



HAL
open science

Group Lasso for generalized linear models in high dimension

Mélanie Blazère, Jean-Michel Loubes, Fabrice Gamboa

► **To cite this version:**

Mélanie Blazère, Jean-Michel Loubes, Fabrice Gamboa. Group Lasso for generalized linear models in high dimension. 2014. hal-00850689v3

HAL Id: hal-00850689

<https://hal.science/hal-00850689v3>

Preprint submitted on 22 Jan 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Oracle inequalities for a Group Lasso procedure applied to generalized linear models in high dimension

Mélanie Blazère, Jean-Michel Loubes and Fabrice Gamboa

Abstract

We present a Group Lasso procedure for generalized linear models (GLMs) and we study the properties of this estimator applied to sparse high-dimensional GLMs. Under general conditions on the covariates and on the joint distribution of the pair covariates, we provide oracle inequalities promoting group sparsity of the covariables. We get convergence rates for the prediction and estimation error and we show the ability of this estimator to recover good sparse approximation of the true model. Then we extend this procedure to the case of an Elastic net penalty. At last we apply these results to the so-called Poisson regression model (the output is modeled as a Poisson process whose intensity relies on a linear combination of the covariables). The Group Lasso method enables to select few groups of meaningful variables among the set of inputs.

Keywords: Generalized linear model, high dimension, sparse model, groups of variables, Group Lasso, oracle inequalities.

1 Introduction

Handling high dimensional data is nowadays required for many practical applications ranging from astronomy, economics, industrial problems to biomedicine. Being able to extract information from these large data sets has been at the heart of statistical studies over the last decades and many papers have extensively studied this setting in a lot of fields ranging from statistical inference to machine learning. We refer for instance to references therein.

For high-dimensional data, classical methods based on a direct minimization of the empirical risk can lead to over fitting. Actually adding a complexity penalty enables to avoid it by selecting fewer coefficients. Using an ℓ_0 penalty leads to sparse solutions but the usually non-convex minimization problem turns out to be extremely difficult to handle when the number of parameters becomes large. Hence the ℓ_1 type penalty has been introduced to overcome this issue. On the one hand this penalty achieves sparsity of an estimated parameter vector and on the other hand it requires only convex optimization type calculations which are computationally feasible even for high dimensional data. The use of a ℓ_1 type penalty, first proposed in [25] by Tibshirani, is now a well established procedure which has been studied in a large variety of models. We refer for example to [2], [5], [26], [8], [24], [32] and [4].

Group sparsity can be promoted by imposing a ℓ_2 penalty to individual groups of variables and then a ℓ_1 penalty to the resulting block norms. Yuan and Lin [31] proposed an extension of the Lasso in the case of linear regression and presented an algorithm when the model matrices in each group are orthonormal. This extension, called the Group Lasso, encourages blocks sparsity. Wei and Huang [29] studied the properties of the Group Lasso for linear regression, Nardi and Rinaldo [18] established asymptotic properties and Lounici, Pontil, van de Geer and Tsybakov [14] stated oracle inequalities in linear Gaussian noise under group sparsity. Meir, van de Geer and Bühlmann [17] considered the Group Lasso in the case of logistic regression and Zhang and Huang [32] studied the benefit of group sparsity. Another important reference is the work of Negahban, Ravikumar, Wainwright and Yu [20]. In this last paper a unified framework for the study of rates of convergence in high dimensional setting is provided under two key assumptions (restricted strong convexity and decomposability). This work and these two assumptions will be discussed in more details in Section 3.

In this paper we focus on the Group Lasso penalty to select and estimate parameters in the generalized linear model. One of the application is the Poisson regression model. More precisely, we consider the generalized linear model introduced by McCullagh and Nelder [16]. Let F be a distribution on \mathbb{R} and let (X, Y) be a pair of random variables with $X \in \mathbb{R}^p$ and $Y \in \mathbb{R}$. The conditional law of $Y|X = x$ is modeled by a distribution from the exponential family and the canonical parameter is the linear predictor. Thus the conditional distribution of the observations given $X = x$ is $P(Y|\beta^*, x) = \exp(y\beta^{*T}x - \psi(\beta^{*T}x))$, where $\beta^{*T}x$ satisfies $\int \exp(y\beta^{*T}x)F(dy) < \infty$ and ψ is the normalized function. Notice that $\mathbb{E}(Y|X) \stackrel{\text{a.s.}}{=} \psi'(\beta^{*T}X)$ in other words $\beta^{*T}X \stackrel{\text{a.s.}}{=} h(\mathbb{E}(Y|X))$ where $h = \psi'^{-1}$ is the so-called link function. Some common examples of generalized linear models are the Poisson regression for count data, logistic and probit regression for binary data or multinomial regression for categorical data. The quantities of interest that we would like to estimate are the component $(\beta_j^*)_{1 \leq j \leq p}$ of β^* and for a given x we may also wish to predict the response $Y|X = x$. A natural field of applications for such models is given by genomics and Poisson regression type models. In particular thousands of variables such as expressions of genes and bacterias can be measured for each animal (mices) in a (pre-)clinical study thanks to the developpement of micro arrays (see, for example, [3] and [7]). A typical goal is to classify their health status, e.g. healthy or diseased, based on their bio-molecular profile, i.e. the thousands of bio-molecular variables measured for each individual (see for instance [9], [22] and [21]).

The paper falls into the following parts. In Section 2 we describe the model and the Group Lasso estimator for generalized linear models. In Section 3 we present the main results on coefficients estimation and prediction error and in Section 4 we consider the model in the particular case of a Poisson regression. We also study here the general model in the case of a mixture of an ℓ_1 and ℓ_2 penalty. The Group Lasso estimator is then used on simulated data sets in Section 5 and its performances are compared to those of the Lasso estimator. The Appendix is devoted to the proof of the main theorems.

2 Sparse Variable selection for generalized linear models

2.1 The model

The Exponential family on the real line is a unified family of distributions parametrized by a one dimensional parameter and is widely used for practical modelling. Let F be a probability distribution on \mathbb{R} not concentrated on a point and

$$\Theta := \left\{ \theta \in \mathbb{R} : \int \exp(\theta x)F(dx) < \infty \right\}.$$

Define

$$M(\theta) := \int \exp(\theta x)F(dx) \quad (\theta \in \Theta)$$

and

$$\psi(\theta) := \log(M(\theta)) \quad (\theta \in \Theta).$$

Let $P(y; \theta) = \exp(\theta y - \psi(\theta))$ with $\theta \in \Theta$. The densities of probability (related to a mesure adapted to the continuous or discrete case) $\{P(\cdot; \theta) : \theta \in \Theta\}$ is called the exponential family. Θ is the natural parameter space and θ is called the canonical parameter. The exponential family includes most of the commonly used distributions like normal, gamma, poisson or binomial distributions [16].

We consider a pair of random variables (X, Y) where $Y \in \mathbb{R}$ and $X \in \mathbb{R}^p$ such that the conditional distribution $Y|X = x$ is $P(Y|\beta^*, x) = \exp(y\beta^{*T}x - \psi(\beta^{*T}x))$, with $\beta^{*T}x \in \Theta$. Our aim is to estimate the components $(\beta_j^*)_{1 \leq j \leq p}$ of β^* in order to predict the response $Y|X = x$ conditionally on a given value of x . We assume that

- (H.1): the variable X is almost surely bounded by a constant L i.e. there exists a constant $L > 0$ such that $\|X\|_\infty \leq L$ a.s.
- (H.2): for all $x \in [-L, L]^p$, $\beta^{*T}x \in \text{Int}(\Theta)$

and we consider

$$\Lambda = \{\beta \in \mathbb{R}^p : \forall x \in [-L, L]^p, \beta^T x \in \Theta\}.$$

Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be i.i.d copies of (X, Y) where $X_i = (X_{i,1}, \dots, X_{i,p})^T$ for $i = 1, \dots, n$. We consider the case of high-dimensional regression i.e $p \gg n$. The log-likelihood for a generalized linear model is given by

$$\mathcal{L}(\beta) = \sum_{i=1}^n [Y_i \beta^T X_i - \psi(\beta^T X_i)].$$

We denote the generalized linear model loss function by

$$l(\beta) =: l(\beta; x, y) =: -y f_\beta(x) + \psi(f_\beta(x)), \quad (1)$$

where $f_\beta(x) := \beta^T x$. Notice that this function is convex in β (as ψ is convex). The associated risk is defined by $\mathbb{P}l(\beta) =: \mathbb{E}l(\beta; Y, X)$ and the empirical risk by $\mathbb{P}_n l(\beta)$ where

$$\mathbb{P}_n l(\beta) := \frac{1}{n} \sum_{i=1}^n [-Y_i \beta^T X_i + \psi(\beta^T X_i)].$$

Obviously

$$\beta^* = \operatorname{argmin}_{\beta \in \Lambda} \mathbb{P}l(\beta).$$

2.2 Group Lasso for generalized linear model

Assume that X is structured into G_n groups each of size d_g for $g \in \{1, \dots, G_n\}$. For $i = 1, \dots, n$ we set

$$X_i = (X_i^1, \dots, X_i^g, \dots, X_i^{G_n})^T$$

where

$$X_i^g = (X_{i,1}^g, \dots, X_{i,d_g}^g)^T$$

and $\sum_{g=1}^{G_n} d_g = p$. This decomposition is often natural in biology and micro arrays data when the covariates are genes expression (see for instance [23] and [30]). We allow the number of groups to increase with the sample size n , so we can consider the case where $G_n \gg n$. Define $d_{\max} := \max_{g \in \{1, \dots, G_n\}} d_g$ and $d_{\min} := \min_{g \in \{1, \dots, G_n\}} d_g$. For $\beta \in \mathbb{R}^p$ we denote by β^g the sub-vector of β whose indexes correspond to the index set of the g^{th} group of X .

Let us consider the Group Lasso estimator which achieves group sparsity and is obtained as the solution of the convex optimization problem

$$\hat{\beta}_n = \operatorname{argmin}_{\beta \in \Lambda} \left\{ \mathbb{P}_n l(\beta) + \sum_{g=1}^{G_n} s(d_g) \|\beta^g\|_2 \right\}$$

where r_n is the tuning parameter.

Here $\|\cdot\|_2$ refers to the Euclidian norm and s is a given function. An increase in r_n leads to a diminution of the β^g to zero, this means that some blocks become simultaneously zero and groups of predictors drop out of the model. Typically we choose $s(d_g) := \sqrt{d_g}$ to penalize more heavily groups of large size. Notice that if all the groups are of size one then we recover the Lasso estimator. The Group Lasso achieves variables selection and estimation simultaneously as the Lasso does. The penalty function is the sum of the ℓ_2 norm of the groups of variables. Thus the Group Lasso estimator acts like the Lasso at the group level [18], [14], [31]. Actually the objective function above is the sum of a particular loss function (which is convex and based on the exponential family) with a weighted regularizer. This type of convex optimization problem is referred as a regularized M-estimator in the paper of Negahban et al. [20]. We can also notice that the penalty norm satisfies the decomposability condition defined in this last paper.

We study estimation and prediction properties of the Group Lasso in high dimensional settings when the number of groups exceeds the sample size i.e. $G_n \gg n$. Define $H^* = \{g : \beta^{*g} \neq 0\}$ the index set of the groups for which the corresponding sub vectors of β^* are non-zero and $m^* := |H^*|$.

Such a set characterizes the sparsity of the model. In the following H^{*c} denotes the index set of the groups which are not in H^* . We can notice that H^* and m^* depend on n but for simplicity we do not specify this dependency. In general, it will be hopeless to estimate all unknown parameters from data except if we make the assumption that the true parameter is group sparse. In this paper we consider that β^* is partitioned into a number of groups, in correspondance with the partition of X , only few of which are relevant. The index of group sparsity m^* will be discussed more deeply after Theorem 3.7. We also assume (H.3): there exists a constant $B > 0$ such that $\sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2 \leq B$.

3 Main Results

3.1 Bounds for estimation and prediction error

Under some assumptions on the value of the parameter r_n and on the covariance matrix of X we are going to show the ability of this estimator to recover good sparse approximation of the true model. To prove oracle inequalities for the Group Lasso applied to generalized linear model we need to state concentration inequalities for the empirical process $\mathbb{P}_n(l(\beta))$ for $\beta \in \Lambda$. This step is essential to compute an appropriate lower bound for the regularization parameter that ensures good statistical properties of the estimator with high probability. Notice that this step requires a boundedness assumption on the $(X_{i,j})$ (cf. proof of Proposition 3.1).

To state concentration inequalities we first break down the empirical process into a linear part and a part which depends on the normalized parameter ψ

$$(\mathbb{P}_n - \mathbb{P})(l(\beta)) = (\mathbb{P}_n - \mathbb{P})(l_l(\beta)) + (\mathbb{P}_n - \mathbb{P})(l_\psi(\beta))$$

where $l_l(\beta) := l_l(\beta, x, y) = -y\beta^T x$ and $l_\psi(\beta) := l_\psi(\beta, x) = \psi(\beta^T x)$. Define

$$\mathcal{A} = \bigcap_{g=1}^{G_n} \{L_g \leq r_n/2\}$$

where

$$L_g := \left\| \frac{1}{\sqrt{d_g n}} \sum_{i=1}^n (Y_i X_i^g - \mathbb{E}(Y X^g)) \right\|_2$$

for all $g \in \{1, \dots, G_n\}$ and

$$\mathcal{B} = \left\{ \sup_{\beta: \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq M} |\nu_n(\beta, \beta^*)| \leq \frac{r_n}{2} \right\}$$

where

$$\nu_n(\beta, \beta^*) := \frac{(\mathbb{P}_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta))}{\sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 + \varepsilon_n}$$

with $M = 8B + \varepsilon_n$ and $\varepsilon_n = \frac{1}{n}$. We assume that G_n and n are such that $\frac{\log(2G_n)}{n} \leq 1$. The next proposition shows that the event $\mathcal{A} \cap \mathcal{B}$ occurs with high probability for some suitable values of the tuning parameter.

Proposition 3.1. *Let*

$$r_n \geq AKL \left\{ C_{L,B} \vee \max_{\{|x| \leq L\kappa_n\} \cap \Theta} |\psi'(x)| \right\} \sqrt{\frac{2 \log(2G_n)}{n}}$$

where K is a universal constant, $A > \sqrt{2}$, $\kappa_n := 17B + \frac{2}{n}$ and $C_{L,B}$ is defined in Lemma 3.2. We have

$$\mathbb{P}(\mathcal{A} \cap \mathcal{B}) \geq 1 - (C + 2d_{max})(2G_n)^{-A^2/2}.$$

where C is a universal constant.

The proof rests on concentration inequalities and is detailed in the appendix. Indeed inferring a bound for the probability of the events \mathcal{A} and \mathcal{B} is equivalent to prove concentration inequalities for the linear and non linear part of the empirical process. A concentration inequality for the linear part is derived from Bernstein inequality once the following lemma has been proved. This lemma provides moment bounds for Y .

Lemma 3.2. *Let (X, Y) a pair of random variables whose conditional distribution is $P(Y; \beta^* | X = x) = \exp(y\beta^{*T}x - \psi(\beta^{*T}x))$ and assume assumptions (H.1-3) are fulfilled. For all $k \in \mathbb{N}^*$ there exists a constant $C_{L,B}$ (which depends only on L and B) such that $\mathbb{E}(|Y|^k) \leq k!(C_{L,B})^k$.*

The boundedness assumption on the components of X is required to prove this lemma. Then, for the non linear part of the empirical process, we use again this assumption to show that we can restrict the study of ψ to a suitable compact set. Since ψ is lipchitzian on this compact set, concentration results for lipchitzian loss functions (see [12]) allow to bound the probability of the event \mathcal{B} .

Thus, on the event \mathcal{A} which occurs with high probability (see Proposition 3.1), we have an upper bound for the linear part of the empirical process $(\mathbb{P}_n - \mathbb{P}) \left(l_l(\beta^*) - l_l(\hat{\beta}_n) \right)$.

Proposition 3.3. *On the event \mathcal{A}*

$$(\mathbb{P}_n - \mathbb{P}) \left(l_l(\beta^*) - l_l(\hat{\beta}_n) \right) \leq \frac{r_n}{2} \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2.$$

Proof. We have

$$\begin{aligned} & (\mathbb{P}_n - \mathbb{P}) \left(l_l(\beta^*) - l_l(\hat{\beta}_n) \right) \\ &= \sum_{g=1}^{G_n} (\hat{\beta}_n^g - \beta^{*g})^T \left[\frac{1}{n} \sum_{i=1}^n Y_i X_i^g - \mathbb{E}(Y X^g) \right] \\ &\leq \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \left\| \frac{1}{\sqrt{d_g n}} \sum_{i=1}^n (Y_i X_i^g - \mathbb{E}(Y X^g)) \right\|_2. \end{aligned}$$

The last bound is obtained by using Cauchy-Schwarz inequality. Then on the event \mathcal{A} the proposition follows. \square

Therefore the difference between the linear part of the empirical process and its expectation is bounded from above by the tuning parameter multiplied by the norm (associated to the Group Lasso penalty) of the difference between the estimated parameter and the true parameter β^* . We can state a similar result for the non linear part of the empirical process, the key of the proof is based on the fact that the estimator $\hat{\beta}_n$ is in a neighborhood of the target parameter β^* on the event $\mathcal{A} \cap \mathcal{B}$.

Lemma 3.4. *On the event $\mathcal{A} \cap \mathcal{B}$ we have $\sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \leq M$, where we recall that $M = 8B + \varepsilon_n$ and $\varepsilon_n = \frac{1}{n}$.*

Then the next proposition provides an upper bound for $(\mathbb{P}_n - \mathbb{P}) \left(l_\psi(\beta^*) - l_\psi(\hat{\beta}_n) \right)$ and is directly involved by the definition of \mathcal{B} and Lemma 3.4.

Proposition 3.5. *On the event $\mathcal{A} \cap \mathcal{B}$*

$$\begin{aligned} & (\mathbb{P}_n - \mathbb{P}) \left(l_\psi(\beta^*) - l_\psi(\hat{\beta}_n) \right) \\ &\leq \frac{r_n}{2} \left(\sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + \varepsilon_n \right). \end{aligned}$$

Once concentration of the loss function around its mean is stated, we have to ensure that the loss function is not too flat in such a way that if the loss difference $l(\hat{\beta}_n) - l(\beta^*)$ converges to zero then $\hat{\beta}_n$ converges to β^* . In this paper such a property holds assuming that the covariance matrix satisfies a Group Stabil condition (see condition below). This condition is closely related to the one of Negahban et al. [20] called Restricted Strong Convexity property. Notice that in our analysis the boundedness of the covariates is required to prove such a property. In fact, thanks to the boundedness of the covariates, we first bound from below the mean deviation of the loss function by a quadratic function (see Proposition A.1 in Appendix A). This first step enables to relate the deviation in the loss to the deviation of the estimated parameter from the true one. Then the Group Stabil condition implies that the loss function satisfies a local strong convexity property. This has been characterized as a common step for convergence of M -estimator (see [20]). Notice that the boundedness assumption on the components of X is not required to obtain a kind of strong convexity. For example, as stated by Negahban et al., if the covariates have sub-gaussian tails and if the covariance matrix is positive definite then the loss function satisfies a form of restricted strong convexity property with high probability. However, as noticed above the boundedness assumption is necessary to establish Proposition 3.1.

Therefore the key condition to derive oracle inequalities rests on the correlation between the covariates i.e. on the behavior of the Gram matrix $\frac{1}{n} \sum_{i=1}^n X_i X_i^T$ which is necessarily singular when $p > n$. Meier, van de Geer and Bühlmann [17] proved that the group Lasso is consistent in the particular case of logistic regression and gave bounds for the prediction error under the assumption that $\mathbb{E}(XX^T)$ is non singular. In this paper we give sharp bounds for estimation and prediction errors for generalized linear models using a weaker condition similar to the restricted eigenvalue condition (RE) of Bickel, Ritov and Tsybakov [2]. This condition is quite weaker than the one of Bunea, Tsybakov and Wegkamp [6]. In their article Bickel, Ritov and Tsybakov also give several sufficient conditions for RE (which are easier to check). Here we use a condition which is a group version of the Stabil Condition first introduced by Bunea [5] for logistic regression in the case of an ℓ_1 penalty. This condition is similar (within a constant ε) to the condition used by Lounici, Pontil, van de Geer and Tsybakov [14] to state oracle inequalities for linear regression. For $c_0, \varepsilon > 0$, we define the restricted set as

$$S(c_0, \varepsilon) = \left\{ \delta : \sum_{g \in H^{*c}} \sqrt{d_g} \|\delta^g\|_2 \leq c_0 \sum_{g \in H^*} \sqrt{d_g} \|\delta^g\|_2 + \varepsilon \right\}.$$

On this set we assume that the covariance matrix satisfies the Group Stabil condition defined below. This condition ensures local strong convexity in a neighborhood of β^* .

Let $\Sigma := \mathbb{E}(XX^T)$ be the $p \times p$ covariance matrix.

Definition : Group Stabil Condition

Let $c_0, \varepsilon > 0$ be given. Σ satisfies the Group Stabil condition $GS(c_0, \varepsilon, k)$ if there exists $0 < k < 1$ such that

$$\delta^T \Sigma \delta \geq k \sum_{g \in H^*} \|\delta^g\|_2^2 - \varepsilon$$

for any $\delta \in S(c_0, \varepsilon)$.

Before going any further we have to define two norms that are used to control the estimation error. For all $z \in \mathbb{R}^p$, let

$$\|z\|_R := \sum_{g=1}^{G_n} \sqrt{d_g} \|z^g\|_2$$

and

$$\|z\|_{2,1} := \sum_{g=1}^{G_n} \|z^g\|_2.$$

The norm $\|\cdot\|_R$ is called the regularizer norm.

We are now able to state the main result of this paper which provides meaningful bounds for the estimation and prediction error when the true model is sparse and $\log(G_n)$ is small as compared to n .

Let $\gamma^* := \sum_{g \in H^*} d_g$. We recall that $\kappa_n := 17B + \frac{2}{n}$.

Theorem 3.6. *Assume condition $GS(3, \frac{1}{2n}, k)$ is fulfilled. Let*

$$r_n \geq AKL \left\{ C_{L,B} \vee \max_{\{|x| \leq L\kappa_n\} \cap \Theta} |\psi'(x)| \right\} \sqrt{\frac{\log(2G_n)}{n}}$$

where $A > \sqrt{2}$, K is a universal constant and $C_{L,B}$ is defined in Lemma 3.2. Then, with probability at least $1 - (C + 2d_{max})(2G_n)^{-A^2/2}$ (where C is given in Proposition 3.1), we have

$$\|\hat{\beta}_n - \beta^*\|_R \leq \frac{4}{c_n k} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n},$$

$$\|\hat{\beta}_n - \beta^*\|_{2,1} \leq \frac{4}{c_n k \sqrt{d_{min}}} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n \sqrt{d_{min}}}$$

and

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{16}{c_n^2 k} r_n^2 \gamma^* + \frac{2r_n + 1}{2c_n n}.$$

where

$$c_n := \min_{\{|x| \leq L(9B + \frac{1}{n})\} \cap \Theta} \left\{ \frac{\psi''(x)}{2} \right\}.$$

Notice that $c_n > 0$ since the measure associated to the distribution F is not concentrated on a point.

These results are similar to those of Nardi and Rinaldino [18] who proved asymptotic properties of the Group Lasso estimator for linear models. We can notice that if $\gamma^* = O(1)$ then the bound on the estimation error is of the order $O\left(\sqrt{\frac{\log G_n}{n}}\right)$ and the Group Lasso estimator still remains consistent for the $\ell_{2,1}$ -estimation error and for the ℓ_2 -prediction error under the Group Stabil condition if the number of groups increases almost as fast as $O(\exp(n))$. The term $\sqrt{\log G_n}$ is the price to pay for having a large number of factors and not knowing where are the non zero ones.

Since $\|\hat{\beta}_n - \beta^*\|_2 \leq \|\hat{\beta}_n - \beta^*\|_{2,1}$, we also have

$$\|\hat{\beta}_n - \beta^*\|_2 \leq \frac{4}{c_n k \sqrt{d_{min}}} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n \sqrt{d_{min}}}.$$

However, a sharper bound for the ℓ_2 -norm of the estimation error holds but under a stronger assumption than $GS(3, \frac{1}{2n}, k)$.

Theorem 3.7. *Assume $GS(2m^*, 3, \frac{1}{2n}, k')$ i.e there exists $0 < k' < 1$ such that*

$$\delta^T \Sigma \delta \geq k' \sum_{g \in J} \|\delta^g\|_2^2 - \frac{1}{2n}$$

for any δ such that $\sum_{g \in J^c} \sqrt{d_g} \|\delta^g\|_2 \leq 3 \sum_{g \in J} \sqrt{d_g} \|\delta^g\|_2 + \frac{1}{2n}$ and J such that $|J| \leq 2m^*$. Then we have with probability at least $1 - (C + 2d_{max})(2G_n)^{-A^2/2}$

$$\|\hat{\beta}_n - \beta^*\|_2^2 \leq 10 \frac{d_{max}}{d_{min}} \left\{ \frac{16}{k'^2 c_n^2} r_n^2 \gamma^* + \frac{2r_n + 1}{2k' c_n n} + \frac{1}{2k' n} \right\}.$$

Proof. Let $J^* = H^* \cup I$ and I is the set of indice corresponding to the m^* largest values of $\sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2$ in H^{*c} . We can prove (see proof of Theorem 3.1 in [14] with $\lambda_g = r_n \sqrt{d_g}$) that

$$\sum_{g \in J^{*c}} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 \leq 9 \frac{d_{max}}{d_{min}} \sum_{g \in J^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 \quad (2)$$

and

$$\sum_{g \in J^{*c}} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \leq 3 \sum_{g \in J^*} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + \frac{1}{2n}.$$

Therefore on one hand, using assumption $GS(2m^*, 3, \frac{1}{2n}, k')$, we deduce

$$k' \sum_{g \in J^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 \leq \mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 + \frac{1}{2n} \quad (3)$$

and on the other hand, by the same arguments as those used in the proof of Theorem 3.6 to state Equation (14), we have

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{16}{c_n^2 k'} r_n^2 \gamma^* + \frac{2r_n + 1}{2nc_n}. \quad (4)$$

From equation (3) and equation (4) we conclude

$$\sum_{g \in J^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 \leq \frac{16}{k'^2 c_n^2} r_n^2 \gamma^* + \frac{2r_n + 1}{2k' c_n n} + \frac{1}{2k' n}. \quad (5)$$

Finally inequalities (2) and (5) conclude the proof. \square

The convergences rates obtained in Theorem 3.6 and Theorem 3.7 are exactly of the same order as the ones stated by Lounici and al. [14] for the Group Lasso in a Gaussian setting. The oracle inequality stated in Theorem 3.7 shows that the l_2 -estimation error is bounded by $O\left(\gamma^* \frac{\log G_n}{n}\right)$ under $GS(2m^*, 3, \frac{1}{2n}, k')$. Therefore, in the case of finite size groups, m^* still could be much larger than $\log G_n$ and the estimator remains consistent. We can also notice that the number of samples required in order that the prediction and estimation error (with respect to the l_2 norm) goes to zero is almost of the order of $O(\gamma^* \log G_n)$.

As mentioned above the group structured norm satisfies the decomposability property (see [20]). Furthermore, under the assumptions made, the loss function satisfies a local restricted strong convexity property. According to Negahban et al., these properties are two important conditions that ensure good statistical properties of M-estimators. This is especially true for the Group Lasso applied to generalized linear model as shown in Theorem 3.6 and Theorem 3.7. Indeed, these two theorems demonstrate the ability of the Group Lasso to recover good approximation of the true model for sparse generalized linear models under the Group Stabil condition.

3.2 Lasso for generalized linear models

When each group is of size one we recover the Lasso estimator

$$\hat{\beta}_n = \underset{\beta \in \Lambda}{\operatorname{argmin}} \{ \mathbb{P}_n l(\beta) + 2r_n \|\beta\|_1 \}$$

where $\|\beta\|_1 = \sum_{j=1}^n |\beta_j|$. Thus following step by step the proof of Theorem 3.6 we can easily deduce bounds for estimation and prediction error for the Lasso estimator in the case of generalized linear models. Notice that the l_2 -estimation error of the Lasso applied to GLMs was first studied by Negahban and al. (see [19] and [20]). The Lasso is a special case of the Group Lasso where $\gamma^* = s^*$ with $s^* := |I^*| = |\{j : \beta_j^* \neq 0\}|$ and $G_n = p$. We still consider high-dimensional data i.e. $n \ll p$ and sparsity assumption on the target β^* i.e. $s^* \ll p$ and we assume (H.1-3) except that for (H.3) we consider the l_1 norm i.e. $\|\beta^*\|_1 \leq B$. The condition GS in this case requires the existence of $0 < k < 1$ such that $\delta^T \Sigma \delta \geq k \sum_{j \in I^*} \delta_j^2 - \varepsilon$ for any $\delta \in S(c_0, \varepsilon) = \left\{ \delta \in \mathbb{R}^p : \sum_{j \in I^{*c}} |\delta_j| \leq c_0 \sum_{j \in I^*} |\delta_j| + \varepsilon \right\}$. We recover the Stabil condition $St(c_0, \varepsilon, k)$ of [5]. We also define condition $St(2s^*, c_0, \varepsilon, k')$ i.e there exists $0 < k' < 1$ such that $\delta^T \Sigma \delta \geq k' \sum_{j \in J} \delta_j^2 - \varepsilon$ for any δ which satisfies $\sum_{j \in J^c} |\delta_j| \leq c_0 \sum_{j \in J} |\delta_j| + \varepsilon$ and J such that $|J| \leq 2s^*$.

Theorem 3.8. Assume condition $St(3, \frac{1}{2n}, k)$ is fulfilled. Let

$$r_n \geq AKL \left\{ C_{L,B} \vee \max_{\{|x| \leq Lk_n\} \cap \Theta} |\psi'(x)| \right\} \sqrt{\frac{\log(2p)}{n}}$$

where $A > \sqrt{2}$ and $C_{L,B}$ depends only on L and B . We have, with probability at least $1 - C(2p)^{-A^2/2}$ (where C is a universal constant),

$$\|\hat{\beta}_n - \beta^*\|_1 \leq \frac{4}{c_n k} r_n s^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n}$$

and

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{16}{c_n^2 k} r_n^2 s^* + \frac{2r_n + 1}{2nc_n}.$$

Furthermore if $St(2s^*, 3, \frac{1}{2n}, k')$ holds then we have

$$\|\hat{\beta}_n - \beta^*\|_2^2 \leq 10 \frac{d_{max}}{d_{min}} \left\{ \frac{16}{k'^2 c_n^2} r_n^2 s^* + \frac{2r_n + 1}{2k' c_n n} + \frac{1}{2k' n} \right\},$$

with $c_n := \min_{\{|x| \leq L(9B + \frac{1}{n})\} \cap \Theta} \left\{ \frac{\psi''(x)}{2} \right\}.$

Proof. The proof of Theorem 3.8 follows the same guidelines as the one for the Group Lasso. The main difference comes from concentration inequalities for the linear and log Laplace transform part of the loss function, leading to simpler bounds. \square

This result extends the one of Bunea [5] for logistic regression to generalized linear model and states convergence rates for the estimation and prediction error. The error bounds presented in Theorem 3.8 are of the same order as the ones stated by Bickel, Ritov and Tsybakov [2] in their analysis of the properties of the Lasso for standard linear models. We can also notice that in [19] Negahban and al. obtained bounds for the l_2 -estimation error of the Lasso applied to GLMs of the same order as the one we get but under stronger conditions. In fact they assume that the distribution of the response Y based on a predictor X is given by

$$P(Y|X; \beta^*) = \exp(Y\beta^{*T}X - \psi(\beta^{*T}X))$$

with $|X| \leq A$ and $|Y| \leq B$ and that ψ'' is bounded from below on a suitable set. In our work we do not make these two last assumptions. The restricted eigenvalue property they use is also slightly stronger than $St(3, \varepsilon_n, k)$. In addition we establish oracle inequalities for the prediction error.

The bounds in Theorem 3.8 are meaningful if r_n is small (in particular if $n \gg \log(p)$) and s^* is small. Indeed, the bound on the l_1 -estimation error is of the order of $O\left(s^* \sqrt{\frac{\log p}{n}}\right)$. We can also notice that the minimum number of samples required to make the l_2 -estimation and prediction error decrease to zero is of the order of $O(s^* \log p)$.

Then a relevant issue is the benefit of the Group Lasso over the Lasso. To understand better the power of a group structured estimator when the covariates have a group structure we refer to Huang and Zhang [10]. In this paper the authors investigate the benefit of sparsity with group structure compared to the usual Lasso. They develop a concept called strong group sparsity which means that the signal β^* is efficiently covered by grouping. They showed that the Group Lasso is better than the usual Lasso for strongly group sparse data in the case of a standard linear model (see also [15]). Comparing Theorem 3.6 and Theorem 3.8 we see that an important improvement over the Lasso is given when the number of non-zero groups is much smaller than the total number of non-zero coefficients. Actually when the number of covariates increases this estimator removes completely the effect of the number of predictors. Indeed if we assume that the maximum size of the groups is finite then the estimation error for the Group Lasso depends only on the total number of groups G_n and on the number of significant groups m^* . Besides the prediction error

for the Lasso is of the order of $O(s^* \frac{\log p}{n})$ whereas it is of the order of $O(m^* \frac{\log G_n}{n})$ for the Group Lasso. Therefore if β^* has a group structure this is a meaningful improvement when p is large and the number of groups G_n is much more smaller. Notice this comparison must be tempered by the fact that the two estimators require different conditions on the covariance matrix. In fact for a same data set there is not a condition which is weaker and implies the other. We also refer to the simulations in Section 5 that show some of the benefit of the Group Lasso over the Lasso when the covariates have a group structure in term of estimation and prediction error.

4 Applications and extensions

4.1 Group Lasso for Poisson regression

Sparse logistic regression has been widely studied in the literature (see, for example, [5] and [17]) but not the Poisson one for a sparse model. This last model is also very useful for many practical applications. For instance it is used to model count data and contingency tables. Poisson models are a special case of generalized linear models where the conditionnal law of Y given $X = x$ has a Poisson distribution with parameter $\lambda^*(x) := \exp(\beta^{*T} x)$. Therefore the conditional mean for Poisson regression is modeled by $\mathbb{E}(Y|X = x) = \exp(\beta^{*T} x)$. Thus the normalized function ψ is the exponential function and is defined on \mathbb{R} . For this special link function we are going to specify the constants which appear in Theorem 3.6. The log-likelihood based on the observations is given by

$$\mathcal{L}(\beta) = \sum_{i=1}^n [Y_i \beta^T X_i - \exp(\beta^T X_i) - \log(Y_i!)]$$

and thus the Poisson loss function is defined by

$$l(\beta) =: l(\beta; x, y) =: -y\beta^T x + \exp(\beta^T x).$$

It is formula (1) with $\psi = \exp$. In the particular case of Poisson regression the conditional law is defined by $Y|X \sim \mathcal{P}(\lambda^*(X))$ and the higher moments for a Poisson distribution are given by

$$\mathbb{E}(Y^k | X) = \sum_{l=1}^k (\lambda^*(X))^l S_{l:k}$$

where $S_{l:k} = \frac{1}{l!} \sum_{i=0}^l (-1)^{l-i} \binom{l}{i} i^k \geq 0$ is the number of partitions of a set with l members into k undersets. The number $\sum_{l=1}^k S_{l:k} := B_k$ is called the k^{th} Bell number and this number satisfies the relation $B_{n+1} = \sum_{k=0}^n \binom{n}{k} B_k$ (see for example [11]). So we can easily prove by induction that $B_k \leq k!$ for all $k \geq 1$. Then we have on the event $\{0 \leq \lambda^*(X) < 1\}$

$$\mathbb{E}(Y^k) = \mathbb{E}\left(\sum_{i=1}^k (\lambda^*(X))^i S_{i:k}\right) \leq \sum_{i=1}^k S_{i:k} \leq k!$$

and on the event $\{1 \leq \lambda^*(X)\}$ using (H.1) combined with (H.3) we find

$$\mathbb{E}(Y^k) = \sum_{i=1}^k S_{i:k} \mathbb{E}\left((\lambda^*(X))^i\right) \leq k!(e^{LB})^k$$

Because $e^{LB} \geq 1$ we deduce that for all $k \geq 1$

$$\mathbb{E}|Y|^k \leq k!(e^{LB})^k.$$

Thus for Poisson regression we have $C_{L,B} = e^{LB}$. Besides $\max_{|x| \leq L\kappa_n} \{|\psi'(x)|\} = e^{L\kappa_n}$ and $\min_{|x| \leq L(9B + \frac{1}{n})} \left\{\frac{\psi''(x)}{2}\right\} = \frac{1}{2}e^{-L(9B + \frac{1}{n})}$ where we recall that $\kappa_n := 17B + \frac{2}{n}$. Therefore, in the case of Poisson regression, Theorem 3.6 becomes

Corollary 4.1. Let $c_n := \frac{1}{2}e^{-L(9B+\frac{1}{n})}$. Assume that condition $GS(3, \frac{1}{2n}, k)$ holds. If

$$r_n \geq AKLe^{L(17B+\frac{2}{n})} \sqrt{\frac{\log(2G_n)}{n}}$$

with $A > \sqrt{2}$ then, with probability at least $1 - (C + 2d_{max})(2G_n)^{-A^2/2}$, we have

$$\begin{aligned} \|\hat{\beta}_n - \beta^*\|_R &\leq \frac{4}{c_n k} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n}, \\ \|\hat{\beta}_n - \beta^*\|_{2,1} &\leq \frac{4}{c_n k \sqrt{d_{min}}} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n \sqrt{d_{min}}} \end{aligned}$$

and

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{16}{c_n^2 k} r_n^2 \gamma^* + \frac{2r_n + 1}{2nc_n}.$$

Furthermore if condition $GS(2m^*, 3, \frac{1}{2n}, k')$ holds, then

$$\|\hat{\beta}_n - \beta^*\|_2^2 \leq \frac{d_{max}}{d_{min}} \left[160 \frac{1}{k'^2 c_n^2} r_n^2 \gamma^* + 5 \frac{2r_n + 1}{k' c_n n} + 10 \frac{1}{2k' n} \right].$$

4.2 Elastic net for generalized linear models

The most difficult part of the proof of Theorem 3.6 is to prove Proposition 3.1. Once this proposition has been proved it becomes easy to generalize the results presented above to any standard penalization using a modified version of the condition GS (depending on the norm we use). For example we can replace the ℓ_1 norm by a combination of ℓ_1 and ℓ_2 norms. It is the so-called Elastic net introduced by Zou, Trevor and Hastie [33] in the frame of linear regression. They showed that this estimator outperforms the Lasso in many situations for real world data and simulations. It is an alternative to the Group Lasso (the Elastic net has a behaviour similar to the one of the Group Lasso estimator). As the Lasso does, the Elastic net encourages sparsity and group selection but contrary to the Lasso when the sample size n is smaller than p the Elastic net can select more than n significant variables. This estimator is the solution of a convex optimization problem. Zou, Trevor and Hastie in [33] proposed an algorithm to solve this problem. The Elastic net estimator for the generalized linear model is defined by

$$\hat{\beta}_n = \operatorname{argmin}_{\beta \in \Lambda} \left\{ \mathbb{P}_n l(\beta) + 2r_n \|\beta\|_1 + t_n \|\beta\|_2^2 \right\} \quad (6)$$

where r_n and t_n are the penalty parameters. Theorem 4.2 is an extension of the results first proved by Bunea [5] in the special case of logistic regression. Let $2t_n B = r_n$. We have the following theorem

Theorem 4.2. Assume condition $St(4, \frac{1}{2n}, k)$ holds. Let

$$r_n \geq AKL \left\{ C_{L,B} \vee \max_{\{|x| \leq L(17B+\frac{2}{n})\} \cap \Theta} |\psi'(x)| \right\} \sqrt{\frac{\log(2p)}{n}}$$

where $A > \sqrt{2}$ and $C_{L,B}$ depends only on L and B . Then, with probability at least $1 - C(2p)^{-A^2/2}$ (where C is a universal constant), we have

$$\|\hat{\beta}_n - \beta^*\|_1 \leq \frac{(2.5)^2}{t_n + c_n k} r_n s^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n}$$

and

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{2(2.5)^2}{c_n k (t_n + c_n k)} r_n^2 s^* + \frac{2r_n + 3}{2c_n n}.$$

where $c_n := \min_{\{|x| \leq L(9B+\frac{1}{n})\} \cap \Theta} \left\{ \frac{\psi''(x)}{2} \right\}$.

We can notice that thanks to the ℓ_2 penalty the bound for the ℓ_1 and ℓ_2 errors are less sensitive to small value of k and small value of c_n (which can appears when L or B are large).

5 Simulations

We are going to compare the performances of the Lasso and Group Lasso for Poisson Regression on simulated data sets. Computations have been performed using R. We use the package `grplasso` developed by Meir, van de Geer and Buhlmann [17] for the Group Lasso and the package `glmnet` developed by Friedman, Hastie and Tibshirani [8] for the Lasso. The function `glmnet` fits the entire lasso regularization path for some generalized linear model via penalized maximum likelihood. We use this function in the particular case of Poisson regression. The function `grplasso` fits the solution of a Group Lasso problem for a model of type `grpl.model` which generates models to be used for Group Lasso algorithm and identify the exponential family of the response and the link function which is used. Here we consider the function `PoissReg()` which generates a Poisson model.

We simulate 100 data sets for each simulation and we ran the Lasso and Group Lasso on these data sets. Each data set X is cut into three separate subsets: a training data set, a validation data set and a test data set. We simulate responses via the model $Y \sim \mathcal{P}(\exp(X\beta^*))$ where $\beta^* = (\beta^{*1}, \dots, \beta^{*G})$ with g groups with non zero coefficients among the G groups. The training data set (X_{train} of size n_{train}) is used to fit the model (we estimate the target β^* for a sequence of the tuning parameter λ and denote by β_λ the estimate of β^* obtained for such a parameter for the Lasso and the Group Lasso estimator). Then, we use the validation data (X_{valid} of size n_{valid}) to evaluate the performance of the fitted model according to a specific loss function. We define the optimal tuning parameter as the one for which the deviation from the fitted mean to the response is minimal i.e. $\lambda_{\text{opt}} \in \underset{\lambda}{\operatorname{argmin}} \left\{ \frac{1}{n_{\text{valid}}} \|Y - \exp(X_{\text{valid}}\beta_\lambda)\|_2^2 \right\}$. From that we determine the model with the parameter vector $\beta_{\lambda_{\text{opt}}}$. Then we compute the hits i.e the number of correctly identified relevant variables, the false positives i.e the number of non significant variables chosen as relevant and the degree of freedom i.e the total number of variables selected in the model. Finally, we estimate the performance of the selected model by computing the coefficients estimation error $\|\beta^* - \beta_{\lambda_{\text{opt}}}\|_1$ and the prediction error $\|X_{\text{test}}\beta^* - X_{\text{test}}\beta_{\lambda_{\text{opt}}}\|_2$ on the test data (X_{test} of size n_{test}). We ran the Lasso and Group Lasso on these data sets.

Eight models are considered in the simulations. For each simulation $n_{\text{train}} = 50$, $n_{\text{valid}} = 50$, $n_{\text{test}} = 100$. To compare the Lasso and Group Lasso we use a random design matrix where the predictors are simulated as followed according to a uniform distribution to have bounded predictors.

1.
 - $X_i = U_1 + \varepsilon_i$ for $1 \leq i \leq 10$ with $U_1 \sim U([0, 1])$
 - $X_i = U_2 + \varepsilon_i$ for $11 \leq i \leq 20$ with $U_2 \sim U([0, 1])$
 - $X_i = U_i$ for the last 100 variables $U_i \sim U([-0.1, 0.1])$

with ε_i i.i.d $\sim U([0, 0.01])$.

The covariates within the first two blocks are highly correlated (~ 0.8) and there are small correlations between the blocks. The target is

$$\beta^* = (\underbrace{0.3, \dots, 0.3}_{10}, \underbrace{0.2, \dots, 0.2}_{10}, \underbrace{0, \dots, 0}_{10}, \underbrace{0, \dots, 0}_{10}).$$

2. The simulation is the same as the first one except that ε_i i.i.d $\sim U([0, 1])$. Thus there are small correlations within and between groups (~ 0.5).
3. The simulation is the same as the first one except that ε_i i.i.d $\sim U([0, 1.2])$. Thus there are very small correlations within and between groups (~ 0.2).

For all the following simulations the non zero groups are generated in the same way as in the second example. For $j = 1, \dots, m^*$ and $i = 1, \dots, d_j$, $X_i = U_j + \varepsilon_i$ where $U_j \sim U([0, 1])$ and ε_i i.i.d $\sim U([0, 1])$. The non influential groups are simulated according to $X_i = U_i$ with $U_i \sim U([-0.1, 0.1])$. In the two following simulations we increase the size of the non zero groups

4.

$$\beta^* = (\underbrace{0.3, \dots, 0.3}_{20}, \underbrace{0.2, \dots, 0.2}_{20}, \underbrace{0, \dots, 0}_{20}, \dots, \underbrace{0, \dots, 0}_{20}).$$

5.

$$\beta^* = (\underbrace{0.3, \dots, 0.3}_{5}, \underbrace{0.2, \dots, 0.2}_{5}, \underbrace{0, \dots, 0}_{5}, \dots, \underbrace{0, \dots, 0}_{5}).$$

In the last three simulations it is the number of the non zero groups which is increased.

6.

$$\beta^* = (\underbrace{0.2, \dots, 0.2}_{10}, \underbrace{0.2, \dots, 0.2}_{10}, \underbrace{0, \dots, 0}_{10}, \dots, \underbrace{0, \dots, 0}_{10}).$$

7.

$$\beta^* = (\underbrace{0.2, \dots, 0.2}_{10}, \dots, \underbrace{0.2, \dots, 0.2}_{10}, \underbrace{0, \dots, 0}_{10}, \dots, \underbrace{0, \dots, 0}_{10}).$$

8.

$$\beta^* = (\underbrace{0.2, \dots, 0.2}_{10}, \dots, \underbrace{0.2, \dots, 0.2}_{10}, \underbrace{0, \dots, 0}_{10}, \dots, \underbrace{0, \dots, 0}_{10}).$$

The results of the eight simulations are reported in Table I hereafter where p is the number of covariates, s^* the number of significant covariates, G the number of groups, m^* the number of non zero groups and v is the size of the non zero groups. The Group Lasso seems to perform better than the Lasso to include the relevant predictors into the model particularly when there are high within group correlations. The Lasso tends to select fewer variables among the influential ones (the Lasso selects only some variables from the groups of highly correlated predictors) than the Group Lasso does in the case of highly correlated covariates and when the size or the number of the nonzero groups is large. The Group Lasso succeeds in including the true significant groups in most of the cases. On the contrary the Group Lasso estimator tends to add more irrelevant covariates than the Lasso does in particular when the number or the size of the non influential groups is large. We can expect such a result because the Group Lasso estimator includes not single variable one after another but groups of variables (when one of the covariates is included in the model all the others which belongs to the same group are also included in the model). Thus the Group Lasso selects models that are larger than the true model. However when the Group Lasso selects a number of covariates which is the same than the number of significant covariates it is, with high probability, the correct groups which are included. We have also measured the performances of the Lasso and Group Lasso in terms of estimation error and prediction error. The Group Lasso seems to perform better than the Lasso in term of estimation error and of prediction error in most of the cases and the improvement is particularly meaningful for prediction error. We can also notice that the performances of the two estimators decreases with the increase of G and m^* . To conclude, the Group Lasso estimator seems to perform better than the Lasso to include the hits in the model and in terms of prediction and estimation error when the covariates are structured into groups and in particular in the case of high correlations within groups.

To better understand the difference of behaviour between the Lasso and the Group Lasso in terms of variables selection we also provide some kind of ROC curves. These curves are created by plotting the fraction of true positives (of the estimated parameter) out of the positive coefficients of β^* for discretized values of the penalty parameter going from zero (completely dense solution) up to the value where the first group of covariates enters the model (completely sparse solution). To ensure that we well detect each new inclusion of covariates, the discretization involves 10000

Simulation	1	2	3	4	5	6	7	8
s^*/p	20/120	20/120	20/120	40/240	10/60	20/120	40/140	60/160
G	12	12	12	12	12	12	14	16
m^*	2	2	2	2	2	2	4	6
v	10	10	10	20	5	10	10	10
Mean Hit lasso (%)	86, 9	99, 4	99, 95	66, 27	87, 6	95, 9	78, 79	54, 16
Mean Hit group lasso (%)	100	100	100	98, 5	100	100	100	98, 8
Mean False positive lasso (%)	15, 02	10, 61	9, 25	2, 39	47, 38	18, 96	9, 33	8, 87
Mean False positive group lasso (%)	28, 7	36, 1	33, 9	74	34, 2	29, 6	61, 9	37, 9
Mean Nonzero lasso	32, 4	30, 49	29, 24	31, 29	32, 45	38, 14	40, 85	41, 37
Mean Nonzero gp lasso	48, 7	56, 1	53, 9	187, 4	27, 1	49, 6	101, 9	97, 2
Mean Prediction error lasso	0, 19	0, 15	0, 18	6, 73	0, 15	0, 17	3, 22	17, 97
Mean Prediction error gp lasso	0, 021	0, 007	0, 004	1, 98	0, 03	0, 016	0, 012	1, 32
Mean Estimation error lasso	6, 36	2, 84	2, 08	9, 79	15, 31	6, 37	8, 45	16, 28
Mean Estimation error gp lasso	5, 54	2, 66	1, 64	18, 65	4, 22	3, 40	2, 77	12, 76

Table 1: Results of the simulations for the eight models

values of the penalty parameter. We plot the curves for each model considered. Three models are described below and the respective ROC curves are shown in Figures 1, 2 and 3.

For each model we have

- for $j = 1, \dots, m^*$ and $i = 1, \dots, d_j$, $X_i = U_j + \varepsilon_i$ where ε_i i.i.d $\sim U([-1, 1])$ and $U_j \sim U([0, 1])$ for simulation 1 and $U_j \sim U([-1, 1])$ for simulation 2 and 3 .
- for the non significant groups $X_i = U_i$ with $U_i \sim U([-0.1, 0.1])$ for simulation 1 and 2 and $U_i \sim U([-0.01, 0.01])$ for simulation 3.

1. For the first simulation the target is

$$\beta^* = \underbrace{(0.2, \dots, 0.2, \dots, 0.2, \dots, 0.2)}_{10}, \underbrace{(0, \dots, 0, \dots, 0, \dots, 0)}_{50}.$$

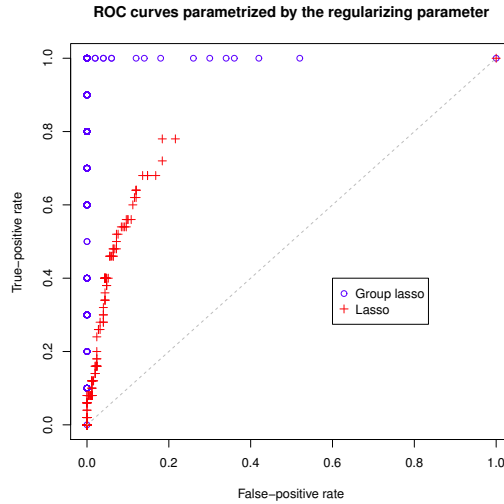


Figure 1: Simulation 1

2. For the second simulation the target is

$$\beta^* = (\underbrace{0.2, \dots, 0.2}_{5}, \dots, \underbrace{0.2, \dots, 0.2}_{5}, \underbrace{0, \dots, 0}_{5}, \dots, \underbrace{0, \dots, 0}_{5}).$$

10
100

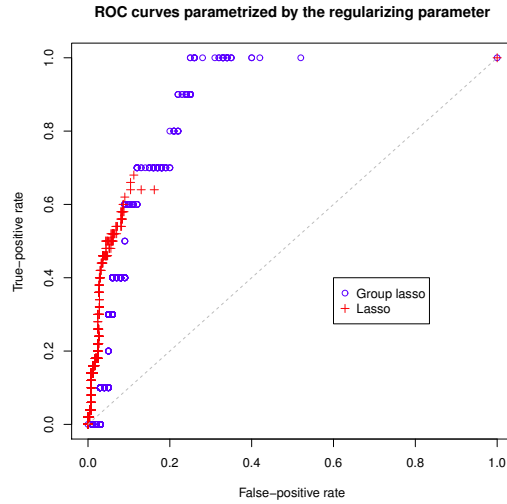


Figure 2: Simulation 2

3. For the third simulation the target is

$$\beta^* = (\underbrace{0.2, \dots, 0.2}_{5}, \dots, \underbrace{0.2, \dots, 0.2}_{5}, \underbrace{0, \dots, 0}_{5}, \dots, \underbrace{0, \dots, 0}_{5}).$$

10
100

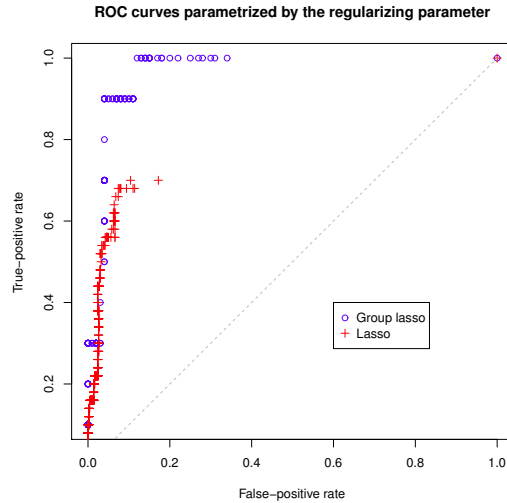


Figure 3: Simulation 3

The simulations clearly illustrate that contrary to the Lasso the Group Lasso includes not single covariates one after another but all the covariates which belong to the same group at the same time. The Group Lasso tends to include in priority the significant groups and then the non significant groups in such a way that for some values of the penalty parameter the rate of perfect selection can

be very close to one and even equal to one in some particular cases (see for instance Figure 1). We can notice that the ROC curves for the Lasso follow at the beginning those of the Group Lasso and then fall below. The break in the curves of the Lasso illustrate a special feature of this estimator. Actually the number of covariates included in the model for the Lasso is restricted by the sample size whereas for the Group Lasso the restriction is given by the number of groups. That is why for the Lasso the number of true positives but also the number of false positives are smaller than those of the Group Lasso (once approximately 100 covariates are added in the model the Lasso stop including other ones). We can also notice that the Group Lasso is less stable than the Lasso in terms of variables selection. In fact a little change in the penalty parameter can substantially increase the number of false positives. This particular feature of the Group Lasso due to a group structure penalty is reflected in Table I with a rate of false positives larger for the Group Lasso in most of the cases. Thus Figures 1, 2 and 3 confirm what has been previously deduced from the results presented in Table I. To conclude the Group Lasso seems to be more reliable than the Lasso to include the variables of interest but in return it tends to incorporate more false positives.

6 Conclusion

We consider the generalized linear model in high dimensional settings and use the Group Lasso estimator to estimate the regression parameter β^* when the covariates are naturally structured into groups of variables and the true parameter is group sparse (just a few variables are relevant to explain the response). Under some assumptions on the sparsity of β^* , on the correlations between the groups of covariates and on the tuning parameter of the estimator we state general oracle inequalities for estimation and error prediction for the Group Lasso estimator applied to generalized linear models. In the particular case of groups of size one, we provide original inequalities for the Lasso in the case of generalized linear models extending the results of Bunea [5] for logistic regression. Furthermore, we extend these results to other penalties such as the Elastic net. Then we compare the performances of the Lasso to the ones of the Group Lasso on simulated data. We show the improvement in terms of variables selection and prediction error of the Group Lasso compared to the Lasso when the covariates are structured into groups. Moreover we illustrate on these simulated data the impact of the total number of groups, of the number of non zero groups and of the size of the groups on the performances of the Group Lasso. The conclusion is that the Group Lasso estimator behave well in a high dimensional setting under sparsity and group correlations assumptions. However the main drawback of the Group Lasso is that we need an a priori knowledge on the groups and it is not always possible.

A Proof of Theorem 3.6

A.1 The main steps of the proof

The proof follows the guidelines in [4] or [13]. Using the mere definition of $\hat{\beta}_n$, we have

$$\mathbb{P}_n l(\hat{\beta}_n) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g\|_2 \leq \mathbb{P}_n l(\beta^*) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2. \quad (7)$$

Hence we get

$$\begin{aligned} & \mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g\|_2 \\ & \leq (\mathbb{P}_n - \mathbb{P}) \left(l(\beta^*) - l(\hat{\beta}_n) \right) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2. \end{aligned} \quad (8)$$

We decompose the empirical process into a linear part and a part which depends on the normalized parameter ψ .

$$(\mathbb{P}_n - \mathbb{P}) \left(l(\beta^*) - l(\hat{\beta}_n) \right)$$

$$= (\mathbb{P}_n - \mathbb{P}) \left(l_l(\beta^*) - l_l(\hat{\beta}_n) \right) + (\mathbb{P}_n - \mathbb{P}) \left(l_\psi(\beta^*) - l_\psi(\hat{\beta}_n) \right)$$

where

$$l_l(\beta) := l_l(\beta, x, y) = -y\beta^T x$$

and

$$l_\psi(\beta) := l_\psi(\beta, x) = \psi(\beta^T x).$$

From Proposition 3.3 and Proposition 3.5 and by adding $r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2$ to both sides of the inequality (8) we find, on $\mathcal{A} \cap \mathcal{B}$, that

$$\begin{aligned} & r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + \mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) \\ & \leq 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \left(\|\hat{\beta}_n^g - \beta^{*g}\|_2 + \|\beta^{*g}\|_2 - \|\hat{\beta}_n^g\|_2 \right) + \frac{r_n}{2} \varepsilon_n. \end{aligned}$$

If $g \notin H^*$ then $\|\hat{\beta}_n^g - \beta^{*g}\|_2 + \|\beta^{*g}\|_2 - \|\hat{\beta}_n^g\|_2 = 0$ and otherwise $\|\beta^{*g}\|_2 - \|\hat{\beta}_n^g\|_2 \leq \|\hat{\beta}_n^g - \beta^{*g}\|_2$. So the last inequality can be bounded by

$$4r_n \sum_{g \in H^*} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + \frac{r_n}{2} \varepsilon_n. \quad (9)$$

By the definition of β^* we have $\mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) > 0$ and therefore

$$\sum_{g \notin H^*} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \leq 3 \sum_{g \in H^*} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + \frac{\varepsilon_n}{2}$$

i.e $\hat{\beta}_n - \beta^* \in S(3, \frac{\varepsilon_n}{2})$. The next proposition provides a lower bound for $\mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right)$.

Proposition A.1. *On the event $\mathcal{A} \cap \mathcal{B}$ we have*

$$\mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) \geq c_n \mathbb{E} \left[\left(f_{\hat{\beta}_n}(X) - f_{\beta^*}(X) \right)^2 \right]$$

with $c_n := \min_{|x| \leq L(9B + \frac{1}{n})} \left\{ \frac{\psi''(x)}{2} \right\}$.

Proof.

$$\begin{aligned} \mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) &= -\mathbb{E} \left[\mathbb{E}(Y|X) \left(f_{\hat{\beta}_n}(X) - f_{\beta^*}(X) \right) \right] \\ &\quad + \mathbb{E} \left[\psi'(f_{\beta^*}(X)) \left(f_{\hat{\beta}_n}(X) - f_{\beta^*}(X) \right) \right] \\ &\quad + \mathbb{E} \left[\frac{\psi''(f_{\hat{\beta}_n}(X))}{2} \left(f_{\hat{\beta}_n}(X) - f_{\beta^*}(X) \right)^2 \right] \end{aligned}$$

where $\tilde{\beta}^T X$ is an intermediate point between $\hat{\beta}_n^T X$ and $\beta^{*T} X$ given by a second order Taylor expansion of ψ . Since $\psi'(f_{\beta^*}(X)) = \mathbb{E}(Y|X)$ we find

$$\mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) = \mathbb{E} \left[\frac{\psi''(f_{\tilde{\beta}}(X))}{2} \left(f_{\hat{\beta}_n}(X) - f_{\beta^*}(X) \right)^2 \right].$$

Besides we have

$$\begin{aligned} |\tilde{\beta}^T X| &\leq |\hat{\beta}_n^T X - \beta^{*T} X| + |\beta^{*T} X| \\ &\leq \sum_{g=1}^{G_n} \|\beta^{*g} - \hat{\beta}_n^g\|_2 \|X^g\|_2 + \sum_{g=1}^{G_n} \|\beta^{*g}\|_2 \|X^g\|_2. \end{aligned}$$

Applying (H.1), we find

$$\|X^g\|_2 \leq L\sqrt{d_g}.$$

Then using Lemma 3.4 and (H.3) we find

$$|\tilde{\beta}^T X| \leq LM + LB \quad \text{a.s.} \quad (10)$$

Furthermore, β^* and $\hat{\beta}_n$ belongs to Λ which is a convex set. Therefore $\tilde{\beta} \in \Lambda$ and $\tilde{\beta}^T X \in \Theta$ a.s. Thus we conclude

$$\mathbb{P}\left(l(\hat{\beta}_n) - l(\beta^*)\right) \geq c_n \mathbb{E}\left[\left(f_{\hat{\beta}_n}(X) - f_{\beta^*}(X)\right)^2\right]$$

where $c_n := \min_{\{|x| \leq L(M+B)\} \cap \Theta} \left\{ \frac{\psi''(x)}{2} \right\}$. □

From Proposition A.1 and (9) we deduce that

$$\begin{aligned} r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + c_n \mathbb{E}\left(\hat{\beta}_n^T X - \beta^{*T} X\right)^2 \\ \leq 4r_n \sum_{g \in H^*} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + \frac{r_n}{2} \varepsilon_n. \end{aligned} \quad (11)$$

Then the end of the proof is similar to the end of the proof of Theorem 2.4 in [5]. Let Σ be the $p \times p$ covariance matrix whose entries are $\mathbb{E}(X_k X_j)$. We have

$$\mathbb{E}\left(\hat{\beta}_n^T X - \beta^{*T} X\right)^2 = (\hat{\beta}_n - \beta^*)^T \Sigma (\hat{\beta}_n - \beta^*).$$

Because condition $GS(3, \frac{\varepsilon_n}{2}, k)$ is satisfied we have

$$c_n (\hat{\beta}_n - \beta^*)^T \Sigma (\hat{\beta}_n - \beta^*) \geq c_n k \sum_{g \in H^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 - \frac{\varepsilon_n}{2}.$$

Then by using Cauchy-Schwarz inequality in (11) we find

$$\begin{aligned} r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + c_n k \sum_{g \in H^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 \\ \leq 4r_n \sqrt{\sum_{g \in H^*} d_g} \sqrt{\sum_{g \in H^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2} + (r_n + 1) \frac{\varepsilon_n}{2}. \end{aligned}$$

Now the fact that $2xy \leq tx^2 + y^2/t$ for all $t > 0$ leads to the following inequality

$$\begin{aligned} r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 + c_n k \sum_{g \in H^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 \\ \leq 4tr_n^2 \gamma^* + \frac{1}{t} \sum_{g \in H^*} \|\hat{\beta}_n^g - \beta^{*g}\|_2^2 + (r_n + 1) \frac{\varepsilon_n}{2}. \end{aligned} \quad (12)$$

Replacing t by $\frac{1}{c_n k}$ in (12) we obtain

$$\sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \leq \frac{4}{c_n k} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{\varepsilon_n}{2}.$$

i.e

$$\|\hat{\beta}_n - \beta^*\|_R \leq \frac{4}{c_n k} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{\varepsilon_n}{2}.$$

What is more

$$\sqrt{d_{\min}} \|\hat{\beta}_n - \beta^*\|_{2,1} \leq \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2.$$

This yields

$$\|\hat{\beta}_n - \beta^*\|_{2,1} \leq \frac{4}{c_n k \sqrt{d_{\min}}} r_n \gamma^* + \left(1 + \frac{1}{r_n}\right) \frac{1}{2n \sqrt{d_{\min}}}. \quad (13)$$

A similar argument could be made to prove

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{16}{c_n^2 k} r_n^2 \gamma^* + \frac{2r_n + 1}{c_n} \frac{\varepsilon_n}{2}. \quad (14)$$

Finally we conclude the proof using Proposition 3.1.

A.2 Proof of Proposition 3.1

Proof. Let $A > \sqrt{2}$. We recall that we have made the assumption $\frac{\log(2G_n)}{n} \leq 1$. We deduce Proposition 3.1 from the two following lemmas.

Lemma A.2. *Let*

$$r_n \geq \left(8\sqrt{2}ALC_{L,B} \sqrt{\frac{\log(2G_n)}{n}} \right) \vee \left(16A^2LC_{L,B} \frac{\log(2G_n)}{n} \right)$$

with $A > 1$. Then

$$\mathbb{P}\{\mathcal{A}\} \geq 1 - 2d_{\max}(2G_n)^{1-A^2}.$$

Lemma A.3. *Let*

$$r_n \geq 20AL \left(\max_{\{|x| \leq L\kappa_n\} \cap \Theta} |\psi'(x)| \right) \sqrt{\frac{2 \log 2G_n}{n}}$$

where $A \geq 1$. Then

$$\mathbb{P}(\mathcal{B}) \geq 1 - C(2G_n)^{-A^2/2}$$

where we recall $\kappa_n := 17B + \frac{2}{n}$. We can notice that $\mathbb{P}(\mathcal{B}) \xrightarrow[n \rightarrow \infty]{} 1$.

Thus if

$$r_n \geq AKL \left\{ C_{L,B} \vee \max_{\{|x| \leq L\kappa_n\} \cap \Theta} |\psi'(x)| \right\} \sqrt{\frac{2 \log(2G_n)}{n}}$$

with K choosen such that

$$r_n \geq \max(C_1, C_2, C_3)$$

where

$$C_1 := 8\sqrt{2}ALC_{L,B} \sqrt{\frac{\log(2G_n)}{n}}$$

$$C_2 := 16A^2LC_{L,B} \frac{\log(2G_n)}{n}$$

and

$$C_3 := 20AL \left(\max_{\{|x| \leq L\kappa_n\} \cap \Theta} |\psi'(x)| \right) \sqrt{\frac{2 \log 2G_n}{n}}$$

then $\mathbb{P}(\mathcal{A} \cap \mathcal{B}) \geq 1 - (2d_{\max} + C)(2G_n)^{-A^2/2}$. \square

A.3 Proofs of the technical lemmata

Proof of Lemma 3.2

Proof. Let $\theta \in \text{Int}(\Theta)$ and Y_θ be a real random variable with density $\exp(\theta y - \psi(\theta))F(dy)$. First we prove that for all $k \in \mathbb{N}$, there exists a constant $C_\theta > 0$ which depends on θ such that

$$\mathbb{E}|Y_\theta|^k \leq k!C_\theta^k.$$

The k^{th} absolute moment of Y_θ is the k^{th} derivate of $H_\theta := \mathbb{E}(e^{s|Y_\theta|})$ at 0. Let $s \in \mathbb{C}$ be given. We have

$$H_\theta(s) = \frac{M_+(s + \theta) + M_-(\theta - s)}{M(\theta)}$$

where we define M_+ and M_- as

$$M_+ : \begin{cases} \mathbb{C} & \longrightarrow \mathbb{C} \\ z & \longmapsto \int_{y \geq 0} e^{yz} F(dy) \end{cases}$$

and

$$M_- : \begin{cases} \mathbb{C} & \longrightarrow \mathbb{C} \\ z & \longmapsto \int_{y < 0} e^{yz} F(dy). \end{cases}$$

M is analytic on $\Omega_\Theta := \{z \in \mathbb{C} : \text{Re}(z) \in \text{Int}(\Theta)\}$ and so does M_+ and M_- . Therefore $H_\theta : s \mapsto \mathbb{E}(e^{s|Y_\theta|})$ is analytic on $U_\theta := \{s \in \mathbb{C} : \text{Re}(s + \theta) \in \text{Int}(\Theta) \text{ and } \text{Re}(\theta - s) \in \text{Int}(\Theta)\}$. Since $\theta \in \text{Int}(\Theta)$, H_θ is analytic at the point 0 and hence the function is also holomorphic in a neighborhood of 0. We recall the following result for analytic functions (see [28]).

Theorem: *If f is holomorphic on a domain Ω of \mathbb{C} then $f \in C^\infty(\Omega)$ and if in addition Ω is simply connected then for all contour γ around $z \in \Omega$ we have*

$$f^{(n)}(z) = \frac{n!}{2i\pi} \int_\gamma \frac{f(v)}{(v - z)^{n+1}} dv.$$

Using the previous theorem on H_θ at 0 and taking γ as a circle with radius R centered in 0 (such that H_θ is holomorphic on $D(0, R)$, of course R depends on θ) we obtain that for all $k \in \mathbb{N}$

$$|H_\theta^{(k)}(0)| \leq \frac{k!}{2\pi} \left| \int_\gamma \frac{H_\theta(v)}{v^{k+1}} dv \right| \leq \frac{k!}{R^k} \sup_{|z| \leq R} |H_\theta(z)|. \quad (15)$$

and the result follows with $C_\theta := \max\left(1, \frac{1}{R} \sup_{|z| \leq R} |H_\theta(z)|\right)$. Thanks to assumption (H.2) we can apply (15) with $\theta = \beta^{*T} X$ and find that for all $k \in \mathbb{N}$,

$$\mathbb{E}(|Y|^k | X) \leq k! (C_{\beta^{*T} X})^k.$$

Finally (H.1) combined with (H.3) leads to

$$\mathbb{E}(|Y|^k) \leq k! \left(\sup_{|\theta| \leq LB} C_\theta \right)^k$$

and the result follows with $C_{L,B} := \sup_{|\theta| \leq LB} C_\theta$. \square

Proof of Lemma 3.4

Proof. The proof is based on the convexity of the loss function and of the penalty, the main idea of the proof is similar to the one used by Bühlmann and van de Geer [4] for the Lasso to show consistency of the excess of risk. Define $t := \frac{M}{M + \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2}$ and $\tilde{\beta} := t\hat{\beta}_n + (1-t)\beta^*$.

Notice $\sum_{g=1}^{G_n} \sqrt{d_g} \|\tilde{\beta}^g - \beta^{*g}\|_2 \leq M$. By convexity of $\beta \rightarrow l_\psi(\beta)$ and $\beta \rightarrow \|\beta\|_2$ combined with the fact that $\tilde{\beta}_n$ satisfies (7) we find

$$\begin{aligned} & \mathbb{P} \left(l(\tilde{\beta}) - l(\beta^*) \right) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\tilde{\beta}^g\|_2 \\ & \leq (\mathbb{P}_n - \mathbb{P}) \left(l(\beta^*) - l(\tilde{\beta}) \right) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2. \end{aligned}$$

On the event $\mathcal{A} \cap \mathcal{B}$ we have

$$\begin{aligned} & \mathbb{P} \left(l(\tilde{\beta}) - l(\beta^*) \right) + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\tilde{\beta}^g\|_2 \\ & \leq r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\tilde{\beta}^g - \beta^{*g}\|_2 + r_n \frac{\varepsilon_n}{2} + 2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2. \end{aligned}$$

Because $\mathbb{P} \left(l(\tilde{\beta}) - l(\beta^*) \right) \geq 0$, by adding to both sides of the inequality $2r_n \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2$ and by using the triangular inequality we have

$$\sum_{g=1}^{G_n} \sqrt{d_g} \|\tilde{\beta}^g - \beta^{*g}\|_2 \leq \frac{\varepsilon_n}{2} + 4 \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^{*g}\|_2.$$

Therefore, using (H.3), we have

$$\sum_{g=1}^{G_n} \sqrt{d_g} \|\tilde{\beta}^g - \beta^{*g}\|_2 \leq \frac{\varepsilon_n}{2} + 4B = \frac{M}{2}.$$

i.e.

$$t \sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \leq \frac{M}{2}$$

and then the definition of t leads to

$$\sum_{g=1}^{G_n} \sqrt{d_g} \|\hat{\beta}_n^g - \beta^{*g}\|_2 \leq M.$$

□

Proof of Lemma A.2

Proof. We have

$$\begin{aligned} \mathbb{P}(\mathcal{A}^c) & \leq \sum_{g=1}^{G_n} \mathbb{P} \left\{ \left\| \frac{1}{n} \sum_{i=1}^n (Y_i X_i^g - \mathbb{E}(Y X^g)) \right\|_2 > \frac{r_n^2}{4} d_g \right\} \\ & \leq \sum_{g=1}^{G_n} \sum_{j=1}^{d_g} \mathbb{P} \left\{ \frac{1}{n} \left| \sum_{i=1}^n (Y_i X_{i,j}^g - \mathbb{E}(Y X_j^g)) \right| > \frac{r_n}{2} \right\}. \end{aligned} \quad (16)$$

For $j = 1, \dots, d_g$ and $i = 1, \dots, n$, let

$$W_{ij}^g := Y_i X_{i,j}^g - \mathbb{E}(Y_i X_j^g).$$

The random variables $\{W_{ij}^g\}_{i=1, \dots, n}$ are independant, identically distributed and centered and for all $m \geq 2$

$$\mathbb{E} |W_{ij}^g|^m \leq \sum_{k=0}^m \binom{m}{k} \mathbb{E} |Y_i X_{i,j}^g|^k (\mathbb{E} |Y_i X_{i,j}^g|)^{m-k}$$

By using Jensen inequality we obtain

$$\mathbb{E}|W_{ij}^g|^m \leq 2^m \max_{k=1, \dots, m} \left\{ \mathbb{E}|Y_i X_{i,j}|^k \mathbb{E}|Y_i X_{i,j}|^{m-k} \right\}.$$

For all $k \in \mathbb{N}$, by (H.1) and Lemma 3.2 we have

$$\mathbb{E}|Y_i X_{i,j}|^k \leq L^k k! (C_{L,B})^k.$$

Therefore $\mathbb{E}|W_{ij}^g|^m \leq m!(2LC_{L,B})^m$. Hence the conditions are satisfied to apply Bernstein concentration inequality [1] with $K = 2LC_{L,B}$ and $\sigma^2 = 8(LC_{L,B})^2$. Thus we obtain

$$\begin{aligned} & \mathbb{P} \left(\frac{1}{n} \left| \sum_{i=1}^n W_{ij}^g \right| > r_n/2 \right) \\ & \leq 2 \left(\exp \left(\frac{-nr_n}{16LC_{L,B}} \right) + \exp \left(\frac{-nr_n^2}{32(2LC_{L,B})^2} \right) \right). \end{aligned} \quad (17)$$

Finally, from (16) and (17), we deduce that $\mathbb{P}(\mathcal{A}^c)$ is bounded by

$$2d_{\max} G_n \left(\exp \left(\frac{-nr_n}{16LC_{L,B}} \right) + \exp \left(\frac{-nr_n^2}{32(LC_{L,B})^2} \right) \right).$$

Therefore if

$$r_n \geq A^2 16LC_{L,B} \frac{\log(2G_n)}{n} \vee A8\sqrt{2}LC_{L,B} \sqrt{\frac{\log(2G_n)}{n}}$$

with $A > 1$ then

$$\mathbb{P} \{ \mathcal{A}^c \} \leq 2d_{\max} (2G_n)^{1-A^2}.$$

□

Proof of Lemma A.3

Proof. The proof rests on the following Lemma

Lemma A.4. *Let $R > 0$ be given. Define*

$$Z_R := \sup_{\sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq R} \{ |(\mathbb{P}_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta))| \}.$$

If $A \geq 1$ then

$$\mathbb{P} \left(Z_R \geq A5DLR \sqrt{\frac{2 \log 2G_n}{n}} \right) \leq (2G_n)^{-A^2} \quad (18)$$

where $D := \max_{\{|x| \leq L(R+B)\} \cap \Theta} \{ |\psi'(x)| \}$.

Proof. Let β satisfy $\sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq R$. Notice that if we change X_i by X'_i while keeping the others fixed then Z_R is modified of at most $\frac{2}{n} LR \exp(L(R+B))$. To see this let

$$\mathbb{P}_n = \frac{1}{n} \sum_{j=1}^n 1_{X_j, Y_j}$$

and

$$\mathbb{P}'_n = \frac{1}{n} \sum_{j=1, j \neq i}^n 1_{X_j, Y_j} + 1_{X'_i, Y'_i}$$

then we have

$$(\mathbb{P}_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta)) - (\mathbb{P}'_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta))$$

$$\begin{aligned}
&= \frac{1}{n} \left(l_\psi(\beta^*, X_i) - l_\psi(\beta, X_i) - l_\psi(\beta^*, X'_i) + l_\psi(\beta, X'_i) \right) \\
&\leq \frac{1}{n} |\psi'(\tilde{\beta}^T X_i)| |\beta^{*T} X_i - \beta^T X_i| + \frac{1}{n} |\psi'(\tilde{\beta}^T X'_i)| |\beta^{*T} X'_i - \beta^T X'_i|
\end{aligned}$$

with $\tilde{\beta}^T X_i$ which is an intermediate point between $\beta^T X_i$ and $\beta^{*T} X_i$ (using a first order Taylor expansion of the exponential function). Then, using the same argument as for (10), we have

$$|\tilde{\beta}^T X_i| \leq LR + LB.$$

Therefore

$$\begin{aligned}
&(\mathbb{P}_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta)) - (\mathbb{P}'_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta)) \\
&\leq \frac{2}{n} LR \max_{\{|x| \leq L(R+B)\} \cap \Theta} |\psi'(x)| = \frac{2}{n} LRD.
\end{aligned}$$

We can apply McDiarmid's inequality also called the bounded difference inequality.

Theorem. *Let A a set. Assume $g : A^N \rightarrow \mathbb{R}$ is a function that satisfies the bounded difference inequality*

$$\sup_{x_1, \dots, x_n, x'_i \in A} |g(x_1, \dots, x_n) - g(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)| \leq c_i.$$

Let X_1, \dots, X_n be independent random variables all taking values in the set A . Then for all $t > 0$,

$$\mathbb{P}\{g(X_1, \dots, X_n) - \mathbb{E}g(X_1, \dots, X_n) \geq t\} \leq e^{-2t^2 / \sum_{i=1}^n c_i^2}.$$

We can apply McDiarmid's inequality to Z_R and obtain

$$\mathbb{P}(Z_R - \mathbb{E}Z_R \geq u) \leq \exp\left(-\frac{nu^2}{2R^2L^2(D)^2}\right).$$

Therefore if $r_n \geq ADLR\sqrt{\frac{2\log 2G_n}{n}}$ with $A > 0$ then

$$\mathbb{P}(Z_R - \mathbb{E}Z_R \geq r_n) \leq (2G_n)^{-A^2}. \quad (19)$$

Now we have to bound the mean $\mathbb{E}Z_R$.

Lemma A.5.

$$\mathbb{E}Z_R \leq 4RLD\sqrt{\frac{2\log(2G_n)}{n}}.$$

Proof. First let us introduce two theorems. Let X_1, \dots, X_n independent random variables with values in some space \mathcal{X} and \mathcal{F} a class of real-valued functions on \mathcal{X} .

Theorem: Symmetrization theorem [27]. *Let $\epsilon_1, \dots, \epsilon_n$ be Rademacher sequence independent of X_1, \dots, X_n and $f \in \mathcal{F}$. Then*

$$\begin{aligned}
&\mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n \{f(X_i) - \mathbb{E}(f(X_i))\} \right| \right) \\
&\leq 2\mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n \epsilon_i f(X_i) \right| \right).
\end{aligned}$$

Theorem: Contraction principle [12]. *Let x_1, \dots, x_n elements of \mathcal{X} and $\epsilon_1, \dots, \epsilon_n$ be Rademacher sequence. Consider Lipschitz functions g_i . Then for any function f and h in \mathcal{F} , we have*

$$\begin{aligned}
&\mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n \epsilon_i \{g_i(f(x_i)) - g_i(h(x_i))\} \right| \right) \\
&\leq 2\mathbb{E} \left(\sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n \epsilon_i (f(x_i) - h(x_i)) \right| \right)
\end{aligned}$$

Let $\epsilon_1, \dots, \epsilon_n$ a Rademacher sequence independant of X_1, \dots, X_n and let $\mathcal{S}_R := \left\{ \beta \in \mathbb{R}^p : \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq R \right\}$. Then by the Symmetrization theorem and the Contraction theorem (ψ is D -lipschitz on the compact set \mathcal{S}_R) we have

$$\begin{aligned} \mathbb{E}Z_R &\leq 4D\mathbb{E} \left(\sup_{\beta \in \mathcal{S}_R} \frac{1}{n} \sum_{i=1}^n \left| \epsilon_i (\beta^{*T} X_i - \beta^T X_i) \right| \right) \\ &\leq 4DR\mathbb{E} \left(\max_{g \in \{1, \dots, G_n\}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \frac{\|X_i^g\|_2}{\sqrt{d_g}} \right| \right), \end{aligned}$$

by Holder inequality for the last bound. Now we are going to use the following theorem which is a consequence of Hoeffding inequality.

Theorem. Let X_1, \dots, X_n be independent random variables on \mathcal{X} and f_1, \dots, f_n real-valued functions on \mathcal{X} which satisfies for all $j = 1, \dots, p$ and all $i = 1, \dots, n$

$$\mathbb{E}f_j(X_i) = 0, \quad |f_j(X_i)| \leq a_{ij}.$$

Then

$$\mathbb{E} \left(\max_{1 \leq j \leq p} \left| \sum_{i=1}^n f_j(X_i) \right| \right) \leq \sqrt{2 \log(2p)} \max_{1 \leq j \leq p} \sqrt{\sum_{i=1}^n a_{ij}^2}.$$

By applying this theorem we obtain

$$\mathbb{E} \left(\max_{g \in \{1, \dots, G_n\}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \frac{\|X_i^g\|_2}{\sqrt{d_g}} \right| \right) \leq L \sqrt{\frac{2 \log(2G_n)}{n}}.$$

Thus

$$\mathbb{E}Z_R \leq 4RLD \sqrt{\frac{2 \log(2G_n)}{n}}. \quad (20)$$

□

So we can conclude from (19) and (20) that if $A \geq 1$ then

$$\mathbb{P} \left(Z_R \geq A5DLR \sqrt{\frac{2 \log 2G_n}{n}} \right) \leq (2G_n)^{-A^2} \quad (21)$$

for all $R > 0$.

□

Split up

$$\left\{ \beta \in \mathbb{R}^p : \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq M \right\},$$

where $M = 8B + \varepsilon_n$, into two sets which are

$$\begin{aligned} E_1 &= \left\{ \beta : \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq \varepsilon_n \right\}, \\ E_2 &= \left\{ \beta : \varepsilon_n \leq \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq M \right\} \\ &\subseteq \bigcup_{j=1}^{j_n} \left\{ \beta : 2^{j-1} \varepsilon_n < \sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 \leq 2^j \varepsilon_n \right\} \end{aligned}$$

where $j_n := \lceil \log_2(nM) \rceil + 1$ is the smaller integer such that $2^{j_n} \varepsilon_n \geq M$. We recall that

$$\nu_n(\beta, \beta^*) := \frac{(\mathbb{P}_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta))}{\sum_{g=1}^{G_n} \sqrt{d_g} \|\beta^g - \beta^{*g}\|_2 + \varepsilon_n}$$

and to simplify notation let

$$\alpha_n(\beta, \beta^*) := (\mathbb{P}_n - \mathbb{P})(l_\psi(\beta^*) - l_\psi(\beta))$$

and

$$\Phi(t) := \max_{\{|x| \leq t\} \cap \Theta} |\psi'(x)|.$$

Let $A \geq 1$. We recall that $\kappa_n := 17B + \frac{2}{n} = 2M + B$. On the event E_1 ,

$$\begin{aligned} & \mathbb{P} \left(\sup_{\beta \in E_1} |\nu_n(\beta, \beta^*)| \geq A10L\Phi(L\kappa_n) \sqrt{\frac{2 \log(2G_n)}{n}} \right) \\ & \leq \mathbb{P} \left(\sup_{\beta \in E_1} |\alpha_n(\beta, \beta^*)| \geq A10L\Phi(L\kappa_n) \varepsilon_n \sqrt{\frac{2 \log(2G_n)}{n}} \right) \\ & \leq \mathbb{P} \left(\sup_{\beta \in E_1} |\alpha_n(\beta, \beta^*)| \geq A5L\Phi(L(\varepsilon_n + B)) \varepsilon_n \sqrt{\frac{2 \log(2G_n)}{n}} \right) \end{aligned}$$

given that $2M \geq \varepsilon_n$. From Lemma A.4 with $R = \varepsilon_n$ we deduce

$$\begin{aligned} & \mathbb{P} \left(\sup_{\beta \in E_1} |\nu_n(\beta, \beta^*)| \geq A10L\Phi(L\kappa_n) \sqrt{\frac{2 \log(2G_n)}{n}} \right) \\ & \leq (2G_n)^{-A^2}. \end{aligned} \tag{22}$$

On the event E_2 , using the same type of argument as for (22) with $R = 2^j \varepsilon_n$ (given that $2M \geq 2^j \varepsilon_n$) for all $j = 1, \dots, j_n$, we find

$$\begin{aligned} & \mathbb{P} \left(\sup_{\beta \in E_2} |\nu_n(\beta, \beta^*)| \geq A10L\Phi(L\kappa_n) \sqrt{\frac{2 \log(2G_n)}{n}} \right) \\ & \leq j_n (2G_n)^{-A^2}. \end{aligned}$$

Finally we have

$$\leq C' (2G_n)^{-\frac{A^2}{2}} \tag{23}$$

where C' is a constant (because $j_n = \lfloor \log_2(nM) \rfloor + 1$ and $n \ll G_n$) and the result of Lemma A.3 follows from (22) and (23) with $C = 1 + C'$. \square

B Proof of Theorem 4.2

The main step of the proof are the same as for the Lasso.

Proof. By the same arguments as the ones used to prove (8) we have

$$\begin{aligned} & \mathbb{P} \left(l(\beta^*) - l(\hat{\beta}_n) \right) + 2r_n \|\hat{\beta}_n\|_1 + t_n \|\hat{\beta}_n\|_2^2 \\ & \leq (\mathbb{P}_n - \mathbb{P}) \left(l(\beta^*) - l(\hat{\beta}_n) \right) + 2r_n \|\beta\|_1 + t_n \|\beta^*\|_2^2. \end{aligned} \tag{24}$$

The upper bound of $(\mathbb{P}_n - \mathbb{P}) \left(l(\beta^*) - l(\hat{\beta}_n) \right)$, of $(\mathbb{P}_n - \mathbb{P}) \left(l_\psi(\beta^*) - l_\psi(\hat{\beta}_n) \right)$ and the lower bound of $\mathbb{P} \left(l(\beta^*) - l(\hat{\beta}_n) \right)$ remains the same as those presented in the proof of Theorem 3.8 (see Proposition 3.3, Proposition 3.5 and Proposition A.1). Once these three propositions are proved, the rest of the proof is similar to the one for logistic regression presented in [5]. On the event $\mathcal{A} \cap \mathcal{B}$ (which occurs

with probability at least $1 - 2(2p)^{1-A^2} - C'(2p)^{-A^2/2} \geq 1 - C(2p)^{-A^2/2}$ by adding $r_n \|\hat{\beta}_n - \beta^*\|_1$ and $t_n \sum_{j \in I^*} (\beta_j^* - \hat{\beta}_j)^2$ to both sides of the inequality (24), we have

$$\begin{aligned} & r_n \|\hat{\beta}_n - \beta^*\|_1 + \mathbb{P} \left(l(\beta^*) - l(\hat{\beta}_n) \right) + t_n \sum_{j \in I^*} (\beta_j^* - \hat{\beta}_j)^2 \\ & \leq 2r_n \|\hat{\beta}_n - \beta^*\|_1 + 2r_n \|\beta^*\|_1 - 2r_n \|\hat{\beta}_n\|_1 + t_n \sum_{j \in I^*} (\beta_j^* - \hat{\beta}_j)^2 \\ & \quad - t_n \|\hat{\beta}_n\|_2^2 + t_n \|\beta^*\|_2^2 + \frac{r_n}{2} \varepsilon_n \end{aligned} \quad (25)$$

with $\varepsilon_n = \frac{1}{n}$. On one hand, by the same argument as for (9) we get

$$2r_n \|\hat{\beta}_n - \beta^*\|_1 + 2r_n \|\beta^*\|_1 - 2r_n \|\hat{\beta}_n\|_1 \leq 4r_n \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j|$$

and on the other hand

$$\begin{aligned} & t_n \sum_{j \in I^*} (\beta_j^* - \hat{\beta}_j)^2 - t_n \|\hat{\beta}_n\|_2^2 + t_n \|\beta^*\|_2^2 \\ & \leq 2t_n \sum_{j \in I^*} \beta_j^{*2} - 2t_n \sum_{j \in I^*} \beta_j^* \hat{\beta}_j \\ & \leq r_n \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j|. \end{aligned}$$

Therefore inequality (25) can be bounded by

$$\begin{aligned} & r_n \|\hat{\beta}_n - \beta^*\|_1 + \mathbb{P} \left(l(\hat{\beta}_n) - l(\beta^*) \right) + t_n \sum_{j \in I^*} (\beta_j^* - \hat{\beta}_j)^2 \\ & \leq 4r_n \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j| + r_n \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j| + \frac{r_n}{2} \varepsilon_n. \end{aligned}$$

Since $\mathbb{P} \left(l(\beta^*) - l(\hat{\beta}_n) \right) \geq 0$ we have

$$r_n \|\hat{\beta}_n - \beta^*\|_1 \leq 5r_n \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j| + \frac{r_n}{2} \varepsilon_n$$

and then

$$\sum_{j \in I^{*c}} |\beta_j^* - \hat{\beta}_j| \leq 4 \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j| + \frac{\varepsilon_n}{2}.$$

Thus $(\hat{\beta}_n - \beta^*) \in S(4, \frac{\varepsilon_n}{2})$. Using Proposition A.1 in the case of groups of size one we find

$$\begin{aligned} & r_n \|\hat{\beta}_n - \beta^*\|_1 + t_n \sum_{j \in I^*} (\beta_j^* - \hat{\beta}_j)^2 + c_n \mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \\ & \leq 5r_n \sum_{j \in I^*} |\beta_j^* - \hat{\beta}_j| + \frac{r_n}{2} \varepsilon_n, \end{aligned}$$

with $c_n := \min_{\{|x| \leq L(9B + \frac{1}{n})\} \cap \Theta} \left\{ \frac{\psi''(x)}{2} \right\}$. The rest of the proof follows the guidelines of the proof of (13) and (14) and leads to

$$\|\hat{\beta}_n - \beta^*\|_1 \leq \frac{(2.5)^2 r_n s^*}{t_n + c_n k} + \left(1 + \frac{1}{r_n}\right) \frac{\varepsilon_n}{2}.$$

and

$$\mathbb{E} \left(\hat{\beta}_n^T X - \beta^{*T} X \right)^2 \leq \frac{2(2.5)^2}{c_n k (t_n + c_n k)} r_n^2 s^* + \frac{2r_n + 3}{2nc_n}.$$

□

References

- [1] G. Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57:33–45, 1962.
- [2] P.J. Bickel, Y. Ritov, and A.B. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- [3] P.O. Brown, D. Botstein, et al. Exploring the new world of the genome with dna microarrays. *Nature genetics*, 21:33–37, 1999.
- [4] P. Bühlmann and S. Van De Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, 2011.
- [5] F. Bunea. Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1 + \ell_2$ penalization. *Electronic Journal of Statistics*, 2:1153–1194, 2008.
- [6] F. Bunea, A.B. Tsybakov, and M.H. Wegkamp. Aggregation for gaussian regression. *The Annals of Statistics*, 35(4):1674–1697, 2007.
- [7] D.J. Duggan, M. Bittner, Y. Chen, P. Meltzer, J.M. Trent, et al. Expression profiling using cdna microarrays. *Nature genetics*, 21:10–14, 1999.
- [8] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.
- [9] R.A. Heller, M. Schena, A. Chai, D. Shalon, T. Bedilion, J. Gilmore, D.E. Woolley, and R.W. Davis. Discovery and analysis of inflammatory disease-related genes using cdna microarrays. *Proceedings of the National Academy of Sciences*, 94(6):2150–2155, 1997.
- [10] Junzhou Huang and Tong Zhang. The benefit of group sparsity. *The Annals of Statistics*, 38(4):1978–2004, 2010.
- [11] J. Katriel. On a generalized recurrence for bell numbers. *Journal of Integer Sequences*, 11(2):3, 2008.
- [12] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer, 1991.
- [13] J.M. Loubes and S. Van De Geer. Adaptive estimation with soft thresholding penalties. *Statistica Neerlandica*, 56(4):453–478, 2002.
- [14] K. Lounici, M. Pontil, S. Van De Geer, and A.B. Tsybakov. Oracle inequalities and optimal inference under group sparsity. *The Annals of Statistics*, 39(4):2164–2204, 2011.
- [15] Karim Lounici, Massimiliano Pontil, Alexandre B Tsybakov, and Sara van de Geer. Taking advantage of sparsity in multi-task learning. *arXiv preprint arXiv:0903.1468*, 2009.
- [16] P. McCullagh and J.A. Nelder. Generalized linear models. monographs on statistics and applied probability 37. *Chapman Hall, London*, 1989.
- [17] L. Meier, S. Van De Geer, and P. Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- [18] Y. Nardi and A. Rinaldo. On the asymptotic properties of the group lasso estimator for linear models. *Electronic Journal of Statistics*, 2:605–633, 2008.
- [19] Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *NIPS Conference, Vancouver, Canada*.
- [20] Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.

- [21] P.J. Park, L. Tian, and I.S. Kohane. Linking gene expression data with patient survival times using partial least squares. *Bioinformatics*, 18:S120–S127, 2002.
- [22] J. Quackenbush. Microarray analysis and tumor classification. *New England Journal of Medicine*, 354(23):2463–2472, 2006.
- [23] M.R. Segal, K.D. Dahlquist, and B.R. Conklin. Regression approaches for microarray data analysis. *Journal of Computational Biology*, 10(6):961–980, 2003.
- [24] B. Tarigan and S.A. Van De Geer. Classifiers of support vector machine type with ℓ_1 complexity regularization. *Bernoulli*, 12(6):1045–1076, 2006.
- [25] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [26] S.A. Van de Geer. High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, 36(2):614–645, 2008.
- [27] A. Van der Vaart and J. Wellner. *Weak convergence and empirical processes: with applications to statistics*. Springer, 1996.
- [28] C. Wagschal. *Fonctions holomorphes, équations différentielles: exercices corrigés*. Hermann, 2003.
- [29] F. Wei and J. Huang. Consistent group selection in high-dimensional linear regression. *Bernoulli: official journal of the Bernoulli Society for Mathematical Statistics and Probability*, 16(4):1369, 2010.
- [30] M. West, C. Blanchette, H. Dressman, E. Huang, S. Ishida, R. Spang, H. Zuzan, J.A. Olson, J.R. Marks, and J.R. Nevins. Predicting the clinical status of human breast cancer by using gene expression profiles. *Proceedings of the National Academy of Sciences*, 98(20):11462–11467, 2001.
- [31] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2005.
- [32] Cun-Hui Zhang and Jian Huang. The sparsity and bias of the lasso selection in high-dimensional linear regression. *The Annals of Statistics*, 36(4):1567–1594, 2008.
- [33] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.