

Outlier Detection Using Nonconvex Penalized Regression

Yiyuan She and Art B. Owen

Author's Footnote:

Art B. Owen is Professor, Department of Statistics, Stanford University. Mailing address: Department of Statistics, Stanford University, Stanford, CA 94305. (email: owen@stanford.edu). Yiyuan She is Assistant Professor, Department of Statistics, Florida State University. Mailing address: Department of Statistics, Florida State University, Tallahassee, FL 32306-4330 (email: yshe@stat.fsu.edu). This work was partially supported by NSF grants DMS-0604939 and DMS-0906056. The authors are grateful to the anonymous referees and the associate editor for their careful comments and useful suggestions. We also thank a reviewer for suggesting a comparison to the elastic net. The authors thank J. Wisnowski and Professor J. Simpson for sharing the S-PLUS code in Wisnowski, Montgomery & Simpson (2001). We thank Len Stefanski for bringing McCann & Welsch (2007) to our attention.

Abstract

This paper studies the outlier detection problem from the point of view of penalized regressions. Our regression model adds one mean shift parameter for each of the n data points. We then apply a regularization favoring a sparse vector of mean shift parameters. The usual L_1 penalty yields a convex criterion, but we find that it fails to deliver a robust estimator. The L_1 penalty corresponds to soft thresholding. We introduce a thresholding (denoted by Θ) based iterative procedure for outlier detection (Θ -IPOD). A version based on hard thresholding correctly identifies outliers on some hard test problems. We find that Θ -IPOD is much faster than iteratively reweighted least squares for

large data because each iteration costs at most $O(np)$ (and sometimes much less) avoiding an $O(np^2)$ least squares estimate. We describe the connection between Θ -IPOD and M -estimators. Our proposed method has one tuning parameter with which to both identify outliers and estimate regression coefficients. A data-dependent choice can be made based on BIC. The tuned Θ -IPOD shows outstanding performance in identifying outliers in various situations in comparison to other existing approaches. This methodology extends to high-dimensional modeling with $p \gg n$, if both the coefficient vector and the outlier pattern are sparse.

1 Introduction

Outliers are a pervasive problem in statistical data analysis. Nonrigorously, outliers refer to one or more observations that are different from the bulk of the data. Hampel et al. (1986) estimate that a routine data set may contain about 1-10% (or more) outliers. Unfortunately, outliers often go unnoticed (Rousseeuw & Leroy 1987), although they may have serious effects in estimation, inference, and model selection (Weisberg 1985). Perhaps the most popular statistical modeling method is ordinary least squares (OLS) regression. OLS is very sensitive to outliers — a single unusual observation may break it down completely. Our goal in this work is outlier identification for regression models, together with robust coefficient estimation.

We consider the linear regression model given by $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$ where $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a fixed matrix of predictors, $\boldsymbol{\beta} \in \mathbb{R}^p$ is a fixed (unknown) coefficient vector and $\boldsymbol{\epsilon} \in \mathbb{R}^n$ is a random error vector. The i 'th case is written $y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \epsilon_i$.

Suspected outliers are most commonly found by looking at residuals $r_i = y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}}$ where $\hat{\boldsymbol{\beta}}$ is the OLS estimate of $\boldsymbol{\beta}$. It is well-known that such raw residuals can fail to detect outliers at leverage points. A better way to detect an outlier is the *leave-one-out* approach (Weisberg 1985). If the i 'th case is suspected to be an outlier, then

we compute the externally studentized residual

$$t_i = \frac{y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}}_{(i)}}{\hat{\sigma}_{(i)} (1 + \mathbf{x}_i^\top (\mathbf{X}_{(i)}^\top \mathbf{X}_{(i)})^{-1} \mathbf{x}_i)^{1/2}} \quad (1.1)$$

where $\mathbf{X}_{(i)}$, $\hat{\boldsymbol{\beta}}_{(i)}$ and $\hat{\sigma}_{(i)}$ are the predictor matrix, coefficient estimate and scale estimate respectively, based on $n - 1$ observations, leaving out the i 'th. Large values $|t_i| > \eta$ are then taken to suggest that observation i is an outlier. The threshold $\eta = 2.5$ (Rousseeuw & Leroy 1987) is a reasonable choice. If $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ then $t_i \sim t_{(n-p-1)}$ and we can even attach a significance level to $|t_i|$. After removing an apparent outlier from the data, one then looks for others.

Studentized residuals, and other leave-one-out methods such as Cook's distance and DFFITS, are simple and effective when there is only one outlier. When there are multiple outliers, these simple methods can fail. Two phenomena have been remarked on. In *masking*, when an outlying i 'th case has been left out, the remaining outliers cause either a large value of $\hat{\sigma}_{(i)}$ or a small value of $|y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}}_{(i)}|$, or both, and as a result observation i does not look like an outlier. Therefore, multiple outliers may mask each other and go undetected. In *swamping*, the effect of outliers is to make $|y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}}_{(i)}|$ large for a non-outlying case i . Swamping could lead one to delete good observations and becomes more serious in the presence of multiple outliers.

In this paper we take the studentized residual as our starting point. The t -test for whether observation i' is an outlier is the same as testing whether the parameter γ is zero in the regression $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \gamma \mathbf{1}_{i=i'} + \boldsymbol{\epsilon}$. Because we don't know which observations might be outliers, we use a model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\gamma} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}), \quad (1.2)$$

in which the parameter γ_i is nonzero when observation i is an outlier. This mean-shift

model allows any combination of observations to be outliers. It has $n + p$ regression parameters and only n data points. Our approach is to fit (1.2) imposing sparsity on $\boldsymbol{\gamma}$ in order to avoid the trivial estimate $\hat{\boldsymbol{\gamma}} = \mathbf{y}$ and to get a meaningful estimate of $\boldsymbol{\beta}$. The resulting algorithm is called thresholding (denoted by Θ) based iterative procedure for outlier detection, or Θ -IPOD for short.

All of our proposals (apart from one exception noted where it arises) require a preliminary robust regression to be run. This practice is in line with the best current robust regression methods. The preliminary regression supplies a robust estimate of $\boldsymbol{\beta}$, and usually a robust estimate of σ as well. The robust regression methods that we compare to are known to outperform the preliminary regressions that they use as starting points. We will compare Θ -IPOD to those best performing methods.

The rest of the paper is organized as follows. Section 2 surveys the literature on robust regression. Section 3 develops the soft-IPOD algorithm which fits (1.2) using an L_1 penalty on $\boldsymbol{\gamma}$. This algorithm minimizes a convex criterion, but it is *not* robust. Section 4 develops a family of algorithms replacing soft-thresholding by a general thresholding rule Θ . We find that some nonconvex criteria properly identify multiple outliers in some standard challenging test cases. Section 5 investigates the computational efficiency of Θ -IPOD in comparison to iteratively reweighted least squares (IRLS). IRLS requires a QR decomposition at each iteration, while Θ -IPOD needs only one. As a result, though possibly requiring more iterations, it is much faster in large problems that we investigate. In Section 6, we discuss the important problem of parameter tuning in outlier detection and carry out an empirical study to demonstrate the advantage of our penalized approach. Section 7 extends the technique to high-dimensional data with $p > n$. Our conclusions are in Section 8.

2 Survey of robust regression

Many methods have been developed for multiple outlier identification. Robust or resistant regressions, such as M -estimators (Huber 1981) and Least Trimmed Squares (LTS) (Rousseeuw & Leroy 1987), were proposed to provide a trustworthy coefficient estimate even in the presence of multiple outliers. There are also clustering-based procedures, such as Serbert et al. (1998), and some informal methods based on graphics. Multiple outlier detection procedures usually alternate between *two* steps. One step scans regression output (coefficient and variance estimates) to identify seemingly clean observations. The other fits a linear regression model to those clean observations. The algorithm can be initialized with OLS, but generally it is better to initialize it with something more robust.

If a data set contains more than one outlier, masking may occur and the task of outlier detection is much more challenging. Well-known examples of masking include the Belgian telephone data and the Hertzsprung-Russell star data (Rousseeuw & Leroy 1987), as well as some artificial datasets like the Hawkins-Bradru-Kass (HBK) data (Hawkins, Bradu & Kass 1984) and the Hadi-Simonoff (HS) data (Hadi & Simonoff 1993). The HS and HBK data sets both have swamping effects.

The main challenge in multiple outlier detection is to counter masking and swamping effects. Hadi & Simonoff (1993) describe two broad classes of algorithms — direct methods and indirect methods. The direct procedures include forward search algorithms (Hadi & Simonoff 1997; Peña & Yohai 1995; Atkinson & Riani 2000) and backward selection (Menjoge & Welsch 2010) among others. Indirect methods are those that use residuals from a robust regression estimate to identify the outliers. Examples of indirect methods include Least Median of Squares (LMS) (Rousseeuw 1984), Least Trimmed Squares (LTS) (Rousseeuw & Leroy 1987), S-estimators (Rousseeuw & Yohai 1984), MM-estimators (Yohai 1987), one-step GM estimators (Simpson

et al. 1992) and SIS estimators (Coakley & Hettmansperger 1993). Almost all start with an initial high breakdown point estimate (not necessarily efficient) say from LTS, S, MTS (Nguyen & Welsch 2010), or Peña & Yohai’s (1999) (denoted by PY) fast procedure. Most published examples in the outlier detection literature are in 10 or fewer dimensions. Methods with high breakdown point typically have costs that grow exponentially in the dimension. In practice, when there are a large number of predictors, PY can be applied to provide an initial estimate with certain robustness, although in theory its breakdown point property is not well established (Peña & Yohai 1999).

It is worth mentioning that outlier identification and robust regression are two closely related but not quite identical problems (Rousseeuw & Leroy 1987; Maronna et al. 2006). Even if we perfectly identified the regression coefficient vector, there could still be some overlap between the residual distributions for good and outlying points. That is, given the true regression coefficient β there would be type one and type two errors in trying to identify outliers. On the other hand, if we could perfectly identify all gross outliers, it is a relatively easy task to obtain a robust coefficient estimate, as will be supported by our experiments in Section 6.

3 Soft-IPOD

We will use the mean shift model (1.2) from the introduction which predicts \mathbf{y} by the usual linear model $\mathbf{X}\beta$ plus an outlier term γ . If $\gamma_i = 0$ then the i ’th case is good, and otherwise it is an outlier. Our goals are to find a robust estimate of β as well as to estimate γ thereby identifying which cases are outliers and which are not. We assume that γ is sparse because outliers should not be the norm. We suppose at first that $n > p$ and that $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top$ has full rank p . Section 7 considers the case when β is sparse, too, with p possibly greater than n . Yet the majority of our paper focuses on the outlier problem only. Let \mathbf{H} be the hat matrix defined by

$\mathbf{H} = \mathbf{H}(\mathbf{X}) = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$. The i 'th diagonal entry of \mathbf{H} , denoted h_i , is called the leverage of the i 'th observation.

The assumed sparsity of $\boldsymbol{\gamma}$ motivates using an L_1 -penalized regression to minimize

$$f_{\text{soft}}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \boldsymbol{\lambda}) \equiv \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \boldsymbol{\gamma}\|_2^2 + \sum_{i=1}^n \lambda_i |\gamma_i| \quad (3.1)$$

over $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$. The choice of $\boldsymbol{\lambda} \in [0, \infty)^n$ is discussed below. The function f_{soft} in (3.1) is jointly convex in $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$. Its simple form suggests that an alternating optimization can be applied: given $\boldsymbol{\gamma}$, the optimal $\boldsymbol{\beta}$ is the OLS estimate from regression of $\mathbf{y} - \boldsymbol{\gamma}$ on \mathbf{X} ; given $\boldsymbol{\beta}$, this L_1 -penalized problem is orthogonal and separable in $\boldsymbol{\gamma}$, and the optimal $\boldsymbol{\gamma}$ can be obtained by soft-thresholding. This alternation between updates to $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ closely matches the interplay between detecting outliers and fitting regression parameters mentioned in Section 2. We call the resulting algorithm soft-IPOD because it is based on soft thresholding. This formulation was earlier used by McCann & Welsch (2007) and Gannaz (2006).

We can derive a principled choice of $\boldsymbol{\lambda}$ by comparing to Donoho & Johnstone's (1994) classical work which states that $\lambda(n) = \sigma\sqrt{2\log n}$ is minimax optimal when σ is known and the predictors are orthogonal so that the residuals are uncorrelated. Our residuals are correlated. For example, at the first step $\mathbf{r} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}_{\text{ols}} \sim \mathcal{N}(\mathbf{H}\boldsymbol{\gamma}, \sigma^2(\mathbf{I} - \mathbf{H}))$. This motivates taking $\lambda_i = \sigma\sqrt{2(1 - h_i)\log n}$ and then Šidák's inequality (Šidák 1967) controls the probability of wrongly declaring $\gamma_i \neq 0$. It is conservative. If there are many outliers, then σ is not easily found. As a result, we prefer to set $\lambda_i = \lambda\sqrt{1 - h_i}$ and tune the regularization parameter λ in a data-dependent way. See Sections 4 and 6 for details.

To initialize the algorithm we must specify either $\boldsymbol{\beta}$ or $\boldsymbol{\gamma}$. Because f_{soft} is convex in $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ the starting point is not crucial. We could use $\boldsymbol{\gamma} = \mathbf{0}$ for example. (For a detailed discussion of the initial estimate in the general situation, see Section 6

and Section 8.) A pseudo-code outline for soft-IPOD is deferred to Section 4. Soft-IPOD corresponds to Algorithm 1 (Θ -IPOD) there, with a soft thresholding rule Θ . Soft-IPOD is guaranteed to converge, even though it is not contractive. See Section 4.2.

Although (3.1) is a well-formulated model and soft-IPOD is computationally efficient, in the presence of multiple outliers with moderate or high leverage values, this method fails to remove masking and swamping effects. Take the artificial HBK data as an illustration. Using a robust estimate $\hat{\sigma}$ from LTS, and $\lambda_i = \hat{\sigma} \sqrt{2(1 - h_i) \log n}$, we obtained $\hat{\gamma} = [0, \dots, 0, -8.6, -9.7, -7.6, -8.4, 0, \dots, 0]^\top$ which identifies cases 11-14 as serious outliers, while the true γ is $[10, \dots, 10, 0, \dots, 0]^\top$ with cases 1-10 being the actual outliers. This erroneous identification is not a matter of parameter tuning. Figure 1 plots the soft-IPOD solution over the relevant range for λ . Whatever value of λ we choose, cases 11-14 are sure to be swamped, and cases 1-10 are very likely to be masked. The L_1 approach is not able to identify the correct outliers without swamping. The McCann-Welsch approach, using LARS, is similar to soft-IPOD, and is not robust. Our extensive experience shows that this L_1 technique hardly works for any benchmark dataset in the outlier detection literature. We suspect that a convex criterion is inherently incompatible with robustness.

Next we delineate the parallels between soft-IPOD and Huber's M-estimate regression which is similarly non-robust. As pointed out by Gannaz (2006) and Antoniadis (2007), there is a connection between the L_1 -penalized regression (3.1) and Huber's M-estimate. Huber's loss function is

$$\rho(t; \lambda) = \begin{cases} \lambda|t| - \lambda^2/2, & \text{if } |t| > \lambda \\ t^2/2, & \text{if } |t| \leq \lambda. \end{cases}$$




Figure 1: The solution path of the L_1 penalized regression on the HBK data. The estimates $\hat{\gamma}_i(\lambda)$ ($i = 1, 2, \dots, 75$) are plotted against the constraint value c , given by $\|\hat{\gamma}\|_1 / \|\mathbf{r}_{ols}\|_1$ (the dual parameter of λ), where \mathbf{r}_{ols} is the OLS residual vector. The top 10 paths correspond to cases 1-10, the true outliers. The bottom four correspond to cases 11-14 which are all good leverage points, but possibly swamped.

Huber’s method with concomitant scale estimation, minimizes

$$l(\boldsymbol{\beta}, \sigma) \equiv \sum_{i=1}^n \rho\left(\frac{y_i - \mathbf{x}_i^\top \boldsymbol{\beta}}{\sigma}; \lambda\right) \sigma + \frac{1}{2} cn \sigma \quad (3.2)$$

over $\boldsymbol{\beta}$ and σ jointly, where $c \geq 0$ and $\lambda > 0$ are given constants. We can prove the following result, which is slightly more general than the version given by Gannaz (2006) and Antoniadis (2007).

Proposition 3.1. *For any $c \geq 0$ and $\lambda > 0$ suppose that we minimize*

$$g(\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma) \equiv \frac{1}{2\sigma} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \boldsymbol{\gamma}\|_2^2 + \frac{1}{2} cn \sigma + \lambda \|\boldsymbol{\gamma}\|_1, \quad (3.3)$$

over $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and σ . Then the minimizing $\boldsymbol{\beta}$ and σ match those from minimizing Huber’s criterion (3.2). For any fixed $\sigma > 0$ minimizing (3.3) over $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ yields the minimizer of (3.2) over $\boldsymbol{\beta}$.

This connection is helpful to understand the inherent difficulty with L_1 -penalized regression described earlier. It is well known that Huber’s method cannot even handle moderate leverage points well (Huber 1981, p. 192) and is prone to masking and swamping in outlier detection. Its break-down point is 0.

A promising way to improve (3.1) is to adopt a different penalty function possibly nonconvex. Our launching point in this paper is, however, the operator Θ in the above iterative algorithm. Substituting an appropriate Θ for soft-thresholding in the IPOD (see Algorithm 1), we may obtain a good estimator of $\boldsymbol{\gamma}$ (and $\boldsymbol{\beta}$ as well).

4 Θ -IPOD

To deal with masking and swamping in the presence of multiple outliers, we consider the iterative procedure described in Section 3 using a general Θ operator, referred to as Θ -IPOD. A Θ -IPOD estimate is a limit point of $(\boldsymbol{\beta}^{(j)}, \boldsymbol{\gamma}^{(j)})$, denoted by $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$. Some-

what surprisingly, simply replacing soft-thresholding by hard-thresholding (henceforth hard-IPOD) resolves the masking and swamping problem, for the challenging HBK example problem. After picking $\hat{\sigma}$ from LTS and using a zero start we obtained

$$\hat{\gamma}_{1:10} = \left(9.7 \quad 10.2 \quad 10.4 \quad 9.7 \quad 10.1 \quad 10.0 \quad 10.8 \quad 10.4 \quad 9.8 \quad 10.1 \right)^\top$$

and $\hat{\gamma}_{11:75} = \mathbf{0}$ which perfectly detects the true outliers. The coefficient estimate $\hat{\beta}$ from hard-IPOD directly gives the OLS estimate computed from the clean observations.

Some questions naturally arise: What optimization problem is Θ -IPOD trying to solve? For an arbitrary Θ , does Θ -IPOD converge at all? Or, under what conditions does Θ -IPOD converge? To answer these questions, we limit our discussions of Θ to thresholding rules.

4.1 IPOD, Penalized Regressions, and M -estimators

We begin by defining the class of threshold functions we will study. It includes well-known thresholds such as soft and hard thresholding, SCAD, and Tukey's bisquare.

Definition 4.1 (Threshold function). A threshold function is a real valued function $\Theta(t; \lambda)$ defined for $-\infty < t < \infty$ with λ as the parameter ($0 \leq \lambda < \infty$) such that

- 1) $\Theta(-t; \lambda) = -\Theta(t; \lambda)$,
- 2) $\Theta(t; \lambda) \leq \Theta(t'; \lambda)$ for $t \leq t'$,
- 3) $\lim_{t \rightarrow \infty} \Theta(t; \lambda) = \infty$, and
- 4) $0 \leq \Theta(t; \lambda) \leq t$ for $0 \leq t < \infty$.

In words, $\Theta(\cdot; \lambda)$ is an odd monotone unbounded shrinkage rule for t , at any λ .

A vector version of Θ is defined componentwise if either t or λ are replaced by vectors. When both t and λ are vectors, we assume they have the same dimension.

Algorithm 1 Θ -IPOD

given $\mathbf{X} \in \mathbb{R}^{n \times p}$, $\mathbf{y} \in \mathbb{R}^n$, $\varepsilon > 0$, $\boldsymbol{\lambda} \in [0, \infty)^n$, robust pilot estimate $\boldsymbol{\beta}^{(0)} \in \mathbb{R}^p$
 $j \leftarrow 0$, $\boldsymbol{\gamma}^{(j)} \leftarrow \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}$, converged \leftarrow FALSE
 $(Q, R) \leftarrow \text{QRdecomp}(\mathbf{X})$
// $\mathbf{X} = QR$, $R \in \mathbb{R}^{p \times p}$ is upper triangular, $Q \in \mathbb{R}^{n \times p}$ and $Q^\top Q = I_p$
while not converged **do**
 $\mathbf{y}^{\text{adj}} \leftarrow \mathbf{y} - \boldsymbol{\gamma}^{(j)}$
 $\boldsymbol{\beta}^{(j)} \leftarrow R^{-1}Q^\top \mathbf{y}^{\text{adj}}$
 $\mathbf{r}^{(j)} \leftarrow \mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(j)}$
 $\boldsymbol{\gamma}^{(j+1)} \leftarrow \Theta(\mathbf{r}^{(j)}; \boldsymbol{\lambda})$
 converged $\leftarrow \|\boldsymbol{\gamma}^{(j+1)} - \boldsymbol{\gamma}^{(j)}\|_\infty < \varepsilon$
 $j \leftarrow j + 1$
end while
deliver $\hat{\boldsymbol{\beta}} = \boldsymbol{\beta}^{(j-1)}$, $\hat{\boldsymbol{\gamma}} = \boldsymbol{\gamma}^{(j)}$

Huber's soft-thresholding rule corresponds to an absolute error, or L_1 penalty. More generally, for any thresholding rule $\Theta(\cdot; \lambda)$, a corresponding penalty function $P = P_\Theta$ can be defined. There may be multiple penalty functions for a given threshold as demonstrated by hard thresholding (4.7) below. The following *three-step construction* finds the penalty with the smallest curvature (Antoniadis 2007; She 2009):

$$\begin{aligned}\Theta^{-1}(u; \lambda) &= \sup\{t : \Theta(t; \lambda) \leq u\}, \\ s(u; \lambda) &= \Theta^{-1}(u; \lambda) - u, \quad \text{and} \\ P(\theta; \lambda) &= P_\Theta(\theta; \lambda) = \int_0^{|\theta|} s(u; \lambda) \, du,\end{aligned}\tag{4.1}$$

where $u \geq 0$ holds throughout (4.1). The constructed penalty P_Θ is nonnegative and is continuous in θ .

Theorem 4.1. *Let Θ be a thresholding rule as given by Definition 4.1 and let P_Θ be the corresponding penalty defined in (4.1). Given $\lambda_i \geq 0$, the objective function is*

defined by

$$f_P(\boldsymbol{\beta}, \boldsymbol{\gamma}) \equiv \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \boldsymbol{\gamma}\|_2^2 + \sum_{i=1}^n P(\gamma_i; \lambda_i), \quad (4.2)$$

where $P(\cdot; \cdot)$ is any function satisfying $P(\theta; \lambda) - P(0; \lambda) = P_\Theta(\theta; \lambda) + q(\theta; \lambda)$ where $q(\cdot; \lambda)$ is nonnegative and $q(\Theta(\theta; \lambda)) = 0$ for all θ . Then the Θ -IPOD iteration sequence $(\boldsymbol{\beta}^{(j)}, \boldsymbol{\gamma}^{(j)})$ satisfies

$$f_P(\boldsymbol{\beta}^{(j)}, \boldsymbol{\gamma}^{(j)}) \geq f_P(\boldsymbol{\beta}^{(j)}, \boldsymbol{\gamma}^{(j+1)}) \geq f_P(\boldsymbol{\beta}^{(j+1)}, \boldsymbol{\gamma}^{(j+1)}), \quad \text{for any } j \geq 0. \quad (4.3)$$

The function q will often be zero, but we use non-zero q below to demonstrate that multiple penalties (infinitely many, as a matter of fact) yield hard thresholding, including the L_0 -penalty. Theorem 4.1 shows that Θ -IPOD converges. Any limit point of $(\boldsymbol{\beta}^{(j)}, \boldsymbol{\gamma}^{(j)})$ must be a stationary point of (4.2). Θ -IPOD also gives a general connection between penalized regression (4.2) and M -estimators. Recall that an M -estimator is defined to be a solution to the score equation

$$\sum_{i=1}^n \psi \left(\frac{y_i - \mathbf{x}_i^\top \boldsymbol{\beta}}{\sigma}; \lambda \right) \mathbf{x}_i = \mathbf{0}, \quad (4.4)$$

where λ is a general parameter of the ψ function. Although $\boldsymbol{\beta}$ and σ can be simultaneously estimated by Huber's Proposal 2 (Huber 1981), a more common practice is to fix σ at an initial robust estimate and then optimize over $\boldsymbol{\beta}$ (Hampel et al. 1986). Unless otherwise specified, we consider equation (4.4) as constraining $\boldsymbol{\beta}$ with σ fixed.

Proposition 4.1. *For any thresholding rule $\Theta(\cdot; \lambda)$, for any Θ -IPOD estimate $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$, $\hat{\boldsymbol{\beta}}$ is an M -estimate associated with ψ , as long as (Θ, ψ) satisfies*

$$\Theta(t; \lambda) + \psi(t; \lambda) = t, \quad \forall t. \quad (4.5)$$

We have mentioned that Huber’s method or soft-IPOD behaves poorly in outlier detection. The problem is that it never rejects gross outliers. To reject gross outliers, re-descending ψ -functions are advocated, corresponding to a class of thresholdings offering little shrinkage for large components, or using nonconvex penalties for solving the sparsity problem (4.2). The differences between Θ -IPOD and the corresponding M -estimator are as follows:

- (i) M -estimators focus on robust estimation of β . For an explicit sparse γ estimate, a cutoff value is usually needed for the residuals. Minimizing f_P from equation (4.2) directly yields a sparse $\hat{\gamma}$ for outlier detection and a robust $\hat{\beta}$.
- (ii) Instead of designing a robust loss function in M -estimators, (4.2) considers a penalty function; λ is not a criterion (loss) parameter but a regularization parameter that we will tune in a data-dependent way.

Figure 2 illustrates some of the better known threshold functions along with their corresponding penalties, ψ -functions and loss functions. In this article, we make use of the following formulas

$$\Theta_{\text{soft}}(x; \lambda) = \begin{cases} 0, & \text{if } |x| \leq \lambda \\ x - \text{sgn}(x)\lambda, & \text{if } |x| > \lambda, \end{cases} \quad (4.6)$$

$$\Theta_{\text{hard}}(x; \lambda) = \begin{cases} 0, & \text{if } |x| \leq \lambda \\ x, & \text{if } |x| > \lambda, \end{cases} \quad (4.7)$$

for soft and hard thresholding, respectively. The usual penalty for hard thresholding is $P(x; \lambda) = \max(\lambda|x| - x^2/2, \lambda^2/2)$. Theorem 4.1 justifies use of the L_0 -penalty

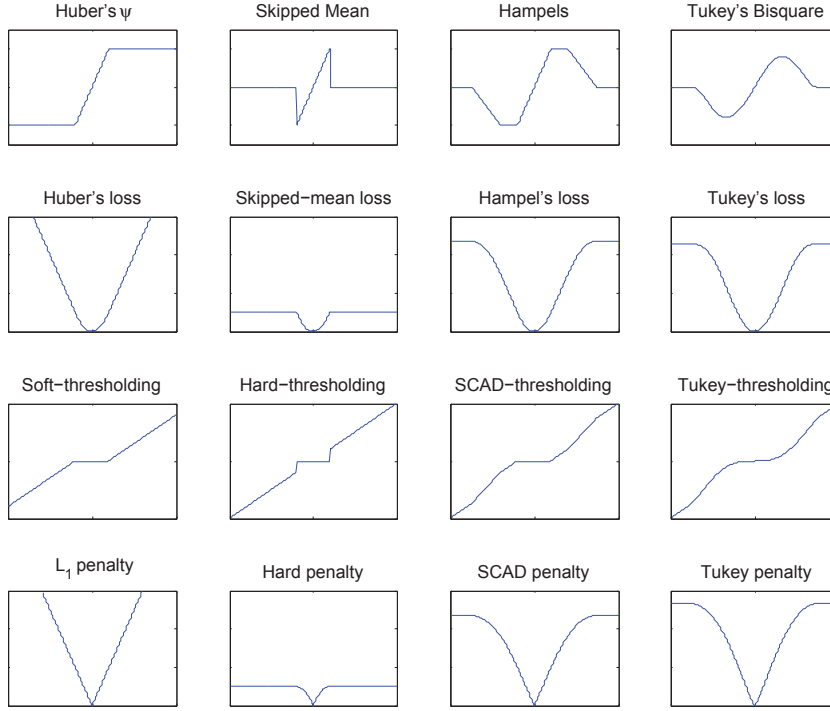


Figure 2: The ψ functions (top row), loss functions (second row), thresholding functions (third row), and penalty functions (bottom row) for some well known thresholding rules. SCAD-thresholding (Antoniadis & Fan 2001) is a special case of the Hampel rule.

$P(x; \lambda) = \lambda^2/2 \cdot I_{x \neq 0}$ using

$$q(x; \lambda) = \begin{cases} \frac{(\lambda - |x|)^2}{2}, & \text{if } 0 < |x| < \lambda \\ 0, & \text{if } x = 0 \text{ or } |x| \geq \lambda. \end{cases}$$

4.2 Θ -IPOD and TISP

Algorithm 1 can be simplified. The Θ -IPOD iteration can be carried out without recomputing $\beta^{(j)}$ at each iteration. We need only update γ via

$$\gamma^{(j+1)} = \Theta(\mathbf{H}\gamma^{(j)} + (\mathbf{I} - \mathbf{H})\mathbf{y}; \lambda), \quad (4.8)$$

at each iteration, where $\lambda_i = \lambda\sqrt{1 - h_i}$. The multiplication $(\mathbf{I} - \mathbf{H})\mathbf{y}$ in (4.8) can also be precomputed. After getting the final $\hat{\gamma}$, we can estimate β by OLS. The resulting

Algorithm 2 Simplified Θ -IPOD

given $\mathbf{X} \in \mathbb{R}^{n \times p}$, $\mathbf{y} \in \mathbb{R}^n$, $\lambda > 0$, $\varepsilon > 0$, and threshold $\Theta(\cdot; \cdot)$
 $j \leftarrow 0$, $\boldsymbol{\gamma}^{(j)} \leftarrow \boldsymbol{\gamma}^{(0)}$, converged \leftarrow FALSE
 $\mathbf{H} \leftarrow \text{HatMatrix}(\mathbf{X})$, $\mathbf{r} \leftarrow \mathbf{y} - \mathbf{H}\mathbf{y}$
while not converged **do**
 $\boldsymbol{\gamma}^{(j+1)} \leftarrow \Theta(\mathbf{H}\boldsymbol{\gamma}^{(j)} + \mathbf{r}; \boldsymbol{\lambda})$
converged $\leftarrow \|\boldsymbol{\gamma}^{(j+1)} - \boldsymbol{\gamma}^{(j)}\|_\infty < \varepsilon$
 $j \leftarrow j + 1$
end while
deliver $\hat{\boldsymbol{\gamma}} = \boldsymbol{\gamma}^{(j)}$, $\hat{\boldsymbol{\beta}} = \text{OLScoef}(\mathbf{X}, \mathbf{y} - \hat{\boldsymbol{\gamma}})$

\mathbf{X} must have full rank. $\boldsymbol{\gamma}^{(0)}$ can be $\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(0)}$ for a robust pilot estimate $\boldsymbol{\beta}^{(0)}$.

simplified Θ -IPOD algorithm is given in Algorithm 2. For any given thresholding rule Θ , let $f_P(\boldsymbol{\gamma}) \equiv \frac{1}{2}\|(\mathbf{I} - \mathbf{H})(\mathbf{y} - \boldsymbol{\gamma})\|_2^2 + \sum_{i=1}^n P(\gamma_i; \lambda_i)$, where P can be any penalty function satisfying the conditions in Theorem 4.1. Then $f_P(\boldsymbol{\gamma}^{(j+1)}) \leq f_P(\boldsymbol{\gamma}^{(j)})$ for the Θ -IPOD iterates $\boldsymbol{\gamma}^{(j)}$, $j \geq 0$.

Setting up for Algorithm 2 costs $O(np^2)$ for a dense regression. The dominant cost in a given iteration comes from computing $\mathbf{H}\boldsymbol{\gamma}^{(j)}$. Given a QR decomposition of \mathbf{H} we can compute that matrix product in $O(np)$ work as $Q(Q^T\boldsymbol{\gamma}^{(j)})$. The cost of the update could be even less if $\boldsymbol{\gamma}^{(j)}$ has fewer than p nonzero entries and one maintains the dense matrix \mathbf{H} .

To give (4.8) another explanation, suppose the spectral decomposition of the hat matrix \mathbf{H} is given by $\mathbf{H} = \mathbf{U}\mathbf{D}\mathbf{U}^T$. Define an index set $c = \{i : D_{ii} = 0\}$ and let \mathbf{U}_c be formed by taking the corresponding columns of \mathbf{U} . Then a *reduced model* can be obtained from the mean shift outlier model (1.2)

$$\tilde{\mathbf{y}} = \mathbf{A}\boldsymbol{\gamma} + \boldsymbol{\epsilon}', \quad \boldsymbol{\epsilon}' \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I}_{(n-p) \times (n-p)}), \quad (4.9)$$

where $\tilde{\mathbf{y}} = \mathbf{U}_c^T\mathbf{y}$ and $\mathbf{A} = \mathbf{U}_c^T \in \mathbb{R}^{(n-p) \times n}$. The term $\mathbf{X}\boldsymbol{\beta}$ has disappeared from the reduced model. For a special \mathbf{U}_c , the vector $\tilde{\mathbf{y}}$ has the BLUS residuals of Theil (1965). The regression model (4.9) is a $p > n$ sparsity problem. It is also a wavelet ap-

proximation problem that Antoniadis & Fan (2001) studied in the context of wavelet denoising, because \mathbf{A} satisfies $\mathbf{A}\mathbf{A}^\top = \mathbf{I}$. Furthermore, the regularized one-step estimator (ROSE) (Antoniadis & Fan 2001) is the first step of Θ -IPOD.

We can build a connection between the reduced model and simplified Θ -IPOD. She (2009) proposed a class of thresholding-based iterative selection procedures (TISP) for model selection and shrinkage. Θ -TISP for solving the sparsity problem (4.9) is given by

$$\boldsymbol{\gamma}^{(j+1)} = \Theta(\mathbf{A}^\top \tilde{\mathbf{y}}/k_0^2 + (\mathbf{I} - \mathbf{A}^\top \mathbf{A}/k_0^2)\boldsymbol{\gamma}^{(j)}; \boldsymbol{\lambda}/k_0^2) \quad (4.10)$$

with k_0 equal to the largest singular value of the Gram matrix $\mathbf{A}^\top \mathbf{A} = \mathbf{H}$. The iteration (4.10) reduces exactly to (4.8). Therefore, all TISP studies apply to the Θ -IPOD algorithm. For example, the TISP convergence theorem can be used to establish a version of Theorem 4.1, and the nonasymptotic probability bounds for sparsity recovery reveal masking and swamping errors. In particular, the advocated hard-thresholding-like Θ in TISP corresponds to a redescending ψ in our outlier identification problem.

The simplified procedure (4.8) is easy to implement and is computationally efficient, because the iteration does not involve complicated operations like matrix inversion. Model (4.9) is simpler than the original (1.2) because $\boldsymbol{\beta}$ does not appear and all observations are *clean*. They have non-outlying errors because we have moved the outlier variables into the regression. Using this characterization of $\hat{\boldsymbol{\gamma}}$, it is not difficult to show that the IPOD-estimate $\hat{\boldsymbol{\beta}}$ satisfies the regression, scale, and affine equivariant properties (Hampel et al. 1986) desirable for a good robust regression estimator:

- (i) $\hat{\boldsymbol{\beta}}(\mathbf{X}, \mathbf{y} + \mathbf{X}\boldsymbol{\eta}) = \hat{\boldsymbol{\beta}}(\mathbf{X}, \mathbf{y}) + \boldsymbol{\eta}, \quad \forall \boldsymbol{\eta} \in \mathbb{R}^p;$
- (ii) $\hat{\boldsymbol{\beta}}(\mathbf{X}, c\mathbf{y}) = c\hat{\boldsymbol{\beta}}(\mathbf{X}, \mathbf{y}), \quad \forall c \in \mathbb{R};$

(iii) $\hat{\beta}(\mathbf{X}\mathbf{C}, \mathbf{y}) = \mathbf{C}^{-1}\hat{\beta}(\mathbf{X}, \mathbf{y})$, for any nonsingular \mathbf{C} .

We are making here a mild assumption that the initial robust estimate is equivariant, as LTS, S and PY (Peña & Yohai 1999) are, and we're ignoring the possible effect of convergence criteria on equivariance.

It remains to select λ . The clean dataset $(\mathbf{A}, \tilde{\mathbf{y}})$ in (4.9) has uncorrelated noise contamination, so we can use it to tune the parameters via BIC and related methods; see Section 6 for details.

5 Θ -IPOD vs. IRLS

We consider some computational issues in this section. As Proposition 4.1 suggests, Θ -IPOD solves an M -estimation problem. The standard fitting algorithm for M -estimates is the well-known iteratively re-weighted least squares (IRLS). Let $w(t; \lambda) = \psi(t; \lambda)/t$, taking $0/0 = 0$ if necessary. The ψ -equation (4.4) that defines an M -estimator can be rewritten as $\sum_{i=1}^n w(r_i/\sigma; \lambda) r_i \mathbf{x}_i = \mathbf{0}$, where $r_i = y_i - \mathbf{x}_i^\top \beta$. Accordingly, an M -estimate corresponds to a weighted LS estimate as is well known. These multiplicative weights can help downweight the bad observations. Iteratively updating the weights yields the IRLS algorithm, which is the most common method for computing M -estimates. Model (1.2) indicates that M -estimation can also be characterized through **additive** effects on all observations.

We performed a simulation to compare the speed of Θ -IPOD to that of IRLS. The observations were generated according to $\mathbf{y} = \mathbf{X}\beta + \gamma + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. The predictor matrix \mathbf{X} is constructed as follows: first let $\mathbf{X} = \mathbf{U}\Sigma^{1/2}$, where $U_{ij} \stackrel{\text{iid}}{\sim} U(-15, 15)$ and $\Sigma_{ij} = \rho^{1_{i \neq j}}$ with $\rho = 0.5$; then modify the first O rows to be leverage points given by $L \cdot [1, \dots, 1]$. We consider all nine cases with $L \in \{15, 20, 30\}$ and $O \in \{5, 10, 20\}$. Three more cases correspond to additive outliers at O points that were not leverage points (no rows of \mathbf{X} changed). The shift vector is given by $\gamma = (\{8\}^O, \{0\}^{n-O})^\top$.

Given n and p , we report the total cost of computing all M -estimates for these different combinations of the number of outliers (O) and leverage value (L), each combination simulated 10 times. The scale parameter (λ) decreased from $\|(\mathbf{I} - \mathbf{H})\mathbf{y} ./ \sqrt{\text{diag}(\mathbf{I} - \mathbf{H})}\|_\infty$ to 0.5 with fixed step size -0.1 , where $./$ stands for elementwise division. The upper limit is the largest possible standardized residual. Empirically, the lower bound 0.5 yields approximately half of γ_i nonzero. Also for $\lambda < 0.5$ IRLS often encountered a singular WLS during the iteration, or take exceptionally long time to converge, and so could not be compared to Θ -IPOD. We used IRLS (with fixed $\sigma = 1$, as an oracle would have) and simplified Θ -IPOD. The common convergence criterion was $\|\boldsymbol{\gamma}^{(j+1)} - \boldsymbol{\gamma}^{(j)}\|_\infty < 10^{-4}$. We studied all sample sizes $n \in \{30, 50, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ and we took $p = n/10$. The CPU times (in seconds) are plotted against the sample size in Figure 3.

The speed advantage of Θ -IPOD over IRLS in this simulation tended to increase with n and p . At $n = 1000$ the hard-IPOD algorithm required about 2 or 3 times (averaged over all lambda values) as many iterations as IRLS but was about 10 times faster than IRLS, due to faster iterations. For small n , we saw little speed difference. As remarked above, IRLS was sometimes unstable, with singular WLS problems arising for redescending ψ , when fewer than p weights were nonzero. Such cases cause only a small problem for Θ -IPOD: the update $\mathbf{H}\boldsymbol{\gamma}^{(j)}$ cannot then take advantage of sparsity, but it is very stable.

Finally, it is worth mentioning that although we attained impressive computational gain over the popular IRLS, we consider speed to be secondary compared to robustness (and the speed advantage of Θ -IPOD can be moot when the preliminary method is very expensive).




Figure 3: Computational time comparison between Θ -IPOD and IRLS for problems of different sizes. The IRLS computational times are represented by red squares, and the Θ -IPOD computational times by green circles. The plots, from left to right, top to bottom, correspond to the examples in Figure 2 respectively.

6 Parameter Tuning in Outlier Detection

The parameter λ in an M -estimators (4.4) is often chosen to be a constant (for all n), based on either efficiency or breakdown of the estimator. The value $2.5\hat{\sigma}$ is popular (Rousseeuw & Leroy 1987; Wilcox 2005; Maronna et al. 2006). But as mentioned in Section 3, even with no outliers and $\mathbf{X} = \mathbf{I}$, a constant λ independent of n is far from optimal (Donoho & Johnstone 1994). It is also hard, in robust regressions, to select the cutoff value η at which to identify outliers, because the residual distri-

bution is usually unknown. Gervini & Yohai (2002) base an asymptotically efficient choice for η on a Kolmogorov-Smirnov statistic, but they need to assume that the standardized robust residuals are IID $\mathcal{N}(0, 1)$.

For Θ -IPOD, these two issues correspond to the problem of tuning λ . In most M -estimators of practical interest, the ψ function satisfies $\psi(t/\sigma; \lambda) \propto \psi(t; \lambda\sigma)$. Then we may set $\sigma = 1$ and carefully tune λ . Even with this simplification, a *data-dependent* λ is difficult to choose using cross-validation on the (\mathbf{X}, \mathbf{y}) data. The reason is that a large prediction error can indicate either a suboptimal β or an outlier. Fortunately, turning to the reduced model (4.9) greatly mitigates the problem because all (transformed) observations are *clean*. Then BIC can be applied to $(\mathbf{A}, \tilde{\mathbf{y}})$ if the proportion of outliers is not large.

Specifically, because all candidate estimates lie along the Θ -IPOD solution path, we design a local BIC to apply BIC on a proper local interval of the degrees of freedom (DF). First we generate the hard-IPOD solution path by decreasing λ from $\|(\mathbf{I} - \mathbf{H})\mathbf{y}\cdot/\sqrt{\text{diag}(\mathbf{I} - \mathbf{H})}\|_\infty$ to 0. Given λ and the corresponding estimate $\hat{\gamma}(\lambda)$, let $nz(\lambda) = \{i : \hat{\gamma}_i(\lambda) \neq 0\}$. We rely on model (4.9) and study its variable selection to give the correct form of BIC. For hard-IPOD, $\hat{\gamma}_{nz}$ is an OLS estimate with one parameter per detected outlier and the degrees of freedom are given by $\text{DF}(\lambda) = |nz(\lambda)|$. We use BIC with a slight modification:

$$\text{BIC}^*(\lambda) = m \log(\text{RSS}/m) + k(\log(m) + 1), \quad (6.1)$$

where $m = n - p$, $\text{RSS} = \|\tilde{\mathbf{y}} - \mathbf{A}\hat{\gamma}\|_2^2 = \|(\mathbf{I} - \mathbf{H})(\mathbf{y} - \hat{\gamma})\|_2^2$ and $k(\lambda) = \text{DF}(\lambda) + 1$. We have taken $k = \text{DF} + 1$ to account for the noise scale parameter, and BIC^* uses $\log(m) + 1$ instead of BIC 's $\log(m)$ because we have found this change to be better empirically. Note that the sample size m here is smaller than the dimension of γ . According to Chen & Chen (2008), similar modifications are necessary to preserve

the model selection properties of BIC for problems with fewer observations than predictors. We would like to apply BIC^* on a proper local interval of DF given by $[\nu_L, \nu_U]$. We take $\nu_U \leq n/2$, assuming that the proportion of outliers is under 50%. The $(\text{DF}(\lambda), \text{BIC}^*(\lambda))$ curve sometimes has narrow local minima near the ends of the λ range. To counter that effect we fit a smoothing spline to the set of data points $(\text{DF}(\lambda), \text{BIC}^*(\lambda))$ and chose the local minimum with the largest neighborhood. The neighborhood size can be determined using the local maxima of the smoothing spline. Of course there may be other reasonable ways to counter that problem.

We carried out simulation experiments to test the performance of the tuned Θ -IPOD. The matrix \mathbf{X} was generated the same way as in Section 5 using dimension $p \in \{15, 50\}$, $n = 1000$ observations of which the first $O \in \{200, 150, 100, 50, 10\}$ were outliers at the highly leveraged location given by $L \in \{15, 20\}$ times a vector of 1s. The outliers were generated by a mean shift $\boldsymbol{\gamma} = (\{5\}^O, \{0\}^{n-O})^\top$ added to \mathbf{y} . Because Θ -IPOD is affine, regression, and scale equivariant and the outliers are at a single x value, we may set $\boldsymbol{\beta} = \mathbf{0}$ without loss of generality. The intercept term is always included in the modeling.

Five outlier detection methods were considered for comparison: hard-IPOD (tuned), MM-estimator, Gervini & Yohai's (2002) fully efficient one-step procedure (denoted by GY), the compound estimator S1S, and the LTS. (The direct procedures, such as Hadi & Simonoff (1997), Peña & Yohai (1995), behaved poorly and their results were not reported.) The S-PLUS Robust library provides implementations of MM, GY, and LTS; for the implementation of S1S, we refer to Wilcox (2005). Robust also provides a default initial estimate with high breakdown point (not necessarily efficient), which is used in the first three algorithms in our experiments: when $p = 15$, it is the S-estimate computed via random resampling; when $p = 50$, it is the estimate from the fast PY procedure. Since it is well known that the initial estimate is outper-

formed by MM in terms of estimation efficiency and robustness, both theoretically and empirically (Yohai 1987; Salibián-Barrera & Yohai 2006; Peña & Yohai 1999), its results are not reported. (All routines are available for the R language (R Development Core Team 2007) as well. The package `robust` provides an R version of the Insightful `Robust` library.)

All methods apart from Θ -IPOD require a cutoff value to identify which residuals are outliers. We applied the fully efficient procedure which performs at least as well as the fixed choice of $\eta = 2.5$ in various situations (Gervini & Yohai 2002).

Each model was simulated 100 times. We report outlier identification results for each algorithm using three benchmark measures:

- M** the mean masking probability (fraction of undetected true outliers),
- S** the mean swamping probability (fraction of good points labeled as outliers),
- JD** the joint outlier detection rate (fraction of simulations with 0 masking).

In outlier detection, masking is more serious than swamping. The former can cause gross distortions while the latter is often just a matter of lost efficiency.

Ideally, $M \approx 0$, $S \approx 0$, and $JD \approx 100\%$. JD is the most important measure on easier problems while M makes the most sense for hard problems. The simulation results are summarized in Tables 1 and 2. Figures 4 and 5 present M and JD for $p = 15$ and 50 respectively. While our main purpose is to identify outliers, a robust coefficient estimate $\hat{\beta}$ can be easily obtained from Θ -IPOD. The MSE in β for $p = 15$ is shown in Figure 6. All methods had small slope errors, though Θ -IPOD and GY performed best. The results for $p = 50$ (not shown) are similar.

MM and GY are two standard methods provided by the S-PLUS `Robust` library. Nevertheless, seen from the tables, the MM-estimator, though perhaps most popular in robust analysis, does not yield good identification results when the outliers have high leverage values and the number of outliers is not small (for example, $O = 200$,




Figure 4: Masking (M) and joint detection (JD) results when $p = 15$.




Figure 5: Masking (M) and joint detection (JD) when $p = 50$.

	O = 200			O = 100			O = 50			O = 20			O = 10		
	JD	M	S	JD	M	S	JD	M	S	JD	M	S	JD	M	S
No L :															
IPOD	43	0.4	2.1	38	0.6	1.6	47	0.8	1.2	61	0.9	0.9	94	0.6	0.7
MM	0	11.9	0.0	0	8.8	0.0	0	6.9	0.1	9	5.2	0.1	70	3.4	0.3
GY	0	12.7	0.0	0	9.1	0.0	0	7.2	0.1	6	5.5	0.1	66	3.9	0.3
S1S	0	34.3	0.0	0	14.3	0.1	0	8.9	0.1	3	6.6	0.1	54	5.5	0.2
LTS	0	66.3	0.0	0	15.4	0.1	2	3.7	0.3	28	2.3	0.5	88	1.2	0.7
L=15:															
IPOD	51	0.4	2.2	49	0.5	1.6	55	0.6	1.2	63	0.8	0.9	92	0.8	0.7
MM	0	57.1	1.2	0	9.7	0.1	0	7.0	0.1	8	5.1	0.1	68	3.3	0.3
GY	0	18.5	0.0	0	10.6	0.1	0	7.6	0.1	7	5.4	0.1	67	3.6	0.3
S1S	0	100.0	0.0	0	58.9	0.3	0	15.1	0.2	1	8.9	0.2	59	5.1	0.2
LTS	0	99.9	0.1	0	49.6	0.1	4	3.4	0.3	40	1.8	0.5	85	1.5	0.6
L=20:															
IPOD	49	0.4	2.1	49	0.6	1.6	52	0.7	1.2	63	0.9	0.9	92	0.8	0.7
MM	0	96.0	1.3	0	20.0	0.2	0	7.3	0.1	7	5.3	0.1	68	3.4	0.3
GY	0	24.1	0.1	0	12.2	0.1	0	8.2	0.1	8	5.6	0.1	67	3.6	0.3
S1S	0	100.0	0.1	0	95.5	0.1	0	20.1	0.2	1	11.0	0.2	58	5.7	0.2
LTS	0	100.0	0.5	0	89.5	0.1	4	4.0	0.3	41	2.1	0.5	86	1.4	0.7

Table 1: Outlier identification results on simulated data with $p = 15$. Five methods are compared: hard-IPOD (tuned), MM-estimator, Gervini & Yohai’s fully efficient one-step procedure, S1S, and LTS. Leverage L and number of outliers O are described in the text as are the quality measures JD (joint detection), M (masking) and S (swamping).

$L = 15, 20$). GY improves MM a lot in this situation and gives similar results otherwise. The experiments also show that S1S behaves poorly in outlier detection and LTS works better in the presence of $\leq 5\%$ outliers. Unfortunately, all four methods have high masking probabilities and very low joint identification rates, which become worse for large p . The Θ -IPOD method dominates them significantly for masking and joint detection. **Recall that each experiment was repeated 100 times. With the standard errors available, we can perform a paired t -test for comparing say the masking probability between Θ -IPOD and other methods. We found the t -statistics are mostly greater than 3 and become huge when O increases. In particular, LTS vs. Θ -IPOD yields the smallest t -statistics (~ 1.2) when $O = 10$ but also the largest (~ 1700) when $O = 200$. While**

	O = 200			O = 100			O = 50			O = 20			O = 10		
	JD	M	S	JD	M	S	JD	M	S	JD	M	S	JD	M	S
No L :															
IPOD	32	0.6	2.4	35	0.7	1.7	40	1.0	1.3	50	1.3	0.9	90	1	0.7
MM	0	13.9	0.0	0	9.2	0.1	0	6.4	0.1	15	4.0	0.3	85	1.5	0.6
GY	0	15.0	0.0	0	10.0	0.1	0	7.0	0.1	12	4.4	0.2	82	2	0.5
S1S	0	43.8	0.1	0	17.8	0.1	0	9.3	0.1	5	6.1	0.2	68	3.7	0.4
LTS	0	62.8	0.0	0	20.9	0.1	1	4.9	0.3	20	2.3	0.7	87	1.4	1.1
L=15:															
IPOD	44	0.5	2.4	39	0.7	1.7	47	0.9	1.3	60	1.1	0.9	94	0.6	0.7
MM	0	75.1	2.0	0	9.6	0.1	0	6.2	0.1	18	3.6	0.3	86	1.4	0.6
GY	0	20.3	0.1	0	10.5	0.1	0	6.7	0.1	13	3.9	0.2	81	1.9	0.5
S1S	0	100.0	0.0	0	60.8	0.3	0	15.1	0.2	5	7.5	0.2	68	3.7	0.4
LTS	0	100.0	0.7	0	54.4	0.1	2	4.0	0.3	40	2.0	0.7	90	1	1.0
L=20:															
IPOD	41	1.5	2.4	38	1.8	1.7	49	0.9	1.3	60	1.2	0.9	93	0.7	0.7
MM	0	97.1	1.3	0	21.2	0.3	1	6.4	0.1	19	3.7	0.3	87	1.3	0.6
GY	0	27.4	0.1	0	12.4	0.1	1	7.2	0.1	12	4.1	0.2	82	1.9	0.5
S1S	0	100.0	0.3	0	99.6	0.1	0	20.2	0.2	2	10.1	0.3	63	4.1	0.4
LTS	0	100.0	1.4	0	98.8	0.4	1	4.6	0.3	43	1.8	0.8	89	1.1	1.1

Table 2: Outlier identification results on the simulation data where $p = 50$. Methods, conditions and results are the same as in Table 1.

Θ -IPOD does much better on masking it does slightly worse on swamping. This is an acceptable tradeoff because masking is far more important.

7 Outlier Detection with $p > n$

Here we extend outlier detection to problems with $p > n$ including of course some high-dimensional problems. Needless to say, this context is more challenging but has diverse modern applications in signal processing, bioinformatics, finance, and neuroscience among others. Performing the task of outlier identification for data with $p > n$ or even $p \gg n$ goes beyond the traditional robust analysis which requires a large number of observations relative to the dimensionality.

The large- p version of Θ -IPOD is indeed possible under our general framework (1.2). If $\beta \in \mathbb{R}^p$ is also assumed to be *sparse* in the mean shift outlier model, one can




Figure 6: Coefficient estimation errors when $p = 15$.

directly work in an augmented data space $\mathbf{y} \sim \mathbf{B}\boldsymbol{\xi}$ with $\mathbf{B} = [\mathbf{X} \ \mathbf{I}]$ and $\boldsymbol{\xi} = \begin{bmatrix} \boldsymbol{\beta} \\ \boldsymbol{\gamma} \end{bmatrix}$ to study this large- p sparsity problem. Concretely, the TISP technique can be used to derive the Θ -IPOD iteration $\boldsymbol{\xi}^{(j+1)} = \Theta(\mathbf{B}^\top \mathbf{y} / k_0^2 + (\mathbf{I} - \mathbf{B}^\top \mathbf{B} / k_0^2) \boldsymbol{\xi}^{(j)}; \lambda / k_0^2)$, or

$$\begin{aligned} \boldsymbol{\beta}^{(j+1)} &\leftarrow \Theta \left(\boldsymbol{\beta}^{(j)} + \frac{1}{k_0^2} \mathbf{X}^\top \mathbf{y} - \frac{1}{k_0^2} \mathbf{X}^\top (\mathbf{X} \boldsymbol{\beta}^{(j)} + \boldsymbol{\gamma}^{(j)}); \frac{\lambda}{k_0^2} \right) \quad \text{and} \\ \boldsymbol{\gamma}^{(j+1)} &\leftarrow \Theta \left(\left(1 - \frac{1}{k_0^2} \right) \boldsymbol{\gamma}^{(j)} + \frac{1}{k_0^2} \mathbf{y} - \frac{1}{k_0^2} \mathbf{X} \boldsymbol{\beta}^{(j)}; \frac{\lambda}{k_0^2} \right). \end{aligned} \quad (7.1)$$

The convergence of (7.1) is guaranteed even for $n < p$ (She 2009), as long as $[\mathbf{X} \ \mathbf{I}]$ is scaled properly, which here amounts to taking $k_0 \geq \sqrt{\|\mathbf{X}\|_2^2 + 1}$. Let $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\gamma}}$ denote the estimates from (7.1). The nonzero components of $\hat{\boldsymbol{\beta}}$ locate the relevant predictors while $\hat{\boldsymbol{\gamma}}$ identifies the outliers and measures their outlyingness. (It is straightforward to apply (7.1) to *any* p for simultaneous variable selection and outlier detection.)

Not surprisingly, L_1 (or soft-thresholding) fails again for this challenging sparsity problem and thresholdings corresponding to redescending ψ 's should be used in (7.1). (It is natural to consider the elastic net (Zou & Hastie 2005) for this problem as well. But the elastic net has a convex criterion, and it can be shown to have breakdown zero.) Since it is usually unknown how severe the outliers are, one may adopt the

following hybrid of hard-thresholding and ridge-thresholding

$$\Theta(x; \lambda, \eta) = \begin{cases} 0, & \text{if } |x| < \lambda \\ \frac{x}{1+\eta}, & \text{if } |x| \geq \lambda \end{cases} \quad (7.2)$$

referred to as the hybrid-thresholding in She (2009), or hard-ridge thresholding in this paper. The corresponding penalty from the three-step construction is

$$P_{\Theta}(x; \lambda, \eta) = \begin{cases} -\frac{1}{2}x^2 + \lambda|x|, & \text{if } |x| < \frac{\lambda}{1+\eta} \\ \frac{1}{2}\eta x^2 + \frac{1}{2}\frac{\lambda^2}{1+\eta}, & \text{if } |x| \geq \frac{\lambda}{1+\eta}. \end{cases}$$

Setting $q(x; \lambda, \eta) = \frac{1+\eta}{2}(|x| - \lambda)^2 1_{0 < |x| < \lambda}$, we get an alternative penalty

$$P(x; \lambda, \eta) = \frac{1}{2} \frac{\lambda^2}{1+\eta} 1_{x \neq 0} + \frac{1}{2} \eta x^2. \quad (7.3)$$

From (7.3), we see that hard-ridge thresholding successfully fuses the L_0 -penalty and the L_2 -penalty (ridge-penalty) via Θ -thresholding. The L_0 portion induces sparsity, while the L_2 portion shrinks β . The difference from hard thresholding is as follows. When hard thresholding makes $\hat{\gamma}_i \neq 0$ we get $\hat{\gamma}_i = r_i$, removing any influence of observation i on $\hat{\beta}$. With hard-ridge thresholding, the influence of the observation can be removed partially, with the extent of removal controlled by η . This is helpful for nonzero but small $|\gamma_i|$ corresponding to mild outlyingness. In addition, the L_2 shrinkage also plays a role in estimation and prediction.

The large p setting brings a computational challenge. Because the large- p sparse Θ -IPOD can be very slow, we used the following proportional Θ -IPOD to screen out some ‘nuisance dimensions’ with coefficients being exactly zero (due to our sparsity assumption). Concretely, at each update in (7.1), λ was chosen to get precisely αn nonzero components in the new $(\beta^{(j+1)}, \gamma^{(j+1)})$. We used $\alpha = 0.75$ though other

choices could be made. Proportional Θ -IPOD yields at most αn candidate predictors to have nonzero coefficients, making the problem simpler. Next, we run (7.1) with just those predictors, getting a solution path in λ and choosing λ by BIC^* . In proportional Θ -IPOD, a variable that gets killed at one iteration may reappear in the fit at a later iteration. This is quite different from independent screenings based on marginal statistics like FDR (Benjamini & Hochberg 1995) or SIS (Fan & Lv 2008).

For our example we used the proportional hard-ridge IPOD. Specifically, for η in a small grid of values, we ran proportional Θ -IPOD using the hard-ridge thresholding function Θ . We chose η by BIC^* .

We perform the joint robust variable selection and outlier detection on the sugar data of Brown et al. (1998). The data set concerns NIR spectroscopy of compositions of three sugars in aqueous solution. We concentrate on glucose (sugar 2) as the response variable. The predictors are second derivative spectra of 700 absorbances at frequencies corresponding to wavelengths of 1100nm to 2498nm in steps of 2nm. There are 125 samples for model training. A test set with 21 test samples is available. These 21 samples were specially designed to be difficult to predict, with compositions outside the range of the other 125. See Brown (1993) for details.

Brown et al. (1998) used an MCMC Bayesian approach for variable selection, but only 160 equally spaced wavelengths were analyzed due to the computational cost. We used all 700 wavelengths in fitting our sparse robust model that also estimates outliers. The problem size is $125 \times (700 + 125)$. After the screening via proportional hard-ridge-IPOD, we ran hard-ridge-IPOD and tuned it as follows. First we set $\lambda = 0$ which turns η into an ordinary ridge parameter. We fit that ridge parameter η , obtaining η^* by minimizing BIC^* of Section 6. Then for each η in the grid $\{0.5\eta^*, 0.05\eta^*, 0.005\eta^*\}$, we found $\lambda(\eta)$ to minimize BIC^* , and finally chose among the three $(\lambda(\eta), \eta)$ combinations to minimize BIC^* . We think there is no reason to use $\eta > \eta^*$ and a richer grid than

the one we used would be reasonable, but we chose a small grid for computational reasons. Our experience is that even small η values improve prediction.




Figure 7: Sparse coefficients and outlier shifts on the sugar data.

Figure 7 plots the final estimates $\hat{\beta}$ and $\hat{\gamma}$. From the left panel, we see that the estimated model is sparse. It selects only 15 of the wavelengths. The estimate $\hat{\gamma}$ shown in right side of Figure 7 suggests that observation 99 might be an outlier. We found $\hat{\gamma}_{99} = 1.7706$ and $r_{99} = y_{99} - \mathbf{x}_{99}^T \hat{\beta} = 1.7711$, which indicates that this observation is (almost) unused in the model fitting. Our model has good prediction performance; the mean-squared error is 0.219 on the test data, improving the reported MCMC results by about 39%. The robust residual plot is shown in Figure 8.

8 Discussion

The main contribution of this paper is to consider the class of Θ -estimators (She 2009), under the mean shift outlier model assumption, defined by the fixed point $\gamma = \Theta(\mathbf{H}\gamma + (\mathbf{I} - \mathbf{H})\mathbf{y}; \lambda)$ which can be directly computed by Θ -IPOD. With a good design of Θ , we successfully identified the outliers as well as estimating the coefficients robustly. This technique is associated with M -estimators, but gives a new characterization in form of penalized regressions. Furthermore, we successfully generalized this penalized/thresholding methodology to high-dimensional problems to accommodate and identify gross outliers in variable selection and coefficient estimation.




Figure 8: Residual plot from hard-ridge-IPOD on the sugar data. The apparent outlier, observation 99, is shown in red star.

When outliers are also leverage points, the Gram matrix of the reduced model (4.9), i.e., $\mathbf{A}^\top \mathbf{A} = \mathbf{I} - \mathbf{H}$, may demonstrate high correlation between clean observations and outliers for small h_i . It is well known from recent advances of the lasso (e.g. the irrepresentable conditions (Zhao & Yu 2006) and the sparse Riesz condition (Zhang & Huang 2008)) that the convex L_1 -penalty encounters great trouble in this situation. Rather, nonconvex penalties which correspond to redescending ψ 's must be applied, with a high breakdown point initial estimate obtained by, say, the fast-LTS (Rousseeuw & Van Driessen 2002) or fast-S (Salibián-Barrera & Yohai 2006) when p is small, or the fast PY procedure (Peña & Yohai 1999) when p is large.

In this framework, determining the efficiency parameter in M -estimators and choosing a cutoff value for outlier identification are both accomplished by tuning the choice of $\boldsymbol{\lambda}$. Our experience shows that adopting an appropriate data-dependent choice of $\boldsymbol{\lambda}$ is crucial to guarantee good detection performance.

We close with a note of caution that on an important issue for the outlier detection literature. We have found that our robust regression algorithms work better than competitors on our simulated data sets and that they give the right answers on some

small well studied real data sets. But our method and the others we study all rely on a preliminary robust fit. The most robust preliminary fits have a cost that grows exponentially with dimension and for them $p = 15$ is already large. In high dimensional problems, the preliminary fit of choice is seemingly the PY procedure. It performed well in our examples, but it does not necessarily have high breakdown. Should the preliminary method fail, followup methods like Θ -IPOD or the others, may or may not correct it. We have seen hard-IPOD work well from a non-robust start (e.g. the HBK problem) but it would not be reasonable to expect this will always hold. We expect that Θ -IPOD iterations will benefit from improvements in preliminary robust fitting methods for high dimensional problems. For low dimensional problems robust preliminary methods like LTS are fast enough.

A Proofs

• Proof of Proposition 3.1

Since both Huber's method and the L_1 -penalized regression are convex, it is sufficient to consider the KKT equations. In this proof, we define

$$\Theta(t; \lambda) = \begin{cases} t - \lambda, & \text{if } t > \lambda \\ 0, & \text{if } |t| < \lambda \\ -t + \lambda, & \text{if } t < -\lambda \end{cases}, \quad \psi(t; \lambda) = \begin{cases} \lambda, & \text{if } t > \lambda \\ t, & \text{if } |t| < \lambda \\ -\lambda, & \text{if } t < -\lambda \end{cases}.$$

Obviously, $\Theta(t; \lambda) + \psi(t; \lambda) = t$. For simplicity, we use the same symbols for the vector versions of the two functions: $\Theta(\mathbf{t}; \lambda) = [\Theta(t_i, \lambda)]_{i=1}^n$, $\psi(\mathbf{t}; \lambda) = [\psi(t_i, \lambda)]_{i=1}^n$, $\forall \mathbf{t} \in \mathbb{R}^n$.

The KKT equations for Huber's estimate $(\hat{\boldsymbol{\beta}}, \hat{\sigma})$ are given by

$$\mathbf{X}^\top \psi \left(\frac{\mathbf{y} - \mathbf{X}\boldsymbol{\beta}}{\sigma}; \lambda \right) = \mathbf{0} \quad (\text{A.1})$$

$$\frac{\partial}{\partial \sigma} \left(\sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{X}_i^\top \boldsymbol{\beta}}{\sigma}; \lambda \right) \sigma \right) + nc = 0 \quad (\text{A.2})$$

Let $\hat{\mathbf{r}} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}$, $\hat{G} = \{i : |\hat{r}_i| < \lambda\hat{\sigma}\}$, and $\hat{O} = \{i : |\hat{r}_i| > \lambda\hat{\sigma}\}$. Since

$$\frac{\partial}{\partial \sigma} (2\rho(t/\sigma; \lambda)\sigma) = \begin{cases} -t^2/\sigma^2, & \text{if } |t| < \lambda\sigma \\ -\lambda^2, & \text{if } |t| > \lambda\sigma \end{cases}$$

(A.2) becomes $nc = \lambda^2|\hat{O}| + \sum_{i \in \hat{G}} \hat{r}_i^2/\hat{\sigma}^2$. To summarize, $(\hat{\boldsymbol{\beta}}, \hat{\sigma})$ satisfies

$$0 = \mathbf{X}^\top \psi \left(\frac{\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}}{\hat{\sigma}}; \lambda \right) \quad \text{and} \quad (\text{A.3})$$

$$\hat{\sigma}^2 = \|\hat{\mathbf{r}}_{\hat{G}}\|_2^2 / (cn - \lambda^2|\hat{O}|). \quad (\text{A.4})$$

Next, the joint KKT equations for L_1 -penalized regression estimates $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}, \hat{\sigma})$ are

$$\mathbf{X}^\top \mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}^\top (\mathbf{y} - \hat{\boldsymbol{\gamma}})$$

$$\hat{\boldsymbol{\gamma}} = \Theta(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}; \lambda\hat{\sigma}) \quad \text{and}$$

$$0 = \|\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\gamma}}\|_2^2 \cdot \left(-\frac{1}{\hat{\sigma}^2}\right) + nc,$$

from which it follows that $\mathbf{X}^\top \psi((\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})/\hat{\sigma}; \lambda) = \mathbf{0}$, and $\hat{\sigma}^2 = \|\hat{\mathbf{r}} - \hat{\boldsymbol{\gamma}}\|_2^2 / nc$. But

$\|\hat{\mathbf{r}} - \hat{\boldsymbol{\gamma}}\|_2^2 = \sum_{i \in \hat{G}} \hat{r}_i^2 + \lambda^2 \hat{\sigma}^2 |\hat{O}|$, so we obtain

$$\mathbf{0} = \mathbf{X}^\top \psi \left(\frac{\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}}{\hat{\sigma}}; \lambda \right), \quad \text{and}$$

$$\hat{\sigma}^2 = \|\hat{\mathbf{r}}_{\hat{G}}\|_2^2 / (cn - \lambda^2|\hat{O}|),$$

which are exactly the same as (A.3) and (A.4). Therefore, (3.3) leads to Huber's method with joint scale estimation. When σ is fixed, it is not difficult to show that the two minimizations still yield the same β -estimate.

Finally, Huber notices his method behaves poorly even for moderate leverage points and considers an improvement of using $c_1\psi(\hat{r}_i/(c_2\sigma))$ to replace $\psi(\hat{r}_i/\sigma)$ (Huber 1981). He claims, based on heuristic arguments, that $c_1 = c_2 = \sqrt{1 - h_i}$ is a good choice. This is quite natural as seen from our new characterization. The regression matrix for γ in (3.3) is actually $\mathbf{I} - \mathbf{H}$ with the column norms given by $\sqrt{1 - h_i}$. The weighted L_1 -penalty of $\sum(\lambda\sqrt{1 - h_i}) \times |\gamma_i|$ then leads to Huber's improvement, due to the fact that $K\psi(t/K; \lambda) = \psi(t; K\lambda)$ for $K > 0$. (However it cannot completely avoid masking and swamping unless a redescending ψ is used.) \square

• **Proof of Proposition 4.1**

By definition, for any Θ -IPOD estimate $(\hat{\beta}, \hat{\gamma})$, $\hat{\gamma}$ is a fixed point of $\gamma = \Theta(\mathbf{H}\gamma + (\mathbf{I} - \mathbf{H})\mathbf{y}; \lambda)$, and $\hat{\beta} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top (\mathbf{y} - \hat{\gamma})$. It follows that

$$\begin{aligned} \mathbf{X}^\top \psi(\mathbf{y} - \mathbf{X}\hat{\beta}) &= \mathbf{X}^\top \psi(\mathbf{y} - \mathbf{H}(\mathbf{y} - \hat{\gamma}); \lambda) \\ &= \mathbf{X}^\top (((\mathbf{I} - \mathbf{H})\mathbf{y} + \mathbf{H}\hat{\gamma}) - \Theta((\mathbf{I} - \mathbf{H})\mathbf{y} + \mathbf{H}\hat{\gamma}); \lambda)) \\ &= \mathbf{X}^\top ((\mathbf{I} - \mathbf{H})\mathbf{y} + \mathbf{H}\hat{\gamma} - \hat{\gamma}) \\ &= \mathbf{X}^\top (\mathbf{I} - \mathbf{H}) \cdot (\mathbf{y} - \hat{\gamma}) = \mathbf{0}, \end{aligned}$$

and so $\hat{\beta}$ is an M -estimate associated with ψ . \square

• **Proof of Theorem 4.1**

The second inequality in (4.3) is straightforward from the algorithm design. To show the first inequality is true, it is sufficient to prove the following lemma.

Lemma A.1. *Given a thresholding rule Θ , let P be any function satisfying $P(\theta; \lambda) = P(0; \lambda) + P_\Theta(\theta; \lambda) + q(\theta; \lambda)$ where $q(\cdot; \lambda)$ is nonnegative and $q(\Theta(\theta; \lambda)) = 0$ for all θ . Then, the minimization problem $\min_\theta (t - \theta)^2/2 + P(\theta; \lambda)$ has a unique optimal solution*

$\hat{\theta} = \Theta(t; \lambda)$ for every t at which $\Theta(\cdot; \lambda)$ is continuous.

This is a generalization of Proposition 3.2 in Antoniadis (2007). Note that P (and P_{Θ}) may not be differentiable at 0 and may not be convex.

Without loss of generality, suppose $t > 0$. It suffices to consider $\theta \geq 0$ since $f(\theta) \geq f(-\theta), \forall \theta \geq 0$, where $f(\theta) \equiv (t - \theta)^2/2 + P(\theta; \lambda)$. First, Θ^{-1} given by (4.1) is well-defined. We have

$$\begin{aligned} f(\theta) - f(\Theta(t; \lambda)) &= \int_{\Theta(t; \lambda)}^{\theta} f'(u) du + q(\theta; \lambda) - q(\Theta(t; \lambda)) \\ &= \int_{\Theta(t; \lambda)}^{\theta} (u - t + \Theta^{-1}(u; \lambda) - u) du + q(\theta; \lambda) \\ &= \int_{\Theta(t; \lambda)}^{\theta} (\Theta^{-1}(u; \lambda) - t) du + q(\theta; \lambda). \end{aligned}$$

Suppose $\theta > \Theta(t; \lambda)$. By definition $\Theta^{-1}(\theta; \lambda) \geq t$, and thus $f(\theta) \geq f(\Theta(t; \lambda))$. There must exist some $u \in [\Theta(t; \lambda), \theta)$ s.t. $\Theta^{-1}(u; \lambda) > t$. Otherwise we would have $\Theta(t'; \lambda) > \theta \geq \Theta(t; \lambda)$ for any $t' > t$, and thus $\Theta(\cdot)$ would be discontinuous at t . Noticing that $\Theta^{-1}(\cdot)$ is monotone, $\int_{\Theta(t; \lambda)}^{\theta} (\Theta^{-1}(u; \lambda) - t) du > 0$, or $f(\theta) > f(\Theta(t; \lambda))$. A similar reasoning applies to the case $\theta < \Theta(t; \lambda)$. The proof is now complete. \square

References

- Antoniadis, A. (2007), “Wavelet methods in statistics: Some recent developments and their applications,” *Statistics Surveys*, 1, 16–55.
- Antoniadis, A., & Fan, J. (2001), “Regularization of Wavelets Approximations,” *JASA*, 96, 939–967.
- Atkinson, A. C., & Riani, M. (2000), *Robust diagnostic regression analysis*, New York: Springer-Verlag.
- Benjamini, Y., & Hochberg, Y. (1995), “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289–300.

- Brown, P. (1993), *Measurements, Regression and Calibration*, Oxford: Oxford University Press.
- Brown, P. J., Vannucci, M., & Fearn, T. (1998), “Multivariate Bayesian Variable Selection and Prediction,” *JRSSB*, 60(3), 627–641.
- Chen, J., & Chen, Z. (2008), “Extended Bayesian information criterion for model selection with large model space,” *Biometrika*, 95, 759–771.
- Coakley, C. W., & Hettmansperger, T. P. (1993), “A bounded influence, high breakdown, efficient regression estimator,” *J. Amer. Statist. Assoc.*, 88(423), 872–880.
- Donoho, D., & Johnstone, I. (1994), “Ideal Spatial Adaptation via Wavelet Shrinkages,” *Biometrika*, 81, 425–455.
- Fan, J., & Lv, J. (2008), “Sure independence screening for ultrahigh dimensional feature space,” *Journal Of The Royal Statistical Society Series B*, 70(5), 849–911.
- Gannaz, I. (2006), Robust estimation and wavelet thresholding in partial linear models,, Technical report, University Joseph Fourier, Grenoble, France.
- Gervini, D., & Yohai, V. J. (2002), “A class of robust and fully efficient regression estimators,” *Ann. Statist.*, 30(2), 583–616.
- Hadi, A. S., & Simonoff, J. S. (1993), “Procedures for the identification of multiple outliers in linear models,” *J. Amer. Statist. Assoc.*, 88(424), 1264–1272.
- Hadi, A. S., & Simonoff, J. S. (1997), “A More Robust Outlier Identifier for Regression Data,” *J. Amer. Statist. Assoc.*, pp. 281–282.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., & Stahel, W. A. (1986), *Robust Statistics: The Approach Based on Influence Functions* Wiley.
- Hawkins, D. M., Bradu, D., & Kass, G. V. (1984), “Location of several outliers in multiple-regression data using elemental sets,” *Technometrics*, 26(3), 197–208.
- Huber, P. J. (1981), *Robust Statistics*, New York: John Wiley & Sons.
- Maronna, R. A., Martin, D. R., & Yohai, V. J. (2006), *Robust Statistics: Theory and Methods*, Chichester: Wiley.

- McCann, L., & Welsch, R. E. (2007), “Robust variable selection using least angle regression and elemental set sampling,” *Computational Statistics & Data Analysis*, 52(1), 249–257.
- Menjoge, R. S., & Welsch, R. E. (2010), “A diagnostic method for simultaneous feature selection and outlier identification in linear regression,” *Computational statistics and data analysis*, 54, 3181–3193.
- Nguyen, T., & Welsch, R. (2010), “Outlier detection and least trimmed squares approximation using semi-definite programming,” *Computational Statistics and Data Analysis*, 54(12), 3212 – 3226.
- Peña, D., & Yohai, V. (1999), “A fast procedure for outlier diagnostics in large regression problems,” *J. Amer. Statist. Assoc.*, 94(446), 434–445.
- Peña, D., & Yohai, V. J. (1995), “The detection of influential subsets in linear regression by using an influence matrix,” *J. Roy. Statist. Soc. Ser. B*, 57(1), 145–156.
- R Development Core Team (2007), “R: A Language and Environment for Statistical Computing,” <http://www.R-project.org>. ISBN 3-900051-07-0.
- Rousseeuw, P. J. (1984), “Least median of squares regression,” *J. Amer. Statist. Assoc.*, 79(388), 871–880.
- Rousseeuw, P. J., & Leroy, A. M. (1987), *Robust regression and outlier detection*, Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics, New York: John Wiley & Sons Inc.
- Rousseeuw, P., & Van Driessen, K. (2002), “Computing LTS Regression for Large Data Sets,” *Estadística*, 54, 163–190.
- Rousseeuw, P., & 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.
- Salibian-Barrera, M., & Yohai, V. J. (2006), “A fast algorithm for S-regression estimates,” *J. Computat. Graphic. Statist*, 15, 414–427.
- Serbert, D., Montgomery, D., & Rollier, D. (1998), “A clustering algorithm for identifying multiple outliers in linear regression,” *Computational Statistics & Data Analysis*, 27, 461–484.

- She, Y. (2009), “Thresholding-based Iterative Selection Procedures for Model Selection and Shrinkage,” *Electronic Journal of Statistics*, 3, 384–415.
- Simpson, D. G., Ruppert, D., & Carroll, R. J. (1992), “On one-step GM estimates and stability of inferences in linear regression,” *J. Amer. Statist. Assoc.*, 87(418), 439–450.
- Theil, H. (1965), “The analysis of disturbances in regression analysis,” *Journal of the American Statistical Association*, 60(312), 1067–1079.
- Šidák, Z. (1967), “Rectangular Confidence Regions for the Means of Multivariate Normal Distribution,” *JASA*, 62, 626–633.
- Weisberg, S. (1985), *Applied Linear Regression*, 2nd edn, New York: Wiley.
- Wilcox, R. R. (2005), *Introduction to Robust Estimation and Hypothesis Testing*, 2nd edn, San Diego, CA, USA: Academic Press.
- Wisnowski, J. W., Montgomery, D. C., & Simpson, J. R. (2001), “A Comparative analysis of multiple outlier detection procedures in the linear regression model,” *Computational Statistics & Data Analysis*, 36(3), 351–382.
- Yohai, V. J. (1987), “High breakdown-point and high efficiency robust estimates for regression,” *Ann. Statist.*, 15(2), 642–656.
- Zhang, C.-H., & Huang, J. (2008), “The sparsity and bias of the Lasso selection in high-dimensional linear regression,” *Ann. Statist.*, 36, 1567–1594.
- Zhao, P., & Yu, B. (2006), “On Model Selection Consistency of Lasso,” *Journal of Machine Learning Research*, 7, 2541–2563.
- Zou, H., & Hastie, T. (2005), “Regularization and Variable Selection via the Elastic Net,” *JRSSB*, 67(2), 301–320.