



**HAL**  
open science

# An overview of variable selection procedures using regularization paths in high-dimensional Gaussian linear regression

Perrine Lacroix, Mélina Gallopin, Marie-Laure Martin

► **To cite this version:**

Perrine Lacroix, Mélina Gallopin, Marie-Laure Martin. An overview of variable selection procedures using regularization paths in high-dimensional Gaussian linear regression. 2024. hal-03366851

**HAL Id: hal-03366851**

**<https://hal.science/hal-03366851v1>**

Preprint submitted on 14 Feb 2024

**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.

---

# AN OVERVIEW OF VARIABLE SELECTION PROCEDURES USING REGULARIZATION PATHS IN HIGH-DIMENSIONAL GAUSSIAN LINEAR REGRESSION

---

## **Perrine Lacroix**

Université Paris-Saclay, CNRS, Laboratoire de Mathématiques d'Orsay, 91405, Orsay, France.  
Université Paris-Saclay, CNRS, INRAE, Université Evry, Institute of Plant Sciences Paris-Saclay (IPS2),  
91190, Gif sur Yvette, France.  
Université Paris Cité, Institute of Plant Sciences Paris-Saclay (IPS2), 91190, Gif sur Yvette, France.  
perrine.lacroix@universite-paris-saclay.fr

## **Mélina Gallopin**

Université Paris-Saclay, CNRS, CEA, Institut de Biologie Intégrative de la Cellule (UMR 9198), 91405, Orsay, France  
melina.gallopin@universite-paris-saclay.fr

## **Marie-Laure Martin**

Université Paris-Saclay, CNRS, INRAE, Université Evry, Institute of Plant Sciences Paris-Saclay (IPS2),  
91190, Gif sur Yvette, France.  
Université Paris Cité, Institute of Plant Sciences Paris-Saclay (IPS2), 91190, Gif sur Yvette, France.  
Université Paris-Saclay, AgroParisTech, INRAE, UMR MIA Paris-Saclay, 91120, Palaiseau, France  
marie\_laure.martin@agroparistech.fr

## **ABSTRACT**

Current high-throughput technologies provide a large amount of variables to describe a phenomenon. Only a few variables are generally sufficient to answer the question. Identify them in a high-dimensional Gaussian linear regression model is the one of the most-used statistical methods. In this article, we describe step-by-step the variable selection procedures built upon regularization paths. Regularization paths are obtained by combining a regularization function and an algorithm. Then, they are combined either with a model selection procedure using penalty functions or with a sampling strategy to obtain the final selected variables. We perform a comparison study by considering three simulation settings with various dependency structures on variables. In all the settings, we evaluate (i) the ability to discriminate between the active variables and the non-active variables along the regularization path (pROC-AUC), (ii) the prediction performance of the selected variable subset (MSE) and (iii) the relevance of the selected variables (recall, specificity, FDR). From the results, we provide recommendations on strategies to be favored depending on the characteristics of the problem at hand. We obtain that the regularization function Elastic-net provides most of the time better results than the  $\ell_1$  one and the lars algorithm has to be privileged as the GD one. ESCV provides the best prediction performances. Bolasso and the knockoffs method are judicious choices to limit the selection of non-active variables while ensuring selection of enough active variables. Conversely, the data-driven penalties considered in this review are not to be favored. As for Tigress and LinSelect, they are conservative methods.

**Keywords** Variable selection · Gaussian linear regression · High-dimension · Regularization path · Comparison study

## 1 Introduction

Recent scientific advances allow us to have access to large-scale data: the size of the data sets is exploding, as well as the complexity of each of them. For instance, in genomics, to describe the molecular activities, microarrays and RNA-sequencing technologies provide a quantification of the expression of all the genes simultaneously, making it easier to study their interactions. Studies on genetic associations are diversifying, considering a wide range of phenotypes, including gene expression or proteomic and metabolomic data. Other fields are concerned with such massive data sets, such as medicine with the development of imaging, treatment or disease monitoring, or market strategy. In a statistical point of view, the number of parameters to estimate explodes and reduction of dimension is required to select only relevant variables and summarize the redundant information for a given model. In this review, we focus on the variable selection procedures in high-dimensional linear Gaussian regression models. The considered dataset with a number of variables  $p$  close to or slightly higher than the number of observations  $n$  is a real challenge since it hampers the use of the traditional estimation methods. A regularization of the cost function is required so that only a subset of variables is selected to explain the response variable.

In most reviews on variable selection in high-dimensional Gaussian linear regression, a focus is done on the construction of the regularization path. It is based on the minimization of a cost function penalized by a regularization function and provides an order on variables. [1] provides a meticulous theoretical analysis of the  $\ell_1$  regularization function. In particular, for a given number of active variables, the author discusses the choice of the number of observations to ensure asymptotic properties to recover these active variables. [2] compared several regularization functions in a simulation study by using semi-real datasets. In their simulation design, they considered several numbers of observations, of variables and of active variables. They also modified the signal-to-noise ratio and considered two scenarios of variable correlations. The results of the different regularization functions are inspected with ROC curves and partial ROC curves when  $0.5 \times n$  and  $0.9 \times n$  variables are selected. [3] compared a large set of regularization functions with a simulation design similar to [2]. They evaluated both prediction and variable identification but the main difference with our investigations is that the only considered model selection procedure applied from the regularization path is the cross-validation one. Finally some reviews considered different contexts. [4] were interested in robust variable selection strategies when heavy-tailed errors and outliers in response variables exist. They discussed the different steps from the modification of the least squares function to the choice of the parameters for the model selection through a presentation of algorithms accounting for outliers. [5] considered a variety of models from survival models to generalized linear models, frequently used in biomedical research. [6] considered a wide range of model structures (linear, grouped, additive, partially linear and non-parametric) and discussed three main categories of algorithms for the variable selection.

Our review distinguishes itself from the previous ones since we propose an evaluation of both construction of regularization paths and choice of the final selected variables. This leads to 33 combinations. Moreover, for model selection procedure, we add in this review non-asymptotic methods which are generally not considered. To construct the regularization path, we test two regularization functions (Lasso [7] and Elastic-Net [8]) combined with two algorithms (LARS [9] and gradient descent algorithm [10]). Each regularization path provides a collection of variable subsets. To choose one of them, we compare model selection and variable identification approaches. On the one hand, the model selection uses penalization criteria of the least squares (eBIC [11], data-driven calibration strategies [12, 13, 14, 15] and LinSelect [16, 17]). On the other hand, the variable identification methods (ESCV [18], Bolasso [19], Stability Selection [20], Tigress [21] and the knockoffs method [22]) use sampling strategies to stabilize the selected variable subset while limiting the selection of non-active variables. Methods based on multiple testing procedures [23] and Bayesian approaches are not included in this review. We refer the readers to [24] for an empirical comparison of frequentist and Bayesian points of view. Lastly, we assume no prior knowledge between interactions, spatial localization and chronological information and refer to [25, 26] for such approaches.

After a description of the methods and some of their theoretical properties, we compare them in a simulation study by considering three settings. In the first one, the variables are independent and are drawn from a Gaussian distribution. It allows a comparison in the theoretical framework used to develop them. In the second setting, two structures of the correlation between variables are considered to evaluate how dependencies usually observed affect the methods. Observations are generated according to a Gaussian linear model, the most favorable case where assumptions broadly hold. Finally, the third setting mimics the biological complexity of transcription factor regulations. Observations are generated using the *FRANK* algorithm [27].

In these three settings, performances of the methods are evaluated for their prediction performance and for their ability to identify the active variables. To discriminate active variables to the others, we use the pROC-AUC metric. We use the mean squared errors (MSE) to measure the prediction performance, and the recall, specificity and false discovery rate (FDR) metrics to assess the quality of the selected variables in terms of active variables. As [2, 3], we notice that there is no unambiguous winner among all the studied approaches. Our goal is to provide recommendations for

a judicious choice of a method according to the application. In particular, Elastic-Net should be preferred to Lasso for the regularization function, as well as the lars algorithm. Moreover, to ensure a prediction ability, ESCV and the knockoffs seem to be the most judicious choices. If the goal is to recover the active variables, ESCV and eBIC are preferable whereas Bolasso, the knockoffs and LinSelect should be privileged to limit the non-active variables in the selected subset.

The rest of the paper is organized as follows. Section 2 describes the statistical framework and all the methods we compare in this review. Section 3 presents the simulation settings, technical aspects about the implementation of methods and the evaluation metrics we consider. Section 4 is devoted to all the results. Section 5 offers practical recommendations to choice the best methods. Lastly, a discussion is provided in Section 6.

## 2 Methods

### 2.1 Statistical framework

For the sequel, the norms  $|\cdot|_0$ ,  $|\cdot|_1$  and  $\|\cdot\|$  are respectively the usual norms 0, 1 and 2 on  $\mathbb{R}^q$ , for any  $q \in \mathbb{N}^*$ . It means that for a vector  $\beta \in \mathbb{R}^q$ ,

$$|\beta|_0 = \sum_{j=1}^q \mathbf{1}_{\{\beta_j \neq 0\}}; \quad |\beta|_1 = \sum_{j=1}^q |\beta_j|; \quad \|\beta\| = \sqrt{\sum_{j=1}^q \beta_j^2}.$$

We consider a Gaussian linear regression model where the response variable  $Y$  is explained by a linear combination of  $p$  variables  $X = (X_1, \dots, X_p)$ :

$$Y = X\beta^* + \varepsilon.$$

The parameter  $\beta^* \in \mathbb{R}^p$  is the vector of unknown parameters and  $\varepsilon$  follows a centered Gaussian distribution with an unknown variance denoted  $\sigma^2$ .

To estimate the parameters  $\beta^*$  and  $\sigma^2$ , we observe  $y_i \in \mathbb{R}$  and the variables  $(x_{i1}, \dots, x_{ip}) \in \mathbb{R}^p$  for  $i \in \{1, \dots, n\}$ . We suppose that observations are independent. We consider the high-dimensional framework with  $p \sim n$  or  $p \gg n$ , preventing the traditional least squares estimation. In this context, we assume that only a small number of variables among the  $p$  ones explains the response variable. These variables are associated to a non-zero coefficient in  $\beta^*$  and are named active variables. Under this sparsity assumption, the target is for  $t > 0$ :

$$\min_{\beta \in \mathbb{R}^p: |\beta|_0 \leq t} \|Y - X\beta\|^2,$$

where its associated Lagrangian form is for  $\lambda > 0$ :

$$\min_{\beta \in \mathbb{R}^p} \{ \|Y - X\beta\|^2 + \lambda |\beta|_0 \}. \quad (1)$$

The proof of the equivalence and the link between  $t$  and  $\lambda$  are provided in [7].

Determining the hyperparameter  $\lambda$  is one of the major issues and the challenge lies in its calibration to adjust a trade-off between sparsity and a good linear adjustment. A large value of  $\lambda$  provides a small subset of variables (assumption of sparsity satisfied) but it might correspond to a linear adjustment far from the response variable. A small value of  $\lambda$  provides a large subset of variables (assumption of sparsity not satisfied) but it might correspond to a linear adjustment close to the response variable. Moreover, the criterion (1) being non-convex, the existence and the uniqueness of the solution are not guaranteed. So, as presented in [5], Equation (1) is replaced with the optimization problem:

$$\min_{\beta \in \mathbb{R}^p} \{ \|Y - X\beta\|^2 + \lambda F(\beta) \}, \quad (2)$$

where  $F$  is a continuous and convex regularization function satisfying the existence of a minimum for any  $\lambda$ .

### 2.2 Regularization functions

Several regularization functions exist and we present the most used. The first one is the Lasso regularization [7] with

$$F(\beta) = |\beta|_1.$$

If  $\lambda$  is well chosen, it provides a consistent estimator of  $\beta^*$ . This procedure achieves the best trade-off between regularity (convexity, reasonable computationally solving) and sparsity for independent variables. However, when some variables are correlated, Lasso tends to select randomly only one of them rather than selecting none or all of them. A solution to

this problem is the Adaptive Lasso procedure where each variable is weighted with respect to an initial estimator [28]. Another well-known regularization function is the Ridge regularization [29]:

$$F(\beta) = \|\beta\|^2.$$

In addition to taking variable dependencies into account, Ridge provides a strictly convex and derivable optimization problem with an explicit estimator of  $\beta^*$ . However, this estimator is not sparse. To combine sparsity and correlated variables, the Elastic-Net regularization [8] has been proposed:

$$F(\beta) = (1 - \alpha)\|\beta\|_1 + \alpha\|\beta\|^2,$$

where the parameter  $\alpha$  controls the trade-off.

When prior knowledge on variable dependencies is available, there exist other regularization functions, not considered here: the Group Lasso [30], Overlap Group Lasso [31], Hierarchical Group Lasso [32], double sparse Lasso [33] and fused Lasso [34].

### 2.3 Regularization path construction for Lasso and Elastic-Net

The optimization problem (2) has generally no explicit solution and requires a computational approach. A first algorithm is LARS [9] where one variable is added at each step and each step corresponds to a value of  $\lambda$ . A grid  $\Lambda$  of  $\lambda$  is then naturally provided. Briefly speaking, the first subset contains the variable  $X_j$  which has the largest absolute correlation with  $Y$ . The second subset contains exactly two variables:  $X_j$  and the variable which is the most correlated with the residuals of the regression of  $Y$  on  $X_j$ . LARS provides an exact solution of the optimization problem with nested subsets. It is an important property for theoretical considerations. A second algorithm is based on the gradient descent method [10]. This algorithm constructs a regular grid  $\Lambda$  of a given size by starting with the largest  $\lambda$  corresponding to the first nonempty variable subset. Then, a variable subset is obtained for each  $\lambda$  of this grid by solving (2) with the cyclic coordinate descent method. In contrast to LARS, the gradient descent method provides a proxy of the optimization problem with independent solutions along the grid. Consequently, the gradient descent method may provide a richer collection with several models of the same dimension.

Whatever the choice of both regularization function and algorithm, a collection of variable subsets  $(m_\lambda)_{\lambda \in \Lambda}$  is obtained at the end of the regularization path construction. Each  $m_\lambda$  is associated with an estimator of  $\beta^*$ , however this estimator being known to be biased [35], it is commonly replaced with the least-squares estimator calculated on the variable subset  $m_\lambda$  and is denoted  $\hat{\beta}_\lambda$  [36]. The number of non-zero coefficients of  $\hat{\beta}_\lambda$  or in an equivalent manner, the number of variables in  $m_\lambda$  is denoted  $D_\lambda$ . The next step consists in selecting a variable subset among the collection either with a model selection criterion or with a variable identification procedure.

### 2.4 Model selection

To select the best subset  $m_{\hat{\lambda}}$ , model selection approaches consist of minimizing a penalized loss function in  $\lambda \in \Lambda$ :

$$\gamma(m_\lambda) + \text{pen}(n, p, D_\lambda). \quad (3)$$

The loss function  $\gamma(m_\lambda)$ , quantifying the quality of the model fit, is either the least-squares function  $\|Y - X\hat{\beta}_\lambda\|^2$  or the deviance  $-2\log(L(Y, X; \hat{\beta}_\lambda, \hat{\sigma}_\lambda^2))$ , where  $L$  is the likelihood function calculated with  $\hat{\beta}_\lambda$  and  $\hat{\sigma}_\lambda^2$ , the empirical estimators associated to the  $m_\lambda$ . The penalty function  $\text{pen}(n, p, D_\lambda)$  accounts for the model complexity and the characteristics of the sample: higher the penalty values, smaller the number of selected variables and farther the linear combination  $X\hat{\beta}_\lambda$  to the response variable  $Y$ .

**Asymptotic criteria.** The first criteria are asymptotic: their properties are verified when the sample size  $n$  tends to infinity. In this review, we focus on the more recent asymptotic criterion, called eBIC [11], used to get a consistent estimator by penalizing the deviance by:

$$\text{pen}_{\text{eBIC}}(n, p, D_\lambda) = D_\lambda \log(n) + 2\delta \log\left(\binom{p}{D_\lambda}\right), \quad (4)$$

where  $\delta$  is a value in  $[0, 1]$ .

**Non-asymptotic criteria.** In a practical consideration, having guarantees for  $n$  going to infinity has no sense [17] and applying criteria with properties confirmed for any fixed sample  $n$  size is more relevant. Introduced by Birgé and Massart [37], the goal of non-asymptotic criteria is to achieve the risk oracle:

$$\inf_{\lambda \in \Lambda} \mathbb{E}[\|X\beta^* - X\hat{\beta}_\lambda\|^2],$$

and instead of getting asymptotic equality of the kind

$$\mathbb{P} \left( \lim_{n \rightarrow +\infty} \frac{\mathbb{E}[ \|X\beta^* - X\hat{\beta}_\lambda\|^2 ]}{\inf_{\lambda \in \Lambda} \mathbb{E}[ \|X\beta^* - X\hat{\beta}_\lambda\|^2 ]} = 1 \right) = 1,$$

they get an inequality holding for any value of  $n$ :

$$\mathbb{E}[ \|X\beta^* - X\hat{\beta}_\lambda\|^2 ] \leq C_n \inf_{\lambda \in \Lambda} \{ \mathbb{E}[ \|X\beta^* - X\hat{\beta}_\lambda\|^2 ] \} + R_n, \quad (5)$$

where  $C_n \approx 1$  at least for  $n$  large and  $R_n$  is small comparable to the risk oracle. The selected model is the minimizer of Equation (3) where the loss function is the least-squares function and two penalty functions, which do not require the knowledge of the variance, are available.

The first penalty is a data-driven penalty [12]:

$$\text{pen}_{\text{Data-driven}}(n, p, D_\lambda) = 2\kappa D_\lambda \left( 2.5 + \log \left( \frac{p}{D_\lambda} \right) \right), \quad (6)$$

where the constant 2.5 has been proposed in a context of changepoint detection in a signal [13]. The constant  $\kappa$  is calibrated from the sample. For that, two strategies are proposed. The first one is the slope heuristics: assuming that the least-squares function is linear in  $D_\lambda \left( 2.5 + \log \left( \frac{p}{D_\lambda} \right) \right)$  as soon as  $D_\lambda$  is large enough (see Figure 2 of [14]), the constant  $\kappa$  is then equal to the estimated slope. The second strategy is the dimension jump: assuming the existence of  $\kappa^*$  such that for all the values smaller than  $\kappa^*$ , the associated model has a very high dimension, whereas for all the values greater than  $\kappa^*$ , the associated model has a reasonable dimension (see Figure 1 of [14]), the constant  $\kappa$  is then equal to the estimated  $\kappa^*$ . For more practical and theoretical details, we refer the reader to [14] and to the survey [15].

The second penalty function is LinSelect proposed in [16] and generalized for a high dimensional context in [17]. It is built from the empirical estimator of the variance onto each  $m_\lambda$ :

$$\text{pen}_{\text{LinSelect}}(n, p, D_\lambda) = 1.1 \times \frac{n - D_\lambda}{n - D_\lambda - 1} \Psi(D_\lambda + 1, n - D_\lambda - 1, e^{-L_\lambda}), \quad (7)$$

where the  $L_\lambda$  are weights satisfying some properties and the function  $\Psi[D, N, q]$  is the unique solution of the equation:

$$\phi[D, N, \Psi(D, N, q)] = q,$$

where  $\phi[D, N, x]$  is defined for  $x \geq 0$ :

$$\phi[D, N, x] = \frac{1}{D} \mathbb{E} \left[ \max \left( 0, \chi_D^2 - x \frac{\chi_N^2}{N} \right) \right],$$

for  $\chi_D^2$  and  $\chi_N^2$  two independent  $\chi^2$  random variables with degrees of freedom  $D$  and  $N$  respectively.

## 2.5 Variable identification

The high-dimensional framework usually leads to unstable results: addition, suppression, or modification of some observations could radically change the selected variable subset. For prediction, different sets of variables can give the same prediction performances. However, when the objective is the identification of the active variables, this instability is a drawback.

To circumvent this sampling uncertainty, the idea is to work with perturbed datasets generated from the original sample. Cross-validation [38, 39] is commonly proposed. It consists in splitting  $K$  times the original sample into a training set and a testing set. The training set is used to calculate an estimator  $\hat{\beta}_\lambda^k$  and the testing set is used to evaluate the mean squared error. The selected model is the one which minimizes the mean squared error in  $\lambda$ . Applying cross-validation in a high-dimensional context is computationally expensive and known to be unstable. An alternative is ESCV [18] which estimates the instability along the regularization path with the  $K$  perturbed datasets and selects the variable subset which minimizes the instability in  $\lambda$ :

$$\frac{\frac{1}{K} \sum_{k=1}^K \|X\hat{\beta}_\lambda^k - \frac{1}{K} \sum_{\ell=1}^K X\hat{\beta}_\lambda^\ell\|^2}{\| \frac{1}{K} \sum_{\ell=1}^K X\hat{\beta}_\lambda^\ell \|^2}.$$

Sampling strategy is also a solution. Two widely used approaches are Bolasso [19] and Stability Selection [20]. They mainly differ in their sampling strategy: Bolasso generates datasets of  $n$  data uniformly chosen with replacement among the original sample, whereas Stability Selection generates datasets of  $\lfloor \frac{n}{2} \rfloor$  distinct data randomly chosen and includes also the complement of each generated dataset in the sampling strategy to limit the subsampling effects. In Stability Selection, they also propose a random perturbation in the Lasso regularization:

$$F(\beta) = \sum_{j=1}^p \frac{|\beta_j|}{w_j},$$

where  $w_j \sim \mathcal{U}([\theta, 1])$  with  $\theta > 0$ . Sampling strategies get an occurrence frequency of each variable and those having the highest occurrence frequencies are retained to constitute the final variable subset. Tigress method [21] modifies the calculation of the occurrence frequency by averaging on the grid.

The last type of variable identification method is the knockoffs method [22], which controls the False Discovery Rate (FDR). This method starts with the construction of a matrix  $\tilde{X}$  such that  $\tilde{X}$  and  $X$  have the same covariance structure with  $\tilde{X}_j$  the least correlated to  $X_j$ . It is done through linear algebra tools [22, 40]. Then, a regularization path is constructed on the augmented matrix  $X\tilde{X}$  of size  $n \times 2p$  where the active variables are expected to be selected very earlier than their copy. Let denote

$$W_j = \max \left( Z_j, \tilde{Z}_j \right) \times \text{sign} \left( Z_j - \tilde{Z}_j \right), \quad (8)$$

where  $Z_j$  and  $\tilde{Z}_j$  correspond to the largest  $\lambda$  for which  $X_j$  and  $\tilde{X}_j$  are selected respectively. A positive value of  $W_j$  states that  $X_j$  is selected before its copy  $\tilde{X}_j$  and a large positive value indicates that  $X_j$  is selected rapidly. Let  $q$  be the target FDR, the final variable subset is composed by the  $X_j$  such that  $W_j \geq T$  with:

$$T = \min \left\{ t \in \{|W_j|, j = 1, \dots, p\} \setminus \{0\}, \frac{1 + \#\{j : W_j \leq -t\}}{\min(1, \#\{j : W_j \geq t\})} \leq q \right\}.$$

### 3 Comparison study

#### 3.1 Three simulation settings

The design of the simulation study is composed of three simulation settings to study the behavior of the methods with respect to the dependency structures between the variables.

**Simulation under independent design.** This is the simplest setting where the high-dimensional framework is the single handicap [41, 3]. The matrix  $X$  is simulated by concatenation of  $p$  independent standard Gaussian vectors of size  $n$ . The number of non-zero coefficients of the vector  $\beta^*$  is generated from a uniform variable on integers between 10 and 15. Theirs values are generated from a uniform distribution between 0.5 and 2 and the response variable  $Y$  is defined by  $Y = X\beta^* + \varepsilon$ , where  $\varepsilon \sim \mathcal{N}(0, I_n)$ . The independent design is a benchmark in this comparison study since the statistical properties of each method must be verified on the independent structure.

**Simulation under a Gaussian graphical model.** An equivalence exists between the network inference by Gaussian Graphical model and support estimation in Gaussian linear regressions [42]. An edge between the nodes  $i$  and  $j$  in the network states that  $X_i$  and  $X_j$  are dependent and either  $X_j$  is on the support of the regression when  $X_i$  is the response variable; or  $X_i$  is on the support of the regression when  $X_j$  is the response variable. In this direction, datasets with dependency structure are simulated using a Gaussian Graphical model. More precisely, a sample of size  $n$  is generated from a  $(p + 1)$  multivariate centered Gaussian distribution with covariance matrix  $\Sigma$ , where the dependency structure is encoded in the precision matrix  $\Sigma^{-1}$  [43, 44]. The response variable  $Y$  is chosen as a column of the  $(p + 1)$  multivariate centered Gaussian and the remaining columns constitute the matrix  $X$  of size  $n \times p$ . It differs from the previous papers [45, 46, 47, 3], where the response variable is simulated once the matrix  $X$  is fixed, but our simulation choice has been motivated by applications such that regulatory network inference where transcription factors can be response variables and also active variables.

We consider two graph patterns:

- *Cluster*: the precision matrix is simulated as a block diagonal matrix with  $B$  blocks of equal size. The response variable  $Y$  is defined as the first variable.
- *Scale-free*: a few variables have a lot of neighbors in the network while all the others have a few ones. We consider two response variables corresponding to the variables having the highest and the smallest number of neighbors. These simulation designs are called *scale-free-max* and *scale-free-min* respectively.

**Simulation under a dynamical process.** It is the most realistic setting, based on the algorithm *FRANK* [27] which simulates large networks with characteristics of gene regulatory networks. In this algorithm, variables are categorized into a set of transcription factors that activate or inhibit a set of target genes and the *FRANK* data are generated from a dynamic process and deviate from the statistical model assumptions, especially the Gaussian distribution. We use *FRANK* with only transcription factor variables in order to compare the results with those from the other settings. We consider as response variables the variables having the highest and the smallest number of neighbors. These simulation designs are called *FRANK-max* and *FRANK-min* respectively.

**Simulation parameters.** For all the settings, we set  $n = 150$  and  $p = 199$ . We generate 100 samples of size  $2n$  to create a training set of size  $n$  for the estimation and a validation set of size  $n$  to evaluate the methods. For each sample, before being used, the variables are centered and scaled. To generate data from Gaussian graphical model, we use the function *huge.generator* from the R package *huge* (version 1.3.4.1). For the cluster design, the block number  $B$  equals 5 and the probability of connection within a component is set to the default value 0.3. For the *FRANK* algorithm, we use the online version available on the website <https://m2sb.org/?page=FRANK> with  $p + 1$  transcription factors and  $2n$  observations. The number of eigenvalues of the matrix on the unit circle is fixed to 2 and the minimum and maximum of sparsity are set to 1 and 50. Other parameters are set to default values. In subsection 4.8, we study the impact of the high-dimension by increasing the initial dataset from the independent design and from the Gaussian graphical model design to sizes  $n = 300, 600, 1200$ .

### 3.2 Investigated methods and their parameters

A total of 16 methods of model selection are defined by the combination of a regularization function (Lasso or Elastic-net) with an algorithm (LARS or the gradient descent method) and a penalty function (eBIC, LinSelect or the 2 data-driven penalties). A total of 17 methods of variable identification are defined: when the sampling strategy is performed before the definition of the grids, 8 methods are defined by the combination of the sampling strategy (Bolasso or Stability Selection) with a regularization function (Lasso or Elastic-net) and an algorithm (LARS or the gradient descent method). When the grid  $\Lambda$  is fixed, the sampling strategy is performed for each  $\lambda$  of the grid, which implies using the gradient descent algorithm, hence 4 methods are defined. Furthermore, we include Tigress, the knockoffs method and ESCV. For the last two, a gradient descent algorithm is used with either Lasso or Elastic-net.

For the LARS algorithm, we use the function *enet* of the R package *elasticnet* (version 1.1.1). The maximal number of steps to define the grid size is the default value  $50 \times \min(p, n - 1)$ . For the gradient descent method, we use the function *glmnet* of the R package *glmnet* (version 3.0) and set the grid size at 1000. Both functions propose the Lasso and elastic-net regularization functions. We set  $\alpha = 0.5$  for Elastic-Net regularization.

To perform model selection, eBIC is implemented with  $\delta = 1$ . LinSelect is implemented in the function *tuneLasso* of the R package *LINselect* (version 1.1.3). The data-driven penalties are calculated by using the function *capushe* of the R package *capushe* (version 1.1.1). The parameters are set to the default values except for the minimum percentage of points for the plateau selection set to 0.1.

To perform variable identification, the function *escv.glmnet* of the R package *HDCI* (version 1.0.2) is used for the ESCV strategy, with a number of groups  $K = 10$ . Bolasso and Stability Selection are implemented with 100 samples. We do not investigate the impact of the presence of a random perturbation in the LASSO regularization function for Stability Selection and set  $\theta = 1$ . A variable is selected when its occurrence frequency is higher than 0.8. For Tigress, we use the function *tigress* of the R package *tigress* (version 0.1.0) with 50 steps for the LARS algorithm. For the knockoffs method, we use the function *knockoff.filter* with option *create.second\_order* of the R package *knockoff* (version 0.3.2), we calculate the  $W_j$  with the function *stat.lasso\_lambdasmax* and set the FDR to 0.1.

### 3.3 Evaluation metrics

First, we evaluate the performance of regularization path constructions. We use the partial area under the receiver operating characteristic curve (pROC-AUC) where the x-axis is the proportion of selected non-active variables among the non-active variables and the y-axis is the proportion of selected active variables among the active variables. A high value of pROC-AUC states that the regularization path is able to discriminate the active variables from the others. Since the lengths of the regularization paths differ according to the chosen regularization and algorithm, the pROC-AUC are calculated by truncating the values of the x-axis at the largest value common to all the regularization paths, to compare them fairly.



Second, we evaluate the prediction performance of the methods by calculating the mean squared errors (MSE) on each validation set  $(\tilde{Y}, \tilde{X})$ :

$$\frac{1}{n} \sum_{i=1}^n \left( \tilde{Y}_i - (\tilde{X} \hat{\beta}_{\tilde{\lambda}})_i \right)^2, \quad (9)$$

where  $\hat{\beta}_{\tilde{\lambda}}$  is the estimator of  $\beta^*$  calculated on the associated training set. As data are centered and scaled, a MSE value lower than 1 means that the method has a prediction ability: the selected variables predict  $Y$  better than the empty set.

Finally, we evaluate the variable identification by using three metrics: the recall (the proportion of the selected active variables among the active variables), the specificity (the proportion of the non-active variables not selected among the non-active variables) and the false discovery proportion (the proportion of selected non-active variables among the selected variables). By averaging on the 100 simulated samples, the false discovery proportion becomes the False Discovery Rate (FDR). While the recall and the specificity respectively control the number of active variables which are selected and the number of non-active variables which are non selected, the False Discovery Rate evaluates the quality of the selected variables subset. So, considering each of these metrics when assessing method performance yields different information. Since the objective is to limit the selection of non-active variables while selecting as many active variables as possible, recall and specificity are expected to be close to 1 while the FDR is expected to be low or slightly smaller than the threshold fixed by the knockoffs method.

## 4 Results

The first six subsections are dedicated to the results obtained for the independent design and when samples are generated from a Gaussian graphical model when  $n = 150$ . In our simulation settings, the sizes of the estimated supports of  $\beta^*$  do not take extreme values: the sparsity hypothesis is respected (we refer to Table 1 summarizing the active variable number per simulation setting). The mean of the active variable numbers for the independent design is close to the one for the cluster design (around 12, relatively small). As the independent design is a benchmark in this study, comparison of results from independent and cluster designs allows evaluating the impact of the presence of a dependency structure on variables. The scale-free design allows us to investigate the method behavior with respect to the number of active variables. In scale-free-min, the support size always equals 1. In contrast, for scale-free-max, the support size is 31.41 on average with a high standard deviation (9.70) to fully investigate the methods. Subsection 4.7 is devoted to the FRANK data. As these datasets deviate from the statistical model assumptions, we analyse performances of the methods in an independent way. Lastly, we discuss the behavior of the methods for different values of  $n$  to investigate the impact of the high-dimension.

For the sequel, the notation *ind* is the diminutive of the independent design. The notations *GD*, *E-Net* and  $\ell_1$  denote respectively the gradient descent algorithm, the Elastic-Net regularization function and the Lasso regularization function. The slope heuristics method and the dimension jump are named *slope* and *jump* respectively. Bolasso and Stability Selection are named *bol* and *ss* respectively. Lastly, *grid* and *sub* denote respectively the strategy when grids are first generated and the strategy when samples are first generated. Throughout this section, we mainly discuss the median of the evaluation metrics obtained from the 100 simulated samples of each scenario.

### 4.1 Size of the selected variable subsets

Tables 2 and 3 summarize the number of selected variables by the different methods.

The best combination to get closer to the number of active variables at the step of the regularization path construction is the lars algorithm with the E-Net regularization and when samples are first generated. For the independent setting, we observe that eBIC provides the closest number to the size of the support among the model selection methods. Linselect generally selects the empty set. Concerning the data-driven penalties, the number of selected variables is always significantly higher than the number of active variables. Compared to the slope heuristics strategy, the dimension jump strategy selects less variables and seems to be less sensitive to the choice of the algorithm and the regularization function. Of note, the slope heuristic based on lars and E-Net shows similar results to dimension jump. Among the variable identification methods, ESCV and the knockoffs method provide a number of selected variable very close to the number of active variables, whereas all the other methods select less variables than expected. Tigress is very conservative with less than 5 selected variables in average. Bolasso and Stability Selection with lars and bolasso always select the empty set.

When a dependency structure is introduced, for the model selection methods, we observe that eBIC becomes as conservative as Linselect. The dimension jump still has to be preferred to the slope heuristic. All the variable

identification methods show difficulties to manage the dependency between the variables. They select very small subsets of variables in the cluster design and favor the empty subset for the scale-free-min. For the scale-free-max design, ESCV, the knockoff method, bolasso and Stability Selection based on lars and E-Net when samples are first generated (*sub*) outperform.

#### 4.2 Area under the ROC curve (pROC-AUC)

Figure 1 summarizes the pROC-AUC values of the different regularization paths. The target median value is as high as possible.

For the independent setting, the highest median values are obtained with the lars algorithm: almost 0.99 for the E-Net regularization but the median value decreases around 0.73 for the  $\ell_1$  regularization. By contrast, the GD algorithm discriminates less well the active and non active variables since the median values are around 0.52 for both regularization functions.

For correlation structure settings, the highest median values are obtained with the lars algorithm combined with the E-Net regularization: 0.84 and 0.98 for cluster and scale-free-max designs respectively and 1 for scale-free-min design. Lars combined with the  $\ell_1$  regularization does not stand out from the GD algorithm: 0.55, 0.64 and 0.73 for cluster, scale-free-max and scale-free-min designs, whereas the values of the GD algorithm are around 0.55, 0.55 and 0.72.

In conclusion, whatever the settings, lars algorithm with E-Net regularization is the best combination to discriminate the active variables from the non-active ones.

#### 4.3 Mean squared errors (MSE)

Figure 2 and 3 summarize the MSE values of the different methods. The target median value is 0.

For the independent setting, all the model selection strategies are predictive but differences can be observed: the median value is around 0.07 for eBIC and the data-driven penalties and increases to around 0.96 for LinSelect. For the variable identification methods, the smallest median value is around 0.07 for ESCV and the knockoffs method. In contrast, Tigress has a median value ten times higher. Concerning the Bolasso and Stability Selection, the smallest median values are respectively 0.27 and 0.17 and are obtained with the lars algorithm combined with the E-Net regularization when samples are first generated (*sub*).

For correlation structure settings, among the model selection methods, the data-driven penalties are only predictive for the scale-free-max design with median values between 0.30 and 0.45. For the cluster and scale-free-min designs, the median values of eBIC and LinSelect are elevated, higher than 0.92 but lower than 1, whatever the algorithm and the regularization function. These methods provide median values between 0.6 and 1 for the scale-free-max design. Among the variable identification methods, the median values are also larger than 0.92 but lower than 1 for the cluster and the scale-free-min designs, indicating that having prediction performances is also difficult for these methods. In contrast, for the scale-free-max design, the median values are around 0.33 for ESCV and the knockoffs method; 0.5 and 0.44 for Bolasso and Stability Selection when samples are first generated (*sub*) and with lars algorithm combined with E-Net regularization. The other methods provide median values between 0.6 and 1.

To conclude, when the variables are independent, the best methods are eBIC, ESCV and the knockoffs method, whatever the algorithm and the regularization function used. When a dependency structure exists, the predictive performances are deteriorated. For the cluster and scale-free-min designs, the methods are not really predictive since the smallest values of MSE are higher than 0.92. For scale-free-max design, the best methods are ESCV, the knockoffs method and the data-driven penalty with the dimension jump strategy when lars algorithm and E-Net penalty are combined.

#### 4.4 Recall

Figures 4 and 5 summary the recall values of the different methods. The target median value is 1.

For the independent setting, all the model selection methods except LinSelect, have a median value equals 1, meaning that all the active variables are selected. Of note, eBIC combined with the lars algorithm and the E-Net regularization shows more variability than the other methods. Concerning LinSelect, the median value is always equal to 0.07. Among the variable identification methods, ESCV and the knockoffs method median values equal 1 whatever the algorithm and the regularization used. In contrast, Bolasso and Stability Selection results depend on the algorithm and the regularization used. The highest median values are respectively 0.57 and 0.71 and are obtained with the lars algorithm combined with the E-Net regularization function and when samples are first generated. If E-Net is replaced by the  $\ell_1$  regularization, the median values go down to 0. Finally Tigress obtains a median value around 0.18.

When a dependency structure exists, among the model selection methods, the data-driven penalties have a better recall than eBIC and Linselect for the three designs. However, their results depend on the combination of the algorithm and the regularization function used: the slope heuristics strategy combined with the GD algorithm and E-Net is preferable to the dimension jump strategy for the cluster design, but for the scale-free-max the dimension jump combined with lars and E-Net is slightly better than the slope heuristics strategy combined with lars and E-Net. Of note, the values obtained by the slope heuristics have higher variability than those obtained by dimension jump. Among the variable identification methods, the results depend strongly on the design. For the cluster design, the best median values are obtained with ESCV and Stability Selection combined both with the lars algorithm and the E-Net regularization and when the samples are generated first (*sub*). Nevertheless, we can observe that the values remain low, around 0.2 and there exists a large variability across the 100 samples. For the scale-free-max design, the knockoffs method and ESCV provide the best results with median values both close to 0.75, whatever the algorithm and the regularization function. The other methods are clearly worst. For the scale-free-min design, only ESCV has a median value of 1, the others have a median value of 0.

To conclude, when variables are independent, all the model selection methods except LinSelect, ESCV and the knockoffs method recover the active variables. When a dependency structure exists, the results are deteriorated. For the model selection method, the results of the data-driven penalties are better for the three designs but the choice of the algorithm and regularization function depend on the design. An important instability is observed for Bolasso and Stability Selection with respect to the choices of regularization function, the algorithm and the sampling strategy.

#### 4.5 Specificity

Figures 6 and 7 summary the specificity values of the different methods. The target median value is 1.

For the independent setting, only the data-driven penalties show a specificity different from 1. The results obtained by the different algorithms and regularization functions suggest that the reason is the construction of the regularization path since lars combined with E-Net provides the best specificity. But even with this combination, for both data-driven penalties, some non active variables are selected and it might explained by the size of the estimated support which is larger than the size of the true support. All the variable identification methods have a specificity very close to 1.

We draw the same conclusions when a correlation structure exists. Of note, with slope heuristics strategies, values decrease with more variability.

To conclude, the dependency structure between variables and the choice of both algorithm and regularization function do not impact the specificity. All the methods except the data-driven penalties have a specificity very close to 1. For the data-driven penalties, the estimated support being larger than the true support (see subsection 4.1): it suggests that the penalty term (6) is too small.

#### 4.6 FDP and FDR

Figures 8 and 9 summary the FDP values of the different methods. The target median value is low but not equal 0 or slightly smaller than the threshold fixed by the knockoffs method.

For the independent setting, the median values are generally smaller for the  $\ell_1$  than E-Net penalty, whatever the choice of the algorithm used. Among the model selection methods, LinSelect has a FDP equals 0 because the method generally selects no variable. The data-driven penalties have a high FDP. As already observed for the specificity, it could be due to a too small penalty term (6) which selects too much variables. Only eBIC has a reasonable estimated FDR at 0.06 and 0.17 with the  $\ell_1$  and E-Net regularization respectively. For all the variable identification methods, the median value of the FDP is equal to 0 except the knockoffs method where the median value equals 0.08 for an expected control at 0.1.

Unsurprisingly, the presence of a dependency structure impacts negatively the FDP values, especially for the data-driven penalties and ESCV. The  $\ell_1$  penalty generally provide smaller median values than E-Net penalty. Among the model selection methods, for eBIC and Linselect, the median value is always equal to 0. If we calculate an estimated FDR from only the samples where the selected subset is non empty, we get an estimated FDR higher than 0.8 for both methods and for the cluster and scale-free-min designs. Only for the scale-free-max design, eBIC achieves an estimated FDR around 0.11 whereas Linselect achieves 0.34. The data-driven penalty strategies provide high median values. Among the variable identification methods, for the cluster design, only ESCV and the Stability Selection when the lars algorithm is combined with E-Net and when the samples are generated first provide a median value of FDP different from 0. However there exists a large variability and it leads to an estimated FDR higher than 0.24. For scale-free-max, only the knockoffs method has a median value different from 0, leading to an estimated FDR at 0.08. For the scale-free-min, the methods generally select the empty subset and if it is not the case, the selected subset contains non-active variables leading to a high estimated FDR.

To conclude, in the independent setting, eBIC has a reasonable estimated FDR and the knockoffs method gets an estimated FDR close to the threshold provided as an input of the method. When a dependency structure exists, eBIC, Bolasso, Tigress en the knockoffs method should be preferred. Conversely, the data-driven penalty strategies always provide high median values. Because of the variability of the results, the  $\ell_1$  regularization is preferred over the E-Net one.

#### 4.7 Results from the FRANK designs

This subsection is devoted to analyze performances of the methods from the FRANK data. These datasets are interesting for an application point of view since the dependency structure of the data is close to a gene regulatory network. A complementary study shows that the distribution of the FRANK data is far from the Gaussian distribution. Thus, the FRANK data deviate from the statistical model assumptions and this subsection investigates the impact of the non-Gaussian assumption. We recall that we consider the variables having either the highest (setting denoted by *FRANK-max*) or the smallest (setting denoted by *FRANK-min*) number of neighbors as response variables. We only present the conclusions. Graphical representations are available at <https://sites.google.com/view/placroix/research>.

##### 4.7.1 Area under the ROC curve (pROC-AUC):

Similarly with scenarios from independent and Gaussian models, the combination of the E-Net regularization with the lars algorithm achieves the highest value of pROC-AUC. However, these values equal 0.5 for FRANK-max and 0.57 for FRANK-min, significantly smaller values compared to other scenarios. Median values from other combinations are smaller than 0.3.

To conclude, the quality of the obtained regularization paths are clearly deteriorated on FRANK data.

##### 4.7.2 MSE:

For FRANK-max design, all the model selection methods lead to a median value larger than 1, meaning that they are not predictive. As for the variable identification methods, the median values are strictly smaller than 1 but strictly larger than 0.99 except for ESCV with values between 0.97 and 0.99.

For FRANK-min design, among the model selection methods, eBIC and LinSelect are the best methods, with values between 0.97 and 0.99. Those from the data-driven penalties are always larger than 1. ESCV is the best method among the variable identification ones, with median value between 0.97 and 0.99. The others provide a median value between 0.99 and 1.

To conclude, the prediction performances are deteriorated on the FRANK designs compared to the independent and Gaussian models. The model selection methods lose even their prediction performances.

##### 4.7.3 Recall:

For FRANK-max design, all the median values are null except for ESCV, the dimension jump and slope heuristics with values respectively smaller than 0.06, 0.08 and 0.24.

For FRANK-min design, values are null for all the methods.

To conclude, sensibility is drastically deteriorated on FRANK data meaning that all the methods fail to select the active variables.

##### 4.7.4 Specificity:

For FRANK-max design, concerning the model selection methods, median values for eBIC and LinSelect are larger than 0.99 while those for the data-driven penalties decrease until 0.93 and 0.7 for respectively the dimension jump and the slope heuristics strategies. As for the identification methods, the median values equal 1 except for ESCV where values are larger than 0.96 and Bolasso with lars, E-Net and when samples are first generated where values are larger than 0.98.

For FRANK-min design, the same conclusions are observed except for the slope heuristics where median values decrease until 0.84.

To conclude, the specificity values remain high on FRANK data for all methods and the data-driven penalty strategies are the worst ones.

#### 4.7.5 FDP and FDR:

For FRANK-max design, concerning the model selection methods, the median values equal 1 for eBIC and LinSelect, meaning that the selected variables are all non-active. Median values equal 0.9 and 0.88 for the slope heuristics and the dimension jump respectively. Concerning the identification methods, median values equal 0.8 for ESCV and Stability Selection with lars, E-Net and when samples are first generated, and 0.6 for Bolasso with lars, E-Net and when samples are first generated. All the other combinations provide a null median value.

For FRANK-min design, all the model selection methods achieve 1. As for the variable identification methods, all the median values are null except for ESCV, Bolasso and Stability Selection with lars, E-Net and when samples are first generated with values equal 1, 0.7 and 1 respectively.

To conclude, the FDR values are drastically deteriorated on FRANK data with values close to 1 in most cases.

### 4.8 Impact of the high-dimension

This subsection looks at the impact of the high-dimension. If the results of a method remain similar whereas the sample size  $n$  increases, then it ensures that the high-dimension is not the reason of poor performances. For a neutral comparison, results of this subsection are obtained from the same datasets used from Subsections 4.1 to 4.6 that we complete to get samples of size  $n = 300, 600$  and  $1200$ . We only present the conclusions. Graphical representations are available at <https://sites.google.com/view/placroix/research>.

#### 4.8.1 Area under the ROC curve (pROC-AUC)

In the independent setting, the pROC-AUC equals 1 for lars when  $n$  increases. For the GD algorithm, the result is less expected since the pROC-AUC increases from 0.5 to 0.9 when  $n$  goes from 150 to 300 and then decreases up to 0.76 when  $n$  equals 600 or 1200. The difference between the two regularization functions with lars algorithm that we observed at  $n = 150$  does not exist when  $n$  increases.

When a dependency structure exists, lars combined with E-Net remains the best but the differences between the four combinations between an algorithm and a regularization function become negligible.

#### 4.8.2 MSE

For the independent design, prediction performances of LinSelect improve drastically since values are below 0.1 when  $n = 600$ . When  $n = 1200$ , all the model selection methods provide median values smaller than 0.3. For the variable identification methods, the median values decrease with  $n$  except for Bolasso and Stability Selection when the grids are generated first. For  $n = 300$ , the knockoffs method has a median value close to 1 and then values decrease with  $n$ . As soon as  $n = 600$ , all the methods are similar with median values lower than 0.2 except for Tigress which still remains less predictive. We observe that the combination between lars and E-Net is recommended for Bolasso and Stability Selection.

When a dependency structure exists, approximately no change is observed for cluster and scale-free-min designs. Slight decreases in values are observed but median values remain close to 1. For scale-free-max design, increasing the value of  $n$  changes the best method which become eBIC instead of the data-driven penalty strategy with the dimension jump strategy. All model selection methods give values between 0.2 and 0.3 when  $n = 600$ . Concerning the variable identification methods, only Bolasso and Stability Selection used with lars and E-Net show a decreasing with  $n$  up to a median value indicating predictive performances (at most 0.4). Tigress and the knockoffs remain with high values (respectively 0.8 and around 1).

#### 4.8.3 Recall:

For the independent design, all the model selection methods have a recall equals 1 as soon as  $n = 600$ . Among the variable identification methods, it is not verified for Tigress and for Bolasso and Stability Selection where values are between 0.7 and 1 when samples are first generated but lower than 0.5 when grids are first generated. Of note that the knockoffs methods show an instability with respect to  $n$ .

When a dependency structure exists, the recall is generally higher for the model selection methods than for the variable identification methods. Linselect does not manage properly the recall in the cluster design. Concerning the variable identification methods, they behave differently in designs reflecting an instability: the knockoffs method remains at 0 until  $n = 1200$  for cluster and scale-free-min but finishes at 1 for the scale-free-max; ESCV values achieve 1 for scale-free-min, 0.9 for scale-free-max with E-Net but 0.1 for the  $\ell_1$  regularization, and 0.3 for cluster. As for

Bolasso and Stability Selection, median values are 0.9 when grids are first generated but only 0.6 when samples are first generated. Lastly, Tigress is not sensitive, with values always lower than 0.2.

#### 4.8.4 Specificity

In the independent design, only the data-driven penalty strategies do not have a high median value when  $n = 150$ , but when  $n$  increases, 1 is achieved. Of note, the LinSelect performances are deteriorated as soon as  $n = 600$  and are unstable according to the choice of algorithm and regularization function. For the variable identification methods, the median values remain at 1 for all the methods when  $n$  increases.

When a dependency structure exists, the GD algorithm is better than the lars one. We observe that the specificity decreases when  $n$  increases for the data-driven penalties. One possible explanation is that the size of the selected subsets tend to be larger than the set of the active variables when  $n$  increases and consequently some non-active variables are selected. As soon as  $n = 600$ , eBIC has a decreasing specificity for the cluster design, even if the median size of the estimated support is smaller than the true support. Linselect has also a decreasing specificity in the scale-free-max design as soon as  $n = 600$  and the explanation is the number of selected variables which is really too high compared to the number of active variables. For the variable identification methods, the specificity is always 1 except for Bolasso and Stability Selection in the cluster and scale-free-min designs and for the knockoffs method when  $n = 1200$  in the scale-free-max design.

#### 4.8.5 FDP and FDR

In the independent design, the median value decreases when  $n$  increases for all the methods except for Linselect which selects too much variables. The median values of the variable identification methods tend to 0 when  $n$  increases.

When a dependency structure exists, Linselect has surprisingly a median value increasing with  $n$  except for the scale-free-min design. This is also the case for the data-driven penalties. eBIC is the model selection method having the best estimated FDR whatever the designs. The median values of the variable identification methods remain very low when  $n$  increases, surprisingly except for Stability Selection in scale-free-min and cluster designs.

## 5 Take home message

As preliminary recommendations, Tables 2 and 3 suggest that the E-Net regularization, the lars algorithm and the strategy of samples first generated have to be privileged to get a size of selected variables set close to the size of the set of active variables. The considered data-driven penalties provide a too large set of selected variables. Conversely, Tigress and LinSelect are both conservative leading to an almost empty set of selected variables.

### 5.1 Recommendation per method

The structure of dependencies between variables in a dataset is often unknown. Therefore, we summarize in the following table the best combination of the algorithm, the regularization function and eventually the strategy of the procedure for each of the variable selection method, regardless of the dependency between the variables:

method	regularization function	algorithm	strategy
<b>path:</b>	E-Net	lars	
<b>eBIC:</b>	$\ell_1$	lars	
<b>data-driven penalties:</b>	<i>to avoid</i>		
<b>LinSelect:</b>	indifferent	indifferent	
<b>ESCV:</b>	indifferent	indifferent	
<b>The knockoffs:</b>	indifferent	indifferent	
<b>Tigress:</b>	$\ell_1$ (per default)	lars (per default)	
<b>Bolasso:</b>	E-Net	lars	samples first generated
<b>Stability Selection:</b>	<i>to avoid</i>		

More precisely, the data-driven penalties have to be avoided due to bad performances and instabilities in results. Among the variable identification methods, we suggest to use Bolasso instead of Stability Selection since (i) Stability

Selection gives similar or poorer results compared to Bolasso, (ii) Stability Selection results often depend on the variable dependency structure and are unstable, (iii) the Stability Selection method is processed on samples of size  $\lfloor \frac{n}{2} \rfloor$ , which is even smaller than  $p$ , providing additional estimation problem due to the high-dimension.

## 5.2 Recommendation per metric

In this subsection, we summarize metric per metric the best methods as well as those to avoid, regardless of the dependency between the variables.

In the following Table are summarizing methods getting the best performances to discriminate the active variables and the non-actives ones (high value of pROC-AUC):

<b>pROC-AUC:</b>	
Best regularization function:	E-Net
Best algorithm:	lars

In the following Table are summarizing methods getting the best or the worst prediction performances (small value of MSE):

<b>MSE:</b>	
Best methods:	ESCV, knockoffs, LinSelect, eBIC
Best regularization function:	E-Net
Best algorithm:	lars
methods to avoid:	data-driven penalties

More precisely, ESCV and knockoffs have the smallest values of MSE and LinSelect and eBIC have also reasonably small MSE values, making them predictive methods. Using the data-driven penalties has to be avoided when the dependency structure is unknown since the MSE values are larger than 1 on the cluster and scale-free-min designs, which is undesirable.

In the following Table are summarizing the best and the worst methods to provide variable sets with enough active variables while taking into account that the non-active variables are not selected (high value of both recall and specificity):

<b>recall and specificity:</b>	
Best methods:	ESCV, eBIC
Best regularization function:	E-Net
Best algorithm :	GD, lars
methods to avoid:	Tigress, LinSelect, data-driven penalties

Recall and specificity metrics have to be considered together. Indeed, only controlling the recall would lead to a variable set containing all the active variables but also many non-active ones. Conversely, only controlling the specificity would lead to a variable set with a few non-active variables but many active variables would not be selected. Thus, only controlling specificity leads to a conservative method with a small recall, while controlling recall leads to a sensitive method with a small specificity. Controlling both metrics simultaneously should provide a variable set closer to the active variable set. Table above provides methods with high values for both metrics. More precisely, Tigress and LinSelect are set apart since their recall values are smaller than 0.5. In contrast, the data-driven penalties are set apart since their specificity values are the smallest with respect to all the others. The remaining variable identification methods have the same performances in terms of specificity metric and ESCV provides the highest recall value. In the same way, the remaining model selection methods have the same performances in terms of the recall metric and eBIC provides the highest specificity value. These values are just below those of ESCV but eBIC recall values are just above

ESCV recall ones, which makes eBIC and ESCV competitive methods for control of both recall and specificity.

In the following Table are summarizing the best and the worst methods to provide variable sets with a small number of non-active variables (small value of FDR):

<b>FDR:</b>	
Best methods:	Bolasso, the knockoffs
Best regularization function:	$\ell_1$
Best algorithm:	GD, lars
methods to avoid:	ESCV, data-driven penalties

More precisely, Bolasso and the knockoffs have FDP values smaller than 0.1, the threshold usually fixed in practice. LinSelect and Tigress are set apart since they often select no variable.

## 6 Discussion

High-dimensional Gaussian linear regression is commonly used to model interactions between entities when the number of variables is close to or larger than the number of observations. This framework raises many methodological questions and this review aims at highlighting the performances of some methods according to some different objectives. This work is focused on variables selection methods from regularization path constructions. We propose an evaluation of both the regularization path construction and the choice of the final selected variables. The first step is based on the minimization of the least-squares penalized by a regularization function and the main question is the choice of the regularization function. The second step is based on selection of the final variable subset either by penalized criterion minimization with either asymptotic or non-asymptotic properties (model selection methods), or by data sampling strategies (variable identification methods). To evaluate the different methods in a fair way, we simulated different settings, each one having its own characteristics. The independent setting is irrelevant for most of the applications since having completely independent variables is rare in practice. However, this setting is commonly used to develop the statistical methods and is a benchmark to evaluate performances of any method. The settings based on the Gaussian graphical model generating correlated variables are the most favorable case where assumptions of our model broadly hold. The FRANK setting provides a more realistic framework (it is based on a dynamic process to generate a gene regulatory network) but deviates from the statistical model's assumptions, especially the Gaussian distribution. Lastly, the methods are evaluated for different performances: the ability to discriminate the active variables from the others through the pROC-AUC, the prediction performances through the MSE, the ability to recover the active variables through the recall, the ability to not select the non-active variables through the specificity and the quality of the selected variables subset through the FDR.

The impact of dependency structure between variables is evaluated by comparing results on the independent case with those on the three other designs. Our results show significant degradation of performances for all the metrics on the three last settings in comparison with the independent one. This proves that controlling the variable dependency structure is an important assumption for the statistical procedure. Our observations suggest that when dependence between variables exists, the methods work better when the support is large enough (scale-free-max).

One of the most striking conclusions is that choosing the best method depends on the metric to control. Analyses are performed by studying the medians and the variability of the results metric by metric. The first step of the statistical framework is to order the variables through regularization path constructions. According to the simulation study, the combination of the lars algorithm and the E-Net regularization is the best one to discriminate the active variables and the non-active ones. In a complementary study, we observe that some non-active variables appear at the beginning of the regularization paths and their number is reduced when  $n$  increases. So, the high-dimensional context impacts the variable selection procedure at the step of the regularization paths constructions. To define a final variable subset from the regularization paths, ESCV and the knockoffs are the best methods for prediction performances. This is surprising since model selection procedures are constructed for theoretical guarantees on the predictive risk. Among them, eBIC and LinSelect are competitive and provide slightly higher values of the MSE compared with ESCV and the knockoffs. To recover all active variables, all the model selection methods excepted LinSelect are to be favored: they provide a high recall. ESCV is the best method among the variable identification ones. All variable identification methods are specific and eBIC is the best method among the model selection ones for specificity. So, to control both recall and specificity, ESCV and eBIC are the best choices. Lastly, to ensure that the selected variables are active variables, a



small value of FDR is expected. In this direction, Bolasso and the knockoffs get a FDR value smaller than 0.1, the common threshold used in practice. Tigress and LinSelect are conservative providing often an empty set of variables.

As general conclusions, while methods based on asymptotic properties are commonly used in literature, this review emphasizes that the non-asymptotic ones should also be considered. We often observe bad performances with the data-driven penalties. One reason may be that data-driven methods are based on a heuristic whereas LinSelect, the other non-asymptotic model selection procedure, was constructed from an oracle inequality. Moreover, we also want to point out that for the data-driven penalty methods, the penalty shape and the multiplicative constant 2.5 in (6) are derived from a variable selection procedure in a context of changepoint detection [13], so in a different statistical framework from the one we are studying in this review. These quantities may be not adapted in a high-dimensional Gaussian linear regression and more work is required to propose other calibrations adapted in high-dimensional regression. More generally, the choice of parameters (setting to default values) for each method can also be discussed.

A statistical framework always requires some assumptions. For the high-dimensional Gaussian linear regression, data are assumed to be distributed according to a Gaussian distribution, observations are supposed to be independent, variable dependencies are well controlled and the set of active variables is supposed to be small. Experimental datasets generally do not verify all these assumptions. In this work, we just investigate the impact of the high-dimension and the relaxation of the Gaussian assumption but some other relaxations can be studied in future works. We observe that the MSE continue to be high for the data-driven penalties and Tigress, as well as the recall remaining small, even if  $n$  increases. So, high-dimension is not the reason to explain these bad performances. Surprisingly, FDR values drastically increase with  $n$  for LinSelect, eBIC and Stability Selection and the specificity values decrease for LinSelect with  $n$ . Since our study shows deterioration of values for all metrics in the FRANK data, the Gaussian distribution seems to be an important assumption for the considered statistical model. Conclusions could indicate the importance of data transformations and preprocessing steps if the Gaussian assumption is not satisfied (see [48]). Of note, in a complementary study, the *shrinkage* transformation (available in the R function *huge.generator* from the R package *huge*) is tested on FRANK-max and FRANK-min to replace the classical normalization per variable. The MSE values are slightly improved but all the other metrics are slightly deteriorated (boxplots are provided in supplementary material available in <sup>1</sup>).

## 7 Supplementary data

The scripts are available from <sup>2</sup>.

Boxplots of results in Subsection 4.8 and 4.7 are provided in supplementary material available in <sup>3</sup>.

## Acknowledgments

We are grateful to the INRAE MIGALE bioinformatics facility (MIGALE, INRAE, 2020. Migale bioinformatics Facility, doi: 10.15454/1.5572390655343293E12) for providing computing resources. IPS2 benefits from the support of the LabEx Saclay Plant Sciences-SPS (ANR-17-EUR-0007). This work was supported by a public grant as part of the Investissement d'avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH.

## References

- [1] M.J. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using  $l_1$ -constrained quadratic programming. *IEEE Transactions on Information Theory*, 55(5):2183–2202, 2009.
- [2] P. Bühlmann and J. Mandozzi. High-dimensional variable screening and bias in subsequent inference, with an empirical comparison. *Computational Statistics*, 29(3):407–430, 2014.
- [3] F. Wang, S. Mukherjee, S. Richardson, and S.M. Hill. High-dimensional regression in practice: an empirical study of finite-sample prediction, variable selection and ranking. *Statistics and computing*, 30(3):697–719, 2020.
- [4] C. Wu and S. Ma. A selective review of robust variable selection with applications in bioinformatics. *Briefings in bioinformatics*, 16(5):873–883, 2015.
- [5] S. Vinga. Structured sparsity regularization for analyzing high-dimensional omics data. *Briefings in Bioinformatics*, 22(1):77–87, 2021.

<sup>1</sup><https://sites.google.com/view/placroix/research>

<sup>2</sup>[https://forgemia.inra.fr/GNet/high-dimensional\\_regression\\_comparison](https://forgemia.inra.fr/GNet/high-dimensional_regression_comparison)

<sup>3</sup><https://sites.google.com/view/placroix/research>

- [6] Loann David Denis Desboulets. A review on variable selection in regression analysis. *Econometrics*, 6(4):45, 2018.
- [7] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [8] 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.
- [9] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- [10] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.
- [11] J. Chen and Z. Chen. Extended bayesian information criteria for model selection with large model spaces. *Biometrika*, 95(3):759–771, 2008.
- [12] L. Birgé and P. Massart. Minimal penalties for gaussian model selection. *Probability theory and related fields*, 138(1-2):33–73, 2007.
- [13] E. Lebarbier. Detecting multiple change-points in the mean of gaussian process by model selection. *Signal processing*, 85(4):717–736, 2005.
- [14] J.P. Baudry, C. Maugis, and B. Michel. Slope heuristics: overview and implementation. *Statistics and Computing*, 22(2):455–470, 2012.
- [15] S. Arlot. Minimal penalties and the slope heuristics: a survey. *arXiv preprint arXiv:1901.07277*, 2019.
- [16] Y. Baraud, C. Giraud, and S. Huet. Gaussian model selection with an unknown variance. *The Annals of Statistics*, 37(2):630–672, 2009.
- [17] C. Giraud, S. Huet, and N. Verzelen. High-dimensional regression with unknown variance. *Statistical Science*, 27(4):500–518, 2012.
- [18] C. Lim and B. Yu. Estimation stability with cross-validation (escv). *Journal of Computational and Graphical Statistics*, 25(2):464–492, 2016.
- [19] F. Bach. Bolasso: model consistent lasso estimation through the bootstrap. In *Proceedings of the 25th international conference on Machine learning*, pages 33–40, 2008.
- [20] N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [21] A.C. Haury, F. Mordelet, P. Vera-Licona, and J.P. Vert. Tigress: trustful inference of gene regulation using stability selection. *BMC systems biology*, 6(1):145, 2012.
- [22] R. Barber and E.J. Candès. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085, 2015.
- [23] M. Kos and M. Bogdan. On the asymptotic properties of slope. *Sankhya A*, 82(2):499–532, 2020.
- [24] G. Celeux, M. El Anbari, J.M. Marin, and C.P. Robert. Regularization in regression: comparing bayesian and frequentist methods in a poorly informative situation. *Bayesian Analysis*, 7(2):477–502, 2012.
- [25] H.D. Bondell and B.J. Reich. Consistent high-dimensional bayesian variable selection via penalized credible regions. *Journal of the American Statistical Association*, 107(500):1610–1624, 2012.
- [26] Z. Razaghi-Moghadam and Z. Nikoloski. Supervised learning of gene-regulatory networks based on graph distance profiles of transcriptomics data. *NPJ systems biology and applications*, 6(1):1–8, 2020.
- [27] C. Carré, A. Mas, and G. Krouk. Reverse engineering highlights potential principles of large gene regulatory network design and learning. *NPJ systems biology and applications*, 3(1):1–15, 2017.
- [28] H. Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.
- [29] A Hoerl and R Kennard. Ridge regression, in ‘encyclopedia of statistical sciences’, vol. 8, 1988.
- [30] 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, 2006.
- [31] L. Jacob, G. Obozinski, and J.P. Vert. Group lasso with overlap and graph lasso. In *Proceedings of the 26th annual international conference on machine learning*, pages 433–440, 2009.

- [32] P. Zhao, G. Rocha, and B. Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *The Annals of Statistics*, 37(6A):3468–3497, 2009.
- [33] Zhifan Li, Yanhang Zhang, and Jianxin Yin. Sharp minimax optimality of lasso and slope under double sparsity assumption, 2023.
- [34] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- [35] N. Meinshausen. Relaxed lasso. *Computational Statistics & Data Analysis*, 52(1):374–393, 2007.
- [36] Pierre Connault. *Calibration d’algorithmes de type Lasso et analyse statistique de données métallurgiques en aéronautique*. PhD thesis, Paris 11, 2011.
- [37] L. Birgé and P. Massart. Gaussian model selection. *Journal of the European Mathematical Society*, 3(3):203–268, 2001.
- [38] D. Allen. The relationship between variable selection and data augmentation and slow feature analysis. *Technometrics*, 16:125–127, 1974.
- [39] M. Stone. Cross-validators choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133, 1974.
- [40] E. Jean Candès, Y. Fan, L. Janson, and J. Lv. *Panning for gold: Model-free knockoffs for high-dimensional controlled variable selection*. Department of Statistics, Stanford University, 2016.
- [41] J. Fan, Y. Fan, and E. Barut. Adaptive robust variable selection. *Annals of statistics*, 42(1):324, 2014.
- [42] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, 34(3):1436–1462, 2006.
- [43] S.L. Lauritzen. *Graphical Models*. Oxford University Press, 1996.
- [44] Irene Córdoba, Gherardo Varando, Concha Bielza, and Pedro Larrañaga. On generating random gaussian graphical models. *International Journal of Approximate Reasoning*, 125:240–250, 2020.
- [45] H. Zou and M. Yuan. Regularized simultaneous model selection in multiple quantiles regression. *Computational Statistics & Data Analysis*, 52(12):5296–5304, 2008.
- [46] L. Wang, Y. Wu, and R. Li. Quantile regression for analyzing heterogeneity in ultra-high dimension. *Journal of the American Statistical Association*, 107(497):214–222, 2012.
- [47] B. Peng and L. Wang. An iterative coordinate descent algorithm for high-dimensional nonconvex penalized quantile regression. *Journal of Computational and Graphical Statistics*, 24(3):676–694, 2015.
- [48] H. Liu, F. Han, M. Yuan, J. Lafferty, and L. Wasserman. High-dimensional semiparametric gaussian copula graphical models. *The Annals of Statistics*, 40(4):2293–2326, 2012.

## 8 Appendix: Boxplots for $n = 150$ per metric

In this appendix, table of the number of active variables per simulation setting and table of the number of selected variables per method and per simulation setting are provided when  $n = 150$ . Then, boxplots of results from scenarios independent, cluster, scale-free-max and scale-free-min when  $n = 150$  are provided. They are arranged by metric. For boxplots of results when  $n = 300, 600, 1200$  as well as for boxplots of results from the FRANK scenarios, the reader is invited to consult the supplementary material available in <sup>4</sup>.

The *glm*, *enet* and *lasso* denote respectively the gradient descent algorithm, the Elastic-Net regularization function and the Lasso regularization function. The slope heuristics method and the dimension jump are named respectively by *slope* and *jump*. Bolasso and Stability Selection are named respectively by *b* and *ss*. Lastly, *grid* and *sub* denote respectively the strategy when grids are first generated and the strategy when samples are first generated.

Active variable number	ind	cluster	scale_free_max	scale_free_min
mean	12.59	11.63	31.41	1
( <i>sd</i> )	(1.76)	(2.75)	(9.70)	(0)

Table 1: Active variable number

Selected variables number	ind	cluster	scale_free_max	scale_free_min
GD_E-Net_ebic (mean)	14.34	1.27	7.53	1.03
( <i>sd</i> )	(2.84)	(0.63)	(2.30)	(0.17)
GD_E-Net_slope (mean)	36.80	62.69	60.63	57.99
( <i>sd</i> )	(16.64)	(34.98)	(21.47)	(36.28)
GD_E-Net_jump (mean)	25.83	23.37	42.20	24.38
( <i>sd</i> )	(12.23)	(17.13)	(20.19)	(19.70)
GD_E-Net_linselect (mean)	1.07	1.32	1.59	1.04
( <i>sd</i> )	(0.33)	(1.12)	(1.51)	(0.20)
GD_ℓ <sub>1</sub> _ebic (mean)	13.50	1.27	7.67	1.03
( <i>sd</i> )	(2.21)	(0.63)	(2.40)	(0.17)
GD_ℓ <sub>1</sub> _slope (mean)	35.69	62.94	62.61	59.67
( <i>sd</i> )	(16.12)	(35.49)	(21.30)	(35.06)
GD_ℓ <sub>1</sub> _jump (mean)	28.23	23.85	42.49	24.12
( <i>sd</i> )	(11.93)	(17.61)	(20.16)	(18.67)
GD_ℓ <sub>1</sub> _linselect (mean)	1.07	1.30	1.53	1.04
( <i>sd</i> )	(0.33)	(1.02)	(1.49)	(0.20)
lars_E-Net_ebic (mean)	14.80	1.25	7.43	1.02
( <i>sd</i> )	(3.76)	(0.61)	(2.43)	(0.14)
lars_E-Net_slope (mean)	20.81	23.15	45.33	22.09
( <i>sd</i> )	(10.43)	(15.93)	(17.30)	(14.17)
lars_E-Net_jump (mean)	23.47	23.34	43.78	22.86
( <i>sd</i> )	(13.39)	(13.89)	(15.88)	(13.69)
lars_E-Net_linselect (mean)	1.04	1.23	1.55	1.03
( <i>sd</i> )	(0.28)	(0.76)	(1.51)	(0.17)
lars_ℓ <sub>1</sub> _ebic (mean)	13.50	1.26	7.71	1.02
( <i>sd</i> )	(2.21)	(0.61)	(2.36)	(0.14)
lars_ℓ <sub>1</sub> _slope (mean)	61.65	54.34	67.72	53.51
( <i>sd</i> )	(29.42)	(32.72)	(24.92)	(33.65)
lars_ℓ <sub>1</sub> _jump (mean)	32.30	25.13	44.72	23.84
( <i>sd</i> )	(15.33)	(19.54)	(22.14)	(18.35)
lars_ℓ <sub>1</sub> _linselect (mean)	1.04	1.26	1.46	1.03
( <i>sd</i> )	(0.28)	(0.81)	(1.44)	(0.17)

Table 2: Selected variables number for model selection methods

<sup>4</sup><https://sites.google.com/view/placroix/research>

Selected variables number	ind	cluster	scale_free_max	scale_free_min
GD_bol_grid_E-Net (mean) ( <i>sd</i> )	4.23 (1.28)	0.20 (0.45)	1.00 (0.86)	0.01 (0.10)
GD_bol_grid_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	4.08 (1.32)	0.10 (0.36)	0.46 (0.59)	0.00 (0.00)
GD_bol_sub_E-Net (mean) ( <i>sd</i> )	2.48 (1.02)	0.17 (0.43)	0.42 (0.57)	0.01 (0.10)
GD_bol_sub_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	2.34 (0.99)	0.08 (0.31)	0.23 (0.42)	0.00 (0.00)
GD_escv_E-Net (mean) ( <i>sd</i> )	11.61 (4.42)	6.81 (8.71)	22.42 (11.18)	4.64 (10.37)
GD_escv_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	11.59 (4.41)	6.86 (8.62)	22.42 (11.13)	5.07 (10.46)
GD_knockoffs_E-Net (mean) ( <i>sd</i> )	12.11 (4.68)	0.00 (0.00)	25.74 (8.44)	0.00 (0.00)
GD_knockoffs_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	12.11 (4.68)	0.00 (0.00)	25.74 (8.44)	0.00 (0.00)
GD_ss_grid_E-Net (mean) ( <i>sd</i> )	6.18 (1.74)	0.49 (0.70)	2.84 (1.27)	0.19 (0.42)
GD_ss_grid_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	6.16 (1.70)	0.33 (0.59)	1.73 (1.10)	0.09 (0.29)
GD_ss_sub_E-Net (mean) ( <i>sd</i> )	3.89 (1.23)	0.34 (0.59)	1.64 (1.05)	0.15 (0.39)
GD_ss_sub_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	3.81 (1.20)	0.26 (0.50)	1.04 (0.90)	0.07 (0.26)
lars_bol_sub_E-Net (mean) ( <i>sd</i> )	7.15 (1.35)	1.48 (1.65)	9.19 (2.21)	0.43 (0.54)
lars_bol_sub_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	0.00 (0.00)	0.00 (0.00)	0.04 (0.20)	0.00 (0.00)
lars_ss_sub_E-Net (mean) ( <i>sd</i> )	9.11 (1.35)	2.79 (1.97)	12.64 (2.25)	0.99 (0.80)
lars_ss_sub_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	0.00 (0.00)	0.00 (0.00)	0.16 (0.42)	0.00 (0.00)
lars_tigress_ℓ <sub>1</sub> (mean) ( <i>sd</i> )	2.14 (0.89)	0.31 (0.56)	0.40 (0.59)	0.14 (0.35)

Table 3: Selected variables number for variable identification methods

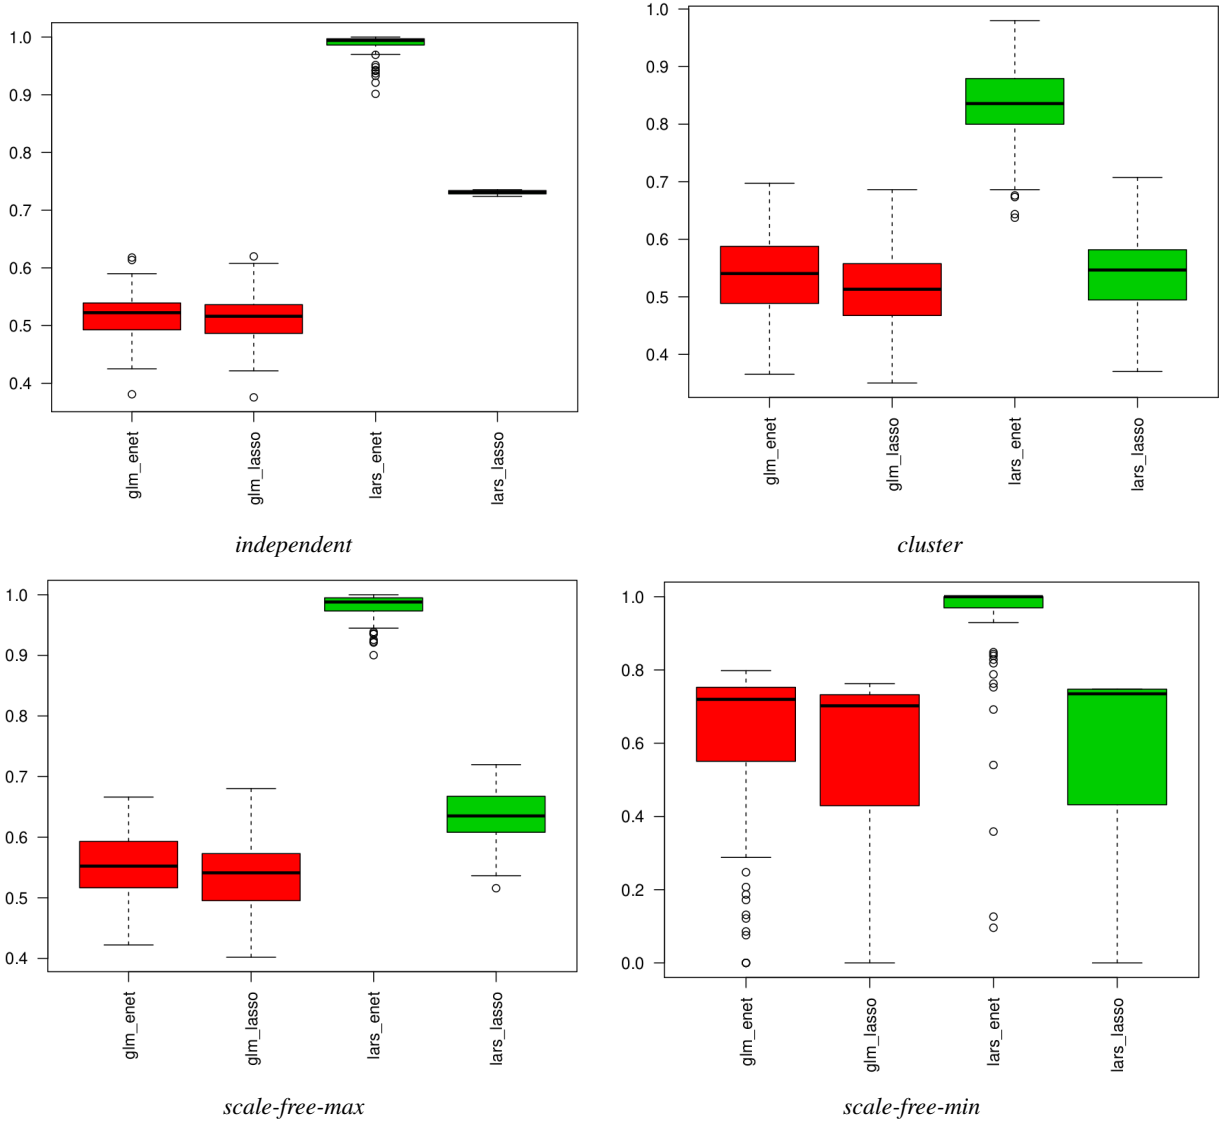


Figure 1: Boxplots of the pROC-AUC values for  $n = 150$ .




Figure 2: Boxplots of the MSE values for model selection procedures and for  $n = 150$ .




Figure 3: Boxplots of the MSE values for variable identification procedures and for  $n = 150$ .



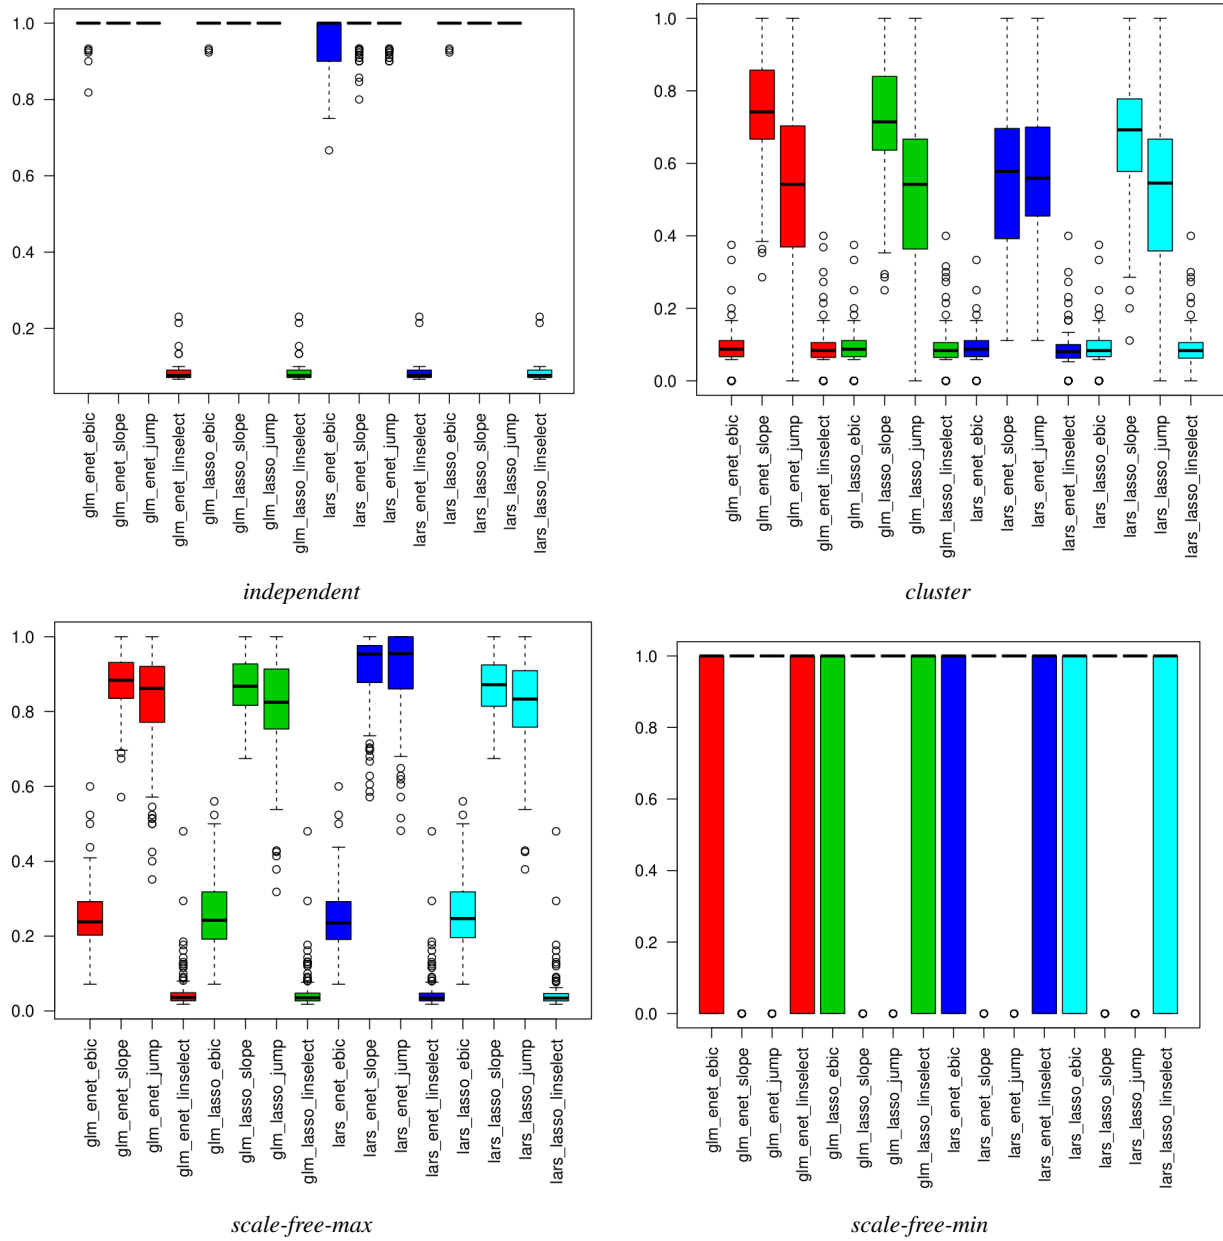


Figure 4: Boxplots of the recall values for model selection procedures and for  $n = 150$ .




Figure 5: Boxplots of the recall values for variable identification procedures and for  $n = 150$ .




Figure 6: Boxplots of the specificity values for model selection procedures and for  $n = 150$ .

# Variable selection in high-dimensional regression




Figure 7: Boxplots of the specificity values for variable identification procedures and for  $n = 150$ .

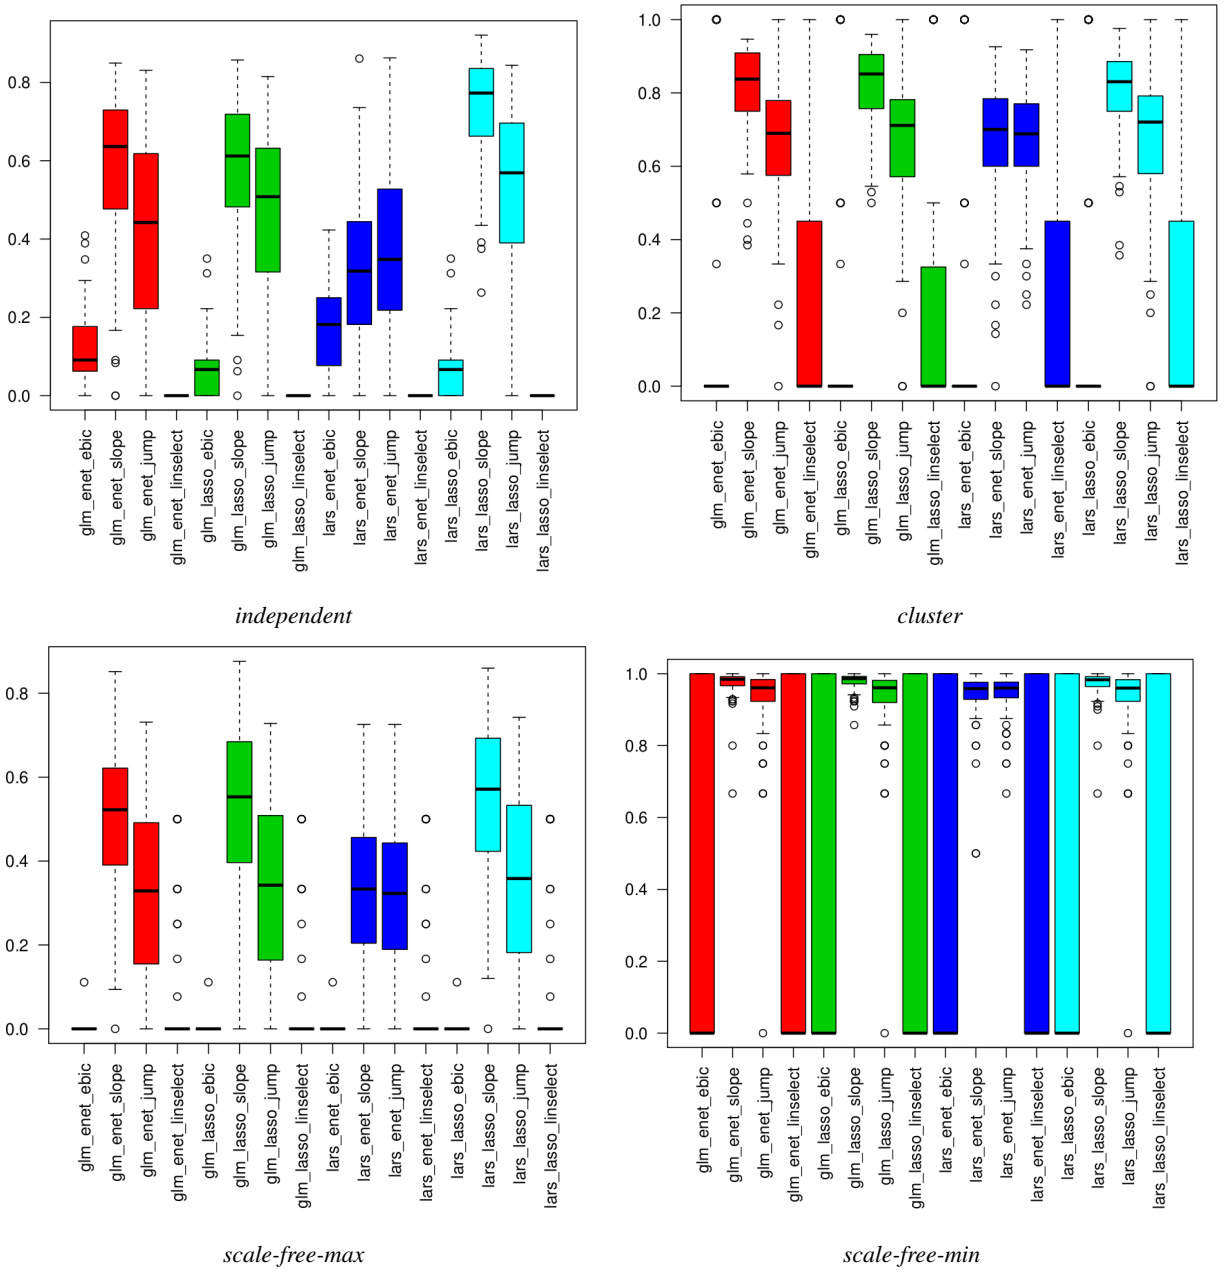


Figure 8: Boxplots of the FDP values for model selection procedures and for  $n = 150$ .




Figure 9: Boxplots of the FDP values for variable identification procedures and for  $n = 150$ .