

Robust penalized least squares of depth trimmed regression for high-dimensional data

YIJUN ZUO

Department of Statistics and Probability

Michigan State University, East Lansing, MI 48824, USA

zuo@msu.edu

December 24, 2023

Abstract

Challenges with data in the big-data era include (i) the dimension p is often larger than the sample size n (ii) outliers or contaminated points are frequently hidden and more difficult to detect. Challenge (i) renders most conventional methods inapplicable. It, thus, attracts tremendous attention from statistics, computer science, industry, and bio-medical communities. Numerous penalized regression methods have been introduced as modern methods for analyzing high-dimensional data. Disproportionate attention has been paid to the challenge (ii) though. Penalized regression methods can do their job very well and are expected to handle the challenge (ii) simultaneously. The fact is most of them can break down by a single outlier (or single adversary contaminated point) as revealed in this article. The latter systematically examines a large class of penalized regression methods in terms of their robustness and provides quantitative assessment and reveals that most of leading penalized regression methods can break down by a single outlier. Consequently, a novel robust penalized regression method based on the least squares of depth trimmed residuals is proposed and studied carefully. Experiments with simulated and real data reveal that the newly proposed method can outperform some leading competitors in terms of estimation and prediction accuracy in the cases considered.

AMS 2000 Classification: Primary 62J05, 62G36; Secondary 62J99, 62G99

Key words and phrase: penalized regression, robustness of regularized regression estimators, least squares of depth trimmed residuals, finite sample prediction error bound, approximate computation algorithms.

Running title: robust penalized depth trimmed regression

1 Introduction

Least squares regression, the classic approach Consider the classic linear regression model,

$$y_i = (1, \mathbf{x}_i')\boldsymbol{\beta}_0 + e_i := \mathbf{w}_i'\boldsymbol{\beta}_0 + e_i, \quad (1)$$

where random variables y_i and $e_i \in \mathbb{R}$, random vector $\mathbf{x}_i \in \mathbb{R}^{p-1}$ ($p \geq 2$), and $\boldsymbol{\beta}_0 \in \mathbb{R}^p$ is an unknown parameter of interest, ' stands for the transpose. One wants to estimate the $\boldsymbol{\beta}_0$ based on a given sample $\mathbf{z}^{(n)} := \{(\mathbf{x}_i', y_i)', i \in \{1, \dots, n\}\}$ from a parent model $y = \mathbf{w}'\boldsymbol{\beta}_0 + e$.

Call the difference between y_i (observed value) and $\mathbf{w}_i'\boldsymbol{\beta}$ (predicted value), r_i , for a candidate coefficient vector $\boldsymbol{\beta}$ (which is often suppressed).

$$r_i := r_i(\boldsymbol{\beta}) = y_i - (1, \mathbf{x}_i')\boldsymbol{\beta} := y_i - \mathbf{w}_i'\boldsymbol{\beta}. \quad (2)$$

To estimate $\boldsymbol{\beta}_0$, the classic *least squares* (LS) estimator is the minimizer of the sum of the squared residuals (SSR): $\hat{\boldsymbol{\beta}}_{ls} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \sum_{i=1}^n r_i^2$. Alternatively, one can replace the square by the absolute value to obtain the *least absolute deviations* (lad) estimator (aka, L_1 , in contrast to the L_2 (LS), estimator). A straightforward derivation leads to

$$\hat{\boldsymbol{\beta}}_{ls} = (\mathbf{X}_n' \mathbf{X}_n)^{-1} \mathbf{X}_n' \mathbf{Y}_n. \quad (3)$$

where $\mathbf{Y}_n = (y_1, \dots, y_n)'$, $\mathbf{X}_n = (\mathbf{w}_1, \dots, \mathbf{w}_n)'$ and the columns of \mathbf{X}_n are assumed to be linearly independent (i.e. \mathbf{X}_n has a full rank p ($\leq n$)).

The LS estimator is popular in practice across a broader spectrum of disciplines due to its (i) great computability (with the computation formula); and (ii) optimal properties (the best linear unbiased estimator (BLUE) and the uniformly minimum variance unbiased estimator (UMVUE), page 186 of [46] when the i.i.d. error e_i follows a normal $\mathcal{N}(0, \sigma^2)$).

It, however, can behave badly when the error distribution is slightly departed from the normal distribution, particularly when the errors are heavy-tailed or contain outliers.

Penalized regression, the state of the art In modern applied data analysis, the number of variables often is even larger than the number of observations. Traditional methods such as the LS can then no longer be applied due to the design matrix \mathbf{X} being less than p rank ($n < p$), hence LS estimator is no longer unique and its variance is large if \mathbf{X} is close to collinear. Furthermore, models that include the full set of explanatory variables often have poor prediction performance as they tend to have large variance while large models are in general difficult to interpret. *Ridge regression*, minimizing SSR, subject to a constraint $\sum_{i=1}^p |\beta_i|^2 < t$

$$\hat{\boldsymbol{\beta}}_{ridge}(\lambda) := \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \left\{ \sum_{i=1}^n r_i^2 + \lambda \sum_{i=1}^p |\beta_i|^2 \right\}, \quad (4)$$

first proposed in [23, 24], is a useful tool for improving prediction in regression situations with highly correlated predictors and tackling the non-inverse issue,

$$\hat{\boldsymbol{\beta}}_{ridge}(\lambda) = (\mathbf{X}_n' \mathbf{X}_n + \lambda \mathbf{I}_{d \times d})^{-1} \mathbf{X}_n' \mathbf{Y}_n, \quad (5)$$

its variance is smaller than that of the LS estimator. Therefore, better estimation can be achieved on the average in terms of mean squared error (MSE) with a little sacrifice of bias, and predictions can be improved overall.

Bridge regression, introduced in [15], generalized the ridge regression, minimizes SSR subject to a constraint $\sum_{i=1}^p |\beta_i|^\gamma \leq t$ with $\gamma \geq 0$,

$$\hat{\beta}_{bridge}(\lambda, \gamma) := \arg \min_{\beta \in \mathbb{R}^p} \left\{ \sum_{i=1}^n r_i^2 + \lambda \sum_{i=1}^p |\beta_i|^\gamma \right\}. \quad (6)$$

Ridge regression ($\gamma = 2$) and subset selection ($\gamma = 0$) are special cases.

Least absolute shrinkage and selection operator (lasso) was introduced in [51], minimizing SSR subject to a constraint $\sum_{j=1}^p |\beta_j| \leq t$, is a special case of the bridge with $\gamma = 1$. As pointed out in [51], the lasso shrinks the LS estimator $\hat{\beta}_{ls}$ towards 0 and potentially sets $\hat{\beta}_j = 0$ for some j . That is, it performs as a variable selection operator simultaneously.

Other remarkable approaches of regularized regression for high dimensional sparse data include, among others, **(i)** *elastic nets*, introduced in [59], a generalization of the ridge and lasso models, which combines the two penalties and yields

$$\hat{\beta}_{enet}(\lambda_1, \lambda_2) := \arg \min_{\beta \in \mathbb{R}^p} \left\{ \sum_{i=1}^n r_i^2 + \lambda_1 \sum_{i=1}^p |\beta_i| + \lambda_2 \sum_{i=1}^p \beta_i^2 \right\}. \quad (7)$$

(ii) *square-root lasso*, introduced in [4], to avoid the pre-estimation of standard deviation σ of the error term in lasso and achieve a better performance, is defined as

$$\hat{\beta}_{sqr-lasso} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \left(\sum_{i=1}^n r_i^2 \right)^{1/2} + \lambda \sum_{i=1}^p |\beta_i| \right\}. \quad (8)$$

(iii) *slope* (Sorted L-One Penalized Estimation) introduced in [6], aim to control the false discover rate (FDR),

$$\hat{\beta}_{slope} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \sum_{i=1}^n r_i^2 + \sum_{i=1}^p \lambda_i |\beta_{(i)}| \right\}, \quad (9)$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0$ and $|\beta_{(1)}| \geq |\beta_{(2)}| \geq \dots \geq |\beta_{(p)}|$.

Strong connections between some modern methods and a method called *least angle regression* (lar) was revealed in [11] where they developed an algorithmic framework that includes all of these methods (lasso, boosting, forward stagewise regression) and provided a fast implementation, for which they used the term ‘lars’. Lars is a promising technique/algorithm for variable selection applications, offering a nice alternative to stepwise regression. For an excellent review on lars and lasso, see [22].

Other outstanding penalized regression estimators include, among others, SCAD [13], [14] and MCP [57]. It is not our goal to review all existing penalized/regularized regression estimators in the literature. For a detailed account about lasso and its variants, refer to Table 6 of [16] or Fig. 1 of [54], and [58] and references therein.

The penalized regression estimators above improve prediction accuracy meanwhile enhance the interpretability of the model. They, however, pay the price of inducing a little bit of bias in addition to the lack of robustness. There are numerous published articles related to

lasso and regularized regression in the literature. However, there are disproportionately few addressing the robustness of the estimators. Are they as robust as supposed (or expected)? Or rather can they resist the influence of just a single contaminated point (or outlier) that is typically buried in high-dimensional data?

Robust versions of the lasso (or ridge) estimators have been sporadically considered in the literature. The LS in lasso (or ridge), is replaced by M-estimators, as in [55] and [32]; or replaced by a Huber-type loss function, as in [37] and [49]; or by lads, as in [53];

$$\hat{\beta}_{lad-lasso} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \sum_{i=1}^n |r_i| + \lambda \sum_{i=1}^p |\beta_i| \right\}, \quad (10)$$

or replacing correlations in lars by a robust type of correlation, as in [27] (Rlars); or by S- ([41]) and MM- ([56]) estimators, as in [34] for ridge regression (Rrr); or by the least trimmed squares (LTS) ([40]), as in [2]. The LTS is defined as

$$\hat{\beta}_{lts}^n := \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^h (r^2)_{i:n}, \quad (11)$$

where $(r^2)_{1:n} \leq (r^2)_{2:n} \leq \dots, (r^2)_{n:n}$ are the ordered squared residuals, $\lceil n/2 \rceil \leq h \leq n$, and $\lceil \cdot \rceil$ is the ceiling function. The authors of [2] replaced SSR by the loss of LTS and defined

$$\hat{\beta}_{lts-lasso} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \sum_{i=1}^h (r^2)_{i:n} + h\lambda \sum_{i=1}^p |\beta_i| \right\}, \quad (12)$$

The idea of [2] has extended to logistic regression with elastic net penalty in [29], and penalized weighted M-type estimators for the logistic regression have also been studied in [5].

Most estimators above (except Rlars, Rrr, and $\hat{\beta}_{lts-lasso}$), like both L_1 and L_2 (LS) estimators, unfortunately, have the worst case 0% asymptotic breakdown point (i.e., one bad point can ruin (break down) the estimator), in sharp contrast to the 50% of the least squares of trimmed (LST) estimator (see Section 3.1 of [65] or Section 3 here). [29] and [5] both assert their estimators are robust, but no qualitative robustness assessment of their estimators has been established yet. The same situation with the estimator in [45].

Now let us take a close look at the three exceptions above. The main drawback of the Rlars is the lack of a natural definition or a clear objective function, as commented in [2]. The main focus of [34] is robustifying ridge regression (Rrr).

Only $\hat{\beta}_{lts-lasso}$ in [2] has an *established* high finite sample breakdown point (see Section 3 for definition). Their result, though covers the lasso-type estimators, does not cover the elastic nets and other estimators; the authors failed to (i) explain why their estimator can have a breakdown point higher than 50% and (ii) study the properties (such as equivariance and consistency) of their estimator. Furthermore, the LTS is notorious for its inefficiency (i.e., usually has a large variance). On the other hand, the LST introduced in [65] is as robust as LTS and can outperform the LTS (especially in efficiency) as demonstrated in [65].

Based on observations above, questions we want to address are: **(i)** Can one replace the loss of LS with that of the LST in the penalized regression? How does the resulting estimator perform? Is it more robust, compared with existing ones? **(ii)** Can one provide a more general breakdown robustness assessment that covers a large general class of regularized regression estimators and explain why a breakdown point higher than 50% can be **(iii)** Besides robustness, what are other desirable properties for a regression estimator?

The main contributions of this article include **(i)** it examines the robustness for a large class of (penalized) regression estimators (including all leading ones, e.g. the L_1 and L_2 , the ridge, the bridge, the lasso, the elastic net, the slope, the lad-lasso, and the sqrt-lasso) and proves that they can break down by a single adversary contaminating point; **(ii)** it, hence, introduces a novel and robust penalized least squares of depth trimmed estimator ($\hat{\beta}_{lst-enet}$) that has a much higher breakdown point robustness and outperforms leading competitors; and **(iii)** it proposes an efficient computational algorithm for the estimator and vets its performance for simulated and real high-dimensional data and reveals that $\hat{\beta}_{lst-enet}$ can outperform its competitors in most cases considered; **(iv)** it discusses three desired equivariance properties for regression estimators.

Section 2 establishes a robust result for a large class of general regularized regression estimators and reveals that most of leading estimators (including bridge, lasso, slope, and enet) has the worst breakdown point robustness. Section 3 invokes the least squares of depth trimmed (LST) procedure introduced in [65] and reviews its robust property meanwhile introduces a class of penalized regression estimators based on the LST and studies their properties including, existence and uniqueness, robustness, and equivariance. Section 4 is devoted to the establishment of the finite sample prediction error bound and estimator consistency. Section 5 addresses the computation issue of $\hat{\beta}_{lst-enet}$. Section 6 consists of simulation/comparison study and real data application of five competing methods. Section 7 ends the article with some concluding discussions. Proofs of main results are deferred to an Appendix.

2 Robustness of a class of penalized regression estimators

Are the existing numerous penalized regression methods mentioned above as robust as they are expected or believed to be? Or rather can they resist the influence of just a single outlier (or adversary single-point contamination)? We now formally address this issue.

2.1 A robustness measure

In the finite sample practice, the most prevailing quantitative measure of the robustness of any regression or location estimators is the *finite sample breakdown point*, introduced in [10].

Definition [10] The finite sample *replacement breakdown point* (RBP) of a regression estimator \mathbf{t} at a given un-contaminated sample $\mathbf{z}^{(n)} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n\}$, where $\mathbf{z}_i := (\mathbf{x}'_i, y_i)'$, is

$$\text{RBP}(\mathbf{t}, \mathbf{z}^{(n)}) = \min_{1 \leq m \leq n} \left\{ \frac{m}{n} : \sup_{\mathbf{z}_m^{(n)}} \|\mathbf{t}(\mathbf{z}_m^{(n)}) - \mathbf{t}(\mathbf{z}^{(n)})\|_2 = \infty \right\}, \quad (13)$$

where $\mathbf{z}_m^{(n)}$ stands for an arbitrary contaminated sample by replacing m original sample points in $\mathbf{z}^{(n)}$ with arbitrary points in \mathbb{R}^p and $\|\mathbf{x}\|_q = (\sum_{i=1}^n x_i^q)^{1/q}$ is the ℓ_q -norm for vector $\mathbf{x} \in \mathbb{R}^n$. \square

Namely, the RBP of an estimator is the minimum replacement fraction that could drive the estimator beyond any bounds. It turns out that both L_1 (least absolute deviations) and L_2 (least squares) estimators have RBP $1/n$ (or 0%) whereas LST can have $(\lfloor n/2 \rfloor - p + 2)/n$ (or 50%) (see Section 3), the highest possible asymptotic value for any regression equivariant estimators (see pages 124-125 of Rousseeuw and Leroy (1987) [42] or Section 3), where $\lfloor \cdot \rfloor$ is the floor function. We now present a general RBP result for a large class of penalized regression estimators.

2.2 A general result on penalized regression estimators

Theorem 2.1 For any given data set $\mathbf{z}^{(n)} = \{(\mathbf{x}'_i, y_i)', i \in \{1, \dots, n\}\}$ in \mathbb{R}^p ($p \geq 2$), let $\hat{\boldsymbol{\beta}}^*(\lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)})$ be the penalized regression estimator, which minimizes the objective

$$O(\boldsymbol{\beta}, \lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)}) := f\left(\frac{1}{n} \sum_{i=1}^n \mathcal{L}(r_i)\right) + g(\boldsymbol{\beta}, \lambda_1, \lambda_2, \gamma), \quad (14)$$

where $\lambda_i, \gamma \geq 0$, the combined penalty function $g(\boldsymbol{\beta}, \lambda_1, \lambda_2, \gamma) \geq 0$ and $f(x)$ is continuous, strictly increasing over $[0, \infty)$, $f(0) = 0$ and $f(x) \rightarrow \infty$ as $x \rightarrow \infty$, the loss function $\mathcal{L}(x)$ is non-negative, non-decreasing over $[0, \infty)$, $\mathcal{L}(0) = 0$ and $\mathcal{L}(x) \rightarrow \infty$ when $x \rightarrow \infty$. Then

$$\text{RBP}(\hat{\boldsymbol{\beta}}^*(\lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)}), \mathbf{z}^{(n)}) = 1/n.$$

Proof: see the Appendix. \square

Remarks 2.1

(i) Conditions on $\mathcal{L}(x)$ are relative loose, they hold automatically if $\mathcal{L}(x)$ is non-negative, non-decreasing, and convex in x and $\mathcal{L}(0) = 0$. The $\mathcal{L}(x)$ in theorem covers almost all loss functions (with two exceptions out of fifteen) in Table 6 of [16]. The penalty function $g(\boldsymbol{\beta}, \lambda_1, \lambda_2, \gamma)$ covers almost all existing ones including, among others, $\lambda_1 \|\mathbf{D}_1 \boldsymbol{\beta}\|_\gamma + \lambda_2 \|\mathbf{D}_2 \boldsymbol{\beta}\|_2^2$, with \mathbf{D}_i being penalty matrices. Typical f includes the identity or square-root function (covering the square-root lasso [4]). To the best of our knowledge, a similar general result, judging the robustness of a large class of penalized regression estimators, never appeared in the literature before.

(ii) The RBP result in the theorem is very general since the loss function covers most of the existing loss functions in the machine learning and AI literature, e.g., the most popular ones: negative log-likelihood; the ℓ_1 , the ℓ_2 , or any ℓ_q ($q \geq 1$) loss, Huber loss, and the loss of the lasso and most of its variants (see Table 6 of [16]). The penalty format covers most of the existing ones in the literature (indeed, it covers all twenty-five penalty functions listed in Table 1 of [54]).

(iii) The great generality of the result in the theorem implies that most of the existing penalized regression (and the classic L_1 and L_2) estimators are not robust. In fact, they can break down with just one single outlier which often buries in high dimensional data. \square

Now that most of the existing penalized regression estimators can be broken down by a single outlier (or single-point contamination). Furthermore, existing robust penalized regression estimators are most ad hoc, e.g., Rlars of [27] is for robustifying lars, and Rrr of [34] is for robustifying ridge regression, and [5] is mainly for robustifying the penalized logistic regression estimators. Only $\hat{\beta}_{lts-lasso}$ of [2] and $\hat{\beta}_{enetLTS}$ of [29] that employed the loss of LTS to replace the SSR in lasso have really high breakdown robustness meanwhile do the variable selection job. But the major drawback of the LTS is its inefficiency (it has a larger variance) as demonstrated in [65] (also Sections 3 and 6) and Figures 2 and 3 of [29].

A natural question is: can one construct a penalized regression estimator that is robust against the outliers or contamination meanwhile is more efficient (i.e., with a smaller variance than that of the LTS)? In the sequel, we achieve this goal by adopting the loss of least squares of depth trimmed (LST) introduced in [65] and applying it to the penalized regression setting.

3 A class of penalized regression estimators based on the LST

3.1 Least Squares of depth trimmed (LST) estimator

Definition 3.1 To robustify the LS, least trimmed squares (LTS) estimator was introduced in [40]. The procedure orders the squared residuals and then trims the larger ones and keeps at least $h \geq \lceil n/2 \rceil$ squared residuals, the minimizer of the sum of those *trimmed squared residuals* is called an LTS estimator as defined in (11). $\hat{\beta}_{lts}^n$ is highly robust but is not very efficient, as reported in [35] (page 132) or in [48] having just 7% or 8% asymptotic efficiency. A more efficient competitor, least squares of trimmed (LST) estimator, is introduced in [65], overcoming LTS drawback while sharing its high robustness and fast computation advantages.

Depth (or outlyingness)-based trimming scheme trims points that lie on the outskirts (i.e. points that are less deep, or outlying). The outlyingness (or, equivalently, depth) of a point \mathbf{x} is defined to be (strictly speaking, $\text{depth} = 1/(1 + \text{outlyingness})$) in Zuo (2003) and Zuo and Serfling (2000))

$$O(\mathbf{x}, \mathbf{x}^{(n)}) = |\mathbf{x} - \text{Med}(\mathbf{x}^{(n)})| / \text{MAD}(\mathbf{x}^{(n)}), \quad (15)$$

where $\mathbf{x}^{(n)} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is a data set in \mathbb{R}^1 , $\text{Med}(\mathbf{x}^{(n)}) = \text{median}(\mathbf{x}^{(n)})$ is the median of the data points, and $\text{MAD}(\mathbf{x}^{(n)}) = \text{Med}(\{|\mathbf{x}_i - \text{Med}(\mathbf{x}^{(n)})|, i \in \{1, 2, \dots, n\}\})$ is the median of absolute deviations to the center (median). It is readily seen that $O(\mathbf{x}, \mathbf{x}^{(n)})$ is a generalized standard deviation, or equivalent to the one-dimensional projection depth (see Zuo and Serfling (2000) and Zuo (2003, 2006) for a high dimensional version). For the notion of outlyingness, cf. Stahel (1981), Donoho (1982), and Donoho and Gasko (1992).

For a given constant α (hereafter assume $\alpha \geq 1$), β , and $\mathbf{z}^{(n)}$, define a set of indexes

$$I(\beta) = \left\{ i : O(\mathbf{r}_i, \mathbf{r}^{(n)}) \leq \alpha, i \in \{1, \dots, n\} \right\}. \quad (16)$$

where $\mathbf{r}^{(n)} = \{r_1, r_2, \dots, r_n\}$ and r_i is defined in (2). Namely, the set of subscripts so that the outlyingness (see (15)) (or depth) of the corresponding residuals are no greater (or less) than α (or $1/(1 + \alpha)$). It depends on $\mathbf{z}^{(n)}$ and α , which are suppressed in the notation. For a fixed constant α in the depth trimming scheme, consider the quantity

$$Q^n(\beta) := Q(\mathbf{z}^{(n)}, \beta, \alpha) := \sum_{i=1}^n r_i^2 \mathbb{1} \left(O(r_i, \mathbf{r}^{(n)}) \leq \alpha \right) = \sum_{i \in I(\beta)} r_i^2, \quad (17)$$

where $\mathbb{1}(A)$ is the indicator of A (i.e., it is one if A holds and zero otherwise). Namely, residuals with their outlyingness (or equivalently reciprocal of depth minus one) greater than α will be trimmed. When there is a majority ($\geq \lfloor (n+1)/2 \rfloor$) identical r_i s, we define $\text{MAD}(\mathbf{r}^{(n)}) = 1$ (since those r_i lie in the deepest position (or are the least outlying points)).

Minimizing $Q(\mathbf{z}^{(n)}, \beta, \alpha)$, one gets the *least sum of squares of trimmed* (LST) estimator,

$$\hat{\beta}_{lst}^n := \hat{\beta}_{lst}(\mathbf{z}^{(n)}, \alpha) = \arg \min_{\beta \in \mathbb{R}^p} Q(\mathbf{z}^{(n)}, \beta, \alpha). \quad (18)$$

Differences between LST and LTS Compared with the LTS defined in (11), it is readily seen that both estimators trim residuals. However, there are two essential differences: **(i)** the trimming schemes are different. The LTS employs a rank-based trimming scheme that focuses only on the relative position of points (squared residuals) with respect to others and ignores the magnitude of the point and the relative distance between points whereas the LST exactly catches these two important attributes. It orders data (residuals) from a center (the median) outward and trims the points (residuals) that are far away from the center. This is known as depth-based trimming. **(ii)** Besides the trimming scheme difference, there is another difference between the LTS and the LST, that is, the order of trimming and squaring. In the LTS, squaring is first, followed by trimming whereas, in the LST, the order is reversed.

All the differences lead to an unexpected performance difference in the LTS and the LST as demonstrated in the illustrating Figure 1 (see Ex 1.1 of [65]) and Table 1.

Existence and uniqueness of $\hat{\beta}_{lst}^n$ have been addressed in [65], it is also equivariant (see [65]).

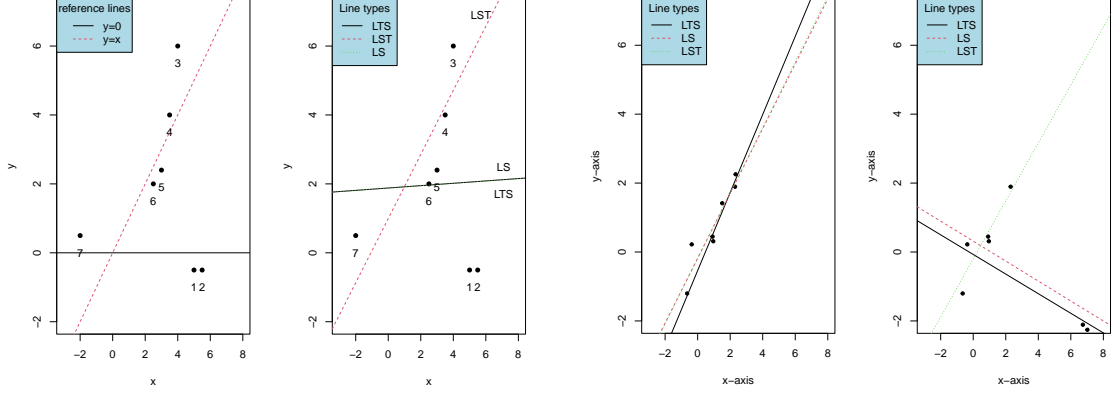
Equivariance A regression estimator \mathbf{t} is called *regression, scale, and affine equivariant* if, respectively (see page 116 of [42]) with $N = \{1, 2, \dots, n\}$

$$\begin{aligned} \mathbf{t}(\{(\mathbf{w}'_i, y_i + \mathbf{w}'_i \mathbf{b})', i \in N\}) &= \mathbf{t}(\{(\mathbf{w}'_i, y_i)', i \in N\}) + \mathbf{b}, \quad \forall \mathbf{b} \in \mathbb{R}^p \\ \mathbf{t}(\{(\mathbf{w}'_i, s y_i)', i \in N\}) &= s \mathbf{t}(\{(\mathbf{w}'_i, y_i)', i \in N\}), \quad \forall s \in \mathbb{R}^1 \\ \mathbf{t}(\{(A' \mathbf{w}_i)', y_i), i \in N\}) &= A^{-1} \mathbf{t}(\{(\mathbf{w}'_i, y_i)', i \in N\}), \quad \forall \text{ nonsingular } A \in \mathbb{R}^{p \times p}. \end{aligned}$$

Now with the measure of robustness (presented in the last section), naturally one wants to ask the question: is $\hat{\beta}_{lst}^n$ theoretically more robust than the LS estimator $\hat{\beta}_{ls}^n$?

Robustness and efficiency of LST

We shall say $\mathbf{z}^{(n)}$ is *in general position* when any p of observations in $\mathbf{z}^{(n)}$ gives a unique determination of β . In other words, any $(p-1)$ dimensional subspace of the space $(\mathbf{x}', y)'$



(a) Left panel: plot of seven artificial points and two reference lines ($y = 0$ and $y = x$), which line would you pick? Sheerly based on the trimming scheme and objective function value, if one uses the number $h = \lfloor n/2 \rfloor + \lfloor (p+1)/2 \rfloor$ given on page 132 of [42], that is, employing four smallest squared residuals, then the LTS prefers $y = 0$ to $y = x$ whereas the LST reverses the preference. Right panel: the same seven points are fitted by the LTS, the LST, and the LS (benchmark). A solid black line is the LTS given by `ltsReg`. Red dashed line is given by the LST, and green dotted line is given by the LS - which is identical to the LTS line in this case.

(b) Left panel: plot of seven highly correlated normal points (with mean being the zero vector and covariance matrix with diagonal entries being one and off-diagonal entries being 0.88) and three lines given by the LST, the LTS, and the LS. The LS line is identical to the LTS line again. Right panel: the LTS line (solid black) and the LST line (dashed red), and the LS (dotted green) for the same seven highly correlated normal points but with two points contaminated nevertheless. The LS line is parallel to the LTS line due to the attributes in the R function `ltsReg` that is based on [43]).

Figure 1: (a) Performance difference between the LST and the LTS. (b) Performance difference between the LST and the LTS when there are contaminated points (x -axis leverage points).

contains at most p observations of $\mathbf{z}^{(n)}$. When the observations come from continuous distributions, the event ($\mathbf{z}^{(n)}$ being in general position) happens with probability one.

Lemma 3.1 [65] For $\hat{\beta}_{lst}^n$ defined in (18) with $(\alpha \geq 1)$ and $\mathbf{z}^{(n)}$ in general position, we have

$$\text{RBP}(\hat{\beta}_{lst}^n, \mathbf{z}^{(n)}) = \begin{cases} \lfloor (n+1)/2 \rfloor / n, & \text{if } p = 1, \\ (\lfloor n/2 \rfloor - p + 2) / n, & \text{if } p > 1. \end{cases} \quad (19)$$

□

Namely, the LST shares the best 50% asymptotic breakdown value of the LTS meanwhile it is much more efficient than the LTS as demonstrated in Table 1 below (also see [65]).

Here we generate $R = 1000$ samples $\{\mathbf{z}_i = (\mathbf{x}'_i, y_i)', i \in \{1, \dots, n\}\}$ with various ns from the normal distribution, where mean being a zero-vector in \mathbb{R}^{20} , and covariance being a 20 by 20 matrix with diagonal entries being 1 and off-diagonal entries being 0.9. Then $\varepsilon\%$ of each sample are contaminated by $m = \lceil n\varepsilon \rceil$ points, we select m points of $\{\mathbf{z}_i, i \in \{1, \dots, n\}\}$

randomly and replace them by $(7, 7, \dots, 7, -7)'$ then apply the three methods, the LTS, the LST, and the LS, to the contaminated data sets, we calculate the empirical mean squared error (EMSE) for estimator \mathbf{t} , $\text{EMSE} := \sum_{i=1}^R \|\mathbf{t}_i - \boldsymbol{\beta}_0\|^2 / R$. Here \mathbf{t}_i is the realization of \mathbf{t} obtained from the i th sample with size n and dimension $p = 20$. If \mathbf{t} is regression equivariant, then we can assume (w.l.o.g.) that the true parameter $\boldsymbol{\beta}_0 = \mathbf{0} \in \mathbb{R}^p$. Meanwhile, we also compute, $\text{SVAR}(\mathbf{t}) := \sum_{i=1}^R \|\mathbf{t}_i - \bar{\mathbf{t}}\|^2 / (R-1)$, the sample variance of \mathbf{t} , where $\bar{\mathbf{t}}$ is the mean of all $\{\mathbf{t}_i, i \in \{1, 2, \dots, R\}\}$, and obtain the finite sample relative efficiency (RE) of a procedure (denoted by P) with respect to the LS by the ratio of $\text{RE}(P) := \text{SVAR}(\hat{\boldsymbol{\beta}}_{ls}) / \text{SVAR}(\hat{\boldsymbol{\beta}}_P)$. At the same time, we record the total time (TT) consumed by different procedures for all replications. The performance of LTS, LST, and LS is assessed by the four criteria: EMSE, SVAR, TT, and RE. Results are listed in Table 1 (note that both LTS and LST can resist up to 50% contamination without breakdown).

procedure	EMSE	SVAR	TT	RE	EMSE	SVAR	TT	RE
		n=100	$\varepsilon = 0\%$			n=200	$\varepsilon = 0\%$	
LTS	0.5282	0.4764	153.82	0.5164	0.2172	0.1651	265.98	0.6510
LST	0.2969	0.2447	10.911	1.0053	0.1594	0.1073	31.690	1.0019
LS	0.2979	0.2460	1.3254	1.0000	0.1596	0.1075	1.5022	1.0000
		n=100	$\varepsilon = 30\%$			n=200	$\varepsilon = 30\%$	
LTS	142.73	142.54	891.18	0.0864	31.662	31.627	855.82	0.1583
LST	0.4394	0.3875	25.787	31.768	0.2128	0.1606	60.492	31.182
LS	12.410	12.311	1.3358	1.0000	5.1072	5.1072	1.5138	1.0000

Table 1: EMSE, SVAR, TT (seconds), and RE for LTS, LST, and LS based on 1000 Gaussian samples for various n s and contamination rates. LST is computed by `lstReg` with $\alpha = 3$, LTS was computed by `ltsReg`.

Inspecting the table reveals some shocking findings (i) in the pure Gaussian sample setting, LST can outperform both LTS and LS in terms of EMSE and SVAR, this is even more striking in the contaminated sample setting; (ii) in terms of speed, no one can run faster than the LS, but LST (R-based) running faster than LTS (Fortran-based) meanwhile with a smaller EMSE and SVAR in all cases considered; (iii) LTS performs unexpectedly disappointing not only always being the slowest but also always having the worst EMSE and SVAR, robust LTS is ironically inferior to non-robust LS for the contaminated samples (RE is just 8.6%).

3.2 A class of penalized regression estimators based on the LST

Definition 3.2 Now that we have a much more robust regression estimator than the LS, which turns out to be more efficient than the LTS. It is quite natural to replace the SSR in (7) by the Q^n defined in (17), and minimize it, subject to two constraints: ℓ_γ -constraint

$\sum_{i=1}^p |\beta_i|^\gamma \leq t_1$, $t_1 \geq 0$, $\gamma \geq 1$; and ℓ_2 -constraint $\sum_{i=1}^p \beta_i^2 \leq t_2$, $t_2 \geq 0$, the minimizer is

$$\hat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma) := \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n r_i^2 w_i + \lambda_1 \sum_{i=1}^p |\beta_i|^\gamma + \lambda_2 \sum_{i=1}^p \beta_i^2 \right\}, \quad (20)$$

where $\lambda_i := \lambda(t_i) \geq 0$, α , $\gamma \geq 1$, and $w_i := w_i(\beta) := w_i(\beta, \alpha, r_i, \mathbf{z}^{(n)}) = \mathbb{1}(O(r_i, r^{(n)}) \leq \alpha)$.

Before studying its robustness, we address existence and uniqueness of $\hat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$.

Existence and uniqueness

Existence and uniqueness are implicitly assumed for many other penalized regression estimators in the literature. We formally address them for $\hat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$ below.

Theorem 3.1

- (i) $\hat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$ in (20) always exists;
- (ii) $\hat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$ in (20) is unique provided that (a) $\lambda_1 > 0$ and $\gamma > 1$ or (b) $\lambda_2 > 0$.

Proof: see the Appendix. □

The proof of above theorem needs the following result.

Lemma 3.2 Let $S \subset \mathbb{R}^p$ be an open set and $f(\mathbf{x}): \mathbb{R}^p \rightarrow \mathbb{R}^1$ be strictly convex over S and continuous over \bar{S} (the closure of S). Let \mathbf{x}^* be the global minimum of $f(\mathbf{x})$ over S and \mathbf{y} be a point on the boundary of S , then $f(\mathbf{y}) > f(\mathbf{x}^*)$.

Proof: see the Appendix. □

Remarks 3.1

- (i) Note that $\hat{\beta}_{lst-enet}^n(\alpha, 0, 0, \gamma) = \hat{\beta}_{lst}^n$. A sufficient condition for its uniqueness is $\mathbf{C}_n := \mathbf{X}_n' \text{diag}(w_1, \dots, w_n) \mathbf{X}_n / n$ being invertible. That is, the rank of \mathbf{X}_n and the matrix formed by any its $k := \sum_i w_i$ sub-rows is p (see [65]). However, in many applied data set cases, the number of variables (p) is even larger than the number of observations (n), we must have $\text{rank} < p$. So it might not be unique. However, if (a) or (b) in (ii) of the theorem holds, then the strictly convexity guarantees the uniqueness of $\hat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$.
- (ii) The uniqueness of $\hat{\beta}_{lst-enet}^n(\infty, \lambda_1, 0, \gamma)$ (here $\alpha = \infty$ is in the sense that $\alpha \rightarrow \infty$), that is, the uniqueness of $\hat{\beta}_{bridge}^n(\lambda_1, \gamma)$ has been intensively discussed in the literature, see e.g., in the Theorems 1 and 2 of [17], it was shown that $\hat{\beta}_{lst-enet}^n(\infty, \lambda_1, 0, \gamma)$ is unique if $\lambda_1 > 0$ and $\gamma > 1$ plus some condition on the Hessian matrix of SSR; in their Lemma 2 ($\gamma = 1$) of [59], authors showed that it is not unique when there is repeated row of \mathbf{X}_n , it is unique with probability one under the some assumption on predictor variables, argued in [52] and in [1] ($\gamma = 1$). Also see section 2.6 of [21]. □

The most relevant question now is: Is $\widehat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$ much more robust than the existing ones? Or rather, what is its RBP? The next result covers both the LST and the LTS based regularized estimators and provides an affirmative answer to the question.

Theorem 3.2 Let $\widehat{\beta}(\lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)})$ be the penalized regression estimator which minimizes the objective function

$$Q(\beta, \lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)}) := \frac{1}{n} \sum_{i=1}^n r_i^2 w_i + \lambda_1 \sum_{i=1}^p |\beta_i|^\gamma + \lambda_2 \sum_{i=1}^p \beta_i^2, \quad (21)$$

where $w_i \in \{0, 1\}$ is an indicator function: $\mathbb{1}(r_i^2 \leq r_{h:n}^2)$ or $\mathbb{1}(O(r_i, r^{(n)}) \leq \alpha)$ and $\sum_{i=1}^n w_i = k$ ($\lceil n/2 \rceil \leq k \leq n$), $\lambda_i \geq 0$, and $\lambda_1 + \lambda_2 > 0$, $1 \leq \gamma \leq 2$. Then

$$\text{RBP}(\widehat{\beta}(\lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)}), \mathbf{z}^{(n)}) = (n - k + 1)/n.$$

Proof: see the Appendix. □

Remarks 3.2

(i) The theorem covers the special case, the main result (Theorem 1) of [2] with $k = h$ (the default h value is $\lfloor (n + p + 1)/2 \rfloor$) for the $\widehat{\beta}_{lts-lasso}^n$ (incidentally, the boldfaced $\boldsymbol{\rho}$ there is defined to be a vector in \mathbb{R}^n , Theorem 1 fails to explain how to order n vectors in \mathbb{R}^n). The theorem covers the RBP of $\widehat{\beta}_{lst-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma)$ in (19) for any $\alpha \geq 1$, the RBP reaches its highest value $(\lfloor n/2 \rfloor + 1)/n$ when $\alpha = 1$ ($k = \lfloor (n + 1)/2 \rfloor$ in this case). It also for the first time provides the RBP for the reweighted sparse-LTS (or enet-LTS) in [2] (or [29]) with $k = n_w$ there. The theorem tells that ridge, bridge, lasso, and enet all have the lowest RBP $1/n$. The square loss in (21) can be extended to be the more general loss \mathcal{L} in Theorem 2.1.

(ii) Note that the RBP result can be higher than $(\lfloor (n - p)/2 \rfloor + 1)/n$, the latter is the upper bound for any *regression equivariant* estimator (see Theorem 4 on page 125 of [42]). The main reason for this is that the estimator violates the regression equivariance. Without the regression equivariance, any constant vector will have the best possible RBP (100%), but it is not a good estimator at all. Note that the RBP definition in [2] and [28], promoting a constant vector as the most robust estimator, is different from the traditional one. □

Equivariance

Among *regression, scale, and affine equivariance*, the three desired properties (discussed in Sec.3), the regression equivariance is the most fundamental (see [62]), it demands that if one shifts response variable y up and down, then the regression line (or hyperplane) should shift accordingly up and down. The LS estimator and all its robust alternatives mentioned so far satisfy the three properties. But this is not the case for most of regularized regression estimators. In fact,

Theorem 3.3 Among three equivariant properties, only scale equivariance is possessed by $\widehat{\beta}_{ridge}$ in (4), the $\widehat{\beta}_{sqr-lasso}$ in (8), the $\widehat{\beta}_{lad-lasso}$ in (10), and $\widehat{\beta}_{lts-enet}^n(\alpha, 0, \lambda_2, \gamma)$ in (20) among all penalized regression estimators discussed previously.

Proof: Scale equivariance of the $\hat{\beta}_{ridge}$, the $\hat{\beta}_{sqr-lasso}$, the $\hat{\beta}_{lad-lasso}$, the $\hat{\beta}_{lad-lasso}$, and $\hat{\beta}_{lts-enet}^n(\alpha, 0, \lambda_2, \gamma)$ is trivial verification in light of (4), (8), (10), and (20). For other properties and penalized estimators, it suffices to show that regression equivariance is violated.

When y_i is shifted to $y_i + \mathbf{w}_i' \mathbf{b}$, if the regression coefficients β is also shifted to $\beta + \mathbf{b}$, then SSR is invariant whereas the constraint or penalty still on β . \square

Remarks 3.3

(i) There has been an abundance of theoretical and computational work on the generalized lasso and its variants and its special cases. Among hundreds, if not thousands, publications on penalized regression in the literature, very few addressed equivariance. Exceptions are [36], and [28]. It is admitted in [36] that their shooting S-estimator fails to meet the regression equivariance. [28] asserted that via transformation and re-transformation, their estimator enjoys the three equivariance properties.

(ii) Standardizing y and \mathbf{x} columns are common practice in the literature for many computational algorithms for regularized estimators. This, however, amounts to assuming implicitly that these estimators meet the three equivariance properties. Furthermore, centering the observations of y and \mathbf{x} might spread the contamination or outlyingness. \square

4 Finite sample prediction error bounds—consistency

In this section we assume that the true model is $\mathbf{Y} = \mathbf{X}\beta_0 + \mathbf{e}$ where $\mathbf{Y} = (y_1, \dots, y_n)'$, $\mathbf{X}_{n \times p} = (\mathbf{w}_1, \dots, \mathbf{w}_n)'$, and $\mathbf{e} = (e_1, \dots, e_n)'$ with y_i , e_i , and \mathbf{w}_i defined in (1) and (2). We investigate the difference between $\mathbf{X}\hat{\beta}_{lts-enet}^n$ and $\mathbf{X}\beta_0$ (prediction error). Write $\hat{\beta}^n$ for $\hat{\beta}_{lts-enet}^n$ for simplicity.

In light of w_i in (20), the index set $I(\beta)$ in (16) is the same as $\{i : w_i = 1\}$, the scalar $w_i \in \{0, 1\}$ in (20) is different from the vector \mathbf{w}_i above. Write $D(\beta) = \text{diag}(w_1, \dots, w_n)$ with w_i defined in (20). Let A be a n by n symmetric positive semidefinite matrix, a norm (or seminorm) induced by A is $\|\mathbf{x}\|_A^2 = \mathbf{x}' A \mathbf{x}$ for any $\mathbf{x} \in \mathbb{R}^n$. Although $\hat{\beta}^n$ provides predictions for all i , but we just employed residuals r_i with $i \in I(\hat{\beta}^n)$ in (20), so instead of looking at $\|\mathbf{X}(\hat{\beta}^n - \beta_0)\|^2$, we will focus on the squared prediction error $\|\mathbf{X}(\hat{\beta}^n - \beta_0)\|_{D(\hat{\beta}^n)}^2$.

Lemma 4.1 Assume that β_0 is the true parameter of the model in (1), $\hat{\beta}^n := \hat{\beta}_{lts-enet}^n$ is defined in (20). We have

$$\begin{aligned} \|\mathbf{X}(\hat{\beta}^n - \beta_0)\|_{D(\hat{\beta}^n)}^2 &\leq \frac{2}{n} \mathbf{e}' D(\hat{\beta}^n) \mathbf{X}(\hat{\beta}^n - \beta_0) + \frac{1}{n} (\|\mathbf{e}\|_{D(\beta_0)}^2 - \|\mathbf{e}\|_{D(\hat{\beta}^n)}^2) \\ &\quad + \lambda_1 \|\beta_0\|_\gamma^\gamma + \lambda_2 \|\beta_0\|_2^2 - \lambda_1 \|\hat{\beta}^n\|_\gamma^\gamma - \lambda_2 \|\hat{\beta}^n\|_2^2. \end{aligned} \quad (22)$$

Proof: see the Appendix. \square

Write $(e_1^*, \dots, e_n^*) := (\mathbf{e}^*)'$ with $e_i^* = e_i \mathbf{1}(i \in I(\hat{\beta}^n))$. Define two sets

$$\mathcal{S}_1 := \left\{ \max_{1 \leq j \leq p} 2|(\mathbf{e}^*)' \mathbf{x}^{(j)}|/n \leq q_1 \right\}, \quad \mathcal{S}_2 := \left\{ \|\mathbf{e}\|_{D^*}^2/\sigma^2 - N_d \leq q_2 \right\},$$

where $\mathbf{x}^{(j)}$ is the j th column of the fixed design matrix $\mathbf{X}_{n \times p}$, $D^* = D(\beta_0) - D(\hat{\beta}^n)$, a diagonal matrix with $D^*(i, i) = \mathbf{1}(D(\beta_0)(i, i) = 1 \text{ and } D(\hat{\beta}^n)(i, i) = 0)$ and $N_d = |I(\beta_0)| - |I(\beta_0) \cap I(\hat{\beta}^n)|$. Note that $\mathbf{e}'D(\hat{\beta}^n) = (e_1^*, \dots, e_n^*) := (\mathbf{e}^*)'$. Assume hereafter that $\max_{1 \leq j \leq p} \|\mathbf{x}^{(j)}\|_2 \leq c_x$ for a constant c_x .

In the classical setting e_i in (1) is assumed $N(0, \sigma^2)$, it is needed for the second result below, but for the first, it can be relaxed to be a sub-Gaussian variable. For the definition of the latter, we refer to Definition 1.2 of [39] and/or Theorem 2.1.1 of [38].

Lemma 4.2 (i) Let e_i s in (1) be independent sub-Gaussian variables that have variance proxy σ^2 , then $(\mathbf{e}^*)'\mathbf{x}^{(j)}/c_x$ is a sub-Gaussian variable with variance proxy σ^2 . (ii) Let e_i in (1) be i.i.d. $N(0, \sigma^2)$, then $\|\mathbf{e}\|_{D^*}^2/\sigma^2$ follows a χ^2 distribution with N_d degrees of freedom.

Proof: see the Appendix. \square

Lemma 4.3 Assume that e_i s in (1) are i.i.d. $N(0, \sigma^2)$ and other assumptions in Lemmas 4.1-4.2 hold, for any $\delta \in (0, 1)$ let

$$q_1 = \frac{4c_x\sigma}{n} \left(2\sqrt{p} + \sqrt{2\log(2/\delta)} \right); \quad q_2 = 2\sqrt{\log(2/\delta)} \left(\sqrt{|I(\beta_0)|} + \sqrt{\log(2/\delta)} \right),$$

then

$$P(\mathcal{S}_1) \geq 1 - \delta/2; \quad P(\mathcal{S}_2) \geq 1 - \delta/2. \quad (23)$$

Proof: see the Appendix. \square

In light of all Lemmas we are in the position to present the main result.

Theorem 4.1 Set γ in (20) to be one and assume that the assumptions in Lemma 4.3 hold. For any $\delta \in (0, 1)$, selecting $\lambda_1 \geq q_1$. Then with probability at least $1 - \delta$, one has

$$\|\mathbf{X}(\hat{\beta}^n - \beta_0)\|_{D(\hat{\beta}^n)}^2 \leq 2\lambda_1\sqrt{p}\|\beta_0\|_2 + \lambda_2\|\beta_0\|_2^2 + \frac{\sigma}{n}(q_2 + N_d). \quad (24)$$

Remarks 5.1

(i) If select $\lambda_1 \geq q_1$ and in the order of $O(\sqrt{p}/n)$ and $\sqrt{\lambda_2} \leq \lambda_1\sqrt{p}$ and if $\|\beta_0\|_2$ is in the order less than $O(n/p)$ (e.g., $o(n/p)$), then one obtains the consistency if $N_d = o(n)$ since $q_2 = O(n^{1/2})$.

Theorem 4.1 and (23) certainly provide a finite sample squared perdition bound, but the assumption of $N_d = o(n)$ above is too arbitrary, it can be dropped nevertheless. Set α in (19) to be one, then $K := |I(\beta_0)| = |I(\hat{\beta}^n)| = \lfloor (n+1)/2 \rfloor$. Treat the two parts of D^* separately, notice that both $\|\mathbf{e}\|_{D(\beta_0)}^2/\sigma^2$ and $\|\mathbf{e}\|_{D(\hat{\beta}^n)}^2/\sigma^2$ have a χ^2 distribution with the same degrees of freedom K . Write $\frac{1}{n}(\|\mathbf{e}\|_{D(\beta_0)}^2 - \|\mathbf{e}\|_{D(\hat{\beta}^n)}^2)$ in (21) as $\frac{\sigma^2}{n} \left((\|\mathbf{e}\|_{D(\beta_0)}^2/\sigma^2 - K) + (K - \|\mathbf{e}\|_{D(\hat{\beta}^n)}^2/\sigma^2) \right)$. Apply the exponential tail bounds on page 1325 of [30], the upper bound in RHS of (23) becomes $2\lambda_1\sqrt{p}\|\beta_0\|_2 + \lambda_2\|\beta_0\|_2^2 + 2\sigma^2 q_2/n$. The consistency is obtained without $N_d = o(n)$ assumption.

(ii) In above discussions, we treat the unknown σ as known. It appears in q_1 and in the upper bound of (23). In practice, we have to estimate it by an estimator, say $\hat{\sigma}$ so that $P(\hat{\sigma} \geq \sigma)$ with high probability (say, $1 - \delta/3$, in this case, if we change $\delta/2$ and $\log(2/\delta)$ in Lemma 4.3 to $\delta/3$ and $\log(3/\delta)$ respectively, then Theorem 4.1 still holds). Such an estimator $\hat{\sigma}$ has been given on page 104 of [7].

(iii) One limitation of Theorem 4.1 is that the design matrix is fixed. For the general random design X case, one can treat it following the approaches of [3] and [18]. \square

5 Computation algorithm

Re-parametrizations

(i) Following the notation used in (34), we note that the objective function on the RHS of (20) can be written as (also see the proof of Lemma 4.1)

$$O_n(\beta, \lambda_1, \lambda_2, \gamma) = \frac{1}{n} \|\mathbf{Y} - \mathbf{X}\beta\|_{D(\beta)}^2 + \lambda_1 \sum_{j=1}^p |\beta_j|^\gamma + \lambda_2 \|\beta\|_2^2,$$

where $D(\beta) = \text{diag}(w_1, \dots, w_n)$ with w_i defined in (20). Now for every $\lambda_2 > 0$, if we write $\mathbf{X}_{(n+p) \times p}^* = (1 + \lambda_2)^{-1/2} (\mathbf{X}'_{n \times p}, \sqrt{\lambda_2} \mathbf{I}_{p \times p})'$, $\mathbf{Y}_{(n+p) \times 1}^* = (\mathbf{Y}'_{n \times 1}, \mathbf{0}'_{p \times 1})'$, $\beta_{p \times 1}^* = (1 + \lambda_2)^{1/2} \beta$. Then $D^*(\beta^*)_{(n+p) \times p} := (D(\beta^*), \mathbf{I}_{p \times p})' = (D(\beta), \mathbf{I}_{p \times p})'$. If let $\lambda_1^* := \lambda_1 / (1 + \lambda_2)^{1/2}$, we have

$$O_n(\beta^*, \lambda_1, \lambda_2, \gamma) = O_n(\beta, \lambda_1, \lambda_2, \gamma) = \frac{1}{n} \|\mathbf{Y}^* - \mathbf{X}^* \beta^*\|_{D^*(\beta^*)}^2 + \lambda_1^* \sum_{j=1}^p |\beta_j^*|^\gamma,$$

An ℓ_1 -type penalized regression with an objective function much resembling that of a lasso-type problem (especially when in the $\gamma = 1$ case). Denote the minimizer of the objective function above by $\hat{\beta}^*$. It can be computed via the approach for lasso such as the lars algorithm of [11].

(ii) Alternatively, if we set $\lambda^* = \lambda_1 + \lambda_2$ and $\alpha^* = \lambda_2 / (\lambda_1 + \lambda_2)$ (note that $\lambda_1 + \lambda_2 > 0$, otherwise we have a non-penalized problem addressed in [65]), then we have

$$O_n(\beta, \lambda_1, \lambda_2, \gamma) = O_n(\beta, \alpha^*, \lambda^*, \gamma) := \frac{1}{n} \|\mathbf{Y} - \mathbf{X}\beta\|_{D(\beta)}^2 + \lambda^* \left((1 - \alpha^*) \sum_{j=1}^p |\beta_j|^\gamma + \alpha^* \|\beta\|_2^2 \right). \quad (25)$$

Note that $\alpha^* \in [0, 1)$ (a pure ridge regression case is excluded) and $\lambda^* \in (0, \lambda_0]$ for some λ_0 (which is set to be $\max_{1 \leq j \leq p} |2\mathbf{Y}'\mathbf{x}^{(j)}|/n$ as in the literature, see e.g., [2] and Section 2.12 of [7]). *Boundedness of parameters* is the advantage of this formulation. For a given data set $\mathbf{z}^{(n)} = \{(\mathbf{x}'_i, y_i)', i \in \{1, \dots, n\}\}$, we now present the outline our approximate algorithm (AA) for $\hat{\beta}_{lst-enet}^n$.

Pseudocode for computing $\hat{\beta}_{lst-enet}^n$ (lst-enet)

(1) Sample two indices $\{i, j\}$ (two points) and obtain at least p β s: β^k ($k \in \{1, \dots, p\}$) using the algorithm AA1 in [65] and obtain index sets $I(\beta^k) := \{i : w_i := w_i(\beta^k) = 1\}$.

(2) For each β^k , employing the strategy below select a pair (α^*, λ^*) with respect to sub-data sets $(D(\beta^k)\mathbf{X}, D(\beta^k)\mathbf{Y})$.

(3) Based on the sub-data sets $(D(\beta^k)\mathbf{X}, D(\beta^k)\mathbf{Y})$ obtain solution $\hat{\beta}^k$ via LARS algorithm (limited the total steps to 900)

(4) Evaluate $O_n(\beta, \alpha^*, \lambda^*, \gamma)$ with respect to β^k and $\hat{\beta}^k$ ($k \in \{1, \dots, p\}$). Update $\hat{\beta}_{lst-enet}^n$ (initially it is a $\mathbf{0}$ vector) to be the one that minimizes

$$O_n(\beta, \alpha^*, \lambda^*, \gamma) = \frac{1}{n} \sum_{i \in I(\beta)} w_i r_i^2(\beta) + \lambda^* \left((1 - \alpha^*) \sum_{j=1}^p |\beta_j|^\gamma + \alpha^* \|\beta\|_2^2 \right).$$

(5) Repeat (1)-(4) 50 times and output the one that has the minimum objective value.

In algorithm above, (α^*, λ^*) is assumed to be selected. Now we address the issue how to choose this pair. Obviously, we can search among a finite grids over the region $[0, 1) \times (0, \lambda_0]$.

Choice of the penalty/tuning parameters via cross-validation.

We first pick a (relatively small) grid of values for λ^* , say from 0 (excluded) to λ_0 with $\lambda_0/10$ as the step so that there are 10 equal spaced grid points. For the estimation of λ_0 , one can see [11] and [2], or (ii) of Re-parametrization above.

For each λ^* , we will select an α value among 10 equal spaced grid points over $[0, 1)$ via five-fold cross-validation (CV). In k -fold cross-validation, the data are split randomly in k blocks (folds) of approximately equal size. Each block is left out once to fit the model, and the left-out block is used as test data (see Section 7.10.1 of [20]).

The CV is a popular method for estimating the prediction error and comparing different models (see [59] and [20]). The popular **R** package *glmnet* can be used to select the parameters as did in [29], which automatically checks the model quality for a sequence of values for α , taking the mean squared error as an evaluation criterion.

We use a 5-fold CV via our own developed program to avoid the drawback of the *glmnet* which often leads to the error message “from glmnet C++ code (error code 7777); All used predictors have zero variance” (this especially is true under the adversary contamination scenario). The latter leads to a problem for evaluating the performance of the procedure *enetLTS* of [29] when the contamination at 10% level in next section. We have to drop *enetLTS* in the comparison in that situation. We will run ten times of our 5-fold CV, then the pair (λ^*, α) with the minimum averaged CV error will be the final chosen pair (λ^*, α^*) .

The lars algorithm can be used to fit a linear model based on the $k - 1$ blocks to obtain a $\hat{\beta}(\lambda^*, \alpha)$. Other algorithms, such as coordinate descent algorithms (including Fu’s shooting algorithm ([17]) (see 2.11.1 of [7]) can be employed to speed up the computation.

6 Illustration examples and comparison

6.1 Simulation

All R code for simulation and examples as well as figures in this article (downloadable via <https://github.com/left-github-4-codes/lst-enet>) were run on a desktop Intel(R)Core(TM) 21 i7-2600 CPU @ 3.40 GHz.

Five regularized regression procedures We like to compare the performance of our procedure lst-enet with leading regularized regression procedures including lasso, lars, enet, and enetLTS. lasso will be computed via **R** package “lars”, it can be obtained via “elasticnet”. The latter package is mainly for the enet whereas the former mainly focuses on lars. Though lasso could be obtained via “glmnet” but due to the contamination scenario, the glmnet often does not work. Unfortunately, enetLTS employing glmnet in its CV calculation, it can not hand the model $y = \mathbf{w}'\boldsymbol{\beta}_0 + e$ appeared in (1) (an error message “glmnet fails at standardization step”). We use an alternative model given below

Simulation designs To copy with the situation above, we simulate data from the true model: $y = \mathbf{X}\boldsymbol{\beta}_0 + \sigma e$, $e \sim N(0, 1)$, where the true unknown parameter $\boldsymbol{\beta}_0$ is assumed to be a p -dimensional vector with the first $p_1 := \lceil 6\% * p \rceil$ components are threes and the rest $p_2 := p - p_1$ components are zeros. σ is set to be 0.5 but could be changed to other values (leading to different signal-to-noise ratio).

Design I: take sample of \mathbf{X} from $N(\mathbf{0}, \sigma \mathbf{I}_{p \times p})$ and e from $N(0, 1)$. **Design II:** take sample from $\mathbf{X} \sim N(\mathbf{0}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma}(i, j) = \rho_1^{|i-j|}$, $1 \leq i, j \leq p_1$, $\boldsymbol{\Sigma}(i, j) = \rho_2^{|i-j|}$, $p_1 < i, j \leq p$, $\rho_1 = 0.95$, $\rho_2 = 0.05$, all other entries of $\boldsymbol{\Sigma}$ are zeros and $e \sim N(0, 1)$. We take $n \in \{50, 100\}$ samples from the \mathbf{X} and e above and calculate the response $y_i = \mathbf{X}_i\boldsymbol{\beta}_0 + \sigma e_i$, $i \in \{1, 2, \dots, n\}$.

Contamination levels and schemes Let ε be the contamination level, when $\varepsilon = 0$ there is no contamination, an ideal situation (and not realistic). Consider the scenario $\varepsilon \in \{0, 0.05, 0.1, 0.2\}$ (i.e., 0%, 5%, 10%, 20% contamination). Let $m = \lfloor \varepsilon * n \rfloor$, sample m indices from $\{1, \dots, n\}$.

Contamination Scheme I: add 20 to the corresponding m components of (e_1, \dots, e_n) , compute $y_i = \mathbf{X}_i\boldsymbol{\beta}_0 + \sigma e_i$, $i \in \{1, 2, \dots, n\}$, and add 20 (component-wise) to the corresponding m rows of $(\mathbf{X}_1, \dots, \mathbf{X}_n)'$. **Scheme II:** add 20 to the corresponding m components of (e_1, \dots, e_n) , compute $y_i = \mathbf{X}_i\boldsymbol{\beta}_0 + \sigma e_i$, $i \in \{1, 2, \dots, n\}$. Replace the corresponding m rows of $(\mathbf{X}_1, \dots, \mathbf{X}_n)'$ by a p -vector with its first component being 10^4 and the rest are zeros, do the same for the corresponding m components of (y_1, \dots, y_n) but with a scalar 10^{10} .

Four performance criteria The first measure is the estimation error, or L2-error/L2-loss between the true parameter $\boldsymbol{\beta}_0$ and the estimator $\hat{\boldsymbol{\beta}}_P$ via procedure P and is defined as:

$$\text{L2-error}(\boldsymbol{\beta}_0, \hat{\boldsymbol{\beta}}_P) := \|\boldsymbol{\beta}_0 - \hat{\boldsymbol{\beta}}_P\|_2^2, \quad (26)$$

where $\|\mathbf{a} - \mathbf{b}\|_2$ is the ℓ_2 -norm between the two p -dimensional vectors.

On the other hand, one has to take the performance measure into the context of the sparsity model consideration. In the following we introduce the *true sparsity discovery rate*

(TSDR) and the *false sparsity discovery rate* (FSDR). For notation simplicity, we denote the unknown parameter by β^0 (assume it has at least one zero coordinate), an estimator by $\hat{\beta}^P$.

$$\text{TSDR}(\beta^0, \hat{\beta}^P) := \frac{\sum_{i=1}^p \mathbb{1}(\beta_i^0 = 0, \hat{\beta}_i^P = 0)}{\sum_{i=1}^p \mathbb{1}(\beta_i^0 = 0)}, \quad (27)$$

namely, the fraction of correctly detecting/discovering the zero coordinates of the true parameter β^0 . The higher the TSDR, the better the $\hat{\beta}^P$.

$$\text{FSDR}(\beta^0, \hat{\beta}^P) := \frac{\sum_{i=1}^p \mathbb{1}(\beta_i^0 \neq 0, \hat{\beta}_i^P = 0)}{\sum_{i=1}^p \mathbb{1}(\beta_i^0 \neq 0)}, \quad (28)$$

namely, the fraction of falsely detecting/discovering as zero coordinate for the true parameter β^0 . The lower the FSDR, the better the $\hat{\beta}^P$.

The fourth performance measure is a popular one, it is (square-)root of mean squared (prediction) error (RMSE) on testing data. That is, for a given data set, one first partitions data into training and testing two parts (we take the ratio 7:3 for partition). Then fit the model and get estimator based on the training data and using the testing data to get the RMSE. Testing data sets are often assumed to be clean (have no contamination or outliers) in the literature. This, however, is not realistic in practice.

Let \mathbf{X}_{test} , y_{test} be the testing data and $\hat{\beta}_P$ be the estimator obtained from the training data. Then

$$\text{RMSE}(\hat{\beta}_P) := \left(\text{mean}((y_{test} - \mathbf{X}_{test}\hat{\beta}_P)^2) \right)^{1/2}. \quad (29)$$

The four performance measures above were discussed in the literature before, all are hoped to be small except the TSDR which is hoped to be as high as possible. All (but RMSE) depend on the unknown parameter β^0 .

Tuning parameter choices In the following simulation, the tuning parameter α in (17), or (20) is set to be 3 and h in enetLTS is set to be the default value $\lfloor 0.75(n+1) \rfloor$.

Example 6.1 We first consider $\varepsilon = 0$. For simplicity, data are generated according to design **I** and set $n = 100$, $p = 50$ (low dimension case) or $n = 50$, $p = 300$ (high dimension and sparsity case). We generated 50 samples for \mathbf{X} and e and obtained corresponding responses y . The simulation results are displayed in Figure 2. For description simplicity, we use hereafter P1 for lst-enet, P2 for lasso, P3 for lars, P4 for enetLTS, P5 for enet in the Figures.

Inspecting Figure 2 reveals that (i) with respect to (w.r.t.) FSDR, all four perform equally well with 0% mis-discovery rate; (ii) w.r.t. TSDR, lst-enet, Lars, and lasso perform stably and at a highest rate while enet with a less stable lower rate but enetLTS performs most unstable with the lowest median rate; (iii) w.r.t. RMSE, lasso and lars are the best followed by lst-enet, enetLTS has the median RMSE that is also among the best but with the widest spread of RMSE while enet has the largest (and wider spread of) RMSE; (iv) w.r.t. L2-error, lst-enet, lars, and lasso are among the best while enetLTS has the worst performance followed by enet. Overall, lst-enet, lars, and lasso are among the best whereas enetLTS performs worst overall followed by enet.

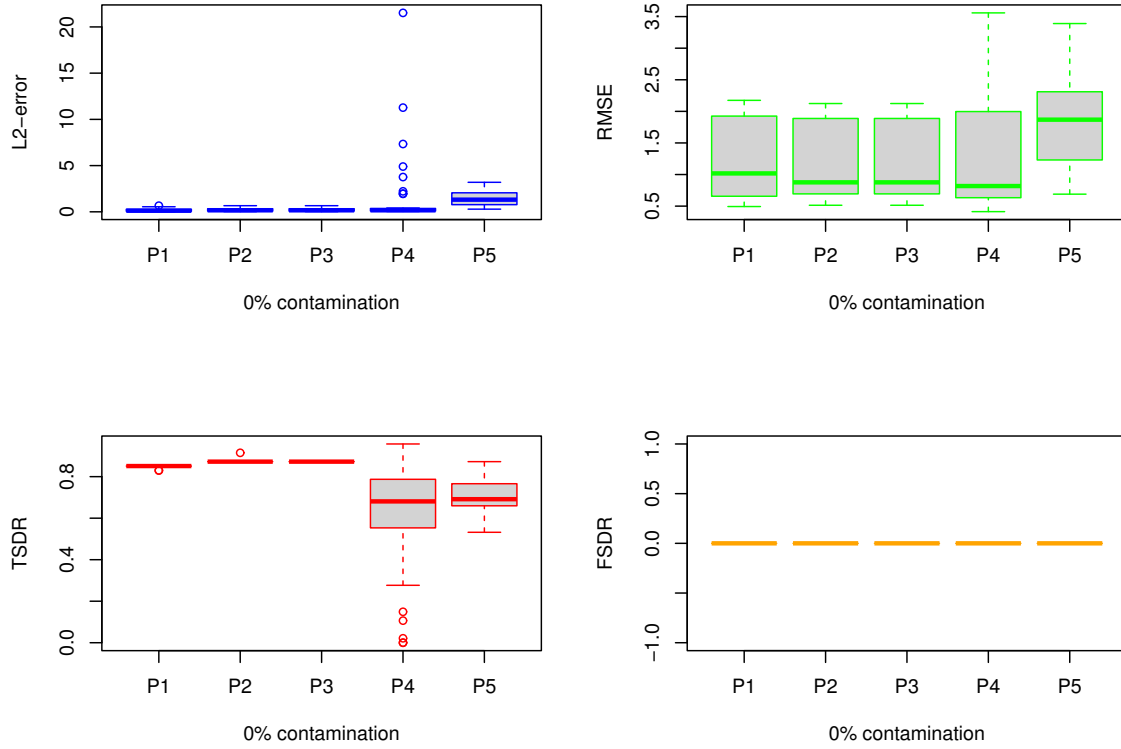


Figure 2: Boxplots for five procedures (P1 stands for lst-enet, P2 for lasso, P3 for lars, P4 for enetLTS, P5 for enet) and 50 samples each with $n = 100$ and $p = 50$ that are generated from design I with 0% contamination rate.

Example 6.2 Perfect normal data are not realistic in practice. We now consider $\varepsilon = 0.05$ (i.e. 5% contamination), all others are the same as Example 7.1 except the contamination scheme II will be adopted (in 7.1 contamination scheme does not matter). We first consider $n = 100$, $p = 50$ (low dimension case) and for simplicity generate data according design I. Performance of five procedures in 50 samples is displayed in Figure 3.

Inspecting the Figure reveals that (i) w.r.t. L2-error, lst-enet, enetLTS and enet are the best performers while lasso and lars are equally dissatisfactory; (ii) w.r.t. RMSE, the situation is the same as in the L2-error case; (iii) w.r.t. TSDR, enet is the best performer (this perhaps is false best since it might assign zero to all components of the estimator $\hat{\beta}$ that could lead to 100% of its FSDR) while enetLTS is the worst; (iv) w.r.t. FSDR, lst-enet, enetLTS are the best performers followed by lasso and lars while enet is the loser. Overall, lst-enet is the only winner.

The simulation study above with 5% contamination is repeated but $n = 50$ and $p = 300$ (high dimensional case) and simulation design II is adopted. Results are displayed in Figure 4.

Reviewing the Figure discovered that (i) w.r.t. L2-error, lst-enet, enetLTS, and enet are the best performers while lasso and lars are disappointing; (ii) w.r.t. RMSE; the situation is almost the same as in L2-error case; (iii) w.r.t. TSDR, lst-enet, lasso, lars and enet are

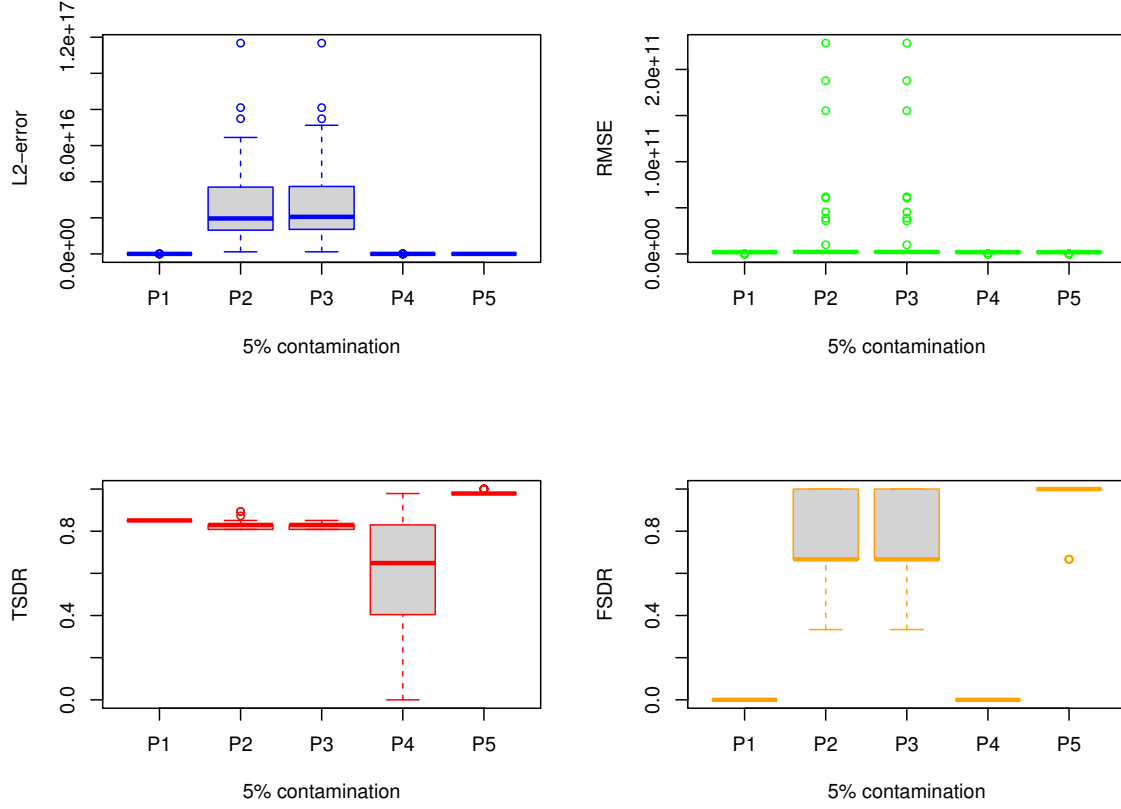


Figure 3: Boxplots for five procedures and 50 samples each with $n = 100$ and $p = 50$ that are generated from design I with 5% contamination rate.

the best performers while enetLTS is the loser; (iv) w.r.t. FSDR, enet is the worst performer (since its FSDR is almost 100%), enetLTS has the lowest median value while it has the widest spread. lst-enet is the second best performer, lasso and lars are disappointed. Overall, lst-enet is the only winner.

Example 6.3 In practice, 10% (or even 20%) contamination is not rare. Next we consider the case $\varepsilon = 0.1$ (i.e., 10% contamination), contamination scheme II will be adopted. Samples of 50 with $n = 50, p = 300$ are generated with simulation design I. Due to the higher level contamination and the usage of **R** package glmnet in its background CV calculation, enetLTS fails to go through the computation we have to drop it in our comparison. Simulation results are displayed in Figure 5.

Inspecting the Figure reveals that (i) w.r.t. L2-error, lst-enet and enet are the best while lasso and lars are inferior; (ii) w.r.t. RMSE, the situation is the same as in L2-error case; (iii) w.r.t. TSDR, enet is the worst performer (it assigns zeros to almost all components of $\hat{\beta}$ that will lead to 100% of its FSDR), lst-enet and lasso are the best performers followed by lars; (iv) w.r.t. FSDR, lst-enet is the best performer, enet is the worst one while lasso and lars perform dissatisfactory. Overall, lst-enet is the winner.

The advantage of lst-enet is even better demonstrated in Figure 6 when $n = 100$ and

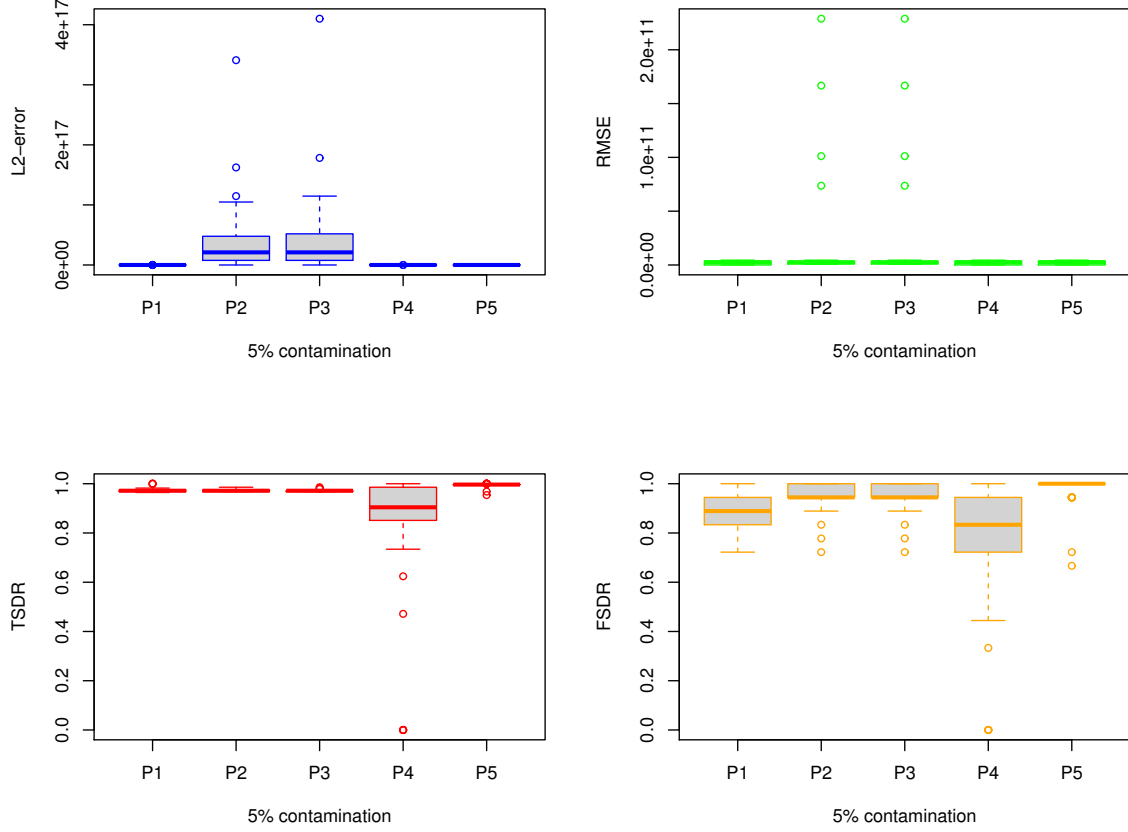


Figure 4: Boxplots for five procedures and 50 samples each with $n = 50$ and $p = 300$ that are generated from design I with 5% contamination rate.

$p = 50$ and $\varepsilon = 0.2$ (i.e., 20% contamination).

6.2 A real data example

Example 6.4 To analyze a realistic dataset with very large number of variables, we consider the well-known cancer data from the National Cancer Institute (NCI60); see [50] for more detail about this dataset. A total of 59 of the human cancer cell-lines ($n = 59$) were assayed for gene expression and protein expression. The data set, downloadable from the CellMiner program package, NCI (<http://discover.nci.nih.gov/cellminer/>) and available from the R package robustHD, has been repeatedly studied in the literature, see e.g., [31].

We process the data set by following the approach in the literature and treat the gene expression microarray data as the predictors \mathbf{X}_{raw} (a 59 by 22283 matrix) and the protein expression data as responses variables \mathbf{Y}_{raw} (a 59 by 162 matrix). Similar to [31] or [2], we order the protein expression variables according to their scale (employing MAD as a scale estimator instead of the standard deviation) and select the one with median MAD, serving

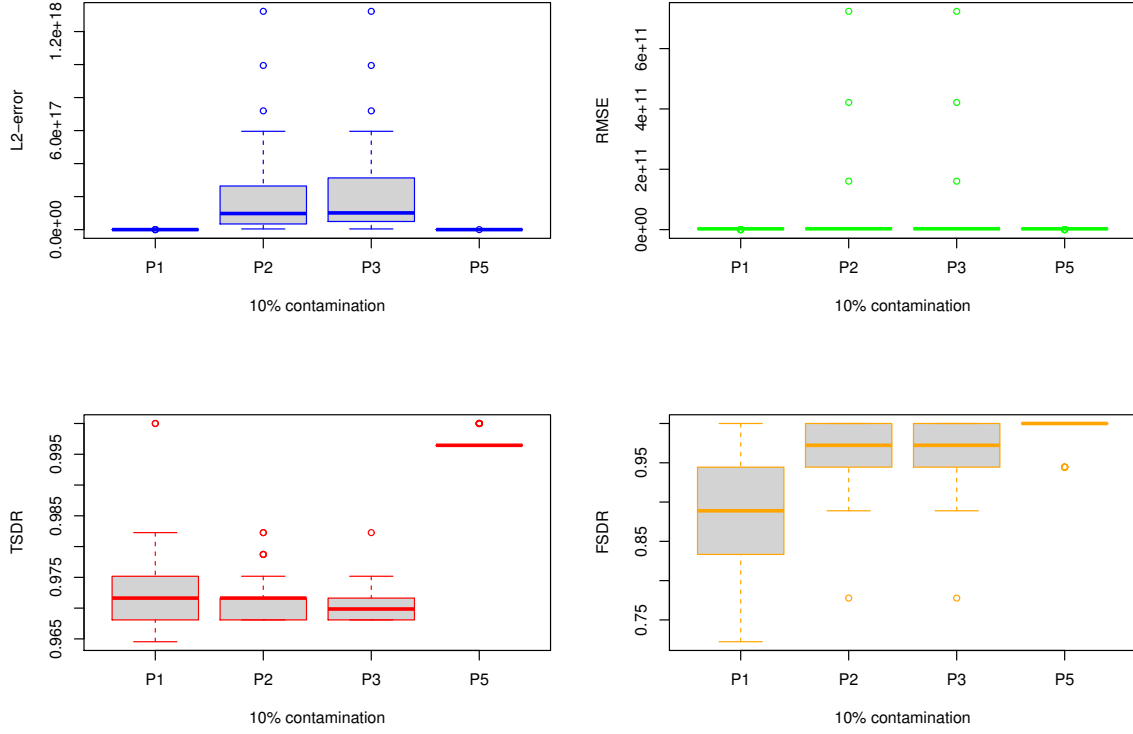


Figure 5: Boxplots for four procedures and 50 samples each with $n = 50$ and $p = 300$ that are generated from design I with 10% contamination rate.

as our dependent variable. It is 75th column of the protein expression data matrix. Denote it by \mathbf{Y} .

Next, we selected out genes using their correlations with \mathbf{Y} . Here we adopt the robust correlation measure in [27]. We obtain 22283 ordered (decreasing) correlations and select top $k1 = 100$ corresponding columns of \mathbf{X}_{raw} and combined with the bottom $1000 - k1$ columns as our final \mathbf{X} , reducing the number of genes from 2,2283 to $p = 1,000$. The number p could easily be changed by adjusting $k1$.

We partition (by rows) $\mathbf{X}_{59 \times 1000}$ and \mathbf{Y} into $x.train$, $y.train$ and $x.test$, $y.test$ according the rate 7 : 3. That is 41 rows of \mathbf{X} and \mathbf{Y} for the training data sets, the rest 18 rows as testing data sets. We do this step 50 times and each time we calculate the RMSE (the only measure that still valid without the given β_0) for the five procedures. The results are displayed in Figure 7, where the fifth performance measure is introduced, that is, the empirical mean squared error defined as:

$$\text{EMSE}(\hat{\beta}_P) := \frac{1}{R} \sum_{i=1}^R (\hat{\beta}_P - \bar{\beta}_P)^2, \quad (30)$$

where R is the replication number, namely, $\frac{R}{R-1} \text{EMSE}$ is the sample variance of $\hat{\beta}_P$. \bar{x} stands for the sample mean of x_i .

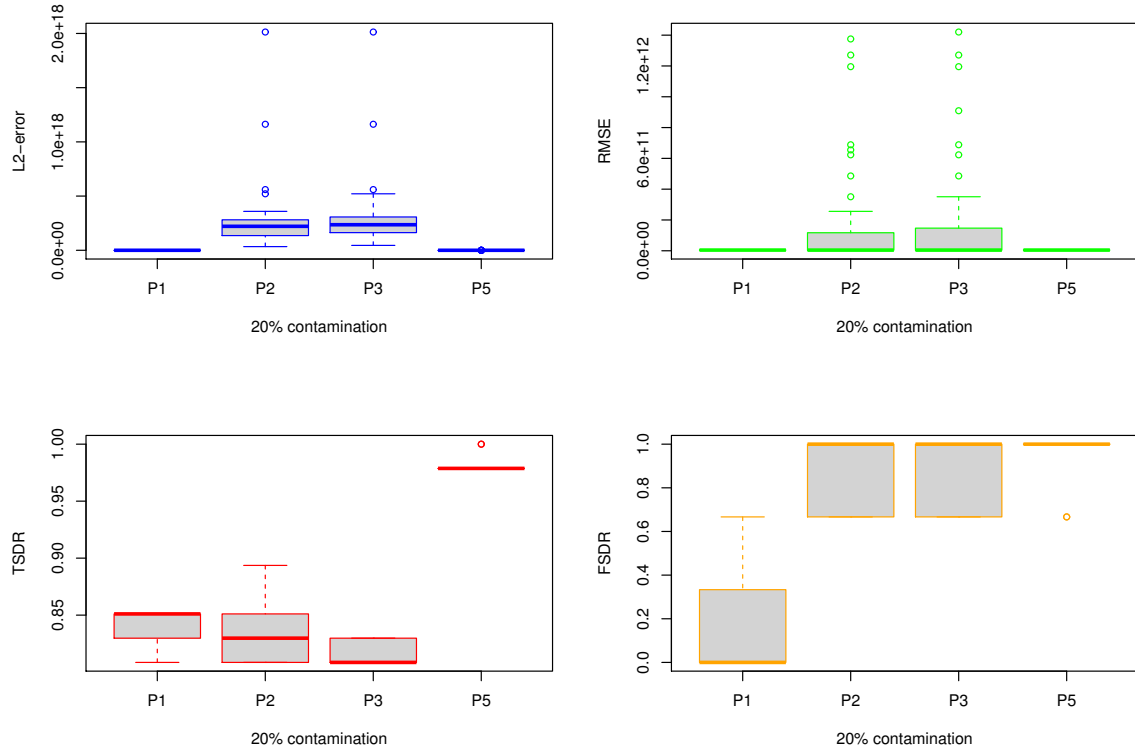


Figure 6: Boxplots for four procedures and 50 samples each with $n = 100$ and $p = 50$ that are generated from design I with 20% contamination rate.




Figure 7: Boxplots for five procedures for real data set NCI60 with 50 times partition of a selected sub-dataset $p = 1000$ and $n = 59$.

Inspecting the Figure reveals that (i) lasso, lars, enetLTS (enetLTS has the wider spread) have the smallest RMSE but their sample variances (EMSE) are among the largest; (ii) RMSE of the lst-enet is the second smallest but it is the most stable estimator with the distinguished smallest sample variance (EMSE) which means that with different training and testing data sets obtained by random partitioning, lst-enet produces very closed solutions; (iii) enetLTS has the lowest median RMSE but its sample variance is the remarkably large; (iv) enet has the categorical largest RMSR while its sample variance is also the largest. Overall, lst-enet is recommended with the rivals enetLTS, lasso and lars.

7 Concluding discussions

The originality of the current work (i) For the first time it discovered that most of the leading penalized regression estimators for high-dimensional sparse data can break down by a single outlier (or contaminating point) (Theorem 2.1). (ii) For the first time it demonstrated that the newly proposed lst-enet estimator not only possesses a high breakdown robustness (Theorem 3.2) but also performs well in simulation studies and a real data example (Section 6), serving as a robust alternative to regularized regression estimators. (iii) the newly proposed lst-enet estimator, its unique existence (Theorem 3.1), and its finite-sample prediction error bound (Theorem 4.1) are original notion and novel results, respectively. (iv) The equivariance discussion and result (Theorem 3.3) are novel and original, and so is the (ii) of Remarks 3.3. (v) Table 1 provides brand new evidence of the superiority of LST over LTS.

Connection to the existing work (i) It adduced the procedure LST introduced in [65] (see Section 3.1 and figure 1, they are quoting results from [65]). (ii) The idea of replacing the square loss of LS in lasso by a loss of robust alternative of LS (e.g. the LTS) appeared in the literature before (see, e.g. [2] and [29]). (iii) A RBP result first appeared in [2] intended for $\hat{\beta}_{lts-lasso}^n$ (Theorem 1, the main result of [2]), however, our Theorem 3.2 covers their Theorem 1 ((a) the loss function in Theorem 3.2, being more general, covers the one in Theorem 1 of [2], it is not vice versa since the loss in [2] involving the sum of h ordered n -vectors but Theorem 1 fails to explain how to order the vectors; (b) the penalty terms in Theorem 3.2 are more general and cover the one in Theorem 1, it is not vice versa).

Robust measure Finite sample breakdown point has been served as a prevailing quantitative robustness measure in finite sample practice, the main advantage/beauty is its non-randomness and probability-free nature that is exactly why it was enthusiastically welcomed and quick became adopted in a broad spectrum of disciplines after its introduction in 1983.

Critics (e.g., [45]), however, would like to have a more complicated version, a version that includes randomness and Orlicz norm. They argued that worst case performance might not be a good robustness measure. On the other hand, it is common practice to use the worst case performance as in the complexity of an algorithm or the safety of a passenger cars case.

Future possible work (a) Further performance measure could be pursued including (i) whether $\hat{\beta}$ performs well on future samples (i.e., whether $E(Y - \mathbf{w}'\hat{\beta})^2$ is small); (ii) whether $\hat{\beta}$ closely approximates the “true” parameter β_0 (i.e., whether $\|\hat{\beta} - \beta_0\|$ is small with high probability); or (iii) whether $\hat{\beta}$ correctly identifies the relevant coordinates of the “true,” sparse parameter β_0 (i.e., whether $(\beta_{0j} = 0) \Leftrightarrow (\hat{\beta}_j = 0)$ with high probability). (b)

Extension of current regression work to a more general setting to cover discriminant analysis, logistic regression, and other topics.

Acknowledgments

The author thanks Prof.s Haolei Weng, Yiyuan She and Wei Shao for insightful comments and stimulating discussions which significantly improved the manuscript.

Appendix: proofs of main results

Proof of Theorem 2.1

Proof: Clearly, it suffices to show that $\text{RBP}(\hat{\beta}^*(\lambda_1, \lambda_2, \gamma, \mathbf{z}^{(n)}), \mathbf{z}^{(n)}) \leq 1/n$. Equivalently, to show that one contaminating point can break down the estimator. Assume, otherwise, one point is not enough to break down the estimator. That is, there exists an M such that

$$\sup_{\mathbf{z}_1^{(n)}} \|\hat{\beta}^*(\lambda_1, \lambda_2, \gamma, \mathbf{z}_1^{(n)})\|_2 < M < \infty, \quad (31)$$

where $\mathbf{z}_1^{(n)}$ stands for any contaminated data set by replacing one point in the original data set $\mathbf{z}^{(n)}$ with an arbitrary point in \mathbb{R}^p . We seek a contradiction now.

Replace $\mathbf{z}_1 = (\mathbf{x}'_1, y_1)'$ in $\mathbf{z}^{(n)} = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$ by $\mathbf{z}_1^* = ((\delta, 0, \dots, 0), \kappa\delta)'$. Denote the contaminated data set by $\mathbf{z}_1^{(n)}$ and the estimator based on it as $\hat{\beta}^* := \hat{\beta}^*(\lambda_1, \lambda_2, \gamma, \mathbf{z}_1^{(n)})$.

Let $M_y = \max_i |y_i|, M_x = \max_i |\mathbf{x}_{i1}|$. Let $\beta_\kappa = (0, \kappa, 0, \dots, 0)' \in \mathbb{R}^p$ and set $\kappa = (\sqrt{p}+1)M+1$. Select a large δ such that $f(\mathcal{L}(\delta)/n) \geq f(\mathcal{L}(M_y + M_x\kappa)) + g(\beta_\kappa, \lambda_1, \lambda_2, \gamma) + 1$. This is possible since $\mathcal{L}(x) \rightarrow \infty$ when $|x| \rightarrow \infty$ and $\mathcal{L}(x)$ is non-decreasing over $(0, \infty)$, $f(x)$ is strictly increasing over $[0, \infty)$ and $f(x) \rightarrow \infty$ as $x \rightarrow \infty$. Then

$$\begin{aligned} O(\hat{\beta}^*) &\leq O(\beta_\kappa) = f\left(\frac{1}{n} \sum_{i=1}^n \mathcal{L}(r_i(\beta_\kappa))\right) + g(\beta_\kappa, \lambda_1, \lambda_2, \gamma) \\ &= f\left(\frac{1}{n} \sum_{i=2}^n \mathcal{L}(r_i(\beta_\kappa))\right) + g(\beta_\kappa, \lambda_1, \lambda_2, \gamma) \quad (\text{since } \mathcal{L}(r_1) = \mathcal{L}(0) = 0) \\ &\leq f\left(\frac{n-1}{n} \mathcal{L}(M_y + M_x\kappa)\right) + g(\beta_\kappa, \lambda_1, \lambda_2, \gamma) \\ &< f\left(\frac{1}{n} \mathcal{L}(\delta)\right) - 1. \end{aligned} \quad (32)$$

On the other hand, for any $\beta = (\beta_1, \beta_2, \dots, \beta_p)' \in \mathbb{R}^p$ such that $(\sqrt{p}+1)\|\beta\|_2 \leq \kappa - 1$, one

has

$$\begin{aligned}
O(\beta) &\geq f\left(\frac{1}{n}\mathcal{L}(y_1 - \mathbf{w}_1\beta)\right) = f\left(\frac{1}{n}\mathcal{L}(\kappa\delta - (\beta_1 + \delta\beta_2))\right) \\
&\geq f\left(\frac{1}{n}\mathcal{L}(\kappa\delta - (\delta|\beta_1| + \delta|\beta_2|))\right) \\
&\geq f\left(\frac{1}{n}\mathcal{L}(\delta(\kappa - (\sqrt{p} + 1)\|\beta\|_2))\right) \geq f\left(\frac{1}{n}\mathcal{L}(\delta)\right),
\end{aligned} \tag{33}$$

where the facts: (i) $|\beta_i| \leq \|\beta\|_i (i = 1, 2)$ and (ii) $\|\beta\|_1 \leq \sqrt{p}\|\beta\|_2$ are utilized.

Combining (32) and (33), leads to the conclusion that

$$\|\hat{\beta}^*(\lambda_1, \lambda_2, \gamma, \mathbf{z}_1^{(n)})\|_2 > \frac{\kappa - 1}{\sqrt{p} + 1} = M,$$

which contradicts (31). \square

Proof of Theorem 3.1

Proof: $\lambda_1 + \lambda_2 = 0$ case has been treated in [65], we treat $\lambda_1 + \lambda_2 > 0$ case here.

(i) Denote the objective function on the RHS of (20) as

$$O_n(\beta, \lambda_1, \lambda_2, \gamma) := O(\beta, \lambda_1, \lambda_2, \alpha, \gamma, \mathbf{z}^{(n)}) = \frac{1}{n} \sum_{i=1}^n r_i^2 w_i + \lambda_1 \sum_{j=1}^p |\beta_j|^\gamma + \lambda_2 \sum_{i=1}^p \beta_i^2. \tag{34}$$

Name the three terms on the RHS above as $g(\alpha, \beta, \mathbf{z}^{(n)})$, $g_1(\lambda_1, \gamma, \beta)$, and $g(\lambda_2, \beta)$, respectively. It is readily seen that the RHS of (20) is equivalent to minimizing $G(\beta, \alpha, \lambda_1, \gamma) := g(\alpha, \beta, \mathbf{Z}^{(n)}) + g_1(\lambda_1, \gamma, \beta)$ subject to $\sum_{i=1}^p \beta_i^2 \leq t_2$, $t_2 \geq 0$

By Lemma 2.2 of [65], $g(\alpha, \beta, \mathbf{z}^{(n)})$ is continuous in β (this is not as obvious as one believed) while the continuity of $g_1(\lambda_1, \gamma, \beta)$ in β is obvious. Therefore we have a continuous function of β , $G(\beta, \alpha, \lambda_1, \gamma)$, which obviously has minimum value over the compact set $\|\beta\|_2 \leq t_2$.

(ii) Follows the approach originated in [65], we partition the parameter space \mathbb{R}^p of β into disjoint open pieces R_{β^k} , $1 \leq k \leq L \leq \binom{n}{\lfloor (n+1)/2 \rfloor}$ and $\cup_{1 \leq k \leq L} \overline{R}_{\beta^k} = \mathbb{R}^p$, where \overline{A} stands for the closure of the set A , and

$$R_{\beta^k} = \{\beta \in \mathbb{R}^p : I(\beta) = I(\beta^k), O_{i_1}(\beta) < O_{i_2}(\beta) \cdots < O_{i_K}(\beta)\}, \tag{35}$$

where $O_i(\beta) := O(r_i, r^{(n)})$ for a given $\mathbf{z}^{(n)}$ and β , i_1, \dots, i_K in $I(\beta)$ and $K = |I(\beta)|$ with w_i defined in (20)

$$I(\beta) = \{i : w_i = 1\}. \tag{36}$$

For any $\beta \in \mathbb{R}^p$, either there is R_η and $\beta \in R_\eta$ or there is R_ξ , such that $\beta \notin R_\eta \cup R_\xi$ and $\beta \in \overline{R}_\eta \cap \overline{R}_\xi$. Now we claim that $\hat{\beta} := \hat{\beta}_{ls-enet}^n(\alpha, \lambda_1, \lambda_2, \gamma) \in R_{\beta^{k_0}}$ for some $1 \leq k_0 \leq L$.

Otherwise, assume that $\hat{\beta} \in \overline{R}_{\beta^{k_0}}$. By Lemma 2.2 of [65], $g(\alpha, \beta, \mathbf{Z}^{(n)})$ (denoted by $Q^n(\beta)$ there) is convex over $R_{\beta^{k_0}}$. Therefore, $O_n(\beta, \lambda_1, \lambda_2, \gamma)$ is strictly convex in β over

$R_{\beta^{k_0}}$. Assume that β^* is the global minimum of $O_n(\beta, \lambda_1, \lambda_2, \gamma)$ over $R_{\beta^{k_0}}$. Then it is obviously that $O_n(\widehat{\beta}, \lambda_1, \lambda_2, \gamma) \leq O_n(\beta^*, \lambda_1, \lambda_2, \gamma)$. But this is impossible in light of Lemma 3.2. The strict convexity of $O_n(\beta, \lambda_1, \lambda_2, \gamma)$ over $R_{\beta^{k_0}}$ guarantees the uniqueness. \square

Proof of Lemma 3.2

Proof: Let $B(\mathbf{x}^*, r)$ be a small ball centered at \mathbf{x}^* with a small radius r and $B(\mathbf{x}^*, r) \subset S$. Let $B^c := S - B(\mathbf{x}^*, r)$ and $\alpha^* = \inf_{\mathbf{x} \in B^c} f(\mathbf{x})$. Then, $\alpha^* > f(\mathbf{x}^*)$ (in light of strict convexity). Since $y \in \overline{S}$, then there is a sequence $\{\mathbf{x}_j\} \in B^c$ such that $\mathbf{x}_j \rightarrow \mathbf{y}$ and $f(\mathbf{x}_j) \rightarrow f(\mathbf{y})$ as $j \rightarrow \infty$. Hence $f(\mathbf{y}) = \lim_{j \rightarrow \infty} f(\mathbf{x}_j) \geq \alpha^* > f(\mathbf{x}^*)$. \square

Proof of Theorem 3.2

Proof: We complete it in two steps.

(i) $m \leq n - k$ contaminating points are not enough to break down the estimator. Let $M_y = \max_i |y_i|$, denote the minimizer of the Q in (21) for the contaminated sample as $\widehat{\beta}$. Then, it is obviously that

$$Q(\widehat{\beta}) \leq Q(\mathbf{0}, \lambda_1, \lambda_2, \gamma, \mathbf{z}_m^{(n)}) = \frac{1}{n} \sum_{i=1}^n r_i^2 w_i \leq \frac{k}{n} (M_y^2),$$

where the last inequality deserves further explanations. Note that there are at least k uncontaminated (original) points. Therefore, in the case that $w_i = \mathbf{1}(r_i^2 \leq r_{h:n}^2)$, the RHS of the above display wants to keep the sum of smallest $k = h$ squared residuals (y_i^2), this sum is certainly no greater than that of $k = h$ squared residuals from the k original points.

Likewise, in the case of $w_i = \mathbf{1}(O(r_i, r^{(n)}) \leq \alpha)$, the RHS wants to keep the sum of squared residuals (y_i^2) from $k = |I(\mathbf{0})|$ points that have the smallest outlyingness no greater than α , which is certainly no greater than the sum of k squared residuals from the k original points.

Assume, w.l.o.g. that $\lambda_1 > 0$ ($\lambda_2 > 0$ is even easier). Consider any β with $\|\beta\|_2 \geq M := ((k+1)M_y^2/n\lambda_1)^{1/\gamma}$, then

$$Q(\beta, \lambda_1, \lambda_2, \gamma, \mathbf{z}_m^{(n)}) > \lambda_1 \sum_{i=1}^p |\beta_i|^\gamma = \lambda_1 \|\beta\|_\gamma^\gamma \geq \lambda_1 \|\beta\|_2^\gamma \geq \frac{k+1}{n} M_y^2,$$

where the fact that $\|x\|_q \leq \|x\|_p$ when $1 \leq p \leq q < \infty$ is invoked.

The two displays above imply that

$$\|\widehat{\beta}(\lambda_1, \lambda_2, \gamma, \mathbf{z}_m^{(n)})\|_2 < M.$$

(ii) $m = n - k + 1$ contaminating points are enough to break down the estimator.

The structure and basic idea of this part is an analogue to that of proof of Theorem 2.1. Assume, otherwise, m points are not enough to break down the estimator. That is, there exists an M such that

$$\sup_{\mathbf{z}_m^{(n)}} \|\widehat{\beta}(\lambda_1, \lambda_2, \gamma, \mathbf{z}_m^{(n)})\|_2 < M < \infty, \quad (37)$$

where $\mathbf{z}_m^{(n)}$ stands for any contaminated data set by replacing m points in the original data set $\mathbf{z}^{(n)}$ with m arbitrary points in \mathbb{R}^p . We seek a contradiction now.

Replacing m original points \mathbf{z}_i s with the point $((\delta, 0, \dots, 0), \delta\kappa)'$. Denote the contaminated data set by $\mathbf{z}_m^{(n)}$ and the estimator based on it as $\hat{\beta} := \hat{\beta}(\lambda_1, \lambda_2, \gamma, \mathbf{z}_m^{(n)})$.

Let $M_y = \max_i |y_i|, M_x = \max_i |\mathbf{x}_{i1}|$. Let $\beta_\kappa = (0, \kappa, 0, \dots, 0)' \in \mathbb{R}^p$ and set $\kappa = M + 1$. Select a large δ such that

$$\delta^2 \geq \frac{\max(k-m, k-1)}{n} (M_y + \kappa M_x)^2 + \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2 + 1,$$

This is possible since $x^2 \rightarrow \infty$ when $|x| \rightarrow \infty$. Note that $k \geq m$.

It is readily seen that all residuals based on β_κ and m contaminated points are zeros. All non-zero residuals correspond to uncontaminated original points. Then in the case that $w_i = \mathbb{1}(r_i^2 \leq r_{k:n}^2)$

$$\begin{aligned} Q(\hat{\beta}) &\leq Q(\beta_\kappa) = \frac{1}{n} \sum_{i=1}^n r_i^2 w_i + \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2 \\ &= \begin{cases} \frac{1}{n} \sum_{i=1}^{k-m} r_{j_i}^2 + \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2, & \text{if } k > m \\ \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2 & \text{else,} \end{cases} \end{aligned} \quad (38)$$

where the last equality is due that fact that the objective function sums the smallest k squared residuals, but among n squared residuals, m of them are zeros.

For the case $w_i = \mathbb{1}(O(r_i, r^{(n)}) \leq \alpha)$, one has

$$\begin{aligned} Q(\hat{\beta}) &\leq Q(\beta_\kappa) = \frac{1}{n} \sum_{i=1}^n r_i^2 w_i + \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2 \\ &= \frac{1}{n} \sum_{i=1}^{k-1} r_{j_i}^2 + \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2, \end{aligned} \quad (39)$$

where the last equality is due the fact that there is at most $n - m = k - 1$ non-zero residuals. Overall we have

$$Q(\hat{\beta}) \leq Q(\beta_\kappa) \leq \frac{\max(k-m, k-1)}{n} (M_y + \kappa M_x)^2 + \lambda_1 \kappa^\gamma + \lambda_2 \kappa^2 \leq \delta^2 - 1. \quad (40)$$

On the other hand, for any β with $\beta_1 \leq \|\beta\|_2 \leq \kappa - 1$, one has

$$Q(\beta) \geq (\kappa\delta - \delta\beta_1)^2 \geq \delta^2, \quad (41)$$

where we utilize the fact that among the k squared residuals, there is at least one residual that is based on a contaminated point since un-contaminated points are at most $k - 1$.

Combining (40) and (41) leads to

$$\|\widehat{\beta}(\lambda_1, \lambda_2, \gamma, \mathbf{z}_m^{(n)})\|_2 > \kappa - 1 = M,$$

which contradicts (37). \square

Proof of Lemma 4.1

Proof: Employing the true model assumption: $\mathbf{Y} = \mathbf{X}\beta_0 + \mathbf{e}$ and that $\widehat{\beta}^n$ is the minimizer of the RHS of (19), this is straightforward by some algebraic derivations. \square

Proof of Lemma 4.2

Proof: According to the Definition 1.2 of [39], each e_i^* is a sub-Gaussian variable. Write $\mathbf{v}^{(j)} = (v_1, \dots, v_n)' := \mathbf{x}^{(j)}/c_x$, then $(\mathbf{e}^*)'\mathbf{x}^{(j)}/c_x = \sum_{i \in I(\widehat{\beta}^n)} v_i e_i$ with $\|\mathbf{v}\|_2 \leq 1$. Following the proof of Theorem 1.6 of [39], one obtains the desired result (i). (ii) follows from the fact that e_i^2/σ^2 has a χ^2 distribution with one degree of freedom. \square

Proof of Lemma 4.3

Proof: First we note that

$$P\left(\max_{1 \leq j \leq p} |(\mathbf{e}^*)'\mathbf{x}^{(j)}|/n > q_1/2\right) \leq P\left(\max_{\|\mathbf{v}\| \leq 1} |\mathbf{v}'\mathbf{e}^*| > nq_1/(2c_x)\right),$$

where $\mathbf{v} \in \mathbb{R}^n$. Now invoking Lemma 5.2 and Theorem 2.2.2 and Remark 2.2.2 of [38] (set $nq_1/(2c_x)$ to be the t in Remark 2.2.2), one gets that

$$P\left(\max_{1 \leq j \leq p} |(\mathbf{e}^*)'\mathbf{x}^{(j)}|/n > q_1/2\right) \leq P\left(\max_{\|\mathbf{c}\| \leq 1} |\mathbf{c}'\mathbf{e}^*| > q_1/2\right) \leq \delta/2,$$

the statement about $P(\mathcal{S}_1)$ follows.

For the statement about $P(\mathcal{S}_2)$, we first invoke Lemma 5.2 and notice that $\|\mathbf{e}\|_{D^*}^2/\sigma^2$ follows a χ^2 distribution with N_d degrees of freedom, then invoke Lemma 1 and Comments on page 1325 of [30], we get

$$P\left(\|\mathbf{e}\|_{D^*}^2/\sigma^2 - N_d \geq 2\sqrt{N_d t} + 2t\right) \leq e^{-t},$$

now if one sets $\delta/2 = e^{-t}$, that is $t = \log 2/\delta$, then one gets that

$$P(\mathcal{S}_2) \geq 1 - \delta/2.$$

The second statement follows. \square

References

- [1] Ali, A. and Tibshirani, R. J. (2019), “The Generalized Lasso Problem and Uniqueness”, *Electron. J. Statist.*, 13(2): 2307-2347.
- [2] Alfons, A., Croux, C., and Gelper, S. (2013), “Sparse least trimmed squares regression for analyzing high-dimensional large data sets”, *Ann. Appl. Stat.*, 7(1): 226-248.

- [3] Bartlett, P.L., Mendelson, S., · Neeman, J. (2011), ℓ_1 -regularized linear regression: persistence and oracle inequalities, *Probab. Theory Relat. Fields*, (2012) 154:193–224 DOI 10.1007/s00440-011-0367-2
- [4] Belloni, A., Chernozhukov, V., and Wang, L. (2011). “Square-root lasso: Pivotal recovery of sparse signals via conic programming”. *Biometrika*, 98(4):791–806.
- [5] Bianco, A. M., Boente, G., Chebi, G. (2022), “Penalized robust estimators in sparse logistic regression”, *TEST* (2022), 31:563–594 <https://doi.org/10.1007/s11749-021-00792-w>.
- [6] Bogdan, M., Van Den Berg, E., Sabatti, C., Su, W., and Candès, E. J. (2015). “SLOPE-adaptive variable selection via convex optimization”. *Ann. Appl. Stat.* 9(3): 1103-1140, DOI: 10.1214/15-AOAS842.
- [7] Bühlmann, P. and Van De Geer, S. (2011), *Statistics for High-Dimensional Data: Methods, Theory and Applications*, Springer-Verlag Berlin Heidelberg
- [8] Donoho, D. L. “Breakdown properties of multivariate location estimators”. PhD Qualifying paper, Harvard Univ. (1982).
- [9] Donoho, D. L., and Gasko, M. (1992), “Breakdown properties of multivariate location parameters and dispersion matrices”, *Ann. Statist.*, **20**, 1803-1827.
- [10] Donoho, D. L., and Huber, P. J. (1983), “The notion of breakdown point”, in: P. J. Bickel, K. A. Doksum and J. L. Hodges, Jr., eds. *A Festschrift for Erich L. Lehmann* (Wadsworth, Belmont, CA) pp. 157-184.
- [11] Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004), “Least angle regression”. *Ann. Statist.*, 32 407–499
- [12] Emmert-Streib, F. and Dehmer, M. (2019), “High-Dimensional LASSO-Based Computational Regression Models: Regularization, Shrinkage, and Selection”, *Mach. Learn. Knowl. Extr.*, 1, 359–383; doi:10.3390/make1010021
- [13] Fan, J. (1997). Comments on «wavelets in statistics: A review» by A. Antoniadis. *Journal of the Italian Statistical Society*, 6(2):131
- [14] Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* 96 1348–1360.
- [15] Frank, I.E. and Friedman, J.H. (1993), “A statistical view of some chemometrics regression tools (with discussion)”, *Technometrics*, 35: 109–148.
- [16] Freijeiro-González, L., Febrero-Bande, M., and González-Manteiga, W. (2022), “A critical review of LASSO and its derivatives for variable selection under dependence among covariates”, *International Statistical Review*, 90(1), 118–145 doi: 10.1111/insr.12469
- [17] Fu, W.J. (1998), “Penalized regression: The Bridge versus the Lasso”, *Journal of Computational and Graphical Statistics*, 7(3): 397–416.
- [18] Guédon, O., Mendelson, S., Pajor, A., and Tomczak-Jaegermann, N. (2007), “Subspaces and orthogonal decompositions generated by bounded orthogonal systems”. *Positivity* 11 269–283.
- [19] Hastie, T. and Efron, B. (2011), “lars: Least angle regression, lasso and forward stagewise”. R package version 0.9-8.

- [20] Hastie, T., Tibshirani, R., and Friedman, J. (2017), The Elements of Statistical Learning Data Mining, Inference, and Prediction, Springer Series in Statistics.
- [21] Hastie, T., Tibshirani, R., and Wainwright, M. J. (2015), Statistical Learning With Sparsity: The Lasso and Generalizations, Boca Raton, FL: CRC Press. [254]
- [22] Hesterberg, T., Choi, N. H., Meier, L., and Fraley, C. (2008), “Least angle and ℓ_1 penalized regression: A review”, *Statistics Surveys* Vol. 2 61–93 ISSN: 1935-7516 DOI: 10.1214/08-SS035
- [23] Hoerl, A.E., and Kennard, R.W. (1970a), “Ridge Regression: Biased Estimation for Nonorthogonal Problems,” *Technometrics*, 12, 55–67.
- [24] Hoerl, A.E., and Kennard, R.W. (1970b), “Ridge Regression: Applications to Nonorthogonal Problems,” *Technometrics*, 12, 69–82.
- [25] Huber, P. J. (1981). Robust Statistics. Wiley, New York.
- [26] James, G., Witten, D., Hastie, T., Tibshirani, R. (2013), An Introduction to Statistical Learning, with Applications in R. Springer New York Heidelberg Dordrecht London
- [27] Khan, J. A., Van Aelst, S. and Zamar, R. H. (2007), “Robust linear model selection based on least angle regression”. *J. Amer. Statist. Assoc.*, 102 1289-1299.
- [28] Kong, D., Bondell, H.D., and Wu, Y. (2018), “Fully efficient robust estimation, outlier detection and variable selection via penalized regression”, *Statistica Sinica*, 28 (2018), 1031-1052.
- [29] Kurnaz FS, Hoffmann I, Filzmoser P (2018), “Robust and sparse estimation methods for high-dimensional linear and logistic regression”, *Chemomet Intell Lab Syst*, 172:211–222,
- [30] Laurent, B. and Massart, P. (2000), “Adaptive estimation of a quadratic functional by model selection”, *Ann. Statist.*, 28(5): 1302-1338, DOI: 10.1214/aos/1015957395
- [31] Lee, D., Lee, W., Lee, Y. and Pawitan, Y. (2011). “Sparse partial least-squares regression and its applications to high-throughput data analysis”. *Chemometrics and Intelligent Laboratory Systems*, 109 1–8.
- [32] Li, G., Peng, H. and Zhu, L. (2011), “Nonconcave penalized M-estimation with a diverging number of parameters”. *Statistica Sinica*, 21 391–419.
- [33] Maglott, D., Ostell, J., Pruitt, K. D. and Tatusova, T. (2005), “Entrez Gene: genecentered information at NCBI”. *Nucleic Acids Research*, 33 D54–D58.
- [34] Maronna, R. A. (2011), “Robust ridge regression for high-dimensional data. *Technometrics*, 53 44–53.”,
- [35] Maronna, R. A., Martin, R. D., and Yohai, V. J. (2006), “Robust Statistics: Theory and Methods”, John Wiley & Sons
- [36] Öllerer, V., Alfons A., and Croux, C. (2016), “The shooting S-estimator for robust regression”, *Comput Stat* (2016) 31:829–844.
- [37] Rosset, S. and Zhu, J. (2004), “Discussion of “Least angle regression” by Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.”, *Ann. Statist.*, 32 469–475.
- [38] Pauwels, E., (2020), Lecture notes: Statistics, optimization and algorithms in high dimension.
- [39] Rigollet, P. and Hütter, J.C. (2017), High Dimensional Statistics, Lecture Notes.

- [40] Rousseeuw, P. J. (1984), “Least median of squares regression”. *J. Amer. Statist. Assoc.*, 79, 871-880.
- [41] Rousseeuw, P. and Yohai, V. (1984), “Robust regression by means of S-estimators”. In *Robust and Nonlinear Time Series Analysis* (Heidelberg, 1983), vol. 26 of *Lecture Notes in Statist.* New York: Springer, pp. 256-272.
- [42] Rousseeuw, P. J. and Leroy, A. M. (1987). *Robust Regression and Outlier Detection*. New York, NY, USA: John Wiley & Sons, Inc.
- [43] Rousseeuw, P. J. and Van Driessen, K. (2006), “Computing LTS Regression for Large Data Sets”, *Data Mining and Knowledge Discovery* 12, 29-45.
- [44] Uma T Shankavaram, William C Reinhold, Satoshi Nishizuka, Sylvia Major, Daisaku Morita, Krishna K Chary, Mark A Reimers, Uwe Scherf, Ari Kahn, Douglas Dolginow, Jeffrey Cossman, Eric P Kaldjian, Dominic A Scudiero, Emanuel Petricoin, Lance Liotta, Jae K Lee, John N Weinstein (2007), “Transcript and protein expression profiles of the NCI-60 cancer cell panel: an integromic microarray study”, *Mol Cancer Ther* . 2007 Mar; 6(3):820-32.
- [45] She, Y., Wang, Z., and Shen, J. (2022), “Gaining Outlier Resistance With Progressive Quantiles: Fast Algorithms and Theoretical Studies”, *J. Amer. Statist. Assoc.*, 2022, 117(539), 1282–1295.
- [46] Shao, J. (2003), *Mathematical Statistics*, Springer Texts in Statistics.
- [47] Stahel, W. A. (1981), *Robuste Schätzungen: Infinitesimale Optimalität und Schätzungen von Kovarianzmatrizen*. Ph.D. dissertation, ETH, Zurich.
- [48] Stromberg, A. J., Hawkins, D. M., and Hössjer, O. (2000), “The Least Trimmed Differences Regression Estimator and Alternatives”, *J. Amer. Statist. Assoc.*, 95, 853-864.
- [49] Sun, Q., Zhou, W.X., and Fan, J.Q. (2020), “Adaptive Huber Regression”, *J. Amer. Statist. Assoc.*, 115:529, 254-265, DOI: 10.1080/01621459.2018.1543124
- [50] Tan, C.S., Salim, A., Ploner, A., Lehtio, J., Chia, K.S., Pawitan, Y., (2009) “Correlating gene and protein expression data using correlated factor analysis”, *BMC Bioinformatics* 10 272, doi:10.1186/1471-2105-10-272.
- [51] Tibshirani, R. (1996), “Regression shrinkage and selection via the lasso”. *Journal of the Royal Statistical Society. Series B.*, 58, 267-288.
- [52] Tibshirani, R. J. (2013), “The lasso problem and uniqueness”, *Electronic Journal of Statistics*, Vol. 7, 1456–1490.
- [53] Wang, H., Li, G. and Jiang, G. (2007). Robust regression shrinkage and consistent variable selection through the LAD-Lasso. *Journal of Business & Economic Statistics*, 25, 347-355.
- [54] Wang, P.; Chen, S.; Yang, S. (2022), “Recent Advances on Penalized Regression Models for Biological Data”, *Mathematics*, 2022(10), 3695. <https://doi.org/10.3390/math10193695>
- [55] Van De Geer, S. (2008). “High-dimensional generalized linear models and the lasso”. *Ann. Statist.*, 36 614–645.
- [56] Yohai, V. J. (1987), “High breakdown-point and high efficiency robust estimates for regression”. *Ann. Statist.*, 15, 642-656.
- [57] Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38, 894-942.

- [58] Zhang, C.-H. and Zhang, T. (2012). A general theory of concave regularization for high-dimensional sparse estimation problems. *Statistical Science*, 27, 576-593
- [59] Zou, H. and Hastie, T. (2005), ‘Regularization and variable selection via the elastic net’, *Journal of the Royal Statistical Society: Series B*, 67(2), 301–320.
- [60] Zuo, Y. (2003) “Projection-based depth functions and associated medians”, *Ann. Statist.*, 31, 1460-1490.
- [61] Zuo, Y. (2006), “Multi-dimensional trimming based on projection depth”, *Ann. Statist.*, 34(5), 2211-2251.
- [62] Zuo, Y. (2024), “Non-asymptotic robustness analysis of regression depth median”, *Journal of Multivariate Analysis*, (in press), arXiv: 2009.00646
- [63] Zuo, Y. (2022), “Non-asymptotic analysis and inference for an outlyingness induced winsorized mean”, *Stat. Papers*, <https://doi.org/10.1007/s00362-022-01353-5>
- [64] Zuo, Y., Serfling, R., (2000), “General notions of statistical depth function”, *Ann. Statist.*, 28, 461-482.
- [65] Zuo, Y. and Zuo, H. (2023), “Least sum of squares of trimmed residuals regression”, *Electronic Journal of Statistics* 17(2), 2416-2446. arXiv:2202.10329