



HAL
open science

Identification of prognostic and predictive biomarkers in high-dimensional data with PPLasso

Wencan Zhu, Céline Lévy-Leduc, Nils Ternès

► **To cite this version:**

Wencan Zhu, Céline Lévy-Leduc, Nils Ternès. Identification of prognostic and predictive biomarkers in high-dimensional data with PPLasso. *BMC Bioinformatics*, 2023, 24 (1), pp.25. <10.1186/s12859-023-05143-0>. <hal-03559682>

HAL Id: hal-03559682

<https://hal.science/hal-03559682v1>

Submitted on 7 Feb 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.



HAL Authorization

IDENTIFICATION OF PROGNOSTIC AND PREDICTIVE BIOMARKERS IN HIGH-DIMENSIONAL DATA WITH PPLASSO

WENCAN ZHU, CÉLINE LÉVY-LEDUC, AND NILS TERNÈS

ABSTRACT. In clinical trials, identification of prognostic and predictive biomarkers is essential to precision medicine. Prognostic biomarkers can be useful for the prevention of the occurrence of the disease, and predictive biomarkers can be used to identify patients with potential benefit from the treatment. Previous researches were mainly focused on clinical characteristics, and the use of genomic data in such an area is hardly studied. A new method is required to simultaneously select prognostic and predictive biomarkers in high dimensional genomic data where biomarkers are highly correlated. We propose a novel approach called PPLasso (Prognostic Predictive Lasso) integrating prognostic and predictive effects into one statistical model. PPLasso also takes into account the correlations between biomarkers that can alter the biomarker selection accuracy. Our method consists in transforming the design matrix to remove the correlations between the biomarkers before applying the generalized Lasso. In a comprehensive numerical evaluation, we show that PPLasso outperforms the traditional Lasso approach on both prognostic and predictive biomarker identification in various scenarios. Finally, our method is applied to publicly available transcriptomic data from clinical trial RV144. Our method is implemented in the PPLasso R package which will be soon available from the Comprehensive R Archive Network (CRAN).

1. INTRODUCTION

With the advancement of precision medicine, there has been an increasing interest in identifying prognostic or predictive biomarkers. A prognostic biomarker is linked to a clinical outcome (disease recurrence, progression of the disease etc.) in the presence or in the absence of a therapy while a predictive biomarker is associated with a response or a lack of response to a specific therapy. Ballman (2015) and Clark (2008) provided a comprehensive explanation and concrete examples to distinguish prognostic from predictive biomarkers, respectively.

Concerning the biomarker selection, the high dimensionality of genomic data is one of the main challenges as explained in Fan and Li (2006). To identify effective biomarkers in high-dimensional settings, several approaches can be considered including hypothesis-based tests described in McDonald (2009), wrapper approaches proposed in Saeys *et al.* (2007), and penalized approaches such as Lasso designed by Tibshirani (1996) among others. Hypothesis-based tests consider each biomarker independently and thus ignore potential correlations between them. Wrapper approaches often show high risk of overfitting and are computationally expensive for high-dimensional data as explained in Smith (2018).

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

More efforts have been devoted to penalized methods given their ability to automatically perform variable selection and coefficient estimation simultaneously as highlighted in Fan and Lv (2009). However, Lasso showed some potential drawbacks when biomarkers are highly correlated. Particularly, when the Irrepresentable Condition (IC) proposed by Zhao and Yu (2006) is violated, Lasso can not guarantee to correctly identify true effective biomarkers. In genomic data, biomarkers are usually highly correlated such that this condition can hardly be satisfied, see Wang *et al.* (2019). Several methods have been proposed to address this issue. For example, Wang and Leng (2016) proposed the HOLP approach which consists in removing the correlation between the columns of the design matrix; Wang *et al.* (2019) proposed to handle the correlation by assigning similar weights to correlated variables in their approach called Precision Lasso; Zhu *et al.* (2021) proposed to remove the correlations by applying a whitening transformation to the data before using the generalized Lasso criterion designed by Tibshirani and Taylor (2011).

The challenge of finding prognostic biomarkers has been extensively explored with previously introduced methods, however, the discovery of predictive biomarkers has seen much less attention. Limited to binary endpoint, Foster *et al.* (2011) proposed to first predict response probabilities for treatment and use this probability as the response in a classification problem to find effective biomarkers. Tian *et al.* (2012) proposed a new method to detect interaction between the treatment and the biomarkers by modifying the covariates. This method can be implemented on continuous/binary/time-to-event endpoint. Lipkovich *et al.* (2011) proposed a method called SIDES, which adopts a recursive partitioning algorithm for screening treatment-by-biomarker interactions. This method was further improved in Lipkovich and Dmitrienko (2014) by adding another step of preselection on predictive biomarkers based on variable importance. The method was demonstrated with continuous endpoint. More recently, Sechidis *et al.* (2018) applied approaches coming from information theory for ranking biomarkers on their prognostic/predictive strength. Their method is applicable only for binary or time-to-event endpoint. Moreover, all of these methods were assessed under the situation where the sample size is relatively large and the number of biomarkers is limited, which is hardly the case for genomic data.

In the literature mentioned above, the authors focused on one of the problematic of identifying prognostic or predictive biomarkers, but rarely on both. In this paper, we propose a novel approach called PPLasso (Predictive Prognostic Lasso) to simultaneously identify prognostic and predictive biomarkers in a high dimensional setting with continuous endpoints, as presented in Section 2. Extensive numerical experiments are given in Section 3 to assess the performance of our approach and to compare it to other methods. PPLasso is also applied to the clinical trial RV144 in Section 4. Finally, we give concluding remarks in Section 5.

2. METHODS

In this section, we propose a novel approach called PPLasso (Predictive Prognostic Lasso) which consists in writing the identification of predictive and prognostic biomarkers as a variable selection problem in an ANCOVA (Analysis of Covariance) type model **2.1. Statistical modeling.** Let \mathbf{y} be a continuous response or endpoint and t_1, t_2 two treatments. Let also \mathbf{X}_1 (resp. \mathbf{X}_2) denote the design matrix for the n_1 (resp. n_2) patients with treatment t_1 (resp. t_2), each containing measurements on p candidate biomarkers:

$$(1) \quad \mathbf{X}_1 = \begin{bmatrix} X_{11}^1 & X_{11}^2 & \cdots & X_{11}^p \\ X_{12}^1 & X_{12}^2 & \cdots & X_{12}^p \\ \vdots & \vdots & \ddots & \vdots \\ X_{1n_1}^1 & X_{1n_1}^2 & \cdots & X_{1n_1}^p \end{bmatrix}, \mathbf{X}_2 = \begin{bmatrix} X_{21}^1 & X_{21}^2 & \cdots & X_{21}^p \\ X_{22}^1 & X_{22}^2 & \cdots & X_{22}^p \\ \vdots & \vdots & \ddots & \vdots \\ X_{2n_2}^1 & X_{2n_2}^2 & \cdots & X_{2n_2}^p \end{bmatrix}.$$

To take into account the potential correlation that may exist between the biomarkers in the different treatments, we shall assume that the rows of \mathbf{X}_1 (resp. \mathbf{X}_2) are independent centered Gaussian random vectors with a covariance matrix equal to Σ_1 (resp. Σ_2).

To model the link that exists between \mathbf{y} and the different types of biomarkers we propose using the following model:

$$(2) \quad \mathbf{y} = \begin{pmatrix} y_{11} \\ y_{12} \\ \vdots \\ y_{1n_1} \\ y_{21} \\ y_{22} \\ \vdots \\ y_{2n_2} \end{pmatrix} = \mathbf{X} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \beta_{11} \\ \beta_{12} \\ \vdots \\ \beta_{1p} \\ \beta_{21} \\ \beta_{22} \\ \vdots \\ \beta_{2p} \end{pmatrix} + \begin{pmatrix} \epsilon_{11} \\ \epsilon_{12} \\ \vdots \\ \epsilon_{1n_1} \\ \epsilon_{21} \\ \epsilon_{22} \\ \vdots \\ \epsilon_{2n_2} \end{pmatrix},$$

where $(y_{i1}, \dots, y_{in_i})$ corresponds to the response of patients with treatment t_i , i being equal to 1 or 2,

$$\mathbf{X} = \begin{bmatrix} 1 & 0 & X_{11}^1 & X_{11}^2 & \cdots & X_{11}^p & 0 & 0 & \cdots & 0 \\ 1 & 0 & X_{12}^1 & X_{12}^2 & \cdots & X_{12}^p & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & X_{1n_1}^1 & X_{1n_1}^2 & \cdots & X_{1n_1}^p & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 & \cdots & 0 & X_{21}^1 & X_{21}^2 & \cdots & X_{21}^p \\ 0 & 1 & 0 & 0 & \cdots & 0 & X_{22}^1 & X_{22}^2 & \cdots & X_{22}^p \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 1 & 0 & 0 & \cdots & 0 & X_{2n_2}^1 & X_{2n_2}^2 & \cdots & X_{2n_2}^p \end{bmatrix},$$

with α_1 (resp. α_2) corresponding to the effects of treatment t_1 (resp. t_2). Moreover, $\beta_1 = (\beta_{11}, \beta_{12}, \dots, \beta_{1p})'$ (resp. $\beta_2 = (\beta_{21}, \beta_{22}, \dots, \beta_{2p})'$) are the coefficients associated to each of the p biomarkers in treatment t_1 (resp. t_2) group, ' denoting the matrix transposition and $\epsilon_{11}, \dots, \epsilon_{2n_2}$ are standard independent Gaussian random variables independent of \mathbf{X}_1 and

\mathbf{X}_2 . When t_1 stands for the standard treatment or placebo, prognostic (resp. predictive) biomarkers are defined as those having non-zero coefficients in β_1 (resp. in $\beta_2 - \beta_1$) and non prognostic (resp. non predictive) biomarkers correspond to the indices having null coefficients in β_1 (resp. in $\beta_2 - \beta_1$). Model (2) can be written as:

$$(3) \quad \mathbf{y} = \mathbf{X}\boldsymbol{\gamma} + \boldsymbol{\epsilon},$$

with $\boldsymbol{\gamma} = (\alpha_1, \alpha_2, \beta_1', \beta_2')'$. The Lasso penalty is a well-known approach to estimate coefficients with a sparsity enforcing constraint allowing variable selection by estimating some coefficients by zero. It consists in minimizing the following penalized least-squares criterion (Tibshirani (1996)):

$$(4) \quad \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\gamma}\|_2^2 + \lambda \|\boldsymbol{\gamma}\|_1,$$

where $\|\mathbf{u}\|_2^2 = \sum_{i=1}^n u_i^2$ and $\|\mathbf{u}\|_1 = \sum_{i=1}^n |u_i|$ for $\mathbf{u} = (u_1, \dots, u_n)$. A different sparsity constraint was applied to β_1 and $\beta_2 - \beta_1$ to allow different sparsity levels. Hence we propose to replace the penalty $\lambda \|\boldsymbol{\gamma}\|_1$ in (4) by

$$(5) \quad \lambda_1 \|\beta_1\|_1 + \lambda_2 \|\beta_2 - \beta_1\|_1.$$

Thus, a first estimator of $\boldsymbol{\gamma}$ could be found by minimizing the following criterion with respect to $\boldsymbol{\gamma}$:

$$(6) \quad \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\gamma}\|_2^2 + \lambda_1 \left\| \begin{bmatrix} \mathbf{0}_{p,1} & \mathbf{0}_{p,1} & D_1 \\ \mathbf{0}_{p,1} & \mathbf{0}_{p,1} & \frac{\lambda_2}{\lambda_1} D_2 \end{bmatrix} \boldsymbol{\gamma} \right\|_1,$$

where $D_1 = [\text{Id}_p, \mathbf{0}_{p,p}]$ and $D_2 = [-\text{Id}_p, \text{Id}_p]$, with Id_p denoting the identity matrix of size p and $\mathbf{0}_{i,j}$ denoting a matrix having i rows and j columns and containing only zeros. However, since the inconsistency of Lasso biomarker selection is originated from the correlations between the biomarkers, we propose to remove the correlation by “whitening” the matrix \mathbf{X} . More precisely, we consider $\tilde{\mathbf{X}} = \mathbf{X}\boldsymbol{\Sigma}^{-1/2}$, where

$$(7) \quad \boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \boldsymbol{\Sigma}_1 & 0 \\ 0 & 0 & 0 & \boldsymbol{\Sigma}_2 \end{bmatrix}$$

and define $\boldsymbol{\Sigma}^{-1/2}$ by replacing in (7) $\boldsymbol{\Sigma}_i$ by $\boldsymbol{\Sigma}_i^{-1/2}$, where $\boldsymbol{\Sigma}_i^{-1/2} = \mathbf{U}_i \mathbf{D}_i^{-1/2} \mathbf{U}_i^T$, \mathbf{U}_i and \mathbf{D}_i being the matrices involved in the spectral decomposition of $\boldsymbol{\Sigma}_i$ for $i = 1$ or 2 . With such a transformation the columns of $\tilde{\mathbf{X}}$ are decorrelated and Model (3) can be rewritten as follows:

$$(8) \quad \mathbf{y} = \tilde{\mathbf{X}}\tilde{\boldsymbol{\gamma}} + \boldsymbol{\epsilon}$$

where $\tilde{\boldsymbol{\gamma}} = \boldsymbol{\Sigma}^{1/2}\boldsymbol{\gamma}$. The objective function (6) thus becomes:

$$(9) \quad L_{\lambda_1, \lambda_2}^{\text{PPLasso}}(\tilde{\boldsymbol{\gamma}}) = \frac{1}{2} \|\mathbf{y} - \tilde{\mathbf{X}}\tilde{\boldsymbol{\gamma}}\|_2^2 + \lambda_1 \left\| \begin{bmatrix} \mathbf{0}_{p,1} & \mathbf{0}_{p,1} & D_1 \\ \mathbf{0}_{p,1} & \mathbf{0}_{p,1} & \frac{\lambda_2}{\lambda_1} D_2 \end{bmatrix} \boldsymbol{\Sigma}^{-1/2} \tilde{\boldsymbol{\gamma}} \right\|_1.$$

2.2. **Estimation of $\tilde{\gamma}$.** Let us define a first estimator of $\tilde{\gamma} = (\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\beta}'_1, \tilde{\beta}'_2)$ as follows:

$$(10) \quad \hat{\tilde{\gamma}}_0(\lambda_1, \lambda_2) = (\hat{\tilde{\alpha}}_1, \hat{\tilde{\alpha}}_2, \hat{\tilde{\beta}}'_{10}, \hat{\tilde{\beta}}'_{20}) = \text{Argmin}_{\tilde{\gamma}} L_{\lambda_1, \lambda_2}^{\text{PPLasso}}(\tilde{\gamma}),$$

for each fixed λ_1 and λ_2 . To better estimate $\tilde{\beta}_1$ and $\tilde{\beta}_2$, a thresholding was applied to $\hat{\tilde{\beta}}_0(\lambda_1, \lambda_2) = (\hat{\tilde{\beta}}_{10}(\lambda_1, \lambda_2)', \hat{\tilde{\beta}}_{20}(\lambda_1, \lambda_2)')$. For K_1 (resp. K_2) in $\{1, \dots, p\}$, let Top_{K_1} (resp. Top_{K_2}) be the set of indices corresponding to the K_1 (resp. K_2) largest values of the components of $|\hat{\tilde{\beta}}_{10}(\lambda_1, \lambda_2)|$ (resp. $|\hat{\tilde{\beta}}_{20}(\lambda_1, \lambda_2)|$), then the estimator of $\tilde{\beta} = (\tilde{\beta}'_1, \tilde{\beta}'_2)$ after the correction is denoted by $\hat{\tilde{\beta}}(\lambda_1, \lambda_2) = (\hat{\tilde{\beta}}_1^{(K_1)}(\lambda_1, \lambda_2), \hat{\tilde{\beta}}_2^{(K_2)}(\lambda_1, \lambda_2))$ where the j th component of $\hat{\tilde{\beta}}_i^{(K_i)}(\lambda_1, \lambda_2)$, for $i = 1$ or 2 , is defined by:

$$(11) \quad \hat{\tilde{\beta}}_{ij}^{(K_i)}(\lambda_1, \lambda_2) = \begin{cases} \hat{\tilde{\beta}}_{i0j}(\lambda_1, \lambda_2), & j \in \text{Top}_{K_i} \\ K_1\text{th largest value of } |\hat{\tilde{\beta}}_{i0j}(\lambda_1, \lambda_2)|, & j \notin \text{Top}_{K_i}. \end{cases}$$

Note that the corrections are only performed on $\hat{\tilde{\beta}}_0$, the estimators $\hat{\tilde{\alpha}}_1$ and $\hat{\tilde{\alpha}}_2$ were not modified.

To illustrate the interest of using a thresholding step, we generated a dataset based on Model 3 with parameters described in Section 3.1 and $p = 500$. Moreover, to simplify the graphical illustrations, we focus on the case where $\lambda_1 = \lambda_2 = \lambda$. Figure 1 displays the estimation error associated to the estimators of $\tilde{\beta}(\lambda)$ before and after the thresholding. We can see from this figure that the estimation of $\tilde{\beta}(\lambda)$ is less biased after the correction. The choice of K_1 and K_2 will be explained in Section 2.4.

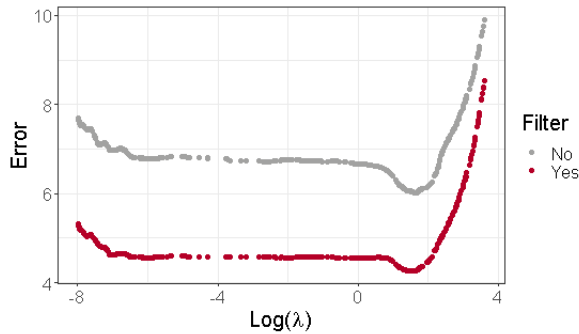


FIGURE 1. Estimation error $\left\| \hat{\tilde{\beta}}_0(\lambda) - \tilde{\beta} \right\|_2$ (gray) and $\left\| \hat{\tilde{\beta}}(\lambda) - \tilde{\beta} \right\|_2$ (red) for all λ .

2.3. **Estimation of γ .** With $\widehat{\beta} = (\widehat{\beta}'_1, \widehat{\beta}'_2)$, the estimators of β_1 and $\beta_2 - \beta_1$ can be obtained by $\widehat{\beta}_{10} = \Sigma_1^{-1/2} \widehat{\beta}_1$ and $(\widehat{\beta}_{20} - \widehat{\beta}_{10}) = \Sigma_2^{-1/2} \widehat{\beta}_2 - \Sigma_1^{-1/2} \widehat{\beta}_1$. As previously, another thresholding was applied to $\widehat{\beta}_{10}$ and $\widehat{\beta}_{20}$: for $i = 1$ or 2 ,

$$(12) \quad \widehat{\beta}_{ij}^{(M_i)}(\lambda_1, \lambda_2) = \begin{cases} \widehat{\beta}_{i0j}(\lambda_1, \lambda_2), & j \in \text{Top}_{M_i} \\ 0, & j \notin \text{Top}_{M_i}, \end{cases}$$

for each fixed λ_1 and λ_2 . The biomarkers with non-zero coefficients in $\widehat{\beta}_1 = \widehat{\beta}_1^{(M_1)}$ (resp. $\widehat{\beta}_2^{(M_2)} - \widehat{\beta}_1^{(M_1)}$) are considered as prognostic (resp. predictive) biomarkers, where the choice of M_1 and M_2 is explained in Section 2.4.

To illustrate the benefits of using an additional thresholding step, we used the dataset described in Section 2.2. Moreover, to simplify the graphical illustrations, we also focus on the case where $\lambda_1 = \lambda_2 = \lambda$. Figure 8 in the Supplementary material displays the number of True Positive (TP) and False Positive (FP) in prognostic and predictive biomarker identification with and without the second thresholding. We can see from this figure that the thresholding stage limits the number of false positives. Note that α_1 and α_2 are estimated by $\widehat{\alpha}_1$ and $\widehat{\alpha}_2$ defined in (10).

2.4. **Choice of the parameters K_1, K_2, M_1 and M_2 .** For each (λ_1, λ_2) and each K_1 , we computed:

$$(13) \quad \widehat{\text{MSE}}_{K_1, K_2}(\lambda_1, \lambda_2) = \|\mathbf{y} - \widetilde{\mathbf{X}} \widehat{\gamma}^{(K_1, K_2)}(\lambda_1, \lambda_2)\|_2^2,$$

where $\widehat{\gamma}^{(K_1, K_2)}(\lambda_1, \lambda_2) = (\widehat{\alpha}_1, \widehat{\alpha}_2, \widehat{\beta}_1^{(K_1)'}, \widehat{\beta}_2^{(K_2)'})$ defined in (10) and in (11). It is displayed in the left part of Figure 2.

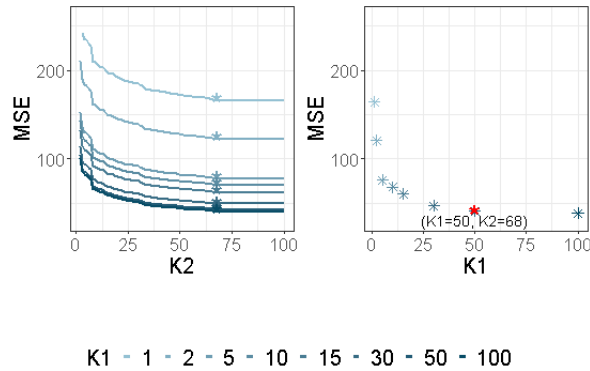


FIGURE 2. Illustration of how to choose K_1 and K_2 ($\delta = 0.95$), final choice is marked with '*'.

For each λ_1, λ_2 and a given $\delta \in (0, 1)$, the parameter \widehat{K}_2 is then chosen as follows for each K_1 :

$$\widehat{K}_2(\lambda_1, \lambda_2) = \text{Argmin} \left\{ K_2 \geq 1 \text{ s.t. } \frac{\widetilde{\text{MSE}}_{(K_1, K_2+1)}(\lambda_1, \lambda_2)}{\widetilde{\text{MSE}}_{(K_1, K_2)}(\lambda_1, \lambda_2)} \geq \delta \right\}.$$

The \widehat{K}_2 associated to each K_1 are displayed with '*' in the left part of Figure 2. Then \widehat{K}_1 is chosen by using a similar criterion:

$$\widehat{K}_1(\lambda_1, \lambda_2) = \text{Argmin} \left\{ K_1 \geq 1 \text{ s.t. } \frac{\widetilde{\text{MSE}}_{(K_1+1, \widehat{K}_2)}(\lambda_1, \lambda_2)}{\widetilde{\text{MSE}}_{(K_1, \widehat{K}_2)}(\lambda_1, \lambda_2)} \geq \delta \right\}.$$

The values of $\widetilde{\text{MSE}}_{(K_1, \widehat{K}_2)}(\lambda_1, \lambda_2)$ are displayed in the right part of Figure 2 in the particular case where $\lambda_1 = \lambda_2 = \lambda$, $\delta = 0.95$ and with the same dataset as the one used in Section 2.2. \widehat{K}_1 is displayed with a red star.

The parameters \widehat{M}_1 and \widehat{M}_2 are chosen in a similar way except that $\widetilde{\text{MSE}}_{K_1, K_2}(\lambda_1, \lambda_2)$ is replaced by $\widetilde{\text{MSE}}_{M_1, M_2}(\lambda_1, \lambda_2)$ where:

$$\widetilde{\text{MSE}}_{M_1, M_2}(\lambda_1, \lambda_2) = \|\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\gamma}}^{(M_1, M_2)}(\lambda_1, \lambda_2)\|_2^2,$$

with $\widehat{\boldsymbol{\gamma}}^{(M_1, M_2)}(\lambda_1, \lambda_2) = (\widehat{\alpha}_1, \widehat{\alpha}_2, \widehat{\boldsymbol{\beta}}_1^{(M_1)'}, \widehat{\boldsymbol{\beta}}_2^{(M_2)'})$ defined in (10) and (12). In the following, $\widehat{\boldsymbol{\gamma}}(\lambda_1, \lambda_2) = \widehat{\boldsymbol{\gamma}}^{(\widehat{M}_1, \widehat{M}_2)}(\lambda_1, \lambda_2)$.

2.5. Estimation of $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_2$. As the empirical correlation matrix is known to be a non accurate estimator of $\boldsymbol{\Sigma}$ when p is larger than n , a new estimator has to be used. Thus, for estimating $\boldsymbol{\Sigma}$ we adopted a cross-validation based method designed by Boileau *et al.* (2021) and implemented in the `cvCovEst` R package. This method chooses the estimator having the smallest estimation error among several compared methods (sample correlation matrix, POET (Fan *et al.* (2013)) and Tapering (Cai *et al.* (2010)) as examples). Since the samples in treatments t_1 and t_2 are assumed to be collected from the same population, $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_2$ are assumed to be equal.

2.6. Choice of the parameters λ_1 and λ_2 . For the sake of simplicity, we limit ourselves to the case where $\lambda_1 = \lambda_2 = \lambda$. For choosing λ we used BIC (Bayesian Information Criterion) which is widely used in the variable selection field and which consists in minimizing the following criterion with respect to λ :

$$\text{BIC}(\lambda) = n \log(\text{MSE}(\lambda)/n) + k(\lambda) \log(n),$$

where n is the total number of samples, $\text{MSE}(\lambda) = \|\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\gamma}}(\lambda)\|_2^2$ and $k(\lambda)$ is the number of non null coefficients in the OLS estimator $\widehat{\boldsymbol{\gamma}}$ obtained by re-estimating only the non null components of $\widehat{\boldsymbol{\beta}}_1$ and $\widehat{\boldsymbol{\beta}}_2 - \widehat{\boldsymbol{\beta}}_1$. The values of the BIC criterion as well as those of the MSE obtained from the dataset described in Section 2.2 are displayed in Figure 3.

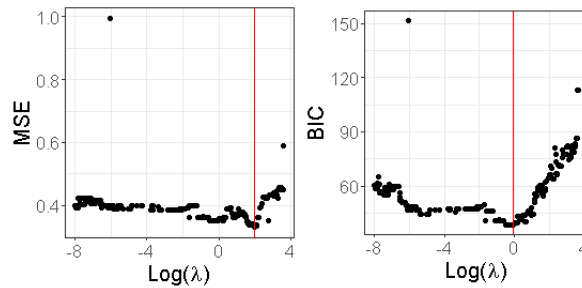


FIGURE 3. MSE and BIC for all λ . The λ minimizing each criterion is displayed with a vertical line.

Table 2 in the supplementary material provides the True Positive Rate (TPR) and False Positive Rate (FPR) when λ is chosen either by minimizing the MSE or the BIC criterion for this dataset. We can see from this table that both of them have TPR=1 (all true positives are identified). However, the FPR based on the BIC criterion is smaller than the one obtained by using the MSE.

3. NUMERICAL EXPERIMENTS

This section presents a comprehensive numerical study by comparing the performance of our method with the Lasso criterion in terms of prognostic and predictive biomarker selection. For the Lasso criterion, in order to directly estimate prognostic and predictive effects, \mathbf{X} and $\boldsymbol{\gamma}$ in Model (3) were replaced by

$$\mathbf{X}^* = \begin{bmatrix} \mathbf{1}_{n_1,1} & \mathbf{0}_{n_1,1} & \mathbf{X}_1 & \mathbf{0}_{n_1,p} \\ \mathbf{0}_{n_2,1} & \mathbf{1}_{n_2,1} & \mathbf{X}_2 & \mathbf{X}_2 \end{bmatrix},$$

and $\boldsymbol{\gamma}^* = (\alpha_1, \alpha_2, \boldsymbol{\beta}_1^*, \boldsymbol{\beta}_2^*)$, respectively, where \mathbf{X}_1 and \mathbf{X}_2 are defined in (1), $\mathbf{0}_{i,j}$ (resp. $\mathbf{1}_{i,j}$) denotes a matrix having i rows and j columns and containing only zeros (resp. ones). Note that this is the modeling proposed by Lipkovich et al. (2017). The sparsity enforcing constraint was put on the coefficients $\boldsymbol{\beta}_1^*$ and $\boldsymbol{\beta}_2^*$ which boils down to putting a sparsity enforcing constraint on $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1$.

3.1. Simulation setting. All simulated datasets were generated from Model (3) where the n_1 (n_2) rows of \mathbf{X}_1 (\mathbf{X}_2) are assumed to be independent Gaussian random vectors with a covariance matrix $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma}_{bm}$, and $\boldsymbol{\epsilon}$ is a standard Gaussian random vector independent of \mathbf{X}_1 and \mathbf{X}_2 . We defined $\boldsymbol{\Sigma}_{bm}$ as:

$$(14) \quad \boldsymbol{\Sigma}_{bm} = \begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{12}^T & \boldsymbol{\Sigma}_{22} \end{bmatrix}$$

where $\boldsymbol{\Sigma}_{11}$ (resp. $\boldsymbol{\Sigma}_{22}$) are the correlation matrix of prognostic (resp. non-prognostic) biomarkers with off-diagonal entries equal to a_1 (resp. a_3). Moreover, $\boldsymbol{\Sigma}_{12}$ is the correlation matrix between prognostic and non-prognostic variables with entries equal to a_2 . In our simulations $(a_1, a_2, a_3) = (0.3, 0.5, 0.7)$, which is a framework proposed by Xue and Qu

(2017). We checked that the Irrepresentable Condition (IC) of Zhao and Yu (2006) is violated and thus the standard Lasso cannot recover the positions of the null and non null variables. For each dataset we assumed randomized treatment allocation between standard and experimental arm with a 1:1 ratio, *i.e.* $n_1 = n_2 = 50$. We further assume a relative treatment effect of 1 ($\alpha_1 = 0$ and $\alpha_2 = 1$). The number of biomarkers p varies from 200 to 2000. The number of active biomarkers was set to 10 (*i.e.* 5 purely prognostic biomarkers with $\beta_{1j} = \beta_{2j} = b_1 = 1$ ($j = 1, \dots, 5$) and 5 biomarkers both prognostic and predictive with $\beta_{1j} = b_1$ and $\beta_{2j} = b_2 = 2$ ($j = 6, \dots, 10$)).

3.2. Evaluation criteria. We considered several evaluation criteria to assess the performance of the methods in selecting the prognostic and predictive biomarkers: the TPR_{prog} as the true positive rate (*i.e.* rate of active biomarkers selected) and FPR_{prog} the false positive rate (*i.e.* rate of inactive biomarkers selected) of the selection of prognostic biomarkers, and similarly for predictive biomarkers with TPR_{pred} and FPR_{pred} . We further note TPR_{all} and FPR_{all} the criterion of overall selection among all candidate biomarkers regardless their prognostic or predictive effect. The objective of the selection is to maximize the TPR_{all} and minimize the FPR_{all} . All metrics were calculated by averaging the results of 100 replications for each scenario.

3.3. Biomarker selection results. For the proposed method, different results were presented. $\text{PPLasso}_{\text{oracle}}$ (*resp.* $\text{PPLasso}_{\text{est}}$) corresponds to the results of the method by considering the true (*resp.* estimated) matrix Σ_{bm} . For estimating Σ_{bm} , we used the approach explained in Section 2.5. Two choices of λ are also presented: “optimal” and “min(bic)”. The former gives the optimal selection that maximizes $(\text{TPR}_{\text{all}} - \text{FPR}_{\text{all}})$ which is also the choice used for Lasso in these simulations. The latter minimizes the BIC criterion defined in Section 2.6.

Figure 4 shows the selection performance of PPLasso and Lasso in the simulation scenario presented in Section 3.1. PPLasso achieved to select all prognostic factors (TPR_{prog} almost 1) even for large p , with limited false positive prognostic biomarkers selected. As compared to the optimal λ maximizing $(\text{TPR}_{\text{all}} - \text{FPR}_{\text{all}})$, the one selected with the BIC tends to select some false positives (average: 32.87 ($\text{FPR}_{\text{prog}} = 0.173$) for $p = 200$ and 10.39 ($\text{FPR}_{\text{prog}} = 0.005$) for $p = 2000$). The results obtained from the oracle and estimated Σ_{bm} are comparable. Selection performance of predictive biomarkers is slightly lowered as compared to prognostic biomarkers. Even if the false positive selection is quite similar between prognostic and predictive biomarkers, PPLasso missed some true predictive biomarkers when λ is selected with the BIC criterion (average $\text{TPR}_{\text{pred}} = 0.982$ and 0.798 for oracle and estimated Σ_{bm} , respectively, with $p = 2000$). In this scenario where the IC is violated, PPLasso globally outperforms Lasso. The Lasso approach failed in selecting all truly prognostic and predictive biomarkers, and the number of missed active biomarkers increased with the dimension p ($\text{TPR}_{\text{prog}} = 0.825$ and 0.448, $\text{TPR}_{\text{pred}} = 0.810$ and 0.454 for $p = 200$ and 2000, respectively).

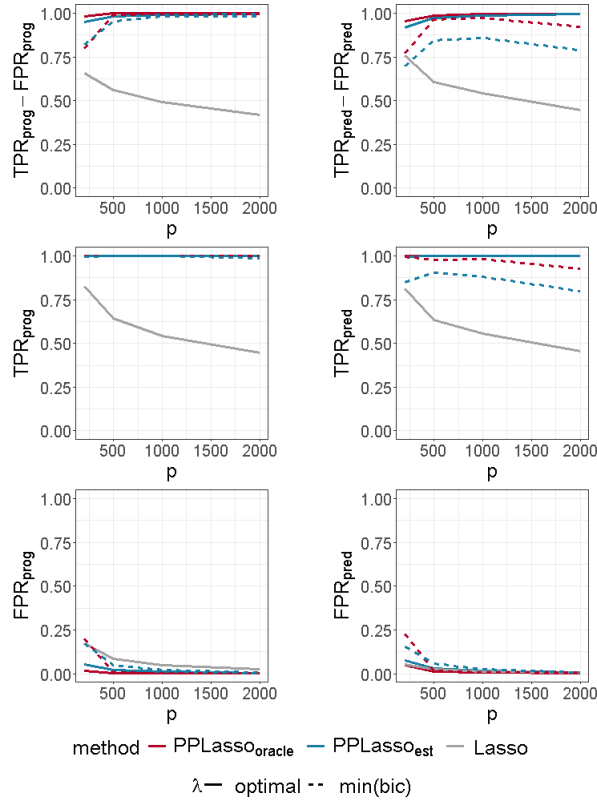


FIGURE 4. (TPR-FPR) and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers.

3.3.1. *Impact of the correlation matrix Σ .* To evaluate the impact of the correlation matrix on the selection performance of the methods, additional scenarios are presented where the IC is satisfied:

- (1) Compound symmetry structure where all biomarkers are equally correlated with a correlation $\rho = 0.5$;
- (2) Independent setting where Σ_{bm} is the identity matrix.

For the scenario with compound symmetry structure displayed in Figure 5, Lasso successfully identified the true prognostic biomarkers (TPR_{prog} close to 1 even for large p) with limited false positive selection. On the other hand, Lasso missed some predictive biomarkers especially when p increases ($\text{TPR}_{\text{pred}} = 0.934$ and 0.502 for $p = 200$ and 2000 respectively). On the contrary, PPLasso successfully identified almost all predictive biomarkers with the optimal choice of λ . Moreover, even when λ is selected by minimizing the BIC criterion ($\text{min}(\text{bic})$), PPLasso_{est} outperformed the Lasso when $p > 500$ with relatively stable TPR_{pred} and FPR_{pred} as p increases.

For the independent setting, as displayed in Figure 6, prognostic biomarkers were globally well identified by all the compared methods with a slightly higher TPR_{prog} for Lasso as

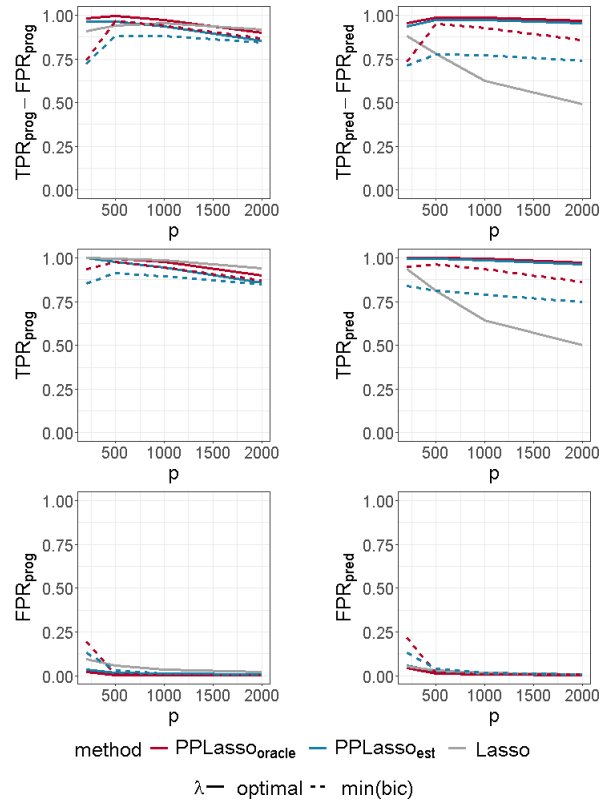


FIGURE 5. (TPR-FPR) and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers for the compound symmetry correlation structure.

compared to PPLasso but also with a slightly higher FPR_{prog} . With regards to predictive biomarkers, PPLasso using Σ_{bm} (oracle) performed also similarly to the Lasso, which is reasonable since no transformation has been used in PPLasso. On the other hand, even if PPLasso with λ selected with “min(bic)” performed similarly with PPLasso with optimal λ for relatively small p , the selection performance is altered for large p and is close to the performance of Lasso with optimal λ selected.

3.3.2. *Impact of the effect size of active biomarkers.* To evaluate the impact of the effect size on biomarker selection performance, the scenario presented in Section 3.1 was considered with different values of b_2 : 1.5, 2 and 2.5.

Since the effect size of prognostic biomarkers did not change the comparison focused on predictive biomarkers. As expected, the reduction of the effect size makes the biomarker selection harder, especially for the Lasso where the predictive biomarker selection is limited when $b_2 = 1.5$: for $p = 2000$, $TPR_{\text{pred}} = 0.454$ (resp. 0.218) for $b_2 = 2$ (resp. 1.5), see Figure 4 and Figure 9 of the supplementary material. The selection performance of PPLasso when λ is selected with min(bic) is also reduced by decreasing b_2 , especially when

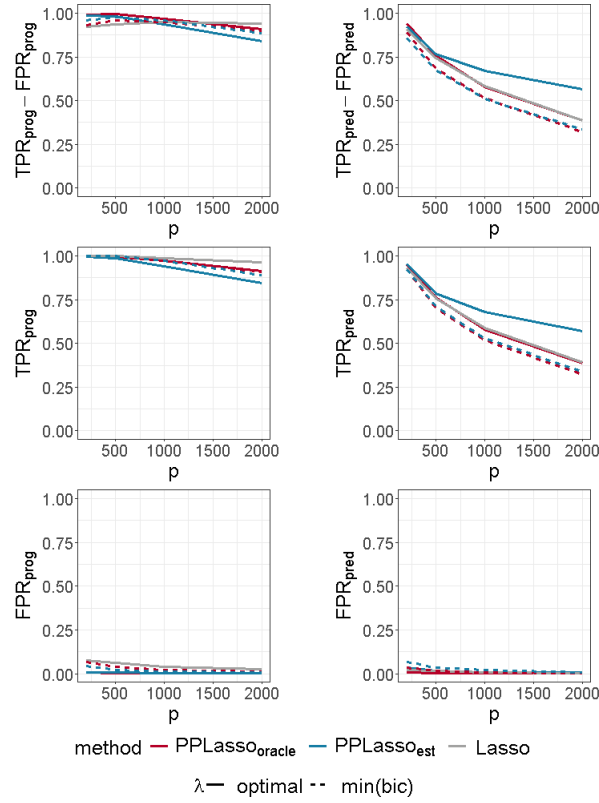


FIGURE 6. (TPR-FPR) and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers (independent setting).

Σ_{bm} is also estimated. Nevertheless, the selection performance remains better than for the Lasso for which the performance displayed are associated to the optimal value of λ . On the other hand, even with limited effect size, PPLasso with optimal λ identified all predictive biomarkers with very limited false positive selection. When b_2 was increased to 2.5, the selection performance for all methods is improved and the results for PPLasso with estimated λ was close to the ones with the optimal λ as displayed in Figure 10 of the supplementary material. As compared with PPLasso, for which the selection performance remains stable as p increases, Lasso is more impacted by the value of p since the true positive selection decreased as p increased ($\text{TPR}_{\text{pred}} = 0.950$ (resp. 0.646) for $p = 200$ (resp. 2000)).

3.3.3. Impact of the number of predictive biomarkers. The impact of the number of true predictive biomarkers was assessed by increasing the number of predictive biomarkers from 5 to 10 in the scenario presented in Section 3.1. When the number of predictive biomarkers increased, the impact on PPLasso is almost negligible. However, for Lasso, we can see from Figure 11 of the supplementary material that it became even harder to identify

predictive biomarkers. TPR_{pred} decreased compared to Figure 4, especially for large p (e.g. $\text{TPR}_{\text{pred}} = 0.124$ for $p = 2000$).

3.3.4. *Impact of the dimension of the dataset.* In this section, we studied a different sample size: $n=50$ with $n_1 = n_2 = 25$ and a different number of biomarkers: $p=5000$.

We can see from Figure 12 of the supplementary material that for $p = 5000$, the selection performance of PPLasso is not altered as compared with $p = 2000$ while the Lasso has more difficulties to identify both prognostic and predictive biomarkers.

When the sample size is smaller ($n=50$), we can see from Figure 13 of the supplementary material that the ability to identify prognostic and predictive biomarkers decreased for all the methods. However, PPLasso still outperformed the Lasso with higher TPR_{prog} and TPR_{pred} and lower FPR_{prog} and FPR_{pred} .

4. APPLICATION TO GENE EXPRESSION DATA IN BREAST CANCER

We applied the previously described methods to publicly available transcriptomic data from the RV144 vaccine trial (Rerks-Ngarm *et al.* (2009)). This trial showed reduced risk of HIV-1 acquisition by 31.2% with vaccination with ALVAC and AIDSVAX as compared to placebo. Transcriptomic profiles of in vitro HIV-1 Env-stimulated peripheral blood mononuclear cells (PBMCs) obtained pre-immunization and 15 days after the immunization (D15) from both 40 vaccinees and 10 placebo recipients were generated to better understand underlying biological mechanisms (Fourati *et al.* (2019), Gene Expression Omnibus accession code: GSE103671).

For illustration purpose, the absolute change at D15 in gene mTOR was considered as the continuous endpoint (response). mTOR plays a key role in mTORC1 signaling pathway which has been shown to be associated with risk of HIV-1 acquisition (Fourati *et al.* (2019), Akbay *et al.* (2020)). The gene expression has been normalized as in the original publication of Fourati *et al.* (2019). After removing non-annotated genes (LOCxxxx and HS.xxxx), the top 2000 genes with the highest empirical variances were included as candidate biomarkers for prognostic and predictive identification from PPLasso and Lasso. The penalty parameter λ for the Lasso was selected through the classical cross-validation approach. For PPLasso, λ was selected based on the criterion described in Section 2.6.

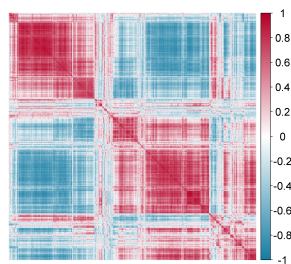


FIGURE 7. Heatmaps of the correlation matrix estimated by the `cvCovEst` R package.

The estimation of Σ was obtained by comparing several candidate estimators from the `cvCovEst` R package and by selecting the estimator having the smallest estimation error. In this application, the combination of the sample covariance matrix and a dense target matrix (*denseLinearShrinkEst*) derived by Ledoit and Wolf (2020) provides the smallest estimation error. Figure 7 displays the estimated Σ and highlights the strong correlation between the genes. Table 3 of the Supplementary material gives details on the compared estimators.

Prognostic and predictive genes selected by PPLasso and Lasso are listed in Table 1. PPLasso selected a slightly higher number of genes as compared to the Lasso (9 vs. 10 and 15 vs. 7 for prognostic and predictive genes, respectively). The intersection between the two methods is moderate (3 prognostic genes, 3 predictive genes, and 1 gene selected as prognostic for Lasso and predictive for PPLasso). Interestingly, some genes selected by both methods such as SLAMF7, TNFRSF6B, TNFRSF18 or NUCKS1 have already been discussed in the HIV-1 field. Moreover, among the predictive genes selected by the PPLasso, some are linked to pathways that have been highlighted as possible target for HIV-1 such as BIRC3 and TLR8.

5. CONCLUSION

We propose a new method named PPLasso to simultaneously identify prognostic and predictive biomarkers. PPLasso is particularly interesting for dealing with high dimensional omics data when the biomarkers are highly correlated, which is a framework that has not been thoroughly investigated yet. From various numerical studies with or without strong correlation between biomarkers, we highlighted the strength of PPLasso in well identifying both prognostic and predictive biomarkers with limited false positive selection. The current method is only dedicated to the analysis of continuous responses through ANCOVA type models. However, it will be the subject of a future work to extend it to other challenging contexts, such as classification or survival analysis.

FUNDING

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

TABLE 1. Selected genes from Lasso and PPLasso.

	prognostic genes	predictive genes
PPLasso	HAPLN3, SLAMF7, GTF3C5, FAM46A, SH3PXD2B, TM4SF1, TNFRSF6B, TNFRSF18, TRPM2	TLR8, YTHDC1, NUCKS1, BIRC3, SLAMF7, NFATC2IP, BOK, MGRN1, KIAA0492, SLC25A36, HMG2, P2RY5, RPL21, MS4A7, RPL12P6
Lasso	DKFZp434K191, NUCKS1, MAFF, SLAMF7, HIST2H2AC, HIST1H4C, IL8, TNFRSF6B, TNFRSF18, SCAND1	DKFZp434K191, YTHDC1, VMO1, BOLA2, HIST1H4C, RPL21, MS4A7

REFERENCES

- Akbay, B., Shmakova, A., Vassetzky, Y., and Dokudovskaya, S. (2020). Modulation of mtorc1 signaling pathway by hiv-1. *Cells*, **9**, 1090.
- Ballman, K. V. (2015). Biomarker: Predictive or prognostic? *Journal of clinical oncology*, **33**(33), 3968–3971.
- Boileau, P., Hejazi, N. S., van der Laan, M. J., and Dudoit, S. (2021). Cross-validated loss-based covariance matrix estimator selection in high dimensions.
- Cai, T., Zhang, C.-H., and Zhou, H. (2010). Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*, **38**.
- Clark, G. (2008). Prognostic factors versus predictive factors: Examples from a clinical trial of erlotinib. *Molecular oncology*, **1**, 406–12.
- Fan, J. and Li, R. (2006). Statistical challenges with high dimensionality: Feature selection in knowledge discovery. *Proc. Madrid Int. Congress of Mathematicians*, **3**.
- Fan, J. and Lv, J. (2009). A selective overview of variable selection in high dimensional feature space. *Stat. Sinica*, **20** **1**, 101–148.
- Fan, J., Liao, Y., and Mincheva, M. (2013). Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society. Series B, Statistical methodology*, **75**.
- Faraway, J. J. (2002). *Practical regression and ANOVA using R*. University of Bath.
- Foster, J., Taylor, J., and Ruberg, S. (2011). Subgroup identification from randomized clinical trial data. *Statistics in medicine*, **30**, 2867–80.
- Fourati, S., Ribeiro, S., Blasco Lopes, F., Talla, A., Lefebvre, F., Cameron, M., Kaewkungwal, J., Pitisuttithum, P., Nitayaphan, S., Rerks-Ngarm, S., Kim, J., Thomas, R., Gilbert, P., Tomaras, G., Koup, R., Michael, N., McElrath, M., Gottardo, R., and Sékaly, R. (2019). Integrated systems approach defines the antiviral pathways conferring protection by the rv144 hiv vaccine. *Nature Communications*, **10**.
- Ledoit, O. and Wolf, M. (2020). The Power of (Non-)Linear Shrinking: A Review and Guide to Covariance Matrix Estimation. *Journal of Financial Econometrics*.
- Lipkovich, I. and Dmitrienko, A. (2014). Strategies for identifying predictive biomarkers and subgroups with enhanced treatment effect in clinical trials using sides. *Journal of biopharmaceutical statistics*, **24**, 130–53.
- Lipkovich, I., Dmitrienko, A., Denne, J., and Enas, G. (2011). Subgroup identification based on differential effect search (sides) – a recursive partitioning method for establishing response to treatment in patient subpopulations. *Statistics in medicine*, **30**, 2601–21.
- Lipkovich, I., Dmitrienko, A., and B. D’Agostino Sr., R. (2017). Tutorial in biostatistics: data-driven subgroup identification and analysis in clinical trials. *Statistics in Medicine*, **36**(1), 136–196.
- McDonald, J. (2009). *Handbook of Biological Statistics 2nd Edition*. Sparky House Publishing Baltimore.
- Rerks-Ngarm, S., Pitisuttithum, P., Nitayaphan, S., Kaewkungwal, J., Chiu, J., Paris, R., Prensri, N., Namwat, C., De Souza, M., Benenson, M., Gurunathan, S., Tartaglia, J., McNeil, J., Francis, D., Stablein, D., Birx, D., Chunsuttiwat, S., Khamboonruang, C.,

- and Kim, J. (2009). Vaccination with alvac and aidsvox to prevent hiv-1 infection in thailand. The New England journal of medicine, **361**, 2209–20.
- Saeyns, Y., Inza, I., and Larranaga, P. (2007). A review of feature selection techniques in bioinformatics. Bioinformatics, **23**(19), 2507–2517.
- Sechidis, K., Papangelou, K., Metcalfe, P. D., Svensson, D., Weatherall, J., and Brown, G. (2018). Distinguishing prognostic and predictive biomarkers: an information theoretic approach. Bioinformatics, **34**(19), 3365–3376.
- Smith, G. (2018). Step away from stepwise. J. Big Data, **5**(32), 1–12.
- Tian, L., Alizadeh, A., Gentles, A., and Tibshirani, R. (2012). A simple method for estimating interactions between a treatment and a large number of covariates. Journal of the American Statistical Association, **109**.
- 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 Taylor, J. (2011). The solution path of the generalized lasso. Ann. Stat., **39**(3), 1335–1371.
- Wang, H., Lengerich, B., Aragam, B., and Xing, E. (2019). Precision lasso: Accounting for correlations and linear dependencies in high-dimensional genomic data. Bioinformatics, **35**(7), 1181–1187.
- Wang, X. and Leng, C. (2016). High dimensional ordinary least squares projection for screening variables. J. R. Stat. Soc. Ser. B (Stat. Methodol.), **78**(3), 589–611.
- Xue, F. and Qu, A. (2017). Variable selection for highly correlated predictors.
- Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. J. Machine Learn. Res., **7**, 2541–2563.
- Zhu, W., Lévy-Leduc, C., and Ternès, N. (2021). A variable selection approach for highly correlated predictors in high-dimensional genomic data. Bioinformatics, **37**(16), 2238–2244.

SUPPLEMENTARY MATERIAL

This supplementary material provides additional numerical experiments, figures and tables for the paper: “Identification of prognostic and predictive biomarkers in high-dimensional data with PPLasso”.

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

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

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

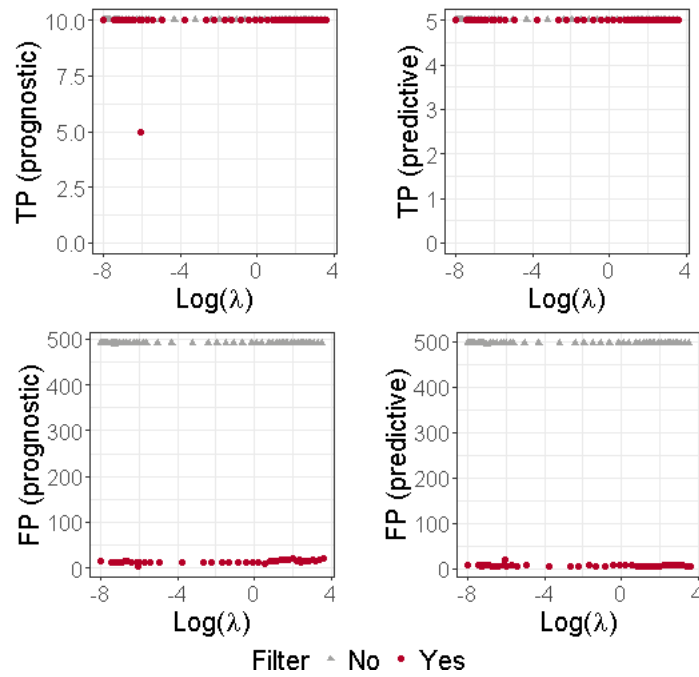


FIGURE 8. Number of True Positives and True Negatives for $\hat{\beta}$ and $\hat{\beta}_0$ on prognostic/predictive biomarkers.

	MSE	BIC
TPR(prognostic)	1.000	1.000
FPR(prognostic)	0.038	0.024
TPR(predictive)	1.000	1.000
FPR(predictive)	0.008	0.006

TABLE 2. TPR and FPR associated to prognostic and predictive biomarker identification with the λ chosen in Figure 3.

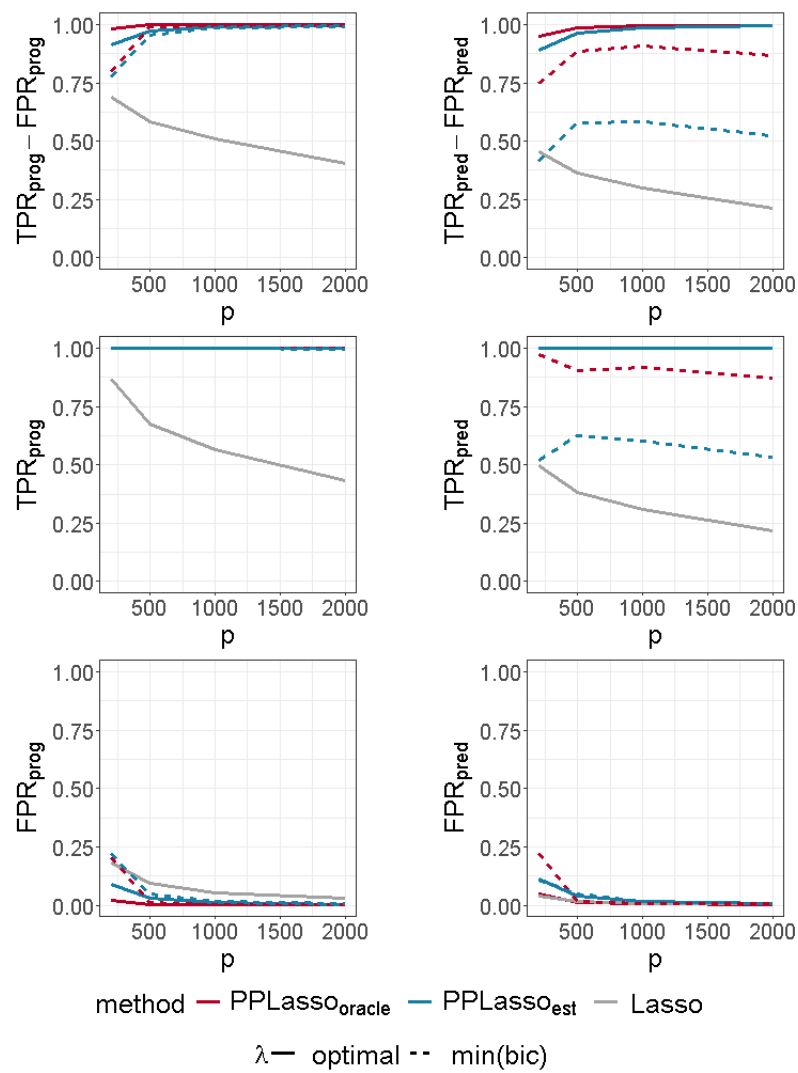


FIGURE 9. (TPR-FPR) and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers ($b_2 = 1.5$).

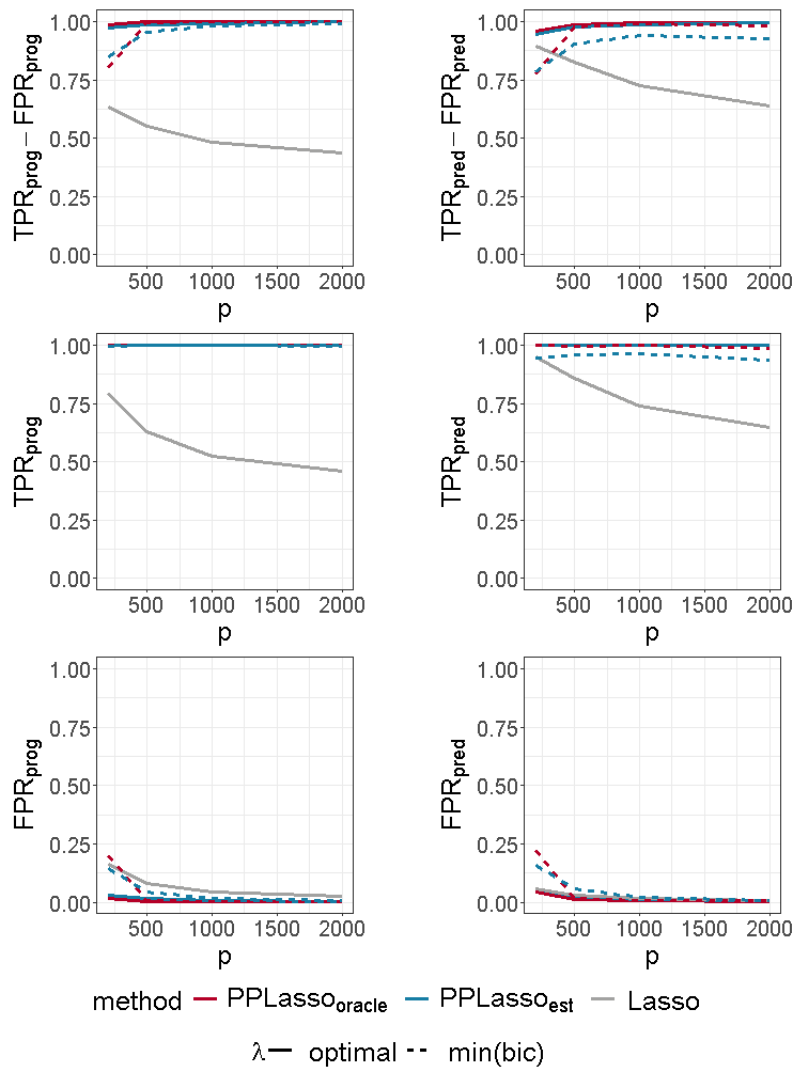


FIGURE 10. $(TPR-FPR)$ and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers ($b_2 = 2.5$).

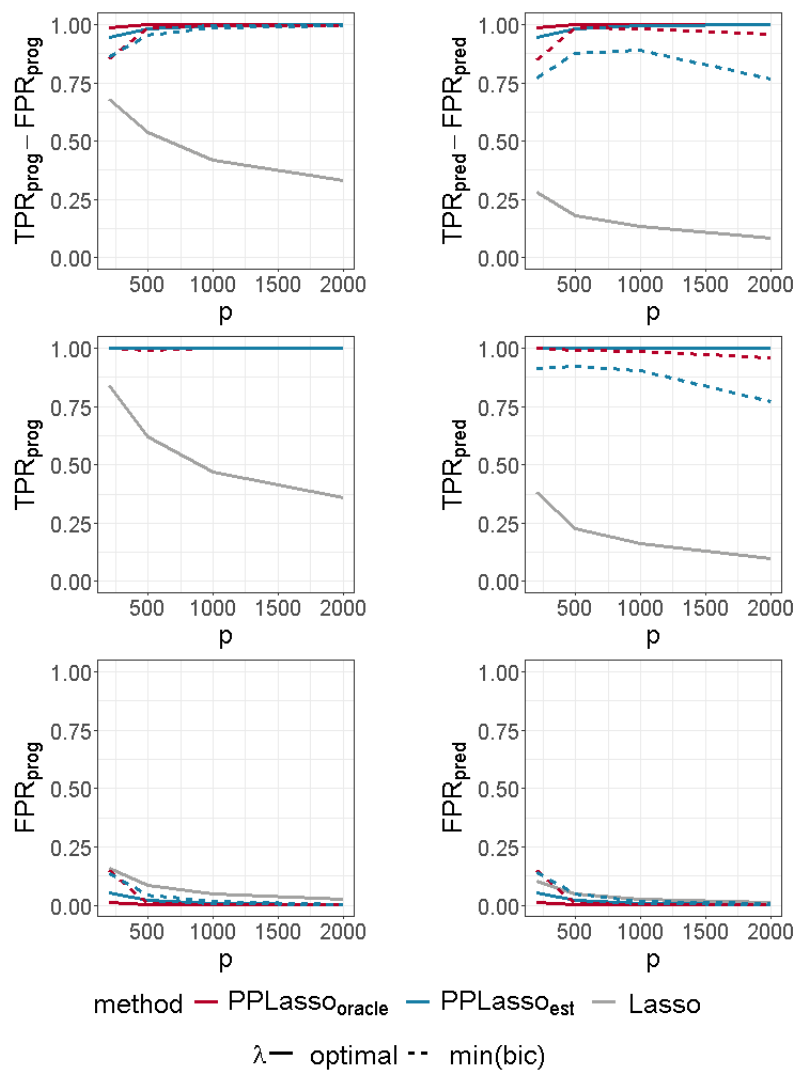


FIGURE 11. (TPR-FPR) and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers (10 predictive biomarkers).

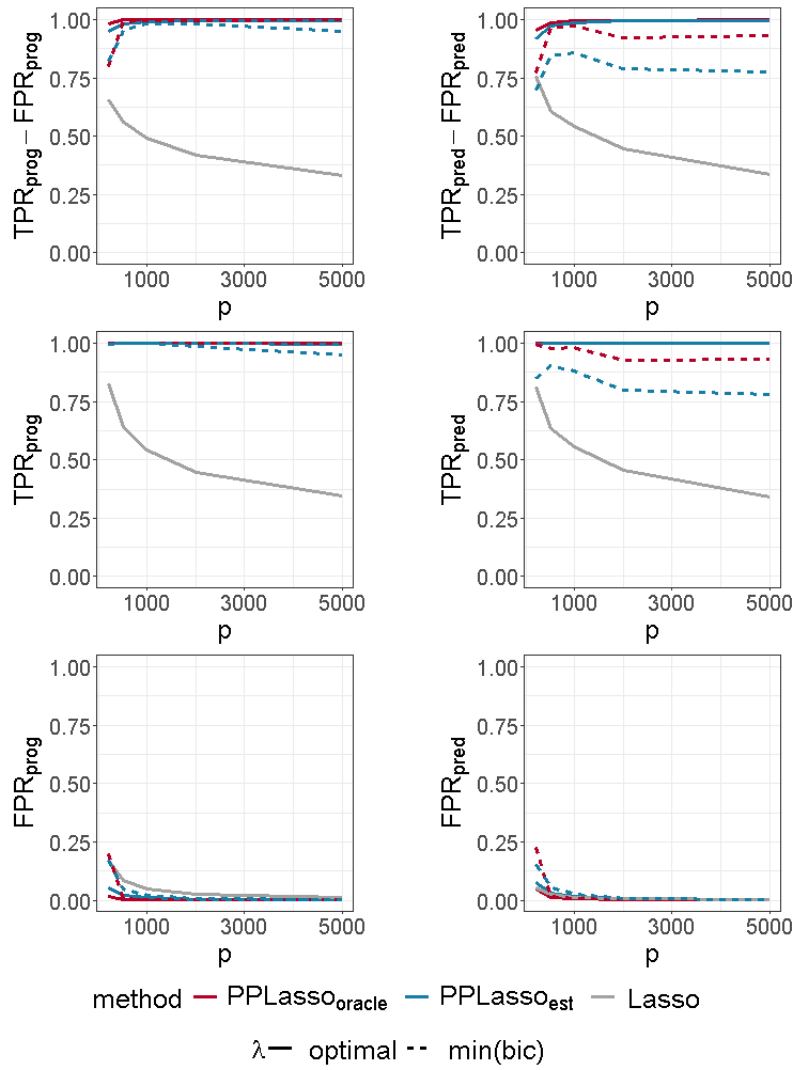


FIGURE 12. $(TPR-FPR)$ and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers (with $p = 5000$).

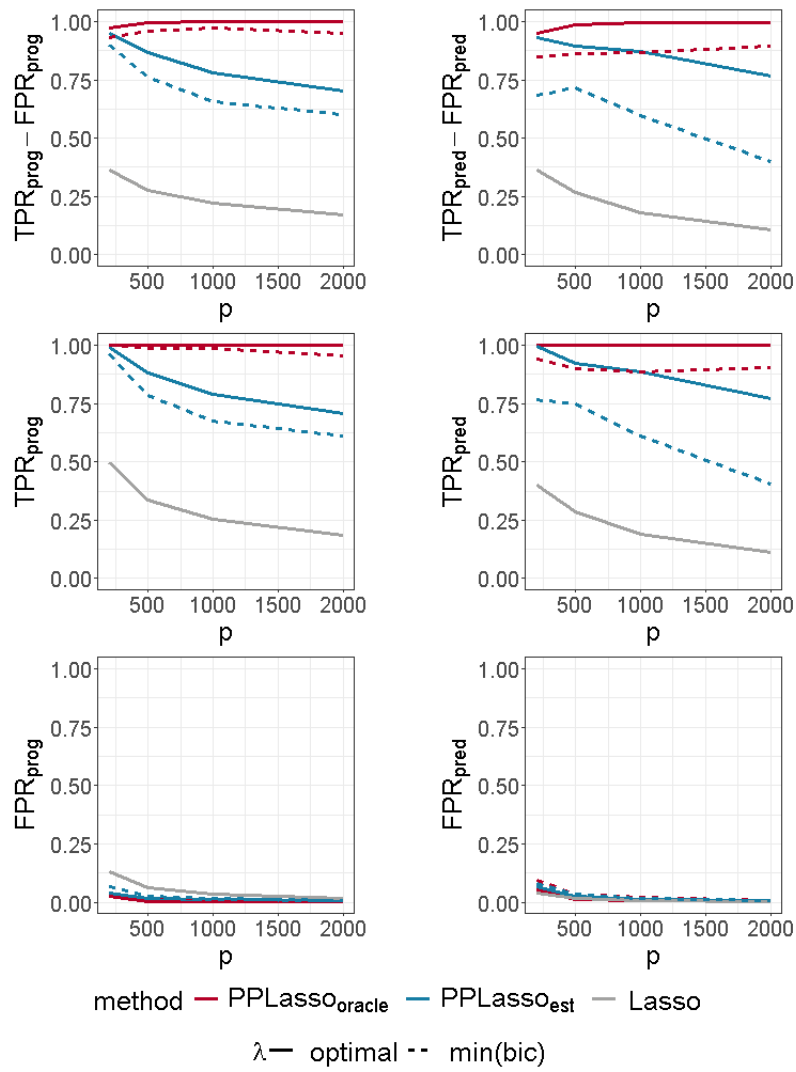


FIGURE 13. $(\text{TPR}-\text{FPR})$ and the corresponding True Positive Rate (TPR) and False Positive Rate (FPR) for prognostic (left) and predictive (right) biomarkers ($n_1 = n_2 = 25$).

Estimator	Hyperparameters	Empirical risk
<i>denseLinearShrinkEst</i>	-	<i>102546</i>
sampleCovEst	-	102547
linearShrinkLWEst	-	103496
poetEst	lambda=0.1, k=2	104522
poetEst	lambda=0.2, k=2	105358
poetEst	lambda=0.1, k=1	105972
poetEst	lambda=0.2, k=1	108222
thresholdingEst	gamma=0.2	137798
thresholdingEst	gamma=0.4	186844

TABLE 3. Empirical risk of tested methods with different hyperparameters.