



HAL
open science

Transductive versions of the LASSO and the Dantzig Selector

Pierre Alquier, Mohamed Hebiri

► **To cite this version:**

Pierre Alquier, Mohamed Hebiri. Transductive versions of the LASSO and the Dantzig Selector. 2009. hal-00390851v2

HAL Id: hal-00390851

<https://hal.science/hal-00390851v2>

Preprint submitted on 6 Jun 2009 (v2), last revised 4 May 2010 (v3)

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.

Transductive versions of the LASSO and the Dantzig Selector

Pierre Alquier and Mohamed Hebiri

Abstract

We consider the linear regression problem, where the number p of covariates is possibly larger than the number n of observations $(x_i, y_i)_{i \leq n}$, under sparsity assumptions. On the one hand, several methods have been successfully proposed to perform this task, for example the LASSO in [Tib96] or the Dantzig Selector in [CT07]. On the other hand, consider new values $(x_i)_{n+1 \leq i \leq m}$. If one wants to estimate the corresponding y_i 's, one should think of a specific estimator devoted to this task, referred in [Vap98] as a "transductive" estimator. This estimator may differ from an estimator designed to the more general task "estimate on the whole domain". In this work, we propose a generalized version both of the LASSO and the Dantzig Selector, based on the geometrical remarks about the LASSO in [Alq08, AH08]. The "usual" LASSO and Dantzig Selector, as well as new estimators interpreted as transductive versions of the LASSO, appear as special cases. These estimators are interesting at least from a theoretical point of view: we can give theoretical guarantees for these estimators under hypotheses that are relaxed versions of the hypotheses required in the papers about the "usual" LASSO. These estimators can also be efficiently computed, with results comparable to the ones of the LASSO.

Contents

1	Introduction	2
2	Preliminaries	4
3	The "easy case": $\text{Ker}(X) = \text{Ker}(Z)$	5
3.1	Definition of the estimators	5
3.2	Theoretical results	7
4	An extension to the general case	9
4.1	General remarks	9
4.2	An example: small labeled dataset, large unlabeled dataset . . .	9
5	Experimental results	11

6	Conclusion	15
7	Proofs	16
7.1	Proof of Propositions 3.1 and 3.2	16
7.2	A useful Lemma	17
7.3	Proof of Theorems 3.3 and 3.4	18
7.4	Proof of Theorem 4.1	21
7.5	Proof of Proposition 4.2	23

1 Introduction

In many modern applications, a statistician often have to deal with very large datasets. Regression problems may involve a large number of covariates p , possibly larger than the sample size n . In this situation, a major issue is dimension reduction, which can be performed through the selection of a small amount of relevant covariates. For this purpose, numerous regression methods have been proposed in the literature, ranging from the classical information criteria such as AIC [Aka73] and BIC [Sch78] to the more recent sparse methods, known as the LASSO [Tib96], and the Dantzig Selector [CT07]. Regularized regression methods have recently witnessed several developments due to the attractive feature of computational feasibility, even for high dimensional data (i.e., when the number of covariates p is large). We focus on the usual linear regression model:

$$y_i = x_i \beta^* + \varepsilon_i, \quad i = 1, \dots, n, \quad (1)$$

where the design $x_i = (x_{i,1}, \dots, x_{i,p}) \in \mathbb{R}^p$ is deterministic, $\beta^* = (\beta_1^*, \dots, \beta_p^*)' \in \mathbb{R}^p$ is the unknown parameter and $\varepsilon_1, \dots, \varepsilon_n$ are i.i.d. centered Gaussian random variables with known variance σ^2 . Let X denote the matrix with i -th line equal to x_i , and let X_j denote its j -th column, with $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, p\}$. So:

$$X = (x'_1, \dots, x'_n)' = (X_1, \dots, X_p).$$

For the sake of simplicity, we will assume that the observations are normalized in such a way that $X'_j X_j / n = 1$. We denote by Y the vector $Y = (y_1, \dots, y_n)'$. For all $\alpha \leq 1$ and any vector $v \in \mathbb{R}^d$, we set $\|\cdot\|_\alpha$, the norm: $\|v\|_\alpha = (|v_1|^\alpha + \dots + |v_d|^\alpha)^{1/\alpha}$. In particular $\|\cdot\|_2$ is the euclidean norm. Moreover for all $d \in \mathbb{N}$, we use the notation $\|v\|_0 = \sum_{i=1}^d \mathbf{1}(v_i \neq 0)$.

The problem of estimating the regression parameter in the high dimensional setting have been extensively studied in the statistical literature. Among others, the LASSO [Tib96] (denoted by $\hat{\beta}^L$), the Dantzig Selector [CT07] (denoted by $\hat{\beta}^{DS}$) and the non-negative garrote (in Yuan and Lin [YL07], denoted by $\hat{\beta}^{NNG}$) have been proposed to deal with this problem for a large p , even for $p > n$. These estimators give very good practical results. For instance in [Tib96], simulations and tests on real data have been provided for the LASSO. We also refer to [Kol07, Kol09, MVdGB08, vdG08, DT07, CH08] for related work with different

estimators: non-quadratic loss, penalties slightly different from ℓ_1 and random design.

From a theoretical point of view, Sparsity Inequalities (SI) have been proved for these estimators under different assumptions. That is upper bounds of order of $\mathcal{O}(\sigma^2 \|\beta^*\|_0 \log(p)/n)$ for the errors $(1/n) \|X\hat{\beta} - X\beta^*\|_2^2$ and $\|\hat{\beta} - \beta^*\|_2^2$ have been derived, where $\hat{\beta}$ is one of the estimators mentioned above. In particular these bounds involve the number of non-zero coordinates in β^* (multiplied by $\log(p)$), instead of dimension p . Such bounds garanty that under some assumptions, $X\hat{\beta}$ and $\hat{\beta}$ are good estimators of $X\beta^*$ and β^* respectively. According to the LASSO $\hat{\beta}^L$, these SI are given for example in [BTW07, BRT07], whereas [CT07, BRT07] provided SI for the Dantzig Selector $\hat{\beta}^{DS}$. On the other hand, Bunea [Bun08] establishes conditions which ensure $\hat{\beta}^L$ and β^* have the same null coordinates. Analog results for $\hat{\beta}^{DS}$ can be found in [Lou08].

Now, let us assume that we are given additional observations $x_i \in \mathbb{R}^p$ for $n+1 \leq i \leq m$ (with $m > n$), and introduce the matrix $Z = (x'_1, \dots, x'_m)'$. Assume that the objective of the statistician is precisely to estimate $Z\beta^*$: namely, he cares about predicting what would be the labels attached to the additional x_i 's. It is argued in [Vap98] that in such a case, a specific estimator devoted to this task should be considered: the transductive estimator. This estimator differs from an estimator tailored for the estimation of β^* or $X\beta^*$ like the LASSO. Indeed one usually builds an estimator $\hat{\beta}(X, Y)$ and then computes $Z\hat{\beta}(X, Y)$ to estimate $Z\beta^*$. The approach taken here is to consider estimators $\hat{\beta}(X, Y, Z)$ exploiting the knowledge of Z , and then to compute $Z\hat{\beta}(X, Y, Z)$.

Some methods in supervised classification or regression were successfully extended to the transductive setting, such as the well-known Support Vector Machines (SVM) in [Vap98], the Gibbs estimators in [Cat07]. It is argued in the semi-supervised learning literature (see for example [CSZ06] for a recent survey) that taking into account the information on the design given by the new additional x_i 's has a stabilizing effect on the estimator.

In this paper, we study a family of estimators which generalizes the LASSO and the Dantzig Selector. The considered family depends on a $q \times p$ matrix A , with $q \in \mathbb{N}$, whose choice allows to adapt the estimator to the objective of the statistician. The choice of the matrix A allows to cover transductive setting.

The rest of paper is organized as follows. In the next section, we motivate the use of the studied family of estimators through geometrical considerations stated in [AH08]. In Sections 3 and 4, we establish Sparsity Inequalities for these estimators. A discussion on the assumptions needed to prove the SI is also provided. In particular, it is shown that the estimators devoted to the transductive setting satisfy these SI with weaker assumptions than those needed by the LASSO or the Dantzig Selector, when $m > p > n$. That is, when the number of news points is large enough. The implementation of our estimators and some numerical experiments are the purpose of Section 5. The results clearly show that the use of a transductive version of the LASSO may improve the performance of the estimation. All proofs of the theoretical results are postponed to Section 7.

2 Preliminaries

In this section we state geometrical considerations (projections on a confidence region) for the LASSO and the Dantzig Selector. These motivate the introduction of our estimators. Finally we discuss the different objectives considered in this paper.

Let us remind that a definition of the LASSO estimate is given by

$$\arg \min_{\beta \in \mathbb{R}^p} \left\{ \|Y - X\beta\|_2^2 + 2\lambda \|\beta\|_1 \right\}. \quad (2)$$

A dual form (in [OPT00]) of this program is also of interest:

$$\begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|X\beta\|_2^2 \\ s.t. \|X'(Y - X\beta)\|_\infty \leq \lambda; \end{cases} \quad (3)$$

actually it is proved in [Alq08] that any solution of Program 3 is a solution of Program 2 and that the set $\{X\beta\}$ is the same where β is taken among all the solutions of Program 2 or among all the solutions of 3. So both programs are equivalent in terms of estimating $X\beta^*$.

Now, let us remind the definition of the Dantzig Selector:

$$\begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \\ s.t. \|X'(Y - X\beta)\|_\infty \leq \lambda. \end{cases} \quad (4)$$

Alquier [Alq08] observed that both Programs 3 and 4 can be seen as a projection of the null vector $\mathbf{0}_p$ onto the region $\{\beta : \|X'(Y - X\beta)\|_\infty \leq \lambda\}$ that can be interpreted as a confidence region, with confidence $1 - \eta$, for a given λ that depends on η (see Lemma 7.1 here for example). The difference between the two programs is the distance (or semi-distance) used for the projection.

Based on these geometrical considerations, we proposed in [AH08] to study the following transductive estimator:

$$\begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|Z\beta\|_2^2 \\ s.t. \|X'(Y - X\beta)\|_\infty \leq \lambda; \end{cases} \quad (5)$$

that is a projection on the same confidence region, but using a distance adapted to the transductive estimation problem. We proved a Sparsity Inequality for this estimator exploiting a novel sparsity measure.

In this paper, we propose a generalized version of the LASSO and of the Dantzig Selector, based on the same geometrical remark. More precisely for $q \in \mathbb{N}^*$, let A be a $q \times p$ matrix. We propose two general estimators, $\hat{\beta}_{A,\lambda}$ (extension of the LASSO, based on a generalization of Program 2) and $\tilde{\beta}_{A,\lambda}$ (transductive Dantzig Selector, generalization of Program 4). These novel estimators depend on two tuning parameters: $\lambda > 0$ is a regularization parameter, it plays the

same role as the tuning parameter involved in the LASSO, and the matrix A that will allow to adapt the estimator to the objective of the statistician. More particularly, depending on the choice of the matrix A , this estimator can be adapted to one of the following objectives:

- **denoising objective:** the estimation of $X\beta^*$, that is a denoised version of Y . For this purpose, we consider the estimator $\hat{\beta}_{A,\lambda}$, with $A = X$. In this case, the estimator will actually be equal to the LASSO $\hat{\beta}^L$ and $\tilde{\beta}_{A,\lambda}$, with the same choice $A = X$ will be equal to the Dantzig Selector;
- **transductive objective:** the estimation of $Z\beta^*$, by $\hat{\beta}_{A,\lambda}$ or $\tilde{\beta}_{A,\lambda}$, with $A = \sqrt{n/m}Z$. We will refer the corresponding estimators as the "Transductive LASSO" and "Transductive Dantzig Selector";
- **estimation objective:** the estimation of β^* itself, by $\hat{\beta}_{A,\lambda}$, with $A = \sqrt{n}I$. In this case, it appears that both estimators are well defined only in the case $p < n$ and are equal to a soft-thresholded version of the usual least-square estimator.

For both estimators and all the above objectives, we prove SI (Sparsity Inequalities). Moreover, we show that these estimators can easily be computed.

3 The "easy case": $\text{Ker}(X) = \text{Ker}(Z)$

In this section, we deal with the "easy case", where $\text{Ker}(A) = \text{Ker}(X)$ (think of $A = X$, $A = \sqrt{n}I$ or $A = \sqrt{n/m}Z$). This setting is natural at least in the case $p < n$ where both kernels are equal to $\{0\}$ in general. We provide SI (Sparsity Inequality, Theorem 3.3) for the studied estimators, based on the techniques developed in [BRT07].

3.1 Definition of the estimators

Definition 3.1. For a given parameter $\lambda \geq 0$ and any matrix A such that $\text{Ker}(A) = \text{Ker}(X)$, we consider the estimator given by

$$\hat{\beta}_{A,\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ -2Y'X(\widetilde{X'X})^{-1}(A'A)\beta + \beta'(A'A)\beta + 2\lambda\|\Xi_A\beta\|_1 \right\},$$

where $(\widetilde{X'X})^{-1}$ is exactly $(X'X)^{-1}$ if $(X'X)$ is invertible, and any pseudo-inverse of this matrix otherwise, and where Ξ_A is a diagonal matrix whose (j, j) -th coefficient is $\xi_j^{\frac{1}{2}}(A)$ with $\xi_j(A) = \frac{1}{n}[(A'A)(\widetilde{X'X})^{-1}(A'A)]_{j,j}$.

Remark 3.1. Equivalently we have

$$\hat{\beta}_{A,\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \|\tilde{Y}_A - A\beta\|_2^2 + 2\lambda\|\Xi_A\beta\|_1 \right\},$$

where $\tilde{Y}_A = A(\widetilde{X'X})^{-1}X'Y$.

Actually, we are going to consider three particular cases of this estimator in this work, depending on the objective of the statistician:

- **denoising objective:** the LASSO, denoted here by $\hat{\beta}_{X,\lambda}$, given by

$$\begin{aligned}\hat{\beta}_{X,\lambda} &\in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \|Y - X\beta\|_2^2 + 2\lambda \|\beta\|_1 \right\} \\ &= \arg \min_{\beta \in \mathbb{R}^p} \left\{ -2Y'X\beta + \beta'X'X\beta + 2\lambda \|\beta\|_1 \right\}\end{aligned}$$

(note that in this case, $\Xi_X = I$ since X is normalized);

- **transductive objective:** the Transductive LASSO, denoted here by $\hat{\beta}_{\sqrt{n/m}Z,\lambda}$, given by

$$\hat{\beta}_{\sqrt{n/m}Z,\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{n}{m} \left\| \tilde{Y}_Z - Z\beta \right\|_2^2 + 2\lambda \|\Xi_{\frac{n}{m}Z'}\beta\|_1 \right\};$$

- **estimation objective:** $\hat{\beta}_{\sqrt{n}I,\lambda}$, defined by

$$\hat{\beta}_{\sqrt{n}I,\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ n \left\| \tilde{Y}_I - \beta \right\|_2^2 + 2\lambda \|\Xi_{\sqrt{n}I}\beta\|_1 \right\}.$$

Let us give the analogous definition for an extension of the Dantzig Selector.

Definition 3.2. For a given parameter $\lambda > 0$ and any matrix A such that $\text{Ker}(A) = \text{Ker}(X)$, we consider the estimator given by

$$\tilde{\beta}_{A,\lambda} = \begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \\ \text{s.t. } \left\| \Xi_A^{-1} A' A ((X'X)^{-1} X'Y - \beta) \right\|_\infty \leq \lambda. \end{cases} \quad (6)$$

Here again, we are going to consider three cases, for $A = X$, $A = \sqrt{n/m}Z$ and $A = \sqrt{n}I$, and it is easy to check that for $A = X$ we have exactly the usual definition of the Dantzig Selector (Program 4). Moreover, here again, note that we can rewrite this estimator:

$$\tilde{\beta}_{A,\lambda} = \begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \\ \text{s.t. } \left\| \Xi_A^{-1} A' (\tilde{Y}_A - A\beta) \right\|_\infty \leq \lambda. \end{cases}$$

The following proposition provides an interpretation of our estimators when $A = \sqrt{n}I$.

Proposition 3.1. Let us assume that $(X'X)$ is invertible. Then $\hat{\beta}_{\sqrt{n}I,\lambda} = \tilde{\beta}_{\sqrt{n}I,\lambda}$ and this is a soft-thresholded least-square estimator: let us put $\hat{\beta}^{LSE} = (X'X)^{-1} X'Y$ then $\hat{\beta}_{\sqrt{n}I,\lambda}$ is the vector obtained by replacing the j -th coordinate $b_j = \hat{\beta}_j^{LSE}$ of $\hat{\beta}^{LSE}$ by $\text{sgn}(b_j) (|b_j| - \lambda \xi_j(nI)/n)_+$, where we use the standard notation $\text{sgn}(x) = +1$ if $x \geq 0$, $\text{sgn}(x) = -1$ if $x < 0$ and $(x)_+ = \max(x, 0)$.

Proposition 3.2 deals with a dual definition of the estimator $\hat{\beta}_{A,\lambda}$.

Proposition 3.2. *When $\text{Ker}(A) = \text{Ker}(X)$, the solutions β of the following program:*

$$\begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|A\beta\|_2^2 \\ \text{s.t. } \left\| \Xi_A^{-1} A' ((\tilde{Y}_A - A\beta)) \right\|_\infty \leq \lambda \end{cases}$$

all satisfy $X\beta = X\hat{\beta}_{A,\lambda}$ and $A\beta = A\hat{\beta}_{A,\lambda}$.

Proofs can be found in Section 7, page 16.

3.2 Theoretical results

Let us first introduce our main assumption. This assumption is stated with a given $p \times p$ matrix M and a given real number $x > 0$.

Assumption $H(M, x)$: there is a constant $c(M) > 0$ such that, for any $\alpha \in \mathbb{R}^p$ such that $\sum_{j:\beta_j^* = 0} \xi_j(M) |\alpha_j| \leq x \sum_{j:\beta_j^* \neq 0} \xi_j(M) |\alpha_j|$ we have

$$\alpha' M \alpha \geq c(M) n \sum_{j:\beta_j^* \neq 0} \alpha_j^2. \quad (7)$$

First, let us explain briefly the meaning of this hypothesis. In the case, where M is invertible, the condition

$$\alpha' M \alpha \geq c(M) n \sum_{j:\beta_j^* \neq 0} \alpha_j^2$$

is always satisfied for any $\alpha \in \mathbb{R}^p$ with $c(M)$ larger than the smallest eigenvalue of M/n . However, for the LASSO, we have $M = (X'X)$ and M cannot be invertible if $p > n$. Even in this case, Assumption $H(M, x)$ may still be satisfied. Indeed, the assumption requires that Inequality (7) holds only for a small for a small subset of \mathbb{R}^p determined by the condition $\sum_{j:\beta_j^* = 0} \xi_j(M) |\alpha_j| \leq x \sum_{j:\beta_j^* \neq 0} \xi_j(M) |\alpha_j|$. For $M = (X'X)$, this assumption becomes exactly the one taken in [BTW07]. In that paper, the necessity of such an hypothesis is also discussed.

Theorem 3.3. *Let us assume that Assumption $H(A'A, 3)$ is satisfied and that $\text{Ker}(A) = \text{Ker}(X)$. Let us choose $0 < \eta < 1$ and $\lambda = 2\sigma\sqrt{2n \log(p/\eta)}$. With probability at least $1 - \eta$ on the draw of Y , we have simultaneously*

$$\left\| A \left(\hat{\beta}_{A,\lambda} - \beta^* \right) \right\|_2^2 \leq \frac{72\sigma^2}{c(A'A)} \log\left(\frac{p}{\eta}\right) \sum_{j:\beta_j^* \neq 0} \xi_j(A),$$

and

$$\left\| \Xi_A \left(\hat{\beta}_{A,\lambda} - \beta^* \right) \right\|_1 \leq \frac{24\sqrt{2}\sigma}{c(A'A)} \left(\frac{\log(p/\eta)}{n} \right)^{\frac{1}{2}} \sum_{j:\beta_j^* \neq 0} \xi_j(A).$$

In particular, the first inequality gives

- if Assumption $H(X'X, 3)$ is satisfied, with probability at least $1 - \eta$,

$$\frac{1}{n} \left\| X \left(\hat{\beta}_{X,\lambda} - \beta^* \right) \right\|_2^2 \leq \frac{72\sigma^2}{nc(X'X)} \|\beta^*\|_0 \log \left(\frac{p}{\eta} \right);$$

- if Assumption $H(\frac{n}{m}Z'Z, 3)$ is satisfied, and if $\text{Ker}(Z) = \text{Ker}(X)$, with probability at least $1 - \eta$,

$$\frac{1}{m} \left\| Z \left(\hat{\beta}_{Z,\lambda} - \beta^* \right) \right\|_2^2 \leq \frac{72\sigma^2}{nc(\frac{n}{m}Z'Z)} \sum_{j:\beta_j^* \neq 0} \xi_j \left(\sqrt{n/m}Z \right) \log \left(\frac{p}{\eta} \right);$$

- and if $(X'X)$ is invertible, with probability at least $1 - \eta$,

$$\left\| \hat{\beta}_{\sqrt{n}I,\lambda} - \beta^* \right\|_2^2 \leq \frac{72\sigma^2}{nc(nI)} \sum_{j:\beta_j^* \neq 0} \xi_j(nI) \log \left(\frac{p}{\eta} \right).$$

This result shows that each of these three estimators satisfy at least a SI for the task it is designed for. For example, the LASSO is proved to have "good" performance for the estimation of $X\beta^*$ and the Transductive LASSO is proved to have good performance for the estimation of $Z\beta^*$. However we cannot assert that, for example, the LASSO performs better than the Transductive LASSO for the estimation of $Z\beta^*$.

Remark 3.2. For $A = X$, the particular case of our result applied to the LASSO is quite similar to the result given in [BTW07] on the LASSO. Actually, Theorem 3.3 can be seen as a generalization of the result in [BTW07] and it should be noted that the proof used to prove Theorem 3.3 uses arguments introduced in [BTW07].

Remark 3.3. As soon as $A'A$ is better determined than $X'X$, Assumption $H(A, x)$ is less restrictive than $H(X'X, x)$. In particular, in the case where $m > n$, Assumption $H((n/m)Z'Z, x)$ is expected to be less restrictive than Assumption $H(X'X, x)$.

Now we give the analogous result for the estimator $\tilde{\beta}_{A,\lambda}$.

Theorem 3.4. Let us assume that Assumption $H(A'A, 1)$ is satisfied and that $\text{Ker}(A) = \text{Ker}(X)$. Let us choose $0 < \eta < 1$ and $\lambda = 2\sigma\sqrt{2n \log(p/\eta)}$. With probability at least $1 - \eta$ on the draw of Y , we have simultaneously

$$\left\| A \left(\tilde{\beta}_{A,\lambda} - \beta^* \right) \right\|_2^2 \leq \frac{72\sigma^2}{c(A'A)} \log \left(\frac{p}{\eta} \right) \sum_{j:\beta_j^* \neq 0} \xi_j(A),$$

and

$$\left\| \Xi_A \left(\tilde{\beta}_{A,\lambda} - \beta^* \right) \right\|_1 \leq \frac{12\sqrt{2}\sigma}{c(A'A)} \left(\frac{\log(p/\eta)}{n} \right)^{\frac{1}{2}} \sum_{j:\beta_j^* \neq 0} \xi_j(A).$$

4 An extension to the general case

In this section, we only deal with the transductive setting, $A = \sqrt{n/m}Z$. Let us remind that in such a framework, we observe X which consists of some observations x_i associated to labels Y_i in Y , for $i \in \{1, \dots, n\}$. Moreover we have additional observations x_i for $i \in \{n+1, \dots, m\}$ with $m > n$. We also recall that Z contains all the x_i for $i \in \{1, \dots, m\}$ and that the objective is to estimate the corresponding labels Y_i , let us put $\tilde{Y} = (Y_1, \dots, Y_m)'$.

4.1 General remarks

Let us have look at the definition of $\hat{\beta}_{\sqrt{n/m}Z, \lambda}$, for example as given in Remark 3.1:

$$\hat{\beta}_{\sqrt{n/m}Z, \lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{n}{m} \left\| \tilde{Y}_Z - Z\beta \right\|_2^2 + 2\lambda \left\| \Xi_{\frac{n}{m}} Z' Z \beta \right\|_1 \right\},$$

where actually $\tilde{Y}_Z = Z \left(\widetilde{X'X} \right)^{-1} XY$ can be interpreted as a preliminary estimator of \tilde{Y} . Hence, in any case, we propose the following procedure.

Let us assume that, depending on the context, the user has a natural (and not necessary efficient) estimator of $\tilde{Y} = (Y_1, \dots, Y_{n+m})'$. Note this estimator \check{Y} .

Definition 4.1. *The Transductive LASSO is given by:*

$$\hat{\beta}_{\check{Y}, \sqrt{n/m}Z, \lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{n}{m} \left\| \check{Y} - Z\beta \right\|_2^2 + 2\lambda \left\| \Xi_{\frac{n}{m}} Z' Z \beta \right\|_1 \right\},$$

and the Transductive Dantzig Selector is defined as:

$$\tilde{\beta}_{\check{Y}, \sqrt{n/m}Z, \lambda} = \begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \left\| \beta \right\|_1 \\ \text{s.t.} \left\| \frac{n}{m} \Xi^{-1} \frac{1}{\sqrt{n/m}Z} Z' (\check{Y} - Z\beta) \right\|_{\infty} \leq \lambda. \end{cases}$$

In the next subsection, we propose a context where we have a natural estimator \check{Y} and give a SI on this estimator.

4.2 An example: small labeled dataset, large unlabeled dataset

The idea of this example is to consider the case where the examples x_i for $1 \leq i \leq n$ are "representative" of the large populations x_i for $1 \leq i \leq m$.

Consider, $Z = (x'_1, \dots, x'_m)'$ where the x'_i s are the points of interest: we want to estimate $\tilde{Y} = Z\beta^*$. However, we just have a very expensive and noisy procedure, that, given a point x_i , returns $Y_i = x_i\beta^* + \varepsilon_i$, where the ε_i 's are $\mathcal{N}(0, \sigma^2)$ independent random variables. In such a case, the procedure cannot be applied for the whole dataset $Z = (x'_1, \dots, x'_m)'$. We can only make a deal with a "representative" sample of size n . A typical case could be $n < p < m$.

First, let us introduce a slight modification of our main hypothesis. It is also stated with a given $p \times p$ matrix M and a given real number $x > 0$.

Assumption $H'(M, x)$: there is a $c(M) > 0$ such that, for any $\alpha \in \mathbb{R}^p$ such that $\sum_{j:\beta_j^* = 0} |\alpha_j| \leq x \sum_{j:\beta_j^* \neq 0} |\alpha_j|$ we have

$$\alpha' M \alpha \geq c(M) n \sum_{j:\beta_j^* \neq 0} \alpha_j^2.$$

We can now state our main result.

Theorem 4.1. *Let us assume that Assumption $H'((n/m)Z'Z, 1)$ is satisfied. Let us choose $0 < \eta < 1$ and $\lambda_1 = \lambda_2 = 10^{-1} \sigma \sqrt{2n \log(p/\eta)}$. Moreover, let us assume that*

$$\forall u \in \mathbb{R}^p \text{ with } \|u\|_1 \leq \|\beta^*\|_1, \quad \left\| \left((X'X) - \frac{n}{m} (Z'Z) \right) u \right\|_\infty < \frac{\sigma}{10} \sqrt{2n \log\left(\frac{p}{\eta}\right)}. \quad (8)$$

Let $\check{Y}_{\lambda_1} = Z \tilde{\beta}_{X, \lambda_1}$ be a preliminary estimator of \check{Y} , based on the Dantzig Selector given by (6) (with $A = X$). Then define the Transductive LASSO by

$$\hat{\beta}_{\frac{n}{m}Z, 20\lambda_2}^* = \begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \frac{n}{m} \|Z\beta\|_2^2 \\ \text{s.t. } \left\| \frac{n}{m} Z'(\check{Y}_{\lambda_1} - Z\beta) \right\|_\infty \leq 20\lambda_2, \end{cases}$$

and the Transductive Dantzig Selector

$$\tilde{\beta}_{\frac{n}{m}Z, \lambda_2}^* = \begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \\ \text{s.t. } \left\| \frac{n}{m} Z'(\check{Y}_{\lambda_1} - Z\beta) \right\|_\infty \leq \lambda_2. \end{cases}$$

With probability at least $1 - \eta$ on the draw of Y , we have simultaneously

$$\begin{aligned} \frac{1}{m} \left\| Z(\tilde{\beta}_{\frac{n}{m}Z, \lambda_2}^* - \beta^*) \right\|_2^2 &\leq \frac{16\sigma^2}{nc((n/m)Z'Z)} \log\left(\frac{p}{\eta}\right) \|\beta^*\|_0, \\ \left\| \tilde{\beta}_{\frac{n}{m}Z, \lambda_2}^* - \beta^* \right\|_1 &\leq \frac{8\sigma}{c((n/m)Z'Z)} \left(\frac{\log(p/\eta)}{n}\right)^{\frac{1}{2}} \|\beta^*\|_0, \end{aligned}$$

and moreover, if $H'((n/m)Z'Z, 5)$ is also satisfied,

$$\begin{aligned} \frac{1}{m} \left\| Z(\hat{\beta}_{\frac{n}{m}Z, 20\lambda_2}^* - \beta^*) \right\|_2^2 &\leq \frac{88\sigma^2}{nc((n/m)Z'Z)} \log\left(\frac{p}{\eta}\right) \|\beta^*\|_0, \\ \left\| \hat{\beta}_{\frac{n}{m}Z, 20\lambda_2}^* - \beta^* \right\|_1 &\leq \frac{54\sigma}{c((n/m)Z'Z)} \left(\frac{\log(p/\eta)}{n}\right)^{\frac{1}{2}} \|\beta^*\|_0. \end{aligned}$$

First, let us remark that the preliminary estimator \check{Y}_{λ_1} is defined using the Dantzig Selector $\hat{\beta}_{X,\lambda_1}$. We could give exactly the same kind of results using a the LASSO $\hat{\beta}_{X,\lambda_1}$ as a preliminary estimator.

Now, let us give a look at the new hypothesis, Inequality (8). We can interpret this condition as the fact that the x_i 's for $1 \leq i \leq n$ are effectively representative of the wide population: so $X'X/n$ is "not too far" from $Z'Z/m$. We will end this section by a result that proves that this is effectively the case in a typical situation.

Proposition 4.2. *Assume that $m = kn$ for an integer value $k \in \mathbb{N} \setminus \{0, 1\}$. Let us assume that X and Z are build in the following way: we have a population $\chi_1 = (\chi_{1,1}, \dots, \chi_{1,p}) \in \mathbb{R}^p, \dots, \chi_m \in \mathbb{R}^p$ (the points of interest). Then, we draw uniformly without replacement, n of the χ_i 's to be put in X : more formally, but equivalently, we draw uniformly a permutation σ of $\{1, \dots, m\}$ and we put $X = (x'_1, \dots, x'_n)' = (\chi'_{\sigma(1)}, \dots, \chi'_{\sigma(n)})'$ and $Z = (x'_1, \dots, x'_m)' = (\chi'_{\sigma(1)}, \dots, \chi'_{\sigma(m)})'$. Let us assume that for any $(i, j) \in \{1, \dots, m\} \times \{1, \dots, p\}$, $\chi_{i,j}^2 < \kappa$ for some $\kappa > 0$, and that $p \geq 2$. Then, with probability at least $1 - \eta$, for any $u \in \mathbb{R}^p$,*

$$\left\| \left(X'X - \frac{n}{m} Z'Z \right) u \right\|_{\infty} \leq \|u\|_1 \frac{2\kappa k}{k-1} \sqrt{2 \log \frac{p}{\eta}}.$$

In particular, if we have

$$\|u\|_1 \leq \|\beta^*\|_1 \text{ and } \kappa \leq \frac{k-1}{10k} \frac{\sigma}{\|\beta^*\|_1}$$

then we have

$$\left\| \left(X'X - \frac{n}{m} Z'Z \right) u \right\|_{\infty} \leq \sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)}.$$

Let us just mention that the assumption $m = kn$ is not restrictive. It has been introduced for the sake of simplicity.

5 Experimental results

Implementation. Since the paper of Tibshirani [Tib96], several effective algorithms to compute the LASSO have been proposed and studied (for instance Interior Points methods [KKL⁺07], LARS [EHJT04], Pathwise Coordinate Optimization [FHHT07], Relaxed Greedy Algorithms [HCB08]). For the Dantzig Selector, a linear method was proposed in the first paper [CT07]. The LARS algorithm was also successfully extended in [JRL09] to compute the Dantzig Selector.

Then there are many algorithms to compute $\hat{\beta}_{A,\lambda}$ and $\tilde{\beta}_{A,\lambda}$, when $A = X$. Thanks to Proposition 3.1, it is also clear that we can easily find an efficient algorithm for the case $A = \sqrt{n}I$.

The general form of the estimators $\hat{\beta}_{A,\lambda}$ and $\tilde{\beta}_{A,\lambda}$ given by Definitions 3.1

and 3.2, allows to use one of the algorithms mentioned previously to compute our estimator in two cases. For example, from Remark 3.1, we have:

$$\hat{\beta}_{A,\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \left\| \tilde{Y}_A - A\beta \right\|_2^2 + 2\lambda \|\Xi_A \beta\|_1 \right\},$$

then we just have to compute \tilde{Y}_A , to put $B = A\Xi_A^{-1}$, to use any program that computes the LASSO to determine

$$\hat{\gamma} \in \arg \min_{\gamma \in \mathbb{R}^p} \left\{ \left\| \tilde{Y}_A - B\gamma \right\|_2^2 + 2\lambda \|\gamma\|_1 \right\}$$

and then to put $\hat{\beta}_{A,\lambda} = \Xi_A^{-1} \hat{\gamma}$.

In the rest of this section, we compare the LASSO and the transductive LASSO on the classical toy example introduced by Tibshirani [Tib96] and used as a benchmark.

Data description. In the model proposed by Tibshirani, we have

$$Y_i = x_i \beta^* + \varepsilon_i$$

for $i \in \{1, \dots, n\}$, $\beta^* \in \mathbb{R}^p$ and the ε_i are i.i.d. $\mathcal{N}(0, \sigma^2)$. Finally, the $(x_i)_{i \in \{1, \dots, m\}}$ are generated from a probability distribution: they are independent and identically distributed

$$x_i \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho & \dots & \dots & \rho^{p-1} \\ \rho & 1 & \rho & \dots & \rho^{p-2} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \rho^{p-2} & \dots & \rho & 1 & \rho \\ \rho^{p-1} & \dots & \dots & \rho & 1 \end{pmatrix} \right),$$

for a given $\rho \in]-1, 1[$.

As in [Tib96], we set $p = 8$. In a first experiment, we take $(n, m) = (7, 10)$, $\rho = 0.5$, $\sigma = 1$ and $\beta^* = (3, 1.5, 0, 0, 2, 0, 0, 0)$ ("sparse"). Then, in order to check the robustness of the results, we consider successively $\rho = 0.5$ by $\rho = 0.9$ (correlated variables), $\sigma = 1$ by $\sigma = 3$ (noisy case), $\beta^* = (3, 1.5, 0, 0, 2, 0, 0, 0)$ by $\beta^* = (5, 0, 0, 0, 0, 0, 0, 0)$ ("very sparse" case), $(n, m) = (7, 10)$ by $(n, m) = (7, 20)$ (larger unlabeled set), $(n, m) = (20, 30)$ ($p < n$, easy case) and finally $(n, m) = (20, 120)$.

We use the version of the Transductive LASSO proposed in Section 4: for a given λ_1 , we first compute the LASSO estimator $\hat{\beta}_{X,\lambda_1}$. In the sequel, the Transductive LASSO is given by

$$\hat{\beta}^{TL}(\lambda_1, \lambda_2) = \begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \frac{n}{m} \|Z\beta\|_2^2 \\ s.t. \left\| \frac{n}{m} Z'(Z\hat{\beta}_{X,\lambda_1} - Z\beta) \right\|_\infty \leq \lambda_2, \end{cases}$$

for a given λ_2 . We compare this two step procedure with the procedure obtained using the usual LASSO only: $\hat{\beta}^L(\lambda) = \hat{\beta}_{X,\lambda}$ for a given λ that may differ from λ_1 . In both cases, the solutions are computed using PCO algorithm. We compute $\hat{\beta}^L(\lambda)$ and $\hat{\beta}^{TL}(\lambda_1, \lambda_2)$ for $(\lambda, \lambda_1, \lambda_2) \in \Lambda^3$ where $\Lambda^3 = \{1.2^k, k = -50, -49, \dots, 30\}$. In the next subsection, we examine the performance of each estimator according to the value of the regularization parameters.

Results. We illustrate here some of the results obtained in the considered cases.

Case $(n, m) = (7, 10)$, $\rho = 0.5$, $\sigma = 1$ and β^ "sparse":*

We simulated 100 experiments and studied the distribution of

$$PERF(X) = \frac{\min_{(\lambda_1, \lambda_2) \in \Lambda^2} \|X(\hat{\beta}^{TL}(\lambda_1, \lambda_2) - \beta^*)\|_2^2}{\min_{\lambda \in \Lambda} \|X(\hat{\beta}^L(\lambda) - \beta^*)\|_2^2},$$

$$PERF(Z) = \frac{\min_{(\lambda_1, \lambda_2) \in \Lambda^2} \|Z(\hat{\beta}^{TL}(\lambda_1, \lambda_2) - \beta^*)\|_2^2}{\min_{\lambda \in \Lambda} \|Z(\hat{\beta}^L(\lambda) - \beta^*)\|_2^2},$$

and

$$PERF(I) = \frac{\min_{(\lambda_1, \lambda_2) \in \Lambda^2} \|\hat{\beta}^{TL}(\lambda_1, \lambda_2) - \beta^*\|_2^2}{\min_{\lambda \in \Lambda} \|\hat{\beta}^L(\lambda) - \beta^*\|_2^2},$$

over all the experiments.

For example, we plot (Figure 1) the histogram of $PERF(X)$ (actually, the three distributions were quite similar). We observe that in 50% of the simulations, $\min_{(\lambda_1, \lambda_2) \in \Lambda^2} \|X(\hat{\beta}^{TL}(\lambda_1, \lambda_2) - \beta^*)\|_2^2 = \min_{(\lambda_1, 0) \in \Lambda^2} \|X(\hat{\beta}^{TL}(\lambda_1, 0) - \beta^*)\|_2^2 = \min_{\lambda \in \Lambda} \|X(\hat{\beta}^L(\lambda) - \beta^*)\|_2^2$. In these cases, the Transductive LASSO does not improve at all the LASSO. But in the other 50%, the Transductive LASSO actually improves the LASSO, and the improvement is sometimes really important. We give an overview of the results in Table 1.

The other cases :

The following conclusions emerge from the experiments: first, $\beta^* = (5, 0, \dots, 0)$ leads to a more significant improvement of the Transductive LASSO compared to the LASSO (Table 1). This good performance of the Transductive LASSO can also be observed when $(n, m) = (7, 10)$ and $(n, m) = (7, 20)$. However in the case $n > p$ (easy case), i.e., $(n, m) = (20, 30)$ and $(n, m) = (20, 120)$, the improvement of the Transductive LASSO with respect to the LASSO becomes less significant (Table 1).

Finally, ρ and σ have of course a significant influence on the performance of the LASSO. However these parameters do not seem to have any influence on the relative performance of the Transductive LASSO with respect to the LASSO (see for instance the three last rows in Table 1, where we kept $(n, m) = (20, 30)$). Quite surprisingly, the relative performance of both estimators does not strongly depend on the estimation objective β^* , $X\beta^*$ or $Z\beta^*$, but on the particular experiment we deal with. According to the realized study and for all the objectives,

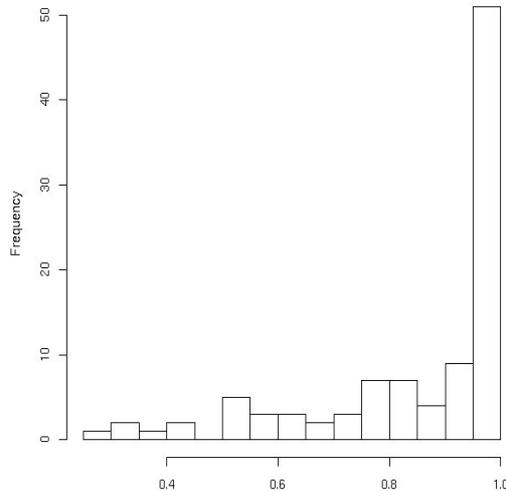


Figure 1: Histogram of $PERF(X)$ with $(n, m) = (7, 10)$, $\rho = 0.5$, $\sigma = 1$ and $\beta^* = (3, 1.5, 0, 0, 2, 0, 0, 0)$.

the Transductive LASSO performs better than the LASSO in about 50% of the experiments. Otherwise, $\lambda_1 = 0$ is the optimal tuning parameter and then, the LASSO and the Transductive LASSO are equivalent.

Also surprising is that as often as not, the minimum in

$$\min_{(\lambda_1, \lambda_2) \in \Lambda^2} \|X(\hat{\beta}^{TL}(\lambda_1, \lambda_2) - \beta^*)\|_2^2 < \min_{(\lambda_1, 0) \in \Lambda^2} \|X(\hat{\beta}^{TL}(\lambda_1, 0) - \beta^*)\|_2^2,$$

does not significantly depend on λ_1 for a very large range of values λ_1 . This is quite interesting for a practitioner as it means that when we use the Transductive LASSO, we deal with only a singular unknown tuning parameter (that is λ_2) and not two.

Discussion on the regularization parameter. Finally, we would like to point out the importance of the tuning parameter λ (in a general term). Figure 2 illustrates a graph of a typical experiment. There are two curves on this graph, that represent the quantities $(1/n)\|X(\hat{\beta}^L(\lambda) - \beta^*)\|_2^2$ and $(1/m)\|Z(\hat{\beta}^L(\lambda) - \beta^*)\|_2^2$ with respect to λ . We observe that both functions do not reach their minimum value for the same value of λ (the minimum is highlighted on the graph by a dot), even if these minimum are quite close.

Since we consider variable selection methods, the identification of the true support $\{j : \beta_j^* \neq 0\}$ of the vector β^* is also in concern. One expects that the estimator $\hat{\beta}$ and the true vector β^* share the same support at least when n is large enough. This is known as the variable selection consistency problem and it has been considered for the LASSO estimator in several works (see

Table 1: Evaluation of the mean ME and the quantile Q_3 of order 0.3 of $PERF(I)$, $PERF(X)$ and $PERF(Z)$. In these experiments, σ always equals 1. The case *sparse* corresponds to $\beta^* = (3, 1.5, 0, 0, 2, 0, 0, 0)$ while the case *very sparse* corresponds to $\beta^* = (5, 0, 0, 0, 0, 0, 0, 0)$.

β^*	(n, m)	ρ	σ	$PERF(I)$		$PERF(X)$		$PERF(Z)$	
				ME	Q_3	ME	Q_3	ME	Q_3
VERY SPARSE	(7, 10)	0.5	1	0.74	0.71	0.76	0.71	0.75	0.70
SPARSE	(7, 10)	0.5	1	0.83	0.76	0.86	0.80	0.88	0.88
SPARSE	(7, 20)	0.5	1	0.84	0.79	0.84	0.81	0.88	0.89
SPARSE	(20, 30)	0.5	1	0.91	0.90	0.93	0.93	0.93	0.95
SPARSE	(20, 30)	0.9	1	0.91	0.93	0.94	0.95	0.93	0.96
SPARSE	(20, 30)	0.5	3	0.90	0.89	0.92	0.92	0.92	0.93

[Bun08, MB06, MY09, Wai06, ZY06]). Recently, [Lou08] provided the variable selection consistency of the Dantzig Selector. Other popular selection procedures, based on the LASSO estimator, such as the Adaptive LASSO [Zou06], the SCAD [FL01], the S-LASSO [Heb08] and the Group-LASSO [Bac08], have also been studied under a variable selection point of view. Following our previous work [AH08], it is possible to provide such results for the Transductive LASSO.

The variable selection task has also been illustrated in Figure 2. We reported the minimal value of λ for which the LASSO estimator identifies correctly the non zero components of β^* . This value of λ is quite different from the values that minimizes the prediction losses. This observation is recurrent in almost all the experiments: the estimation $X\beta^*$, $Z\beta^*$ and the support of β^* are three different objectives and have to be treated separately. We cannot expect in general to find a choice for λ which makes the LASSO, for instance, has good performance for all the mentioned objective simultaneously.

6 Conclusion

In this paper, we propose an extension of the LASSO and the Dantzig Selector for which we provide theoretical results with less restrictive hypothesis than in previous works. These estimators have a nice interpretation in terms of transductive prediction. Moreover, we study the practical performance of the proposed transductive estimators on simulated data. It turns out that the benefit using such methods is emphasized when the model is sparse and particularly when the samples sizes (n labeled points and m unlabeled points) and dimension p are such that $n < p < m$.

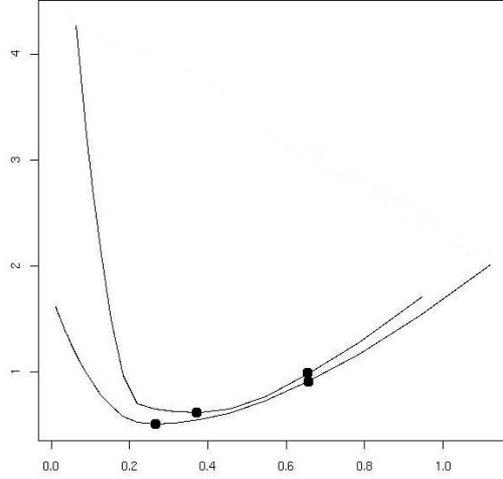


Figure 2: Performance vs. λ .

7 Proofs

In this section, we state the proofs of our main results.

7.1 Proof of Propositions 3.1 and 3.2

Proof of Proposition 3.1. Let us assume that $(X'X)$ is invertible. Then just remark that the criterion minimized by $\hat{\beta}_{\sqrt{n}I, \lambda}$ is just

$$n \left\| \hat{\beta}^{LSE} - \beta \right\|_2^2 + 2\lambda \|\Xi_{nI} \beta\|_1 = \sum_{j=1}^p \left\{ \left[\hat{\beta}_j^{LSE} - \beta_j \right]^2 + \frac{2\lambda \xi_j(\sqrt{n}I)}{n} |\beta_j| \right\}.$$

So we can optimize with respect to each coordinate β_j individually. It is quite easy to check that the solution is, for β_j ,

$$\text{sgn} \left(\hat{\beta}_j^{LSE} \right) \left(\left| \hat{\beta}_j^{LSE} \right| - \frac{\lambda \xi_j(\sqrt{n}I)}{n} \right)_+.$$

The proof for $\hat{\beta}_{\sqrt{n}I, \lambda}$ is also easy as it solves

$$\begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|\beta\|_1 \\ \text{s.t. } \left\| n \Xi_{nI}^{-1} (\hat{\beta}^{LSE} - \beta) \right\|_\infty \leq \lambda. \end{cases}$$

□

Proof of Proposition 3.2. Let us write the Lagrangian of the program

$$\begin{cases} \arg \min_{\beta \in \mathbb{R}^p} \|A\beta\|_2^2 \\ \text{s.t. } \left\| \Xi_A^{-1}(A'A)((\widetilde{X}'\widetilde{X})^{-1}X'Y - \beta) \right\|_\infty \leq \lambda, \end{cases}$$

$$\begin{aligned} \mathcal{L}(\beta, \gamma, \mu) = & \beta(Z'Z)\beta + \gamma' \left[\Xi_A^{-1}(A'A)((\widetilde{X}'\widetilde{X})^{-1}X'Y - \beta) - \lambda E \right] \\ & + \mu' \left[\Xi_A^{-1}(A'A)(\beta - (\widetilde{X}'\widetilde{X})^{-1}X'Y) - \lambda E \right] \end{aligned}$$

with $E = (1, \dots, 1)'$, and for any j , $\gamma_j \geq 0$, $\mu_j \geq 0$ and $\gamma_j \mu_j = 0$. Any solution $\underline{\beta} = \underline{\beta}(\gamma, \mu)$ must satisfy

$$0 = \frac{\partial \mathcal{L}}{\partial \underline{\beta}}(\underline{\beta}, \lambda, \mu) = 2\underline{\beta}(A'A) + (\gamma - \mu)\Xi_A^{-1}(A'A)$$

so

$$(A'A)\underline{\beta} = (A'A)\Xi_A^{-1} \frac{\mu - \gamma}{2}.$$

Note that the conditions $\gamma_j \geq 0$, $\mu_j \geq 0$ and $\gamma_j \mu_j = 0$ means that there is a $\zeta_j \in \mathbb{R}$ such that $\zeta_j = \xi_j^{\frac{1}{2}}(A)(\mu_j - \gamma_j)/2$, $|\zeta_j| = \xi_j^{\frac{1}{2}}(A)(\gamma_j + \mu_j)/2$, and so $\gamma_j = 2(\zeta_j/\xi_j^{\frac{1}{2}}(A))_-$ and $\mu_j = 2(\zeta_j/\xi_j^{\frac{1}{2}}(A))_+$, where $(a)_+ = \max(a; 0)$ and $(a)_- = \max(-a; 0)$. Let also ζ denote the vector which j -th component is exactly ζ_j , we obtain

$$(A'A)\underline{\beta} = (A'A)\zeta,$$

or, using the condition $\text{Ker}(A) = \text{Ker}(X)$, $X\underline{\beta} = X\zeta$ and $A\underline{\beta} = A\zeta$. This leads to

$$\mathcal{L}(\underline{\beta}, \gamma, \mu) = -2Y'X(\widetilde{X}'\widetilde{X})^{-1}(A'A)\zeta + \zeta'(A'A)\zeta + 2\lambda\|\Xi_A\zeta\|_1,$$

and note that the first order condition also implies that γ and μ (and so ζ) maximize \mathcal{L} . This ends the proof. \square

7.2 A useful Lemma

The following lemma will be used in the proofs of Theorems 3.3 and 3.4.

Lemma 7.1. *Let us put $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)'$. If $\text{Ker}(A) = \text{Ker}(X)$ we have, with probability at least $1 - \eta$,*

$$\forall j \in \{1, \dots, p\}, \left| \left[A'A(\widetilde{X}'\widetilde{X})^{-1}X'\varepsilon \right]_j \right| \leq \xi_j(A)\sigma \sqrt{2n \log \frac{p}{\eta}},$$

or, in other words,

$$\left\| \Xi_A^{-1}(A'A)((\widetilde{X}'\widetilde{X})^{-1}X'Y - \beta^*) \right\|_\infty \leq \sigma \sqrt{2n \log \frac{p}{\eta}}.$$

Proof of the lemma. By definition, $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$ and so

$$(A'A)(\widetilde{X'X})^{-1}X'\varepsilon \sim \mathcal{N}(0, \sigma^2(A'A)(\widetilde{X'X})^{-1}(A'A)).$$

So, for all j , $[(A'A)(\widetilde{X'X})^{-1}X'\varepsilon]_j$ comes from a $\mathcal{N}(0, \sigma^2 \xi_j^2(A))$ distribution. This implies the first point, the second one is trivial using $Y = X\beta^* + \varepsilon$. \square

7.3 Proof of Theorems 3.3 and 3.4

Proof of Theorem 3.3. By definition of $\hat{\beta}_{A,\lambda}$ we have

$$\begin{aligned} & -2Y'X(\widetilde{X'X})^{-1}(A'A)\hat{\beta}_{A,\lambda} + \left(\hat{\beta}_{A,\lambda}\right)'(A'A)\hat{\beta}_{A,\lambda} + 2\lambda\|\Xi_{A'}\hat{\beta}_{A,\lambda}\|_1 \\ & \leq 2Y'X(\widetilde{X'X})^{-1}(A'A)\beta^* + (\beta^*)'(A'A)\beta^* + 2\lambda\|\Xi_A\beta^*\|_1. \end{aligned}$$

Since $Y = X\beta^* + \varepsilon$, we obtain

$$\begin{aligned} & 2(\beta^*)'X'X(\widetilde{X'X})^{-1}(A'A)\left(\beta^* - \hat{\beta}_{A,\lambda}\right) + \left(\hat{\beta}_{A,\lambda}\right)'(A'A)\hat{\beta}_{A,\lambda} - (\beta^*)'(A'A)\beta^* \\ & + 2\varepsilon'X(\widetilde{X'X})^{-1}(A'A)\left(\beta^* - \hat{\beta}_{A,\lambda}\right) \leq 2\lambda\|\Xi_A\beta^*\|_1 - 2\lambda\|\Xi_A\hat{\beta}_{A,\lambda}\|_1. \end{aligned}$$

Now, if $\text{Ker}(X) = \text{Ker}(A)$ then we have $X'X(\widetilde{X'X})^{-1}(A'A) = (A'A)$ and then the previous inequality leads to

$$\begin{aligned} & \left(\beta^* - \hat{\beta}_{A,\lambda}\right)'(A'A)\left(\beta^* - \hat{\beta}_{A,\lambda}\right) \\ & \leq 2\varepsilon'X(\widetilde{X'X})^{-1}(A'A)\left(\hat{\beta}_{A,\lambda} - \beta^*\right) + 2\lambda\|\Xi_A\beta^*\|_1 - 2\lambda\|\Xi_A\hat{\beta}_{A,\lambda}\|_1. \quad (9) \end{aligned}$$

Now we have to work on the term $2\varepsilon'X(\widetilde{X'X})^{-1}(A'A)\left(\hat{\beta}_{A,\lambda} - \beta^*\right)$. Note that

$$\begin{aligned} 2\varepsilon'X(\widetilde{X'X})^{-1}(A'A)\left(\hat{\beta}_{A,\lambda} - \beta^*\right) & = 2\sum_{j=1}^p\left(\hat{\beta}_{A,\lambda} - \beta^*\right)_j\left[(A'A)(\widetilde{X'X})^{-1}X'\varepsilon\right]_j \\ & \leq 2\sum_{j=1}^p\left|\left(\hat{\beta}_{A,\lambda} - \beta^*\right)_j\right|\left|\left[(A'A)(\widetilde{X'X})^{-1}X'\varepsilon\right]_j\right| \\ & \leq 2\sigma\sqrt{2n\log\left(\frac{p}{\eta}\right)}\sum_{j=1}^p\xi_j^{\frac{1}{2}}(A)\left|\left(\hat{\beta}_{A,\lambda}\right)_j - \beta_j^*\right| \end{aligned}$$

with probability at least $1 - \eta$, by Lemma 7.1. We plug this result into Inequality (9) (and replace λ by its value $2\sigma\sqrt{2n\log(p/\eta)}$) to obtain

$$\left(\beta^* - \hat{\beta}_{A,\lambda}\right)'(A'A)\left(\beta^* - \hat{\beta}_{A,\lambda}\right)$$

$$\leq 2\sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)} \sum_{j=1}^p \xi_j^{\frac{1}{2}}(A) \left\{ \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right| + 2 \left(|\beta_j^*| - \left| \left(\hat{\beta}_{A,\lambda} \right)_j \right| \right) \right\}$$

and then

$$\begin{aligned} & \left(\beta^* - \hat{\beta}_{A,\lambda} \right)' (A' A) \left(\beta^* - \hat{\beta}_{A,\lambda} \right) \\ & + 2\sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)} \sum_{j=1}^p \xi_j^{\frac{1}{2}}(A) \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right| \\ & \leq 4\sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)} \sum_{j=1}^p \xi_j^{\frac{1}{2}}(A) \left\{ \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right| + |\beta_j^*| - \left| \left(\hat{\beta}_{A,\lambda} \right)_j \right| \right\} \\ & = 4\sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)} \sum_{j:\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) \left\{ \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right| + |\beta_j^*| - \left| \left(\hat{\beta}_{A,\lambda} \right)_j \right| \right\} \\ & \leq 8\sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)} \sum_{j:\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right|. \end{aligned} \quad (10)$$

This implies, in particular, that $\beta^* - \hat{\beta}_{A,\lambda}$ is an admissible vector α in Assumption $H(A' A, 3)$ because

$$\sum_{j=1}^p \xi_j^{\frac{1}{2}}(A) \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right| \leq 4 \sum_{j:\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right|.$$

On the other hand, thanks to Inequality (10), we have

$$\begin{aligned} & \left(\beta^* - \hat{\beta}_{A,\lambda} \right)' (A' A) \left(\beta^* - \hat{\beta}_{A,\lambda} \right) \\ & \leq 6\sigma \sqrt{2n \log \left(\frac{p}{\eta} \right)} \sum_{j:\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) \left| \left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right| \\ & \leq 6\sigma \sqrt{2n \sum_{j:\beta_j^* \neq 0} \left[\left(\hat{\beta}_{A,\lambda} \right)_j - \beta_j^* \right]^2 \sum_{j:\beta_j^* \neq 0} \xi_j(A) \log \left(\frac{p}{\eta} \right)} \\ & \leq 6\sigma \sqrt{\frac{2}{c(A' A)} \left(\beta^* - \hat{\beta}_{A,\lambda} \right)' (A' A) \left(\beta^* - \hat{\beta}_{A,\lambda} \right) \sum_{j:\beta_j^* \neq 0} \xi_j(M) \log \left(\frac{p}{\eta} \right)}, \end{aligned} \quad (11)$$

where we used Assumption $H(A' A, 3)$ for the last inequality. Then

$$\left(\beta^* - \hat{\beta}_{A,\lambda} \right)' (A' A) \left(\beta^* - \hat{\beta}_{A,\lambda} \right) \leq 72 \frac{\sigma^2}{c(A' A)} \log \left(\frac{p}{\eta} \right) \sum_{j:\beta_j^* \neq 0} \xi_j(A). \quad (12)$$

A similar reasoning as in (11) leads to

$$\begin{aligned} & 2\sigma \sqrt{2n \log\left(\frac{p}{\eta}\right)} \sum_{j=1}^p \xi_j^{\frac{1}{2}}(A) \left| (\hat{\beta}_{A,\lambda})_j - \beta_j^* \right| \\ & \leq 8\sigma \sqrt{\frac{2}{c(A'A)} (\beta^* - \hat{\beta}_{A,\lambda})' (A'A) (\beta^* - \hat{\beta}_{A,\lambda}) \sum_{j:\beta_j^* \neq 0} \xi_j(M) \log\left(\frac{p}{\eta}\right)}. \end{aligned}$$

Finally, combine this last inequality with (12) to obtain the desired bound for $\|\Xi_A(\beta^* - \hat{\beta}_{A,\lambda})\|_1$. This ends the proof. \square

Proof of Theorem 3.4. We have

$$\begin{aligned} (\tilde{\beta}_{A,\lambda} - \beta^*)' (A'A) (\tilde{\beta}_{A,\lambda} - \beta^*) &= [\Xi_A(\tilde{\beta}_{A,\lambda} - \beta^*)]' \Xi_A^{-1} (A'A) (\tilde{\beta}_{A,\lambda} - \beta^*) \\ &\leq \|\Xi_A(\tilde{\beta}_{A,\lambda} - \beta^*)\|_1 \|\Xi_A^{-1} (A'A) (\tilde{\beta}_{A,\lambda} - \beta^*)\|_\infty \\ &\leq \|\Xi_A(\tilde{\beta}_{A,\lambda} - \beta^*)\|_1 \left\{ \|\Xi_A^{-1} (A'A) ((\widetilde{X}'\widetilde{X})^{-1} X'Y - \beta^*)\|_\infty \right. \\ &\quad \left. + \|\Xi_A^{-1} (A'A) ((\widetilde{X}'\widetilde{X})^{-1} X'Y - \tilde{\beta}_{A,\lambda})\|_\infty \right\}, \quad (13) \end{aligned}$$

by the constraint in the definition on $\tilde{\beta}_{A,\lambda}$ we have

$$\|\Xi_A^{-1} (A'A) ((\widetilde{X}'\widetilde{X})^{-1} X'Y - \tilde{\beta}_{A,\lambda})\|_\infty \leq \lambda,$$

while Lemma 7.1 implies that for $\lambda = 2\sigma \sqrt{2n \log(p/\eta)}$ we have

$$\|\Xi_A^{-1} (A'A) ((\widetilde{X}'\widetilde{X})^{-1} X'Y - \beta^*)\|_\infty \leq \frac{\lambda}{2},$$

with probability at least $1 - \eta$; and so:

$$(\tilde{\beta}_{A,\lambda} - \beta^*)' (A'A) (\tilde{\beta}_{A,\lambda} - \beta^*) \leq \frac{3\lambda}{2} \|\Xi_A(\tilde{\beta}_{A,\lambda} - \beta^*)\|_1.$$

Moreover note that, by definition,

$$\begin{aligned} 0 &\leq \|\Xi_A \beta^*\|_1 - \|\Xi_A \tilde{\beta}_{A,\lambda}\|_1 \\ &= \sum_{\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) |\beta_j^*| - \sum_{\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) |(\tilde{\beta}_{A,\lambda})_j| - \sum_{\beta_j^* = 0} \xi_j^{\frac{1}{2}}(A) |(\tilde{\beta}_{A,\lambda})_j| \\ &\leq \sum_{\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) |\beta_j^* - (\tilde{\beta}_{A,\lambda})_j| - \sum_{\beta_j^* = 0} \xi_j^{\frac{1}{2}}(A) |\beta_j^* - (\tilde{\beta}_{A,\lambda})_j|, \end{aligned}$$

this implies that $\beta^* - (\tilde{\beta}_{A,\lambda})$ is an admissible vector in the relation that defines Assumption $H(A'A, 1)$. Let us combine this result with Inequality (13), we obtain

$$\begin{aligned}
(\tilde{\beta}_{A,\lambda} - \beta^*)'(A'A)(\tilde{\beta}_{A,\lambda} - \beta^*) &\leq \frac{3\lambda}{2} \|\Xi_A(\beta^* - \tilde{\beta}_{A,\lambda})\|_1 \\
&\leq 3\lambda \sum_{\beta_j^* \neq 0} \xi_j^{\frac{1}{2}}(A) \left| \beta_j^* - (\tilde{\beta}_{A,\lambda})_j \right| \\
&\leq 3\lambda \sqrt{\left(\sum_{\beta_j^* \neq 0} \xi_j(A) \right) \left(\sum_{\beta_j^* \neq 0} \left| \beta_j^* - (\tilde{\beta}_{A,\lambda})_j \right|^2 \right)} \\
&\leq 3\lambda \left(\sum_{\beta_j^* \neq 0} \xi_j(A) \right)^{\frac{1}{2}} \sqrt{\frac{1}{nc(A'A)} (\tilde{\beta}_{A,\lambda} - \beta^*)'(A'A)(\tilde{\beta}_{A,\lambda} - \beta^*)}. \quad (14)
\end{aligned}$$

So we have,

$$(\tilde{\beta}_{A,\lambda} - \beta^*)'(A'A)(\tilde{\beta}_{A,\lambda} - \beta^*) \leq 9\lambda^2 \frac{1}{nc(A'A)} \sum_{\beta_j^* \neq 0} \xi_j(A),$$

and as a consequence, Inequality (14) gives the upper bound on $\|\Xi_A(\tilde{\beta}_{A,\lambda} - \beta^*)\|_1$, and this ends the proof. \square

7.4 Proof of Theorem 4.1

Proof of Theorem 4.1. The proof is almost the same as in the previous case. For the sake of simplicity, let us write $\tilde{\beta}^*$ instead of $\tilde{\beta}^*_{\sqrt{n/m}Z, \lambda_2}$ and the same for $\hat{\beta}^*$. We first give a look at the Dantzig Selector:

$$\begin{aligned}
\frac{n}{m} (\tilde{\beta}^* - \beta^*)' Z' Z (\tilde{\beta}^* - \beta^*) &\leq \|\tilde{\beta}^* - \beta^*\|_1 \left\| \frac{n}{m} Z' Z (\tilde{\beta}^* - \beta^*) \right\|_\infty \\
&\leq \|\tilde{\beta}^* - \beta^*\|_1 \left\{ \left\| \frac{n}{m} Z' (Z\tilde{\beta}^* - \check{Y}_{\lambda_1}) \right\|_\infty + \left\| \frac{n}{m} Z' (Z\beta^* - \check{Y}_{\lambda_1}) \right\|_\infty \right\} \\
&\leq \|\tilde{\beta}^* - \beta^*\|_1 \left\{ \left\| \frac{n}{m} Z' (Z\tilde{\beta}^* - \check{Y}_{\lambda_1}) \right\|_\infty + \|X'(X\beta^* - Y)\|_\infty \right. \\
&\quad \left. + \|X'(X\tilde{\beta}_{X,\lambda_1} - Y)\|_\infty + \left\| \left(\frac{n}{m} Z' Z - X'X \right) (\beta^* - \tilde{\beta}_{X,\lambda_1}) \right\|_\infty \right\}. \quad (15)
\end{aligned}$$

By Lemma 7.1, for $\lambda_1 = 10^{-1}\sigma\sqrt{2n\log(p/\eta)}$ we have

$$\|X'Y - X'X\beta^*\|_\infty \leq 10\lambda_1,$$

with probability at least $1 - \eta$. On the other hand, we have

$$\|\beta^* - \tilde{\beta}_{X,\lambda_1}\|_1 \leq \|\beta^*\|_1 + \|\tilde{\beta}_{X,\lambda_1}\|_1 \leq 2\|\beta^*\|_1,$$

by definition of the Dantzig Selector. Then, let $u = (\beta^* - \tilde{\beta}_{X,\lambda_1})/2$ and use Inequality (8) for this specific u . This ensures that

$$\left\| \left(\frac{n}{m} Z'Z - X'X \right) (\beta^* - \tilde{\beta}_{X,\lambda_1}) \right\|_\infty \leq 2\lambda_1. \quad (16)$$

The definition of the Dantzig Selector also implies that

$$\left\| X' (X\tilde{\beta}_{X,\lambda_1} - Y) \right\|_\infty \leq \lambda_1,$$

and finally the definition of the estimator leads to

$$\left\| \frac{n}{m} Z' (Z\tilde{\beta}^* - \check{Y}_{\lambda_1}) \right\|_\infty \leq \lambda_2 = \lambda_1,$$

and as a consequence, Inequality (15) becomes

$$\frac{n}{m} (\tilde{\beta}^* - \beta^*)' Z'Z (\tilde{\beta}^* - \beta^*) \leq 14\lambda_1 \|\tilde{\beta}^* - \beta^*\|_1.$$

Using the fact that $\|\tilde{\beta}^*\|_1 \leq \|\beta^*\|_1$ gives

$$\begin{aligned} \frac{n}{m} (\tilde{\beta}^* - \beta^*)' Z'Z (\tilde{\beta}^* - \beta^*) &\leq 14\lambda_1 \|\tilde{\beta}^* - \beta^*\|_1 \leq 28\lambda_1 \sum_{\beta_j^* \neq 0} |\beta_j^* - (\tilde{\beta}^*)_j| \\ &\leq 28\lambda_1 \sqrt{|\{j : \beta_j^* \neq 0\}| \left(\sum_{\beta_j^* \neq 0} |\beta_j^* - (\tilde{\beta}^*)_j|^2 \right)} \\ &\leq 28\lambda_1 |\{j : \beta_j^* \neq 0\}|^{\frac{1}{2}} \sqrt{\frac{1}{nc(n/m(Z'Z))} \frac{n}{m} (\tilde{\beta}^* - \beta^*)' Z'Z (\tilde{\beta}^* - \beta^*)}. \quad (17) \end{aligned}$$

To establish the last inequality, we used Assumption $H'((n/m)Z'Z, 1)$. Then we have,

$$\frac{n}{m} (\tilde{\beta}^* - \beta^*)' Z'Z (\tilde{\beta}^* - \beta^*) \leq 28^2 \lambda_1^2 |\{j : \beta_j^* \neq 0\}| \frac{1}{nc(n/m(Z'Z))},$$

This inequality, combined with (17), end the proof for the Dantzig Selector.

Now, let us deal with the LASSO case. The dual form of the definition of the estimator leads to

$$\begin{aligned} -2 \frac{n}{m} \check{Y}_{\lambda_1} Z \hat{\beta}^* + \frac{n}{m} (\hat{\beta}^*)' Z'Z \hat{\beta}^* + 40\lambda_2 \|\hat{\beta}^*\|_1 \\ \leq -2 \frac{n}{m} \check{Y}_{\lambda_1} Z \beta^* + \frac{n}{m} (\beta^*)' Z'Z \beta^* + 40\lambda_2 \|\beta^*\|_1 \end{aligned}$$

and so

$$\begin{aligned} & -2\frac{n}{m}\tilde{\beta}_{X,\lambda_1}Z'Z\hat{\beta}^* + \frac{n}{m}(\hat{\beta}^*)'Z'Z\hat{\beta}^* + 40\lambda_2\|\hat{\beta}^*\|_1 \\ & \leq -2\frac{n}{m}\tilde{\beta}_{X,\lambda_1}Z'Z\beta^* + \frac{n}{m}(\beta^*)'Z'Z\beta^* + 40\lambda_2\|\beta^*\|_1. \end{aligned}$$

As a consequence,

$$\begin{aligned} & \frac{n}{m}(\hat{\beta}^* - \beta^*)'Z'Z(\hat{\beta}^* - \beta^*) \\ & \leq 2\frac{n}{m}(\hat{\beta}^* - \beta^*)'Z'Z(\tilde{\beta}_{X,\lambda_1} - \beta^*) + 40\lambda_2(\|\beta^*\|_1 - \|\hat{\beta}^*\|_1). \end{aligned}$$

Now, we try to upper bound $(\hat{\beta}^* - \beta^*)'Z'Z(\tilde{\beta}_{X,\lambda_1} - \beta^*)$. We remark that

$$\begin{aligned} & \frac{n}{m}(\hat{\beta}^* - \beta^*)'Z'Z(\tilde{\beta}_{X,\lambda_1} - \beta^*) \leq \|\hat{\beta}^* - \beta^*\|_1 \left\| \frac{n}{m}(Z'Z)(\tilde{\beta}_{X,\lambda_1} - \beta^*) \right\|_\infty \\ & \leq \|\hat{\beta}^* - \beta^*\|_1 \left[\left\| \left(\frac{n}{m}Z'Z - X'X \right) (\tilde{\beta}_{X,\lambda_1} - \beta^*) \right\|_\infty \right. \\ & \quad \left. + \left\| X'X(\tilde{\beta}_{X,\lambda_1} - \beta^*) \right\|_\infty \right] \leq 13\lambda_1 \|\hat{\beta}^* - \beta^*\|_1, \end{aligned}$$

where we used (16) and the fact that

$$\left\| X'X(\tilde{\beta}_{X,\lambda_1} - \beta^*) \right\|_\infty \leq \left\| X'(X\tilde{\beta}_{X,\lambda_1} - Y) \right\|_\infty + \|X'\varepsilon\|_\infty \leq \lambda_1 + 10\lambda_1 = 11\lambda_1.$$

Then we have

$$\begin{aligned} & \frac{n}{m}(\hat{\beta}^* - \beta^*)'Z'Z(\hat{\beta}^* - \beta^*) \\ & \leq 26\lambda_1 \|\hat{\beta}^* - \beta^*\|_1 + 40\lambda_2(\|\beta^*\|_1 - \|\hat{\beta}^*\|_1), \end{aligned}$$

and so

$$\begin{aligned} & \frac{n}{m}(\hat{\beta}^* - \beta^*)'Z'Z(\hat{\beta}^* - \beta^*) + 14\lambda_1 \|\hat{\beta}^* - \beta^*\|_1 \\ & \leq 40\lambda_1 \left(\|\hat{\beta}^* - \beta^*\|_1 + \|\beta^*\|_1 - \|\hat{\beta}^*\|_1 \right). \end{aligned}$$

Up to a multiplying constant, the rest of the proof of Theorem 4.1 is the same as the last lines in the proof of Theorem 3.3. Then we omit it here. \square

7.5 Proof of Proposition 4.2

Proof of Proposition 4.2. First, let us remark that

$$\begin{aligned} \left\| \left(X'X - \frac{n}{m} Z'Z \right) u \right\|_{\infty} &= n \sup_{1 \leq i \leq p} \sum_{j=1}^p u_j \left(\frac{X'_i X_j}{n} - \frac{Z'_i Z_j}{m} \right) \\ &\leq n \|u\|_1 \sup_{1 \leq i, j \leq p} \left| \frac{X'_i X_j}{n} - \frac{Z'_i Z_j}{m} \right|. \end{aligned}$$

Now, using the "exchangeable-distribution inequality" in [Cat07] we obtain, for a given pair (i, j) , for any $\tau > 0$, with probability at least $1 - \eta$,

$$\begin{aligned} \frac{X'_i X_j}{n} - \frac{Z'_i Z_j}{m} &\leq \frac{\tau k^2}{2n(k+1)^2} \left(\frac{1}{m} \sum_{k=1}^m X_{i,k}^2 X_{j,k}^2 \right) + \frac{\log \frac{1}{\eta}}{\tau} \\ &\leq \frac{\tau k^2 \kappa^2}{2n(k+1)^2} + \frac{\log \frac{1}{\eta}}{\tau} = \frac{\kappa k}{k-1} \sqrt{\frac{2 \log \frac{1}{\eta}}{n}}, \end{aligned}$$

for $\tau = (\log(1/\eta)(k-1)2n/k\kappa^2)^{1/2}$ and so, by a union bound argument, with probability at least $1 - \eta$, for any pair (i, j) ,

$$\left| \frac{X'_i X_j}{n} - \frac{Z'_i Z_j}{m} \right| \leq \frac{\kappa k}{k-1} \sqrt{\frac{2 \log \frac{2p^2}{\eta}}{n}} \leq \frac{2\kappa k}{k-1} \sqrt{\frac{2 \log \frac{p}{\eta}}{n}},$$

(where we used $p \geq 2$). □

References

- [AH08] P. Alquier and M. Hebiri. Generalization of l1 constraint for high-dimensional regression problems. Preprint Laboratoire de Probabilités et Modèles Aléatoires (n. 1253), arXiv:0811.0072, 2008.
- [Aka73] H. Akaike. Information theory and an extension of the maximum likelihood principle. In B. N. Petrov and F. Csaki, editors, *2nd International Symposium on Information Theory*, pages 267–281. Budapest: Akademia Kiado, 1973.
- [Alq08] P. Alquier. Lasso, iterative feature selection and the correlation selector: Oracle inequalities and numerical performances. *Electron. J. Stat.*, pages 1129–1152, 2008.
- [Bac08] F. Bach. Consistency of the group lasso and multiple kernel learning. *J. Mach. Learn. Res.*, 9:1179–1225, 2008.
- [BRT07] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of lasso and dantzig selector. Submitted to the *Ann. Statist.*, 2007.
- [BTW07] F. Bunea, A. Tsybakov, and M. Wegkamp. Aggregation for Gaussian regression. *Ann. Statist.*, 35(4):1674–1697, 2007.

- [Bun08] F. Bunea. *Consistent selection via the Lasso for high dimensional approximating regression models*, volume 3. IMS Collections, 2008.
- [Cat07] O. Catoni. *PAC-Bayesian Supervised Classification (The Thermodynamics of Statistical Learning)*, volume 56 of *Lecture Notes-Monograph Series*. IMS, 2007.
- [CH08] C. Chesneau and M. Hebiri. Some theoretical results on the grouped variables lasso. *Mathematical Methods of Statistics*, 17(4):317–326, 2008.
- [CSZ06] O. Chapelle, B. Schölkopf, and A. Zien. *Semi-supervised learning*. MIT Press, Cambridge, MA, 2006.
- [CT07] E. Candès and T. Tao. The dantzig selector: statistical estimation when p is much larger than n . *Ann. Statist.*, 35, 2007.
- [DT07] A. Dalalyan and A.B. Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. *COLT 2007 Proceedings. Lecture Notes in Computer Science 4539 Springer*, pages 97–111, 2007.
- [EHJT04] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004. With discussion, and a rejoinder by the authors.
- [FHHT07] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Ann. Appl. Statist.*, 1(2):302–332, 2007.
- [FL01] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96(456):1348–1360, 2001.
- [HCB08] C. Huang, G. L. H. Cheang, and A. Barron. Risk of penalized least squares, greedy selection and l1 penalization for flexible function libraries. preprint, 2008.
- [Heb08] M. Hebiri. Regularization with the smooth-lasso procedure. Preprint LPMA, 2008.
- [JRL09] G. James, P. Radchenko, and J. Lv. Dasso: Connections between the dantzig selector and lasso. *JRSS (B)*, 71:127–142, 2009.
- [KKL⁺07] S. J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. An interior-point method for large-scale l1-regularized least squares. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):606–617, 2007.
- [Kol07] V. Koltchinskii. Dantzig selector and sparsity oracle inequalities. Preprint, 2007.

- [Kol09] V. Koltchinskii. Sparse recovery in convex hulls via entropy penalization. *Annals of Statistics*, 37(3):1332–1359, 2009.
- [Lou08] K. Lounici. Sup-norm convergence rate and sign concentration property of Lasso and Dantzig estimators. *Electron. J. Stat.*, 2:90–102, 2008.
- [MB06] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *Ann. Statist.*, 34(3):1436–1462, 2006.
- [MVdGB08] L. Meier, S. Van de Geer, and P. Bühlmann. High-dimensional additive modeling. To appear in the *Annals of Statistics*, 2008.
- [MY09] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.*, 37(1):246–270, 2009.
- [OPT00] M. Osborne, B. Presnell, and B. Turlach. On the LASSO and its dual. *J. Comput. Graph. Statist.*, 9(2):319–337, 2000.
- [Sch78] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464, 1978.
- [Tib96] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.
- [Vap98] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1998.
- [vdG08] S. van de Geer. High-dimensional generalized linear models and the lasso. *Ann. Statist.*, 36(2):614–645, 2008.
- [Wai06] M. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using l_1 -constrained quadratic programming. Technical report n. 709, Department of Statistics, UC Berkeley, 2006.
- [YL07] M. Yuan and Y. Lin. On the non-negative garrotte estimator. *J. R. S. S. (B)*, 69(2):143–161, 2007.
- [Zou06] H. Zou. The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429, 2006.
- [ZY06] P. Zhao and B. Yu. On model selection consistency of Lasso. *J. Mach. Learn. Res.*, 7:2541–2563, 2006.