

Rank-based Lasso - efficient methods for high-dimensional robust model selection

Wojciech Rejchel

*Faculty of Mathematics and Computer Science
Nicolaus Copernicus University
Chopina 12/18, 87-100 Toruń, Poland*

WREJCHEL@GMAIL.COM

Małgorzata Bogdan

*Faculty of Mathematics and Computer Science
University of Wrocław
Joliot-Curie 15, 50-383 Wrocław, Poland
and
Department of Statistics
Lund University
Tycho Brahes väg 1, Lund, Sweden*

MALGORZATA.BOGDAN@UWR.EDU.PL

Editor: Garvesh Raskutti

Abstract

We consider the problem of identifying significant predictors in large data bases, where the response variable depends on the linear combination of explanatory variables through an unknown monotonic link function, corrupted with the noise from the unknown distribution. We utilize the natural, robust and efficient approach, which relies on replacing values of the response variables by their ranks and then identifying significant predictors by using well known Lasso. We provide new consistency results for the proposed procedure (called „RankLasso”) and extend the scope of its applications by proposing its thresholded and adaptive versions. Our theoretical results show that these modifications can identify the set of relevant predictors under a wide range of data generating scenarios. Theoretical results are supported by the simulation study and the real data analysis, which show that our methods can properly identify relevant predictors, even when the error terms come from the Cauchy distribution and the link function is nonlinear. They also demonstrate the superiority of the modified versions of RankLasso over its regular version in the case when predictors are substantially correlated. The numerical study shows also that RankLasso performs substantially better in model selection than LADLasso, which is a well established methodology for robust model selection.

Keywords: Lasso, Model Selection, Ranks, Single Index Model, Sparsity, U -statistics

1. Introduction

Model selection is a fundamental challenge when working with large-scale data sets, where the number of predictors exceeds significantly the number of observations. In many practical problems finding a small set of significant predictors is at least as important as accurate estimation or prediction. Among many approaches to high-dimensional model selection one

can distinguish a large group of methods based on penalized estimation (Hastie et al., 2001; Bühlmann and van de Geer, 2011). Under the linear regression model

$$Y_i = \beta' X_i + \varepsilon_i, \quad i = 1, \dots, n,$$

where $Y_i \in \mathbb{R}$ is a response variable, $X_i \in \mathbb{R}^p$ is a vector of predictors, $\beta \in \mathbb{R}^p$ is the vector of model parameters and ε_i is a random error, the penalized model selection approaches usually recommend estimating the vector of regression coefficients β by

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^n (Y_i - \beta' X_i)^2 + Pen(\beta), \quad (1)$$

where $\sum_{i=1}^n (Y_i - \beta' X_i)^2$ is the quadratic loss function measuring the model fit and $Pen(\beta)$ is the penalty on the model complexity. The main representative of these methods is Lasso (Tibshirani, 1996), which uses the l_1 -norm penalty. The properties of Lasso in model selection, estimation and prediction are deeply investigated, e.g. in Meinshausen and Bühlmann (2006); Zhao and Yu (2006); Zou (2006); van de Geer (2008); Bickel et al. (2009); Ye and Zhang (2010); Bühlmann and van de Geer (2011); Huang and Zhang (2012); Su et al. (2017); Tardivel and Bogdan (2018). These articles discuss the properties of Lasso in the context of linear or generalized linear models and their results hold under specific assumptions on the relationship between the response and explanatory variables and/or the distribution of the random errors. However, it is quite common that a complex data set does not satisfy these assumptions or they are difficult to verify. In such cases it is advised to use „robust” methods of model selection.

In this paper we consider the single index model

$$Y_i = g(\beta' X_i, \varepsilon_i), \quad i = 1, \dots, n, \quad (2)$$

where g is unknown monotonic link function. Thus, we suppose that predictors influence the response variable through the link function g of the scalar product $\beta' X_i$. However, we make no assumptions on the form of the link function g (except being monotonic) nor on the distribution of the error term ε_i . Specifically, we do not assume the existence of the expected value of ε_i .

The goal of model selection is the identification of the set of relevant predictors

$$T = \{1 \leq j \leq p : \beta_j \neq 0\}. \quad (3)$$

The literature on the topic of robust model selection is quite considerable and the comprehensive review can be found e.g. in Wu and Ma (2015). Many of the existing methods suppose that the linear model assumption is satisfied and consider the robustness with respect to the noise. Here the most popular approaches rely on replacing the regular quadratic loss function with the loss function, which is more robust with respect to outliers, like e.g. the absolute value or Huber loss functions (Huber, 1964). Model selection properties of the penalized regression procedures with such robust loss functions were investigated, among others, in Wang et al. (2007); Gao and Huang (2010); Belloni and Chernozhukov (2011); Wang et al. (2012); Wang (2013); Fan et al. (2014); Peng and Wang (2015); Zhong et al. (2016); Avella-Medina and Ronchetti (2018). Among these methods one can mention the

approach of Johnson and Peng (2008); Johnson (2009), where the loss function is expressed in terms of residual ranks. On the other hand, the issues of model selection in misspecified models were discussed in e.g. Lu et al. (2012); Lv and Liu (2014), while robustness with respect to the unknown link function g in the single index model (2) was discussed, for instance, in Kong and Xia (2007); Zeng et al. (2012); Alquier and Biau (2013); Plan and Vershynin (2016); Cheng et al. (2017). In particular, Plan and Vershynin (2016) proposed a procedure, which can estimate the parameter β with accuracy to the multiplicative constant, when predictors X are Gaussian and $\mathbb{E}Y^2 < \infty$. Their approach was extended to incorporate different loss functions (Genzel, 2017), non-Gaussian predictors (Yang et al., 2017a,b; Wei et al., 2019) and to high-dimensional varying index coefficient models (Na et al., 2019). The extensions in (Yang et al., 2017a,b; Wei et al., 2019; Na et al., 2019) are based on Stein’s lemma (Stein, 1972; Stein et al., 2004) and the proposed solutions depend on the distribution of predictors, which is assumed to be known. Additionally, all these works require some moment assumptions on Y , which precludes application of this methodology, when the errors have a heavy-tailed distribution, like e.g., Cauchy or some cases of log-normal distribution.

Penalized robust model selection procedures for single-index models with heavy tailed errors were developed e.g. in Zhu and Zhu (2009); Song and Ma (2010); Wang and Zhu (2015); Zhong et al. (2016); Rejchel (2017b,a), where their desired statistical properties are confirmed. However, the application of procedures based on robust loss functions (e.g. piecewise-linear) in the context of the analysis of large data sets is often limited due to their computational complexity and/or the need of the development of dedicated optimization algorithms. For instance, in Sections 3 and 4 we consider the Least Absolute Deviation Lasso (LADLasso) estimator (Wang et al., 2007; Belloni and Chernozhukov, 2011; Fan et al., 2014), which turns out to be computationally very slow even for moderate dimension experiments.

In the current paper we consider an alternative approach for identifying important predictors in the single index model. Our method does not require knowledge of the distribution of predictors or any moment assumptions on the error distribution. Moreover, it is computationally fast and can work efficiently with complex high-dimensional data sets. Our procedure is very simple and relies on replacing actual values of the response variables Y_i by their centred ranks. Ranks R_i are defined as

$$R_i = \sum_{j=1}^n \mathbb{I}(Y_j \leq Y_i), \quad i = 1, \dots, n, \quad (4)$$

where $\mathbb{I}(\cdot)$ is the indicator function. Next, we identify significant predictors by simply solving the following Lasso problem;

$$\mathbf{RankLasso:} \quad \hat{\theta} = \arg \min_{\theta \in \mathbb{R}^p} Q(\theta) + \lambda |\theta|_1, \quad (5)$$

where

$$Q(\theta) = \frac{1}{2n} \sum_{i=1}^n (R_i/n - 0.5 - \theta' X_i)^2. \quad (6)$$

This procedure does not require any dedicated algorithm and can be executed using efficient implementations of Lasso in „R” (R Development Core Team, 2017) packages: „lars” (Efron

et al., 2004) or „glmnet” (Friedman et al., 2010). Technically, RankLasso was introduced before in Zhu and Zhu (2009); Wang and Zhu (2015), who used a slightly more complicated definition, which makes it difficult to notice the relationship with ranks.

Replacing values of response variables by their ranks is a well-known approach in non-parametric statistics and leads to robust procedures. The premier examples of a rank approach are the Wilcoxon test, that is a widely used alternative to the Student’s t-test, or the Kruskal-Wallis ANOVA test. While the rank tests often have a low power for a small number of observations, they can achieve high efficiency for large sample sizes. As shown in Žak et al. (2007); Bogdan et al. (2008), this carries over to the high efficiency of identifying important predictors in the sparse high-dimensional regression models, where the number of true nonzero regression coefficients is much smaller than the sample size n .

The methodology proposed in Žak et al. (2007); Bogdan et al. (2008) relies on minimization of the rank version of the Bayesian Information Criterion, which in principle is N-P hard. While the heuristics based on the greedy search algorithms can in some cases identify approximately optimal models, they are not reliable in the case when predictors are highly correlated. Instead, RankLasso is based on a convex optimization algorithm, which can be easily solved even when $p \gg n$ and the explanatory variables are highly correlated.

One of the disadvantages of the rank approach is the loss of information about the shape of the link function. Therefore, RankLasso cannot be directly used to build the predictive model for the response variable. In the case when errors are subgaussian, such predictive models can be constructed using e.g. the estimation method for the single-index model proposed in Balabdaoui et al. (2019), which can be also extended to the general error distribution by replacing the L_2 -loss with the robust loss functions (say, the L_1 or the Huber loss functions). However, this method can handle only a small number of predictors. In this article we demonstrate that significant predictors can be appropriately fished out from the large data base by simple modifications of RankLasso.

Specifically, in Subsection 2.2 we provide the definition of the parameter θ^0 , that is estimated by RankLasso, and discuss its relationship to the true vector of regression coefficients β . It turns out that under certain standard assumptions, the support of θ^0 coincides with the support of β and the methods based on ranks can identify the set of relevant predictors. However, similarly as in the case of regular Lasso, RankLasso can identify the true model only under very restrictive ”irrepresentable conditions” on the correlations between predictors and the sparsity of the vector of regression coefficients, see e.g. Zhao and Yu (2006); van de Geer and Bühlmann (2009); Wang and Zhu (2015). An intuitive explanation of these problems with model selection relates to the role of a tuning parameter λ . Namely, to obtain good model selection properties the parameter λ needs to be sufficiently large to discard irrelevant predictors. However, large λ leads to a large bias of Lasso estimators. In the result a non-explained effect of relevant predictors is intercepted by even slightly correlated variables, which leads to early false discoveries along the Lasso path and substantial difficulties with identification of the true model, see e.g. Su et al. (2017). This problem can be solved by using smaller value of the tuning parameter. An illustration of this phenomenon can be found in Weinstein et al. (2020), where it is shown that the ordering of estimated regression coefficients does not remain constant along the Lasso path. It turns out that the small mean squared error of Lasso estimates is typically obtained for a relatively small value of λ , where the coefficients corresponding to true discoveries

become larger than the ones for false discoveries, which appeared earlier on the Lasso path (Weinstein et al., 2020). Therefore, while regular Lasso usually cannot identify the true model for any selection of the tuning parameter λ , there often exists a range of λ values, which provide a good separation between the estimated regression coefficients of true and false predictors, see e.g. Ye and Zhang (2010); Bühlmann and van de Geer (2011); Tardivel and Bogdan (2018); Weinstein et al. (2020). In Weinstein et al. (2020) a detailed power comparison of two model selection procedures:

- **Lasso:** select $j : \hat{\beta}_j^{Lasso}(\lambda) \neq 0$,
- **Thresholded Lasso:** select $j : |\hat{\beta}_j^{Lasso}(\lambda)| > t$ for some threshold $t > 0$

is performed in the situation when the covariates are independent gaussian variables. This analysis shows that appropriately thresholded Lasso with the tuning parameter selected by cross-validation yields much higher power than regular Lasso for the same expected number of false discoveries and can identify the true model under much weaker regularity assumptions. Also, cross-validated Lasso estimators are often good candidates for the first step estimates for adaptive Lasso (Zou, 2006).

In this paper we use the above ideas and extend the scope of applications of RankLasso by proposing its thresholded and adaptive versions. In the case of standard Lasso similar modifications were introduced and discussed e.g. in Zou (2006); Candès et al. (2008); Zhou (2009); Tardivel and Bogdan (2018); Weinstein et al. (2020). We prove that the proposed modifications of RankLasso are model selection consistent in the model (2) under much weaker conditions than the ones provided in Wang and Zhu (2015) for regular RankLasso. More specifically, our results show that the modifications of RankLasso can identify the true model for any unknown monotonic link function and any unknown distribution of the error term and under much weaker restrictions on the design matrix and the signal sparsity than in the case of regular RankLasso. These theoretical results require a substantial modification of the proof techniques as compared to the similar results for regular Lasso. It is related to the fact that ranks are dependent, so (6) is a sum of dependent random variables. In Subsection 2.4 we describe how this problem can be overcome with the application of the theory of U -statistics. We also present extensive numerical results illustrating that the modifications of RankLasso can indeed properly identify relevant predictors, when the link function is not linear, error terms come from, say, the Cauchy distribution and predictors are substantially correlated. Specifically, it can be observed that, contrary to regular RankLasso, the proposed modifications can control the number of false discoveries and achieve a high power under strongly correlated designs. These results also show that RankLasso compares favorably with LADLasso, which is a well established methodology for robust model selection (Wang et al., 2007; Belloni and Chernozhukov, 2011; Fan et al., 2014).

The paper is organized as follows: in Section 2 we present theoretical results on the model selection consistency of RankLasso and its modifications. In Subsection 2.2 we discuss the relationship between β and the parameter estimated by RankLasso. We show that our approach is able to identify the support of β in the single index model. In Subsection 2.3 we consider properties of estimators in the high-dimensional scenario, where the number of predictors can be much larger than the sample size. We establish nonasymptotic bounds on the estimation error and separability of RankLasso. We use these results to prove model

selection consistency of thresholded and weighted RankLasso. In Subsection 2.4 we briefly draw a road map to the proofs of main results. Sections 3 and 4 are devoted to experiments that illustrate the properties of rank-based estimators on simulated and real data sets, respectively. The paper is concluded in Section 5. The proofs of main and auxiliary results are relegated to the appendix. We also place in the appendix results for the low-dimensional case, where the number of predictors is fixed and the sample size diverges to infinity. In this case we provide the necessary and sufficient conditions for model selection consistency of RankLasso and much weaker sufficient conditions for thresholded and weighted versions of RankLasso.

2. Model selection properties of RankLasso and its modifications

In this section we provide theoretical results concerning model selection as well as estimation properties of RankLasso and its thresholded and weighted versions. We start with specifying the assumptions on our model.

2.1 Assumptions and notation

Consider the single index model (2). In this paper we assume that the design matrix X and the vector of the error terms ε satisfy the following assumptions.

Assumption 1 *We assume that $(X_1, \varepsilon_1), \dots, (X_n, \varepsilon_n)$ are i.i.d. random vectors such that the distribution of X_1 is absolutely continuous and X_1 is independent of the noise variable ε_1 . Additionally, we assume that $\mathbb{E}X_1 = 0$, $H = \mathbb{E}X_1X_1'$ is positive definite and $H_{jj} = 1$ for $j = 1, \dots, p$.*

The single index model (2) does not allow to estimate an intercept and can identify β only up to a multiplicative constant, because any shift or scale change in $\beta'X_i$ can be absorbed by g . However, in many situations RankLasso can properly identify the support T of β . In this paper we will prove this fact under the following assumption.

Assumption 2 *We assume that for each $\theta \in \mathbb{R}^p$ the conditional expectation $\mathbb{E}(\theta'X_1|\beta'X_1)$ exists and*

$$\mathbb{E}(\theta'X_1|\beta'X_1) = d_\theta\beta'X_1$$

for a real number $d_\theta \in \mathbb{R}$.

Assumption 2 is a standard condition in the literature on the single index model or on the model misspecification, see e.g. Brillinger (1983); Ruud (1983); Li and Duan (1989); Zhu and Zhu (2009); Wang and Zhu (2015); Zhong et al. (2016); Kubkowski and Mielniczuk (2017). It is always satisfied in the simple regression models (i.e. when $X_1 \in \mathbb{R}$), which are often used for initial screening of explanatory variables, see e.g. Fan and Lv (2008). It is also satisfied whenever X_1 comes from the *elliptical distribution*, like the multivariate normal distribution or multivariate t -distribution. The interesting paper Hall and Li (1993) advocates that Assumption 2 is a nonrestrictive condition when the number of predictors is large, which is the case that we focus on in the paper. In the experimental section of this article we show that RankLasso, proposed here, is able to identify the support of β also

when the columns of the design matrix contain genotypes of independent Single Nucleotide Polymorphisms, whose distribution is not symmetric and clearly does not belong to the elliptical distribution.

The identifiability of the support of β by the rank procedure requires also the assumptions on the monotonicity of the link function g and the cumulative distribution function of Y_1 . The following Assumption 3, which combines Assumptions 1 and 2 and the monotonicity assumptions, will be used in most of theoretical results in this article.

Assumption 3 *We assume that the design matrix and the error term satisfy Assumptions 1 and 2, the cumulative distribution function F of the response variable Y_1 is increasing and g in (2) is increasing with respect to the first argument.*

In this paper we will use the following notation:

- $X = (X_1, X_2, \dots, X_n)'$ is the $(n \times p)$ -matrix of predictors,
- $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$,
- $Z_i = (X_i, Y_i), i = 1, \dots, n$,
- $T' = \{1, \dots, p\} \setminus T$ is a complement of T ,
- X_T is a submatrix of X , with columns whose indices belong to the support T of β , see (3),
- θ_T is a restriction of a vector $\theta \in \mathbb{R}^p$ to the indices from T ,
- p_0 is the number of elements in T ,
- the l_q -norm of a vector is defined as $|\theta|_q = \left(\sum_{j=1}^p |\theta_j|^q \right)^{1/q}$ for $q \in [1, \infty]$.

2.2 Identifying the support of β

RankLasso does not estimate β , but the vector

$$\theta^0 = \arg \min_{\theta \in \mathbb{R}^p} \mathbb{E} Q(\theta), \quad (7)$$

where $Q(\theta)$ is defined in (6). Since H is positive definite, the minimizer θ^0 is unique and is given by the formula

$$\theta^0 = \frac{1}{n^2} H^{-1} \left(\mathbb{E} \sum_{i=1}^n R_i X_i \right). \quad (8)$$

Now, using the facts that

$$\sum_{i=1}^n R_i X_i = \sum_{i=1}^n \sum_{j=1}^n \mathbb{I}(Y_j \leq Y_i) X_i = \sum_{i \neq j} \mathbb{I}(Y_j \leq Y_i) X_i + \sum_{i=1}^n X_i \quad (9)$$

and that $EX_i = 0$, we can write

$$\theta^0 = \frac{n-1}{n} H^{-1} \mu, \quad (10)$$

where $\mu = \mathbb{E} [\mathbb{I}(Y_2 \leq Y_1) X_1]$ is the expected value of the U -statistic

$$A = \frac{1}{n(n-1)} \sum_{i \neq j} \mathbb{I}(Y_j \leq Y_i) X_i. \quad (11)$$

In the next theorem we state the relation between θ^0 and β .

Theorem 1 Consider the model (2). If Assumptions 1 and 2 are satisfied, then

$$\theta^0 = \gamma_\beta \beta$$

with

$$\gamma_\beta = \frac{\frac{n-1}{n} \beta' \mu}{\beta' H \beta} = \frac{\frac{n-1}{n} \text{Cov}(F(Y_1), \beta' X_1)}{\beta' H \beta}, \quad (12)$$

where F is a cumulative distribution function of a response variable Y_1 .

Additionally, if F is increasing and g is increasing with respect to the first argument, then $\gamma_\beta > 0$, so the signs of β coincide with the signs of θ^0 and

$$T = \{j : \beta_j \neq 0\} = \{j : \theta_j^0 \neq 0\}. \quad (13)$$

We can apply Theorem 1 to the additive model $Y_i = g_1(\beta' X_i) + \varepsilon_i$ with an increasing function g_1 . Then, under Assumptions 1 and 2, $\theta^0 = \gamma_\beta \beta$ with $\gamma_\beta > 0$ if, for example, the support of the noise variable is a real line. Moreover, since the procedure based on ranks is invariant with respect to increasing transformations of response variables the same applies to the model $Y_i = g_2(\beta' X_i + \varepsilon_i)$ with an increasing function g_2 .

2.3 High-dimensional scenario

In this subsection we consider properties of the RankLasso estimator and its modifications in the case where the number of predictors can be much larger than the sample size. To obtain the results of this subsection we need the additional condition:

Assumption 4 We suppose that the vector of significant predictors $(X_1)_T$ is subgaussian with the coefficient $\tau_0 > 0$, i.e. for each $u \in \mathbb{R}^{p_0}$ we have $\mathbb{E} \exp(u'(X_1)_T) \leq \exp(\tau_0^2 u'u/2)$. Moreover, the irrelevant predictors are univariate subgaussian, i.e. for each $a \in \mathbb{R}$ and $j \notin T$ we have $\mathbb{E} \exp(aX_{1j}) \leq \exp(\tau_j^2 a^2/2)$ for positive numbers τ_j . Finally, we denote $\tau = \max(\tau_0, \tau_j, j \notin T)$.

We need subgaussianity of the vector of predictors to obtain exponential inequalities in the proofs of the main results in this subsection. This condition is a standard assumption while working with random predictors (Raskutti et al., 2010; Huang et al., 2013; Bühlmann and van de Geer, 2015) in high-dimensional models.

2.3.1 ESTIMATION ERROR AND SEPARABILITY OF RANKLASSO

Model selection consistency of RankLasso in the high-dimensional case was proved in Wang and Zhu (2015, Theorem 2.1). However, this result requires the stringent irrepresentable condition. Moreover, it is obtained under the polynomial upper bound on the dependency of p on n and provides only a rough guidance of selection of the tuning parameter λ . In our article we concentrate on estimation consistency of RankLasso, which paves the way for model selection consistency of the weighted and thresholded versions of this method. Compared to the asymptotic results of Wang and Zhu (2015) our results are stated in the form of non-asymptotic inequalities, they do not require the irrepresentable condition, allow for the exponential increase of p as a function of n and provide a precise guidance on selection of regularization parameter λ .

We start with introducing the cone invertibility factor (CIF), that plays an important role in investigating properties of estimators based on the Lasso penalty (Ye and Zhang, 2010). In the case $n > p$ one usually uses the minimal eigenvalue of the matrix $X'X/n$ to express the strength of correlations between predictors. Obviously, in the high-dimensional scenario this value is equal to zero and the minimal eigenvalue needs to be replaced by some other measure of predictors interdependency, which would describe the potential of consistent estimation of model parameters.

Let T be the set of indices corresponding to the support of the true vector β and let θ_T and $\theta_{T'}$ be the restrictions of the vector $\theta \in \mathbb{R}^p$ to the indices from T and T' , respectively. Now, for $\xi > 1$ we consider a cone

$$\mathcal{C}(\xi) = \{\theta \in \mathbb{R}^p : |\theta_{T'}|_1 \leq \xi |\theta_T|_1\} .$$

In the case when $p \gg n$ three different characteristics measuring the potential for consistent estimation of the model parameters have been introduced:

- the restricted eigenvalue (Bickel et al., 2009):

$$RE(\xi) = \inf_{0 \neq \theta \in \mathcal{C}(\xi)} \frac{\theta' X' X \theta / n}{|\theta|_2^2} ,$$

- the compatibility factor (van de Geer, 2008):

$$K(\xi) = \inf_{0 \neq \theta \in \mathcal{C}(\xi)} \frac{p_0 \theta' X' X \theta / n}{|\theta_T|_1^2} ,$$

- the cone invertibility factor (CIF, (Ye and Zhang, 2010)): for $q \geq 1$

$$\bar{F}_q(\xi) = \inf_{0 \neq \theta \in \mathcal{C}(\xi)} \frac{p_0^{1/q} |X' X \theta / n|_\infty}{|\theta|_q} .$$

Relations between the above quantities are discussed, for instance, in van de Geer and Bühlmann (2009); Ye and Zhang (2010); Huang et al. (2013). The fact that these conditions are much weaker than irrepresentable conditions is also established there.

In this article we will use CIF, since this factor allows for a sharp formulation of convergency results for all l_q norms with $q \geq 1$, see Ye and Zhang (2010, Section 3.2). The population (non-random) version of CIF is given by

$$F_q(\xi) = \inf_{0 \neq \theta \in \mathcal{C}(\xi)} \frac{p_0^{1/q} |H\theta|_\infty}{|\theta|_q} ,$$

where $H = EX_1 X_1'$. The key property of the random and the population versions of CIF, $\bar{F}_q(\xi)$ and $F_q(\xi)$, is that, in contrast to the smallest eigenvalues of matrices $X'X/n$ and H , they can be close to each other in the high-dimensional setting, see Huang et al. (2013, Lemma 4.1) or van de Geer and Bühlmann (2009, Corollary 10.1). This fact is used in the proof of Theorem 2 (given below).

In the simulation study in Section 3 we consider predictors, which are independent or equi-correlated, i.e. $H_{jj} = 1$ and $H_{jk} = b$ for $j \neq k$ and $b \in [0, 1)$. In this case the smallest eigenvalue of H is $1 - b$. For $\xi > 1$ and $q \geq 2$ we can bound CIF from below by

$$F_q(\xi) \geq (1 + \xi)^{-1} p_0^{1/q-1/2} (1 - b) ,$$

which illustrates that CIF diminishes with the increase of p_0 and b . Also, in the case when $H = I$ and $q = \infty$, $F_q(\xi) = 1$, independently of ξ or p_0 .

The next result describes the estimation accuracy of RankLasso.

Theorem 2 *Let $a \in (0, 1)$, $q \geq 1$ and $\xi > 1$ be arbitrary. Suppose that Assumptions 3 and 4 are satisfied. Moreover, suppose that*

$$n \geq \frac{K_1 p_0^2 \tau^4 (1 + \xi)^2 \log(p/a)}{F_q^2(\xi)} \quad (14)$$

and

$$\lambda \geq K_2 \frac{\xi + 1}{\xi - 1} \tau^2 \sqrt{\frac{\log(p/a)}{\kappa n}}, \quad (15)$$

where K_1, K_2 are universal constants and κ is the smallest eigenvalue of the correlation matrix between the true predictors $H_T = (H_{jk})_{j,k \in T}$. Then there exists a universal constant $K_3 > 0$ such that with probability at least $1 - K_3 a$ we have

$$|\hat{\theta} - \theta^0|_q \leq \frac{4\xi p_0^{1/q} \lambda}{(\xi + 1) F_q(\xi)}. \quad (16)$$

Besides, if X_1 has a normal distribution $N(0, H)$, then κ and τ can be dropped in (14) and (15).

In Theorem 2 we provide the bound for the estimation error of RankLasso. This result also provides the conditions for the estimation consistency of RankLasso, which can be obtained by replacing a with a sequence a_n , that decreases not too fast and selecting the minimal $\lambda = \lambda_n$ satisfying the condition (15). The consistency holds in the high-dimensional scenario, i.e. the number of predictors can be significantly greater than the sample size. Indeed, the consistency in the l_∞ -norm holds e.g. when $p = \exp(n^{\alpha_1})$, $p_0 = n^{\alpha_2}$, $a = \exp(-n^{\alpha_1})$, where $\alpha_1 + 2\alpha_2 < 1$, and λ is equal to the right-hand side of the inequality (15), provided that $F_\infty(\xi)$ and κ are bounded from below (or slowly converging to 0) and τ is bounded from above (or slowly diverging to ∞).

Theorem 2 is an analog of Theorem 3 in Ye and Zhang (2010), which refers to the linear model with noise variables having a finite variance. Similar results for quantile regression, among others for LADLasso, can be found in Belloni and Chernozhukov (2011, Theorem 2) and Wang (2013, Theorem 1). These results do not require existence of the noise variance, but impose some restrictions on the conditional distribution of Y_1 given X_1 , cf. Belloni and Chernozhukov (2011, Assumption D.1) or Wang (2013, the condition (13)).

In Alquier and Biau (2013); Plan and Vershynin (2016); Yang et al. (2017a,b); Wei et al. (2019) the results similar to Theorem 2 were obtained also in the context of the single index model. However, all these articles impose some moment restrictions on Y and thus cannot handle heavy-tailed noise distributions. Besides, in Alquier and Biau (2013) the link function are predictors are assumed to be bounded, while the methods analyzed in Plan and Vershynin (2016); Yang et al. (2017a,b); Wei et al. (2019) require the knowledge of the predictors' distribution. In the case of RankLasso we were able to prove Theorem 2 under a non-restrictive Assumption 2 on the distribution of predictors and without assumptions on existence of moments of Y or the noise variable ε .

The following corollary is an easy consequence of Theorem 2. It states that under assumptions of Theorem 2 RankLasso can asymptotically separate relevant and irrelevant predictors.

Corollary 3 *Suppose that conditions of Theorem 2 are satisfied for $q = \infty$. Let $\theta_{min}^0 = \min_{j \in T} |\theta_j^0|$. If $\theta_{min}^0 \geq \frac{8\xi\lambda}{(\xi+1)F_\infty(\xi)}$, then*

$$P\left(\forall_{j \in T, k \notin T} |\hat{\theta}_j| > |\hat{\theta}_k|\right) \geq 1 - K_3 a.$$

The separation of predictors by RankLasso given in Corollary 3 is a very important property. It will be used to prove model selection consistency of the thresholded and weighted RankLasso in the next part of the paper.

Finally, we discuss the condition of Corollary 3 that θ_{min}^0 cannot be too small, i.e. $\theta_{min}^0 \geq \frac{8\xi\lambda}{(\xi+1)F_\infty(\xi)}$. Using Theorem 1 we know that $\theta^0 = \gamma_\beta \beta$ and $\gamma_\beta > 0$, so this condition refers to the strength of the true parameter β and requires that

$$\min_{j \in T} |\beta_j| \geq \frac{8\xi\lambda}{\gamma_\beta(\xi+1)F_\infty(\xi)}. \quad (17)$$

Compared to the similar condition for regular Lasso in the linear model, the denominator contains an additional factor γ_β . This number is usually smaller than one, so RankLasso needs larger sample size to work well. This phenomenon is typical for the single-index model, where the similar restrictions hold for competitive methods like e.g. LADLasso.

Below we provide a simplified version of Theorem 2, formulated under the assumption that $F_\infty(\xi)$ and κ are lower bounded and τ is upper bounded. This formulation will be used in the following subsection to increase the transparency of the results on model selection consistency of weighted and thresholded RankLasso.

Corollary 4 *Let $a \in (0, 1)$ be arbitrary. Suppose that Assumptions 3 and 4 are satisfied. Moreover, assume that there exist $\xi_0 > 1$ and constants $C_1 > 0$ and $C_2 < \infty$ such that $\kappa \geq C_1$, $F_\infty(\xi_0) \geq C_1$ and $\tau \leq C_2$. If*

$$n \geq K_1 p_0^2 \log(p/a)$$

and

$$\lambda \geq K_2 \sqrt{\frac{\log(p/a)}{n}},$$

then

$$P\left(|\hat{\theta} - \theta^0|_\infty \leq 4\lambda/C_1\right) \geq 1 - K_3 a, \quad (18)$$

where the constants K_1 and K_2 depend only on ξ_0, C_1, C_2 and K_3 is a universal constant provided in Theorem 2.

2.3.2 MODIFICATIONS OF RANKLASSO

The main drawback of RankLasso considered in Subsection 2.3.1 is that it can recover the true model only if the restrictive irrepresentable condition is satisfied. If this condition does not hold, then RankLasso can achieve a high power only by including a large number of irrelevant predictors. In Theorems 5 and 7 we state that this problem can be overcome by the application of weighted or thresholded versions of RankLasso. In both cases we rely on the initial RankLasso estimator $\hat{\theta}$ of θ^0 , which is estimation consistent under the assumptions of Theorem 2 or Corollary 4. Theorems 5 and 7 are stated under simplified assumptions of Corollary 4. We have decided to establish them in these versions to make this subsection more communicable.

First, we consider thresholded RankLasso, which is denoted by $\hat{\theta}^{th}$ and defined as

$$\hat{\theta}_j^{th} = \hat{\theta}_j \mathbb{I}(|\hat{\theta}_j| \geq \delta), \quad j = 1, \dots, p, \quad (19)$$

where $\hat{\theta}$ is the RankLasso estimator given in (5) and $\delta > 0$ is a threshold. Theorem 5 provides the conditions under which this procedure is model selection consistent.

Theorem 5 *We assume that Corollary 4 holds and that the sample size and the tuning parameter λ for RankLasso are selected according to Corollary 4. Moreover, suppose that $\theta_{min}^0 = \min_{j \in T} |\theta_j^0|$ is such that it is possible to select the threshold δ so as*

$$\theta_{min}^0/2 \geq \delta > K_4 \lambda,$$

where $K_4 = 4/C_1$ is the constant from (18). Then it holds

$$P\left(\hat{T}^{th} = T\right) \geq 1 - K_3 a,$$

where K_3 is the universal constant from Theorem 2 and $\hat{T}^{th} = \{1 \leq j \leq p : \hat{\theta}_j^{th} \neq 0\}$ is the estimated set of relevant predictors by thresholded RankLasso.

Remark 6 *Theorem 5 illustrates that thresholded RankLasso has the potential for identifying the support of β under rather mild regularity conditions. This means that under these conditions the sequence of nested models based on the ranking provided by RankLasso estimates contains the true model. In the simulation study in Section 3 we select the threshold such that thresholded RankLasso returns the same number nonzero coefficients as weighted RankLasso, whose consistency is proved in Theorem 7 (below). However, there exists a variety of other methods, which could be used to select one of these nested models, including e.g. the rank version of the modified Bayesian Information Criterion proposed and discussed in Žak et al. (2007); Bogdan et al. (2008). Another plausible approach for identifying true predictors while controlling the false discovery rate at any given level could rely on the application of the knockoffs methodology of Barber and Candès (2015); Candès et al. (2018), based on the RankLasso Coefficient Difference statistics. Concerning the selection of the value of the tuning parameter λ , the results of Weinstein et al. (2020) for regular Lasso and our simulations for RankLasso, suggest that selection of the tuning parameter λ by cross-validation usually leads to the satisfactory performance of thresholded versions of*

Lasso. Here one could also mention the work of Pokarowski and Mielniczuk (2015) in the context of regular multiple regression, who suggest selecting λ such that the related sequence of nested regression models allows to obtain a minimal value of the respective information criterion.

Next, we consider the weighted RankLasso that minimizes

$$Q(\theta) + \lambda_a \sum_{j=1}^p w_j |\theta_j|, \quad (20)$$

where $\lambda_a > 0$ and weights are chosen in the following way: for arbitrary number $K > 0$ and the RankLasso estimator $\hat{\theta}$ from the previous subsection we have $w_j = |\hat{\theta}_j|^{-1}$ for $|\hat{\theta}_j| \leq \lambda_a$, and $w_j \leq K$, otherwise.

The next result describes properties of the weighted RankLasso estimator.

Theorem 7 *We assume that Corollary 4 holds and that the sample size and the tuning parameter λ for RankLasso are selected according to Corollary 4. Let $\lambda_a = K_4 \lambda$, where $K_4 = 4/C_1$ is from (18). Additionally, we suppose that the signal strength and sparsity satisfy $\theta_{min}^0/2 > \lambda_a$ and $p_0 \lambda \leq K_5$, where K_5 is sufficiently small constant. Then with probability at least $1 - K_6 a$ there exists a global minimizer $\hat{\theta}^a$ of (20) such that $\hat{\theta}_T^a = 0$ and*

$$|\hat{\theta}_T^a - \theta_T^0|_1 \leq K_7 p_0 \lambda, \quad (21)$$

where K_6 and K_7 are the constants depending only on K_1, \dots, K_5 and the constant K , that is used in the definition of weights.

In Fan et al. (2014, Corollary 1) the authors considered weighted Lasso with the absolute value loss function in the linear model. Thus, this procedure is robust with respect to the distribution of the noise variable. However, working with the absolute value loss function they need that, basically, the density of the noise is Lipschitz in a neighbourhood of zero (Fan et al., 2014, Condition 1). Our thresholded and weighted RankLasso does not require such restrictions. Besides, Theorems 5 and 7 confirm that the proposed procedures works well in model selection in the single index model (2).

It can be seen in the proof of Theorem 7 that K_7 is an increasing function of K , that occurs in the construction of weights. It is intuitively clear, because weights $w_j \leq K$ usually correspond (by Corollary 3) to significant predictors. Therefore, increasing K we shrink coordinates of the estimator, so the bias increases. This fact is described in (21).

2.4 Road map to proofs of main results

In the paper we study properties of the RankLasso estimator and its thresholded and weighted modifications. These estimators are obtained by minimization of the risk $Q(\theta)$ defined in (6) and the penalty. Therefore, the analysis of model selection properties of rank-based estimators is based on investigating these two terms. The penalty term can be handled using standard methods for regular Lasso. However, the analysis of the risk $Q(\theta)$ is different, because it is the sum of dependent variables. In this subsection we show that

the theory of U -statistics (Hoeffding, 1948; Serfling, 1980; de la Peña and Giné, 1999) plays a prominent role in studying properties of the risk $Q(\theta)$.

Consider the key object that is the derivative of the risk at the true point θ^0

$$\nabla Q(\theta^0) = -\frac{1}{n^2} \sum_{i=1}^n R_i X_i + \frac{1}{2n} \sum_{i=1}^n X_i + \frac{1}{n} \sum_{i=1}^n X_i X_i' \theta^0. \quad (22)$$

The second and third terms in (22) are sums of independent random variables, but the first one is a sum of dependent random variables. Using (9) we can express (22) as

$$-\frac{n-1}{n} A + \frac{n-2}{2n^2} \sum_{i=1}^n X_i + \frac{1}{n} \sum_{i=1}^n X_i X_i' \theta^0 \quad (23)$$

where a U -statistic A is defined in (11) and has the kernel

$$f(z_i, z_j) = \frac{1}{2} [\mathbb{I}(y_j \leq y_i) x_i + \mathbb{I}(y_i \leq y_j) x_j]. \quad (24)$$

To handle (23) we will use tools for sums of independent random variables as well as the U -statistics theory. Namely, we use exponential inequalities for sums of independent and unbounded random variables from van de Geer (2016, Corollary 8.2) and its version for U -statistics given in Lemma 14 in the appendix.

3. Simulation study

In this section we present results of the comparative simulation study verifying the properties of RankLasso and its thresholded and adaptive versions in model selection.

We consider the moderate dimension setup, where the number of explanatory variables p increases with n according to the formula $p = 0.01n^2$. More specifically, we consider the following pairs (n, p) : (100, 100), (200, 400), (300, 900), (400, 1600). For each of these combinations we consider three different values of the sparsity parameter $p_0 = \#\{j : \beta_j \neq 0\} \in \{3, 10, 20\}$.

In three of our simulation scenarios the rows of the design matrix are generated as independent random vectors from the multivariate normal distribution with the covariance matrix Σ defined as follows

- for the *independent* case $\Sigma = I$,
- for the *correlated* case $\Sigma_{ii} = 1$ and $\Sigma_{ij} = 0.3$ for $i \neq j$.

In one of the scenarios the design matrix is created by simulating the genotypes of p independent Single Nucleotide Polymorphisms (SNPs). In this case the explanatory variables can take only three values: 0 for the homozygote for the minor allele (genotype $\{a, a\}$), 1 for the heterozygote (genotype $\{a, A\}$) and 2 for the homozygote for the major allele (genotype $\{A, A\}$). The frequencies of the minor allele for each SNP are independently drawn from the uniform distribution on the interval (0.1, 0.5). Then, given the frequency π_j for j -th SNP, the explanatory variable X_{ij} has the distribution: $P(X_{ij} = 0) = \pi_j^2$, $P(X_{ij} = 1) = 2\pi_j(1 - \pi_j)$ and $P(X_{ij} = 2) = (1 - \pi_j)^2$.

The full description of the simulation scenarios is provided below:

- **Scenario 1**

$$Y = X\beta + \varepsilon,$$

where X matrix is generated according to the *independent* case, $\beta_1 = \dots = \beta_{p_0} = 3$ and the elements of $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)$ are independently drawn from the standard Cauchy distribution,

- **Scenario 2** - the regression model, values of regression coefficients and ε are as in Scenario 1, design matrix contains standardized versions of genotypes of p independent SNPs,
- **Scenario 3** - the regression model, values of regression coefficients and ε are as in Scenario 1 and the design matrix X is generated according to the *correlated* case,
- **Scenario 4** - the design matrix X is generated according to the *correlated* case and the relationship between Y_i and $\beta'X_i$ is non-linear:

$$Y_i = \exp(4 + 0.05\beta'X_i) + \varepsilon_i$$

and $\varepsilon_1, \dots, \varepsilon_n$ are independent random variables from the standard Cauchy distribution.

In our simulation study we compare five different statistical methods:

- **rL**: RankLasso defined in (5) with $\lambda := \lambda_{rL}$ and

$$\lambda_{rL} = 0.3 \sqrt{\frac{\log p}{n}}. \quad (25)$$

- **arL**: adaptive RankLasso (20), with $\lambda_a = 2\lambda_{rL}$ and weights

$$w_j = \begin{cases} \frac{0.1\lambda_{rL}}{|\hat{\theta}_j|} & \text{when } |\hat{\theta}_j| > 0.1\lambda_{rL}, \\ |\hat{\theta}_j|^{-1} & \text{otherwise,} \end{cases}$$

where $\hat{\theta}$ is the RankLasso estimator computed above. If $\hat{\theta}_j = 0$, then $|\hat{\theta}_j|^{-1} = \infty$ and j^{th} explanatory variable is removed from the list of predictors before running weighted RankLasso,

- **thrL**: thresholded RankLasso, where the tuning parameter for RankLasso is selected by cross-validation and the threshold is selected in such a way that the number of selected predictors coincides with the number of predictors selected by adaptive RankLasso,
- **LAD**: LADLasso, defined as

$$\arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n |Y_i - \theta'X_i| + \lambda_{LAD} \sum_{j=1}^p |\theta_j|$$

with $\lambda_{LAD} = 1.5 \sqrt{\frac{\log p}{n}}$,

- **cv**: regular Lasso with the tuning parameter selected by cross-validation.

The values of the tuning parameters for RankLasso and LADLasso were selected empirically so that both methods perform comparatively well for $p_0 = 3$ and $n = 200, p = 400$.

We compare the quality of the above methods by performing 200 replicates of the experiment, where in each replicate we generate the new realization of the design matrix X and the vector of random noise ε . We calculate the following statistical characteristics:

- **FDR**: the average value of $FDP = \frac{V}{\max(R, 1)}$, where R is the total number of selected predictors and V is the number of selected irrelevant predictors,
- **Power**: the average value of $TPP = \frac{S}{p_0}$, where $S = R - V$ is the number of properly identified relevant predictors,

- **NMP**: the average value of Numbers of Misclassified Predictors, i.e false positives or false negatives, which equals $V + p_0 - S$.

In Table 1 we compare the average times needed to invoke RankLasso using the *glmnet* package (Friedman et al., 2010) in the *R* software (R Development Core Team, 2017) and robust LADLasso using the *R* package *MTE* (Qin et al., 2017). It can be seen that LADLasso becomes prohibitively slow, when the number of columns of the design matrix exceeds 1000. For $p = 1600$ computing the LADLasso estimator takes more than 30 seconds and is over 3000 times slower than calculating RankLasso.

Table 1: Ratio of times needed to perform LADLasso (in *MTE*) and RankLasso (in *glmnet*)

dimension	t(LAD)/t(rL)
$n = 100, p = 100$	5.32
$n = 200, p = 400$	95.2
$n = 300, p = 900$	655
$n = 400, p = 1600$	3087

Figure 1 illustrates the average number of falsely classified predictors for different methods and under different simulation scenarios. In the case when predictors are independent, RankLasso satisfies assumptions of Wang and Zhu (2015, Theorem 2.1) and its NMP decreases with $p = 0.01n^2$. The same is true for LADLasso. As shown in plots of FDR and Power (Figures 3 and 4 in the appendix), LADLasso becomes more conservative than RankLasso for larger values of p_0 , which leads to slightly larger values of NMP. We can also observe that for independent predictors, the adaptive and thresholded versions perform similarly to the standard version of RankLasso. As expected, regular cross-validated Lasso performs very badly, when the error terms come from the Cauchy distribution. It suffers both from the loss of power and large FDR values. Also, it is interesting to observe that the first two rows in Figure 1 do not differ significantly, which shows that the performance of RankLasso for the realistic independent SNP data is very similar to its performance when the elements of the design matrix are drawn from the Gaussian distribution.

The behaviour of RankLasso changes significantly in the case when predictors are correlated. Namely, NMP of RankLasso increases with p . On the other hand, NMP of both adaptive and thresholded versions of RankLasso decrease with p , so these two methods are able to find the true model consistently. In the case when the relationship between the median value of the response variable and the predictors is linear, LADLasso has a larger power and a similar FDR to RankLasso (see Figures 3 and 4 in the appendix). In the result it returns substantially more of false positives and its NMP is slightly larger than the NMP of RankLasso. In the last row of Figure 1 we can observe that the lack of linearity has a negligible influence on the performance of RankLasso but substantially affects LADLasso, which now has a smaller power and larger FDR than RankLasso.

As shown in Figure 1, in the case of correlated predictors thresholded RankLasso is systematically better than adaptive RankLasso, even though both methods always select the same number of predictors. To explain this phenomenon, in Figure 2 we present the relationship between the number of false discoveries (FD) and the number of true positives

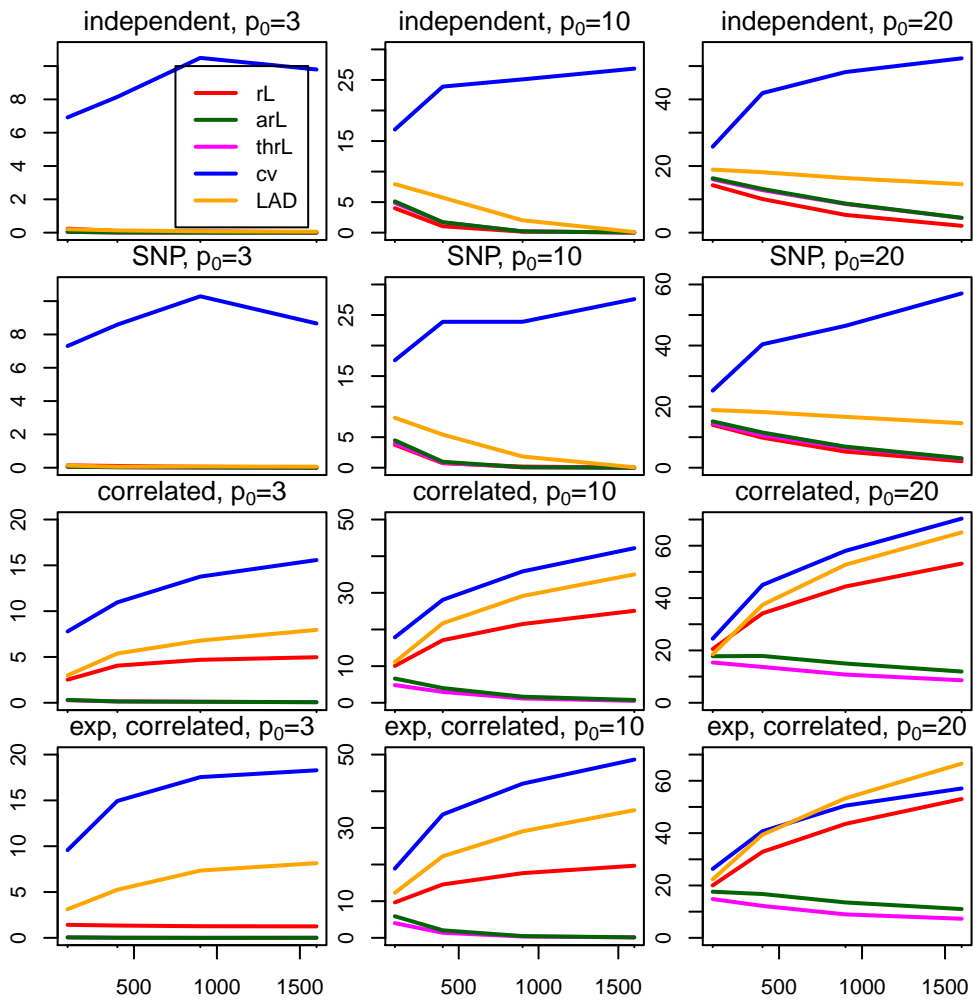


Figure 1: Plots of NMP (average number of misclassified predictors) as the function of p .

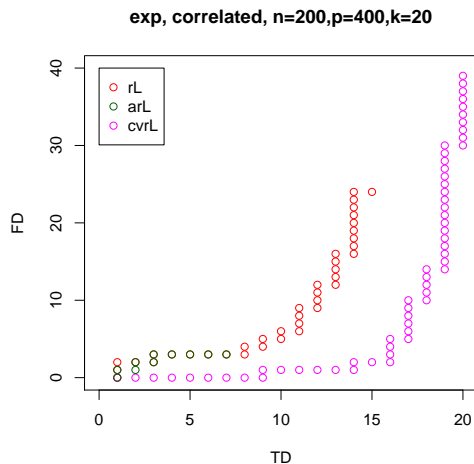


Figure 2: Plots of the number of false discoveries (FD) vs the number of true positives (TP).

(TP) along the decreasingly ordered sequence of absolute values of RankLasso estimates for one realization of the experiment. We consider RankLasso with the value of the tuning parameter selected as in (25) and by cross-validation (denoted by „cvrL”) and the adaptive RankLasso.

Figure 2 illustrates substantial problems with the ordering of predictors by the estimates provided by RankLasso with the relatively large value of λ from equation (25). Among three predictors with the largest absolute values of estimates of regression coefficients, 2 are false positives. The final estimate of RankLasso identifies 15 true predictors and returns 25 false discoveries. Adaptive RankLasso does not provide a substantially better ordering of estimates of regression coefficients but stops at the point when the number of false discoveries drastically increases. In the result, adaptive RankLasso returns 7 true positives and 3 false discoveries. Comparing cross-validated RankLasso to its more conservative version, we can observe that it selects many more regressors - it identifies all 20 of true predictors and returns 40 false discoveries. However, we can also observe that cross-validated RankLasso provides a much better ordering of the estimates of regression coefficients. The FD-TP curve of cross-validated RankLasso is substantially below the FD-TP curve corresponding to RankLasso with larger value of λ . Among 14 regressors with the largest absolute values of regression coefficient provided by cross-validated RankLasso there is only one false positive. Thus, when selection is stopped by matching the adaptive RankLasso, we identify 9 true positives and only 1 false discovery.

4. Analysis of real data

In this section we apply rank methods and their competitors for identifying relationships between expressions of different genes. Variations in gene expression levels may be related to phenotypic variations such as susceptibility to diseases and response to drugs.

Identifying the relationships between these gene expressions facilitates understanding the genetic-pathways and identifying regulatory genes influencing the disease processes. In the considered data set gene expression was interrogated in lymphoblastoid cell lines of 210 unrelated HapMap individuals The International HapMap Consortium (2005) from four populations (60 Utah residents with ancestry from northern and western Europe, 45 Han Chinese in Beijing, 45 Japanese in Tokyo, 60 Yoruba in Ibadan, Nigeria) Stranger et al. (2007). The data set can be found at <ftp://ftp.sanger.ac.uk/pub/genevar/> and was previously studied e.g. in Bradic et al. (2011); Fan et al. (2014). In our analysis we will concentrate on four genes. First of them is the gene CCT8, which was analyzed previously in Bradic et al. (2011). This gene is within the Down syndrome critical region on human chromosomen 21, on the minus strand. The over-expression of CCT8 may be associated with Down syndrome phenotypes. We also consider gene CHRNA6, which was previously investigated in Fan et al. (2014) and is thought to be related to activation of dopamine releasing neurons with nicotine. Since the data on expression levels of these two genes contained only few relatively small outliers, we additionally considered genes PRAME and Hs.444277-S, where the influence of outliers is more pronounced. The boxplots of these gene expressions can be found in Figure 5 in the appendix.

We start with preparing the data set using three pre-processing steps as in Wang et al. (2012): we remove each probe for which the maximum expression among 210 individuals is smaller than the 25-th percentile of the entire expression values, we remove any probe for which the range of the expression among 210 individuals is smaller than 2 and finally we select 300 genes, whose expressions are the most correlated to the expression level of the analyzed gene.

Next, the data set is divided into two parts: the training set with randomly selected 180 individuals and the test set with remaining 30 individuals. Five procedures from Subsection 3 are used to select important predictors using the training set and their accuracy is evaluated using the test set. As a measure of accuracy we cannot use the standard mean square prediction error, because in the single index model (2) the link function g is unknown. Since the link g is increasing wrt the first variable, we can expect that the ordering between values of the response variables Y_i should be well predicted by the ordering between scalar products $\beta'X_i$. Moreover, from Theorem 1 we know that $\theta^0 = \gamma_\beta\beta$ for the positive multiplicative number γ_β , so the ordering between Y_i should be also well predicted by the ordering between $(\theta^0)'X_i$. Therefore, as a accuracy measure of estimators we use the *ordering prediction quality* (OPQ), which is defined as follows: let $T = n_t(n_t - 1)/2$ be the number of different two-element subsets from the test set. The subset $\{i, j\}$ from the test set is properly ordered, if the sign of $Y_i^{test} - Y_j^{test}$ coincides with the sign of $\hat{\theta}'X_i^{test} - \hat{\theta}'X_j^{test}$, where $\hat{\theta}$ is some estimator of θ^0 based on the training set. Let P denote the number of two-element subsets from the test-set, that are properly ordered. The *ordering prediction quality* is defined as

$$OPQ = \frac{P}{T} . \quad (26)$$

Tables 2 and 3 report the average number of selected predictors and the average values of OPQ over 200 random splits into the training and the test sets. These values were calculated for all five model selection methods considered in the simulation study.

Table 2: Average number of selected predictors (SP)

SP	rL	arL	rLth	LADLasso	cv
CCT8	16	6	6	14	25
CHRNA6	19	7	7	8	52
PRAME	16	6	6	6	0
Hs.444277-S	4	2	2	0	0

Table 3: Average values of the Ordering Prediction Quality (OPQ, (26))

OPQ	rL	arL	rLth	LADLasso	cv
CCT8	0.74	0.73	0.74	0.73	0.76
CHRNA6	0.68	0.66	0.64	0.66	0.71
PRAME	0.65	0.62	0.61	0.60	0.08
Hs.444277-S	0.59	0.59	0.58	0.07	0.03

We can observe that for CCT8 the numbers of selected predictors and the prediction accuracy of RankLasso and LADLasso are similar. The thresholded and adaptive RankLasso provide a similar prediction accuracy with much smaller number of predictors. Interestingly, regular cross-validated Lasso yields the best Ordering Prediction Quality, which however requires 4 times as many predictors as thresholded or adaptive RankLasso. For CHRNA6 we see that the support of RankLasso is substantially larger than for LADLasso and adaptive and thresholded RankLasso, which results in slightly better prediction accuracy. Again, the best prediction is obtained from regular cross-validated Lasso, which however uses much larger number of predictors.

The performance of regular Lasso drastically deteriorates for the remaining two genes, whose expressions contain substantially larger outliers. Here regular cross-validated Lasso in most of the cases is not capable of identifying any predictors. In the case of the gene Hs.444277-S the same is true about LADLasso. In the case of the PRAME gene, the highest prediction accuracy is provided by regular RankLasso, which however requires almost three times as many predictors as adaptive or thresholded Lasso. In the case of Hs.444277-S RankLasso identifies 4 predictors, while its modified versions select only 2 genes. These simple models still allow to predict the ordering of gene expressions of Hs.444277-S with accuracy close to 60%.

5. Discussion

Lasso is a well established method for estimation of parameters in the high dimensional regression models. It is also well understood that it can recover the support of the vector of regression coefficients only under very stringent conditions relating the sparsity of this vector and the structure of correlations between columns in the design matrix. This phenomenon can be well explained using the theory of Approximate Message Algorithms (AMP), see e.g. Bayati and Montanari (2012); Su et al. (2017); Wang et al. (2017); Weinstein et al. (2020), which allows to predict the mean squared error of Lasso estimators as the function of the tuning parameter λ . Interestingly, this error tends to take very large values for large

values of λ , which leads to early false discoveries on the Lasso path, see Su et al. (2017). Smaller values of λ typically yield smaller mean squared error and better ordering of Lasso estimates, so the estimated regression coefficients of these early false discoveries can get smaller values than those by true discoveries. Therefore, thresholded and adaptive Lasso can recover the true model under much weaker assumptions than regular Lasso.

In this article we show that the similar phenomenon holds for RankLasso, which can be used to identify predictors in the single-index model, with unknown monotonic link function and unknown error distribution. Our theoretical and empirical results illustrate that the thresholded and adaptive versions of RankLasso can properly identify the predictors even when the link function is non-linear, predictors are highly correlated and the error comes from the Cauchy distribution. When the identified model contains a small number of predictors, the link function can be further estimated using the estimation method of Balabdaoui et al. (2019) and its extensions based on the robust loss functions.

While our results demonstrate clearly a potential of the modified versions of RankLasso, there still remain open questions related e.g. to the choice of the optimal tuning parameter λ and the optimal threshold for the thresholded version or the optimal selection of weights in the weighted version. In the future we plan to extend our method by exploring different approaches to threshold selection, including modifications of the knockoff methodology of Barber and Candès (2015) and Candès et al. (2018). Moreover, combining the rank-based approach with the choice of the threshold based on the information criterion as in Bogdan et al. (2008) or Pokarowski and Mielniczuk (2015) seems also to be an interesting problem to investigate.

The theoretical analysis of RankLasso provided in this paper and the results of Kos and Bogdan (2020) on the asymptotic FDR control of the Sorted L-One Penalized Estimator (SLOPE, Bogdan et al. (2015)) for the regular multiple regression, pave the way for construction of the rank version of SLOPE, so as to obtain the asymptotic FDR control in the single index model in the case when regressors are independent random variables. Concerning the adaptive version of RankLasso or RankSLOPE, it would be of interest to develop an adaptive selection of weights in the spirit of Spike and Slab Lasso (Ročková and George, 2018) or the adaptive Bayesian version of SLOPE (Jiang et al., 2019).

Finally, a single index model (2), which is studied in the paper, can be generalized to

$$Y_i = g(\beta_1' X_i, \dots, \beta_d' X_i, \varepsilon_i),$$

where $\beta_1, \beta_2, \dots, \beta_d$ are p -dimensional vectors and $d < p$. Such model was studied, for instance, in Cohen et al. (2012). The question arises whether the rank-based approach can properly identify true predictors also for the case $d > 1$. A similar problem in the context of misspecified binary regression was considered in Kubkowski and Mielniczuk (2018), where it is shown that the vector of parameters estimated by logistic regression can be expressed as a linear combination of vectors β_1, \dots, β_d under a natural extension of our Assumption 2. The analysis from Kubkowski and Mielniczuk (2018) seems to be a good starting point in establishing a similar relation for our parameter θ_0 given in (7). If this holds, then the theoretical results concerned with the identification of significant predictors by RankLasso should be possible to obtain by a relatively straightforward extension of our proof techniques.

Acknowledgments

We would like to thank Patrick Tardivel for helpful comments. We gratefully acknowledge the grant of the Wroclaw Center of Networking and Supercomputing (WCSS), where most of the computations were performed. Małgorzata Bogdan is supported by Polish National Science Center grants no. 2016/23/B/ST1/00454. We would like also to thank the Associate Editor and two reviewers for their comments, which helped us improve the manuscript.

Appendix

In Section A of the appendix we provide results for Rank-Lasso and its modifications in the low-dimensional scenario. Besides, additional results of numerical experiments are in Section E. The proofs of results obtained in the main paper are given in Sections B and C. Finally, proofs of results from Section A in the appendix are given in Section D.

Appendix A. Low-dimensional scenario

In this section we consider properties of rank estimators in the case where the number of predictors is fixed. In the first part we focus on RankLasso and in the second part we study thresholded and weighted RankLasso.

We assume, without loss of generality, that $T = \{1, \dots, p_0\}$ for some $0 < p_0 < p$, so the response variable Y depends on first p_0 predictors. RankLasso estimates the set T by

$$\hat{T} = \{1 \leq j \leq p : \hat{\theta}_j \neq 0\}.$$

The results, that we obtain in this subsection, are asymptotic, so we can replace the true parameter θ^0 in (10) by

$$\theta^* = \frac{n}{n-1} \theta^0 = H^{-1} \mu. \quad (27)$$

Obviously, it does not change the set of relevant predictors T . We also decompose the matrix $H = \mathbb{E}X_1 X_1'$ as

$$H = \begin{pmatrix} \widehat{H_1}^{p_0 \times p_0} & \widehat{H_2}^{p_0 \times (p-p_0)} \\ H_2' & H_3 \end{pmatrix},$$

so the matrix H_1 describes correlations between relevant predictors and the matrix H_2 contains correlations between relevant and irrelevant predictors.

A.0.1 MODEL SELECTION CONSISTENCY OF RANKLASSO

The next result provides sufficient and necessary conditions for RankLasso to be model selection consistent. They are similar to the results proved in Zou (2006, Theorem 1) and Zhao and Yu (2006, Theorem 1), which concern model selection in the linear model. Theorem 8 extends these results to the single index model (2), which does not require any assumptions on the form of the link function (except being monotonic) nor the distribution of the noise variable.

Theorem 8 *Suppose that Assumption 3 is satisfied, $\mathbb{E}|X_1|^4 < \infty$ and $\lambda \rightarrow 0$, $\sqrt{n}\lambda \rightarrow \infty$.*

(a) *If $\lim_{n \rightarrow \infty} P(\hat{T} = T) \rightarrow 1$, then*

$$|H_2' H_1^{-1} \text{sign}(\theta_T^*)|_\infty \leq 1, \quad (28)$$

where θ^* is defined in (27).

(b) *If the inequality*

$$|H_2' H_1^{-1} \text{sign}(\theta_T^*)|_\infty < 1 \quad (29)$$

holds, then $\lim_{n \rightarrow \infty} P(\hat{T} = T) \rightarrow 1$.

The sufficiency of (29) for model selection consistency of RankLasso was established in Wang and Zhu (2015, Corollary 2.1). In Theorem 8 we strenghten this result by showing that it is almost the necessary condition. The condition (29), called the irrepresentable condition (Zhao and Yu, 2006), is restrictive and satisfied only in some very special cases, like when predictors are independent or when the correlations between „neighboring” variables decay exponentially with their distance. Therefore, RankLasso usually is not consistent in model selection. However, as shown in the following Lemma 9, it can consistently estimate θ^* under much weaker assumptions. This result will be crucial for Subsection A.0.2, where we establish model selection consistency of the thresholded and weighted versions of RankLasso under such weaker assumptions. The next fact is a generalization of Knight and Fu (2000, Theorem 2).

Lemma 9 *Suppose that Assumption 1 is satisfied and $\mathbb{E}|X_1|^4 < \infty$. Let a_n be a sequence such that $a_n \rightarrow 0$, $\frac{1}{a_n \sqrt{n}} \rightarrow b \in [0, \infty)$, $\frac{\lambda}{a_n} \rightarrow c \in [0, \infty)$. Then the RankLasso estimator $\hat{\theta}$ in (5) satisfies*

$$\frac{1}{a_n} (\hat{\theta} - \theta^*) \rightarrow_d \arg \min_{\theta} V(\theta),$$

where

$$V(\theta) = \frac{1}{2} \theta' H \theta + b \theta' W + c \sum_{j \in T} \theta_j \text{sign}(\theta_j^*) + c \sum_{j \notin T} |\theta_j|$$

and W has a normal $N(0, D)$ distribution with the matrix D given in Lemma 17 in Section D.

A.0.2 MODIFICATIONS OF RANKLASSO

In this subsection we introduce two modifications of RankLasso and study their properties in the low-dimensional case.

First of these modifications, the weighted RankLasso estimator, is an analogue of the adaptive Lasso, which was proposed in Zou (2006). The main idea of this approach relies on the application of different weights for different predictors, depending on the value of some initial estimator $\tilde{\theta}$ of θ^* . This estimator needs to be \sqrt{n} -consistent, i.e. it satisfies

$$\sqrt{n} (\tilde{\theta} - \theta^*) = O_P(1) . \quad (30)$$

In particular, according to Lemma 9, $\tilde{\theta}$ can be chosen as the RankLasso estimator with the regularization parameter that behaves as $O(1/\sqrt{n})$. Then, the weighted RankLasso estimator $\hat{\theta}^a$ is obtained as

$$\hat{\theta}^a = \arg \min_{\theta \in \mathbb{R}^p} Q(\theta) + \lambda \sum_{j=1}^p w_j |\theta_j|, \quad (31)$$

where $w_j = |\tilde{\theta}_j|^{-1}$, $j = 1, \dots, p$ and $Q(\theta)$ is given in (6).

Let \hat{T}^a denote a set $\{j \in \{1, \dots, p\} : \hat{\theta}_j^a \neq 0\}$. The properties of $\hat{\theta}^a$ are described in the next theorem.

Theorem 10 *Consider the weighted RankLasso estimator (31) with $\tilde{\theta}$ satisfying (30). Suppose that Assumption 3 is satisfied and $\mathbb{E}|X_1|^4 < \infty$. If $n\lambda \rightarrow \infty$ and $\sqrt{n}\lambda \rightarrow c \in [0, \infty)$, then*

(a) $\lim_{n \rightarrow \infty} \mathbb{P}(\text{sign}(\hat{\theta}^a) = \text{sign}(\beta)) = 1$, where the equality of signs of two vectors is understood componentwise,

(b) $\sqrt{n}(\hat{\theta}_T^a - \theta_T^*) \rightarrow_d N(-H_1^{-1}\bar{c}, H_1^{-1}D_1H_1^{-1})$, where $\bar{c} = c\left(\frac{1}{\theta_1^*}, \dots, \frac{1}{\theta_{p_0}^*}\right)$, θ^* is defined in (27) and the matrix D_1 is the $(p_0 \times p_0)$ upper-left submatrix of the matrix D defined in Lemma 17 in Section D.

Now, we introduce the second modification, which is thresholded RankLasso. This estimator is denoted by $\hat{\theta}^{th}$ and defined in (19).

Theorem 11 *Suppose that Assumption 3 is satisfied and $\mathbb{E}|X_1|^4 < \infty$. If $\sqrt{n}\lambda \rightarrow 0$, $\delta \rightarrow 0$ and $\sqrt{n}\delta \rightarrow \infty$, then*

(a) $\lim_{n \rightarrow \infty} \mathbb{P}(\text{sign}(\hat{\theta}^{th}) = \text{sign}(\beta)) = 1$, where the equality of signs of two vectors is understood componentwise,

(b) $\sqrt{n}(\hat{\theta}_T^{th} - \theta_T^*) \rightarrow_d N(0, (H^{-1}DH^{-1})_1)$, where θ^* is defined in (27) and $(H^{-1}DH^{-1})_1$ is the $(p_0 \times p_0)$ upper-left submatrix of $H^{-1}DH^{-1}$.

Theorems 10 and 11 state that weighted and thresholded RankLasso behave almost like the oracle. They are asymptotically able to identify the support and recognize the signs of coordinates of the true parameter β . Moreover, they estimate nonzero coordinates of θ^* with the standard \sqrt{n} -rate. The crucial fact is that these theorems hold even when the irrepresentable condition is not satisfied. Thus, both modifications of RankLasso allow to identify the true model under much weaker assumptions than vanilla RankLasso.

Theorems 10 and 11 work in the single index model (2) and they do not require any assumptions on the distribution of the noise variables or the form of the increasing link function g . Comparing to other theoretical results concerning model selection with the robust loss functions, like Wang et al. (2007, Theorem), Johnson and Peng (2008, Theorem 2.1), Song and Ma (2010, Theorem 4.2), Rejchel (2017b, Theorem 4.1), Avella-Medina and Ronchetti (2018, Theorem 2), the assumptions of Theorems 10 and 11 are slightly stronger. Specifically, in Theorems 10 and 11 the standard condition on the existence of the second moment of predictors is replaced by the assumption on the existence of the fourth moment.

This results from the fact that we work with the nonlinear model and the quadratic loss function. Apart from computational efficiency, application of the quadratic loss function allows us to solve the theoretical issues related to the dependency between ranks. The stronger assumption on the moments of predictors seems to be a relatively small prize for the gain in computational complexity, which allows to handle large data sets. Moreover, according to the simulation study reported in Section 3 of the main paper, for such large data sets our method has substantially better statistical properties than LADLasso, which is a popular methodology for robust model selection.

Appendix B. Results from Subsection 2.2

Notice that for $Q(\theta)$ defined in (6) we have

$$Q(\theta) = \frac{1}{2n} \sum_{i=1}^n (R_i/n - \theta' X_i)^2 + \theta' \bar{X}/2 - \frac{n+1}{4n} + 1/8.$$

Therefore, due to the fact that predictors X_i are centred we will consider $Q(\theta)$ without subtracting 0.5, that is

$$Q(\theta) = \frac{1}{2n} \sum_{i=1}^n (R_i/n - \theta' X_i)^2$$

in all proofs in this appendix. It will simplify notations.

Proof [Proof of Theorem 1] We start with proving the first part of the theorem. Argumentation is similar to the proof of Li and Duan (1989, Theorem 2.1), but it has to be adjusted to ranks, which are not independent random variables (as distinct from Y_1, \dots, Y_n). Obviously, we have

$$\mathbb{E}Q(\theta) = \frac{1}{2n^3} \sum_{i=1}^n \mathbb{E}R_i^2 - \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}R_i \theta' X_i + \frac{1}{2n} \sum_{i=1}^n \mathbb{E}(\theta' X_i)^2.$$

Vectors $(X_1, Y_1), \dots, (X_n, Y_n)$ are i.i.d. and X_i are centred, so for all $i \neq 1$

$$\begin{aligned} \mathbb{E}R_i \theta' X_i &= \mathbb{E}\mathbb{I}(Y_1 \leq Y_i) \theta' X_i + \sum_{j \neq \{1, i\}} \mathbb{E}\mathbb{I}(Y_j \leq Y_i) \theta' X_i \\ &= \mathbb{E}\mathbb{I}(Y_i \leq Y_1) \theta' X_1 + \sum_{j \neq \{1, i\}} \mathbb{E}\mathbb{I}(Y_j \leq Y_1) \theta' X_1 = \mathbb{E}R_1 \theta' X_1. \end{aligned}$$

Moreover, ranks R_1, \dots, R_n have the same distribution, so $\sum_{i=1}^n \mathbb{E}R_i^2 = n\mathbb{E}R_1^2$. Therefore, we obtain that $\mathbb{E}Q(\theta) = \frac{1}{2}\mathbb{E}\left(\frac{R_1}{n} - \theta' X_1\right)^2$. Using Jensen's inequality and Assumption 2 we have

$$\begin{aligned} \mathbb{E}Q(\theta) &= \frac{1}{2}\mathbb{E}\mathbb{E}\left[\left(\frac{R_1}{n} - \theta' X_1\right)^2 \mid \beta' X_i, \varepsilon_i, i = 1, \dots, n\right] \\ &\geq \frac{1}{2}\mathbb{E}\left[\mathbb{E}\left(\frac{R_1}{n} - \theta' X_1 \mid \beta' X_i, \varepsilon_i, i = 1, \dots, n\right)\right]^2 = \frac{1}{2}\mathbb{E}\left[\frac{R_1}{n} - \mathbb{E}(\theta' X_1 \mid \beta' X_1)\right]^2 \\ &= \frac{1}{2}\mathbb{E}\left(\frac{R_1}{n} - d_\theta \beta' X_1\right)^2 \geq \min_{d \in \mathbb{R}} \mathbb{E}Q(d\beta). \end{aligned}$$

Obviously, we have $\min_d \mathbb{E}Q(d\beta) = \mathbb{E}Q(\gamma_\beta\beta)$, where γ_β is defined in (12). Since θ^0 is the unique minimizer of $\mathbb{E}Q(\theta)$, we obtain the first part of the theorem.

Next, we establish the second part of the theorem. Denote $Z = \beta' X_1$ and $\varepsilon = \varepsilon_1$. It is clear that $\gamma_\beta > 0$ is equivalent to $Cov(Z, F(g(Z, \varepsilon))) > 0$. This covariance can be expressed as

$$\mathbb{E}ZF(g(Z, \varepsilon)) = \mathbb{E}h(\varepsilon), \quad (32)$$

where $h(a) = \mathbb{E}[ZF(g(Z, \varepsilon)) | \varepsilon = a] = \mathbb{E}ZF(g(Z, a))$ for arbitrary a . This fact simply follows from $\mathbb{E}Z = 0$ and independence between Z and ε . If F is increasing and g is increasing with respect to the first variable, then $h(a) > 0$ for arbitrary a by Lemma 12 given below. Obviously, it implies that (32) is positive. ■

The following result was used in the proof of Theorem 1. It is a simple and convenient adaptation of a well-known fact concerning covariance of nondecreasing functions (Thorisson, 1995). Its proof follows Kubkowski (2019, Lemma A.44).

Lemma 12 *Let U be a random variable that is not concentrated at one point, i.e. $P(U = u) < 1$ for each $u \in \mathbb{R}$. Moreover, let $f, h : \mathbb{R} \rightarrow \mathbb{R}$ be increasing functions. Then $Cov(f(U), h(U)) > 0$.*

Proof For all real $a \neq b$ we have $[f(a) - f(b)][h(a) - h(b)] > 0$, because f, h are increasing. Let V be an independent copy of U . Then $P(U \neq V) > 0$ and we obtain

$$\begin{aligned} 0 &< \mathbb{E}[f(U) - f(V)][h(U) - h(V)] \mathbb{I}(U \neq V) \\ &= \mathbb{E}[f(U) - f(V)][h(U) - h(V)] \\ &= 2\mathbb{E}f(U)h(U) - 2\mathbb{E}f(U)\mathbb{E}h(U) \\ &= 2Cov(f(U), h(U)). \end{aligned}$$

■

Appendix C. Results from Subsection 2.3

To prove Theorem 2 we need three auxiliary results: Lemma 13, Lemma 14 and Lemma 15. The first one is borrowed from van de Geer (2016, Corollary 8.2), while the second one is its adaptation to U -statistics.

Lemma 13 *Suppose that Z_1, \dots, Z_n are i.i.d. random variables and there exists $L > 0$ such that $C^2 = \mathbb{E} \exp(|Z_1|/L)$ is finite. Then for arbitrary $u > 0$*

$$P\left(\frac{1}{n} \sum_{i=1}^n (Z_i - \mathbb{E}Z_i) > 2L \left(C \sqrt{\frac{2u}{n}} + \frac{u}{n}\right)\right) \leq \exp(-u).$$

Lemma 14 *Consider a U -statistic*

$$U = \frac{1}{n(n-1)} \sum_{i \neq j} h(Z_i, Z_j)$$

with a kernel h based on i.i.d. random variables Z_1, \dots, Z_n . Suppose that there exists $L > 0$ such that $C^2 = \mathbb{E} \exp(|h(Z_1, Z_2)|/L)$ is finite. Then for arbitrary $u > 0$

$$P\left(U - \mathbb{E}U > 2L\left(C\sqrt{\frac{6u}{n}} + \frac{3u}{n}\right)\right) \leq \exp(-u).$$

Proof Let $g(z_1, z_2) = h(z_1, z_2) - \mathbb{E}h(Z_1, Z_2)$ and \tilde{U} be a U -statistic with a kernel g . Using Hoeffding's decomposition we can represent every U -statistic as an average of (dependent) averages of independent random variables (Serfling, 1980), i.e.

$$\tilde{U} = \frac{1}{n!} \sum_{\pi} \frac{1}{N} \sum_{i=1}^N g(Z_{\pi(i)}, Z_{\pi(N+i)}), \quad (33)$$

where $N = \lfloor \frac{n}{2} \rfloor$ and the first sum on the right-hand side of (33) is taken over all permutations π of a set $\{1, \dots, n\}$. Take arbitrary $s > 0$. Then using Jensen's inequality and the fact that Z_1, \dots, Z_n are i.i.d. we obtain

$$\begin{aligned} \mathbb{E} \exp(s\tilde{U}) &\leq \frac{1}{n!} \sum_{\pi} \mathbb{E} \exp\left[\frac{s}{N} \sum_{i=1}^N g(Z_{\pi(i)}, Z_{\pi(N+i)})\right] \\ &= \mathbb{E} \exp\left[\frac{s}{N} \sum_{i=1}^N g(Z_i, Z_{N+i})\right]. \end{aligned} \quad (34)$$

We have the average of N -i.i.d. random variables in (34), so we can repeat argumentation from the proof of van de Geer (2016, Corollary 8.2). Finally, we use the simple inequality $N \geq n/3$ for $n \geq 2$. \blacksquare

Lemma 15 *Suppose Assumptions 3 and 4 are satisfied. For arbitrary $j = 1, \dots, p$ and $u > 0$ we have*

$$P\left(\frac{1}{n} \sum_{i=1}^n X_{ij} X'_i \theta^0 - \frac{n-1}{n} \mu_j > 5 \frac{\tau^2}{\sqrt{\kappa}} \left(2\sqrt{\frac{2u}{n}} + \frac{u}{n}\right)\right) \leq \exp(-u). \quad (35)$$

Besides, if X_1 has a normal distribution $N(0, H)$, then we can drop τ and κ in (35).

Proof Fix $j = 1, \dots, p$ and $u > 0$. Recall that $H\theta^0 = \frac{n-1}{n}\mu$ by (10). We work with an average of i.i.d. random variables, so we can use Lemma 13. We only have to find $L, C > 0$ such that

$$\mathbb{E} \exp(|X_{1j} X'_1 \theta^0|/L) \leq C^2.$$

For each positive number a, b, s we have the inequality $ab \leq \frac{a^2}{2s^2} + \frac{b^2 s^2}{2}$. Applying this fact and the Schwarz inequality we obtain

$$\mathbb{E} \exp(|X_{1j} X'_1 \theta^0|/L) \leq \sqrt{\mathbb{E} \exp\left(\frac{X_{1j}^2}{s^2 L}\right) \mathbb{E} \exp\left(\frac{s^2 (X'_1 \theta^0)^2}{L}\right)} \quad (36)$$

and the number s will be chosen later. The variable X_{1j} is subgaussian, so using Baraniuk et al. (2011, Lemma 7.4) we can bound the first expectation in (36) by $\left(1 - \frac{2\tau^2}{s^2L}\right)^{-1/2}$ provided that $s^2L > 2\tau^2$. The second expectation in (36) can be bounded using subgaussianity of the vector X_1 in the following way

$$\mathbb{E} \exp\left(\frac{s^2(X_1'\theta^0)^2}{L}\right) \leq \left(1 - \frac{2s^2\tau^2|\theta^0|_2^2}{L}\right)^{-1/2},$$

provided that $2s^2\tau^2|\theta^0|_2^2 < L$. From Theorem 1 we obtain two equalities $\theta^0 = \gamma_\beta\beta$ and $\gamma_\beta = \frac{\frac{n-1}{n} \mathbb{E}\mathbb{I}(Y_2 \leq Y_1)\beta'X_1}{\beta'H\beta}$. Recall that κ is the smallest eigenvalue of the matrix H_T . Therefore, we obtain a bound

$$|\theta^0|_2^2 = \gamma_\beta^2|\beta_T|_2^2 \leq \kappa^{-1},$$

because

$$\mathbb{E}\mathbb{I}(Y_2 \leq Y_1)\beta'X_1 \leq \sqrt{\beta_T'H_T\beta_T}.$$

Taking $L = 2.2\tau^2/\sqrt{\kappa}$ and $s^2 = \sqrt{\kappa}$ we obtain $C \leq 2$, that finishes the proof of the first part of the lemma.

Next, we assume that $X_1 \sim N(0, H)$. Therefore, $X_{1j} \sim N(0, 1)$ and $(\theta^0)'X_1 \sim N(0, (\theta^0)'H\theta^0)$. The argumentation is as above with $s^2 = 1$. We only use the inequality $(\theta^0)'H\theta^0 \leq 1$ and the equality

$$\mathbb{E} \exp\left(\frac{(X_1'\theta^0)^2}{L}\right) = \left(1 - 2(\theta^0)'H\theta^0/L\right)^{-1/2},$$

provided that $L > 2(\theta^0)'H\theta^0$. Therefore, we can take $L = 2.2$. ■

Lemma 16 *Suppose that Assumption 4 and (14) are satisfied. Then for arbitrary $a \in (0, 1)$, $q \geq 1$, $\xi > 1$ with probability at least $1 - 2a$ we have $\bar{F}_q(\xi) \geq F_q(\xi)/2$.*

Proof

Fix $a \in (0, 1)$, $q \geq 1$, $\xi > 1$. We start with considering the l_∞ -norm of the matrix

$$\left|\frac{1}{n}X'X - H\right|_\infty = \max_{j,k=1,\dots,p} \left|\frac{1}{n} \sum_{i=1}^n X_{ij}X_{ik} - \mathbb{E}X_{1j}X_{1k}\right|.$$

Fix $j, k \in \{1, \dots, p\}$. Using subgaussianity of predictors, Lemma 13 and argumentation similar to the proof of Lemma 15 we have for $u = \log(p^2/a)$

$$P\left(\left|\frac{1}{n} \sum_{i=1}^n X_{ij}X_{ik} - \mathbb{E}X_{1j}X_{1k}\right| > K_2\tau^2\sqrt{\frac{\log(p^2/a)}{n}}\right) \leq \frac{2a}{p^2},$$

where K_2 is an universal constant.

Therefore, using union bounds we obtain

$$P\left(\left|\frac{1}{n}X'X - H\right|_\infty > K_2\tau^2\sqrt{\frac{\log(p^2/a)}{n}}\right) \leq 2a. \quad (37)$$

Obviously, we have $|(X'X/n - H)\theta|_\infty \leq |X'X/n - H|_\infty |\theta|_1$ and for each $\theta \in \mathcal{C}(\xi)$ and $q > 1$ we obtain $|\theta|_1 \leq (1 + \xi)|\theta_T|_1 \leq (1 + \xi)p_0^{1-1/q}|\theta_T|_q \leq (1 + \xi)p_0^{1-1/q}|\theta|_q$. Therefore, for each $\theta \in \mathcal{C}(\xi)$

$$\begin{aligned} \frac{p_0^{1/q}|X'X\theta/n|_\infty}{|\theta|_q} &\geq \frac{p_0^{1/q}|H\theta|_\infty}{|\theta|_q} - \frac{p_0^{1/q}|X'X/n - H|_\infty|\theta|_1}{|\theta|_q} \\ &\geq \frac{p_0^{1/q}|H\theta|_\infty}{|\theta|_q} - (1 + \xi)p_0|X'X/n - H|_\infty. \end{aligned}$$

Taking infimum and using (37), we have the following probabilistic inequality

$$\bar{F}_q(\xi) \geq F_q(\xi) - K_2(1 + \xi)p_0\tau^2\sqrt{\frac{\log(p^2/a)}{n}}.$$

To finish the proof we use (14) with K_1 being sufficiently large. \blacksquare

Proof [Proof of Theorem 2]

Let $a \in (0, 1)$ be arbitrary. The main part of the proof is to show that with high probability

$$|\hat{\theta} - \theta^0|_q \leq \frac{2\xi p_0^{1/q}\lambda}{(\xi + 1)\bar{F}_q(\xi)}. \quad (38)$$

Then we apply Lemma 16 to obtain (16).

Thus, we focus on proving (38). Denote $\Omega = \{|\nabla Q(\theta^0)|_\infty \leq \frac{\xi-1}{\xi+1}\lambda\}$. We start with lower bounding probability of Ω . For A defined in (11) and every $j = 1, \dots, p$ we obtain

$$\nabla_j Q(\theta^0) = \left[\frac{1}{n} \sum_{i=1}^n X_{ij} X_i' \theta^0 - \frac{n-1}{n} \mu_j \right] + \frac{n-1}{n} [\mu_j - A^j] - \frac{1}{n^2} \sum_{i=1}^n X_{ij}, \quad (39)$$

so if we find probabilistic bounds for each term on the right-hand side of (39), then using union bounds we get the bound for $|\nabla Q(\theta^0)|_\infty$. Consider the middle term in (39). By (24) we apply Lemma 14 with $h(z_1, z_2) = \frac{1}{2} [\mathbb{I}(y_2 \leq y_1)x_{1j} + \mathbb{I}(y_1 \leq y_2)x_{2j}]$. Variables X_{1j} and X_{2j} are i.i.d., so for arbitrary $L > 0$ we have

$$\mathbb{E} \exp(|h(Z_1, Z_2)|/L) \leq [\mathbb{E} \exp(|X_{1j}|/(2L))]^2. \quad (40)$$

Using the fact that the variable X_{1j} is subgaussian we bound (40) by $4 \exp\left(\frac{\tau^2}{4L^2}\right)$. Taking $L = \tau$ and $u = \log(p/a)$ in Lemma 14 we obtain for some universal constant K_1

$$P\left(A^j - \mu_j > K_1\tau\sqrt{\frac{\log(p/a)}{n}}\right) \leq \frac{a}{p}.$$

The third term in (39) can be handled similarly using Lemma 13. To obtain the bound for the first term in (39) we apply Lemma 15. Taking these results together and using union bounds we obtain that $P(\Omega) \geq 1 - K_2a$ provided that λ satisfies (15).

In further argumentation we consider only the event Ω . Besides, we denote $\tilde{\theta} = \hat{\theta} - \theta^0$, where $\hat{\theta}$ is a minimizer of a convex function (5), that is equivalent to

$$\begin{cases} \nabla_j Q(\hat{\theta}) = -\lambda \text{sign}(\hat{\theta}_j) & \text{for } \hat{\theta}_j \neq 0, \\ |\nabla_j Q(\hat{\theta})| \leq \lambda & \text{for } \hat{\theta}_j = 0, \end{cases} \quad (41)$$

where $j = 1, \dots, p$.

First, we prove that $\tilde{\theta} \in \mathcal{C}(\xi)$. Here our argumentation is standard (Ye and Zhang, 2010). From (41) and the fact that $|\tilde{\theta}|_1 = |\tilde{\theta}_T|_1 + |\tilde{\theta}_{T'}|_1$ we can calculate

$$\begin{aligned} 0 &\leq \tilde{\theta}' X' X \tilde{\theta} / n = \tilde{\theta}' \left[\nabla Q(\hat{\theta}) - \nabla Q(\theta^0) \right] \\ &= \sum_{j \in T} \tilde{\theta}_j \nabla_j Q(\hat{\theta}) + \sum_{j \in T'} \hat{\theta}_j \nabla_j Q(\hat{\theta}) - \tilde{\theta}' \nabla Q(\theta^0) \\ &\leq \lambda \sum_{j \in T} |\tilde{\theta}_j| - \lambda \sum_{j \in T'} |\hat{\theta}_j| + |\tilde{\theta}|_1 |\nabla Q(\theta^0)|_\infty \\ &= [\lambda + |\nabla Q(\theta^0)|_\infty] |\tilde{\theta}_T|_1 + [|\nabla Q(\theta^0)|_\infty - \lambda] |\tilde{\theta}_{T'}|_1. \end{aligned}$$

Thus, using the fact that we consider the event Ω we get

$$|\tilde{\theta}_{T'}|_1 \leq \frac{\lambda + |\nabla Q(\theta^0)|_\infty}{\lambda - |\nabla Q(\theta^0)|_\infty} |\tilde{\theta}_T|_1 \leq \xi |\tilde{\theta}_T|_1.$$

Therefore, from the definition of $\bar{F}_q(\xi)$ we have

$$|\hat{\theta} - \theta^0|_q \leq \frac{p_0^{1/q} |X' X (\hat{\theta} - \theta^0) / n|_\infty}{\bar{F}_q(\xi)} \leq p_0^{1/q} \frac{|\nabla Q(\hat{\theta})|_\infty + |\nabla Q(\theta^0)|_\infty}{\bar{F}_q(\xi)}.$$

Using (41) and the fact, that we are on Ω , we obtain (38).

The case $X_1 \sim N(0, H)$ is a consequence of the analogous part of Lemma 15. \blacksquare

Proof [Proof of Corollary 3] The proof is a simple consequence of the bound (16) with $q = \infty$ obtained in Theorem 2. Indeed, for arbitrary predictors $j \in T$ and $k \notin T$ we obtain

$$|\hat{\theta}_j| \geq |\theta_j^0| - |\hat{\theta}_j - \theta_j^0| \geq \theta_{min}^0 - |\hat{\theta} - \theta^0|_\infty > \frac{4\xi\lambda}{(\xi + 1)F_\infty(\xi)} \geq |\hat{\theta}_k - \theta_k^0| = |\hat{\theta}_k|. \quad \blacksquare$$

Proof [Proof of Theorem 5] The proof is a simple consequence of the uniform bound (18) from Corollary 4. Indeed, for an arbitrary $j \notin T$ we obtain

$$|\hat{\theta}_j| = |\hat{\theta}_j - \theta_j^0| \leq K_4 \lambda < \delta,$$

so $j \notin \hat{T}^{th}$. Analogously, if $j \in T$, then

$$|\hat{\theta}_j| \geq |\theta_j^0| - |\hat{\theta}_j - \theta_j^0| \geq 2\delta - K_4 \lambda > \delta.$$

■

Proof [Proof of Theorem 7] First, we define a function

$$\Gamma^a(\theta) = Q(\theta) + \lambda_a \sum_{j=1}^p w_j |\theta_j|. \quad (42)$$

Next, we fix $a \in (0, 1)$ and set, for simplicity, $\xi_0 = 3$. Consider the event $\Omega = \{|\nabla Q(\theta^0)|_\infty \leq \lambda/2\}$. We know from the proof of Theorem 2 that $P(\Omega) \geq 1 - K_3 a$ and the inequality (18) is satisfied. The proof of Theorem 7 consists of two steps. In the first one we show that with high probability there exists a minimizer of the function

$$g(\theta_T) = \Gamma^a(\theta_T, 0)$$

that is close to θ_T^0 in the l_1 norm. We denote this minimizer by $\hat{\theta}_T^a$. In the second part of the proof we obtain that the vector $(\hat{\theta}_T^a, 0)$, that is $\hat{\theta}_T^a$ augmented by $(p - p_0)$ zeros, minimizes the function (42).

First, consider vectors $v \in \mathbb{R}^{p_0}$ having a fixed common l_1 -norm and a sphere

$$\{\theta_T = \theta_T^0 + p_0 \lambda v\}. \quad (43)$$

Suppose that $|v|_1$ is sufficiently large. We take arbitrary θ_T from the sphere (43) and calculate that

$$Q(\theta_T, 0) - Q(\theta^0) = \frac{1}{2} p_0^2 \lambda^2 v' \frac{1}{n} X_T' X_T v + p_0 \lambda v' [\nabla Q(\theta^0)]_T.$$

Let $\hat{\kappa}$ be the minimal eigenvalue of the matrix $\frac{1}{n} X_T' X_T$. Then we have $v' \frac{1}{n} X_T' X_T v \geq \hat{\kappa} |v|_1^2 / p_0$. Besides, on the event Ω we obtain

$$|v' [\nabla Q(\theta^0)]_T| \leq |v|_1 |\nabla Q(\theta^0)]_T|_\infty \leq \lambda |v|_1 / 2.$$

Proceeding analogously to the proof of Lemma 16 we can show that $\hat{\kappa} \geq \kappa/2$ with probability close to one. Therefore, we obtain

$$Q(\theta_T, 0) - Q(\theta^0) \geq \kappa p_0 \lambda^2 |v|_1^2 / 4 - p_0 \lambda^2 |v|_1 / 2. \quad (44)$$

Next, we work with the penalty term and obtain

$$\left| \lambda_a \sum_{j=1}^{p_0} w_j [|\theta_j^0 + p_0 \lambda v_j| - |\theta_j^0|] \right| \leq \lambda_a p_0 \lambda \sum_{j=1}^{p_0} w_j |v_j|. \quad (45)$$

Moreover, for $j \in T$ we have from Corollary 4 that

$$|\hat{\theta}_j| \geq |\theta_j^0| - |\hat{\theta}_j - \theta_j^0| \geq \theta_{min}^0 - K_4 \lambda > \lambda_a,$$

so $w_j \leq K$. Therefore, the right-hand side of (45) is bounded by $K \lambda \lambda_a p_0 |v|_1$. Combining it with (44) we get

$$g(\theta_T) - g(\theta_T^0) \geq p_0 \lambda^2 |v|_1 (\kappa |v|_1 / 4 - 1/2 - K_4 K). \quad (46)$$

The right-hand side of (46) is positive, because the norm $|v|_1$ can be taken sufficiently large, K, K_4 are constants and κ is lower bounded by a constant. Therefore, the convex function $g(\theta_T)$ takes on a sphere (43) values larger than in the center θ_T^0 . So, there exists a minimizer inside this sphere.

Next, we show that the random vector $(\hat{\theta}_T^a, 0)$ minimizes (42) with high probability, so we have to prove that the event

$$\{|\nabla_j Q(\hat{\theta}_T^a, 0)| \leq w_j \lambda_a \text{ for every } j \notin T\} \quad (47)$$

has probability close to one. By Corollary 4 we have for $j \notin T$

$$|\hat{\theta}_j| = |\hat{\theta}_j - \theta_j^0| \leq K_4 \lambda.$$

Therefore, the corresponding weight $w_j \geq \lambda_a^{-1}$. We can also calculate that

$$\nabla Q(\theta_T, 0) = \frac{1}{n} X' X_T \theta_T - \left[\frac{n-1}{n} A + \frac{1}{n^2} \sum_{i=1}^n X_i \right],$$

so we obtain the inequality

$$\left| \left[\nabla Q(\hat{\theta}_T^a, 0) \right]_{T'} \right|_{\infty} \leq \left| \frac{1}{n} X'_{T'} X_T (\hat{\theta}_T^a - \theta_T^0) \right|_{\infty} + \left| \left[\nabla Q(\theta^0) \right]_{T'} \right|_{\infty}. \quad (48)$$

Consider the event $\Omega = \{|\nabla Q(\theta^0)|_{\infty} \leq \lambda/2\}$ that has probability close to one by the proof of Theorem 2. Then the second term on the right-hand side of (48) can be bounded by $\lambda/2$. The former one can be decomposed as

$$\begin{aligned} \left| \frac{1}{n} X'_{T'} X_T (\hat{\theta}_T^a - \theta_T^0) \right|_{\infty} &\leq \left| \left(\frac{1}{n} X'_{T'} X_T - H_2' \right) (\hat{\theta}_T^a - \theta_T^0) \right|_{\infty} + \left| H_2' (\hat{\theta}_T^a - \theta_T^0) \right|_{\infty} \\ &\leq \left| \frac{1}{n} X'_{T'} X_T - H_2' \right|_{\infty} \left| \hat{\theta}_T^a - \theta_T^0 \right|_1 + |H_2'|_{\infty} \left| \hat{\theta}_T^a - \theta_T^0 \right|_1. \end{aligned} \quad (49)$$

The expression $|H_2'|_{\infty}$ is bounded by one, so from the first part of the proof we can bound, with high probability, the second term in (49) by $K_6 p_0 \lambda$. The l_{∞} -norm in the former expression can be bounded, with probability close to one, by $K_7 \sqrt{\frac{\log(p/a)}{n}}$ as in the proof of Lemma 16. Therefore, we have just proven that with probability close to one

$$\left| \left[\nabla Q(\hat{\theta}_T^a, 0) \right]_{T'} \right|_{\infty} \leq K_8 p_0 \lambda.$$

Combining it with the fact that $w_j \geq \lambda_a^{-1}$ we obtain that the event (47) has probability close to one, because from assumptions of the theorem $p_0 \lambda \leq K_5$ for K_5 small enough. \blacksquare

Appendix D. Results from Section A

We start with the proof of Lemma 9. Then we state Lemma 17 that is also needed in proofs of Theorems 8, 10 and 11. We will use the following notation

$$\Gamma(\theta) = Q(\theta) + \lambda|\theta|_1. \quad (50)$$

Proof [Proof of Lemma 9] Let $a := a_n$ be a fixed sequence such that $a \rightarrow 0$. We can calculate that for every θ

$$Q(\theta^* + a\theta) - Q(\theta^*) = -\frac{a}{n^2}\theta' \left(\sum_{i=1}^n R_i X_i \right) + a\theta' \left(\frac{X'X}{n} \right) \theta^* + \frac{a^2}{2}\theta' \left(\frac{X'X}{n} \right) \theta. \quad (51)$$

Using (9) we obtain that the right-hand side of (51) is

$$\frac{a^2}{2}\theta' \left(\frac{X'X}{n} \right) \theta - a\theta' \left[\frac{n-1}{n}A + \bar{X}/n - \left(\frac{X'X}{n} \right) \theta^* \right],$$

where A is defined in (11). Therefore, we have

$$\begin{aligned} \frac{1}{a^2} [Q(\theta^* + a\theta) - Q(\theta^*)] &= \frac{1}{2}\theta' \left(\frac{X'X}{n} \right) \theta \\ &- \frac{\theta'}{\sqrt{na}} \left[\frac{n-1}{n}\sqrt{n}A + \bar{X}/\sqrt{n} - \left(\frac{X'X}{\sqrt{n}} \right) \theta^* \right]. \end{aligned}$$

Using LLN, Lemma 17 (given below) and Slutsky's theorem we get that

$$\frac{1}{a^2} [Q(\theta^* + a\theta) - Q(\theta^*)] \rightarrow_{f-d} \frac{1}{2}\theta^T H \theta + b\theta' W, \quad (52)$$

where \rightarrow_{f-d} is the finite-dimensional convergence in distribution and $W \sim N(0, D)$. Next, we consider the penalty term and notice that

$$\frac{\lambda}{a^2} \sum_{j=1}^p (|\theta_j^* + a\theta_j| - |\theta_j^*|) \rightarrow c \sum_{j \in T} \theta_j \text{sign}(\theta_j^*) + c \sum_{j \notin T} |\theta_j|. \quad (53)$$

Thus, from (52) and (53) we have the convergence of convex functions

$$\frac{1}{a^2} [\Gamma(\theta^* + a\theta) - \Gamma(\theta^*)] \rightarrow_{f-d} V(\theta), \quad (54)$$

where the function $\Gamma(\theta)$ is defined in (50). The function on the left-hand side of (54) is minimized by $\frac{1}{a}(\hat{\theta} - \theta^*)$ and the convex function on the right-hand side of (54) has a unique minimizer. Thus $\frac{1}{a}(\hat{\theta} - \theta^*) \rightarrow_d \arg \min_{\theta} V(\theta)$, see Geyer (1996). \blacksquare

Lemma 17 *Suppose that Assumption 1 is satisfied and $\mathbb{E}|X_1|^4 < \infty$. Then*

$$\sqrt{n}[A - \mu] - \sqrt{n} \left[\frac{X'X}{n} \theta^* - \mu \right] \rightarrow_d N(0, D),$$

where D is stated precisely in the proof below.

Proof Consider two U -statistics. The first one is A that is defined in (11). The second U -statistic is

$$B = \frac{1}{n} \sum_{i=1}^n X_i X_i' \theta^*$$

and is of the order one. Besides, we have $\mathbb{E}B = H\theta^* = \mu$ by (27). Using Hoeffding (1948, Theorem 7.1) we obtain convergence in distribution in \mathbb{R}^{2p}

$$\sqrt{n} \begin{bmatrix} A - \mu \\ B - \mu \end{bmatrix} \rightarrow_d N(0, \Sigma)$$

for the matrix

$$\Sigma = \begin{pmatrix} \overbrace{\Sigma_1}^{p \times p} & \overbrace{\Sigma_2}^{p \times p} \\ \Sigma_2 & \Sigma_3 \end{pmatrix},$$

where for $j, k = 1, \dots, p$ and the function f in (24) we have

$$(\Sigma_1)_{jk} = 4Cov(\tilde{f}_j(Z_1), \tilde{f}_k(Z_1)),$$

where $\tilde{f}(z_1) = \mathbb{E}[f(Z_1, Z_2)|Z_1 = z_1]$ and $\tilde{f}_j(z_1)$ is its j -th coordinate. The entries of the matrix Σ_3 are

$$(\Sigma_3)_{jk} = Cov(X_{1j} X_1' \theta^*, X_{1k} X_1' \theta^*)$$

and

$$(\Sigma_2)_{jk} = 2Cov(\tilde{f}_j(Z_1), X_{1k} X_1' \theta^*).$$

Next, define $(p \times 2p)$ -dimensional matrix M in the following way: for $j = 1, \dots, p$ put $M_{j,j} = 1$ and $M_{j,p+j} = -1$, and zeros elsewhere. Then

$$\sqrt{n} [A - \mu] - \sqrt{n} \left[\frac{X'X}{n} \theta^* - \mu \right] = M \sqrt{n} \begin{bmatrix} A - \mu \\ B - \mu \end{bmatrix} \rightarrow_d N(0, M \Sigma M').$$

■

Now we prove main results of Section A.

Proof [Proof of Theorem 8] From Lemma 9 for $a = \lambda$ we obtain

$$\lambda^{-1} (\hat{\theta} - \theta^*) \rightarrow_d \arg \min_{\theta} V_2(\theta), \quad (55)$$

where

$$V_2(\theta) = \frac{1}{2} \theta' H \theta + \sum_{j \in T} \theta_j \text{sign}(\theta_j^*) + \sum_{j \notin T} |\theta_j|. \quad (56)$$

We start with the case (b). Let $\eta = \arg \min_{\theta} V_2(\theta)$. We know that η is nonrandom and the function $V_2(\theta)$ is strictly convex. Therefore, using (29) we have

$$\eta = (-H_1^{-1} \text{sign}(\theta_T^*), 0). \quad (57)$$

For fixed $j \in T$ we have

$$\lambda^{-1} \left(\hat{\theta}_j - \theta_j^* \right) \rightarrow_P \eta_j,$$

so $\mathbb{P}(j \notin \hat{T}) = \mathbb{P}(\hat{\theta}_j = 0) \rightarrow 0$. Thus, $\mathbb{P}(T \subset \hat{T}) \rightarrow 1$. Next, we show that $\mathbb{P}(\hat{T} \subset T) \geq 1 - \sum_{j \notin T} \mathbb{P}(j \in \hat{T}) \rightarrow 1$. Consider fixed $j \notin T$ and an event $\{j \in \hat{T}\}$. Recall that $\hat{\theta}$ minimizes

the convex function Γ , so $0 \in \partial\Gamma(\hat{\theta})$, where $\partial\Gamma$ denotes a subgradient of the convex function Γ . The function $Q(\theta)$ is differentiable, so $\partial\Gamma(\hat{\theta}) = \nabla Q(\hat{\theta}) + \lambda\partial|\hat{\theta}|$. Therefore, we have

$$0 = \nabla^j Q(\hat{\theta}) + \lambda \text{sign}(\hat{\theta}_j), \quad (58)$$

where $\nabla^j Q(\hat{\theta})$ is the j -th partial derivative $Q(\theta)$ at $\hat{\theta}$. From (58) we have

$$\lambda^{-1} \left| \nabla^j Q(\hat{\theta}) \right| = 1. \quad (59)$$

We can calculate that

$$\nabla Q(\hat{\theta}) = -\frac{n-1}{n}A - \bar{X}/n + \frac{X'X}{n} \left(\hat{\theta} - \theta^* \right) + \frac{X'X}{n}\theta^*,$$

that gives us

$$\begin{aligned} \lambda^{-1} \nabla Q(\hat{\theta}) &= -\frac{1}{\sqrt{n}\lambda} \left[\frac{n-1}{n} \sqrt{n}A + \bar{X}/\sqrt{n} - \sqrt{n} \frac{X'X}{n} \theta^* \right] \\ &+ \lambda^{-1} \frac{X'X}{n} (\hat{\theta} - \theta^*). \end{aligned} \quad (60)$$

Therefore, using LLN, Lemma 17, (55) and Slutsky's theorem the left-hand side of (59) tends in probability to $\left| (H\eta)_j \right|$. Recall that we consider the event $\{j \in \hat{T}\}$ for $j \notin T$, so we have the inequality

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left(\hat{\theta}_j \neq 0 \right) \leq \mathbb{I} \left(\left| (H\eta)_j \right| = 1 \right),$$

since η is not random. However, from (29) and (57) we obtain

$$\left| (H\eta)_j \right| = \left| (H_2^T H_1^{-1} \text{sign}(\theta_{\mathcal{A}}^*))_j \right| < 1.$$

Therefore, probability $\mathbb{P}(\hat{\theta}_j \neq 0)$ tends to zero that finishes the proof of consistency in model selection.

The proof of the claim (a) is similar to the proof of Zou (2006, Theorem 1, scenario (3)) and uses properties of the function $V_2(\theta)$ in (56). We state it here for completeness of the paper. For each $j \notin T$ we have $\lambda^{-1} \hat{\theta}_j \rightarrow_P \eta_j$, so using model selection consistency of the estimator $\hat{\theta}$ we obtain $\eta_{T^c} = 0$. Moreover, the function $V_2(\theta)$ is convex and η is its minimizer, so $0 \in \partial V_2(\eta)$. Thus, if $j \in T$ we have $0 = (H\eta)_j + \text{sign}(\theta_j^*)$. Therefore, we get

$$\eta_T = -H_1^{-1} \text{sign}(\theta_T^*), \quad (61)$$

because $\eta_{T^c} = 0$. While for $j \notin T$ we obtain that $\left| (H\eta)_j \right| \leq 1$, so component-wise

$$\left| H_2' \eta_T \right| \leq 1. \quad (62)$$

Claim (a) follows from combination of (61) and (62). \blacksquare

In the next result we use epi-convergence in distribution, which is an useful tool in the optimization theory. The following definitions are borrowed from Geyer (1994) and Pflug (1995).

Definition 1 (Epi-convergence) *A sequence f_n of lower semicontinuous (lsc) functions from \mathbb{R}^p to $\bar{\mathbb{R}}$ epi-converges to a lsc function f , if for all x and $x_n \rightarrow x$ the following inequality holds $\liminf_n f_n(x_n) \geq f(x)$ and there exists $x_n \rightarrow x$ such that $\limsup_n f_n(x_n) \leq f(x)$.*

Definition 1 is equivalent to the following one, which is useful to define epi-convergence in distribution.

Definition 2 (Epi-convergence) *A sequence f_n of lower semicontinuous (lsc) functions from \mathbb{R}^p to $\bar{\mathbb{R}}$ epi-converges to a lsc function f , if for each compact set K the following inequality holds $\liminf_n \inf_{x \in K} f_n(x) \geq \inf_{x \in K} f(x)$ and for each open set O we have $\limsup_n \inf_{x \in O} f_n(x) \leq \inf_{x \in O} f(x)$.*

Definition 3 (Epi-convergence in distribution) *Let Z_n and Z be random processes, which are lsc on \mathbb{R}^p . We say that Z_n epi-converges in distribution to Z , if for all k , all closed rectangles R_1, \dots, R_k with open interiors R_1^o, \dots, R_k^o and any real numbers a_1, \dots, a_k we have*

$$\begin{aligned} & P \left(\inf_{u \in R_1} Z(u) > a_1, \dots, \inf_{u \in R_k} Z(u) > a_k \right) \\ & \leq \liminf_n P \left(\inf_{u \in R_1} Z_n(u) > a_1, \dots, \inf_{u \in R_k} Z_n(u) > a_k \right) \\ & \leq \limsup_n P \left(\inf_{u \in R_1^o} Z_n(u) \geq a_1, \dots, \inf_{u \in R_k^o} Z_n(u) \geq a_k \right) \\ & \leq P \left(\inf_{u \in R_1^o} Z(u) \geq a_1, \dots, \inf_{u \in R_k^o} Z(u) \geq a_k \right). \end{aligned}$$

Proof [Proof of Theorem 10] We define a function

$$\Gamma^a(\theta) = Q(\theta) + \lambda \sum_{j=1}^p \frac{|\theta_j|}{|\hat{\theta}_j|}.$$

Let us start with the claim (b). Repeating the same arguments as in the proof of Lemma 9 (for $a = \frac{1}{\sqrt{n}}$) we obtain for every θ

$$nQ \left(\theta^* + \frac{\theta}{\sqrt{n}} \right) - nQ(\theta^*) \rightarrow_{f-d} \frac{1}{2} \theta' H \theta + \theta' W,$$

which using convexity implies weak convergence of the stochastic process

$$\left\{ nQ \left(\theta^* + \frac{\theta}{\sqrt{n}} \right) - nQ(\theta^*) : \theta \in K \right\} \rightarrow_d \left\{ \frac{1}{2} \theta' H \theta + \theta' W : \theta \in K \right\} \quad (63)$$

for every compact set K in \mathbb{R}^p , see Arcones (1998). Now consider the penalty term and use similar arguments to that in the proof of Zou (2006, Theorem 2) to obtain that if $\theta_j^* \neq 0$, then

$$n\lambda \left(\frac{\left| \theta_j^* + \frac{\theta_j}{\sqrt{n}} \right| - |\theta_j^*|}{|\tilde{\theta}_j|} \right) = \sqrt{n}\lambda\sqrt{n} \left(\frac{\left| \theta_j^* + \frac{\theta_j}{\sqrt{n}} \right| - |\theta_j^*|}{|\tilde{\theta}_j|} \right) \rightarrow_P c \frac{\theta_j}{\theta_j^*},$$

because $\sqrt{n}\lambda \rightarrow c$, $\tilde{\theta}_j \rightarrow_P \theta_j^*$ and $\sqrt{n} \left[\left| \theta_j^* + \frac{\theta_j}{\sqrt{n}} \right| - |\theta_j^*| \right] \rightarrow \text{sign}(\theta_j^*)\theta_j$. However, if $\theta_j^* = 0$, then

$$n\lambda \left(\frac{\left| \theta_j^* + \frac{\theta_j}{\sqrt{n}} \right| - |\theta_j^*|}{|\tilde{\theta}_j|} \right) = \sqrt{n}\lambda \frac{|\theta_j|}{|\tilde{\theta}_j|} \rightarrow_P \begin{cases} 0, & \theta_j = 0 \\ \infty, & \theta_j \neq 0, \end{cases}$$

since $\sqrt{n}\tilde{\theta}_j = O_P(1)$ and $n\lambda \rightarrow \infty$. Therefore, we obtain that for every θ

$$n\lambda \sum_{j=1}^p \frac{\left| \theta_j^* + \frac{\theta_j}{\sqrt{n}} \right| - |\theta_j^*|}{|\tilde{\theta}_j|} \rightarrow_P \begin{cases} c \sum_{j \in T} \frac{\theta_j}{\theta_j^*}, & \theta = (\theta_1, \dots, \theta_{p_0}, 0, \dots, 0) \\ \infty, & \text{otherwise.} \end{cases}$$

Since we have infinity in the last limit we cannot use arguments based on uniform convergence on compacts as we have done in the proof of Lemma 9. Here we should follow epi-convergence results (Geyer, 1994; Pflug, 1995; Zou, 2006) that combined with convergence (63) give us that

$$nQ \left(\theta^* + \frac{\theta}{\sqrt{n}} \right) - nQ(\theta^*) + n\lambda \sum_{j=1}^p \frac{\left| \theta_j^* + \frac{\theta_j}{\sqrt{n}} \right| - |\theta_j^*|}{|\tilde{\theta}_j|} \quad (64)$$

epi-converges in distribution (see Definition 3) to $V_3(\theta)$, where

$$V_3(\theta) = \begin{cases} \frac{1}{2}\theta_T' H_1 \theta_T + \theta_T' (W_T + \bar{c}), & (\theta_1, \dots, \theta_{p_0}, 0, \dots, 0) \\ \infty, & \text{otherwise} \end{cases}$$

and $W_T \sim N(0, D_1)$. Furthermore, $[-H_1^{-1}(W_T + \bar{c}), 0]'$ is the unique minimizer of the function V_3 , so epi-convergence in (64) implies convergence of minimizers (Geyer, 1994)

$$\sqrt{n} \left(\hat{\theta}_T^a - \theta_T^* \right) \rightarrow_d -H_1^{-1}(W_T + \bar{c}) \quad \text{and} \quad \sqrt{n} \left(\hat{\theta}_{T'}^a - \theta_{T'}^* \right) \rightarrow_d 0, \quad (65)$$

where $T' = \{p_0 + 1, \dots, p\}$ is the complement of T . It finishes the proof of the second claim.

Next, we go to the claim (a). We prove only that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\hat{T}^a = T \right) = 1,$$

because the equality of the signs of relevant predictors follows simply from Theorem 1 and estimation consistency stated in the claim (b) of this theorem. The reasoning is similar to the proof of Theorem 8(b). Let us start with fixed $j \in T$, then $\mathbb{P}(j \notin \hat{T}^a) = \mathbb{P}(\hat{\theta}_j^a = 0) \rightarrow 0$ by the second claim of the theorem. Next recall that $\hat{\theta}^a$ minimizes the convex function

$\Gamma^a(\theta)$, so $0 \in \partial\Gamma^a(\hat{\theta}^a)$ and $\partial\Gamma^a(\hat{\theta}^a) = \nabla Q(\hat{\theta}^a) + \partial \left(\lambda \sum_{j=1}^p \frac{|\hat{\theta}_j^a|}{|\hat{\theta}_j|} \right)$. If we consider fixed $j \notin T$ and an event $\{j \in \hat{T}^a\}$, then we have

$$0 = \nabla^j Q(\hat{\theta}^a) + \lambda \frac{\text{sign}(\hat{\theta}_j^a)}{|\hat{\theta}_j|}. \quad (66)$$

From (66) we have

$$\sqrt{n} \left| \nabla^j Q(\hat{\theta}^a) \right| = \frac{n\lambda}{\sqrt{n}|\hat{\theta}_j|}. \quad (67)$$

The right-hand side of (67) tends to infinity in probability, because $n\lambda \rightarrow \infty$ and its denominator is bounded in probability. If we can show that the left-hand side of (67) is bounded in probability, then probability of the considered event $\{j \in \hat{T}^a\}$ tends to zero and it finishes the proof. Notice that

$$\sqrt{n} \nabla Q(\hat{\theta}^a) = - \left[\frac{n-1}{n} \sqrt{n} A + \bar{X} / \sqrt{n} - \sqrt{n} \frac{X'X}{n} \theta^* \right] \quad (68)$$

$$+ \frac{X'X}{n} \sqrt{n} (\hat{\theta}^a - \theta^*). \quad (69)$$

The term on the right-hand side of (68) is bounded in probability by Lemma 17. Using LLN, (65) and Slutsky's theorem we can also bound (69) in probability. ■

Proof [Proof of Theorem 11] Using Lemma 9 with $a = 1/\sqrt{n}$ we obtain

$$\sqrt{n} (\hat{\theta} - \theta^*) \rightarrow_d -H^{-1}W, \quad (70)$$

because $\sqrt{n}\lambda \rightarrow 0$. Fix $j \notin T$, so $\theta_j^* = 0$. Then we have from (70) and $\sqrt{n}\delta \rightarrow \infty$ that $\delta^{-1}\hat{\theta}_j \rightarrow_P 0$, so $P(\hat{\theta}_j^{th} = 0) = P(|\hat{\theta}_j| < \delta) \rightarrow 1$.

Similarly, take $j \in T$ such that $\theta_j^* > 0$. From (70) we know that $\hat{\theta}_j$ is a consistent estimator of θ_j^* . Therefore, $P(\hat{\theta}_j^{th} > 0) = P(\hat{\theta}_j > \delta)$ tends to one, because $\delta \rightarrow 0$. Argumentation for $j \in T$ such that $\theta_j^* < 0$ is analogous. Using Theorem 1 we finish the proof of the claim (a) of the theorem.

From (70) we have $\sqrt{n} (\hat{\theta}_T - \theta_T^*) \rightarrow_d -(H^{-1}W)_T$. Moreover, we have just proved that $P(\hat{\theta}_T^{th} = \hat{\theta}_T) \rightarrow 1$. It finishes the proof of the claim (b). ■

Appendix E. Additional simulation results

In Figures 3 and 4 we show plots of FDR and Power for estimators considered in Section 3. In Figure 5 we have boxplots for gene expressions from Section 4.

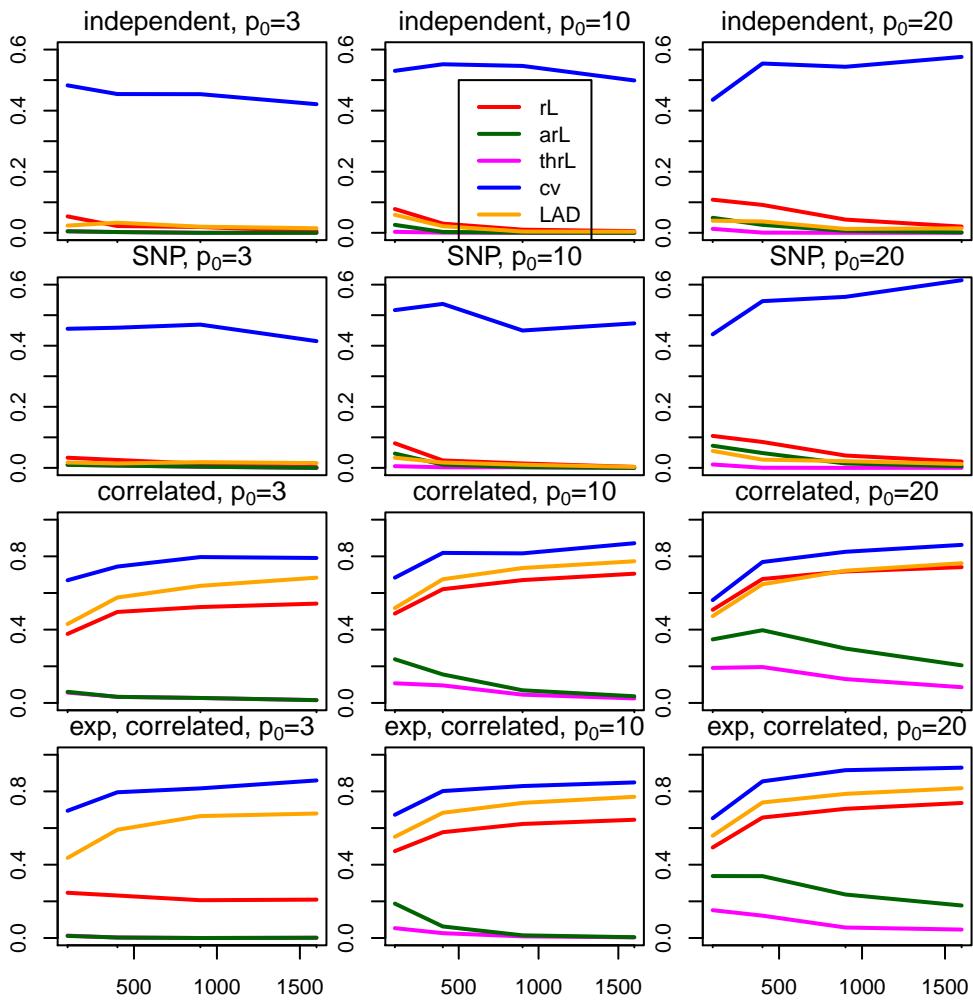


Figure 3: Plots of FDR for different simulation scenarios

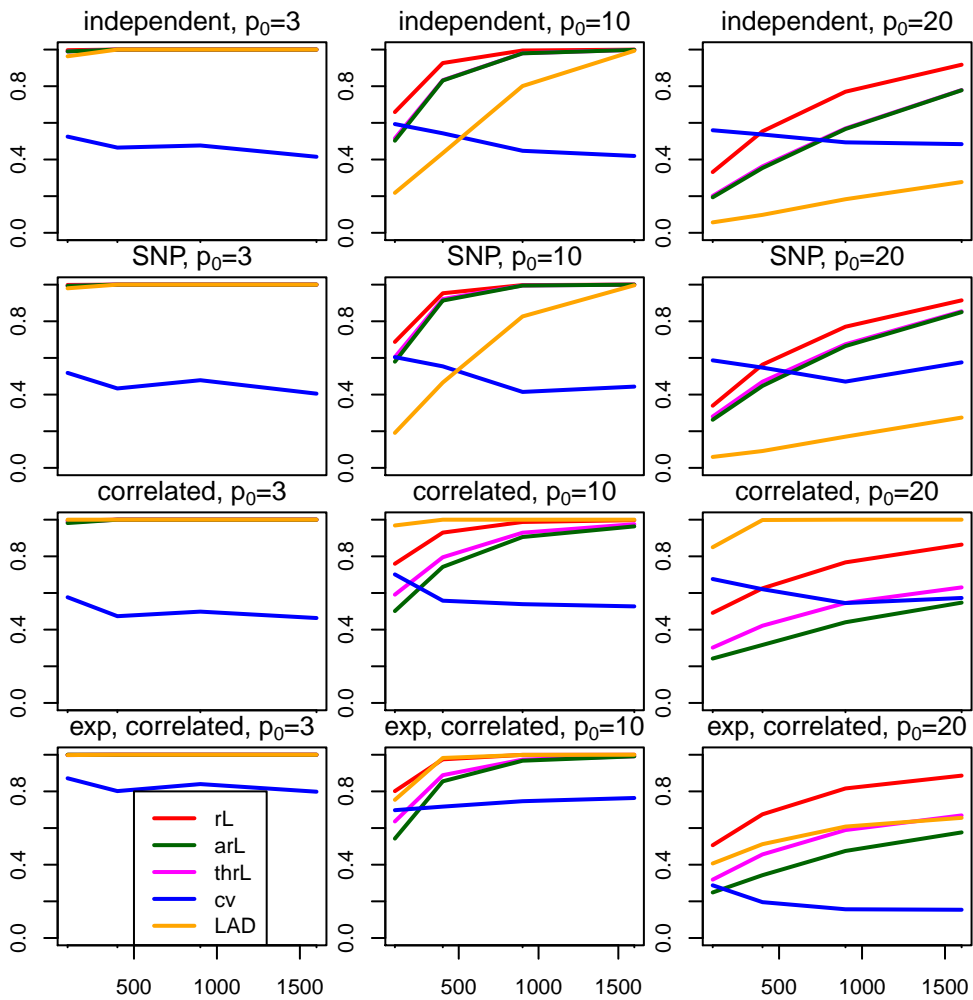



Figure 4: Plots of Power for different simulation scenarios

Figure 5: **Boxplots of gene expressions**

References

- P. Alquier and G. Biau. Sparse single-index model. *Journal of Machine Learning Research*, 14:243–280, 2013.
- M. A. Arcones. Weak convergence of convex stochastic processes. *Statistics and Probability Letters*, 37:171–182, 1998.
- M. Avella-Medina and E. Ronchetti. Robust and consistent variable selection in high-dimensional generalized linear models. *Biometrika*, 105:31–44, 2018.
- F. Balabdaoui, P. Groeneboom, and K. Hendrickx. Score estimation in the monotone single-index model. *Scandinavian Journal of Statistics*, 46:517–544, 2019.
- R. Baraniuk, M. A. Davenport, M. F. Duarte, and C. Hegde. *An Introduction to Compressive Sensing*. Connexions, Rice University, Houston, 2011.
- R. F. Barber and E. J. Candès. Controlling the false discovery rate via knockoffs. *Annals of Statistics*, 43:2055–2085, 2015.
- M. Bayati and A. Montanari. The LASSO risk for Gaussian matrices. *IEEE Transactions on Information Theory*, 58:1997–2017, 2012.
- A. Belloni and V. Chernozhukov. l_1 penalized quantile regression in high dimensional sparse models. *Annals of Statistics*, 39:82–130, 2011.
- P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics*, 37:1705–1732, 2009.

- M. Bogdan, J. K. Ghosh, and M. Żak-Szatkowska. Selecting explanatory variables with the modified version of Bayesian Information Criterion. *Quality and Reliability Engineering International*, 24:627–641, 2008.
- M. Bogdan, E. van den Berg, C. Sabatti, W. Su, and Emmanuel J. Candès. Slope adaptive variable selection via convex optimization. *Annals of Applied Statistics*, 9:1103–1140, 2015.
- J. Bradic, J. Fan, and W. Wang. Penalized composite quasilielihood for ultrahigh-dimensional variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73:325–349, 2011.
- D. R. Brillinger. A generalized linear model with Gaussian regressor variables. In *A Festschrift for Erich Lehmann*, eds. P.J. Bickel, K. Doksum, J.L. Hodges, pages 97–114. Belmont (CA):Wadsworth, 1983.
- P. Bühlmann and S. van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, New York, 2011.
- P. Bühlmann and S. van de Geer. High-dimensional inference in misspecified linear models. *Electronic Journal of Statistics*, 9:1449–1473, 2015.
- E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted l_1 minimization. *Journal of Fourier Analysis and Applications*, 14:877–905, 2008.
- E. J. Candès, Y. Fan, L. Janson, and J. Lv. Panning for gold: Model-free knockoffs for high-dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80:551–577, 2018.
- L. Cheng, P. Zeng, and Y. Zhu. Bs-sim: An effective variable selection method for high-dimensional single index model. *Electronic Journal of Statistics*, 11:3522–3548, 2017.
- A. Cohen, I. Daubechies, R. DeVore, G. Kerkycharian, and D. Picard. Capturing ridge functions in high dimensions from point queries. *Constructive Approximation*, 35:225–243, 2012.
- V. H. de la Peña and E. Giné. *Decoupling: From Dependence to Independence*. Springer-Verlag, New York, 1999.
- B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32:407–499, 2004.
- J. Fan and J. Lv. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70:849–911, 2008.
- J. Fan, Y. Fan, and E. Barut. Adaptive robust variable selection. *Annals of Statistics*, 42:324–351, 2014.
- J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33:1–22, 2010.

- X. Gao and J. Huang. Asymptotic analysis of high-dimensional LAD regression with lasso. *Statistica Sinica*, 20:1485–1506, 2010.
- M. Genzel. High-dimensional estimation of structured signals from non-linear observations with general convex loss functions. *IEEE Transactions on Information Theory*, 63:1601–1619, 2017.
- C. J. Geyer. On the asymptotics of constrained M -estimation. *Annals of Statistics*, 22:1993–2010, 1994.
- C. J. Geyer. On the asymptotics of convex stochastic optimization. *Unpublished manuscript*, 1996.
- P. Hall and K.-C. Li. On almost linearity of low dimensional projections from high dimensional data. *Annals of Statistics*, 21:867–889, 1993.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning; Data Mining, Inference and Prediction*. Springer, New York, 2001.
- W. Hoeffding. A class of statistics with asymptotically normal distribution. *Annals of Mathematical Statistics*, 19:293–325, 1948.
- J. Huang and C.-H. Zhang. Estimation and selection via absolute penalized convex minimization and its multistage adaptive applications. *Journal of Machine Learning Research*, 13:1839–1864, 2012.
- J. Huang, T. Sun, Z. Ying, Y. Yu, and C.-H. Zhang. Oracle inequalities for the lasso in the Cox model. *Annals of Statistics*, 41:1142–1165, 2013.
- P. J. Huber. Robust Estimation of a Location Parameter. *Annals of Statistics*, 53:73–101, 1964.
- W. Jiang, M. Bogdan, J. Josse, B. Miasojedow, V. Ročková, and TraumaBase Group. Adaptive Bayesian SLOPE – high-dimensional model selection with missing values. arXiv:1909.06631, 2019.
- B. A. Johnson. Rank-based estimation in the l_1 regularized partly linear model for censored outcomes with application to integrated analyses of clinical predictors and gene expression data. *Biostatistics*, 10:659–666, 2009.
- B. A. Johnson and L. Peng. Rank-based variable selection. *Journal of Nonparametric Statistics*, 20:241–252, 2008.
- K. Knight and W. Fu. Asymptotics for lasso-type estimators. *Annals of Statistics*, 28:1356–1378, 2000.
- E. Kong and Y. Xia. Variable selection for the single-index model. *Biometrika*, 94:217–229, 2007.
- M. Kos and M. Bogdan. On the asymptotic properties of slope. *Sankhya A*, 82:499–532, 2020.

- M. Kubkowski. Misspecification of binary regression model: properties and inferential procedures. *Unpublished PhD dissertation*, 2019.
- M. Kubkowski and J. Mielniczuk. Active sets of predictors for misspecified logistic regression. *Statistics*, 51:1023–1045, 2017.
- M. Kubkowski and J. Mielniczuk. Projections of a general binary model on a logistic regression. *Linear Algebra and its Applications*, 536:152–173, 2018.
- K.-C. Li and N. Duan. Regression analysis under link violation. *Annals of Statistics*, 17:1009–1052, 1989.
- W. Lu, Y. Goldberg, and J. P. Fine. On the robustness of the adaptive lasso to model misspecification. *Biometrika*, 99:717–731, 2012.
- J. Lv and J. S. Liu. Model selection principles in misspecified models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76:141–167, 2014.
- N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34:1436–1462, 2006.
- S. Na, Z. Yang, Z. Wang, and M. Kolar. High-dimensional varying index coefficient models via Stein identity. *Journal of Machine Learning Research*, 20:1–44, 2019.
- B. Peng and L. Wang. An iterative coordinate descent algorithm for high-dimensional non-convex penalized quantile regression. *Journal of Computational and Graphical Statistics*, 24:676–694, 2015.
- G. Ch. Pflug. Asymptotic stochastic programs. *Mathematics of Operations Research*, 20:769–789, 1995.
- Y. Plan and R. Vershynin. The generalized lasso with non-linear observations. *IEEE Transactions on Information Theory*, 62:1528–1537, 2016.
- P. Pokarowski and J. Mielniczuk. Combined l_1 and greedy l_0 penalized least squares for linear model selection. *Journal of Machine Learning Research*, 16:961–992, 2015.
- Y. Qin, S. Li, Y. Li, and Y. Yu. Penalized Maximum Tangent Likelihood Estimation and Robust Variable Selection. *arXiv:1708.05439*, 2017.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2017. URL <http://www.R-project.org/>.
- G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- W. Rejchel. Oracle inequalities for ranking and U -processes with lasso penalty. *Neurocomputing*, 239:214–222, 2017a.

- W. Rejchel. Model selection consistency of U -statistics with convex loss and weighted lasso penalty. *Journal of Nonparametric Statistics*, 29:768–791, 2017b.
- V. Ročková and E. I. George. The Spike-and-Slab LASSO. *Journal of the American Statistical Association*, 113:431–444, 2018.
- P. A. Ruud. Sufficient conditions for the consistency of maximum likelihood estimation despite misspecification of distribution in multinomial discrete choice models. *Econometrica*, 51:225–228, 1983.
- R. J. Serfling. *Approximation theorems of mathematical statistics*. Wiley, New York, 1980.
- X. Song and S. Ma. Penalized variable selection with U -estimates. *Journal of Nonparametric Statistics*, 4:499–515, 2010.
- C. Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*, pages 583–602, Berkeley, Calif., 1972. University of California Press.
- C. Stein, P. Diaconis, S. Holmes, and G. Reinert. Use of exchangeable pairs in the analysis of simulations. In *Stein Method, Institute of Mathematical Statistics*, 2004.
- B. Stranger, S. Forrest, M. Dunning, C. Ingle, C. Beazley, N. Thorne, R. Redon, C. Bird, A. de Grassi, C. Lee, C. Tyler-Smith, N. Carter, S. Scherer, S. Tavar, P. Deloukas, M. Hurles, and E. Dermitzakis. Relative impact of Nucleotide and Copy Number Variation on Gene Expression Phenotypes. *Science*, 315:848–853, 2007.
- W. J. Su, M. Bogdan, and E. J. Candès. False discoveries occur early on the lasso path. *The Annals of Statistics*, 45:2133–2150, 2017.
- P. J. C. Tardivel and M. Bogdan. On the sign recovery by LASSO, thresholded LASSO and thresholded Basis Pursuit Denoising. *arXiv:1812.05723*, 2018.
- The International HapMap Consortium. A haplotype map of the human genome. *Nature*, 437:1299–1320, 2005.
- H. Thorisson. Coupling methods in probability theory. *Scandinavian Journal of Statistics*, 22:159–182, 1995.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1996.
- S. van de Geer. High-dimensional generalized linear models and the lasso. *Annals of Statistics*, 36:614–645, 2008.
- S. van de Geer. *Estimation and Testing under Sparsity*. Springer, 2016.
- S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.

- H. Wang, G. Li, and G. Jiang. Robust regression shrinkage and consistent variable selection through the LAD-Lasso. *Journal of Business & Economic Statistics*, 25:347–355, 2007.
- L. Wang. The L_1 penalized lad estimator for high dimensional linear regression. *Journal of Multivariate Analysis*, 120:135 – 151, 2013.
- L. Wang, Y. Wu, and R. Li. Quantile regression for analyzing heterogeneity in ultra-high dimension. *Journal of the American Statistical Association*, 107:214–222, 2012.
- S. Wang, H. Weng, and A. Maleki. Which bridge estimator is optimal for variable selection? *arXiv preprint arXiv:1705.08617*, 2017.
- T. Wang and L. Zhu. A distribution-based lasso for a general single-index model. *Science China Mathematics*, 58:109–130, 2015.
- X. Wei, Z. Yang, and Z. Wang. On the statistical rate of nonlinear recovery in generative models with heavy-tailed data. *Proceedings of Machine Learning Research*, 97:6697–6706, 2019.
- A. Weinstein, M. Bogdan, W. Su, R.F. Barber, and E. J. Candès. A power analysis for knockoffs with the lasso coefficient-difference statistic. *arXiv:2007.15346*, 2020.
- C. Wu and S. Ma. A selective review of robust variable selection with applications in bioinformatics. *Briefings in Bioinformatics*, 16:873–883, 2015.
- Z. Yang, K. Balasubramanian, and H. Liu. High-dimensional non-Gaussian single index models via thresholded score function estimation. *Proceedings of Machine Learning Research*, 70:3851–3860, 2017a.
- Z. Yang, K. Balasubramanian, Zhaoran Wang, and H. Liu. Estimating high-dimensional non-gaussian multiple index models via Stein lemma. *Advances in Neural Information Processing Systems 30*, pages 6097–6106, 2017b.
- F. Ye and C. H. Zhang. Rate minimaxity of the lasso and Dantzig selector for the l_q loss in l_r balls. *Journal of Machine Learning Research*, 11:3519–3540, 2010.
- M. Žak, A. Baierl, M. Bogdan, and A. Futschik. Locating multiple interacting quantitative trait loci using rank-based model selection. *Genetics*, 176:1845–1854, 2007.
- P. Zeng, T. He, and Y. Zhu. A lasso-type approach for estimation and variable selection in single index models. *Journal of Computational and Graphical Statistics*, 21:92–109, 2012.
- P. Zhao and B. Yu. On model selection consistency of lasso. *Journal of Machine Learning Research*, 7:2541–2563, 2006.
- W. Zhong, L. Zhu, R. Li, and H. Cui. Regularized quantile regression and robust feature screening for single index models. *Statistica Sinica*, 26:69–95, 2016.
- S. Zhou. Thresholding procedures for high dimensional variable selection and statistical estimation. *Advances in Neural Information Processing Systems*, 22:1436–1462, 2009.

- L.-P. Zhu and L.-X. Zhu. Nonconcave penalized inverse regression in single-index models with high dimensional predictors. *Journal of Multivariate Analysis*, 100:862 – 875, 2009.
- H. Zou. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101:1418–1429, 2006.