

On the sign recovery by LASSO, thresholded LASSO and thresholded Basis Pursuit Denoising

Patrick J C Tardivel, Malgorzata Bogdan

▶ To cite this version:

Patrick J C Tardivel, Malgorzata Bogdan. On the sign recovery by LASSO, thresholded LASSO and thresholded Basis Pursuit Denoising. Scandinavian Journal of Statistics, inPress. hal-01956603v7

HAL Id: hal-01956603 https://hal.science/hal-01956603v7

Submitted on 31 Aug 2021 $\,$

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the sign recovery by LASSO, thresholded LASSO and thresholded Basis Pursuit Denoising

Patrick J.C. Tardivel^{a*} and Małgorzata Bogdan^{a,b},

^a Institute of Mathematics, University of Wrocław, Wrocław, Poland

^b Department of Statistics, Lund University, Lund, Sweden

Abstract

Basis Pursuit (BP), Basis Pursuit DeNoising (BPDN), and LASSO are popular methods for identifying important predictors in the high-dimensional linear regression model $Y = X\beta + \varepsilon$. By definition, when $\varepsilon = 0$, BP uniquely recovers β when $X\beta = Xb$ and $\beta \neq b$ implies $||b||_1 > ||\beta||_1$ (identifiability condition). Furthermore, LASSO can recover the sign of β only under a much stronger irrepresentability condition. Meanwhile, it is known that the model selection properties of LASSO can be improved by hard-thresholding its estimates. This article supports these findings by proving that thresholded LASSO, thresholded BPDN and thresholded BP recover the sign of β in both the noisy and noiseless cases if and only if β is identifiable and large enough. In particular, if X has iid Gaussian entries and the number of predictors grows linearly with the sample size, then these thresholded estimators can recover the sign of β when the signal sparsity is asymptotically below the Donoho-Tanner transition curve. This is in contrast to the regular LASSO, which asymptotically, recovers the sign of β only when the signal sparsity tends to 0. Numerical experiments show that the identifiability condition, unlike the irrepresentability condition, does not seem to be affected by the structure of the correlations in the X matrix.

Keywords: Multiple regression, Basis Pursuit, LASSO, Sparsity, Active set estimation, Sign estimation, Identifiability condition, Irrepresentability condition

1 Introduction

Let us consider the high dimensional linear model

$$Y = X\beta + \varepsilon,\tag{1}$$

^{*}Corresponding author: tardivel@math.uni.wroc.pl

where $X = (X_1|...|X_p)$ is a $n \times p$ design matrix with $n \leq p$, ε is a random vector in \mathbb{R}^n , and $\beta \in \mathbb{R}^p$ is an unknown vector of regression coefficients. The sign vector of β is $S(\beta) = (S(\beta_1), \ldots S(\beta_p)) \in \{-1, 0, 1\}^p$, where for $x \in \mathbb{R}$, $S(x) = \mathbf{1}_{x>0} - \mathbf{1}_{x<0}$. Our main goal is to recover $S(\beta)$. This goal is somewhat more general than the goal of recovering the active set, $\operatorname{supp}(\beta) := \{i \in \{1, \ldots, p\} \mid \beta_i \neq 0\}$. Indeed, given $S(\beta)$ one may recover $\operatorname{supp}(\beta)$ since $\operatorname{supp}(\beta) := \{i \in \{1, \ldots, p\} \mid S(\beta_i) \neq 0\}$ but conversely, given $\operatorname{supp}(\beta)$, one cannot recover $S(\beta)$. Moreover, some theoretical works recommend sign recovery instead of active set recovery [17, 27].

1.1 BP, BPDN and LASSO

First introduced in signal processing [7], BP estimator is a solution of the following optimization problem

$$\hat{\beta}^{\text{BP}} := \operatorname{argmin} \|b\|_1 \text{ subject to } Y = Xb,$$

which can be solved using linear programming. In the noiseless case, when $\varepsilon = 0$ and $Y = X\beta$, we say that β is identifiable with respect to X and the ℓ_1 norm if $\hat{\beta}^{BP}$ uniquely recovers β *i.e.* if $b \neq \beta$, $Xb = X\beta$ implies $\|b\|_1 > \|\beta\|_1$. Note that a geometric characterization of the identifiability condition, depending on X and the sign of β , is given in [26]. This condition implies that β is sparse. Indeed, Lemma 3 in Tardivel et al. [30] shows that $\operatorname{card}\{i \in \{1, \ldots, p\} \mid \hat{\beta}_i^{BP} \neq 0\} \leq n$ i.e. $\hat{\beta}^{BP}$ has at least p - n zeros. Consequently having $k \leq n$, where k is the number of non-zero elements in β , is necessary for the identifiability condition. On the other hand, some assumptions on the sparsity of β guarantee that β is identifiable with respect to the ℓ_1 norm. For example, the identifiability condition holds if the number of non-zero elements of β satisfies the mutual coherence condition [12, 16, 19]. In the special case where the elements of X are iid normal variables, the transition curve $\rho : (0, 1) \mapsto (0, 1)$ [11] characterizes the identifiability condition with respect to the signal sparsity $\xi = k/n$. In particular, it is known that the *identifiability* condition is satisfied with probability converging to 1 if $\xi < \rho(\delta)$ or converging to 0 if $\xi > \rho(\delta)$.

In the case where ε has a continuous distribution, almost surely BP has *n* non-zero components and thus BP cannot recover the sign of β once $k \neq n$. To account for noise, when ε has iid $\mathcal{N}(0, \sigma^2)$ entries for some $\sigma > 0$, Chen and Donoho [7] extended BP by proposing the following estimation algorithm:

$$\widehat{\beta}^{\mathcal{L}} := \underset{b \in \mathbb{R}^p}{\operatorname{argmin}} \ \frac{1}{2} \|Y - Xb\|_2^2 + \lambda \|b\|_1 \text{ with } \lambda = \sigma \sqrt{2\log p}.$$

$$\tag{2}$$

In the seminal work of Tibshirani [31], the algorithm was extended for fitting the general multiple regression models where different values of λ are often more appropriate. The method gained great popularity in the statistical community under the name LASSO (Least Absolute Shrinkage and Selection Operator), while in the signal processing community it is often called Basis Pursuit Denoising (BPDN). In this paper, we will use the term BPDN for a slightly different form of this estimator, where association with BP is even clearer

$$\widehat{\beta}^{\text{BPDN}} := \underset{b \in \mathbb{R}^p}{\operatorname{argmin}} \|b\|_1 \text{ subject to } \|Y - Xb\|_2^2 \le R, \text{where } R \ge 0.$$
(3)

Note that X is fixed but Y is random, and that for a given $\lambda > 0$ for LASSO we cannot choose a fixed R > 0 for BPDN under which the two estimators are almost surely equal. To illustrate this fact, we note that $\hat{\beta}^{\text{BPDN}} = 0$ if and only if $||Y||_2^2 \leq R$, while for LASSO $\hat{\beta}^{\text{L}} = 0$ if and only if $||X'Y||_{\infty} \leq \lambda$.

1.2 Sign recovery by LASSO

Properties of the LASSO sign estimator $S(\hat{\beta}^{L}(\lambda)) := \left\{ S(\hat{\beta}^{L}_{1}(\lambda)), \ldots, S(\hat{\beta}^{L}_{p}(\lambda)) \right\}$ (or properties of the active set estimator $\operatorname{supp}(\hat{\beta}^{L}(\lambda)) := \{i \in \{1, \ldots, p\} \mid \hat{\beta}^{L}_{i}(\lambda) \neq 0\}$) have been intensively studied [14, 21, 22, 35, 39, 40]. According to Theorem 2 of Wainwright [35], the irrepresentability condition is necessary for LASSO to recover $S(\beta)$ with high probability. Indeed, if $\ker(X_{I}) = \{\mathbf{0}\}$, $\|X'_{\overline{I}}X_{I}(X'_{I}X_{I})^{-1}S(\beta_{I})\|_{\infty} > 1$ and both ε and $-\varepsilon$ have the same distribution, then for any choice of the tuning parameter $\lambda > 0$, $\mathbb{P}(S(\hat{\beta}^{L}(\lambda)) = S(\beta)) \leq 1/2$. This result also holds for the noiseless case, in which LASSO does not recover $S(\beta)$. Moreover, Bühlmann and van de Geer [4] (pp. 192-194) have shown that if $\varepsilon = \mathbf{0}$ and the irrepresentability holds strictly (*i.e. if* $\|X'_{\overline{I}}X_{I}(X'_{I}X_{I})^{-1}S(\beta_{I})\|_{\infty} < 1$) then the nonrandom set $\operatorname{supp}(\hat{\beta}^{L}(\lambda))$ is equal to $\operatorname{supp}(\beta)$ as long as the non-zero components of β are sufficiently large. The proof given in [4] can be easily adapted for sign recovery.

Proposition 1, proved in the Appendix, shows that the identifiability condition is weaker than the irrepresentability condition.

Proposition 1 Let X be an $n \times p$ matrix with $p \ge n$ columns in general position. Moreover, let $\beta \in \mathbb{R}^p$, $I := \operatorname{supp}(\beta)$ and suppose that $\operatorname{ker}(X_I) = \{\mathbf{0}\}$. If the irrepresentability condition holds, then the parameter β is identifiable with respect to the ℓ_1 norm.

For the case where the elements of X are iid random variables from the normal distribution, the irrepresentability condition is satisfied with high probability if and only if $k < \frac{n}{2\log p}$ (see [35]). This implies that LASSO cannot recover $S(\beta)$ when $k/n \to \xi > 0$, even if $\xi < \rho(\delta)$.

1.3 Sign recovery by thresholded LASSO

It is well known that LASSO can consistently estimate β under much weaker assumptions than the irrepresentability condition (see, e.g., [23] or [34]). This suggests that an appropriately thresholded version of LASSO can recover $S(\beta)$ under weaker assumptions than the irrepresentability condition [24]. Concerning sign recovery by thresholded BP, first theoretical results were given by Saligrama and Zhao [25]. More recently, Descloux and Sardy [9] show that the stable null space property is a sufficient condition to recover the sign of β by thresholded BP. To our knowledge, the stable null space property has so far been the weakest sufficient condition known to guarantee the recovery of $S(\beta)$. In Theorem 1, we show that "identifiability", a weaker condition than the stable null space property, is sufficient to recover $S(\beta)$ by thresholded BP, thresholded BPDN or thresholded LASSO when the non-zero elements of β are large enough. We also show that "identifiability" is a necessary condition for the recovery of $S(\beta)$ by these methods, regardless of the signal magnitude.

Theorem 1 implies that in the linear sparsity regime of [11] for Gaussian matrices, thresholded BP, thresholded BPDN and thresholded LASSO recover $S(\beta)$ if, asymptotically, $k < n\rho(n/p)$ holds and the signal magnitude tends towards infinity. Table 1 summarizes known conditions for sign or support recovery by LASSO and related methods, when $\varepsilon \neq 0$ and the signal is large enough.

Method	Conditions	Type	Related articles
Regular	Mutual coherence	sufficient	[6, 21]
LASSO	Irrepresentability	necessary and sufficient	[14, 22, 35, 39, 40]
Thresholded	Stable null space	sufficient	[9]
LASSO, BP	Identifiability	necessary and sufficient	current article

Table 1: Some theoretical conditions under which LASSO and related methods can recover the sign or the support of β when $\varepsilon \neq 0$ and non-zero elements of β are large enough. Additional comments related to these conditions are given in Appendix.

1.4 Graphical illustration of the main result

By definition, the irrepresentability condition depends on β by $S(\beta)$. Moreover, as asserted in Proposition 2, the identifiability condition similarly depends only on $S(\beta)$. Thus, the comparison of these two conditions can be performed by considering parameter vectors such that $\beta = S(\beta)$. In Figure 1, we plot the irrepresentability and identifiability curves representing the proportion of sign vectors for which the identifiability condition or the irrepresentability condition is satisfied among k-sparse sign vectors. For each value of k, this proportion was estimated by generating 1000 sign vectors from a uniform distribution on the set $\{u \in \{-1, 0, 1\}^p \mid \operatorname{card}(\operatorname{supp}(u)) = k\}$.

Figure 1 shows a large potential for improvement in sign detection by thresholded LASSO compared to the base LASSO. If the columns of the design matrix are independent Gaussian vectors, then the irrepresentability condition is satisfied with high probability only if $k \leq 10$, which is also a limiting sparsity for the recovery of $S(\beta)$ by LASSO. Instead, the identifiability condition, which provides the limiting sparsity for the recovery of $S(\beta)$ by the thresholded LASSO, holds as long as $k \leq 30$. The simulations are in strong agreement with the asymptotic results in [35] and [13], which predict the upper bounds for k such that the irrepresentability and identifiability conditions hold for Gaussian design matrices with independent entries.

The difference between identifiability and irrepresentability curves becomes even more apparent when there are strong correlations between different columns in the design matrix. As expected, the irrepresentability curve shrinks towards zero. Instead, and rather unexpectedly, the simulated identifiability curve remains intact.



Figure 1: Simulated identifiability and irrepresentability curves for two specific design matrices of dimensions 100×300 . In both cases the rows of X were generated as iid random vectors from the multivariate normal distribution $N(0, \Sigma)$, with $\Sigma = I$ in the left panel and Σ in the second panel having a compound symmetry form with $\Sigma_{ii} = 1$ and $\Sigma_{ij} = 0.9$ for $i \neq j$. The x-axis represents the sparsity k and the y-axis the fraction of sign vectors that satisfy the identifiability condition (respectively, irrepresentability condition). In the left panel, the vertical lines represent the values $k = \frac{100}{2 \log(300)}$ and $k = 100 \times \rho(100/300)$. These values correspond to the asymptotic upper bounds for k so the irrepresentability and identifiability conditions hold for Gaussian design matrices with independent entries.

1.5 Organization of the article

In Section 2, Theorem 1 shows that identifiability is a necessary and sufficient condition for LASSO to separate the non-zero components of β from the noise and to recover asymptotically $S(\beta)$ with thresholded LASSO and thresholded BPDN. Corollary 1 shows that under the asymptotic linear sparsity regime for Gaussian matrices thresholded LASSO and thresholded BPDN can recover the sign of β when asymptotically $k < n\rho(n/p)$.

In Section 3, Proposition 2 shows that the identifiability condition depends only on $S(\beta)$ and not on the size of the non-zero components of β . Here we also introduce the irrepresentability and identifiability curves, which respectively give the fraction of k-sparse sign vectors satisfying the irrepresentability condition and the identifiability condition. Section 4 is devoted to numerical experiments showing that sign estimators derived from the thresholded LASSO and the thresholded BPDN are better than sign estimators derived from LASSO and the adaptive LASSO, and that the knockoff method allows appropriate threshold selection for both methods. The proofs are in the Appendix. In this section, we also formulate Proposition 3 providing a tight upper bound on the probability of recovering $S(\beta)$ through LASSO.

1.6 Notation and assumptions

In this article we always assume that the design matrix X has columns in general position¹ (see, e.g., [32] or the supplementary material for this manuscript). This assumption guarantees that the minimizer of (2) (resp. minimizer of (3)) is unique and that therefore the LASSO estimator (resp. BPDN estimator) is well-defined.

¹Actually, general position is just a sufficient condition under which, for a fixed $X \in \mathbb{R}^{n \times p}$, independently of $Y \in \mathbb{R}^n$ the LASSO minimizer is unique. Recently, this condition was relaxed by Ewald and Schneider [15] to a geometric criterion which is both sufficient and necessary.

This assumption is very weak and generically holds. Indeed, when X is a random matrix such that the entries $(X_{11}, X_{12}, \ldots, X_{np})$ have a density on \mathbb{R}^{np} then, almost surely, X is in general position [32].

Notation used in the following sections is as follows:

- Let I be a subset of $\{1, \ldots, p\}$. We denote by \overline{I} the complement of I, namely $\overline{I} := \{1, \ldots, p\} \setminus I$.
- The notation X_I stands for a matrix whose columns are indexed by the elements of $I: (X_i)_{i \in I}$.
- For $b \in \mathbb{R}^p$, b_I denotes the subvector containing elements of b with indices in I.
- The symbols $supp(b), supp^+(b)$ and $supp^-(b)$ denote, respectively, the sets $\{i \in \{1, ..., p\} \mid b_i \neq 0\}, \{i \in \{1, ..., p\} \mid b_i > 0\}$ and $\{i \in \{1, ..., p\} \mid b_i < 0\}$.
- LASSO and BPDN estimators depend on X, β, ε and the tuning parameter λ > 0 or the regularization parameter R ≥ 0. When appropriate, we use the parentheses to indicate these dependencies. The estimator β indiscriminately represents the LASSO estimator or the BPDN estimator.

To formulate our asymptotic results, we will often consider a sequence of regression parameters $\beta^{(r)} \in \mathbb{R}^p$, $r \in \mathbb{N}$, for which non-zero components tend to infinity in the following way.

Assumption 1

- 1) The sign of $\beta^{(r)}$ is invariant, namely there exists a sign vector $s^0 \in \{-1, 0, 1\}^p$ such that for any $r \in \mathbb{N}, S(\beta^{(r)}) = s^0$.
- 2) The following limit holds $\lim_{r \to +\infty} \min\{|\beta_i^{(r)}|, i \in \operatorname{supp}(s^0)\} = +\infty$.
- **3)** There exists q > 0 such that

$$\forall r \in \mathbb{N}, \frac{\min\{|\beta_i^{(r)}|, i \in \operatorname{supp}(s^0)\}}{\|\beta^{(r)}\|_{\infty}} \ge q.$$

2 Identifiability is a necessary and sufficient condition for the sign recovery

If β does not satisfy the irrepresentability condition, then the LASSO sign estimator $S(\hat{\beta}^{L}(\lambda))$ is very unlikely to recover $S(\beta)$. However, one may relax this condition. In fact, Theorem 1 shows that an appropriately thresholded version of LASSO (or thresholded version of BPDN) recovers $S(\beta)$ if only the non-zero elements of β are sufficiently large and the identifiability condition is satisfied.

Theorem 1 Let X be a $n \times p$ matrix in general position and such that $\operatorname{rank}(X) = n$ and let $\widehat{\beta}$ be the LASSO or BPDN estimator with any fixed value of the tuning parameter $\lambda > 0$ or with any fixed regularization parameter

 $R \ge 0.$

Necessary condition for sign recovery: If $S(\beta)$ is unidentifiable with respect to the ℓ_1 norm, then the sign estimator derived from the thresholded LASSO or thresholded BPDN cannot recover $S(\beta)$. Indeed, for each fixed $\varepsilon \in \mathbb{R}^n$, the sign of at least one non-zero component of β is not correctly estimated by $\hat{\beta}(\varepsilon)$:

$$\exists i \in \operatorname{supp}(\beta) \text{ such that } \widehat{\beta}_i(\varepsilon)\beta_i \leq 0.$$

Sufficient condition for sign recovery: Let $\beta^{(r)}$ be a sequence in \mathbb{R}^p satisfying Assumption 1. If s^0 is identifiable with respect to the ℓ_1 norm, then for any fixed $\varepsilon \in \mathbb{R}^n$ and sufficiently large $r > r_0(\varepsilon)$ the estimator $\widehat{\beta}(\varepsilon, r)$ separates negative components of $\beta^{(r)}$ (i.e $i \in \text{supp}^{-}(\beta^{(r)})$), zero components of $\beta^{(r)}$ (i.e $i \notin \text{supp}(\beta^{(r)})$) and positive components of $\beta^{(r)}$ (i.e $i \in \text{supp}^{+}(\beta^{(r)})$):

i)

$$\forall i \in \operatorname{supp}(\beta^{(r)}), \ \widehat{\beta}_i(\varepsilon, r)\beta_i^{(r)} > 0$$

ii)

$$\max_{i\in\operatorname{supp}^{-}(\beta^{(r)})}\left\{\widehat{\beta}_{i}(\varepsilon,r)\right\} < \min_{i\notin\operatorname{supp}(\beta^{(r)})}\left\{\widehat{\beta}_{i}(\varepsilon,r)\right\} \leq \max_{i\notin\operatorname{supp}(\beta^{(r)})}\left\{\widehat{\beta}_{i}(\varepsilon,r)\right\} < \min_{i\in\operatorname{supp}^{+}(\beta^{(r)})}\left\{\widehat{\beta}_{i}(\varepsilon,r)\right\}.$$

Let us note that the assumptions about X are very weak and hold in general if $n \leq p$. The assumption that rank(X) = n ensures that the BPDN estimator is well-defined for any $R \geq 0$. The general position condition ensures uniqueness of both the LASSO and the BPDN estimator (see, e.g., Proposition 1 in the Supplementary Material).

Theorem 1 emphasises that one cannot recover $S(\beta)$ with a sign estimator derived from LASSO or BPDN if β is not identifiable with respect to the ℓ_1 norm. If β is identifiable with respect to the ℓ_1 norm, Theorem 1 implies that $S(\beta)$ can be recovered by deriving sign estimators from the thresholded LASSO or thresholded BPDN. In Section 4, we show how the appropriate thresholds can be obtained using control variables constructed by the knockoff method (see, e.g., [1, 5]).

In the asymptotic linear sparsity regime for Gaussian matrices, the transition curve $\rho(\cdot)$ described in [11] allows us to characterize the identifiability condition, and so Theorem 1 yields the Corollary 1.

Corollary 1 Let X be a $n \times p_n$ standard Gaussian matrix and let k_n denote the number of non-zero components of $\beta^{(n)} \in \mathbb{R}^{p_n}$.

Necessary condition for sign recovery: If $n/p_n \to \delta \in (0,1)$, $k_n/n \to \xi \in (0,1)$ and $\xi > \rho(\delta)$ then asymptotically the sign of at least one non-zero component of $\beta^{(n)}$ is incorrectly estimated by $\hat{\beta}$:

$$\lim_{n \to +\infty} \mathbb{P}_{X,\varepsilon} \left(\exists i \in \operatorname{supp}(\beta^{(n)}) \text{ such that } \widehat{\beta}_i \beta_i^{(n)} \leq 0 \right) = 1.$$

Sufficient condition for sign recovery: Given n, let $(\beta^{(n,r)})_{r\in\mathbb{N}}$ be a sequence satisfying the Assumption 1 (where q does not depend on n). If $n/p_n \to \delta \in (0,1)$, $k_n/n \to \xi \in (0,1)$ and $\xi < \rho(\delta)$ then the estimator $\widehat{\beta}$ asymptotically separates the negative components of $\beta^{(n,r)}$, zero components of $\beta^{(n,r)}$ and positive components of $\beta^{(n,r)}$:

i)

$$\lim_{n \to +\infty} \lim_{r \to +\infty} \mathbb{P}_{X,\varepsilon} \left(\forall i \in \operatorname{supp}(\beta^{(n,r)}), \ \widehat{\beta}_i \beta_i^{(n,r)} > 0 \right) = 1$$

ii)

$$\lim_{n \to +\infty} \lim_{r \to +\infty} \mathbb{P}_{X,\varepsilon} \left(\max_{i \in \operatorname{supp}^{-}(\beta^{(n,r)})} \left\{ \widehat{\beta}_i \right\} < \min_{i \notin \operatorname{supp}(\beta^{(n,r)})} \left\{ \widehat{\beta}_i \right\} \le \max_{i \notin \operatorname{supp}(\beta^{(n,r)})} \left\{ \widehat{\beta}_i \right\} < \min_{i \in \operatorname{supp}^{+}(\beta^{(n,r)})} \left\{ \widehat{\beta}_i \right\} \right) = 1.$$

3 Identifiability and irrepresentability curves

Given a certain design matrix X, we now define the irrepresentability indicator of the sign vector s. Sign irrepresentability indicator:

$$\Phi_{\rm IC}^X : s \in \{-1, 0, 1\}^p \mapsto \begin{cases} 1 \text{ if } s = (0, \dots, 0) \\ 1 \text{ if } \ker(X_I) = \{\mathbf{0}\} \text{ and } \|X'_{\overline{I}}X_I(X'_IX_I)^{-1}s_I\|_{\infty} \le 1 \text{ where } I := \operatorname{supp}(s) \\ 0 \text{ otherwise} \end{cases}$$

This irrepresentability indicator shows whether the LASSO sign estimator can recover $S(\beta)$. Namely, if $\phi_{\text{IC}}^X(S(\beta)) = 0$ then $S(\beta)$ cannot be recovered using the LASSO sign estimator, even if the non-zero components of β are extremely large.

Proposition 2 shows that the identifiability condition also depends only on $S(\beta)$ and not on the sizes of the non-zero components of β .

Proposition 2 Consider two vectors $b \in \mathbb{R}^p$ and $\tilde{b} \in \mathbb{R}^p$ such that $S(b) = S(\tilde{b})$. Then \tilde{b} is identifiable with respect to the matrix X and the ℓ_1 norm if and only if b is identifiable with respect to the matrix X and the ℓ_1 norm.

Given a certain design matrix X, the sign identifiability indicator is defined as follows.

Sign identifiability indicator:

$$\Phi_{\text{Idtf}}^X : s \in \{-1, 0, 1\}^p \mapsto \begin{cases} 0 \text{ if } s \neq \underset{b \in \mathbb{R}^p}{\operatorname{argmin}} \|b\|_1 \text{ subject to } Xb = Xs \\ 1 \text{ otherwise} \end{cases}$$

This identifiability indicator shows whether the sign estimators obtained by thresholded LASSO and thresholded BPDN can recover $S(\beta)$. Namely, if $\phi_{\text{Idtf}}^X(S(\beta)) = 0$ then the thresholded LASSO (respectively, thresholded BPDN) sign estimator fails to recover $S(\beta)$ even if the non-zero components of β are extremely large.

According to Proposition 2 in the Supplementary Material, β does not satisfy the identifiability condition if the columns $(X_i)_{i \in \text{supp}(\beta)}$ are not linearly independent. Consequently, if $\operatorname{card}(\operatorname{supp}(\beta)) > n$ then $\phi_{\mathrm{IC}}^X(S(\beta)) = \phi_{\mathrm{Idtf}}^X(S(\beta)) = 0$. Let us give some basic properties and comments about the two indicator functions.

- 1. Both ϕ_{IC}^X and ϕ_{Idtf}^X are even.
- 2. Given Proposition 1, for each $s \in \{-1, 0, 1\}^p$, $\Phi^X_{\text{IC}}(s) \le \Phi^X_{\text{Idtf}}(s)$.
- 3. The computation of Φ_{IC}^X requires only a simple matrix calculus; the computation of Φ_{Idtf}^X requires only solving the basis pursuit problem.

3.1 Graphs of identifiability and irrepresentability curves

The number of sign vectors is very large (3^p) and therefore we cannot explicitly specify Φ_{Idtf}^X and Φ_{IC}^X for each sign vector. Instead, we define the identifiability and irrepresentability curves as the following functions of the sparsity k of the vector β , $k = \text{card}(\text{supp}(\beta)) \in \{1, ..., n\}$:

- identifiability curve is defined as $p_{\text{Idtf}}^X(k) := \mathbb{E}_U(\Phi_{\text{Idtf}}^X(U)),$
- irrepresentability curve is defined as $p_{IC}^X(k) := \mathbb{E}_U(\Phi_{IC}^X(U)),$

where U is uniformly distributed on $\{u \in \{-1,0,1\}^p \mid \operatorname{card}(\operatorname{supp}(u)) = k\}$. Figure 1 in the Introduction provides identifiability curves $k \in \{1,\ldots,60\} \mapsto p_{\operatorname{Idtf}}^X(k)$ and irrepresentability curves $k \in \{1,\ldots,60\} \mapsto p_{\operatorname{IC}}^X(k)$ for two specific 100×300 matrices (one generated with iid $\mathcal{N}(0,1)$ entries and the other generated with positively correlated $\mathcal{N}(0,1)$ entries). In addition, for the case where the design matrix X has positively correlated entries, we also consider a situation where U is uniformly distributed on $\{u \in \{0,1\}^p \mid \operatorname{card}(\operatorname{supp}(u)) = k\}$. Specifically, we consider the following setting:

Positively correlated entries and positive components: X is a fixed design matrix with n = 100, p = 300, whose rows are generated by independent draws from the multivariate normal distribution $\mathcal{N}(\mathbf{0}, \Gamma)$, with $\Gamma_{ii} = 1$ for $i \in \{1, \dots, p\}$ and $\Gamma_{ij} = 0.9$ when $i \neq j$ (the same matrix as the one used for the right panel in Figure 1). The distribution of the sign vectors is uniform on $\{u \in \{0, 1\}^p \mid \operatorname{card}(\operatorname{supp}(u)) = k\}$.

Figure 2 shows an interesting behavior of the irrepresentability and identifiability curves $(k \mapsto p_{IC+}^X(k))$ and $k \mapsto p_{Idtf+}^X(k)$ in the above setting. Here we observe that the irrepresentability condition becomes even more stringent than in the case when the distribution of the elements of the sign vector is symmetric. At the same time, the identifiability condition now becomes much weaker and is satisfied under a much wider range of sparsity levels compared to the identifiability curve in the right panel of Figure 1.

Correlated entries and positive components



Figure 2: Graphs of functions $k \mapsto p_{\text{Idtf}+}^X(k)$ and $k \mapsto p_{\text{IC}+}^X(k)$ where X is described in the setting above. The vertical lines represent the values $k = \frac{100}{2\log(300)}$ and $k = 100 \times \rho(100/300)$.

4 Numerical comparisons of sign estimators

Theorem 1 states that the sign estimators provided by thresholded LASSO or thresholded BPDN can recover $S(\beta)$ as long as the identifiability condition is satisfied. Another way to recover $S(\beta)$ is to use a sign estimator provided by the adaptive LASSO proposed in [40]. Indeed, as claimed in [40] or [20], if the weights for adaptive LASSO are based on an accurate estimator for β one obtains a sign estimator that is consistent for $S(\beta)$ under much weaker assumptions than the irrepresentability condition. In this section, we numerically compare the sign estimators obtained by LASSO, thresholded LASSO, thresholded BP and adaptive LASSO. We also include a comparison with LASSO-zero [9], which is a modification of the thresholded BP aimed at controlling the number of false discoveries.

4.1 Selection of the tuning parameter

As explained in [38, 36], a value of the optimal tuning parameter for sign recovery by thresholded LASSO is much smaller than the optimal value of the tuning parameter for LASSO sign estimator. Specifically:

- For the LASSO sign estimator, the tuning parameter must be large enough to prevent the inclusion of false detections.
- For the thresholded LASSO sign estimator, the tuning parameter should be chosen to achieve good separation between zero and non-zero elements of β. This requires good estimation rather than good selection properties. Here, the tuning parameter does not need to be large, as the threshold allows false detections

to be eliminated.

4.1.1 The tuning parameter for the LASSO sign estimator

If $S(\beta)$ satisfies the irrepresentability condition, then Proposition 3 in the Appendix implies that one can choose a tuning parameter λ_L such that, for sufficiently large β , $\mathbb{P}(S(\beta(\lambda_L)) = S(\beta))$ is arbitrarily close to a given value from (0,1) (e.g. 0.95). According to the irrepresentability curve in Figure 1, the irrepresentability condition is satisfied with probability close to 1 if the elements of X are iid $\mathcal{N}(0,1)$ random variables and β contains k = 5 non-zero elements. Therefore, in this setting, we choose λ_L such that the probability of sign recovery converges to 0.95 and the probability of at least one false discovery (Family Wise Error Rate, FWER) converges to 0.05 when the signal magnitude tends to infinity. Since the irrepresentability condition is not satisfied for the remaining scenarios in our simulation study and thus the FWER cannot be controlled at a low level, the performances of LASSO are not reported.

4.1.2 The tuning parameter for the thresholded LASSO sign estimator

The tuning parameter can be chosen using the asymptotic theory of the Approximate Message Passing (AMP) algorithm for LASSO provided, for example, in [2, 28, 38]. In this theory, the elements of the design matrix are iid Gaussian $\mathcal{N}(0,1)$ variables and the components of β are iid random variables from the mixture distribution $\Pi = (1 - \gamma)\delta_0 + \gamma \Pi^*$, where δ_0 is a point mass distribution concentrated at 0 and Π^* is an arbitrary fixed distribution. The asymptotic mean squared error of LASSO is derived under the assumption that the number of observations n and the number of explanatory variables p tend to infinity and that n/p tends to $\delta > 0$. Then the tuning parameter λ_{AMP} is chosen such that the asymptotic mean squared error is minimal (see, e.g., the prescription in [38, 36]). As discussed in [38, 36], such a tuning parameter for any fixed value of Type I error allows maximizing the asymptotic power of the LASSO. In our simulation study, we used this asymptotically optimal $\lambda_{AMP}(k,t)$ for the parameter values: $\delta = n/p = 100/300$, $\gamma = k/p = k/300$ and $\Pi^* = 1/2\delta_t + 1/2\delta_{-t}$, where δ_t is a point mass distribution at t. The values $\lambda_{AMP}(k,t)$ were then used for the simulations in the independent case. In the "correlated" case, these values proved to be sub-optimal, so in this case we additionally obtained results for tuning parameters equal to $0.5\lambda_{AMP}(k,t)$, which provide much better empirical performance.

4.2 Selection of the threshold

We define the thresholded LASSO estimator (or the thresholded BP estimator) to be

$$\forall i \in \{1, \dots, p\}, \widehat{\beta}_i^{\tau} := \widehat{\beta}_i \mathbf{1}_{\{|\widehat{\beta}_i| > \tau\}} .$$

$$\tag{4}$$

Given a threshold $\tau > 0$, we now define FWER as

$$\mathrm{FWER}(\tau) := \mathbb{P}\left(\exists i \notin \mathrm{supp}(\beta), \left|\widehat{\beta}_i^{\tau}\right| \neq 0\right).$$

By taking τ_{α} as the $1 - \alpha$ quantile of the distribution of $\max\left\{|\widehat{\beta}_i|, i \notin \operatorname{supp}(\beta)\right\}$ we would control FWER exactly at the α level. However, τ_{α} cannot be readily determined since β is not known.

To obtain a threshold greater than τ_{α} (and thus control for FWER at the α level), it seems appealing to consider the distribution of the supremum norm of the LASSO estimator (or BP estimator) in the complete null model when $\beta = \mathbf{0}$ [18]. For the BP estimator, Descloux and Sardy [9] propose the threshold $\tau_{\alpha}^{\text{fn}}$ defined as the $1 - \alpha$ quantile of max $\left\{ \left| \hat{\beta}_{1}^{\text{fn}} \right|, \ldots, \left| \hat{\beta}_{p}^{\text{fn}} \right| \right\}$ where $\hat{\beta}^{\text{fn}}$ is the following estimator

 $\widehat{\beta}^{\text{fn}} := \operatorname{argmin} \|\beta\|_1 \text{ subject to } X\beta = \varepsilon, \text{ where } \varepsilon \sim \mathcal{N}_n(0, \sigma^2 I).$

Unfortunately, if the vector β contains some non-zero elements, this intuitive method yields a threshold $\tau_{\alpha}^{\text{fn}}$ smaller than τ_{α} and thus FWER($\tau_{\alpha}^{\text{fn}}$) > α (see also Su et al. [28] for further explanation).

The recently developed knockoff methodology [1, 5] provides control of False Discovery Rate (FDR). This control is achieved by adding additional control variables to the design matrix. Originally designed to control the FDR, the control variables also allow us to approximate the distribution of estimators corresponding to the zero components of β . In this numerical study, we informally use the model-free knockoffs proposed in [5] to approximate a threshold that controls the FWER at a certain level. The approach developed below is suitable for the situation when X is a random matrix whose distribution is invariant to the permutation of the columns. In this setting, we can generate the knockoff variables one by one instead of generating the full $n \times p$ knockoff matrix (see Weinstein et al. [37] for a similar approach). Since adding the controlled variables may change some relevant properties (e.g. the identifiability condition for β), we should ideally add only one knockoff variable at a time when computing LASSO estimates. However, this would lead to a heavy computational burden on the procedure for estimating the relevant threshold. Therefore, in our simulation study, we use model-free knockoffs [5, 37] to generate 30 = p/10 controlled variables. Then Lasso or BP is run on the matrix augmented with these additional columns and the maximum of the absolute values of the regression coefficients over 30 controlled variables is stored. This step is repeated 10 times and the total maximum of p = 300 absolute values of regression coefficients over controlled variables is calculated. The whole procedure is repeated many times (1000 in this case) and the 0.95-quantile of the obtained maxima is used as a threshold to identify zero components of β . To be consistent with the setup of the simulations used to derive the irrepresentability and identifiability curves, we used the same two fixed design matrices X as in Figure 1, while the positions of the k sparse signals and the error terms were randomly generated for each of the 1000 replicates.

4.2.1 LASSO and Adaptive LASSO

In our numerical experiments, we chose the following values of tuning parameters for LASSO and adaptive LASSO:

- For LASSO, we chose $\lambda_L = 81.18$ to control for FWER at the 0.05 level when k = 5 and the covariates are independent.
- For adaptive LASSO, the weights are derived using initial estimates $\hat{\beta}^L(\lambda_{AMP})$, where the tuning parameter is chosen according to the AMP theory, described above. For $i \in \{1, \ldots, p\}$, the weights w_i are defined as $w_i := 1/(|\hat{\beta}_i^L(\lambda_{AMP})| + 10^{-7})$. With these weights and the tuning parameter λ_L described above, the adaptive LASSO has the following expression

$$\widehat{\beta}^{\text{adapt}} := \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \ \frac{1}{2} \|Y - X\beta\|_2^2 + \lambda_L \sum_{i=1}^p w_i |\beta_i|.$$
(5)

In all our simulations, LASSO is computed using glmnet R package.

4.2.2 LASSO-zero from [9]

Given iid $n \times q$ standard Gaussian matrices $G^{(1)}, \ldots, G^{(M)}$, we consider the following basis pursuit minimizers:

$$\forall k \in \{1, \dots, M\}, (\hat{\beta}^{(k)}, \hat{\gamma}^{(k)}) := \underset{b \in \mathbb{R}^p, c \in \mathbb{R}^q}{\operatorname{argmin}} \|b\|_1 + \|c\|_1 \text{ subject to } y = Xb + G^{(k)}c.$$

Now, given a threshold $\tau > 0$, we can define the LASSO-zero estimator $\hat{\beta}_{\tau}^{\text{lass}0(q,M)}$ as follows:

$$\forall i \in \{1, \dots, p\}, \hat{\beta}_i^{med} = \text{median}\{\hat{\beta}_i^{(1)}, \dots, \hat{\beta}_i^{(M)}\} \text{ and } \hat{\beta}_{\tau,i}^{\text{lass0}(q,M)} = \hat{\beta}_i^{med} \mathbf{1}(|\hat{\beta}_i^{med}| > \tau).$$

One can use the knowledge that $\sigma = 1$, to compute the QUT threshold τ as the $1 - \alpha$ quantile of $\|\hat{\beta}^{med}\|_{\infty}$ when $\beta = \mathbf{0}$. Otherwise, one can also pivotize the $\|\hat{\beta}^{med}\|_{\infty}$ statistic to compute this threshold (see [9] for details).

4.3 Numerical comparisons

All numerical experiments are performed with the two 100×300 design matrices X already used in Figure 1 and in Section 3.1. We set $\beta \in \mathbb{R}^p$ such that $k := \operatorname{card}(\operatorname{supp}(\beta))$ with $k \in \{5, 20\}$ and the k elements of $\operatorname{supp}(\beta)$ are drawn without replacement from $\{1, \ldots, p\}$. The non-zero components of β are sampled from the distribution $P(\beta_i = t) = P(\beta_i = -t) = 0.5$ with t > 0. Additionally, when the columns of X are correlated, we consider the constellation in which all non-zero coefficients are equal to t. In all simulations, the error term is generated as $\varepsilon \sim \mathcal{N}(0, Id_n)$.

Figures 3-4 show the comparison between the following sign estimators.

- The sign estimator **L** from LASSO with $\lambda = \lambda_L$.
- The sign estimator **aL** from the adaptive LASSO, described in (5).
- The sign estimator **BPk** from the thresholded BP, where the threshold is given by the "knockoff" method described above.
- The sign estimator **Lk** from the thresholded LASSO with $\lambda = \lambda_{AMP}$ and with a threshold given by the "knockoff" method described above.
- The sign estimator Lks from the thresholded LASSO with $\lambda = 0.5 \lambda_{AMP}$ and with a threshold given by the "knockoff" method described above.
- The sign estimator Lz of LASSO-zero, where M = 30, q = n. The QUT threshold is not data driven and is computed as the 0.95 quantile of ||β^{med}||_∞ when β = 0.
- The sign estimator Lzp of LASSO-zero when M = 30, q = n. The QUT threshold is data-driven and is computed by pivotizing the statistic ||β^{med}||∞, as explained in [9].

We give the curves representing the following statistical properties as a function of t > 0:

- Probability is the proportion of 1000 replicates for which the sign is recovered.
- **FWER** is the proportion of 1000 replicates for which at least one zero component of β is not estimated to be zero.

FWER is directly related to the probability of sign recovery, namely $FWER \ge \alpha$ implies that the probability of sign recovery is less than $1 - \alpha$. The sign estimators **L**, **BPk**, **Lk** and **Lks** were specifically designed to control for *FWER* at the 0.05 level.

If k = 5 and the elements of X are iid normal variables, then the irrepresentability condition holds and LASSO can recover the true model. Figure 3 shows that in this case the upper bound on the probability of recovering the sign given in Proposition 3 is achieved by LASSO and the FWER is controlled. However, for moderately large signals, the probability of recovering $S(\beta)$ by LASSO is much smaller than by the thresholded versions of BP and Lasso, as well as by Lasso-zero. If k = 20, the irrepresentability condition is not satisfied, but the identifiability condition holds, thus the thresholded LASSO and the thresholded BP identify the sign of β when non-zero elements are large and the threshold is properly calibrated.

Figure 3 shows that thresholds for **BP** or **Lz**, which are not data-driven and are computed in the complete null model (i.e. when $\beta = \mathbf{0}$) do not control the FWER when β has many large components (intuitively, when β is far from **0**). Consequently, **BP** or **Lz** cannot recover $S(\beta)$ with a large probability. Instead, our heuristic application of the knockoff method, as well as LASSO-zero when the threshold is data-driven, allow us to almost perfectly control FWER at 0.05. Consequently, when the non-zero components of β are large enough, the sign



Figure 3: This figure shows the FWER and the probability of recovering $S(\beta)$ for each sign estimator and when X contains iid $\mathcal{N}(0,1)$ entries. The graphs on the left show the probability of recovering $S(\beta)$ (on the y-axis) as a function of t (on the x-axis), where t indicates how large the non-zero components of β are. The graphs on the right show the FWER (on the y-axis) as a function of t (on the x-axis).

estimators derived from Lk, BPk, and Lzp recover $S(\beta)$ with probability close to 0.95. Among these methods, the thresholded LASSO estimator Lk with λ selected by AMP theory systematically has the highest probability of recovering $S(\beta)$ for moderately large signals.



Figure 4: This figure shows the FWER and the probability of recovering $S(\beta)$ for each sign estimator when X has strongly correlated columns. The plots on the left show the probability of recovery of $S(\beta)$ (on the y-axis) as a function of t (on the x-axis), where t measures how large the non-zero components of β are. The graphs on the right show the FWER (on the y-axis) as a function of t.

0.0 0.2

2

t: value of the non-null components

3

4

0.2

0.0

2

t: value of the non-null components

3

Figure 4 shows the results of the simulation study for the design matrix X with very highly correlated columns. In this setting, the probability of recovering $S(\beta)$ by LASSO-zero was very close to zero, and for clarity of presentation, LASSO-zero is not included in Figure 4.

Figure 4 confirms the numerical experiments given in Figure 3. In particular, thresholded BP and thresholded LASSO can recover $S(\beta)$ when the threshold is well calibrated and the non-zero components of β are large enough. If the tuning parameter is chosen correctly, thresholded LASSO recovers $S(\beta)$ with a higher probability than thresholded BP. Similar to the case of independent entries, our heuristic knockoff approach allows to control FWER at the 0.05 level.

4.4 Numerical experiments with the design matrix from the riboflavin dataset

In this section, we consider the same simulation setting as in [9], which deals with the riboflavin dataset [3]. The design matrix X contains expression levels of p = 4088 genes for n = 71 Bacillus subtilis bacteria. The matrix is mean-centered (for each $j \in \{1, ..., p\}, \sum_{i=1}^{n} X_{ij} = 0$) and standardized (for each $j \in \{1, ..., p\}, \frac{1}{n-1} \sum_{i=1}^{n} X_{ij}^2 = 1$). Figure 5 shows irrepresentability and identifiability curves of this design matrix.



Identifiability and irrepresentability curves, riboflavin dataset

Figure 5: Graphs of the functions $k \mapsto p_{\text{Idtf}}^X(k)$ and $k \mapsto p_{\text{IC}}^X(k)$ where $k \in \{1, \ldots, 12\}$ and X is the matrix from the riboflavin dataset. The equation of the vertical line is $k = n\rho(n/p)$.

The identifiability and irrepresentability curves shown in Figure 5 confirm the numerical experiments reported in Descloux and Sardy [9]. In particular, the probability of sign recovery by LASSO reported in panel (c) of Figure 6 in [9] is below the irrepresentability curve from Figure 5 and the probability of sign recovery by LASSO-zero is below the identifiability curve.

We also ran the simulations comparing LASSO-zero and thresholded LASSO on the riboflavin dataset. In this study, we set $k := \operatorname{card}(\operatorname{supp}(\beta)) \in \{1, 2, 3\}$ (so that the identifiability condition holds), and the non-zero coefficients are set to 2 with random signs. In addition to the probability of sign recovery and the FWER, we also estimate the False Discovery Rate (FDR) and the True Positive Rate (TPR).

• FDR is the mean of 1000 replicates of the False Discovery Proportion (FDP), where the FDP is defined as follows:

$$FDP = \frac{\operatorname{card}(\{i \notin \operatorname{supp}(\beta) | \hat{\beta}_i \neq 0\})}{\max\{\operatorname{card}(\operatorname{supp}(\hat{\beta})), 1\}}$$

• TPR is the mean of 1000 replicates of the True Positive Proportion (TPP), where the TPP is defined as

follows:

$$TPP = \frac{\operatorname{card}(\operatorname{supp}(\beta) \cap \operatorname{supp}(\hat{\beta}))}{\operatorname{card}(\operatorname{supp}(\beta))}$$

For the thresholded LASSO, we selected the tuning parameter by cross-validation. To estimate the threshold, we used the heuristic knockoff procedure described above, where for each run of LASSO we added q = p/8 = 511 columns to the design matrix X, drawn from the multivariate normal distribution $N(0, \Gamma)$, where $\Gamma_{ii} = 1$ and $\Gamma_{ij} = \frac{1}{p(p-1)(n-1)} \sum_{i \neq j} X'_i X_j = 0.0292.$



Figure 6: This figure provides the probability of sign recovery, FWER, FDR and TPR of LASSO-zero (**Lzp**) and thresholded LASSO (**Lk**) as a function of the sparsity $k \in \{1, 2, 3\}$.

Figure 6 shows that LASSO-zero and thresholded LASSO have FWER above the assumed value of 0.05. This is not surprising since the gene expressions are not exactly normally distributed. However, the FWER of both methods is still quite low (below 0.3) and the FDR is kept close to 0.1. Since the correlations between columns in the design matrix are not too large, the thresholded LASSO and LASSO-zero show similar performance. However, the probability of sign recovery and TPR is systematically larger for thresholded LASSO, which also has smaller FWER and FDR.

5 Conclusion

The focus of this article is on the theoretical properties of sign estimators derived from LASSO, thresholded LASSO, and thresholded BPDN. When $S(\beta)$ is identifiable with respect to the ℓ_1 norm and when non-zero components of β tends to infinity, we have shown that sign estimators derived from thresholded LASSO and thresholded BPDN recover $S(\beta)$. On the other hand, if $S(\beta)$ is not identifiable with respect to the ℓ_1 norm, sign estimators derived from thresholded LASSO and thresholded BPDN cannot recover $S(\beta)$.

We introduced irrepresentability and identifiability curves that provide information about the probability of sign recovery by LASSO or thresholded LASSO and thresholded BPDN as a function of the number of non-zero elements in the vector of regression coefficients.

The performances of the sign estimators derived from LASSO, thresholded LASSO and thresholded BPDN obviously depend on the tuning parameters and the threshold. We have illustrated that AMP theory and the knockoff method are useful to select these parameters. Our simulations show that thresholded LASSO and thresholded BPDN sign estimators outperform the adaptive LASSO and LASSO sign estimators.

Acknowledgments

We would like to thank Emmanuel J. Candès and Wojciech Rejchel for helpful comments. Małgorzata Bogdan's research was supported by NCN grant 2016/23/B/ST1/00454. We gratefully acknowledge the grant from the Wrocław Center of Networking and Supercomputing (WCSS), where most of the computations were performed.

6 Appendix

6.1 Sign recovery with LASSO sign estimator

The upper bound for sign recovery given in Proposition 3 appears in Lemma 3 of Wainwright [35] as a technical result establishing the irrepresentability condition (Theorem 2 in Wainwright [35]). According to Proposition 3, if β is identifiable with respect to the ℓ_1 -norm, this upper bound is asymptotically reached as soon as $\min\{|\beta_i|, i \in$ $\supp(\beta)\}$ tends to $+\infty$.

Lemma 1 (Lemma 3 from [35]) Let $I := \operatorname{supp}(\beta)$ and let $X_I, X_{\overline{I}}$ be matrices whose columns are $(X_i)_{i \in I}$ and $(X_i)_{i \notin I}$, respectively. Let us assume that $\ker(X_I) = \{\mathbf{0}\}$ and let $\zeta_{X,\lambda,S(\beta)} := X'_{\overline{I}}X_I(X'_IX_I)^{-1}S(\beta_I) + \frac{1}{\lambda}X'_{\overline{I}}(Id - X_I(X'_IX_I)^{-1}X'_I)\varepsilon$.

The following upper bound for the sign recovery holds.

$$\mathbb{P}\left(S(\widehat{\beta}^{\mathrm{L}}(\lambda)) = S(\beta)\right) \leq \mathbb{P}\left(\left\|\zeta_{X,\lambda,S(\beta)}\right\|_{\infty} \leq 1\right).$$

Proposition 3 Assume that the assumptions of Lemma 1 hold. Moreover, assume that the sequence $(\beta^{(r)})$ in \mathbb{R}^p satisfies Assumption 1. If s^0 is identifiable with respect to the ℓ_1 -norm, then the following asymptotic results hold.

Sharpness of the upper bound: Asymptotically, the upper bound is obtained.

$$\begin{split} &\lim_{r \to +\infty} \mathbb{P}\left(S(\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = s^{0}\right) \leq \mathbb{P}\left(\left\|\zeta_{X,\lambda,s^{0}}\right\|_{\infty} \leq 1\right), \\ &\lim_{r \to +\infty} \inf \mathbb{P}\left(S(\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = s^{0}\right) \geq \mathbb{P}\left(\left\|\zeta_{X,\lambda,s^{0}}\right\|_{\infty} < 1\right). \end{split}$$

Asymptotic control of FWER: Let us define $\mathbb{P}(\|\zeta_{X,\lambda,s^0}\|_{\infty} < 1) = \gamma$ and $\mathbb{P}(\|\zeta_{X,\lambda,s^0}\|_{\infty} \le 1) = \bar{\gamma}$. The sign of the non-zero elements of $\beta^{(r)}$ is correctly identified with probability converging to 1 and the FWER is controlled at the $1 - \gamma$ level.

$$\begin{split} &\lim_{r \to +\infty} \mathbb{P}\left(\forall i \in I, S(\widehat{\beta}_i^{\mathrm{L}}(\lambda, r)) = s_i^0\right) = 1, \\ &\lim_{r \to +\infty} \sup \mathbb{P}\left(\exists i \notin I, \widehat{\beta}_i^{\mathrm{L}}(\lambda, r) \neq 0\right) \leq 1 - \gamma, \\ &\lim_{r \to +\infty} \inf \mathbb{P}\left(\exists i \notin I, \widehat{\beta}_i^{\mathrm{L}}(\lambda, r) \neq 0\right) \geq 1 - \bar{\gamma}. \end{split}$$

Remark 1 The results in Proposition 3 are quite straightforward when X is orthogonal (i.e. when X'X = I). Indeed, in this case the upper bound is simply the probability that zero components of β are simultaneously estimated to 0, namely $\mathbb{P}(\forall i \notin \operatorname{supp}(\beta), \hat{\beta}_i^{\mathrm{L}}(\lambda) = 0)$

If ε has a covariance matrix $\sigma^2 Id_n$, one can obtain asymptotic results by fixing β and letting σ tend to 0. However, unlike in our asymptotic setting described in Assumption 1, such asymptotic results are very poor and give no information about the FWER. Indeed, when σ tends to 0, the upper bound tends to 0 or 1, depending on whether the irrepresentable condition for β holds or not.

Let us remind that the FWER is equal to $\mathbb{P}(\exists i \notin \operatorname{supp}(\beta), \widehat{\beta}_i^L \neq 0)$. According to Proposition 3, when non-null components of β are infinitely large, the FWER is controlled at level $1 - \mathbb{P}(\|\zeta_{X,\lambda,\operatorname{sign}(\beta)}\|_{\infty} < 1)$. To our knowledge, it is first theoretical result providing a formula for the FWER. Hereafter, let us provide some comments about the FWER control.

To provide a specific value for λ allowing to control the FWER one needs to know the distribution of ε.
For example, when the distribution of ε is known and β = 0 one controls the FWER at level α by taking λ as the 1 − α quantile of ||X'ε||_∞. Let us point out that when ε₁,..., ε_p are iid and the variance σ² of these components is unknown then σ² can be consistently estimate as explained in [10].

- It is easier to control the FWER when X is a random matrix whose distribution is symmetric and invariant by columns permutation than when X is a fixed design matrix. Indeed, when X is random, the distribution of ζ_{X,λ,S(β)} just depends from the sparsity of β and not on S(β).
- Let k ∈ {1,...,p} and let U_k be the set {u ∈ {−1,0,1}^p | card(supp(u)) = k} and let us assume that β ∈ U_k. When X is a random matrix whose distribution is symmetric and invariant by columns permutation by taking λ_α such that P(||ζ_{X,λ_α,u₀}|| < 1) = 1 − α one controls asymptotically the FWER at level α (where u₀ := (1,...,1,0,...,0) ∈ U_k). Such a tuning parameter λ_α is easy to infer by Monte Carlo simulations. When X is fixed, by taking λ_α as follows

$$\frac{1}{\operatorname{card}(U_k)} \sum_{u \in U_k} \mathbb{P}(\|\zeta_{X,\lambda_\alpha,u}\| < 1) = 1 - \alpha,$$
(6)

then the average value of the FWER with respect to $\beta \in U_k$ is $1 - \alpha$ (where non-null components of β are infinitely large). Again λ_{α} is easy to infer. In the numerical experiment, the tuning parameter was selected by solving equation (6).

6.2 Proof of the Proposition 3

First, let us provide lemmas which are useful to prove both Proposition 3 and Theorem 1. Lemma 3 partially proves Proposition 3. Indeed, according to this Lemma, when $(\beta^{(r)})_{r\in\mathbb{N}}$ is a sequence of \mathbb{R}^p satisfying assumptions 1 then the following asymptotic result holds

$$\lim_{r \to +\infty} \mathbb{P}\left(\forall i \in \operatorname{supp}(s^0), S(\widehat{\beta}_i^{\mathrm{L}}(\lambda, r)) = s_i^0 \right) = 1.$$

Lemma 2 Let $(\beta^{(r)})_{r \in \mathbb{N}}$ be a sequence of \mathbb{R}^p satisfying the conditions 1) and 2) of Assumption 1, let us assume that s^0 is identifiable with respect to the ℓ_1 norm and let us set $u_r = \|\beta^{(r)}\|_1$ then

$$\lim_{r \to +\infty} \frac{\widehat{\beta}^{\mathrm{L}}(\varepsilon, r) - \beta^{(r)}}{u_r} = 0$$

Proof: Because $\widehat{\beta}^{L}(\varepsilon, r)$ is the LASSO estimator as defined in (2) then the following inequality occurs

r

$$\frac{1}{2} \|Y - X\widehat{\beta}^{L}(\varepsilon, r)\|_{2}^{2} + \lambda \|\widehat{\beta}^{L}(\varepsilon, r)\|_{1} \leq \frac{1}{2} \|Y - X\beta^{(r)}\|_{2}^{2} + \lambda \|\beta^{(r)}\|_{1}.$$

Since $Y - X\beta^{(r)} = \varepsilon$ one may deduce the following inequalities

$$\lambda \|\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)\|_{1} \leq \frac{1}{2} \|\varepsilon\|_{2}^{2} + \lambda \|\beta^{(r)}\|_{1},$$

$$\Rightarrow \qquad \|\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)/u_{r}\|_{1} \leq \frac{\|\varepsilon\|_{2}^{2}}{2\lambda u_{r}} + 1.$$
(7)

In addition, Cauchy-Schwarz inequality gives the following implications

$$\frac{1}{2} \|\varepsilon + X\beta^{(r)} - X\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)\|_{2}^{2} + \lambda \|\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)\|_{1} \leq \frac{1}{2} \|\varepsilon\|_{2}^{2} + \lambda \|\beta^{(r)}\|_{1},$$

$$\Rightarrow \quad -\|\varepsilon\|_{2} \|X\beta^{(r)} - X\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)\|_{2} + \frac{1}{2} \|X\beta^{(r)} - X\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)\|_{2}^{2} + \lambda \|\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)\|_{1} \leq \lambda \|\beta^{(r)}\|_{1},$$

$$\Rightarrow \quad -\frac{\|\varepsilon\|_{2}}{u_{r}} \left\|X\left(\frac{\widehat{\beta}^{\mathrm{L}}(\varepsilon, r) - \beta^{(r)}}{u_{r}}\right)\right\|_{2} + \frac{1}{2} \left\|X\left(\frac{\widehat{\beta}^{\mathrm{L}}(\varepsilon, r) - \beta^{(r)}}{u_{r}}\right)\right\|_{2}^{2} + \frac{\lambda}{u_{r}} \left\|\frac{\widehat{\beta}^{\mathrm{L}}(\varepsilon, r)}{u_{r}}\right\|_{1} \leq \frac{\lambda}{u_{r}}.$$
(8)

Because u_r tends to $+\infty$ then, according to (7), the sequence $((\widehat{\beta}^{L}(\varepsilon, r) - \beta^{(r)})/u_r)_{r \in \mathbb{N}^*}$ is bounded since the following superior limit is finite

$$\limsup_{r \to +\infty} \left\| \frac{\widehat{\beta}^{\mathrm{L}}(\varepsilon, r) - \beta^{(r)}}{u_r} \right\|_1 \le 2.$$

Consequently, to prove that $\lim_{r\to+\infty} (\widehat{\beta}^{L}(\varepsilon,r) - \beta^{(r)})/u_r = \mathbf{0}$ it is sufficient to show that $\mathbf{0}$ is the unique limit point of this sequence. Let $((\widehat{\beta}^{L}(\varepsilon,\phi(r)) - \beta^{(\phi(r))})/u_{\phi(r)})_{r\in\mathbb{N}^*}$ be a converging subsequence to l (with $\phi: \mathbb{N}^* \to \mathbb{N}^*$ strictly increasing) and without loss of generality, let us assume $\lim_{r\to+\infty} \widehat{\beta}^{L}(\varepsilon,\phi(r))/u_{\phi(r)} = v$ and $\lim_{r\to+\infty} \beta^{(\phi(r))}/u_{\phi(r)} = v'$ so that l = v - v'. By (7) and (8) one may deduce that

$$Xv = Xv' \text{ and } ||v||_1 \le 1.$$

Since, whatever $r \ge 0$, we have $S(\beta^{(\phi(r))}/u_{\phi(r)}) = s^0$ where s^0 is identifiable with respect to the ℓ_1 norm then, according to Proposition 2, one may deduce that $\beta^{(\phi(r))}/u_{\phi(r)}$ is an unitary vector satisfying the identifiability condition. Consequently, $\|v'\|_1 = 1$ and v' is identifiable with respect to the ℓ_1 norm. Consequently, v = v' and thus $l = \mathbf{0}$ is the unique limit point, which implies that

$$\lim_{r \to +\infty} \frac{\widehat{\beta}^{\mathcal{L}}(\varepsilon, r) - \beta^{(r)}}{u_r} = \mathbf{0}.$$

For the proof of Lemma 2, we have not used the third condition of Assumption 1. This condition, under which the smallest non-null component of $\beta^{(r)}$ is not asymptotically infinitely smaller than $\|\beta^{(r)}\|_{\infty}$, is useful to prove Lemma 3.

Lemma 3 Let $(\beta^{(r)})_{r\in\mathbb{N}}$ be a sequence of \mathbb{R}^p satisfying Assumption 1, then

$$\lim_{r \to +\infty} \mathbb{P}(\forall i \in \operatorname{supp}(s^0), S(\widehat{\beta}_i^{\mathrm{L}}(\lambda, r)) = s_i^0) = 1.$$

Proof: Let ε be a fixed vector in \mathbb{R}^p . According to the third condition of Assumption 1 we have $\min\{|\beta_i^{(r)}|, i \in \sup(s^0)\}/\|\beta^{(r)}\|_{\infty} \ge q > 0$, consequently the following inequalities occur

$$\forall i \in \operatorname{supp}(s^0), s_i^0 \frac{\widehat{\beta}_i^{\mathrm{L}}(\varepsilon, \lambda, r) - \beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}} = \frac{s_i^0 \widehat{\beta}_i^{\mathrm{L}}(\varepsilon, \lambda, r)}{\|\beta^{(r)}\|_{\infty}} - \frac{|\beta_i^{(r)}|}{\|\beta^{(r)}\|_{\infty}} \le \frac{s_i^0 \widehat{\beta}_i^{\mathrm{L}}(\varepsilon, \lambda, r)}{\|\beta^{(r)}\|_{\infty}} - q.$$

According to Lemma 2, the following inequality occurs

$$0 = \liminf_{r \to +\infty} s_i^0 \frac{\widehat{\beta}_i^{\mathcal{L}}(\varepsilon, \lambda, r) - \beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}} \leq \liminf_{r \to +\infty} \frac{s_i^0 \widehat{\beta}_i^{\mathcal{L}}(\varepsilon, \lambda, r)}{\|\beta^{(r)}\|_{\infty}} - q$$

Which implies that for r large enough $s_i^0 \widehat{\beta}_i^{\mathrm{L}}(\varepsilon, \lambda, r) > 0$ and thus $S(\widehat{\beta}_i^{\mathrm{L}}(\varepsilon, \lambda, r)) = s_i^0$. When ε is no longer fixed then, for $i \in \mathrm{supp}(s^0)$, almost surely $S(\widehat{\beta}_i^{\mathrm{L}}(r))$ converges to s_i^0 and consequently

$$\lim_{r \to +\infty} \mathbb{P}\left(\forall i \in \operatorname{supp}(s^0), S(\widehat{\beta}_i^{\mathrm{L}}(\lambda, r)) = s_i^0 \right) = 1.$$

- 1		
- 1		1

Proof of Proposition 3:

Let us remind that the vector $\hat{\beta}^{L}(\lambda)$ is the LASSO estimator if and only if the following two inequalities occur simultaneously.

$$X'_{A}(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda)) = \lambda S(\widehat{\beta}^{\mathrm{L}}_{A}(\lambda)), \text{ where } A = \mathrm{supp}(\widehat{\beta}^{\mathrm{L}}(\lambda)),$$
(9)

$$\|X'_{\overline{A}}(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda))\|_{\infty} \leq \lambda.$$
⁽¹⁰⁾

Sharpness of the upper bound) Since the upper bound depends only on s^0 and not on how large the non-null components $\beta^{(r)}$ are then

$$\limsup_{r \to +\infty} \mathbb{P}\left(S(\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = s^{0}\right) \leq \mathbb{P}\left(\left\|\zeta_{X,\lambda,s^{0}}\right\|_{\infty} \leq 1\right).$$

Finally, it must be proven that $\liminf_{r\to+\infty} \mathbb{P}\left(S(\widehat{\beta}^{L}(\lambda,r))=s^{0}\right) \geq \mathbb{P}\left(\left\|\zeta_{X,\lambda,s^{0}}\right\|_{\infty}<1\right)$. Let us remind that $I = \operatorname{supp}(s^{0})$ and let us assume that the following events hold simultaneously

$$X_{I}'(Y - X\widehat{\beta}^{\mathbf{L}}(\lambda)) = \lambda s_{I}^{0} \text{ and } \underbrace{\left\| X_{\overline{I}}'X_{I}(X_{I}'X_{I})^{-1}\lambda s_{I}^{0} + X_{\overline{I}}'\left(Id - X_{I}(X_{I}'X_{I})^{-1}X_{I}'\right)\varepsilon\right\|_{\infty} < \lambda}_{=\left\|\zeta_{X,\lambda,s^{0}}\right\|_{\infty} < 1}.$$
 (11)

We aim to show that the inequalities given above imply that $\widehat{\beta}_{\overline{I}}^{L}(\lambda) = \mathbf{0}$. For convenience, let us set H be the projection matrix $H := X_{I}(X'_{I}X_{I})^{-1}X'_{I}$. When (11) occurs then the following inequalities holds

$$\begin{aligned} \left\| X_{\overline{I}}'H(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda)) + X_{\overline{I}}'(Id - H)\varepsilon \right\|_{\infty} &< \lambda, \\ \left\| X_{\overline{I}}'\left(H(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda)) + (Id - H)\varepsilon \right) \right\|_{\infty} &< \lambda, \\ \left\| X_{\overline{I}}'\left(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda) + X_{\overline{I}}\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda) - HX_{\overline{I}}\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda) \right) \right\|_{\infty} &< \lambda. \end{aligned}$$
(12)

Inequality (12) comes from the following two identities

$$HY = H(X\beta^{(r)}) + H\varepsilon = H(X_I\beta_I^{(r)}) + H\varepsilon = X_I\beta_I^{(r)} + H\varepsilon = X(\beta^{(r)}) + H\varepsilon \text{ and},$$

$$HX\widehat{\beta}^{L}(\lambda) = HX_I\widehat{\beta}^{L}_{I}(\lambda) + HX_{\overline{I}}\widehat{\beta}^{L}_{\overline{I}}(\lambda) = X_I\widehat{\beta}^{L}_{I}(\lambda) + HX_{\overline{I}}\widehat{\beta}^{L}_{\overline{I}}(\lambda) = X\widehat{\beta}^{L}(\lambda) - X_{\overline{I}}\widehat{\beta}^{L}_{\overline{I}}(\lambda) + HX_{\overline{I}}\widehat{\beta}^{L}_{\overline{I}}(\lambda).$$

Let v be the vector $v := X'_{\overline{I}} \left(Y - X \widehat{\beta}^{\mathrm{L}}(\lambda) + X_{\overline{I}} \widehat{\beta}^{\mathrm{L}}_{\overline{I}}(\lambda) - H X_{\overline{I}} \widehat{\beta}^{\mathrm{L}}_{\overline{I}}(\lambda) \right)$. We are going to see that inequality (12) implies that $\widehat{\beta}^{\mathrm{L}}_{\overline{I}}(\lambda) = \mathbf{0}$. Let us assume that $\widehat{\beta}^{\mathrm{L}}_{\overline{I}}(\lambda) \neq \mathbf{0}$ then, on the one hand, the following inequality occurs

$$\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)' v \le \|\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)\|_{1} \|v\|_{\infty} < \lambda \|\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)\|_{1}.$$
(13)

According to (9) the identity $\hat{\beta}_i^{\mathrm{L}}(\lambda)X'_i(Y-X\hat{\beta}^{\mathrm{L}}(\lambda)) = \lambda |\hat{\beta}_i^{\mathrm{L}}(\lambda)|$ occurs. Consequently, on the other hand, the following inequalities hold

$$\begin{aligned} \widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)'v &= \widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)'X'_{\overline{I}}\left(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda) + X_{\overline{I}}\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda) - HX_{\overline{I}}\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)\right), \\ &= \lambda \|\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)\|_{1} + \widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda)'X'_{\overline{I}}(Id - H)X_{\overline{I}}\widehat{\beta}_{\overline{I}}^{\mathrm{L}}(\lambda), \\ &\geq \lambda \|\widehat{\beta}_{\overline{I}}^{L}(\lambda)\|_{1}. \end{aligned}$$
(14)

The last inequality occurs because the projection matrix Id - H is positive semi-definite. Inequalities (13) and (14) provide a contradiction which implies that $\hat{\beta}_{\overline{I}}^{L}(\lambda) = \mathbf{0}$. According to (9), the following implication holds

$$S(\widehat{\beta}_{I}^{\mathrm{L}}(\lambda, r)) = s_{I}^{0} \Rightarrow X_{I}'(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = \lambda s_{I}^{0}.$$

Because s^0 is identifiable with respect to the ℓ_1 norm then, according to Lemma 3, the following convergence in probability occurs

$$\lim_{r \to +\infty} \mathbb{P}(S(\widehat{\beta}_I^{\mathrm{L}}(\lambda, r)) = s_I^0) = \lim_{r \to +\infty} \mathbb{P}(X_I'(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = \lambda s_I^0) = 1.$$
(15)

Using this asymptotic result and since when (11) occurs then $\widehat{\beta}_{\overline{I}}^{L}(\lambda,r) = \mathbf{0}$, one may deduce the following

inequalities

$$\begin{split} \liminf_{r \to +\infty} \mathbb{P} \left(S(\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = s^{0} \right) &= \liminf_{r \to +\infty} \mathbb{P} \left(S(\widehat{\beta}^{\mathrm{L}}_{I}(\lambda, r)) = s^{0}_{I} \text{ and } \widehat{\beta}^{\mathrm{L}}_{\overline{I}}(\lambda, r) = \mathbf{0} \right), \\ &= \liminf_{r \to +\infty} \mathbb{P} (\widehat{\beta}^{\mathrm{L}}_{\overline{I}}(\lambda, r) = \mathbf{0}), \\ &\geq \liminf_{r \to +\infty} \mathbb{P} \left(X'_{I}(Y - X\widehat{\beta}^{\mathrm{L}}(\lambda, r)) = \lambda s^{0}_{I} \text{ and } \|\zeta_{X,\lambda,s^{0}}\|_{\infty} < 1 \right), \\ &\geq \liminf_{r \to +\infty} \mathbb{P} \left(\|\zeta_{X,\lambda,s^{0}}\|_{\infty} < 1 \right). \end{split}$$

Asymptotic full power and asymptotic control of the FWER) According to (15), asymptotically the power is equal to 1, namely $\lim_{r\to+\infty} \mathbb{P}(\forall i \in I, S(\widehat{\beta}_i^{\mathrm{L}}(\lambda, r)) = s_i^0) = 1$. Now let us prove that the FWER is controlled asymptotically. Let us remind that $\mathbb{P}(\|\zeta_{X,\lambda,s^0}\|_{\infty} < 1) = \gamma$ and $\mathbb{P}(\|\zeta_{X,\lambda,s^0}\|_{\infty} \leq 1) = \bar{\gamma}$. Using asymptotic results given above one may deduce the following inequalities.

$$\bar{\gamma} \geq \limsup_{r \to +\infty} \mathbb{P}(S(\widehat{\beta}^{L}(\lambda, r)) = s^{0}),$$

$$\geq \limsup_{r \to +\infty} \mathbb{P}\left(\forall i \in I, S(\widehat{\beta}^{L}_{i}(\lambda, r)) = s^{0}_{i} \text{ and } \forall i \notin I, \widehat{\beta}^{L}_{i}(\lambda, r) = 0\right),$$

$$\geq \limsup_{r \to +\infty} \mathbb{P}(\forall i \notin I, \widehat{\beta}^{L}_{i}(\lambda, r) = 0).$$
(16)

The last inequality comes from (15). Similarly, we have

$$\gamma \le \liminf_{r \to +\infty} \mathbb{P}(\forall i \notin I, \widehat{\beta}_i^L(\lambda, r) = 0).$$
(17)

Consequently, by taking the complement to 1 of the inequalities given in (16) and (17), one may deduce that

$$\liminf_{r \to +\infty} \mathbb{P}(\exists i \notin I, \widehat{\beta}_i^L(\lambda, r) \neq 0) \ge 1 - \bar{\gamma} \text{ and } \limsup_{r \to +\infty} \mathbb{P}(\exists i \notin I, \widehat{\beta}_i^L(\lambda, r) \neq 0) \le 1 - \gamma.$$

Proof of Theorem 1

Lemma 4 provides the same result for BPDN as does Lemma 2 for LASSO. These both lemmas are the keystones to prove Theorem 1.

Lemma 4 Let $(\beta^{(r)})_{r\in\mathbb{N}}$ be a sequence of \mathbb{R}^p satisfying conditions 1) and 2) of Assumption 1, let us assume that s^0 is identifiable with respect to the ℓ_1 norm and let set $u_r = \|\beta^{(r)}\|_1$ then

$$\lim_{r \to +\infty} \frac{\widehat{\beta}^{\mathrm{BPDN}}(\varepsilon,r) - \beta^{(r)}}{u_r} = 0.$$

Proof: Let us define $u(\varepsilon) \in \mathbb{R}^p$ as follows

$$u(\varepsilon) := \underset{b \in \mathbb{R}^p}{\operatorname{argmin}} \|b\|_1 \text{ subject to } Xb = \varepsilon.$$

Because $X(u(\varepsilon)) = \varepsilon$, we have $Y(\varepsilon) = X(\beta^{(r)} + u(\varepsilon))$ and because $\widehat{\beta}^{BPDN}(\varepsilon, r)$ is an admissible point of (3), one deduces the following inequality

$$\left\|\frac{1}{u_r}X\widehat{\beta}^{\text{BPDN}}(\varepsilon,r) - \frac{1}{u_r}X\beta^{(r)}\right\|_2 \le \left\|\frac{1}{u_r}X\widehat{\beta}^{\text{BPDN}}(\varepsilon,r) - \frac{1}{u_r}Y\right\|_2 + \left\|\frac{1}{u_r}Y - \frac{1}{u_r}X\beta^{(r)}\right\|_2 \le \frac{\sqrt{R}}{u_r} + \frac{\|Xu(\varepsilon)\|_2}{u_r}.$$
 (18)

Because $\beta^{(r)} + u(\varepsilon)$ is an admissible point of problem (3) and because $\hat{\beta}^{BPDN}(\varepsilon, r)$ is the minimizer of (3), one may deduce that the following inequalities hold

$$\frac{1}{u_r} \|\widehat{\beta}^{\text{BPDN}}(\varepsilon, r)\|_1 \le \frac{1}{u_r} \|\beta^{(r)} + u(\varepsilon)\|_1 \le 1 + \frac{\|u(\varepsilon)\|_1}{u_r}.$$
(19)

Because u_r tends to $+\infty$ then, according to (19), the sequence $((\widehat{\beta}^{L}(\varepsilon, r) - \beta^{(r)})/u_r)_{r \in \mathbb{N}^*}$ is bounded since the following superior limit is finite

$$\limsup_{r \to +\infty} \left\| \frac{\widehat{\beta}^{\text{BPDN}}(\varepsilon, r) - \beta^{(r)}}{u_r} \right\|_1 \le 2.$$

Consequently, to prove that $\lim_{r\to+\infty} (\widehat{\beta}^{\text{BPDN}}(\varepsilon,r) - \beta^{(r)})/u_r = \mathbf{0}$ it is sufficient to show that $\mathbf{0}$ is the unique limit point of this sequence. Let $((\widehat{\beta}^{\text{L}}(\varepsilon,\phi(r)) - \beta^{(\phi(r))})/u_{\phi(r)})_{r\in\mathbb{N}^*}$ be a converging subsequence to l (with $\phi: \mathbb{N}^* \to \mathbb{N}^*$ strictly increasing) and without loss of generality, let us assume $\lim_{r\to+\infty} \widehat{\beta}^{\text{BPDN}}(\varepsilon,\phi(r))/u_{\phi(r)} = v$ and $\lim_{r\to+\infty} b^{(\phi(r))}/u_{\phi(r)} = v'$ so that l = v - v'. By (18) and (19) one may deduce that

$$Xv = Xv'$$
 and $||v||_1 \le 1$.

Since, whatever $r \ge 0$, we have $S(\beta^{(\phi(r))}/u_{\phi(r)}) = s^0$ where s^0 is identifiable with respect to the ℓ_1 norm then, according to Proposition 2, one may deduce that $\beta^{(\phi(r))}/u_{\phi(r)}$ is an unitary vector satisfying the identifiability condition. Consequently, $\|v'\|_1 = 1$ and v' is identifiable with respect to the ℓ_1 norm. Consequently, v = v' and thus $l = \mathbf{0}$ is the unique limit point, which implies that

$$\lim_{r \to +\infty} \frac{\widehat{\beta}^{\text{BPDN}}(\varepsilon, r) - \beta^{(r)}}{u_r} = \mathbf{0}.$$

Lemma 5 is useful to prove in Theorem 1 that when s^0 is not identifiable then sign estimator derived from thresholded LASSO cannot recover s^0 .

Lemma 5 Let X be a matrix in general position, then the random vector $\hat{\beta}$ is identifiable with respect to X and the ℓ_1 norm.

Proof: Let us remind that when X is in general position then the minimizer $\hat{\beta}$ is unique. Let us assume that $\hat{\beta}$ is not identifiable with respect to X and the ℓ_1 norm, then there exists $b \in \mathbb{R}^p$ such that $Xb = X\hat{\beta}$ and $\|b\|_1 \leq \|\hat{\beta}\|_1$. Consequently, for LASSO, one may deduce that

$$\|Y - Xb\|^2 + \lambda \|b\|_1 \le \|Y - X\widehat{\beta}^{\mathbf{L}}\|^2 + \lambda \|\widehat{\beta}^{\mathbf{L}}\|_1.$$

This inequality contradicts $\hat{\beta}^{L}$ as the unique minimizer of (2). Similarly, when $\hat{\beta}^{BPDN}$ is not identifiable with respect to the ℓ_1 norm then $\hat{\beta}^{BPDN}$ is not the unique minimizer of (3), which provides a contradiction.

For the proofs of Theorem 1 and the proof of Proposition 2 we need to introduce the following inequality which characterizes the identifiability condition [8]. A vector $b \in \mathbb{R}^p$ is identifiable with respect to X and the ℓ_1 norm if and only if the following inequality holds

$$\forall h \in \ker(X) \setminus \{\mathbf{0}\}, \left| \sum_{i \in \operatorname{supp}(b)} S(b)h_i \right| < \sum_{i \notin \operatorname{supp}(b)} |h_i|.$$

$$(20)$$

Proof of Theorem 1:

Necessary condition: Let us assume that $S(\beta)$ is not identifiable with respect to the ℓ_1 norm. Let us show that when the following events hold

$$\operatorname{supp}^{-}(\beta) \subset \operatorname{supp}^{-}(\widehat{\beta}(\varepsilon)) \text{ and } \operatorname{supp}^{+}(\beta) \subset \operatorname{supp}^{+}(\widehat{\beta}(\varepsilon)),$$
(21)

then inequality (20) occurs which contradicts that $S(\beta)$ is not identifiable with respect to the ℓ_1 norm. Let $h \in \ker(X) \setminus \{\mathbf{0}\}$. On the one hand, when (21) occurs, we have

$$\begin{aligned} \left| \sum_{i \in \text{supp}(\beta)} S(\beta_i) h_i \right| &= \left| -\sum_{\text{supp}^-(\beta)} h_i + \sum_{\text{supp}^+(\beta)} h_i \right|, \\ &= \left| -\sum_{i \in \text{supp}^-(\widehat{\beta}(\varepsilon))} h_i + \sum_{i \in \text{supp}^-(\widehat{\beta}(\varepsilon)) \setminus \text{supp}^-(\beta)} h_i + \sum_{i \in \text{supp}^+(\widehat{\beta}(\varepsilon))} h_i - \sum_{i \in \text{supp}^+(\widehat{\beta}(\varepsilon)) \setminus \text{supp}^+(\beta)} h_i \right|, \\ &\leq \left| -\sum_{i \in \text{supp}^-(\widehat{\beta}(\varepsilon))} h_i + \sum_{i \in \text{supp}^+(\widehat{\beta}(\varepsilon))} h_i \right| + \sum_{i \in \text{supp}(\widehat{\beta}(\varepsilon)) \setminus \text{supp}(\beta)} |h_i|. \end{aligned}$$

On the other hand, according to Lemma 5, $\hat{\beta}(\varepsilon)$ is identifiable with respect to the ℓ_1 norm then (20) occurs implying the following inequality

$$\left| -\sum_{i \in \operatorname{supp}^{-}(\widehat{\beta}(\varepsilon))} h_i + \sum_{i \in \operatorname{supp}^{+}(\widehat{\beta}(\varepsilon))} h_i \right| + \sum_{i \in \operatorname{supp}(\widehat{\beta}(\varepsilon)) \setminus \operatorname{supp}(\beta)} |h_i| \\ < \sum_{i \notin \operatorname{supp}(\widehat{\beta}(\varepsilon))} |h_i| + \sum_{i \in \operatorname{supp}(\widehat{\beta}(\varepsilon)) \setminus \operatorname{supp}(\beta)} |h_i| = \sum_{i \notin \operatorname{supp}(\beta)} |h_i|.$$

Consequently the following inequality holds

$$\forall h \in \ker(X) \setminus \{\mathbf{0}\}, \left| \sum_{i \in \operatorname{supp}(\beta)} S(\beta_i) h_i \right| < \sum_{i \notin \operatorname{supp}(\beta)} |h_i|,$$

which, according to (20), contradicts that $S(\beta)$ is not identifiable with respect to the ℓ_1 norm.

Sufficient condition: Let us remind that according to condition **3**) of Assumption 1 the following inequality holds

$$\forall r \in \mathbb{N}, \frac{\min\{|\beta_i^{(r)}|, i \in \operatorname{supp}(s^0)\}}{\|\beta^{(r)}\|_{\infty}} \ge q > 0.$$

According to Lemmas 2 and 4, when s^0 is identifiable with respect to the ℓ_1 norm then

$$\lim_{r \to +\infty} \frac{\widehat{\beta}(\varepsilon, r) - \beta^{(r)}}{\|\beta^{(r)}\|_{\infty}} = 0.$$

Therefore, there exists $r_0(\varepsilon) \ge 0$ such that

$$\forall r \ge r_0(\varepsilon), \left\|\frac{\widehat{\beta}(\varepsilon, r) - \beta^{(r)}}{\|\beta^{(r)}\|_{\infty}}\right\|_{\infty} < q/2 \Leftrightarrow \forall i \in \{1, \dots, p\}, \forall r \ge r_0(\varepsilon), \left|\frac{\widehat{\beta}_i(\varepsilon, r) - \beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}}\right| < q/2.$$

Consequently, when $r \ge r_0(\varepsilon)$, whatever $i \notin \operatorname{supp}(s^0)$ (thus when $\beta_i^{(r)} = 0$) the following inequalities hold

$$\begin{aligned} \forall i \notin \operatorname{supp}(s^0), \left| \frac{\widehat{\beta}_i(\varepsilon, r)}{\|\beta^{(r)}\|_{\infty}} \right| < q/2, \\ \Rightarrow \quad -\|\beta^{(r)}\|_{\infty}q/2 < \min_{i \notin \operatorname{supp}(s^0)} \left\{ \widehat{\beta}_i(\varepsilon, r) \right\} \leq \max_{i \notin \operatorname{supp}(s^0)} \left\{ \widehat{\beta}_i(\varepsilon, r) \right\} < \|\beta^{(r)}\|_{\infty}q/2. \end{aligned}$$

Whatever $i \in \text{supp}^+(s^0)$ (thus when $\beta_i^{(r)} > 0$) the following inequalities hold

$$\begin{aligned} \forall i \in \operatorname{supp}^+(s^0), \frac{\widehat{\beta}_i(\varepsilon, r)}{\|\beta^{(r)}\|_{\infty}} &\geq -\left|\frac{\widehat{\beta}_i(\varepsilon, r) - \beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}}\right| + \frac{\beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}}, \\ \Rightarrow & \min_{i \in \operatorname{supp}^+(s^0)} \left\{\frac{\widehat{\beta}_i(\varepsilon, r)}{\|\beta^{(r)}\|_{\infty}}\right\} &> -q/2 + q = q/2, \\ \Rightarrow & \min_{i \in \operatorname{supp}^+(s^0)} \left\{\widehat{\beta}_i(\varepsilon, r)\right\} > \|\beta^{(r)}\|_{\infty}q/2. \end{aligned}$$

Whatever $i \in \text{supp}^{-}(s^{0})$ (thus when $\beta_{i}^{(r)} < 0$) the following inequalities hold

$$\begin{aligned} \forall i \in \mathrm{supp}^+(s^0), \frac{\widehat{\beta}_i(\varepsilon, r)}{\|\beta^{(r)}\|_{\infty}} &\leq \left| \frac{\widehat{\beta}_i(\varepsilon, r) - \beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}} \right| + \frac{\beta_i^{(r)}}{\|\beta^{(r)}\|_{\infty}} \\ \Rightarrow & \max_{i \in \mathrm{supp}^-(s^0)} \left\{ \frac{\widehat{\beta}_i(\varepsilon, r)}{\|\beta^{(r)}\|_{\infty}} \right\} < q/2 - q = -q/2, \\ \Rightarrow & \max_{i \in \mathrm{supp}^-(s^0)} \left\{ \widehat{\beta}_i(\varepsilon, r) \right\} < -\|\beta^{(r)}\|_{\infty} q/2. \end{aligned}$$

Finally, when $r \geq r_0(\varepsilon)$ we have

i)

$$\operatorname{supp}^{-}(s^{0}) \subset \operatorname{supp}^{-}(\widehat{\beta}_{i}(\varepsilon, r)) \text{ and } \operatorname{supp}^{+}(s^{0}) \subset \operatorname{supp}^{+}(\widehat{\beta}_{i}(\varepsilon, r)).$$

ii)

$$\max_{i \in \operatorname{supp}^{-}(s^{0})} \left\{ \widehat{\beta}_{i}(\varepsilon, r) \right\} < \min_{i \notin \operatorname{supp}(s^{0})} \left\{ \widehat{\beta}_{i}(\varepsilon, r) \right\} \le \max_{i \notin \operatorname{supp}(s^{0})} \left\{ \widehat{\beta}_{i}(\varepsilon, r) \right\} < \min_{i \in \operatorname{supp}^{+}(s^{0})} \left\{ \widehat{\beta}_{i}(\varepsilon, r) \right\}.$$

These achieve the proof of the sufficient condition.

Proof of propositions

The proof of Proposition 1, provided below, is the one reported in the PhD manuscript of Tardivel [29].

Proof of Proposition 1: From Daubechies et al. [8], β is a parameter having a minimal ℓ_1 norm, namely $X\beta = X\gamma \Rightarrow \|\gamma\|_1 \ge \|\beta\|_1$ holds if and only if the following inequality occurs

$$\forall h \in \ker(X), \left| \sum_{i \in I} S(\beta_i) h_i \right| \le \sum_{i \notin I} |h_i|.$$
(22)

We are going to show that when the irrepresentable condition holds for β then the inequality (20) holds.

Let $h \in \ker(X)$ and let us remind that h_I and $h_{\overline{I}}$ denote respectively vectors $(h_i)_{i \in I}$ and $(h_i)_{i \notin I}$. Then the following equality holds

$$\sum_{i \in I} S(\beta_i) h_i = h'_I S(\beta_I) = h'_I X'_I X_I (X'_I X_I)^{-1} S(\beta_I).$$

Because $\mathbf{0} = Xh = X_Ih_I + X_{\overline{I}}h_{\overline{I}}$, one may deduce the following inequalities

$$|h'_{I}S(\beta_{I})| = |h'_{\overline{I}}X'_{\overline{I}}X_{I}(X'_{I}X_{I})^{-1}S(\beta_{I})|,$$

$$\leq ||h_{\overline{I}}||_{1}||X'_{\overline{I}}X_{I}(X'_{I}X_{I})^{-1}S(\beta_{I})||_{\infty}.$$
(23)

Consequently, when the irrepresentable condition holds for β , namely when $\|X'_{\overline{I}}X_{I}(X'_{I}X_{I})^{-1}S(b^{*}_{I})\|_{\infty} \leq 1$, then the inequality (23) gives $|h'_{I}S(\beta_{I})| \leq \|h_{\overline{I}}\|_{1}$. Thus, by the equivalence given in (22), β is a solution of the following basis pursuit problem

minimize
$$\|\gamma\|_1$$
 subject to $X\gamma = X\beta$

Because X is in general position the previous optimisation problem has a unique solution (see, e.g., Proposition 1 in Appendix) thus $X\beta = X\gamma$ and $\gamma \neq \beta$ implies that $\|\gamma\|_1 > \|\beta\|_1$, namely β is identifiable with respect to the ℓ_1 norm.

Let us notice that when the inequality in the irrepresentable condition is strict, Theorem 1 remains true without assuming that X is in general position.

Proof of Proposition 2: Because b is identifiable with respect to the ℓ_1 norm and because $S(\tilde{b}) = S(b)$ implies $\operatorname{supp}(\tilde{b}) = \operatorname{supp}(b)$, then the following inequality holds

$$\forall h \in \ker(X) \setminus \{\mathbf{0}\}, \left| \sum_{i \in \operatorname{supp}(\tilde{b})} S(\tilde{b}_i) h_i \right| < \sum_{i \notin \operatorname{supp}(\tilde{b})} |h_i|.$$

Consequently, according to (20), parameter \tilde{b} is identifiable with respect to the ℓ_1 norm.

6.3 Comparisons of conditions for sign recovery and support recovery

Let $X = (X_1 | \dots | X_p) \in \mathbb{R}^{n \times p}$ and $\beta \in \mathbb{R}^p$.

When ||X₁||₂ = ··· = ||X_p||₂ = 1, we say that β satisfies the mutual coherence condition once the following inequalities occurs

$$\operatorname{card}(\operatorname{supp}(\beta)) < \frac{1+1/M(X)}{2}$$
, where $M(X) = \max_{i \neq j} |X'_i X_j|$.

• We say that β satisfies the irrepresentability condition once

$$\|X'_{\overline{I}}X_I(X'_IX_I)^{-1}\operatorname{sign}(\beta_I)\|_{\infty} < 1, \text{ where } I = \operatorname{supp}(\beta).$$

Moreover, we say that the irrepresentability condition uniformly hold on the support I of β once

$$\forall \theta \in \{-1, 1\}^{|I|}, \|X_{\overline{I}}'X_{I}(X_{I}'X_{I})^{-1}\theta\|_{\infty} < 1.$$

• We say that β satisfies the stable nullspace property once

$$\forall h \in \ker(X) \setminus \{0\}, \sum_{i \in \operatorname{supp}(\beta)} |h_i| < \sum_{i \notin \operatorname{supp}(\beta)} |h_i|.$$

• We remind that β is identifiable with respect to X and the ℓ_1 norm once

$$\forall h \in \ker(X) \setminus \{0\}, \left| \sum_{i \in \operatorname{supp}(\beta)} \operatorname{sign}(\beta_i) h_i \right| < \sum_{i \notin \operatorname{supp}(\beta)} |h_i|$$

Table 2 summarizes comparisons between above conditions

Mutual coherence
$$\stackrel{[33]}{\Longrightarrow}$$
 Irrepresentability
 $\downarrow [19] \qquad \downarrow \qquad Uniform irrepresentability \stackrel{[9]}{\Longrightarrow}$ Stable nullspace
Stable nullspace \implies Identifiability

Table 2: This figure provides implication between above conditions. The implication "stable nullspace \Rightarrow identifiability" is straightforward and the implication "Irrepresentability \Rightarrow Identifiability" is given in Proposition 1. Other implications are proved in [9, 19, 33]. Note that an exhaustive scheme is given in [4, 33].

One may observe that the mutual coherence is a very strong condition once two columns of X are almost equal (namely when M(X) is close to 1). As an example, M(X) = 0.99 when X is the matrix from the riboflavin dataset and thus the mutual coherence condition is extremely strong. Nevertheless, when β has only one non null component, the mutual coherence condition holds and thus β satisfies both the irrepresentability and identifiability conditions as illustrated on Figure 5.

Finally when X is a $n \times p$ standard Gaussian matrix and n < p are both very large then β is identifiable with respect to X and the ℓ_1 norm with a very large probability (resp. very small probability) once card(supp(β)) $< n\rho(n/p)$ (resp. once card(supp(β)) $> n\rho(n/p)$), where $\rho(\cdot)$ is the transition curve of Donoho and Tanner [13], presented in Figure 7.

Supplementary material

We have already said that when X is in general position the minimizer of problem (2) (resp. problem (3)) is unique. Concerning LASSO, a sketch of proof given in Tibshirani [32] shows the uniqueness of the LASSO estimator when X is in general position. In order to provide a self-contained article, we show that when X is in general position, the minimizer of problem (3) is unique when R = 0 as well as when R > 0. We have

Figure 7: This figure illustrates the transition curve of Donoho and Tanner [13].

already stressed that when β is identifiable with respect to the ℓ_1 norm then β is sparse. We show that when the identifiability holds for β then the family $(X_i)_{i \in \text{supp}(\beta)}$ is linearly independent and thus the number of components of β equal to 0 is larger than p - n. Finally, a proof that the stable nullspace property implies the identifiability condition is given

References

- [1] Rina Foygel Barber and Emmanuel J Candès. Controlling the false discovery rate via knockoffs. <u>The</u> Annals of Statistics, 43(5):2055–2085, 2015.
- [2] Mohsen Bayati and Andrea Montanari. The LASSO risk for Gaussian matrices. <u>IEEE Transactions on</u> Information Theory, 58(4):1997–2017, 2012.
- [3] Peter Bühlmann, Markus Kalisch, and Lukas Meier. High-dimensional statistics with a view toward applications in biology. 2014.
- [4] Peter Bühlmann and Sara van de Geer. <u>Statistics for High-Dimensional Data: Methods, Theory and</u> Applications. Springer, 2011.
- [5] Emmanuel Candes, Yingying Fan, Lucas Janson, and Jinchi Lv. Panning for gold: Model-x knockoffs for high-dimensional controlled variable selection. arXiv preprint arXiv:1610.02351, 2016.
- [6] Emmanuel J Candès, Yaniv Plan, et al. Near-ideal model selection by 1 minimization. <u>The Annals of</u> Statistics, 37(5A):2145-2177, 2009.
- [7] Shaobing Chen and David Donoho. Basis pursuit. In <u>Proceedings of 1994 28th Asilomar Conference on</u> Signals, Systems and Computers, volume 1, pages 41–44. IEEE, 1994.

- [8] Ingrid Daubechies, Ronald DeVore, Massimo Fornasier, and C Sinan Güntürk. Iteratively reweighted least squares minimization for sparse recovery. <u>Communications on pure and applied mathematics</u>, 63(1):1–38, 2010.
- [9] Pascaline Descloux and Sylvain Sardy. Model selection with lasso-zero: adding straw to the haystack to better find needles. Journal of Computational and Graphical Statistics, pages 1–29, 2020.
- [10] Lee H Dicker. Variance estimation in high-dimensional linear models. Biometrika, 101(2):269–284, 2014.
- [11] David Donoho and Jared Tanner. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. <u>Philosophical Transactions of the Royal</u> Society A: Mathematical, Physical and Engineering Sciences, 367(1906):4273-4293, 2009.
- [12] David L Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via 1 minimization. Proceedings of the National Academy of Sciences, 100(5):2197–2202, 2003.
- [13] David L Donoho and Jared Tanner. Precise undersampling theorems. <u>Proceedings of the IEEE</u>, 98(6):913– 924, 2010.
- [14] Charles Dossal, Marie-Line Chabanol, Gabriel Peyré, and Jalal Fadili. Sharp support recovery from noisy random measurements by ℓ₁-minimization. <u>Applied and Computational Harmonic Analysis</u>, 33(1):24–43, 2012.
- [15] Karl Ewald and Ulrike Schneider. On the distribution, model selection properties and uniqueness of the lasso estimator in low and high dimensions. Electronic Journal of Statistics, 14(1):944–969, 2020.
- [16] Simon Foucart and Holger Rauhut. <u>A mathematical introduction to compressive sensing</u>, volume 1. Springer, 2013.
- [17] Andrew Gelman and Francis Tuerlinckx. Type s error rates for classical and bayesian single and multiple comparison procedures. Computational Statistics, 15(3):373–390, 2000.
- [18] Caroline Giacobino, Sylvain Sardy, Jairo Diaz-Rodriguez, Nick Hengartner, et al. Quantile universal threshold. Electronic Journal of Statistics, 11(2):4701–4722, 2017.
- [19] Rémi Gribonval and Morten Nielsen. Sparse representations in unions of bases. <u>IEEE Transactions on</u> Information Theory, 49(12):3320–3325, 2003.
- [20] Jian Huang, Shuangge Ma, and Cun-Hui Zhang. Adaptive lasso for sparse high-dimensional regression models. Statistica Sinica, 18(4):1603, 2008.
- [21] Karim Lounici. Sup-norm convergence rate and sign concentration property of lasso and dantzig estimators. Electronic Journal of statistics, 2:90–102, 2008.

- [22] Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. The Annals of Statistics, 34(3):1436–1462, 2006.
- [23] Nicolai Meinshausen and Bin Yu. Lasso-type recovery of sparse representations for high-dimensional data. The Annals of Statistics, 37(1):246–270, 2009.
- [24] Piotr Pokarowski, Wojciech Rejchel, Agnieszka Soltys, Michal Frej, and Jan Mielniczuk. Improving lasso for model selection and prediction. arXiv preprint arXiv:1907.03025, 2019.
- [25] Venkatesh Saligrama and Manqi Zhao. Thresholded basis pursuit: Lp algorithm for order-wise optimal support recovery for sparse and approximately sparse signals from noisy random measurements. <u>IEEE</u> Transactions on Information Theory, 57(3):1567–1586, 2011.
- [26] Ulrike Schneider and Patrick Tardivel. The geometry of uniqueness and model selection of penalized estimators including slope, lasso, and basis pursuit. arXiv preprint arXiv:2004.09106, 2020.
- [27] Matthew Stephens. False discovery rates: a new deal. Biostatistics, 18(2):275–294, 2017.
- [28] Weijie J Su, Małgorzata Bogdan, and Emmanuel J. Candès. False discoveries occur early on the lasso path. The Annals of Statistics, 45(5):2133–2150, 2017.
- [29] Patrick Tardivel. <u>Représentation parcimonieuse et procédures de tests multiples: application à la métabolomique. PhD thesis, Université de Toulouse, Université Toulouse III-Paul Sabatier, 2017.</u>
- [30] Patrick JC Tardivel, Rémi Servien, and Didier Concordet. Sparsest representations and approximations of an underdetermined linear system. Inverse Problems, 34(5):055002, 2018.
- [31] Robert Tibshirani. Regression shrinkage and selection via the lasso. <u>Journal of the Royal Statistical Society</u>. Series B (Methodological), 58(1):267–288, 1996.
- [32] Ryan J Tibshirani. The lasso problem and uniqueness. Electronic Journal of Statistics, 7:1456–1490, 2013.
- [33] Sara A Van de Geer. High-dimensional generalized linear models and the lasso. <u>The Annals of Statistics</u>, pages 614–645, 2008.
- [34] Sara A Van De Geer, Peter Bühlmann, et al. On the conditions used to prove oracle results for the lasso. Electronic Journal of Statistics, 3:1360–1392, 2009.
- [35] Martin J Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using constrained quadratic programming (lasso). IEEE transactions on information theory, 55(5):2183–2202, 2009.
- [36] S. Wang, H. Weng, and A. Maleki. Which bridge estimator is the best for variable selection ? arxiv, 2018.

- [37] Asaf Weinstein, Rina Barber, and Emmanuel J. Candès. A power and prediction analysis for knockoffs with lasso statistics. arXiv:1712.06465, 2017.
- [38] Asaf Weinstein, Weijie J Su, Małgorzata Bogdan, Rina F Barber, and Emmanuel J Candès. A power analysis for knockoffs with the lasso coefficient-difference statistic. arXiv preprint arXiv:2007.15346, 2020.
- [39] Peng Zhao and Bin Yu. On model selection consistency of lasso. <u>The Journal of Machine Learning</u> Research, 7:2541–2563, 2006.
- [40] Hui Zou. The adaptive lasso and its oracle properties. Journal of the American statistical association, 101(476):1418-1429, 2006.