



HAL
open science

A variable selection approach for highly correlated predictors in high-dimensional genomic data

Wencan Zhu, Céline Lévy-Leduc, Nils Ternes

► **To cite this version:**

Wencan Zhu, Céline Lévy-Leduc, Nils Ternes. A variable selection approach for highly correlated predictors in high-dimensional genomic data. *Bioinformatics*, Oxford University Press (OUP), 2021, 37 (16), pp.2238-2244. 10.1093/bioinformatics/btab114. hal-02904344

HAL Id: hal-02904344

<https://hal-agroparistech.archives-ouvertes.fr/hal-02904344>

Submitted on 14 Apr 2022

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.

A VARIABLE SELECTION APPROACH FOR HIGHLY CORRELATED PREDICTORS IN HIGH-DIMENSIONAL GENOMIC DATA

WENCAN ZHU, CÉLINE LÉVY-LEDUC, AND NILS TERNÈS

ABSTRACT. In genomic studies, identifying biomarkers associated with a variable of interest is a major concern in biomedical research. Regularized approaches are classically used to perform variable selection in high-dimensional linear models. However, these methods can fail in highly correlated settings. We propose a novel variable selection approach called WLasso, taking these correlations into account. It consists in rewriting the initial high-dimensional linear model to remove the correlation between the biomarkers (predictors) and in applying the generalized Lasso criterion. The performance of WLasso is assessed using synthetic data in several scenarios and compared with recent alternative approaches. The results show that when the biomarkers are highly correlated, WLasso outperforms the other approaches in sparse high-dimensional frameworks. The method is also successfully illustrated on publicly available gene expression data in breast cancer. Our method is implemented in the WLasso R package which is available from the Comprehensive R Archive Network.

1. INTRODUCTION

The identification of prognostic genomic biomarkers (i.e. biomarkers associated with a variable of interest, for example a clinical endpoint in clinical trials) has become a major concern for the biomedical research field. Indeed, prognostic biomarkers may help to anticipate the prognosis of individual patients and may also be useful to understand a disease at a molecular level and possibly guide for the development of new treatment strategies (Kalia (2015)).

To this end, statistical variable selection approaches are widely used to identify a subset of biomarkers in high-dimensional settings where the number of biomarkers p is much larger than the sample size n . Several reviews focused on this topic (Saeys et al. (2007) and Heinze et al. (2018) for example). Commonly used techniques include hypothesis-based test: t-test (McDonald (2009)), wrapper approaches (Saeys et al. (2007)): forward, backward selection, and penalized approaches: Lasso (Tibshirani (1996)), Elastic-net (Zou and Hastie (2015)), SCAD (Fan and Li (2001)) among others. Hypothesis tests are limited to independently consider associations for each biomarker thus neglecting potential relationships between them. Wrapper approaches often show high risk of overfitting and are computationally expensive for high-dimensional data (Smith (2018)). More efforts have

Key words and phrases. variable selection; highly correlated predictors; genomic data.

been devoted to penalized methods, given the attractive feature of automatically performing variable selection and coefficient estimation simultaneously (Fan and Li (2006)). We shall thus focus on this type of approaches in the following.

Let us consider the following linear regression model:

$$(1) \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}.$$

where $\mathbf{y} = (y_1, \dots, y_n)^T$ is the variable to explain (clinical endpoint), $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_p)$ is the design matrix containing the expression of biomarkers such that the correlation matrix of its columns is $\boldsymbol{\Sigma}$, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is a sparse vector to estimate, namely with a majority of null coefficients, and $\boldsymbol{\epsilon}$ is the error term. The Lasso penalty is a well-known approach to estimate $\boldsymbol{\beta}$ with a sparsity enforcing constraint. It consists in minimizing the following penalized least-squares criterion (Tibshirani (1996)):

$$(2) \quad L_\lambda(\boldsymbol{\beta}) = \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1,$$

where $\|\cdot\|_2$ is the Euclidean norm and $\|\boldsymbol{\beta}\|_1 = \sum_{k=1}^p |\beta_k|$. However, the Lasso has several drawbacks in highly correlated settings (Zou and Hastie (2015)) such as the violation of the Irrepresentable Condition (IC) defined in Zhao and Yu (2006). The authors of this article prove that this condition is necessary and sufficient to recover the support of $\boldsymbol{\beta}$, namely to retrieve the null and non null components in the vector $\boldsymbol{\beta}$ and thus to provide a sign consistent estimator. This condition is defined as follows. Let $S = \{j, \beta_j \neq 0\}$ be the set of active variables, S^c the set of non-active variables and \mathbf{X}_A the submatrix of \mathbf{X} containing only the indices of columns which are in the set A . Then, the design matrix \mathbf{X} satisfies the IC if, for some constant $\eta \in (0, 1]$,

$$(3) \quad \left| (\mathbf{X}_{S^c}^T \mathbf{X}_S (\mathbf{X}_S^T \mathbf{X}_S)^{-1} \text{sign}(\boldsymbol{\beta}_S))_j \right| \leq 1 - \eta, \text{ for all } j,$$

where $\text{sign}(x) = 1$, if $x > 0$, -1 if $x < 0$ and 0 if $x = 0$. Intuitively, this condition means that the correlation between the active and non active explanatory variables is smaller than the correlation between the active explanatory variables. Hence, this condition is most likely to be violated when the correlations between non active and active variables are large. In high-dimensional genomic data, this condition is difficult to guarantee as the correlation between biomarkers is usually high (Michalopoulos et al. (2012)). Wang et al. (2019) tested the irrepresentable condition on several publicly available genomic data and highlighted that the condition is violated in almost all the datasets investigated.

Methods have been proposed to deal with the issue of high correlations between the biomarkers. Preconditioning the Lasso is one of them. It consists in transforming the given data \mathbf{X} and \mathbf{y} before applying the Lasso criterion. For example, Jia and Rohe (2015) and Wang and Leng (2016) proposed to left multiply \mathbf{X} , \mathbf{y} and thus $\boldsymbol{\epsilon}$ in Model (1) by specific matrices to remove the correlations between the columns of \mathbf{X} . A major drawback of the latter, called HOLP (High dimensional Ordinary Least squares Projection), is that the preconditioning step may increase the variance of the error term and thus may alter the variable selection performance. Another recently published method named Precision Lasso (Wang et al. (2019)) proposes to handle the correlation issue by assigning similar weights to correlated variables. This approach revealed better performance than the other

methods when the biomarkers were highly correlated. However, it failed in more favorable cases when the biomarkers are not correlated.

In this paper, we propose an alternative and novel approach, called Whitening Lasso (WLasso), to take into account the correlations that may exist between the predictors (biomarkers). Our method proposes to transform Model (1) in order to remove the correlations existing between the columns of \mathbf{X} and thus to “whiten” them and make the IC valid but without changing the error term ϵ . This prevents us from noise inflation, see (4). Then, the variable (biomarker) selection is performed thanks to the generalized Lasso criterion devised by Tibshirani and Taylor (2011). The full details of our method are provided in Section 2. An extensive simulation study is presented in Section 3 to assess the selection performance of our approach and to compare it to other methods in different settings. WLasso is also applied to a publicly available dataset in breast cancer in Section 4. Finally, we discuss our findings and give concluding remarks in Section 5.

2. METHODS

In this section, we propose a novel variable selection approach called WLasso (Whitening Lasso) which consists in removing the correlations existing between the biomarkers (columns of \mathbf{X}) and in applying the generalized Lasso criterion proposed by Tibshirani and Taylor (2011) for variable selection purpose.

2.1. Model Transformation. Inspired by the literature on preconditioning, we propose to rewrite Model (1) in order to remove the correlation existing between the columns of \mathbf{X} . More precisely, let $\Sigma^{-1/2} := \mathbf{U}\mathbf{D}^{-1/2}\mathbf{U}^T$ where \mathbf{U} and \mathbf{D} are the matrices involved in the spectral decomposition of the symmetric matrix Σ given by: $\Sigma = \mathbf{U}\mathbf{D}\mathbf{U}^T$. We then denote $\tilde{\mathbf{X}} = \mathbf{X}\Sigma^{-1/2}$. Therefore, (1) can be rewritten as follows:

$$(4) \quad \mathbf{y} = \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} + \epsilon,$$

where $\tilde{\boldsymbol{\beta}} = \Sigma^{1/2}\boldsymbol{\beta} := \mathbf{U}\mathbf{D}^{1/2}\mathbf{U}^T\boldsymbol{\beta}$. With such a transformation, since the n rows $\mathbf{x}_1, \dots, \mathbf{x}_n$ of \mathbf{X} are assumed to be independent Gaussian random vectors with a covariance matrix equal to Σ , the covariance matrix of the rows of $\tilde{\mathbf{X}}$ is equal to identity and the columns of $\tilde{\mathbf{X}}$ are thus uncorrelated. The advantage of such a transformation with respect to the preconditioning approach proposed by Wang and Leng (2016) is that the error term ϵ is not modified thus avoiding an increase of the noise which can overwhelm the benefits of a well conditioned design matrix.

To illustrate the benefits of our methodology, observations \mathbf{y} were generated according to Model (1) with $p = 500$, $n = 50$, $\boldsymbol{\beta}$ having 10 non null components which are equal to 2 and with Σ defined by

$$(5) \quad \Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}^T & \Sigma_{22} \end{bmatrix}$$

where Σ_{11} is the correlation matrix of active variables with off-diagonal entries equal to α_1 , Σ_{22} is the one of non active variables with off-diagonal entries equal to α_3 and Σ_{12} is the correlation matrix between active and non active variables with entries equal to α_2 . In

the case where $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, Figure 1 displays the percentage of components j for which the Irrepresentable Condition (3) is not satisfied from 100 replications. We can see from this figure that our approach (WLasso) dramatically improves the number of indices j for which the IC condition is satisfied. The results are even better than those obtained by the transformation proposed by HOLP (Wang and Leng (2016)).

The following illustrations of Section 2 are obtained from observations \mathbf{y} generated according to the previous scenario.

2.2. Estimation of $\tilde{\boldsymbol{\beta}}$. In order to estimate $\tilde{\boldsymbol{\beta}}$ with a sparsity enforcing constraint on $\boldsymbol{\beta}$, we use the generalized Lasso criterion proposed by Tibshirani and Taylor (2011) which consists in minimizing the following criterion with respect to $\boldsymbol{\beta}$:

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\mathbf{D}\boldsymbol{\beta}\|_1,$$

where \mathbf{D} is a specific matrix. Note that this criterion boils down to the classical Lasso criterion if \mathbf{D} is the identity matrix. In Model (4), we thus propose to minimize the following criterion with respect to $\tilde{\boldsymbol{\beta}}$:

$$(6) \quad L_\lambda^{\text{gen}}(\tilde{\boldsymbol{\beta}}) = \|\mathbf{y} - \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}\|_2^2 + \lambda \|\boldsymbol{\Sigma}^{-1/2}\tilde{\boldsymbol{\beta}}\|_1,$$

which guarantees a sparsity enforcing constraint on $\boldsymbol{\beta}$ thanks to the ℓ_1 penalty. We thus obtain

$$\hat{\tilde{\boldsymbol{\beta}}}_0(\lambda) = \text{Argmin}_{\tilde{\boldsymbol{\beta}}} L_\lambda^{\text{gen}}(\tilde{\boldsymbol{\beta}}).$$

To estimate $\tilde{\boldsymbol{\beta}}$, we will not directly use $\hat{\tilde{\boldsymbol{\beta}}}_0(\lambda)$ but the following modified estimator which can be seen as a thresholding of the components of $\hat{\tilde{\boldsymbol{\beta}}}_0(\lambda)$. For K in $\{1, \dots, p\}$, let Top_K be the set of indices corresponding to the K largest values of the components of $|\hat{\tilde{\boldsymbol{\beta}}}_0|$, then the estimator of $\tilde{\boldsymbol{\beta}}$ is $\hat{\tilde{\boldsymbol{\beta}}} = (\hat{\tilde{\boldsymbol{\beta}}}_j^{(K)})$ where $\hat{\tilde{\boldsymbol{\beta}}}_j^{(K)}$ is defined by:

$$(7) \quad \hat{\tilde{\boldsymbol{\beta}}}_j^{(K)}(\lambda) = \begin{cases} \hat{\tilde{\boldsymbol{\beta}}}_{0j}(\lambda), & j \in \text{Top}_K \\ K\text{th largest value of } |\hat{\tilde{\boldsymbol{\beta}}}_{0j}|, & j \notin \text{Top}_K. \end{cases}$$




FIGURE 1. Proportion of components j such that (3) is violated. These results were obtained from 100 replications.

The choice of \widehat{K} is explained in Section 2.4.

Figure 2 displays the average of $(|\widehat{\beta}_j^{(\widehat{K})}(\lambda) - \widetilde{\beta}_j(\lambda)|)_{1 \leq j \leq p}$ for all the values of λ that are considered. We can see from this figure that the thresholding improves the estimation of $\widetilde{\beta}$.

2.3. Estimation of β . As previously, to estimate β , we will first consider $\widehat{\beta}_0 = \Sigma^{-1/2} \widetilde{\beta}$ and then apply a thresholding strategy. Thus, we propose to estimate β by $\widehat{\beta} = (\widehat{\beta}_j^{(\widehat{M})})_{1 \leq j \leq p}$ where $\widehat{\beta}_j^{(\widehat{M})}$ is defined by:

$$(8) \quad \widehat{\beta}_j^{(\widehat{M})}(\lambda) = \begin{cases} \widehat{\beta}_{0j}(\lambda), & j \in \text{Top}_M \\ 0, & j \notin \text{Top}_M. \end{cases}$$

The choice of \widehat{M} is explained in Section 2.4.

As we can see from Figure 3, more true non null (active) components of β (true positive) and more true null (non active) components of β (true negative) can be retrieved with $\widehat{\beta}$ than with $\widehat{\beta}_0$.




FIGURE 2. Boxplots of the average of $(|\widehat{\beta}_{0j}(\lambda) - \widetilde{\beta}_j(\lambda)|)_{1 \leq j \leq p}$ (left) and $(|\widehat{\beta}_j^{(\widehat{K})}(\lambda) - \widetilde{\beta}_j(\lambda)|)_{1 \leq j \leq p}$ for all λ (right) obtained from 100 replications.

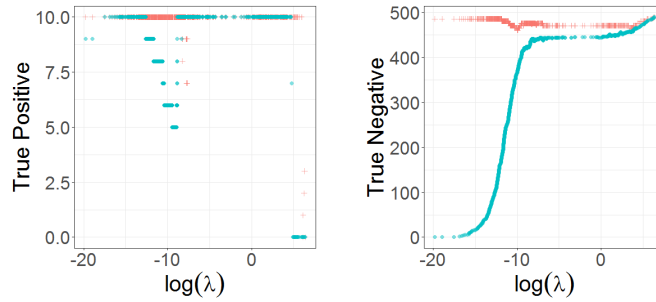


FIGURE 3. Number of True Positive and True Negative for $\widehat{\beta}$ in red and $\widehat{\beta}_0$ in blue for a given vector of observations \mathbf{y} .

2.4. Choice of the parameters. To choose the parameters K and M in (7) and (8) for each λ , we use a strategy based on the Mean Squared Error (MSE). We shall first explain the strategy that we used for choosing \widehat{K} . Let

$$\widetilde{\text{MSE}}_K(\lambda) = \|\mathbf{y} - \widetilde{\mathbf{X}}\widehat{\boldsymbol{\beta}}^{(K)}(\lambda)\|_2^2,$$

where \mathbf{y} , $\widetilde{\mathbf{X}}$ and $\widehat{\boldsymbol{\beta}}^{(K)}(\lambda)$ are defined in (1), (4) and (7), respectively and

$$\widehat{K}(\lambda) = \text{Argmin} \left\{ K \geq 1 \text{ s.t. } \frac{\widetilde{\text{MSE}}_{K+1}(\lambda)}{\widetilde{\text{MSE}}_K(\lambda)} \geq \gamma \right\}, \text{ where } \gamma \in (0, 1).$$

Large values of γ will lead to large values of $\widehat{K}(\lambda)$ and thus to a weak thresholding of the estimator of $\widehat{\boldsymbol{\beta}}$. In practice, as it is shown in Section 3, taking γ in $(0.9, 0.99)$ provides satisfactory and almost similar results.

For the choice of $\widehat{M}(\lambda)$, we use the same procedure except that $\widetilde{\text{MSE}}_K(\lambda)$ is replaced by

$$(9) \quad \text{MSE}_M(\lambda) = \|\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}}^{(M)}(\lambda)\|_2^2,$$

where \mathbf{y} , \mathbf{X} and $\widehat{\boldsymbol{\beta}}^{(M)}(\lambda)$ are defined in (1) and (8), respectively. Both criteria are displayed in Figure 4 for a value of λ which is chosen according to the strategy explained in Section 3.2.

2.5. Estimation of $\boldsymbol{\Sigma}$. Since the matrix $\boldsymbol{\Sigma}$ is unknown in practice, it has to be estimated. In the particular situation where $\boldsymbol{\Sigma}$ has the block structure described in (5), we propose the following strategy. Firstly, we compute the empirical correlation matrix as follows. Let \mathbf{S} be the sample $p \times p$ covariance matrix defined by

$$\mathbf{S} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})', \quad \text{with } \bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i,$$




FIGURE 4. $\widetilde{\text{MSE}}_K(\lambda)$ (left) and $\text{MSE}_M(\lambda)$ (right) for λ chosen thanks to the strategy explained in Section 3.2 for a given vector of observations \mathbf{y} . The vertical dotted lines correspond to $\widehat{K}(\lambda)$ and $\widehat{M}(\lambda)$, respectively.

where \mathbf{x}_i denotes the i th row of \mathbf{X} defined in (1). The corresponding $p \times p$ sample correlation matrix $\mathbf{R} = (R_{i,j})$ is defined by:

$$(10) \quad \mathbf{R}_{i,j} = \frac{S_{i,j}}{\sigma_i \sigma_j}, \quad \forall 1 \leq i, j \leq p,$$

where

$$\sigma_i^2 = \frac{1}{n-1} \sum_{\ell=1}^n (X_{\ell,i} - \bar{X}_i)^2, \quad \text{with } \bar{X}_i = \frac{1}{n} \sum_{\ell=1}^n X_{\ell,i}, \quad \forall 1 \leq i \leq p.$$

Secondly, the two groups (or clusters) of active and non active biomarkers are obtained by using a hierarchical clustering with the complete agglomeration method. Thirdly, the entries of $\hat{\Sigma}$ are computed by averaging the values of \mathbf{R} within the groups. More precisely, let $\rho_{i,j}$ denote the value of the entries in the block having its rows corresponding to Cluster i and its columns to Cluster j . Then, for a given clustering C :

$$(11) \quad \rho_{i,j} = \begin{cases} \frac{1}{\#C(i)\#C(j)} \sum_{k \in C(i), \ell \in C(j)} R_{k,\ell}, & \text{if } C(i) \neq C(j) \\ \frac{1}{\#C(i)(\#C(i)-1)} \sum_{k \in C(i), \ell \in C(i), k \neq \ell} R_{k,\ell}, & \text{if } C(i) = C(j) \end{cases},$$

where $C(i)$ denotes the cluster i , $\#C(i)$ denotes the number of elements in the cluster $C(i)$ and $R_{k,\ell}$ is the (k, ℓ) entry of the matrix \mathbf{R} defined in (10).

We illustrate the performance of our method in Figure 5 in the case where Σ has the structure (5) with $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$. We can see from this figure that the proposed methodology for estimating the correlation coefficients within the blocks of $\hat{\Sigma}$ is very efficient.

2.6. Summary of the W_{Lasso} method. The W_{Lasso} method can be summarized as follows:

- First step: Estimation of the matrix Σ by $\hat{\Sigma}$, see Section 2.5.




FIGURE 5. Estimation of the parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$. The horizontal dotted lines correspond to the true values of the parameters. These results are obtained from 100 replications for each value of p .

- **Second step:** Transformation of Model (1) into Model (4) to remove the correlation existing between the columns of \mathbf{X} , see Section 2.1 where Σ is replaced by $\widehat{\Sigma}$.
- **Third step:** Estimation of $\widetilde{\beta}$ defined in (4), see Section 2.2.
- **Fourth step:** Estimation of β defined in (1), see Section 2.3 and identification of its null and non null components.

3. NUMERICAL EXPERIMENTS

We performed numerical experiments to assess the performance of the WLasso and to compare it with other recent approaches.

All simulated datasets were generated from Model (1) in which the number of predictors (biomarkers) p is equal to 100, 200, 500, 1000 or 2000 and the sample size n is equal to 50 or 100. We randomly chose 10 non null coefficients among the p coefficients of β which correspond to the active biomarkers, thus considering different sparsity levels. The value b of the non null coefficients is equal to either 0.5 or 1 to consider different signal-to-noise ratios.

Regarding the correlation matrix Σ which contains the correlation values between the biomarkers, namely the correlations between the columns of the design matrix \mathbf{X} , several structures were considered:

- Block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$ and $(0.5, 0.7, 0.9)$;
- Independent setting where Σ is the identity matrix.

The results that are presented hereafter are obtained from 100 replications.

3.1. Estimation of Σ . To evaluate the impact of the estimation of Σ , simulations were performed to compare the performance of WLasso when Σ is known and when it is estimated. The results are displayed in Figure 6 for several values of γ (0.9, 0.95, 0.97) which is a parameter appearing in Section 2.4. In the top left part of this figure the largest difference between the True Positive Rate (TPR) and False Positive Rate (FPR) is displayed for several values of p and for $n = 50$. In the top right and bottom parts of the figure, the corresponding TPR and FPR are displayed, respectively. We can see from this figure that for the value of λ maximizing the difference between TPR and FPR and for all the values of γ , all the active variables are properly retrieved without selecting non active variables when Σ is known. In the case where Σ is estimated by using the approach described in Section 2.5, 75% of the active variables are recovered and less than 1% of non active variables are wrongly estimated as active variables. Note that the results displayed in Figure 6 are obtained when $(b, n) = (0.5, 50)$ but we obtained similar conclusions for $(b, n) = (1, 50)$, $(b, n) = (0.5, 100)$ and $(1, 100)$.

3.2. Choice of λ . For tuning the parameter λ involved in our methodology, we propose choosing the value which minimizes $\text{MSE}_{\widehat{M}(\lambda)}(\lambda)$ defined in (9). In Figure 13 of the Supplementary material, we compare the performance of our approach with this choice of λ (solid line) to the optimal one obtained when λ is chosen to yield the largest difference




FIGURE 6. $\max(\text{TPR}-\text{FPR})$ and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, $b = 0.5$ and $n = 50$. Dotted line: Σ , solid line: $\hat{\Sigma}$.

between the TPR and the FPR (dotted line). We can observe from this figure that, for the different values of γ , the TPR is a little bit smaller when λ is estimated but that the FPR is quite similar. Similar results were obtained for $(b, n) = (1, 50)$, $(b, n) = (0.5, 100)$ and $(1, 100)$. However, for larger b or n , the difference between the performance obtained by the optimal choice of λ and by our choice of λ is smaller.

3.3. Comparison with other methods. In this section, we compare our methodology with other approaches: the classical Lasso described in Tibshirani (1996) and two recently proposed methods aiming at handling the correlations between the columns of the design matrix \mathbf{X} : HOLP and Precision Lasso proposed by Wang and Leng (2016) and Wang et al. (2019), respectively. This comparison is performed by computing the TPR and FPR of these approaches for different values of the parameters involved in each of them.

The grid of λ for the classical Lasso and for our approach is provided by the `glmnet` and `genlasso` R packages, respectively. Concerning the Precision Lasso, we found for each value of n and p the λ_{\min} and λ_{\max} leading to p non null estimated coefficients and p null estimated coefficients, respectively. Then, we chose 100 values of λ uniformly distributed in the interval $[\lambda_{\min}, \lambda_{\max}]$ and we used the light implementation of the Precision Lasso. As for HOLP, β is estimated by $\hat{\beta}_{\text{HOLP}} = \mathbf{X}^T(\mathbf{X}\mathbf{X}^T)^{-1}\mathbf{y}$. Then, for each s in $\{1, \dots, p\}$, the components of β which are estimated as non null are the s largest among the $|\hat{\beta}_{\text{HOLP},j}|$, where $\hat{\beta}_{\text{HOLP},j}$ denotes the j th components of $\hat{\beta}_{\text{HOLP}}$. In this case, the parameter controlling

the sparsity level of the estimator of β is s . It has a similar role as λ in the previous approaches.




FIGURE 7. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for WLasso obtained for the λ chosen by the strategy proposed in Section 3.2 (solid line). Results obtained for the optimal choice of λ for WLasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, $b = 0.5$ and $n = 50$.

The corresponding results are displayed in Figures 7 and 8 in the case where $n = 50$ and $b = 0.5$ and Σ has the block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$ and $(0.5, 0.7, 0.9)$, respectively. The top left part of these figures displays the largest difference between TPR and FPR for different values of p , which corresponds to an optimal choice of the parameters. For WLasso, we also display the results obtained when the parameter λ is chosen by using the strategy proposed in Section 3.2, $\gamma = 0.95$ and Σ is estimated using the procedure explained in Section 2.5. The corresponding TPR and FPR for each method are displayed in the top right part and bottom part of the figures, respectively. Note that we also conducted experiments in the case where $b = 1$. Since the conclusions are very similar, the corresponding figures are given in the Supplementary material.

We can see from Figures 7 and 8 that WLasso outperforms the other methods: the TPR is one of the largest while the FPR is the smallest. HOLP has a larger TPR than WLasso. However, the associated FPR is much larger. It has moreover to be noticed that Lasso, HOLP and Precision Lasso are favored with respect to WLasso since their parameters were

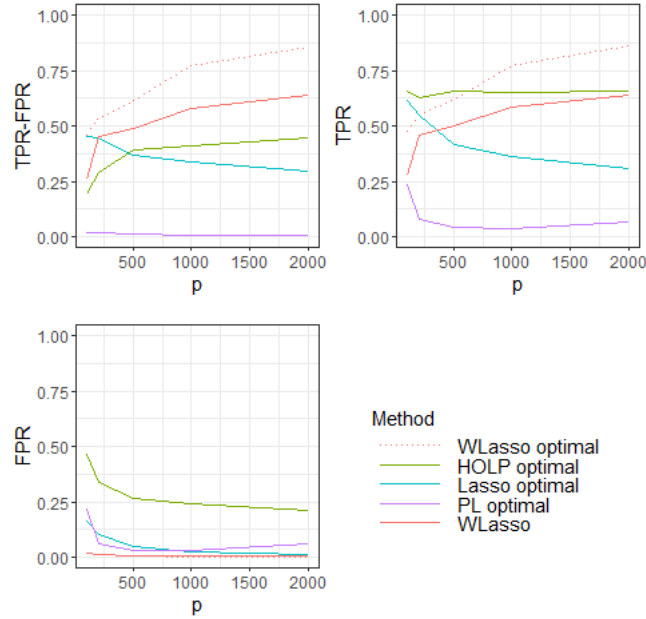


FIGURE 8. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 0.7, 0.9)$, $b = 0.5$ and $n = 50$.

chosen to optimize their performance in terms of TPR and FPR whereas, in Wlasso, the parameter λ was chosen by using the strategy of Section 3.2 and Σ was estimated.

Figure 9 displays the results when the sample size n is increased and equal to 100. We observe from this figure that the overall performance has been improved and that our approach outperforms the others especially in the case where p is large. Similar results are obtained in the case where $b = 1$ and $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 0.7, 0.9)$. We refer the reader to the Supplementary material for further details.

Figure 10 displays the performance of the different methodologies in the case where $\Sigma = \text{Id}$, $n = 50$ and $b = 0.5$, that is in the case where there is no correlation between the biomarkers (columns of \mathbf{X}). We can see from this figure that even in this case, our method, which is designed for handling the correlation between the biomarkers, obtains similar results as the Lasso except for small values of p . However, when λ is chosen to obtain optimal results in terms of the difference between TPR and FPR, our approach achieves the best performance with HOLP. In the case where $n = 100$, our approach obtains the best results, see the Supplementary material.

3.4. Numerical performance. Figure 11 displays the computational times of our approach implemented in the R package `Wlasso` for different values of p and of the parameter

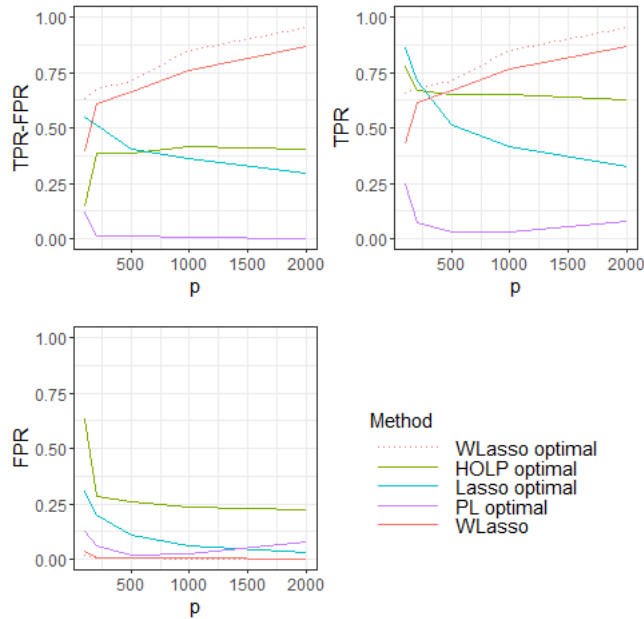


FIGURE 9. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, $b = 0.5$ and $n = 100$.

“maxsteps” (maximum number of steps/ λ s considered in the algorithm) involved in the `genlasso` R package. The timings were obtained on a workstation with 8GB of RAM and Intel Core i5 (2.4GHz) CPU. We can see from this figure that it takes only 6 minutes for processing data with our approach when $n = 50$ and $p = 2000$.

Moreover, we can observe from Figure 12 that the most time consuming step of Wlasso is the one where the generalized Lasso criterion is used (blue part in Figure 12). However, the computational time of this step was divided by two when the parameter “maxsteps” was changed from 2000 (default value) to 500 without changing the variable selection results.

4. APPLICATION TO GENE EXPRESSION DATA IN BREAST CANCER

We applied the previously detailed methods to publicly available data at Gene Expression Omnibus database (www.ncbi.nlm.nih.gov/geo), with accession code GSE2990, see Sotiriou et al. (2006). A total of $n = 189$ tumor samples from patients with breast cancer were available and their microarray data have been collected on 2,283 probes. Expression data has been preprocessed and normalized as in the original publication. A filtering step based on the interquartile range (IQR) was considered to remove some probes as in Gentleman et al. (2005). We removed probes with $\text{IQR} < 1.5$ and those which lack of annotation.

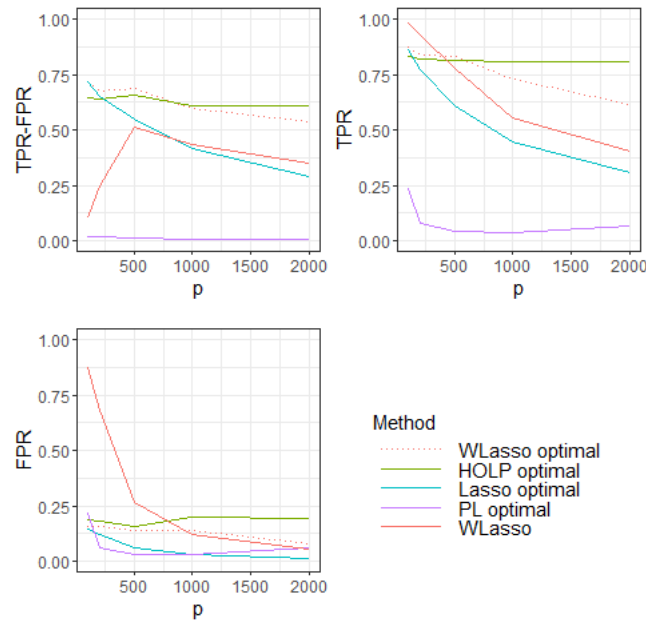


FIGURE 10. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when $\Sigma = \text{Id}$, $b = 0.5$ and $n = 50$.




FIGURE 11. Computational time of our approach Wlasso when the parameter “maxsteps” has the default value, namely 2000 (dotted line) and when maxsteps=500 (solid line).

The remaining $p = 1,112$ probes were then standardized. The goal of the application is to identify genes associated with the development of breast cancer. To this end, the expression of the gene BRCA1 (BREast CANcer gene 1) was considered as the variable \mathbf{y} to explain (response variable). BRCA1 is a well known human tumor suppressor gene that helps to repair damaged DNA, thus the mutation of this gene can notably increase the risk




FIGURE 12. Time allocation for each part of our method Wlasso for $p = 2000$ and two values of the parameter “maxsteps”. The time spent in the first three steps described in Section 2.6 is displayed in the first three bars in red, green, blue, respectively except the correction stage of Steps 3 and 4 (Equations (7) and (8)) which is given in the fourth bar in purple.

of breast cancer, see Rosen et al. (2003). The standardized 1,112 probes were considered as explanatory variables and their correlation is displayed in the Supplementary material.

We compared the performance of the approaches investigated in Section 3 in terms of genes selection. For our approach Wlasso, we used the methodology described in Section 2.5 for estimating the correlation matrix Σ . Assuming that it has the block-wise correlation structure (5), we estimated the coefficients α_1 , α_2 and α_3 by $\hat{\alpha}_1 = 0.17$, $\hat{\alpha}_2 = 0.21$ and $\hat{\alpha}_3 = 0.52$, respectively. As for the Lasso method, the parameter λ was chosen by cross-validation and the number of variables to be selected was fixed to 50 for Precision Lasso and HOLP in order to select approximately the same number of variables as with the Lasso.

Table 1 given in the Supplementary material provides the list of genes corresponding to the selected probes for each method. Unfortunately, HOLP could not provide any results since it requires the computation of the inverse of the matrix $\mathbf{X}\mathbf{X}^T$ which is not invertible in this case. The matrix \mathbf{X}^T is indeed not full rank in this dataset. The database for annotation, visualization and integrated discovery (DAVID) version 6.8 proposed by Dennis et al. (2003) was used to highlight the genes potentially related to the development of breast cancer, which correspond to true positives. Based on this tool, 8 genes were considered as true positives for Wlasso and Lasso which have only one common gene: GSTT1. However, Lasso selected more false positives than Wlasso (44 vs. 26). Among the 50 variables selected by Precision Lasso, 6 are identified as true positives by DAVID including the well-known ESR1 gene in breast cancer. Based on this annotation, Precision Lasso detected less true positives than Lasso. Moreover, Wlasso identified more true positives with less false positives than Precision Lasso. Nevertheless, this application is for illustration purpose only and not for a fair evaluation or comparison of the methods.

5. CONCLUSION

In this paper, we proposed an innovative, efficient and fully data-driven method to deal with the variable selection issue in high-dimensional frameworks where the active variables

are highly correlated with the non-active ones which is implemented in the `WLasso` R package available from the CRAN. The proposed `WLasso` method has been assessed and compared with other methods in a simulation study with several scenarios. In the highly correlated setting, `WLasso` successfully identifies more true positives with limited false positives as compared with the classical Lasso. Contrary to `HOLP`, `WLasso` still works when several columns are linearly dependent and does not suffer from the inflation of noise introduced by the preconditioning. Compared with the recent Precision Lasso approach, which aims to deal with the same issue, `WLasso` obtained better results in terms of selection accuracy in the different settings considered. `WLasso` is also very computationally efficient and demonstrated its abilities to properly identify genes related to breast cancer from a publicly available gene expression dataset. However, the following directions could be considered to improve its performance.

Firstly, the method that we used for estimating Σ could be improved by using more sophisticated approaches such as Perrot-Dockès et al. (2020). Secondly, our way of choosing the parameter λ for the final model selection could also be improved by considering cross-validation or stability selection. Until now, a simple approach has been considered to avoid computation time and performed quite well especially for moderate to high simple size. Thirdly, most of the computational time of our approach is spent in the application of the generalized Lasso criterion. Hence, for an application to genomic datasets having more than twenty thousands of variables, it could be worth speeding it up. This will be the subject of future work.

FUNDING

This work was supported by the Association Nationale Recherche Technologie (ANRT).

REFERENCES

- Dennis, G., B. Sherman, D. Hosack, J. Yang, W. Gao, H. Lane, and R. Lempicki (2003, 02). David: Database for annotation, visualization, and integrated discovery. Genome biology 4, R60.
- Fan, J. and R. Li (2001, 02). Variable selection via nonconcave penalized likelihood and its oracle properties. J. Am. Stat. Assoc. 96, 1348–1360.
- Fan, J. and R. Li (2006, 03). Statistical challenges with high dimensionality: Feature selection in knowledge discovery. Proc. Madrid Int. Congress of Mathematicians 3, 595–622.
- Gentleman, R., V. Carey, W. Huber, R. Irizarry, and S. Dudoit (2005). Bioinformatics and Computational Biology Solutions Using R and Bioconductor (Statistics for Biology and Health). Berlin, Heidelberg: Springer-Verlag.
- Heinze, G., C. Wallisch, and D. Dunkler (2018, 01). Variable selection - a review and recommendations for the practicing statistician. Biometrical J. 60(3), 1–19.
- Jia, J. and K. Rohe (2015). Preconditioning the lasso for sign consistency. Electron. J. Statist. 9(1), 1150–1172.

- Kalia, M. (2015). Biomarkers for personalized oncology: recent advances and future challenges. Metabolism 64(3), S16 – S21. Biomarkers: Current Status and Future Trends.
- McDonald, J. (2009, 01). Handbook of Biological Statistics 2nd Edition. Sparky House Publishing Baltimore.
- Michalopoulos, I., G. Pavlopoulos, A. Malatras, A. Karelas, M. Kostadima, R. Schneider, and S. Kossida (2012, 06). Human gene correlation analysis (hgca): A tool for the identification of transcriptionally co-expressed genes. BMC research notes 5, 265.
- Perrot-Dockès, M., C. Lévy-Leduc, and L. Rajjou (2020). Estimation of large block structured covariance matrices: Application to "multi-omic" approaches to study seed quality. arXiv:1806.10093.
- Rosen, E., S. Fan, R. Pestell, and I. Goldberg (2003, 07). BRCA1 gene in breast cancer. Journal of cellular physiology 196(1), 19–41.
- Saeyns, Y., I. Inza, and P. Larranaga (2007). A review of feature selection techniques in bioinformatics. Bioinformatics 23(19), 2507–2517.
- Smith, G. (2018, 09). Step away from stepwise. J. Big Data 5(32), 1–12.
- Sotiriou, C., P. Wirapati, S. Loi, A. Harris, S. Fox, J. Smeds, H. Nordgren, P. Farmer, V. Praz, B. Haibe-Kains, C. Desmedt, D. Larsimont, F. Cardoso, H. Peterse, D. Nuyten, M. Buyse, M. J. Van de Vijver, J. Bergh, M. Piccart, and M. Delorenzi (2006, 02). Gene Expression Profiling in Breast Cancer: Understanding the Molecular Basis of Histologic Grade To Improve Prognosis. JNCI: Journal of the National Cancer Institute 98(4), 262–272.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. J. R. Stat. Soc. Ser. B (Stat. Methodol.) 58(1), 267–288.
- Tibshirani, R. J. and J. Taylor (2011, 06). The solution path of the generalized lasso. Ann. Stat. 39(3), 1335–1371.
- Wang, H., B. Lengerich, B. Aragam, and E. Xing (2019, 09). Precision lasso: Accounting for correlations and linear dependencies in high-dimensional genomic data. Bioinformatics 35(7), 1181–1187.
- Wang, X. and C. Leng (2016). High dimensional ordinary least squares projection for screening variables. J. R. Stat. Soc. Ser. B (Stat. Methodol.) 78(3), 589–611.
- Zhao, P. and B. Yu (2006, 12). On model selection consistency of lasso. J. Machine Learn. Res. 7, 2541–2563.
- Zou, H. and T. Hastie (2015, 01). Regularization and variable selection via the elastic nets. J. R. Stat. Soc. Ser. B (Stat. Methodol.) 67, 301–320.

SUPPLEMENTARY MATERIAL

This supplementary material provides additional figures and a table for the paper: “A variable selection approach for highly correlated predictors in high-dimensional genomic data”.

Figure 13 illustrates Section 3.2.

Figures 14, 15, 16, 17, 18, 19, 20 and 21 provide similar results as those displayed in Figure 7 of the paper in the following cases:

- Block-wise correlation structure for Σ with $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$ with
 - $b = 1$ and $n = 50$
 - $b = 1$ and $n = 100$
- Block-wise correlation structure for Σ with $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 0.7, 0.9)$ with
 - $b = 1$ and $n = 50$
 - $b = 0.5$ and $n = 100$
 - $b = 1$ and $n = 100$
- Σ is equal to identity with
 - $b = 1$ and $n = 50$
 - $b = 0.5$ and $n = 100$
 - $b = 1$ and $n = 100$

Figures 22 and Table 1 give additional information for the Application Section.




FIGURE 13. (TPR-FPR) and the corresponding True Positive Rate and False Positive Rate for $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, $b = 0.5$ and $n = 50$. Dotted line: optimal choice of λ , solid line: choice of λ explained in Section 3.2 of the paper.




FIGURE 14. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 of the paper (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) of the paper with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, $b = 1$ and $n = 50$.

UMR MIA-PARIS, AGROPARISTECH, INRAE, UNIVERSITÉ PARIS-SACLAY, 75005, PARIS, FRANCE
E-mail address: wencan.zhu@agroparistech.fr

UMR MIA-PARIS, AGROPARISTECH, INRAE, UNIVERSITÉ PARIS-SACLAY, 75005, PARIS, FRANCE
E-mail address: celine.levy-leduc@agroparistech.fr

BIostatistics and Programming Department, SANOFI R&D, 91380 CHILLY MAZARIN, FRANCE
E-mail address: nils.ternes@sanofi.com




FIGURE 15. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.3, 0.5, 0.7)$, $b = 1$ and $n = 100$.

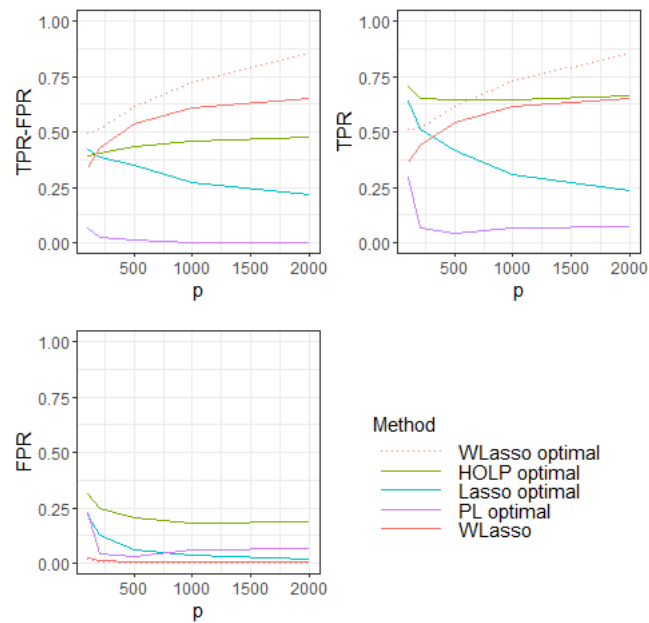


FIGURE 16. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 0.7, 0.9)$, $b = 1$ and $n = 50$.




FIGURE 17. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 of the paper (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) of the paper with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 0.7, 0.9)$, $b = 0.5$ and $n = 100$.

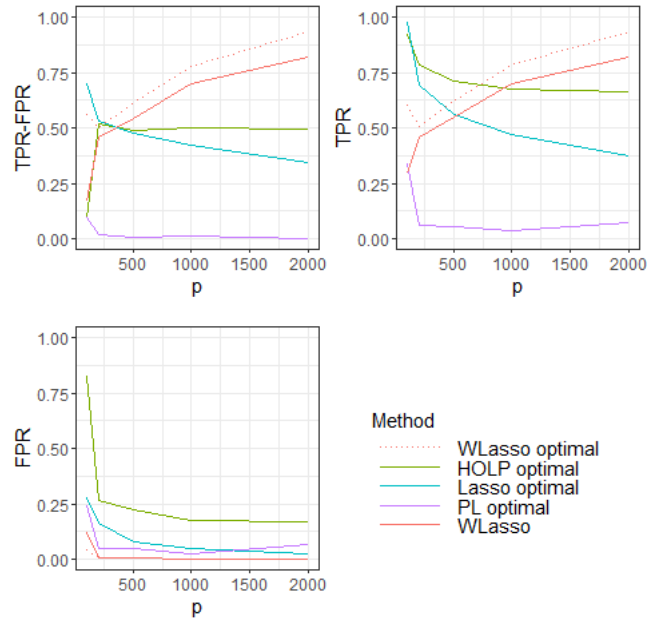


FIGURE 18. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and $(\text{TPR}-\text{FPR})$ for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 of the paper (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when Σ has the block-wise correlation structure defined in (5) of the paper with parameters $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 0.7, 0.9)$, $b = 1$ and $n = 100$.




FIGURE 19. Top left: $\max(\text{TPR-FPR})$ for Lasso, HOLP, Precision Lasso (PL) and TPR-FPR for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 of the paper (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when $\Sigma = \text{Id}$, $b = 1$ and $n = 50$.




FIGURE 20. Top left: $\max(\text{TPR-FPR})$ for Lasso, HOLP, Precision Lasso (PL) and TPR-FPR for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 of the paper (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when $\Sigma = \text{Id}$, $b = 0.5$ and $n = 100$.




FIGURE 21. Top left: $\max(\text{TPR}-\text{FPR})$ for Lasso, HOLP, Precision Lasso (PL) and TPR-FPR for Wlasso obtained for the λ chosen by the strategy proposed in Section 3.2 of the paper (solid line). Results obtained for the optimal choice of λ for Wlasso (dotted line). Corresponding TPR (top right) and FPR (bottom) when $\Sigma = \text{Id}$, $b = 1$ and $n = 100$.




FIGURE 22. Heatmap of correlations between the probes.

TABLE 1. Selected genes from each method, genes related to breast cancer are in bold.

	selected genes
WLasso	TOP2A , NPEPPS, INSIG1, PLK2 , TRIM2, CYP1B1 , KIF5C, ATXN1, PTTG1 , NUCB2, SEMA3C, GSTT1 , LGALS4, KIF11 , NEK2 , RAD51 , DNALI1, AGFG2, IAPP, IFI16, CITED2, SLC19A2, WSB1, RET, FRMD4B, MYBL1, SELENBP1, CALU, ZNF451, IGHM, AZGP1, NCAPG, ANO1, ZNF226
Lasso	CCL5, WSB1, MAPKAPK2, FHL1, TNC, ALCAM, ATP6V1A, RIOK3, RGS2, USP1 , TRIM14, LPL , GGH, GSTT1 , SMN1, MMP11, PSMB9, DST, RAD54L , TFAP2A, DSC2, GABRP, INPP4B, GHR , CD36, PTP4A3, ASS1, H2BFS, ANXA3, CXCL12 , TPX2 , CACNA1D, GOLGA8A, ATG5, SC5D, PTN, C1R, WWP1, DPT, MUC1 , LRRC15, SMA4, CCZ1B, ASNSD1, COPS4, PSD3, VAV3, MS4A4A, KLF13, QPCTL, PLA2G12A, MUM1
PL	H2AFZ, BTG2, SDC1, IGF2R , INSIG1, ALCAM, SDHD, ACADM, FBLN1, SNRPE, N4BP2L2, SLC25A37, IGF1R , PPFIBP1, SYT1, SORBS2, AGA, IFI44L, TFF3, DSC2, MYB, ESR1 , SETBP1, FGFR2 , GOLGA8A, DNAJB6, CD24, HLA-DQB1 , JUP, SULF1, RDX, COL14A1, NFIB, COL6A1, KRT6B, NBPF10, N4BP2L2, DGKH, RGS1, ZSCAN18, LZTFL1, GREM1 , C1orf115, DUSP12, KLHL2, ARMC9, DENND1B, TMPRSS3, WIZ, NTRK2
HOLP	