

# On Lasso refitting strategies

EVGENII CHZHEN<sup>1,\*</sup>, MOHAMED HEBIRI<sup>1,\*\*</sup> and JOSEPH SALMON<sup>2</sup>

<sup>1</sup>LAMA, Université Paris-Est, 5 Boulevard Descartes, 77420, Champs-sur-Marne, France.

E-mail: \*[evgenii.chzhen@univ-paris-est.fr](mailto:evgenii.chzhen@univ-paris-est.fr); \*\*[mohamed.hebiri@u-pem.fr](mailto:mohamed.hebiri@u-pem.fr)

<sup>2</sup>IMAG, Université de Montpellier, Place Eugène Bataillon, 34095 Montpellier Cedex 5, France.

E-mail: [joseph.salmon@umontpellier.fr](mailto:joseph.salmon@umontpellier.fr)

A well-known drawback of  $\ell_1$ -penalized estimators is the systematic shrinkage of the large coefficients towards zero. A simple remedy is to treat Lasso as a model-selection procedure and to perform a second refitting step on the selected support. In this work, we formalize the notion of refitting and provide oracle bounds for arbitrary refitting procedures of the Lasso solution. One of the most widely used refitting techniques which is based on Least-Squares may bring a problem of interpretability, since the signs of the refitted estimator might be flipped with respect to the original estimator. This problem arises from the fact that the Least-Squares refitting considers only the support of the Lasso solution, avoiding any information about signs or amplitudes. To this end, we define a sign consistent refitting as an arbitrary refitting procedure, preserving the signs of the first step Lasso solution and provide Oracle inequalities for such estimators. Finally, we consider special refitting strategies: Bregman Lasso and Boosted Lasso. Bregman Lasso has a fruitful property to converge to the Sign-Least-Squares refitting (Least-Squares with sign constraints), which provides with greater interpretability. We additionally study the Bregman Lasso refitting in the case of orthogonal design, providing with simple intuition behind the proposed method. Boosted Lasso, in contrast, considers information about magnitudes of the first Lasso step and allows to develop better oracle rates for prediction. Finally, we conduct an extensive numerical study to show advantages of one approach over others in different synthetic and semi-real scenarios.

*Keywords:* Bregman; Lasso; linear regression; refitting

## 1. Introduction

Least absolute shrinkage and selection operator (Lasso), introduced by Tibshirani [20], became a popular method in high dimensional statistics due to competitive numerical solvers (as it is a convex program) and fruitful statistical guaranties [2,15]. However, the shrinkage of large magnitudes towards zero, observed in practice, may affect the overall conclusion about the model. Different remedies were proposed to overcome this affect, all of them having their advantages and disadvantages. For instance, one may consider a non-convex penalty instead of the  $\ell_1$  regularization [11,13,25]: this approach increases a computational complexity and might be not applicable in large-scale scenarios. Another way to avoid the underestimation of the coefficients is to perform the second least-squares refitting step based on the first step Lasso solution [1,16]: such an approach brings the problem of interpretability, since the coefficients may switch signs with respect to the original Lasso solution. A lot of theoretical and applied works are devoted to the study of least-squares refitting of an arbitrary first step estimator [1,9,10,16].

Unlike such approaches, we are rather interested in general refitting strategies of the Lasso estimator. We provide a natural definition of what general refitting is, in the sense that we aim at reducing the data-fitting term of the original Lasso estimator. Moreover, since our initial (first step) estimator is the Lasso, our approach allows us to use previous theoretical analysis provided for Lasso to derive guarantees for a wide class of refitting (second step) strategies. In Section 2, we introduce notation, used throughout the article, and the basic Lasso theory is partly covered in Section 3. Readers who are familiar with the Lasso theory may skip Section 3 and proceed to the following sections. Section 4, is concerned with our theoretical framework, where we define a refitting strategy as an estimator which reaches a lower mean square error (MSE), compared to the first step Lasso solution. For this family of refitting strategies, we show that the rates for prediction are bounded by the Lasso rates plus an  $\ell_1$ -norm of difference between Lasso estimator and the refitted vector. Inspired by this result we propose to use an additional information, provided by the Lasso solution, for refitting. It leads to a least-squares refitting with constraints, to avoid an explosion of the refitted coefficients. This estimator can be seen as Boosted Lasso strategy (see Section 4.2 and more particularly Lemma 3) allows us to develop better prediction bounds compared to the classical Lasso bounds.

Additionally, we propose another family of refitting strategies in Section 4.1, which restricts the possibility to switch signs with respect to the first step Lasso solution in addition to lower MSE. For every refitting strategy in this family, we provide a unified oracle inequality stated in Theorem 3 showing minimax rates under the same assumptions as oracle inequalities for Lasso. We introduce Bregman Lasso, which can be seen as a generalization of Bregman Iterations [18, 19], widely used method in compressed sensing settings and has a strong connection with the method proposed by Brinkmann et al. [4]. Analyzing the Bregman Lasso in case of denoising model (Section 4.3.2) we provide useful insights on the proposed method. Additionally, we show that Bregman Lasso is a refitting strategy converging to Sign-Least-Squares Lasso, which can be tracked back in [4]. This estimator restricts the possibility to flip signs, minimizing MSE meanwhile. For Bregman Lasso, we conduct an intensive analysis in the orthogonal design case which is of independent interest and that exhibits some interesting interpretation of this method, and makes some analogies between Bregman Lasso and well-known existing thresholding methods (such as soft/hard/firm-thresholding).

Finally, we conduct an extensive numerical study of different post-Lasso refitting strategies to show advantages of different estimators in various scenarios. Let us conclude this section by summarizing our main contributions, in this paper we aim at:

- defining formally a refitting Lasso estimator,
- introducing specific refitting methods that exploit particular properties, such as preserving the Lasso signs, constraining the coefficients amplitudes, etc.,
- providing oracle inequalities for particular refitting strategies, such as Bregman iteration, Sign-Least-Squares refitting and the Boosted refitting,
- providing further understanding of the Bregman Iterations [18].

Proofs that are not presented in the main body of the article are gathered in the supplemental article [7].

## 2. Framework and notation

The standard Euclidean norm is written  $\|\cdot\|_2$ , the  $\ell_1$ -norm  $\|\cdot\|_1$ , and the  $\ell_\infty$ -norm  $\|\cdot\|_\infty$ . For any integer  $d \in \mathbb{N}$ , we denote by  $[d]$  the set  $\{1, \dots, d\}$  and by  $Q^\top$  the transpose of a matrix  $Q$ , and  $I_d \in \mathbb{R}^{d \times d}$  is the identity matrix of size  $d$ . For two real numbers  $a, b \in \mathbb{R}$  we defined by  $a \vee b$  the maximum between  $a$  and  $b$ . For any vectors  $a, b \in \mathbb{R}^p$  we denote by  $\langle a, b \rangle = a^\top b$  the Euclidean inner product and by  $a \odot b$  the element wise (Hadamard) product of two vectors. Our approach is valid for a broad class of models, but to avoid digression, we study the prediction performance of the Lasso and refitting strategies only for Gaussian linear regression models with deterministic design. More specifically, we consider  $n$  random observations  $y_1, \dots, y_n \in \mathbb{R}$  and fixed covariates  $x_1, \dots, x_n \in \mathbb{R}^p$ . We further assume that there is a regression vector  $\beta^* \in \mathbb{R}^p$  which satisfies the following relation:

$$y = X\beta^* + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I_n), \tag{1}$$

where  $y = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$  is the response vector and  $X = (x_1, \dots, x_n)^\top \in \mathbb{R}^{n \times p}$  the design matrix. We additionally assume, that the columns of  $X$  are normalized in such a way that for all  $j \in [p]$  we have  $\|X_j\|_2^2 = n$ , where  $X_j$  is  $j$ th column of the matrix  $X$ . For any set  $E \subset [p]$ , we denote by  $E^c$  the complement to  $E$  (i.e.,  $E \cup E^c = [p]$ ) and by  $X_E$  the matrix obtained from the matrix  $X$  by erasing all the columns whose indexes are not in  $E$ . Similarly, for any  $\beta \in \mathbb{R}^p$  we write  $\beta_E$  to denote the vector obtained from  $\beta$  by erasing all the components whose indexes are not in  $E$ . For all vectors  $\beta \in \mathbb{R}^p$  we write  $\text{supp}(\beta) \subset [p]$  for the support of the vector  $\beta$ , that is,  $\text{supp}(\beta) = \{j \in [p] : \beta_j \neq 0\}$ . For every real-valued function  $f : \mathbb{R}^p \mapsto \mathbb{R}$  we say that  $g \in \mathbb{R}^p$  is a subgradient of  $f$  at  $x \in \mathbb{R}^p$  if  $f(y) \geq f(x) + \langle g, y - x \rangle$  for all  $y \in \mathbb{R}^p$ . The set of all subgradients of  $f$  at  $x \in \mathbb{R}^p$  is called subdifferential of  $f$  at  $x \in \mathbb{R}^p$  and written as  $\partial f(x)$ . We also remind, that the subdifferential of the  $\ell_1$ -norm  $\partial \|\cdot\|_1$  is a set valued vector function  $\text{sign}(\cdot) = (\text{sign}(\cdot)_1, \dots, \text{sign}(\cdot)_p)^\top$ , defined element-wise by

$$\forall \beta \in \mathbb{R}^p, \forall j \in [p], \quad \text{sign}(\beta)_j = \begin{cases} \{1\}, & \beta_j > 0, \\ \{-1\}, & \beta_j < 0, \\ [-1, 1], & \beta_j = 0. \end{cases} \tag{2}$$

Also, we assume that the unknown vector  $\beta^*$  is sparse, that is,  $\text{supp}(\beta^*) = S$  has small cardinality  $s$  compared to  $n$  and  $p$ . To estimate  $\beta^*$ , we first minimize the negative log-likelihood with  $\ell_1$  penalty [20], which is equivalent for a fixed  $\lambda > 0$  to the following optimization problem

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1. \tag{3}$$

We also remind the Karush–Kuhn–Tucker (KKT) conditions for Eq. (3):

**Lemma 1 (KKT conditions for Lasso).** *The Karush–Kuhn–Tucker conditions [3] for the Lasso problem Eq. (3) read as follows:  $0 \in \lambda \text{sign}(\hat{\beta}) - \frac{1}{n} X^\top (y - X\hat{\beta})$ , or equivalently: there exists  $\hat{\rho} \in \text{sign}(\hat{\beta})$ , such that  $\frac{1}{n} X^\top (y - X\hat{\beta}) = \lambda \hat{\rho}$ .*

### 3. Lasso theory

In this section, we provide one of the classical Lasso oracle inequalities. To this end, we introduce the restricted eigenvalue condition [2], a widely used assumption on the design matrix  $X$ .

**Definition 1 (Bickel et al. [2]).** We say that  $X \in \mathbb{R}^{n \times p}$  satisfies the Restricted Eigenvalue condition  $\text{RE}(c_0, s)$ , where  $c_0 > 0$  and  $s \in [p]$ , if  $\exists \kappa(c_0, s) > 0$  such that for all  $J \subset [p]$  with  $|J| \leq s$  we have for all  $\Delta \in \mathbb{R}^p$

$$\|\Delta_{J^c}\|_1 \leq c_0 \|\Delta_J\|_1 \implies \frac{\|X\Delta\|_2^2}{n\|\Delta_J\|_2^2} \geq \kappa^2(c_0, s).$$

We also state below some classical concentration bound on tail of sup of Gaussian random variables.

**Lemma 2.** Let  $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_n)$  and  $X \in \mathbb{R}^{n \times p}$  be such that  $\forall j \in [p]$  we have  $\|X_j\|_2^2 = n$ , hence with probability at least  $1 - \delta$  we have

$$\|X^\top \varepsilon / n\|_\infty \leq \lambda / 2,$$

where  $\lambda = 2\sigma \sqrt{2 \log(p/\delta)/n}$ .

The following theorem is a starting point of our analysis, whose proof is recalled in the supplemental article [7]. We mention that similar techniques of the proof can be found in [8,14].

**Theorem 1.** If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition  $\text{RE}(3, s)$  and  $\lambda = 2\sigma \sqrt{2 \log(p/\delta)/n}$  for every  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds

$$\frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 \leq \frac{9\lambda^2 s}{4\kappa^2(3, s)},$$

where  $\hat{\beta}$  is a Lasso solution with tuning parameter  $\lambda$ .

### 4. Refitting strategies

Underestimation of large coefficients by Lasso and other  $\ell_1$ -penalized estimators have long been well known by practitioners, and simple remedies have been proposed on case by case analysis. One of such approaches is Least-Squares refitting – widely used in high-dimensional regression to reduce the bias of the coefficients and consists in performing a least-squares re-estimation of the non-zero coefficients of the solution. Such a procedure is theoretically analyzed by Belloni and Chernozhukov [1] applied to an arbitrary first-step estimator. Lederer [16], showed that blind Least-Squares refitting of the Lasso solution is not advised in all possible scenarios and developed a refitting criteria. Deledalle et al. [9,10] are mostly concerned with the practical aspects of

refitting and provide efficient numerical procedures to perform refitting simultaneously along with the original estimator. In contrast, here we are interested in arbitrary refitting of the Lasso solution, which allows to exploit Lasso theory to provide new insights. In this section, we study a general framework for refitting techniques. We define a refitting of a Lasso solution  $\hat{\beta}$  as:

**Definition 2.** Let  $\hat{\beta} = \hat{\beta}(\lambda)$  be a Lasso solution with regularization  $\lambda$ . We call a vector  $\bar{\beta} = \bar{\beta}(\hat{\beta}, X, y)$  a refitting of the Lasso  $\hat{\beta}$  if it reduces the original loss function, namely if:

$$\|y - X\bar{\beta}\|_2 \leq \|y - X\hat{\beta}\|_2. \tag{4}$$

Several known refitting strategies are falling inside of this framework, for instance, Least-Squares Lasso and Relaxed Lasso [17]. Relaxed Lasso was introduced by Meinshausen [17] and is defined for any positive  $\lambda$  and  $\phi \in [0, 1]$  as a solution to the following convex problem:

$$\bar{\beta}^{\lambda, \phi} \in \arg \min_{\beta: \text{supp}(\beta) \subset \text{supp}(\hat{\beta}^\lambda)} \frac{1}{2n} \|y - X\beta\|_2^2 + \phi\lambda \|\beta\|_1, \tag{5}$$

where  $\hat{\beta}^\lambda$  is the first step Lasso solution with parameter  $\lambda$ . To see that the relaxed Lasso in Eq. (5) is a refitting strategy in the sense of Definition 2 it is sufficient to notice that:

$$\begin{aligned} \frac{1}{2n} \|y - X\bar{\beta}^{\lambda, \phi}\|_2^2 + \phi\lambda \|\bar{\beta}^{\lambda, \phi}\|_1 &\leq \frac{1}{2n} \|y - X\hat{\beta}^\lambda\|_2^2 + \phi\lambda \|\hat{\beta}^\lambda\|_1; \\ \frac{1}{2n} \|y - X\hat{\beta}^\lambda\|_2^2 + \lambda \|\hat{\beta}^\lambda\|_1 &\leq \frac{1}{2n} \|y - X\bar{\beta}^{\lambda, \phi}\|_2^2 + \lambda \|\bar{\beta}^{\lambda, \phi}\|_1. \end{aligned}$$

Summing these inequalities and using the fact that  $\lambda - \lambda\phi \geq 0$ , we obtain the required condition in Definition 2.

We emphasize, that one should not expect a superior performance of a refitting from Definition 2, since the only information available for  $\bar{\beta}$  is the mean square error. Additionally notice, that the least-squares solution is obviously a refitting strategy for every Lasso solution, which justifies the previous remark. However, defining and analyzing such a family provides with interesting insights and serves as the step towards more thoughtful refitting strategies which are discussed in Section 4.2.

**Theorem 2.** For  $\lambda = 2\sigma\sqrt{2\log(p/\delta)/n}$ , where  $\delta \in (0, 1)$  is arbitrary, with probability  $1 - \delta$  the following bound holds

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 + \lambda \|\bar{\beta} - \hat{\beta}\|_1,$$

where  $\bar{\beta}$  is a refitting of Lasso solution  $\hat{\beta}$ .

Previous theorem shows, that by controlling the  $\ell_1$ -distance between refitting and Lasso solution, one might obtain satisfying performance, we discuss this idea in Section 4.2. In what follows, we consider specific refitting methods in order to refine prediction error bounds. First,

we consider the sign consistent refitting family that shares the same sign vector as the initial Lasso estimator. Then, we develop the Boosted Lasso, a refitting method that aims at reducing the  $\ell_1$  distance with the initial Lasso estimator. Finally, we introduce the Bregman Lasso, which somehow generalizes the Bregman Iterations [18] by adding some flexibility on the tuning parameters. We also highlight its connections with the Sign-Least-Squares Lasso, a particular sign consistent refitting strategy.

### 4.1. Sign consistent refitting strategies

Previous section is concerned with an arbitrary refitting strategy, which only uses the information about the mean square error of the Lasso solution. Here, we are interested in a more sophisticated family of refitting strategies, which additionally exploits the information provided by the sign of the Lasso solution. Such an approach has some similarities with the methods introduced by Brinkmann et al. [4], even though the authors had a different motivation.

**Definition 3.** Let  $\hat{\beta} = \hat{\beta}(\lambda)$  be a Lasso solution with regularization  $\lambda$ . We call a  $\bar{\beta} = \bar{\beta}(\hat{\beta}, X, y)$  sign-consistent refitting of the Lasso solution  $\hat{\beta}$  if

$$\|y - X\bar{\beta}\|_2 \leq \|y - X\hat{\beta}\|_2 \quad (\text{refitting}), \tag{6}$$

$$\frac{1}{n} X^\top (y - X\hat{\beta}) \in \lambda \text{sign}(\bar{\beta}) \quad (\text{sign-consistency}). \tag{7}$$

**Remark 1.** The sign consistency property in Definition 3 ensures that a sign-consistent refitting vector  $\bar{\beta}$  satisfies for all  $j \in [p]$  the following conditions

- if  $n^{-1} X_j^\top (y - X\hat{\beta}) = \lambda$ , hence  $\bar{\beta}_j \geq 0$ ;
- if  $n^{-1} X_j^\top (y - X\hat{\beta}) = -\lambda$ , hence  $\bar{\beta}_j \leq 0$ ;
- if  $n^{-1} |X_j^\top (y - X\hat{\beta})| < \lambda$ , hence  $\bar{\beta}_j = 0$ .

Notice, that the equations in the previous remark are exactly the first-order optimality conditions for the Lasso problem written component-wise. The next theorem shows that the definition of the sign-refitting allows to develop oracle rates without any additional assumptions, except classical ones used for Lasso bounds.

**Theorem 3.** *If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition RE(3, s) and  $\lambda = 6\sigma\sqrt{2\log(p/\delta)/n}$  for some  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds*

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 + \frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 \leq \lambda^2 s \left( \frac{9}{\kappa^2(3, s)} + \frac{16}{\kappa^2(1, s)} \right),$$

where  $\bar{\beta}$  is a sign-consistent refitting of Lasso solution  $\hat{\beta}$ .

The proof of this theorem is based on Definition 3, to be more precise, we can use an additional equality due to the sign-consistency property in Eq. (7), that is:

$$\frac{1}{n}(\bar{\beta} - \hat{\beta})^\top X^\top (y - X\hat{\beta}) = \lambda \langle \bar{\beta} - \hat{\beta}, \hat{\rho} \rangle = \lambda(\|\bar{\beta}\|_1 - \|\hat{\beta}\|_1).$$

This relation holds, since two estimators  $\hat{\beta}$  (Lasso) and  $\bar{\beta}$  (sign-refitting) share the same subgradient. Hence, we are able to reuse proof techniques similar to the one of Theorem 1 to provide an oracle inequality.

## 4.2. Boosted Lasso

Notice that to prove Theorem 2, we used only the refitting property Definition 2, one of the possible candidates is the following refitting step

$$\bar{\beta} \in \arg \min_{\beta \in \Gamma} \frac{1}{2n} \|y - X\beta\|_2^2, \tag{8}$$

where  $\Gamma = \{\beta \in \mathbb{R}^p : \|\beta - \hat{\beta}\|_1 \leq \hat{s}\lambda\}$  and  $\hat{s} = |\text{supp}(\hat{\beta})|$ . Intuitively, the Lasso coefficients are shrunk towards zero by a value proportional to the tuning parameter  $\lambda$ . Since there are  $\hat{s}$  non-zero coefficient in the Lasso solution, the proposed refitting strategy tries to “unshrink”  $\hat{s}$  non-zero coefficients. As we measure the “shrinkage” factor globally with the  $\ell_1$ -norm, it is natural (inspired by the orthogonal design) to set this factor as  $\hat{s}\lambda$ . This motivates our choice of the feasible set  $\Gamma$ . This estimator is legitimate following results given in [2], Theorem 7.2.

**Theorem 4 (Bickel et al. [2]).** *If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition  $\text{RE}(3, s)$  and  $\lambda = 2\sigma\sqrt{2\log(p/\delta)/n}$  for some  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds*

$$\hat{s} \leq \frac{9s\phi_{\max}}{\kappa^2(3, s)}, \tag{9}$$

where  $\phi_{\max}$  is the maximal eigenvalue of  $X^\top X/n$ .

The proof of this theorem is a direct application of Theorem 7.2 in [2] together with Theorem 1 to bound  $\frac{1}{n}\|X(\beta^* - \hat{\beta})\|_2^2$ , we then omit it here. For the estimator, described in Eq. (8), we can state the following result, which relies on using both Theorem 4 and Theorem 2:

**Corollary 1.** *If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition  $\text{RE}(3, s)$  and  $\lambda = 2\sigma\sqrt{2\log(p/\delta)/n}$  for some  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds*

$$\frac{1}{n}\|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{\lambda^2 s}{\kappa^2(3, s)} \left( \frac{9}{4} + 9\phi_{\max} \right).$$

Note that the control over the magnitudes of the coefficients allows to develop desirable rates for the refitted estimator. The description of the set  $\Gamma$  in the definition of the refitting (8) motivates us to consider the following Boosted Lasso estimator [6].

**Definition 4 (Boosted Lasso).** For any  $\lambda_1, \lambda_2 > 0$  we call  $\hat{\beta}^{\lambda_1, \lambda_2}$  a Boosted Lasso refitting if it is a solution of

$$\hat{\beta}^{\lambda_1, \lambda_2} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda_2 \|\beta - \hat{\beta}^{\lambda_1}\|_1, \tag{10}$$

where  $\hat{\beta}^{\lambda_1}$  is the Lasso solution with tuning parameter  $\lambda_1$ .

**Remark 2.** Indeed, procedure (10) consists in applying the *Twicing* method [22] to the Lasso, which is equivalent to a two-step boosting approach. To see that, let  $\hat{\Delta} = \beta - \hat{\beta}$ . Then, with a change of variable we get:

$$\hat{\Delta} \in \arg \min_{\Delta \in \mathbb{R}^p} \frac{1}{2n} \|(y - X\hat{\beta}^{\lambda_1}) - X\Delta\|_2^2 + \lambda_2 \|\Delta\|_1, \tag{11}$$

and finally  $\hat{\beta}^{\lambda_1, \lambda_2} = \hat{\Delta} + \hat{\beta}^{\lambda_1}$ .

It is known result that in the Lasso case there exists a critical value  $\lambda_{1, \max} = \|X^\top y/n\|_\infty$  such that  $\hat{\beta}^{\lambda_1} = 0$  iff  $\lambda_1 \geq \lambda_{1, \max}$ , due to the previous remark, the Boosted Lasso can be written as a Lasso problem and we can give the result of the same nature.

**Proposition 1.** *If  $\lambda_1 < \lambda_{1, \max}$ , then the solution of Boosted Lasso satisfies:  $\hat{\beta}^{\lambda_1, \lambda_2} = \hat{\beta}^{\lambda_1}$  iff  $\lambda_2 \geq \lambda_1$ . Moreover, if  $\lambda_1 \geq \lambda_{1, \max}$ , then the Boosted Lasso estimator is simply the Lasso estimator with tuning parameter  $\lambda_2$ .*

Control over the  $\ell_1$ -norm of the difference, allows to develop oracle inequality with minimax rates (Corollary 1), similar result but somehow stronger can be shown for the regularized version (Boosted Lasso).

**Lemma 3.** *For any  $\hat{\beta}^{\lambda_1}$  on the event  $\{\|\varepsilon^\top X/n\|_\infty \leq \lambda_1/2\}$  we have:*

$$\frac{1}{n} \|X(\beta^* - \hat{\beta}^{\lambda_1, \lambda_2})\|_2^2 + (2\lambda_2 - \lambda_1) \|\hat{\beta}^{\lambda_1, \lambda_2} - \hat{\beta}^{\lambda_1}\|_1 \leq \frac{1}{n} \|X(\beta^* - \hat{\beta}^{\lambda_1})\|_2^2. \tag{12}$$

Combining Lemma 3 and Theorem 1 we can state the following corollary.

**Corollary 2.** *If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition RE(3, s) and  $\lambda_1 = 2\sigma\sqrt{2\log(p/\delta)}/n$  for some  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds*

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 + (2\lambda_2 - \lambda_1) \|\bar{\beta} - \hat{\beta}\|_1 \leq \frac{9\lambda_1^2 s}{4\kappa^2(3, s)},$$

where  $\bar{\beta} = \hat{\beta}^{\lambda_1, \lambda_2}$  is a boosting refitting of Lasso solution  $\hat{\beta} = \hat{\beta}^{\lambda_1}$ .



The Boosted Lasso is obviously a refitting strategy, and the previous result shows that it is worth refitting in terms of prediction error. Moreover, Proposition 1 and Corollary 2 suggest to select  $\lambda_2 \in (\lambda_1/2, \lambda_1)$ , such a choice allows to improve the Lasso prediction accuracy with high probability. Besides, the result can be applied iteratively as it does not depend on the choice of  $\hat{\beta}^{\lambda_1}$ , hence it works for any iteration step.

**Remark 3.** A possible extension of the Boosted Lasso is the Boosted Support Lasso:

$$\bar{\beta}^{\lambda_1, \lambda_2} \in \arg \min_{\beta: \text{supp}(\beta) \subset \text{supp}(\hat{\beta}^{\lambda_1})} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda_2 \|\beta - \hat{\beta}^{\lambda_1}\|_1, \tag{13}$$

this estimator is inspired by the Relaxed Lasso in Eq. (5) and the Boosted Lasso in Eq. (10). Interestingly, both Lemma 3 and Corollary 2 hold for the Boosted Support Lasso refitting and follow completely identical proof.

### 4.3. Bregman Lasso

We first remind the definition of the Bregman divergence associated with the  $\ell_1$ -norm.

**Definition 5 (Bregman divergence for the  $\ell_1$ -norm).** For any  $z, w \in \mathbb{R}^p$  and any  $\rho \in \partial \|w\|_1$ , the Bregman divergence for the  $\ell_1$ -norm is defined as

$$\mathcal{D}_{\ell_1}^\rho(z, w) = \|z\|_1 - \|w\|_1 - \langle \rho, z - w \rangle \geq 0, \quad \rho \in \partial \|w\|_1. \tag{14}$$

In [18], the authors proposed the Bregman Iterations procedure, originally designed to improve iso-TV results. In Lasso case, Bregman Iterations has the following expression: for a fixed  $\lambda > 0$ , initializing with  $\hat{\rho}_0 = \hat{\beta}_0 = 0_p$ ,

$$\begin{aligned} \hat{\beta}_k &\in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \mathcal{D}_{\ell_1}^{\hat{\rho}_{k-1}}(\beta, \hat{\beta}_{k-1}), \\ \text{s.t. } \hat{\rho}_k &= \hat{\rho}_{k-1} + \frac{1}{\lambda n} X^\top (y - X\hat{\beta}_k), \end{aligned} \tag{15}$$

where  $\mathcal{D}_{\ell_1}^\rho(\cdot, \cdot)$  is the Bregman divergence defined in Eq. (14). The Bregman Iterations can be seen as a discretization of Bregman Inverse Scale Space (ISS), which is analyzed by Osher et al. [19], who provided statistical guarantees for the ISS dynamic. One of the drawbacks of such an approach is the iterative nature of the algorithm: one need to tune the number of iterations  $k$  and the regularization parameter  $\lambda > 0$ . In the recent work by Brinkmann et al. [4], the authors proposed another closely related algorithm, which for a given Lasso solution  $\hat{\beta}^{\lambda_1}$ , performs the following refitting

$$\begin{aligned} \bar{\beta} &\in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 \\ \text{s.t. } \mathcal{D}_{\ell_1}^{\hat{\rho}^{\lambda_1}}(\beta, \hat{\beta}^{\lambda_1}) &= 0. \end{aligned} \tag{16}$$

Unlike the previous approaches, in this section we consider the Lasso solution  $\hat{\beta}^{\lambda_1}$  and the following Bregman Lasso refitting strategy, defined as the following.

**Definition 6 (Bregman Lasso).** For any  $\lambda_1, \lambda_2 > 0$  we call  $\hat{\beta}^{\lambda_1, \lambda_2}$  a Bregman Lasso refitting a solution of

$$\hat{\beta}^{\lambda_1, \lambda_2} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda_2 \mathcal{D}_{\ell_1}^{\hat{\rho}^{\lambda_1}}(\beta, \hat{\beta}^{\lambda_1}), \tag{17}$$

where  $\hat{\beta}^{\lambda_1}$  is a Lasso solution with tuning parameter  $\lambda_1$  and  $\mathcal{D}_{\ell_1}^{\hat{\rho}^{\lambda_1}}(\cdot, \cdot)$  is the Bregman divergence given in Eq. (14).

Considering the particular case of orthogonal design, we show in Section 4.3.2 that Bregman Lasso refitting is a generalization of both approaches. In particular for this design, the Bregman Lasso refitting can recover Bregman Iterations for any  $k$ . Also, Bregman Lasso is computationally more appealing since it only requires evaluating two Lasso problems, while the Bregman Iterations would require  $k$  evaluations.

We start by introducing some basic properties of the Bregman divergence associated with the  $\ell_1$ -norm.

**Lemma 4.** Let  $z, w \in \mathbb{R}^p$ , and denote  $\rho \in \partial \|w\|_1$  a subgradient of the  $\ell_1$ -norm evaluated at  $w \in \mathbb{R}^p$ , therefore the following properties hold independent of the choice of  $\rho$

1.  $z \mapsto \mathcal{D}_{\ell_1}^\rho(z, w)$  is convex for all  $w$ ,
2.  $\mathcal{D}_{\ell_1}^\rho(z, w) = \|z\|_1 - \langle \rho, z \rangle$ ,
3.  $0 \leq \mathcal{D}_{\ell_1}^\rho(z, w) \leq 2\|z\|_1$ ,
4.  $\mathcal{D}_{\ell_1}^\rho(z, w) = \sum_{i=1}^p |z_i| - z_i \rho_i = \sum_{i=1}^p \mathcal{D}_{\ell_1}^{\rho_i}(z_i, w_i)$ , where  $z_i, w_i, \rho_i$  are the  $i$ th components of  $z, w, \rho$ ,
5. if  $\text{sign}(z) = \text{sign}(w)$ , therefore  $\mathcal{D}_{\ell_1}^\rho(z, w) = 0$ .


There is a simple geometrical interpretation of Bregman divergence, associated with the  $\ell_1$ -norm, which is illustrated by Figure 1. According to Property 2 of Lemma 4, Bregman divergence in Eq. (17) can be evaluated as

$$\mathcal{D}_{\ell_1}^{\hat{\rho}^{\lambda_1}}(\beta, \hat{\beta}^{\lambda_1}) = \|\beta\|_1 - \langle \hat{\rho}^{\lambda_1}, \beta \rangle,$$

where  $\hat{\rho}^{\lambda_1}$  is a fixed subgradient of the  $\ell_1$ -norm evaluated at  $\hat{\beta}^{\lambda_1}$ . Since, the subdifferential of the  $\ell_1$ -norm is not uniquely defined when evaluated at zero, it is important to fix the way to pick a subgradient. Possible, and probably the most obvious way to evaluate the subgradient is to write the KKT conditions for Problem (17) and fix  $\hat{\rho}^{\lambda_1}$  as follows

$$\hat{\rho}^{\lambda_1} = \frac{1}{\lambda_1 n} X^\top (y - X\hat{\beta}^{\lambda_1}). \tag{18}$$

Starting from here we stick to this choice of the subgradient.



**Figure 1.** Geometrical interpretation of Bregman divergence. On the left plot the subgradient is uniquely defined and the Bregman divergence equals  $2|z|$ . On the right plot the subgradient is a set, therefore the Bregman divergence can be any number from 0 to  $2|z|$  depending on the choice of  $\rho$ .

**Proposition 2.** *With the choice of the subgradient as in Eq. (18) and setting  $\bar{y} = y + \frac{\lambda_2}{\lambda_1}(y - X\hat{\beta}^{\lambda_1})$ , the Bregman Lasso refitting step can be evaluated as*

$$\hat{\beta}^{\lambda_1, \lambda_2} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|\bar{y} - X\beta\|_2^2 + \lambda_2 \|\beta\|_1. \tag{19}$$

Proposition 2 suggests that one can easily compute the Bregman Lasso refitting based on some Lasso solver using a modified response vector. Moreover, it is important to notice that the Bregman refitting is a particular instance of our proposed framework. In particular the following result holds.

**Proposition 3.** *Bregman Lasso refitting is a refitting strategy in the sense of Definition 2.*

**Theorem 5.** *If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition RE(3, s) and  $\lambda_1 = \frac{1+2r}{r} 2\sigma \sqrt{2 \log(p/\delta)/n}$  for some  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds:*

$$\frac{r+2}{r} \frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 + \frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 \leq \frac{\lambda_1^2 s}{\kappa^2(1, s)} + \frac{9\lambda_1^2 s}{4\kappa^2(3, s)},$$

where  $r = \lambda_2/\lambda_1$  and  $\bar{\beta} = \hat{\beta}^{\lambda_1, \lambda_2}$  is a Bregman refitting of Lasso solution  $\hat{\beta} = \hat{\beta}^{\lambda_1}$ .

The proof can be found in the supplemental article [7], however, let us mention that an important step in the proof is to obtain the following inequality:

$$\|X(\hat{\beta} - \bar{\beta})\|_2^2 + \|y - X\bar{\beta}\|_2^2 \leq \|y - X\hat{\beta}\|_2^2, \tag{20}$$

which quantifies the improvement for the data-fitting term. Several interesting consequences can be mentioned based on Theorem 5: first notice that the oracle inequality depends weakly on

the second parameter  $\lambda_2$ , which indicates that the Bregman refitting is not sensitive to the second parameter; also, the rate on the right hand side  $9\lambda^2s/4\kappa^2(3, s)$  is matching the Lasso upper bound up to a constant term (to the cost of a lower probability).

**Remark 4.** A careful analysis of the proof suggests a stronger result: under the assumptions of Theorem 5, with probability  $1 - \delta$  at least one of the following inequalities holds:

$$\begin{aligned} \frac{r + 2}{r} \frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 + \frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 &\leq \frac{9\lambda_1^2s}{4\kappa^2(3, s)}, \\ \frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 &\leq \frac{9\lambda_1^2s}{4\kappa^2(1, s)}, \end{aligned}$$

where  $r = \lambda_2/\lambda_1$  and  $\bar{\beta} = \hat{\beta}^{\lambda_1, \lambda_2}$  is the Bregman refitting of the Lasso solution  $\hat{\beta} = \hat{\beta}^{\lambda_1}$ . Note that although one cannot tell apart which one of these two inequalities holds, the prediction bounds are improved in both cases. In case the first inequality is true and the classical Lasso bound is tight the Bregman refitting improves on the Lasso. In case the second inequality holds then the Lasso bound from Theorem 1 can be improved using a better constant  $\kappa^2(1, s)$  instead of  $\kappa^2(3, s)$ , which can differ significantly.

### 4.3.1. Geometrical interpretation

It is known that (3) can be equivalently written in the following form

$$\begin{aligned} \hat{\beta}^{T_1} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 \\ \text{s.t. } \|\beta\|_1 \leq T_1, \end{aligned} \tag{21}$$

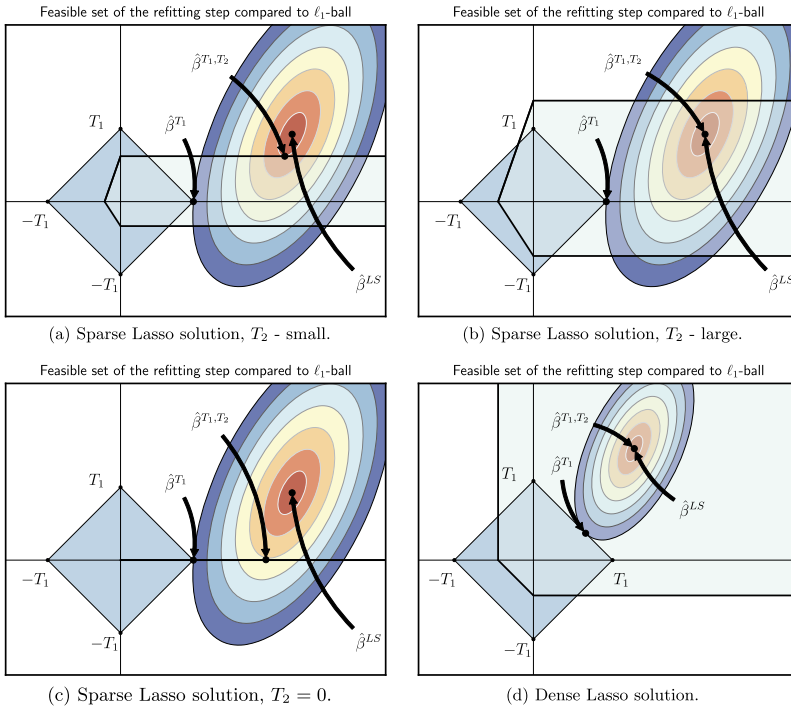
for some  $T_1 = T_1(\lambda_1) \geq 0$ . Similarly, we can write the Bregman Lasso refitting step Eq. (17) in the constrained form as

$$\begin{aligned} \hat{\beta}^{T_1, T_2} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 \\ \text{s.t. } \mathcal{D}_{\ell_1}^{\hat{\beta}^{T_1}}(\beta, \hat{\beta}^{T_1}) \leq T_2, \end{aligned} \tag{22}$$

where  $\hat{\beta}^{T_1}$  is a solution of the constrained version given in Eq. (21). The choice  $T_2 = 0$  is considered in [4], which corresponds to  $\lambda_2$  being large enough in Eq. (17). To provide a geometrical intuition we consider the simple case where  $n = p = 2$ . We denote by  $\hat{\rho}^{T_1}$  the subgradient of the  $\ell_1$ -norm evaluated at the first-step estimator  $\hat{\beta}^{T_1}$ . Assume, that the first-step estimator is given, therefore we consider two principal scenarios (case with negative values could be obtained symmetrically):

- When  $\hat{\beta}_1^{T_1} > 0, \hat{\beta}_2^{T_1} = 0$ , the feasible set of Problem (22) is given by

$$\{(\beta_1, \beta_2)^T \in \mathbb{R} \times \mathbb{R} : |\beta_1| - \beta_1 + |\beta_2| - \hat{\rho}_2^{T_1} \beta_2 \leq T_2\}.$$



**Figure 2.** Feasible set of the first and the refitting steps. Here, the ellipses are levels of the objective function and  $\hat{\beta}^{LS}$  is the least squares estimator. Geometry of the feasible set for the refitting step is described in terms of  $T_2$  and the subgradient  $\hat{\rho}^{T_1}$ .

Notice that any positive value for  $\beta_1$  is admissible in this case. We illustrate this phenomenon on Figure 2(a), Figure 2(b) and Figure 2(c) for various values of  $T_2$ .

- When  $\hat{\beta}_1^{T_1} > 0, \hat{\beta}_2^{T_1} > 0$ , the feasible set of Problem (22) is given by

$$\{(\beta_1, \beta_2)^T \in \mathbb{R} \times \mathbb{R} : |\beta_1| - \beta_1 + |\beta_2| - \beta_2 \leq T_2\}.$$

Notice that every pair  $(\beta_1, \beta_2)^T \in \mathbb{R}^+ \times \mathbb{R}^+$  is inside of the feasible set. We illustrate this phenomenon on Figure 2(d).

In the Lasso case, if  $T_1 = 0$  there is only one trivial solution  $\hat{\beta}^{T_1} = 0$ , this is not the case for the refitting step. Indeed, if  $T_2 = 0$  the feasible set of the refitting step might be non-trivial, depending on the first step Lasso solution. More precisely, if there exists  $j_1 \in [p]$ , such that  $\hat{\beta}_{j_1}^{T_1} \neq 0$  the feasible set of the refitting step always contains  $\{\beta : \beta_{j_1} \leq 0, \beta_{j_2} = 0 \forall j_2 \neq j_1\}$  or  $\{\beta : \beta_{j_1} \geq 0, \beta_{j_2} = 0 \forall j_2 \neq j_1\}$ , depending on  $\text{sign}(\hat{\beta}_{j_1}^{T_1})$ .

4.3.2. Orthogonal design

In this section, we investigate some important properties of the algorithm in Eq. (17) for the denoising model (*i.e.*, we drop the statistical convention and instead assume that  $n = p$  and  $X = I_p$ )  $y = \beta^* + \varepsilon$ , where  $y, \beta^* \in \mathbb{R}^p$ ,  $\varepsilon \sim \mathcal{N}(0, \sigma^2 \text{Id}_p)$ . The following estimators correspond to Eq. (3) and Eq. (17), respectively

$$\hat{\beta}^{\lambda_1} = \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - \beta\|_2^2 + \lambda_1 \|\beta\|_1, \tag{23}$$

$$\hat{\beta}^{\lambda_1, \lambda_2} = \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - \beta\|_2^2 + \lambda_2 (\|\beta\|_1 - \langle \hat{\rho}^{\lambda_1}, \beta \rangle), \tag{24}$$

where the subgradient  $\hat{\rho}^{\lambda_1}$  is given by Eq. (18) and simplifies to

$$\hat{\rho}^{\lambda_1} = \frac{y - \hat{\beta}^{\lambda_1}}{\lambda_1}. \tag{25}$$

We remind that for orthogonal design, the Lasso estimator is simply a soft-thresholding version of the observation  $\hat{\beta}^{\lambda_1} = \text{ST}(y, \lambda_1)$ , where  $\text{ST}(\cdot, \cdot)$  is defined component-wise for any  $j \in [p]$  by:  $\text{ST}(y, \lambda_1)_j = \text{sign}(y_j)(|y_j| - \lambda_1)_+$ .

The next proposition shows that the subgradient follows the signs of  $y$ .

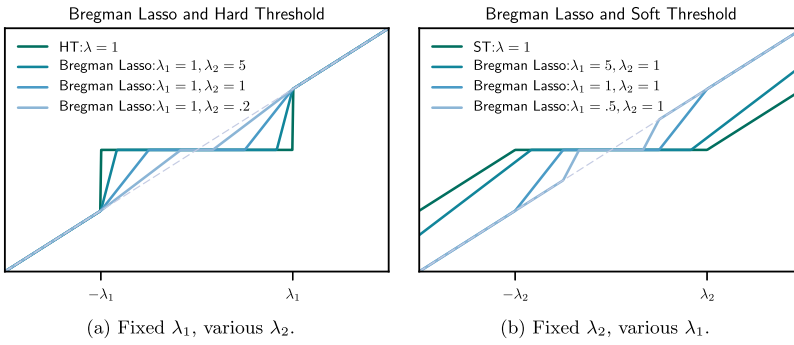
**Proposition 4.** *For all  $\lambda_1 > 0$  we have  $\text{sign}(y_j) = \text{sign}(\hat{\rho}_j^{\lambda_1}), \forall j \in [p]$ .*

The following results states that the Bregman refitting also enjoys close properties to the Lasso but thresholds a translated response vector. This relation can be obviously proved using Eq. (19) of Proposition 2, which establishes that the Bregman Lasso can be seen as Lasso solution applied to a modified signal.

**Proposition 5.** *The solution of the Bregman Lasso refitting step in Eq. (24) relies on the soft-threshold operator and reads:  $\hat{\beta}^{\lambda_1, \lambda_2} = \text{ST}(y + \lambda_2 \hat{\rho}^{\lambda_1}, \lambda_2)$ .*

It is worth mentioning, that in the Lasso case, there exists a so called  $\lambda_{1, \max}$ , which is the smallest value of regularization parameter for which the solution  $\hat{\beta}^{\lambda_1} = 0$  for all  $\lambda_1 > \lambda_{1, \max}$ . However, even though the refitting step can be formulated as Lasso problem, there is not such parameter  $\lambda_{2, \max}$ . It can be counter intuitive on the first sight, but since the parameter  $\lambda_2$  is present inside the data-fitting term, it becomes clear that such extreme value does not exist.

In the sequel, we provide another interpretation of the Bregman Lasso refitting in terms of firm-thresholding operator, introduced and analyzed by Gao and Bruce [12]. We additionally mention, that the firm-thresholding operator is the solution of the least squares problem penalized with MCP regularization [25] (for orthogonal design), which is a non-convex problem. Additionally, the firm-thresholding operator outperforms soft/hard-threshold in terms of bias-variance trade-off, see [12] for theoretical and numerical analysis.



**Figure 3.** The solution of the refitting step is equivalent to the MCP penalty. Extreme cases include both hard and soft threshold operators.

**Proposition 6.** *The solution of the Bregman Lasso refitting step in case of orthogonal design is given by:  $\hat{\beta}^{\lambda_1, \lambda_2} = \text{FT}(y, \lambda_H, 1 + \frac{\lambda_1}{\lambda_2})$ , where  $\lambda_H = (1/\lambda_1 + 1/\lambda_2)^{-1}$  and the firm-thresholding operator FT is defined for  $\mu > 0, \gamma > 1$  component-wise for any  $j \in [p]$  by:*

$$\text{FT}(y, \mu, \gamma)_j = \begin{cases} \frac{\gamma}{\gamma - 1} \text{ST}(y_j, \mu), & |y_j| \leq \mu\gamma, \\ y_j, & |y_j| > \mu\gamma. \end{cases} \tag{26}$$

In the Bregman Lasso (orthogonal) case  $\mu = \lambda_H$  and  $\gamma = 1 + \lambda_1/\lambda_2 = \lambda_1/\lambda_H$ .

**Remark 5.** We notice the following behavior of the solution, described in Proposition 6:

- $\hat{\beta}^{\lambda_1, \lambda_2} \rightarrow \text{HT}(y, \lambda_1)$ , when  $\lambda_2 \rightarrow \infty$ ,
- $\hat{\beta}^{\lambda_1, \lambda_2} \rightarrow \text{ST}(y, \lambda_2)$ , when  $\lambda_1 \rightarrow \infty$ .

These properties are illustrated on Figures 3(a) and 3(b) for several values of  $\lambda_1$  and  $\lambda_2$ .

We would like to emphasize that the Bregman Lasso refitting formulated in the orthogonal case, that is, Eq. (26), coincides with the Firm Thresholding [12]. In the extreme case where  $\lambda_2 \rightarrow \infty$ , one recovers the Hard Thresholding. Note that these two cases are connected to non-convex regularizations corresponding to the MCP [25] and the  $\ell_0$ . For the sake of completeness, we additionally provide the analysis for the Bregman Iterations (orthogonal case) in the form of Eq. (15) to give insights on our motivation to consider two-step iteration with two tuning parameters  $\lambda_1, \lambda_2$ . Similar analysis can be found in [23,24]. Following the proof of Proposition 5 one can prove by induction that the following result holds for the Bregman Iterations defined in Eq. (15).

**Proposition 7.** *For any  $\lambda > 0$  and  $k > 0$ , the Bregman Iteration Eq. (15) is given by*

$$\hat{\beta}_{k+1} = \text{ST}(y + \lambda \hat{\rho}_k, \lambda).$$

We also prove that for every  $k > 1$ , firm-thresholding yields to Bregman Iterations.

**Proposition 8.** *For any  $\lambda > 0$  and  $k > 0$ , the solution generated by the Bregman Iterations Eq. (15) is given by*

$$\hat{\beta}_{k+1} = \text{FT}\left(y, \frac{\lambda}{k+1}, \frac{k+1}{k}\right) = \begin{cases} (k+1) \text{ST}\left(y, \frac{\lambda}{k+1}\right), & \text{if } |y_j| \leq \frac{\lambda}{k}, \\ y_j & \text{if } |y_j| > \frac{\lambda}{k}. \end{cases}$$

Analyzing the previous result helps motivating the introduction of the Bregman Lasso refitting in the form of Eq. (17). Indeed, for orthogonal design, Bregman Lasso refitting generalizes Bregman Iterations Eq. (15) in the sense that for each pair  $(k, \lambda)$  there exists a pair  $(\lambda_1 = \frac{\lambda}{k-1}, \lambda_2 = \lambda)$  such that  $\hat{\beta}_k = \hat{\beta}^{\lambda_1, \lambda_2}$ . In particular, Bregman Lasso refitting includes all possible solution of Bregman Iterations Eq. (15), while the reverse is not true, for instance, when  $k$  is not an integer.

### 5. Sign-Least-Squares Lasso

In this section, we provide some generalizations of the results obtained in Section 4.3.2. We first notice that in the case of orthogonal design, discussed in Section 4.3.2, for a given vector  $y$ , there exists a value of the regularization parameter  $\lambda_2$  such that, the refitting step Eq. (17) is equivalent to hard-thresholding. One might expect that for the case of arbitrary design there exists such a value of  $\lambda_2$  that the refitting step Eq. (17) is equivalent to a Least-Squares refitting on the support obtained via the first Lasso type step as it adds a sign constraints. Yet, the estimator obtained via the refitting step slightly differs from the simple Least-Squares refitting. To state our main result of this section, let us introduce some notation. Consider  $\hat{\rho}^{\lambda_1}$ , defined in Eq. (18), we define the equicorrelation set [21] as

$$E^{\lambda_1} = E(\hat{\rho}^{\lambda_1}) = \{j \in [p] : |\hat{\rho}_j^{\lambda_1}| = 1\}. \tag{27}$$

We will omit  $\lambda_1$  in  $E^{\lambda_1}$  and write  $E$  instead for simplicity. We associate the following Sign-Least-Squares Lasso refitting step with the equicorrelation set given by Eq. (27).

**Definition 7 (Sign-Least-Squares Lasso).** For any  $\lambda_1 > 0$  we call  $\hat{\beta}^{\text{SLS}}$  a Sign-Least-Squares Lasso refitting if it satisfies  $\hat{\beta}_{E^c}^{\text{SLS}} = 0$  and

$$\hat{\beta}_E^{\text{SLS}} \in \arg \min_{\beta_E \in \mathbb{R}^{|E|} : \hat{\rho}_E^{\lambda_1} \circ \beta_E \geq 0} \frac{1}{2n} \|y - X_E \beta_E\|_2^2, \tag{28}$$

where  $|E|$  is the cardinality of  $E$  and  $\hat{\rho}^{\lambda_1}$  is the Lasso subgradient defined in Eq. (18).

Notice, that the refitting is performed on the equicorrelation set and not on the support of the Lasso solution. Possible motivation to consider the equicorrelation set instead of the support can



be described by uniqueness issues of the Lasso [21], while the equicorrelation set (and the signs) is always uniquely defined.

**Proposition 9.** *The Sign-Least-Squares Lasso is a sign consistent refitting strategy of the Lasso solution  $\hat{\beta}^{\lambda_1}$  in the sense of Definition 3.*

**Proof.** Notice that we have the relation  $\{\beta_E \in \mathbb{R}^{|E|} : \hat{\rho}_E^{\lambda_1} \odot \beta_E \geq 0\} = \{\beta_E \in \mathbb{R}^{|E|} : \hat{\rho}^{\lambda_1} \in \text{sign}(\beta)\}$ , and using Remark 1 we conclude.  $\square$

We believe that the introduced Sign-Least-Squares refitting of the Lasso is an interesting and simple alternative to the Least-Squares Lasso. Indeed, both methods achieve similar theoretical guarantees and their computational costs are equivalent (since the estimated supports are generally small). Moreover, the Sign-Least-Squares refitting of the Lasso might be more appealing to practitioners since confusing signs switches are no longer possible.

### 5.1. Bregman Lasso as Sign-Least-Squares Lasso

In this section, we provide connections between the Bregman Lasso and the Sign-Least-Squares Lasso. Since Problem (28) is convex and Slater’s condition is satisfied therefore the KKT conditions state [3] that there exist  $\hat{\beta}_E^{\text{SLS}} \in \mathbb{R}^{|E|}$  and  $\hat{\mu} \in \mathbb{R}_+^{|E|}$ , such that

$$\begin{cases} \frac{1}{n} X_E^\top (y - X_E \hat{\beta}_E^{\text{SLS}}) + \hat{\rho}_E^{\lambda_1} \odot \hat{\mu} = 0, \\ \hat{\mu} \odot \hat{\beta}_E^{\text{SLS}} \odot \hat{\rho}_E^{\lambda_1} = 0, \\ \hat{\mu} \geq 0, \end{cases} \quad (29)$$

so that we can provide the following connection.

**Theorem 6.** *For any signal vector  $y \in \mathbb{R}^n$  and any design matrix  $X \in \mathbb{R}^{n \times p}$  there exists  $\lambda_0 > 0$  such that for any  $\lambda_2 \geq \lambda_0$  the solution of the refitting step Eq. (17) is given by  $\hat{\beta}^{\lambda_1, \lambda_2} = \hat{\beta}^{\text{SLS}}$ .*

The proof of this theorem is below. However we mention that it states that given  $\lambda_1 > 0$ , Bregman Lasso and Sign-Least-Squares Lasso often coincide. The value of the parameter  $\lambda_0$  in Theorem 6 is explicit and can be found at the end of the following proof.

**Proof.** We start by writing the KKT conditions for Bregman Lasso, Eq. (17):

$$\begin{cases} \hat{\rho}_E^{\lambda_1} + \frac{1}{\lambda_2 n} X_E^\top (y - X_E \hat{\beta}_E^{\lambda_1, \lambda_2} - X_{E^c} \hat{\beta}_{E^c}^{\lambda_1, \lambda_2}) \in \text{sign}(\hat{\beta}_E^{\lambda_1, \lambda_2}), \\ \hat{\rho}_{E^c}^{\lambda_1} + \frac{1}{\lambda_2 n} X_{E^c}^\top (y - X_E \hat{\beta}_E^{\lambda_1, \lambda_2} - X_{E^c} \hat{\beta}_{E^c}^{\lambda_1, \lambda_2}) \in \text{sign}(\hat{\beta}_{E^c}^{\lambda_1, \lambda_2}). \end{cases} \quad (30)$$

These are necessary and sufficient conditions for a vector  $\beta$  to be a solution of Bregman Lasso (17). Therefore, it is sufficient to check if the vector defined as  $\hat{\beta}^{\lambda_1, \lambda_2} = \hat{\beta}^{\text{SLS}}$  satisfies Eq. (30)

for some  $\lambda_2$ . Substituting  $\hat{\beta}_E^{\lambda_1, \lambda_2} = \hat{\beta}_E^{\text{SLS}}$  and  $\hat{\beta}_{E^c}^{\lambda_1, \lambda_2} = 0$  we arrive at

$$\begin{cases} \hat{\rho}_E^{\lambda_1} + \frac{1}{\lambda_2 n} X_E^\top (y - X_E \hat{\beta}_E^{\text{SLS}}) \in \text{sign}(\hat{\beta}_E^{\text{SLS}}), \\ \hat{\rho}_{E^c}^{\lambda_1} + \frac{1}{\lambda_2 n} X_{E^c}^\top (y - X_E \hat{\beta}_E^{\text{SLS}}) \in \text{sign}(0). \end{cases} \tag{31}$$

Since  $\hat{\beta}_E^{\text{SLS}}$  satisfies the KKT conditions, given in Eq. (29), this yields to

$$\begin{cases} \hat{\rho}_E^{\lambda_1} \odot \left( \mathbf{1}_{|E|} - \frac{1}{\lambda_2} \hat{\mu} \right) \in \text{sign}(\hat{\beta}_E^{\text{SLS}}), \\ \hat{\rho}_{E^c}^{\lambda_1} + \frac{1}{\lambda_2 n} X_{E^c}^\top (y - X_E \hat{\beta}_E^{\text{SLS}}) \in \text{sign}(0), \end{cases} \tag{32}$$

where  $\mathbf{1}_{|E|} = (1, \dots, 1)^\top \in \mathbb{R}^{|E|}$ . We notice that the second line in Eq. (32) can be satisfied for  $\lambda_2$  large enough, since for each  $j \in E^c$  we have  $|\hat{\rho}_j^{\lambda_1}| < 1$  and  $\hat{\beta}_E^{\text{SLS}}$  does not depend on the value of  $\lambda_2$ . Denote  $\lambda$  the smallest  $\lambda_2$  such that the second condition in Eq. (32) is satisfied. Now, notice that the first line can be studied element-wise, hence we study two distinct cases, where we take  $j \in E$ .

- $\hat{\mu}_j = 0 \implies \hat{\rho}_j^{\lambda_1} \in \text{sign}(\hat{\beta}_E^{\text{SLS}})_j$ , which holds due to the definition of  $\hat{\beta}_E^{\text{SLS}}$ ,
- $\hat{\mu}_j > 0 \xrightarrow{(29)} (\hat{\beta}_E^{\text{SLS}})_j = 0 \implies$  if  $\lambda_2 \geq \hat{\mu}_j$  then  $|\hat{\rho}_j^{\lambda_1}(1 - \frac{\hat{\mu}_j}{\lambda_2})| \leq 1$  and Eq. (32) is satisfied.

Setting  $\lambda_0 = \lambda \vee \hat{\mu}_{\max}$ , where  $\hat{\mu}_{\max} = \max_{j=1}^{|E|} \{\hat{\mu}_j\}$  concludes the proof. □

Combining Theorems 5 and 6 we can state the following corollary, which provides an oracle inequality for the Sign-Least-Squares Lasso.

**Corollary 3.** *If the design matrix  $X \in \mathbb{R}^{n \times p}$  satisfies the restricted eigenvalue condition  $\text{RE}(3, s)$  and  $\lambda = 5\sigma\sqrt{2 \log(p/\delta)/n}$  for some  $\delta \in (0, 1)$ , then with probability  $1 - \delta$  the following bound holds:*

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 + \frac{1}{n} \|X(\beta^* - \hat{\beta})\|_2^2 \leq \frac{\lambda^2 s}{\kappa^2(1, s)} + \frac{9\lambda^2 s}{4\kappa^2(3, s)},$$

where  $\bar{\beta}$  is the Sign-Least-Squares Lasso solution associated with  $\hat{\beta} = \hat{\beta}^\lambda$ .

## 6. Experiments

To evaluate each scenario, we consider the following ‘oracle’ performance measures (in the sense that in practice one can not evaluate them):

- Prediction:  $\|X(\beta^* - \hat{\beta})\|_2^2$ ;
- Estimation:  $\|\beta^* - \hat{\beta}\|_1$ ;

**Table 1.** Description of considered estimators in the experiments. The subgradient  $\hat{\rho}^{\lambda_1}$ , associated with the Lasso solution, is evaluated according to Eq. (18) with  $\hat{\beta}^{\lambda_1} = \hat{\beta}^{\text{Lasso}}$ . For the Relaxed Lasso, the parameter  $\phi$  belongs to the interval (0, 1)

Lasso [20]	$\hat{\beta}^{\text{Lasso}} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \ y - X\beta\ _2^2 + \lambda_1 \ \beta\ _1$
LSLasso [1]	$\hat{\beta}^{\text{LSLasso}} \in \arg \min_{\beta: \text{supp}(\beta) \subset \text{supp}(\hat{\beta}^{\text{Lasso}})} \frac{1}{2n} \ y - X\beta\ _2^2$
SLSLasso [4]	$\hat{\beta}^{\text{SLSLasso}} \in \arg \min_{\beta: \beta \odot \hat{\rho}^{\lambda_1} \geq 0} \frac{1}{2n} \ y - X\beta\ _2^2$
Boosted Lasso	$\hat{\beta}^{\text{Boosted}} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \ y - X\beta\ _2^2 + \lambda_2 \ \beta - \hat{\beta}^{\text{Lasso}}\ _1$
Bregman Lasso	$\hat{\beta}^{\text{Bregman}} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \ y - X\beta\ _2^2 + \lambda_2 (\ \beta\ _1 - \langle \hat{\rho}^{\lambda_1}, \beta \rangle)$
Relaxed Lasso	$\hat{\beta}^{\text{Relaxed}} \in \arg \min_{\beta: \text{supp}(\beta) \subset \text{supp}(\hat{\beta}^{\text{Lasso}})} \frac{1}{2n} \ y - X\beta\ _2^2 + \phi \lambda_1 \ \beta\ _1$

- Predicted sparsity:  $|\text{supp}(\hat{\beta})|$ ;
- True Positive:  $|\{j \in [p] : \beta_j^* \neq 0 \text{ and } \hat{\beta}_j \neq 0\}|$ ;
- False Positive:  $|\{j \in [p] : \beta_j^* = 0 \text{ and } \hat{\beta}_j \neq 0\}|$ ;
- Hamming:  $\frac{1}{p} |\{j \in [p] : \text{False Positive or False Negative or False Sign}\}|$ ,  
 where  $\underbrace{\beta_j^* = 0 \text{ and } \hat{\beta}_j \neq 0}_{\text{False Positive}}, \underbrace{\beta_j^* \neq 0 \text{ and } \hat{\beta}_j = 0}_{\text{False Negative}}, \underbrace{\text{sign}(\beta_j^*) \neq \text{sign}(\hat{\beta}_j)}_{\text{False Sign}}$ .

All except one measures used in our evaluation are standard and are widely considered by various authors. We additionally study Hamming loss, which allows to capture information about miss predicting the sign of the underlying  $\beta^*$ . However, one should keep in mind that the introduced Hamming treats False Positive, False Negative and False Sign mistakes equally. Furthermore, the estimators that we consider are described in Table 1.

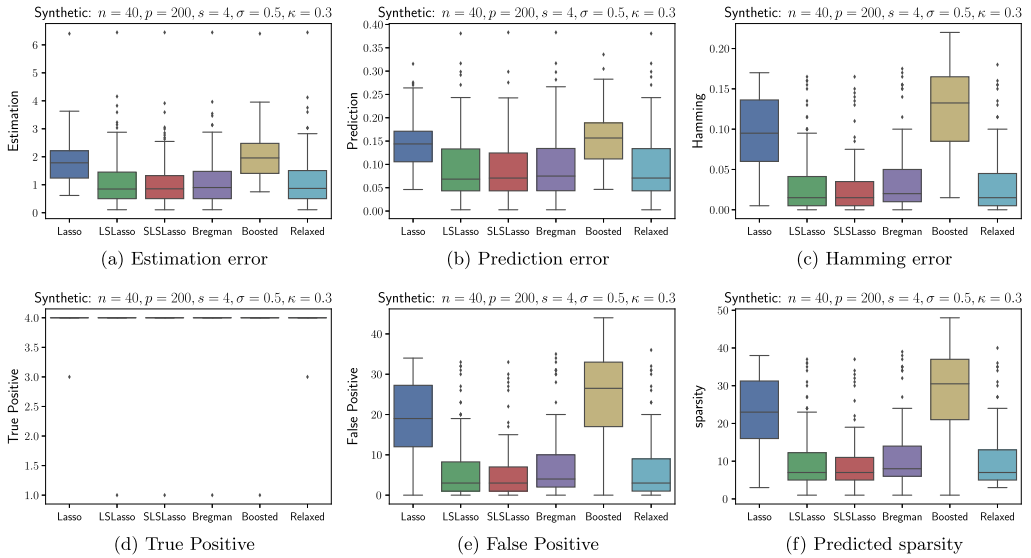
To report our results, we present boxplots for each scenario and performance measure over one hundred experiment replicas. During each simulation run we perform 3 fold cross-validation over a predefined one dimensional grid of 50 points, spread equally on logarithmic scale from  $0.01 \cdot \|X^\top y/n\|_\infty$  to  $\|X^\top y/n\|_\infty$  for  $\lambda_1$ , or two dimensional grid of  $50 \times 50$  points for  $\lambda_1, \lambda_2$  on the same scale. For the Relaxed Lasso the parameter  $\phi$  is chosen over a uniform grid of 50 points laying in the interval (0.001, 0.999). We finally select the tuning parameters achieving the best 3-fold cross-validation performance in terms of MSE, since in practice one does not have access to the underlying  $\beta^*$ .

### 6.1. Synthetic data

In our synthetic experiments, we generated the design matrix  $X \in \mathbb{R}^{n \times p}$  as follows

$$X_j = \sqrt{n} \frac{\kappa \zeta + (1 - \kappa) \xi_j}{\|\kappa \zeta + (1 - \kappa) \xi_j\|_2}, \tag{33}$$

where  $\zeta, \xi_1, \dots, \xi_p \sim \mathcal{N}(0, I_p)$  are independent standard normal vectors. The level of correlations between the covariates is determined by the parameter  $\kappa \in [0, 1]$ . We additionally consider



**Figure 4.** Synthetic dataset: low correlation scenario.

a noise vector  $\varepsilon \sim \mathcal{N}(0, I_n)$  and set the underlying vector as

$$\beta^* = (\underbrace{1, \dots, 1}_s, \underbrace{0, \dots, 0}_{p-s})^\top \in \mathbb{R}^p.$$

Hence, each scenario in the synthetic data section is described by the following parameters: number of observations:  $n$ , number of features:  $p$ , sparsity level of  $\beta^*$ :  $s$ , level of correlations:  $\kappa$ , noise level:  $\sigma$ . We fix  $n = 40, p = 200, s = 4, \sigma = 0.5$  and use the following values of correlations  $\kappa = 0.3; 0.5; 0.7$ , representing low, average and high correlation settings respectively. The results are reported on Figure 4, Figure 5 and Figure 6. We first notice that the Boosted Lasso does not give any significant improvement over the Lasso. Other four refitting strategies can improve the first step Lasso solution. In case of modest correlations inside the design matrix, the improvement can be considered as significant. Moreover, the Sign-Least-Squares Lasso outperforms the Least-Squares Lasso and the Bregman Lasso in average, in addition to the greater interpretability due to the sign preserving properties. The Relaxed Lasso performance is on par with the one of the Least-Squares Lasso. However, the former requires an additional tuning parameter to be calibrated.

### 6.2. Semi-real data

For our experimental study, we generate semi-real datasets, following an approach presented in [5]. The data are generated from the model Eq. (1), where the design matrix  $X$  is obtained from

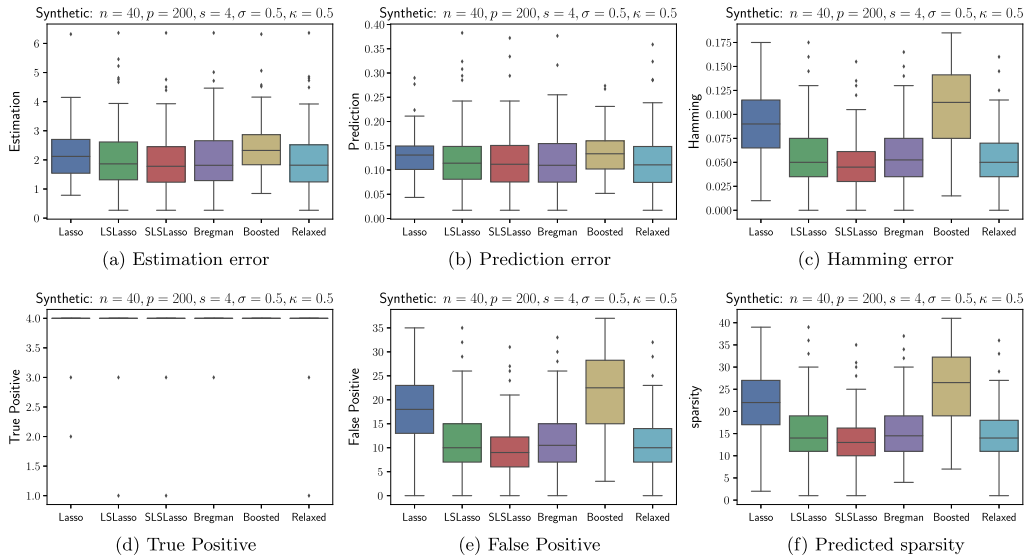


Figure 5. Synthetic dataset: average correlation scenario.

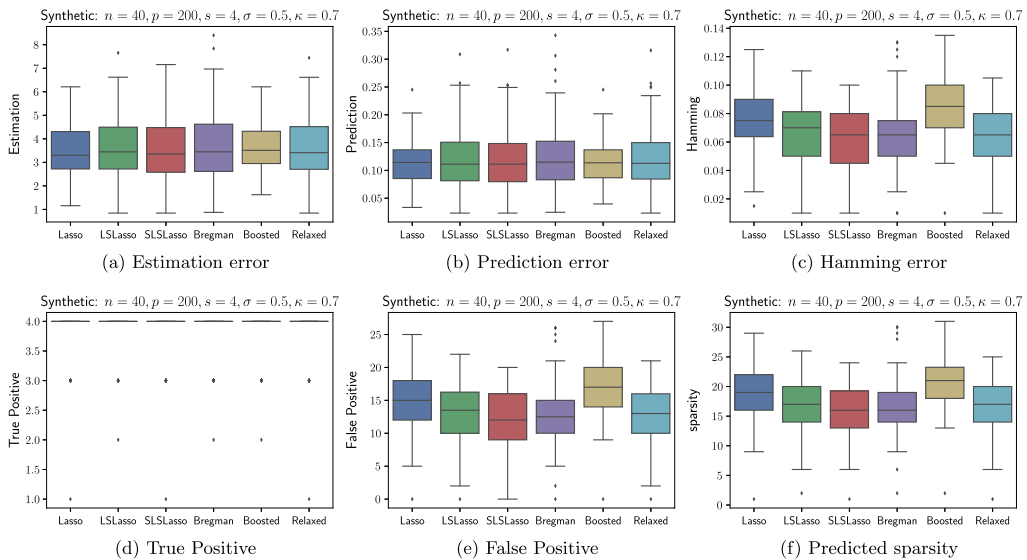


Figure 6. Synthetic dataset: high correlation scenario.

**Table 2.** Description of the parameters for semi-real dataset simulation

Settings		
Signal to noise ration (SNR)	2	8
Number of covariates ( $p$ )	200	1000
Underlying sparsity ( $s$ )	5	20
Correlations settings	Normal	High

the *leukemia* dataset with  $n = 72$ . We consider the following 4 parameters to describe the settings of our experiments:  $p$  – number of covariates,  $s$  – sparsity of the vector  $\beta^*$ , SNR – signal to noise ratio, and the correlation settings – way to generate support of the vector  $\beta^*$ . All the plots are averaged over hundred runs of the simulation process with fixed values of parameters given in Table 2. During each simulation round, we choose the first  $p$  columns from the *leukemia* dataset. Additionally we set the signal to noise ratio, defined as  $\text{SNR} = \frac{\|X\beta^*\|_2}{\sqrt{n\sigma^2}}$ , to control the noise level  $\sigma^2$  in the model Eq. (1). Moreover the true cardinality of  $\beta^*$  is set to  $s$  and each non-zero component of  $\beta^*$  is set to 1 or  $-1$  with equal probabilities. Finally, the support of the vector  $\beta^*$  is formed following two scenarios: normal correlations and high correlations. For normal correlations, we choose randomly  $s$  components out of  $p$  and for high correlations scenario, the first component is chosen randomly and the remaining  $s - 1$  having the highest Pearson correlation with the first one. Additional scenarios can be found in the supplemental article [7], in the main text we provide only three cases due to the space limitation.

On Figure 7 we notice that in the low noise (SNR = 8) and very sparse ( $s = 5$ ) scenario the overall performance is similar to the synthetic dataset, discussed above.

Meanwhile, the conclusion is different for Figure 8, where the noise level (SNR = 2) and the sparsity level ( $s = 20$ ) are high. In this scenario, we observe that simple Lasso outperforms in average all the refitting strategies in terms of estimation error and the TP rate, Boosted Lasso provides with better prediction rate and the Sign-Least-Squares Lasso and the Relaxed Lasso achieves better results in terms of all the other measures. We additionally emphasize, that the Least-Squares Lasso shows large variance and might fail to improve the estimation in some cases.

Similar conclusions as for previous scenario can be made for Figure 9, where the sparsity level ( $s = 20$ ) is still high, but the noise level is reduced (SNR = 8).

We conclude by pointing out that performing or not performing the refitting depends on the measure of interest and on the underlying (unknown in practice) scenario. Boosted Lasso agrees with our theoretical results, as it only improves the Lasso estimator in terms of prediction, but it fails to improve any other measure except True Positive rate (due to higher output sparsity). Least-Squares Lasso may show an undesirable performance in some scenarios and bring the problem of interpretability. Sign-Least-Squares Lasso and Bregman Lasso showed consistent and satisfying performance, however Sign-Least-Squares Lasso outperforms both the Bregman Lasso and the Least-Squares Lasso in average. We additionally emphasize that our results are valid for cross-validation on MSE. Other choices are possible (BIC, AIC,  $AV_p$ ) and may possibly provide with different overall conclusion.

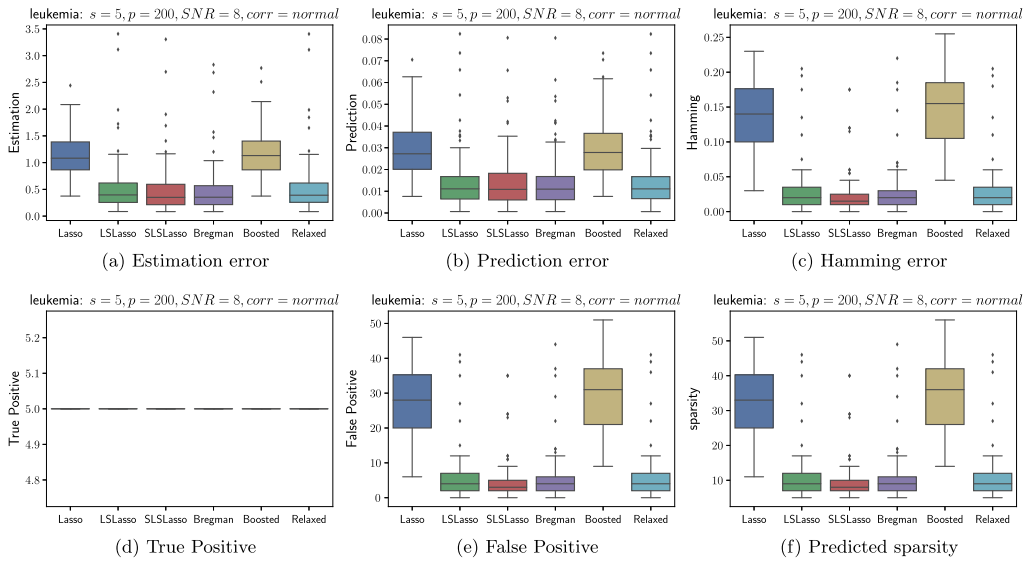


Figure 7. Leukemia dataset with  $p = 200$ ,  $SNR = 8$ ,  $s = 5$ , normal correlations scenario.

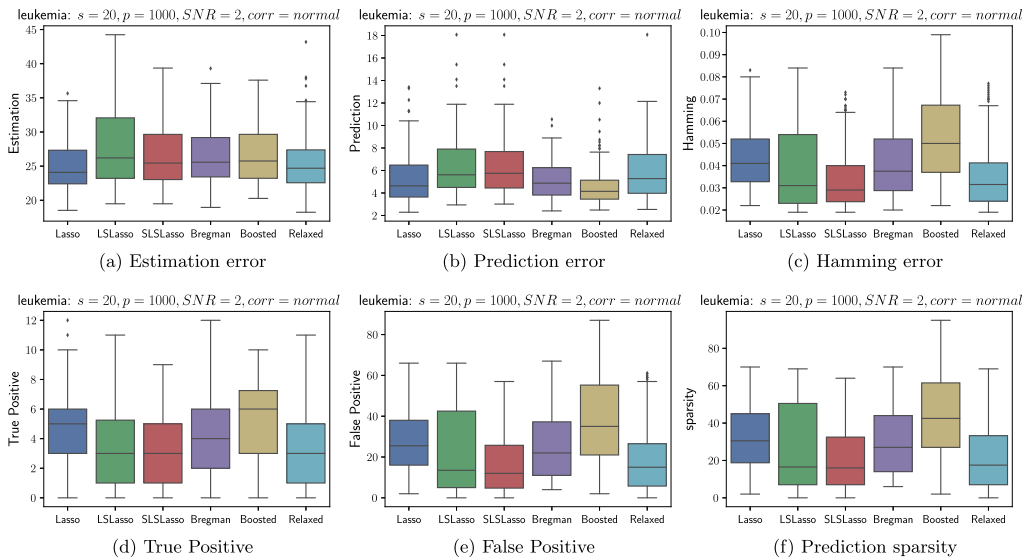
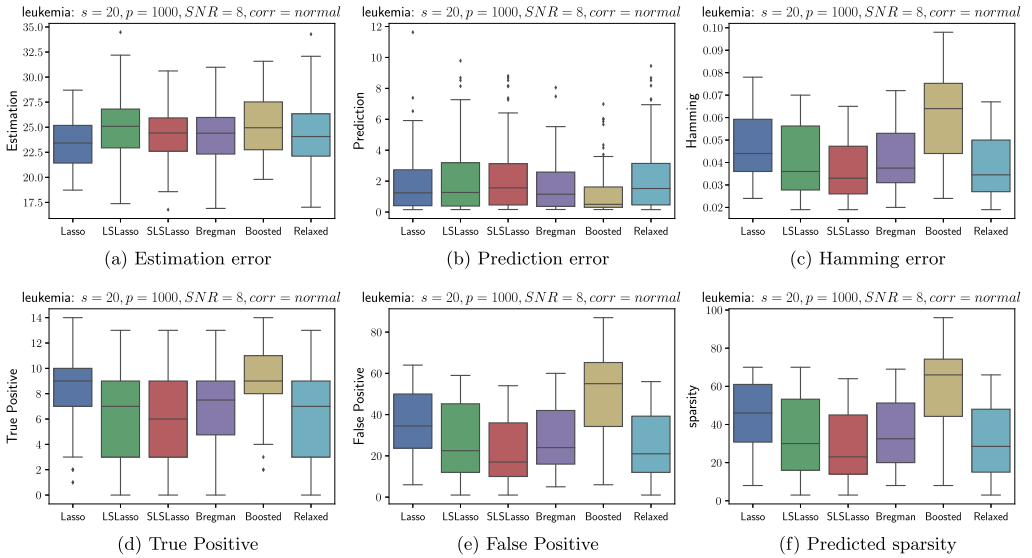


Figure 8. Leukemia dataset with  $p = 1000$ ,  $SNR = 2$ ,  $s = 20$ , normal correlations scenario.



**Figure 9.** *Leukemia* dataset with  $p = 1000$ ,  $SNR = 8$ ,  $s = 20$ , normal correlations scenario.

## Conclusion

In this article, we introduced a simple framework to provide additional statistical guarantees on the refitting and sign-consistent refitting of Lasso solutions. We demonstrated that every sign-consistent refitting strategy satisfies an oracle inequality under the same assumptions as the Lasso bounds. We theoretically analyzed two refitting strategies: Boosted Lasso and Bregman Lasso, which are easy to implement as they require only Lasso solver. It appeared that the Bregman Lasso converges to the Sign-Least-Squares Lasso, a particular refitting strategy with sign preserving properties. Experimental results show the advantages of sign-consistent strategies over the simple Least-Squares Lasso. Possible extension of this work is to consider other families of the refitting strategies by either taking into account an additional information provided by the Lasso or replacing the sign-preserving property. Another interesting road is to use our framework to provide oracle inequalities for estimation error and feature selection.

## Acknowledgments

This work was partially supported by “Laboratoire d’excellence Bézout of Université Paris Est” and by “Chair Machine Learning for Big Data at Télécom ParisTech”.

## Supplementary Material

**Supplement to “On Lasso refitting strategies”** (DOI: [10.3150/18-BEJ1085SUPP](https://doi.org/10.3150/18-BEJ1085SUPP); .pdf). We provide proofs and additional experimental results.



## References

- [1] Belloni, A. and Chernozhukov, V. (2013). Least squares after model selection in high-dimensional sparse models. *Bernoulli* **19** 521–547. [MR3037163](#)
- [2] Bickel, P.J., Ritov, Y. and Tsybakov, A.B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. [MR2533469](#)
- [3] Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge: Cambridge Univ. Press. [MR2061575](#)
- [4] Brinkmann, E.-M., Burger, M., Rasch, J. and Sutour, C. (2017). Bias reduction in variational regularization. *J. Math. Imaging Vision* **59** 534–566. [MR3712429](#)
- [5] Bühlmann, P. and Mandozzi, J. (2014). High-dimensional variable screening and bias in subsequent inference, with an empirical comparison. *Comput. Statist.* **29** 407–430. [MR3261821](#)
- [6] Bühlmann, P. and Yu, B. (2003). Boosting with the  $L_2$  loss: Regression and classification. *J. Amer. Statist. Assoc.* **98** 324–339. [MR1995709](#)
- [7] Chzhen, E., Hebiri, M. and Salmon, J. (2019). Supplement to “On Lasso refitting strategies.” DOI:[10.3150/18-BEJ1085SUPP](#).
- [8] Dalalyan, A.S., Hebiri, M. and Lederer, J. (2017). On the prediction performance of the Lasso. *Bernoulli* **23** 552–581. [MR3556784](#)
- [9] Deledalle, C.-A., Papadakis, N. and Salmon, J. (2015). On debiasing restoration algorithms: Applications to total-variation and nonlocal-means. In *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science* **9087** 129–141. Cham: Springer. [MR3394925](#)
- [10] Deledalle, C.-A., Papadakis, N., Salmon, J. and Vaiter, S. (2017). CLEAR: Covariant LEAst-square refitting with applications to image restoration. *SIAM J. Imaging Sci.* **10** 243–284. [MR3615463](#)
- [11] Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- [12] Gao, H.-Y. and Bruce, A.G. (1997). WaveShrink with firm shrinkage. *Statist. Sinica* **7** 855–874. [MR1488646](#)
- [13] Gasso, G., Rakotomamonjy, A. and Canu, S. (2009). Recovering sparse signals with a certain family of nonconvex penalties and DC programming. *IEEE Trans. Signal Process.* **57** 4686–4698. [MR2722328](#)
- [14] Giraud, C. (2014). *Introduction to High-Dimensional Statistics. Chapman and Hall/CRC Monographs on Statistics and Applied Probability*. Boca Raton: Chapman and Hall/CRC.
- [15] Koltchinskii, V. (2011). *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems. Lecture Notes in Math.* **2033**. Heidelberg: Springer. [MR2829871](#)
- [16] Lederer, J. (2013). Trust, but verify: Benefits and pitfalls of least-squares refitting in high dimensions. Preprint. Available at [arXiv:1306.0113](#).
- [17] Meinshausen, N. (2007). Relaxed Lasso. *Comput. Statist. Data Anal.* **52** 374–393. [MR2409990](#)
- [18] Osher, S., Burger, M., Goldfarb, D., Xu, J. and Yin, W. (2005). An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.* **4** 460–489. [MR2162864](#)
- [19] Osher, S., Ruan, F., Xiong, J., Yao, Y. and Yin, W. (2016). Sparse recovery via differential inclusions. *Appl. Comput. Harmon. Anal.* **41** 436–469. [MR3534446](#)
- [20] Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- [21] Tibshirani, R.J. (2013). The lasso problem and uniqueness. *Electron. J. Stat.* **7** 1456–1490. [MR3066375](#)
- [22] Tukey, J.W. (1977). *Exploratory Data Analysis*. Reading: Addison-Wesley.
- [23] Xu, J. and Osher, S. (2007). Iterative regularization and nonlinear inverse scale space applied to wavelet-based denoising. *IEEE Trans. Image Process.* **16** 534–544. [MR2462742](#)

- [24] Yin, W., Osher, S., Goldfarb, D. and Darbon, J. (2008). Bregman iterative algorithms for  $l_1$ -minimization with applications to compressed sensing. *SIAM J. Imaging Sci.* **1** 143–168. [MR2475828](#)
- [25] Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. [MR2604701](#)

*Received July 2017 and revised June 2018*