(\sqrt{k}\|\beta\|_2 + \sqrt{s}\|\mu\|_2)) \\ &\leq (2\|\beta\|_2 + \|\mu\|_2)(2\|\mu\|_2 + \|\beta\|_2) \\ &\leq 2\|v\|_2^2 + 5\|\mu\|_2\|\beta\|_2 \\ &\leq 5\|v\|_2^2, \end{aligned}$$

where we chose $k' = s' = (1+c_0)^2(k \vee s)$. Combining this last inequality with Equations (C.6) and (C.7), and using again that $m \geq \min_j \Sigma_{jj} \sqrt{n/2}$ with probability greater than $1 - ne^{-n/16}$, lead to

$$(C.10) \quad |\langle X\beta, \mu \rangle| \leq \frac{\sqrt{\lambda_{\max}(\Sigma)}}{\min_j \Sigma_{jj}} (2 + \tau) \sqrt{\frac{2s'}{n}} 5 \|v\|_2^2.$$

Note that with this choice of s' and k' , the assumptions on n and the constant C' defined in the theorem lead to $\binom{p}{k'} \leq (ep/k')^{k'} \leq \exp(n/C')$, and $\binom{n}{s'} \leq (en/s')^{s'} \leq \exp(n/C')$, so we have Equation (C.7) with probability greater than $1 - \exp(-n(\tau^2/2 - 2C'^{-1}))$. With the specific assumption on n in the statement of the theorem, the term in the right part of Equation (C.10) is small enough to obtain:

$$(C.11) \quad 2|\langle Xb, u \rangle| \leq \min \left\{ \frac{\lambda_{\min}(\Sigma)}{256 \times 5 \max_j \Sigma_{jj}}, \frac{1}{16} \right\} \|v\|_2^2$$

Eventually, if $v = [\beta^\top, \mu^\top]^\top \in \mathbb{R}^{p+n} \in \mathcal{C}^p(s, c_0)$, Equation (C.9) still holds, and combining it with Equation (C.7) and Equation (C.6) with $t = 1$ leads

to:

$$(C.12) \quad |\langle X\beta, \mu \rangle| \leq \frac{\sqrt{\lambda_{\max}(\Sigma)}}{\min_j \Sigma_{jj}} (\sqrt{p} + (1 + \tau)\sqrt{s'}) \sqrt{\frac{2}{n}} \|\beta\|_2 \\ \times \left(\|\mu\|_2 + \frac{1}{\sqrt{s'}} (1 + c_0) (\sqrt{p} \|\beta\|_2 + \sqrt{s} \|\mu\|_2) \right).$$

Choosing $s' = (1 + c_0)^2(p \vee s)$,

$$|\langle X\beta, \mu \rangle| \leq \frac{\sqrt{\lambda_{\max}(\Sigma)}}{\min_j \Sigma_{jj}} (2 + \tau) \sqrt{\frac{2s'}{n}} \|\beta\|_2 (2\|\mu\|_2 + \|\beta\|_2) \\ \leq \frac{\sqrt{\lambda_{\max}(\Sigma)}}{\min_j \Sigma_{jj}} (2 + \tau) \sqrt{\frac{2s'}{n}} 2\|v\|_2^2.$$

We conclude as above, thus leading to the second part of the theorem.

C.2. Proof of Theorem 3.4. We will actually show a slightly more general result. Let R be any subset of cardinality r containing the support of the true parameter μ^* and I_R be the matrix obtained by extracting columns with indices in R from the identity matrix. We consider the following minimization:

$$\hat{\beta}, \hat{\mu} = \operatorname{argmin}_{\beta, \mu} \|y - X\beta - I_R\mu\|_2^2 + 2\rho J_{\lambda^{[r]}}(\mu),$$

where $\lambda^{[r]}$ contains the first r terms of the sequence of weights defined in Section 3. Obviously, the theorem will result from the case $P = \{1, \dots, n\}$. Note that $\hat{\mu}$ belongs to \mathbb{R}^r .

Defining $b = \hat{\beta} - \beta^*$ and $u = I_R(\hat{\mu} - \mu_R^*)$ where μ_R^* denotes the vector extracted from μ^* selecting coordinates corresponding to indices in R (note that the eliminated coordinates are zeros), we can apply Lemma A.2 to obtain:

$$\|Xb + u\|_2^2 \leq \varepsilon^\top (Xb + u) + \rho J_{\lambda^{[r]}}(\mu^*) - \rho J_{\lambda^{[r]}}(\hat{\mu}) \\ = \varepsilon^\top (Xb + u) + \rho J_{\lambda}(I_R\mu_R^*) - \rho J_{\lambda}(I_R\hat{\mu}).$$

Note that it is crucial to have $\operatorname{supp}(\mu^*) \subset R$ in order to write $\mu^* = I_R\mu_R^*$. Applying now Lemma A.1 we obtain:

$$(C.13) \quad \|Xb + u\|_n^2 \leq \varepsilon^\top (Xb + u) + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)}),$$

where $\Lambda(s)$ is defined as $\sqrt{\sum_{j=1}^s \lambda_j^2}$. Hence, using Cauchy-Schwarz inequality we get:

$$\|Xb + u\|_2^2 \leq \|X^\top \varepsilon\|_2 \|b\|_2 + \varepsilon^\top u + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)}).$$

Then, by Lemma A.3, with probability greater than $1 - \delta_0/2$ we have (the last inequality is used for the sake of simplicity):

$$\varepsilon^\top u \leq \max(H(u), G(u)) \leq H(u) + G(u),$$

with $H(u)$ and $G(u)$ defined in Lemma A.3. Additionally, $\frac{1}{\sigma^2} \|X^\top \varepsilon\|_2^2$ follows a χ^2 law with p degrees of freedom, so by the third point in Lemma B.2 with $x = Lp$ this provides, choosing $\delta_0 = (s/2n)^s$, that with probability greater than $1 - \frac{1}{2}(s/2n)^s - \exp(-Lp)$:

$$\begin{aligned} \|Xb + u\|_2^2 &\leq c_L \sigma \sqrt{p} \|b\|_2 + H(u) + G(u) + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)}) \\ &\leq c_L \sigma \sqrt{p} \|b\|_2 + \frac{\rho}{2} \sum_{j=1}^n \lambda_j |u|_{(j)} \\ &\quad + \frac{\rho}{2} \sqrt{s \log(2n/s)} \|u\|_2 + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)}) \\ &\leq c_L \sigma \sqrt{p} \|b\|_2 + (2\rho\Lambda(s) \|u\|_2 - \frac{\rho}{2} \sum_{j=s+1}^n \lambda_j |u|_{(j)}), \end{aligned}$$

where $c_L = \sqrt{1 + 2L + 2\sqrt{L}}$ and where we used Equation (C.15) to obtain the last inequality. The fact that the left part of the last inequality is positive gives:

$$\sum_{j=1}^n \lambda_j |u|_{(j)} \leq \sum_{j=s+1}^n \lambda_j |u|_{(j)} + \Lambda(s) \|u\|_2 \leq \frac{2}{\rho} c_L \sigma \sqrt{p} \|b\|_2 + 5\Lambda(s) \|u\|_2,$$

where the left part of the inequality is obtained using Cauchy-Schwarz inequality. Hence,

$$(C.14) \quad \sum_{j=1}^n \frac{\lambda_j}{\lambda_n} |u|_{(j)} \leq \frac{2c_L}{\rho} \sqrt{\frac{p}{\log 2}} \|b\|_2 + 5\sqrt{\frac{s \log(2en/s)}{\log 2}} \|u\|_2,$$

where we used the right part of the inequality (C.15). Choosing $L = 1$ lead to $c_L = \sqrt{5}$, and reminding that $\rho \geq 2(4 + \sqrt{2})$ we conclude that $[b^\top, u^\top]^\top \in \mathcal{C}^p(s_1, 4)$ (see Definition 3.3) with $s_1 = \frac{s \log(2en/s)}{\log 2}$. Therefore, by Condition 3.2 and the definition of κ therein :

$$\begin{aligned} 2 \|Xb + u\|_2^2 &\leq 2\sqrt{5}\sigma\sqrt{p} \|b\|_2 + 4\rho\Lambda(s) \|u\|_2 \\ &\leq \frac{5\sigma^2}{\kappa^2} p + \kappa^2 \|b\|_2^2 + \frac{4\rho^2\Lambda(s)^2}{\kappa^2} + \kappa^2 \|u\|_2^2 \\ &\leq \frac{4\rho^2}{\kappa^2} \Lambda(s)^2 + \frac{5\sigma^2}{\kappa^2} p + \kappa^2 \|v\|_2^2 \\ &\leq \frac{4\rho^2}{\kappa^2} \Lambda(s)^2 + \frac{5\sigma^2}{\kappa^2} p + \|Xb + u\|_2^2. \end{aligned}$$

Thus,

$$\|Xb + u\|_2^2 \leq \frac{4\rho^2}{\kappa^2} \Lambda(s)^2 + \frac{5\sigma^2}{\kappa^2} p,$$

and

$$\|b\|_2^2 + \|u\|_2^2 \leq \frac{4\rho^2}{\kappa^4} \Lambda(s)^2 + \frac{5\sigma^2}{\kappa^4} p.$$

The proof of Theorem 3.4 concludes by the classical inequalities [8]

$$(C.15) \quad s \log\left(\frac{2n}{s}\right) \leq \sum_{j=1}^s \log\left(\frac{2n}{j}\right) = s \log(2n) - \log(s!) \leq s \log\left(\frac{2en}{s}\right).$$

C.3. Proof of Theorem 3.5. As in the previous proof, the more general version still holds and in the same way we obtained (C.13), with the same definition of b and u , we now have:

$$\|Xb + u\|_2^2 \leq \varepsilon^\top (Xb + u) + \nu(\|\beta^*\|_1 - \|\hat{\beta}\|_1) + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)}).$$

With T being the support of the true regression vector β^* we have, using the triangle inequality:

$$\|\beta^*\|_1 - \|\hat{\beta}\|_1 = \|\beta_T^*\|_1 - \|b + \beta^*\|_1 = \|\beta_T^*\|_1 - \|b_T + \beta_T^*\|_1 - \|b_{T^c}\|_1 \leq \|b_T\|_1 - \|b_{T^c}\|_1.$$

Hence we can write:

$$\begin{aligned}
\|Xb + u\|_2^2 &\leq \|X^\top \varepsilon\|_\infty \|b\|_1 + \nu (\|b_T\|_1 - \|b_{T^c}\|_1) + \varepsilon^\top u \\
&\quad + \rho \Lambda(s) \|u\|_2 - \rho \sum_{j=s+1}^n \lambda_j |u|_{(j)} \\
&\leq \|b_T\|_1 (\nu + \|X^\top \varepsilon\|_\infty) - \|b_{T^c}\|_1 (\nu - \|X^\top \varepsilon\|_\infty) + \varepsilon^\top u \\
&\quad + \rho \Lambda(s) \|u\|_2 - \rho \sum_{j=s+1}^n \lambda_j |u|_{(j)}.
\end{aligned}$$

With the choice $\nu = 4\sigma\sqrt{\log p}$ we have $\|X^\top \varepsilon\|_\infty \leq \nu/2$ according to Lemma B.2, with probability greater than $1 - \frac{1}{p}$. Using again Lemma A.3 to bound $\varepsilon^\top u$, we obtain that with probability greater than $1 - \frac{1}{2} \left(\frac{s}{2n}\right)^s - \frac{1}{p}$:

(C.16)

$$\|Xb + u\|_2^2 \leq \|b_T\|_1 (6\sigma\sqrt{\log p}) - \|b_{T^c}\|_1 (2\sigma\sqrt{\log p}) + 2\rho\Lambda(s) \|u\|_2 - \frac{\rho}{2} \sum_{j=s+1}^n \lambda_j |u|_{(j)}.$$

The fact that the left part of the inequality is positive gives:

$$\frac{4}{\rho}\sigma\sqrt{\log p} \|b_{T^c}\|_1 + \sum_{j=s+1}^n \lambda_j |u|_{(j)} \leq \frac{12}{\rho}\sigma\sqrt{\log p} \|b_T\|_1 + 4\Lambda(s) \|u\|_2,$$

and using Cauchy-Schwarz inequality, this leads to:

$$\frac{4}{\rho}\sigma\sqrt{\log p} \|b\|_1 + \sum_{j=1}^n \lambda_j |u|_{(j)} \leq \frac{16}{\rho}\sigma\sqrt{k \log p} \|b\|_2 + 5\Lambda(s) \|u\|_2$$

Eventually we obtain, because $\lambda_n = \sigma\sqrt{\log 2}$ and $\sqrt{\log p} \geq \frac{\rho \log 2}{4}$:

(C.17)

$$\|b\|_1 + \sum_{j=1}^n \frac{\lambda_j}{\lambda_n} |u|_{(j)} \leq \frac{4\sigma\sqrt{\log p}}{\rho\lambda_n} \|b\|_1 + \sum_{j=1}^n \frac{\lambda_j}{\lambda_n} |u|_{(j)} \leq \frac{16\sigma\sqrt{k \log p}}{\rho\lambda_n} \|b\|_2 + \frac{5\Lambda(s)}{\lambda_n} \|u\|_2$$

and the concatenated vector of b and u is therefore in the cone $\mathcal{C}(k_1, s_1, 4)$ with $k_1 = 16k \log p / \log 2$ and $s_1 = s \log(2en/s) / \log 2$. Starting from (C.16),

we obtain, using again κ as the capacity constant in Assumption 3.2:

$$\begin{aligned}
2 \|Xb + u\|_2^2 &\leq \|b_T\|_1 12\sigma\sqrt{\log p} + 4\rho\Lambda(s) \|u\|_2 \\
&\leq 12\sigma\sqrt{k \log p} \|b\|_2 + 4\rho\Lambda(s) \|u\|_2 \\
&\leq \frac{36}{\kappa^2}\sigma^2 k \log p + \kappa^2 \|b\|_2^2 + \frac{4\rho^2}{\kappa^2}\Lambda(s)^2 + \kappa^2 \|u\|_2^2 \\
&\leq \frac{36}{\kappa^2}\sigma^2 k \log p + \frac{4\rho^2}{\kappa^2}\Lambda(s)^2 + \kappa^2 \|v\|_2^2 \\
&\leq \frac{36}{\kappa^2}\sigma^2 k \log p + \frac{4\rho^2}{\kappa^2}\Lambda(s)^2 + \|Xb + u\|_2^2.
\end{aligned}$$

Thus,

$$\|Xb + u\|_2^2 \leq \frac{36}{\kappa^2}\sigma^2 k \log p + \frac{4\rho^2}{\kappa^2}\Lambda(s)^2$$

and using again Assumption 3.2 and the remark after:

$$\|b\|_2^2 + \|u\|_2^2 \leq \frac{36}{\kappa^4}\sigma^2 k \log p + \frac{4}{\kappa^4}\Lambda(s)^2$$

C.4. Proof of Theorem 3.6. In the same way we obtained (C.13), we now have:

$$\|Xb + u\|_2^2 \leq \varepsilon^\top (Xb + u) + \rho(\tilde{\Lambda}(k) \|b\|_2 - \sum_{j=k+1}^p \tilde{\lambda}_j |b|_{(j)}) + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)})$$

We use twice Lemma A.3 to bound $\varepsilon^\top Xb$ and $\varepsilon^\top u$ with $(k/2p)^k$ and $(s/2n)^s$ as respective choices of δ_0 , so that with probability $1 - \frac{1}{2} \left(\frac{s}{2n}\right)^s - \frac{1}{2} \left(\frac{k}{2p}\right)^k$:

$$\begin{aligned}
\|Xb + u\|_2^2 &\leq H(b) + G(b) + H(u) + G(u) \\
&\quad + \rho(\tilde{\Lambda}(k) \|b\|_2 - \sum_{j=k+1}^p \tilde{\lambda}_j |b|_{(j)}) + \rho(\Lambda(s) \|u\|_2 - \sum_{j=s+1}^n \lambda_j |u|_{(j)}) \\
&\leq \frac{\rho}{2} \sum_{j=1}^p \tilde{\lambda}_j |b|_{(j)} + \frac{\rho}{2} \sqrt{k \log(2p/k)} \|b\|_2 + \rho(\tilde{\Lambda}(k) \|b\|_2 - \sum_{j=k+1}^p \tilde{\lambda}_j |b|_{(j)}) \\
&\quad + 2\rho\Lambda(s) \|u\|_2 - \frac{\rho}{2} \sum_{j=s+1}^n \lambda_j |u|_{(j)} \\
&\leq \frac{\rho}{2} 4\tilde{\Lambda}(k) \|b\|_2 - \frac{\rho}{2} \sum_{j=k+1}^p \tilde{\lambda}_j |b|_{(j)} + 2\rho\Lambda(s) \|u\|_2 - \frac{\rho}{2} \sum_{j=s+1}^n \lambda_j |u|_{(j)},
\end{aligned}$$

where we use Equation (C.15) to obtain the last inequality. The left part of the inequality is positive so

$$(C.18) \quad \sum_{j=k+1}^p \tilde{\lambda}_j |b|_{(j)} + \sum_{j=s+1}^n \lambda_j |u|_{(j)} \leq 4\tilde{\Lambda}(k) \|b\|_2 + 4\Lambda(s) \|u\|_2,$$

and

$$(C.19) \quad 2 \|Xb + u\|_2^2 \leq 4\rho\tilde{\Lambda}(k) \|b\|_2 + 4\rho\Lambda(s) \|u\|_2.$$

Equation (C.18) together with the Cauchy-Schwarz inequality lead to

$$(C.20) \quad \sum_{j=1}^p \tilde{\lambda}_j |b|_{(j)} + \sum_{j=1}^n \lambda_j |u|_{(j)} \leq 5\tilde{\Lambda}(k) \|b\|_2 + 5\Lambda(s) \|u\|_2.$$

Combining the equation above with Equation (C.15) show that the concatenated estimator is in $\mathcal{C}(k_1, s_1, 4)$ with s_1 and k_1 as in the statement of the theorem (note that $\tilde{\lambda}_n = \lambda_n = \sigma\sqrt{\log 2}$) and so, noting κ the capacity constant of Assumption 3.2, Equation (C.19) leads to:

$$\begin{aligned} 2 \|Xb + u\|_2^2 &\leq (3 + C)^2 \frac{\tilde{\Lambda}(k)^2}{2\kappa^2} + \kappa^2 \|b\|_2^2 + 4\rho^2 \frac{\Lambda(s)^2}{\kappa^2} + \kappa^2 \|u\|_2^2 \\ &\leq \frac{C'}{\kappa^2} \left(\tilde{\Lambda}(k)^2 + \Lambda(s)^2 \right) + \|Xb + u\|_2^2, \end{aligned}$$

where $C' = 4\rho^2 \vee (3 + C)^2/2$. Finally:

$$\|Xb + u\|_2^2 \leq \frac{C'}{\kappa^2} \left(\tilde{\Lambda}(k)^2 + \Lambda(s)^2 \right),$$

and

$$\|b\|_2^2 + \|u\|_2^2 \leq \frac{C'}{\kappa^4} \left(\tilde{\Lambda}(k)^2 + \Lambda(s)^2 \right).$$

APPENDIX D: PROOF OF THEOREM 4.1

In this section, we give the proof of the asymptotic FDR control presented in Theorem 4.1. In the following, for a given matrix A and a given subset T , A_T denotes the extracted matrix formed by the columns of A with indices in T , whereas $A_{T\cdot}$ denotes the extracted matrix formed by the rows of A with indices in T . For vectors, there is no ambiguity. Moreover, S (of cardinal s) denotes the support of the true parameter μ^* .

We first recall some properties on the dual of the sorted ℓ_1 norm, and also a lemma taken from [42] and stated here without proof:

DEFINITION D.1 ([42]). A vector $a \in \mathbb{R}^n$ is said to majorize $b \in \mathbb{R}^n$ (denoted $b \preceq a$) if they satisfy for all $i \in \{1, \dots, n\}$:

$$|a|_{(1)} + \dots + |a|_{(i)} \geq |b|_{(1)} + \dots + |b|_{(i)}.$$

PROPOSITION D.2 ([5]). Let J_λ be the sorted ℓ_1 norm for a certain non-increasing sequence λ of length n . The unit ball of the dual norm is:

$$\mathcal{C}_\lambda = \{v \in \mathbb{R}^n : v \preceq \lambda\}.$$

LEMMA D.3 ([42], Lemma A.9). Given any constant $\alpha > 1/(1-q)$, suppose $\max\{\alpha s, s+d\} \leq s^* < n$ for any (deterministic) sequence d that diverges to ∞ . Let $\zeta_1, \dots, \zeta_{n-s}$ be i.i.d $\mathcal{N}(0, 1)$. Then

$$(|\zeta|_{(s^*-s+1)}, |\zeta|_{(s^*-s+2)}, \dots, |\zeta|_{(n-s)}) \preceq (\lambda_{s^*+1}^{\text{BH}}, \lambda_{s^*+2}^{\text{BH}}, \dots, \lambda_n^{\text{BH}})$$

with probability approaching one.

We adapt from [42] the definition of a resolvent set below, useful to determine the true support of the mean-shift parameters.

DEFINITION D.4. Let s^* be an integer obeying $s < s^* < n$. The set $S^*(S, s^*)$ is said to be a resolvent set if it is the union of S and the $s^* - s$ indices corresponding to the largest entries of the error term ϵ restricted to \bar{S} .

Let c be any positive constant and fix $s^* \geq s(1+c)/(1-q)$ (q being the target FDR level), so that assumptions of Lemma D.3 are satisfied. For clarity, we denote $S^* = S^*(S, s^*)$. For a resolvent set S^* of cardinality s^* , define the reduced minimization as:

$$(D.1) \quad \beta^{S^*}, \mu^{S^*} = \underset{\beta \in \mathbb{R}^p, \mu \in \mathbb{R}^{s^*}}{\operatorname{argmin}} \left\{ \|y - X\beta - I_{S^*}\mu\|_2^2 + 2\rho J_{\bar{\lambda}}(\beta) + 2\rho J_{\lambda^{[s^*]}}(\mu) \right\},$$

where $\lambda^{[s^*]}$ is the beginning (the first s^* terms) of the sequence of weights in the global problem. Note that a resolvent set contains the support of the true parameter μ^* , so the generalized versions of the main results in Section 3, considered in the proof in Appendix C, hold.

We want to show that the estimator of the unreduced problem $\hat{\mu}$ has null values for coordinates which indices are not in S^* . Precisely, we will show that $\hat{\mu} = I_{S^*}\mu^{S^*}$. The first order conditions for global and reduced minimisation problem above are respectively:

$$(D.2) \quad \begin{cases} X^\top(y - X\hat{\beta} - \hat{\mu}) \in \rho\partial J_{\tilde{\lambda}}(\hat{\beta}) \\ (D.3) \quad y - X\hat{\beta} - \hat{\mu} \in \rho\partial J_{\lambda}(\hat{\mu}) \end{cases}$$

and

$$(D.4) \quad \begin{cases} X^\top(y - X\beta^{S^*} - I_{S^*}\mu^{S^*}) \in \rho\partial J_{\tilde{\lambda}}(\beta^{S^*}) \\ (D.5) \quad I_{S^*}^\top(y - X\beta^{S^*} - I_{S^*}\mu^{S^*}) \in \rho\partial J_{\lambda^{[s^*]}}(\mu^{S^*}) \end{cases}$$

Clearly, Equation (D.4) leads to Equation (D.2) taking $\hat{\beta} = \beta^{S^*}$ and $\hat{\mu} = I_{S^*}\mu^{S^*}$. We must now show that this choice of $\hat{\beta}$ and $\hat{\mu}$ satisfies Equation (D.3).

First, $y - X\hat{\beta} - \hat{\mu}$ must be in the unit ball of the dual norm, that is $y - X\hat{\beta} - \hat{\mu} \preceq \rho\lambda$. Because $y - X\hat{\beta} - \hat{\mu} \in \mathbb{R}^n$ is the concatenation of $I_{S^*}^\top(y - X\hat{\beta} - \hat{\mu})$ and $I_{S^*}^\top(y - X\hat{\beta} - \hat{\mu})$, we must check that S^* satisfy:

$$I_{S^*}^\top(y - X\hat{\beta}^{S^*} - I_{S^*}\hat{\mu}^{S^*}) \preceq \rho\lambda^{-[s^*]},$$

where $\lambda^{-[s^*]}$ is the end of the sequence in the global problem (omitting the first s^* terms). If so, noting that if $a_1 \preceq b_1$ and $a_2 \preceq b_2$ then $a \preceq b$ (with a and b being the respective concatenation of a_1, a_2 and b_1, b_2) and combining it with Equation (D.5) will lead to the belonging at the unit ball of the dual norm.

Equivalently, we must check that

$$y_{\overline{S^*}} - X_{\overline{S^*}, \cdot} \beta^{S^*} \preceq \rho\lambda^{-[s^*]},$$

or also

$$(D.6) \quad X_{\overline{S^*}, \cdot}(\beta^* - \beta^{S^*}) + I_{\overline{S^*}, \varepsilon} \preceq \rho\lambda^{-[s^*]}$$

Lemma D.3, together with the definition of the resolvent set S^* given in Definition D.4, allows us to handle the second term to obtain, with probability tending to one:

$$I_{\overline{S^*}, \varepsilon} \preceq (\lambda^{\text{BH}})^{-[s^*]} \preceq \rho(\lambda^{\text{BH}})^{-[s^*]}.$$

It remains to control the term $X_{\overline{S^*}, \cdot}(\beta^* - \beta^{S^*})$. For our purpose, it is sufficient to show that $\|X_{\overline{S^*}, \cdot}(\beta^* - \beta^{S^*})\|_\infty$ tends to zero when n goes to infinity, because in this case we would have $X_{\overline{S^*}, \cdot}(\beta^* - \beta^{S^*}) \preceq \rho\epsilon(\lambda^{\text{BH}})^{-[s^*]}$ if n is large enough. Thus, let $i \in \{1, \dots, n\}$ and x_i the i^{th} row of X , then we have:

$$|\langle x_i, \beta^* - \beta^{S^*} \rangle| \leq \sum_{j=1}^p |x_{i,j}| |\beta^* - \beta^{S^*}|_j \leq \frac{M}{\sqrt{n}} \|\beta^* - \beta^{S^*}\|_1.$$

Now we distinguish the three cases. For Procedure 3.7, we do not assume sparsity on β so we rely on the Cauchy-Schwarz inequality to obtain

$$|\langle x_i, \beta^* - \beta^{S^*} \rangle| \leq \frac{M}{\sqrt{n}} \sqrt{p} \|\beta^* - \beta^{S^*}\|_2.$$

For Procedure 3.8, Equation (C.17) lead to the bound

$$|\langle x_i, \beta^* - \beta^{S^*} \rangle| \leq \frac{M}{\sqrt{n}} C(k \log p \vee s \log(2en/s)) (\|\beta^* - \beta^{S^*}\|_2 \vee \|\mu^* - \mu^{S^*}\|_2),$$

with C being some positive constant. For Procedure 3.9, Equation (C.20) lead to the bound

$$|\langle x_i, \beta^* - \beta^{S^*} \rangle| \leq \frac{M}{\sqrt{n}} C'(k \log(2ep/k) \vee s \log(2en/s)) (\|\beta^* - \beta^{S^*}\|_2 \vee \|\mu^* - \mu^{S^*}\|_2),$$

with C' being some positive constant.

Therefore the coordinates are uniformly bounded by a quantity tending to zero in each of the three cases of the theorem, thanks to the upper bounds obtained in Section 3. Now it is sufficient to choose n such that $|\langle x_i, \beta^* - \beta^{S^*} \rangle| \leq \rho \epsilon \lambda_n^{\text{BH}}$ (it is important to notice that the right term does not depend on n and equals to $\rho \epsilon \Phi^{-1}(1 - q/2)$) to finally obtain Equation (D.6). Note that Equation (D.6) is the necessary condition for $y - X\beta^{S^*} - I_{S^*}\mu^{S^*}$ to be feasible (meaning in the unit ball \mathcal{C}_λ of the dual norm of J_λ) but this is also sufficient for being in the subdifferential because

$$\partial J_\lambda(x) = \{\omega \in \mathcal{C}_\lambda : \langle \omega, x \rangle = J_\lambda(x)\},$$

and as we have, due to Equation (D.5):

$$\langle I_{S^*}^\top (y - X\beta^{S^*} - I_{S^*}\mu^{S^*}), \mu^{S^*} \rangle = J_{\lambda^{[s^*]}}(\mu^{S^*}),$$

then:

$$\langle y - X\beta^{S^*} - I_{S^*}\mu^{S^*}, I_{S^*}\mu^{S^*} \rangle = J_\lambda(I_{S^*}\mu^{S^*}).$$

Therefore, with probability tending to one, $\hat{\mu} = I_{S^*}\mu^{S^*}$ and in particular

$$(D.7) \quad \text{supp}(\hat{\mu}) \subset S^*.$$

We now show that the support of $\hat{\mu}$ contains the support of μ^* . Considering Equation (D.5) we have in particular the belonging to the unit ball of the dual norm, that is to say:

$$I_{S^*}^\top (y - X\beta^{S^*} - I_{S^*}\mu^{S^*}) \preceq \rho \lambda^{[s^*]}.$$

In particular we have

$$\|I_{S^*}^\top(y - X\beta^{S^*} - I_{S^*}\mu^{S^*})\|_\infty \leq \rho\lambda_1.$$

Having $y = X\beta^* + \mu^* + \varepsilon = X\beta^* + I_{S^*}(\mu^*)_{S^*} + \varepsilon$, the inequality above re-writes as:

$$\|X_{S^*}(\beta^* - \beta^{S^*}) + \mu_{S^*}^* - \mu^{S^*} + I_{S^*}^\top\varepsilon\|_\infty \leq \rho\lambda_1.$$

By the triangle inequality, we obtain:

$$\|\mu_{S^*}^* - \mu^{S^*}\|_\infty \leq \rho\lambda_1 + \|X_{S^*}(\beta^* - \beta^{S^*}) + I_{S^*}^\top\varepsilon\|_\infty \leq \rho\lambda_1 + \|X_{S^*}(\beta^* - \beta^{S^*})\|_\infty + \|I_{S^*}^\top\varepsilon\|_\infty$$

Now, we already said that we have $\|X_{S^*}(\beta^* - \beta^{S^*})\|_\infty \leq \rho\epsilon\lambda_n^{\text{BH}} \leq \rho\epsilon\lambda_1^{\text{BH}}$, and using the standard bound on the norm of a Gaussian noise (see Lemma B.2), we also have, with probability tending to one (precisely with probability $1/n$):

$$\|I_{S^*}^\top\varepsilon\|_\infty \leq \|\varepsilon\|_\infty \leq 2\sigma\sqrt{\log n}.$$

Combining the previous inequalities leads to:

$$\|\mu_{S^*}^* - \mu^{S^*}\|_\infty \leq \rho\lambda_1 + \rho\epsilon\lambda_1^{\text{BH}} + 2\sigma\sqrt{\log n} = \rho(1 + 2\epsilon)\lambda_1^{\text{BH}} + 2\sigma\sqrt{\log n}.$$

A standard bound for the Gaussian quantile function gives $\lambda_1^{\text{BH}} \leq \sigma\sqrt{2\log(2n/q)}$, so with $q \geq 2/n$ (this is quite artificial, q is generally more than 0.01) we obtain:

$$\|(\mu^*)_{S^*} - \mu^{S^*}\|_\infty \leq (1 + \rho(1 + 2\epsilon))2\sigma\sqrt{\log n}.$$

Therefore, because the entries of μ^* are of absolute values greater than the right bound of the above inequality we obtain:

$$S \subset \text{supp}((\mu^*)_{S^*}) \subset \text{supp}(\mu^{S^*}) \subset \text{supp}(\hat{\mu}),$$

and so the Power tends to one.

It remains to show the FDR control, using Equation (D.7). Define the False Discovery Proportion (FDP) as $V/(R \vee 1)$, where R and V are defined in Equation (4.1). Because of the inclusion $S \subset \text{supp}(\hat{\mu})$, the FDP is $(R - s)/R = 1 - s/R$ with probability tending to one. According to Equation (D.7) and the assumption on s^* ,

$$\text{FDP} = 1 - \frac{s}{R} \leq 1 - \frac{s}{s^*} \leq 1 - \frac{1 - q}{1 + c} = \frac{q + c}{1 + c} \leq q + c,$$

with probability tending to one. In expectation, and with n tending to infinity, we obtain:

$$\limsup_{n \rightarrow +\infty} \text{FDR}(\hat{\mu}) \leq q + c,$$

and c being arbitrarily close to zero leads to the conclusion.

APPENDIX E: SUPPLEMENTARY SIMULATIONS

We gather here some extra-simulations in low dimension to complete the ones from Section 5.2 with higher FDR level or/and higher correlation level for the design matrix. As it is our particular interest, we focus on experiments with outliers of weak magnitudes.

Figure 7 below is the same as in Section 5.2 for setting 1, excepted that the correlation of the design matrix is now higher (0.8). Results are similar to those obtained in Section 5.2, that is even with a higher correlation, E-SLOPE is able to make much more discoveries while keeping the FDR under control.

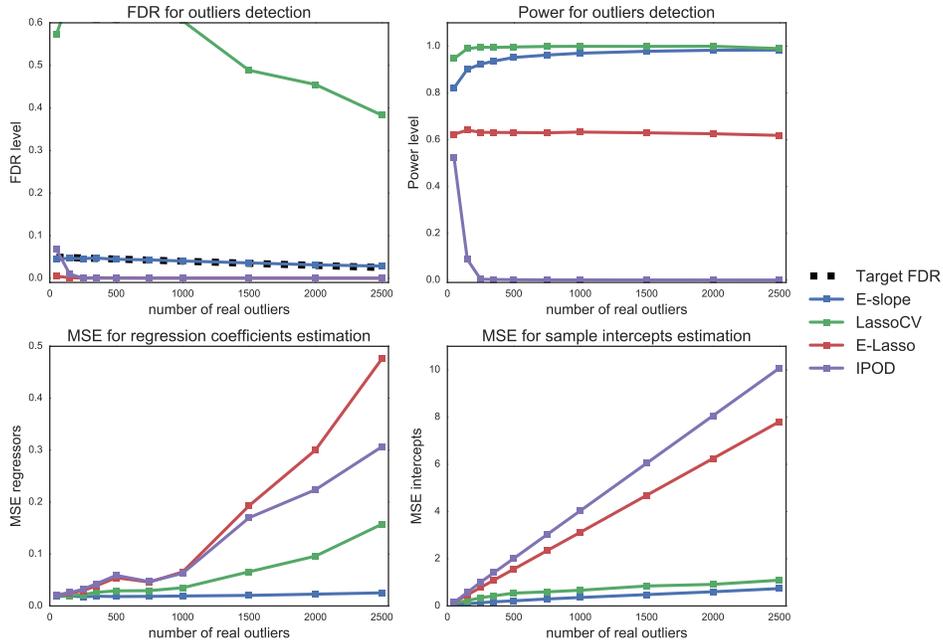


Fig 7: Results for simulation Setting 1 with low-magnitude outliers and correlation $\rho = 0.8$. First row gives the FDR (left) and power (right) of each considered procedure for outliers discoveries. Second row gives the MSE for regressors (left) and intercepts (right). E-SLOPE gives perfect power, is the only one to respect the required FDR, and provides the best MSEs.

Figures 8, 9 below gather the results of simulations in setting 1 of Section 5 with other target FDR (here 10% and 20%) and for both moderate and high correlation (respectively 0.4 and 0.8). E-LASSO and IPOD do not depend on the target FDR so they are not plotted again. The results confirm the

fact that E-SLOPE provides a high power together with a FDR control for a wide range of target FDR level. The left (resp. right) columns contain the results for $\alpha = 10\%$ (resp. $\alpha = 20\%$) as indicated by the straight lines.

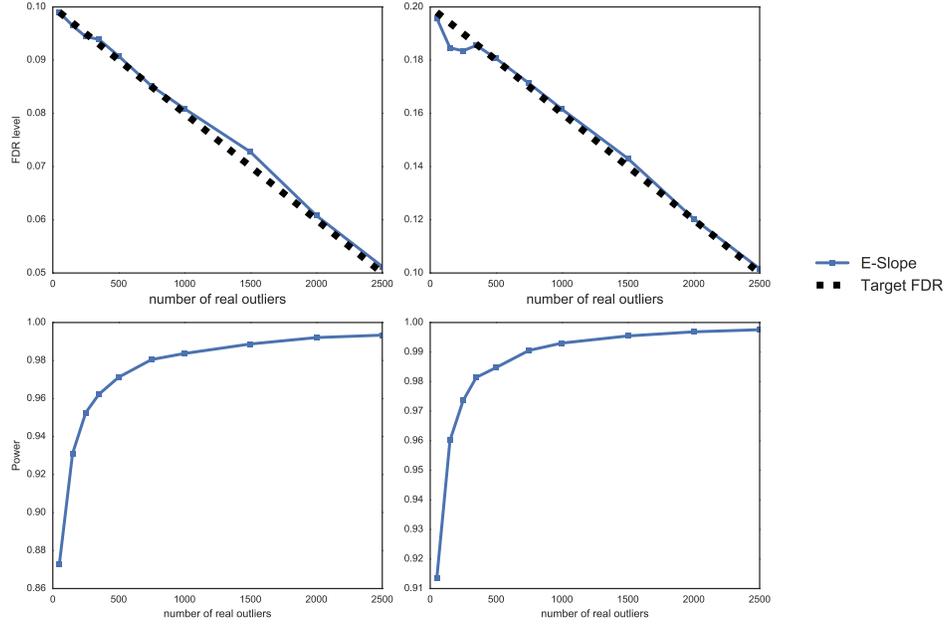


Fig 8: Results for simulation Setting 1 with low-magnitude outliers, correlation $\rho = 0.4$. Left column gives the FDR (top) and power (bottom) for E-SLOPE with target FDR $\alpha = 10\%$. Right column gives the FDR (top) and power (bottom) for E-SLOPE with target FDR $\alpha = 20\%$.

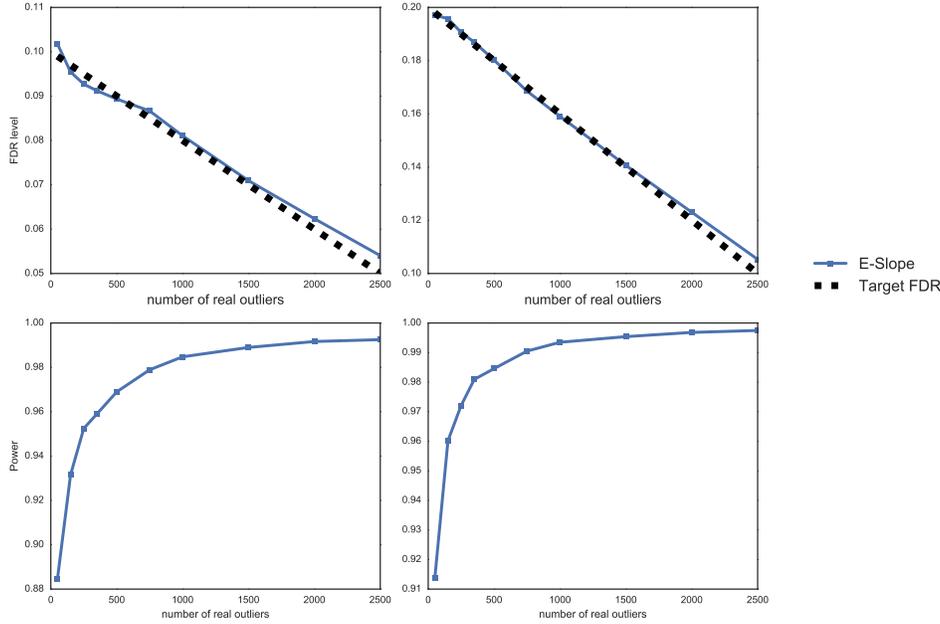


Fig 9: Results for simulation Setting 1 with low-magnitude outliers, correlation $\rho = 0.8$. Left column gives the FDR (top) and power (bottom) for E-SLOPE with target FDR $\alpha = 10\%$. Right column gives the FDR (top) and power (bottom) for E-SLOPE with target FDR $\alpha = 20\%$.

References.

- [1] C. C. Aggarwal and P. S. Yu. Outlier detection for high dimensional data. In *ACM Sigmod Record*, volume 30, pages 37–46. ACM, 2001.
- [2] F. Bach, R. Jenatton, J. Mairal, G. Obozinski, et al. Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning*, 4(1):1–106, 2012.
- [3] F. Bach, R. Jenatton, J. Mairal, G. Obozinski, et al. Structured sparsity through convex optimization. *Statistical Science*, 27(4):450–468, 2012.
- [4] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300, 1995.
- [5] M. Bogdan, E. van den Berg, C. Sabatti, W. Su, and E. J. Candès. Slope - adaptive variable selection via convex optimisation. *Annals of Applied Statistics*, 9(3):1103–1140, 2015.
- [6] M. Bogdan, E. van den Berg, W. Su, and E. J. Candès. Statistical estimation and testing via the ordered ℓ_1 norm. *arXiv:1310.1969*, 2013.
- [7] P. Bühlmann and S. Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- [8] P. C. Bellec, G. Lecué, and A. B. Tsybakov. Slope meets lasso : improved oracle bounds and optimality. *preprint ArXiv*, 2016.
- [9] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal

- reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006.
- [10] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.
- [11] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12(6):805–849, 2012.
- [12] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- [13] S. Chrétien and S. Darses. Sparse recovery with unknown variance: A lasso-type approach. *IEEE Transactions on Information Theory*, 60(7):3970–3988, 2014.
- [14] R. Cook and S. Weisberg. *Residuals and influence in regression*. New York:Chapman and Hall, 1982.
- [15] W. J. Dixon. Analysis of extreme values. *The Annals of Mathematical Statistics*, 21(4):488–506, 1950.
- [16] A. Duval, S. Rolland, A. Compoint, E. Tubacher, B. Iacopetta, G. Thomas, and R. Hamelin. Evolution of instability at coding and non-coding repeat sequences in human msi-h colorectal cancers. *Hum Mol Genet*, 10(5):513–518, 2001.
- [17] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [18] X. Gao and Y. Fang. Penalized weighted least squares for outlier detection and robust regression. ArXiv 2016.
- [19] D. Gervini and V. J. Yohai. A class of robust and fully efficient regression estimators. *Annals of Statistics*, pages 583–616, 2002.
- [20] C. Giraud. *Introduction to High-Dimensional Statistics*. Chapman and Hall/CRC, 2014.
- [21] F. E. Grubbs. Procedures for detecting outlying observations in samples. *Technometrics*, 11(1):1–21, 1969.
- [22] A. Gupta and S. Kohli. An mcdm approach towards handling outliers in web data: a case study using owa operators. *Artificial Intelligence Review*, 46(1):59–82, Jun 2016.
- [23] N. H. Nguyen and T. D. Tran. Robust lasso with missing and grossly corrupted observations. *IEEE transactions on information theory*, 59(4):2036–2058, April 2013.
- [24] A. S. Hadi. A new measure of overall potential influence in linear regression. *Computational Statistics and Data Analysis*, 14:1–27, 1992.
- [25] A. S. Hadi and J. S. Simonoff. Procedures for the identification of multiple outlier in linear models. *Journal of the American Statistical Association*, 88(424):1264–1272, December 1993.
- [26] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. *Robust statistics: the approach based on influence functions*, volume 114. John Wiley & Sons, 2011.
- [27] D. M. Hawkins. *Identification of outliers*, volume 11. Springer, 1980.
- [28] P. J. Huber. The 1972 wald lecture robust statistics: A review. *The Annals of Mathematical Statistics*, pages 1041–1067, 1972.
- [29] P. J. Huber. Wiley series in probability and mathematics statistics. *Robust statistics*, pages 309–312, 1981.
- [30] P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, August 2009.
- [31] E. J. Candès and Y. Plan. Near-ideal model selection by ℓ_1 -minimization. *The Annals of Statistics*, 37(5A):2145–2177, 2009.
- [32] M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery

- using ℓ_1 -constrained quadratic programming (lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, May 2009.
- [33] V. Koltchinskii. Saint flour lectures oracle inequalities in empirical risk minimization and sparse recovery problems. 2009. 2008.
- [34] D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, 47(7):2845–2862, 2001.
- [35] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28(5):1302–1338, 2000.
- [36] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated gaussian design. *Journal of Machine Learning Research*, 11:2241–2259, Aug 2010.
- [37] K. Ro, C. Zou, Z. Wang, G. Yin, et al. Outlier detection for high-dimensional data. *Biometrika*, 102(3):589–599, 2015.
- [38] P. Rousseeuw and V. Yohai. Robust regression by means of s-estimators. In *Robust and nonlinear time series analysis*, pages 256–272. Springer, 1984.
- [39] Y. She and A. B. Owen. Outlier detection using nonconvex penalized regression. *Journal of the American Statistical Association*, 2010.
- [40] A. F. Siegel. Robust regression using repeated medians. *Biometrika*, 69(1):242–244, 1982.
- [41] W. Su, M. Bogdan, and E. J. Candès. False discoveries occur early on the lasso path. *Annals of Statistics*, 45(5):2133–2150, 2017.
- [42] W. Su and E. J. Candès. Slope is adaptive to unknown sparsity and asymptotically minimax. *Annals of Statistics*, 44(3):1038–1068, 2016.
- [43] T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [44] N. Suraweera, B. Iacopetta, A. Duval, A. Compoin, E. Tubacher, and R. Hamelin. Conservation of mononucleotide repeats within 3' and 5' untranslated regions and their instability in msi-h colorectal cancer. *Oncogene*, 20(51):7472–7477, 2001.
- [45] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [46] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58(1):267–288, 1996.
- [47] H. Wang, G. Li, and G. Jiang. Robust regression shrinkage and consistent variable selection through the lad-lasso. *Journal of Business & Economic Statistics*, 25(3):347–355, Jul. 2007.
- [48] T. Wang and Z. Li. Outlier detection in high-dimensional regression model. *Communications in Statistics-Theory and Methods*, 2016. Taylor & Francis.
- [49] S. Weisberg. *Applied linear regression*, volume 528. John Wiley & Sons, 2005.
- [50] C. Yu and W. Yao. Robust linear regression : a review and comparison. *Communications in Statistics - Simulation and Computation*, 2016.
- [51] Y. Zhang, M. J. Wainwright, and M. I. Jordan. Lower bounds on the performance of polynomial-time algorithms for sparse linear regression. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, pages 921–948, 2014.
- [52] H. Zou. The adaptive lasso and its oracles properties. *Journal of the American Statistical Association*, 101(476), 2006.

ALAIN VIROULEAU
CMAP, ECOLE POLYTECHNIQUE
ROUTE DE SACLAY
91128 PALAISEAU CEDEX, FRANCE
E-MAIL: alain.virouleau@polytechnique.edu

STÉPHANE GAÏFFAS
LPMA - UNIV. PARIS DIDEROT
BTIMENT SOPHIE GERMAIN
CASE COURRIER 7012
75205 PARIS CEDEX 13
FRANCE E-MAIL: stephane.gaiffas@polytechnique.edu

AGATHE GUILLOUX
LAMME, UNIVERSITÉ D'EVRY, CNRS
UNIVERSITÉ PARIS-SACLAY
91025, EVRY, FRANCE
E-MAIL: agathe.guilloux@univ-evry.fr

MALGORZATA BOGDAN
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF WROCLAW
WROCLAW, POLAND
E-MAIL: malgorzata.bogdan@uwr.edu.pl