



HAL
open science

On the use of cross-validation for the calibration of the adaptive lasso

Nadim Ballout, Lola Etievant, Vivian Viallon

► To cite this version:

Nadim Ballout, Lola Etievant, Vivian Viallon. On the use of cross-validation for the calibration of the adaptive lasso. *Biometrical Journal*, 2023, 65 (5), pp.e2200047. <10.1002/bimj.202200047>. <hal-04138647>

HAL Id: hal-04138647

<https://hal.science/hal-04138647v1>

Submitted on 23 Jun 2023

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

On the use of cross-validation for the calibration of the adaptive lasso

Nadim Ballout¹, Lola Étievant², Vivian Viallon^{3*}

¹ Univ Lyon, Univ Eiffel, IFSTTAR, Univ Lyon 1, UMRESTTE, UMR_T9405, F-69500 Bron, France.

² Université Claude Bernard Lyon 1, CNRS UMR 5208, Institut Camille Jordan, Lyon, France.

³ Nutrition and Metabolism Branch, International Agency for Research on Cancer (IARC-WHO), Lyon, France;

Abstract

The adaptive lasso refers to a class of methods that use weighted versions of the L_1 -norm penalty, with weights derived from an initial estimate of the parameter vector to be estimated. Irrespective of the method chosen to compute this initial estimate, the performance of the adaptive lasso critically depends on the value of a hyperparameter, which controls the magnitude of the weighted L_1 -norm in the penalized criterion. As for other machine learning methods, cross-validation is very popular for the calibration of the adaptive lasso, that this for the selection of a data-driven optimal value of this hyperparameter. However, the most simple cross-validation scheme is not valid in this context, and a more elaborate one has to be employed to guarantee an optimal calibration. The discrepancy of the simple cross-validation scheme has been well documented in other contexts, but less so when it comes to the calibration of the adaptive lasso, and, therefore, many statistical analysts still overlook it. In this work, we recall appropriate cross-validation schemes for the calibration of the adaptive lasso, and illustrate the discrepancy of the simple scheme, using both synthetic and real-world examples. Our results clearly establish the suboptimality of the simple scheme, in terms of support recovery and prediction error, for several versions of the adaptive lasso, including the popular one-step lasso.

Keywords. adaptive lasso, cross-validation, calibration, hyperparameter, tuning parameter, one-step lasso.

*Corresponding author: viallonv@iarc.fr

1 Introduction

High dimensional data are characterized by a number p of variables larger, or at least not significantly lower, than sample size n . They have become ubiquitous in many fields, including biology, medicine, sociology, and economy (Giraud, 2014). Their analysis raises a number of statistical challenges (Fan and Li, 2006, Hastie et al., 2009), usually summarized under the term “*curse of dimensionality*”. Consequently, it has attracted a lot of attention in the statistical literature over the past decades (Bühlmann and van De Geer, 2011, Donoho et al., 2000, Fan and Li, 2006, Hastie et al., 2009, 2015). In particular, a variety of approaches based on the optimization of penalized versions of the log-likelihood have been developed, to estimate the true parameter vector $\beta^* = (\beta_1^*, \dots, \beta_p^*)^T \in \mathbb{R}^p$ under high-dimensional parametric regression models (Huang et al., 2008a, Tibshirani, 1996). These approaches use a penalty term whose strength is controlled by a hyperparameter – or tuning parameter – and that is added to the loss-function so that the estimation can take advantage of some expected property of the true parameter vector β^* . For example, when β^* is expected to be sparse, popular approaches rely on the use of L_q penalties, for some $q \leq 1$. Among such approaches, the arguably most popular one is the lasso, which uses an L_1 -norm penalty. Extensions such as the group lasso (Jacob et al., 2009), fused lasso (Tibshirani et al., 2005), generalized fused lasso (Viallon et al., 2016), data shared lasso (Ballout et al., 2020, Ballout and Viallon, 2019, Gross and Tibshirani, 2016, Ollier and Viallon, 2017), etc., rely on structured sparsity inducing norms, and can be used when some particular structured sparsity is expected in β^* .

We will here focus on another extension of the lasso, namely the adaptive lasso (Bühlmann and van De Geer, 2011, Zou, 2006). It refers to a class of methods where the L_1 -norm $\|\beta\|_1 = \sum_{j=1}^p |\beta_j|$ used in the standard lasso is replaced by a weighted version $\sum_{j=1}^p w_j |\beta_j|$. Weights $w_j \geq 0$ are typically data-driven, of the form $w_j = 1/(|\tilde{\beta}_j| + \varepsilon)$, with $\varepsilon \geq 0$ and $(\tilde{\beta}_j)_{1 \leq j \leq p}$ some initial estimates of the parameters $(\beta_j^*)_{1 \leq j \leq p}$. Three popular versions of the adaptive lasso will be considered in this work: (i) the original adaptive lasso introduced by Zou (2006), where weights are derived from Ordinary Least Squares (OLS) estimates; (ii) the version proposed by Bühlmann and Meier (2008), where weights are computed from lasso estimates; and (iii) the version proposed by Zhang et al. (2008), where weights are computed from ridge estimates. In the rest of this article, the original adaptive lasso introduced by Zou (2006) will be referred to as the *ols-adaptive lasso*, the one proposed by Zhang et al. (2008) as the *ridge-adaptive lasso*, while we will refer to the method described in Bühlmann and Meier (2008) as the *one-step lasso*, following their terminology.

The adaptive lasso is very popular in practice. First, it can be implemented very easily and efficiently using algorithms originally developed for

the lasso, such as the `glmnet` R package (Friedman et al., 2010). Second, it has been shown to usually outperform the lasso. For example, in the fixed p case, Zou (2006) established that the lasso estimates do not enjoy the asymptotic oracle property (in the sense of Fan and Li (2006)), while the ols-adaptive lasso estimates do under mild conditions on the value of the hyperparameter that controls the strength of the weighted L_1 -norm penalty. In addition, conditions ensuring support recovery in the non-asymptotic framework (which especially allows the study of the $p \gg n$ case) are weaker for the one-step lasso than for the lasso; see, e.g., Corollaries 7.8-7.9 and Section 2.8.3, in Bühlmann and van De Geer (2011).

As for other penalized approaches, the theoretical and empirical performance of the adaptive lasso critically depends on the value of the hyperparameter. Its optimal value involves unknown quantities, such as the variance of the noise under linear regression models, but also unknown constants related to the compatibility and irrepresentability conditions of the design matrix (Bühlmann and van De Geer, 2011). Consequently, the practical selection of the tuning parameter, or calibration, also has attracted a lot of attention in the statistical literature (Arlot, 2019, Chen and Chen, 2008, Chichignoud et al., 2016, Giacobino et al., 2017). Cross-validation is among the most commonly used strategies for the calibration of the adaptive lasso, and more generally of many statistical and machine learning methods; see, e.g., Chapter 7 in Hastie et al. (2009). See also Arlot (2008), Arlot and Celisse (2010). Several versions of cross-validation have been described in the literature, but the K -fold version is arguably the most popular one in practice (Hastie et al., 2009).

Briefly, K -fold cross-validation first consists in partitioning the original sample $D = (y_i, \mathbf{x}_i)_{1 \leq i \leq n}$ into $K \geq 2$ balanced folds $D^{(1)}, \dots, D^{(K)}$, with $D = \cup_{k=1}^K D^{(k)}$. Then, it consists of K steps, where at each step k , (i) the fold $D^{(k)}$ is used as an “independent” test sample, while the remaining $K - 1$ folds $D \setminus D^{(k)}$ are combined and jointly used as the training sample, (ii) the estimator $\hat{\beta}_k$ is computed on the training sample $D \setminus D^{(k)}$, and (iii) its prediction error, say $\text{Pred.Err}(D^{(k)}, \hat{\beta}_k)$, is evaluated on the test sample $D^{(k)}$. The cross-validated prediction error is finally defined as the average of these K prediction errors: $(1/K) \times \sum_k \text{Pred.Err}(D^{(k)}, \hat{\beta}_k)$. We refer to Bates et al. (2021) for additional details on the estimation of prediction error using cross-validation. This cross-validated prediction error can be used to assess the predictive performance of one estimator, or to compare the predictive performance among a set of estimators, or to select the tuning parameter by considering a set of estimators constructed with different values of the tuning parameter (Arlot and Celisse, 2010). However, for the cross-validation to be valid, it is crucial that the full learning algorithm is run on the training sample only, or, in other words, that data from the test sample are not used during the learning stage at all. For example, it is well known that if some variable selection is involved in the learning algorithm under study, this step

has to be performed on the training sample only, and not on the full original sample (Hastie et al., 2009).

The adaptive lasso algorithm is made of two steps, where the first one consists in computing the initial estimates $(\tilde{\beta}_j)_{1 \leq j \leq p}$ (through OLS estimation, the lasso, or ridge for example), and the second one consists in computing the adaptive lasso estimate *per se* with weights derived from these initial estimates. From the principle of the K -fold cross-validation recalled above, both steps have to be performed on the training sample to guarantee the validity of the calibration of the adaptive lasso (Krstajic et al., 2014, Stone, 1974). Although this proper cross-validation scheme was described in the literature (He et al., 2019, Krämer et al., 2009), the discrepancy of a more naive cross-validation scheme, where initial estimates are first derived from the full sample, and then the cross-validation is applied to select the hyperparameter of the adaptive lasso (focusing on the second step only), has not been well described and documented in the literature. In particular, Krämer et al. (2009) do not mention this simple scheme at all, while He et al. (2019) simply note that it leads to possibly large false discovery rates, without stressing that this scheme is fundamentally inappropriate in this case. Likewise, Bühlmann and Meier (2008) recommend the use of cross-validation for the calibration of the tuning parameters in their one-step lasso for both the initial and final estimators (which is fine if a proper cross-validation scheme is used), but they do not stress the inappropriateness of the simple scheme for calibrating the final estimator, nor do they describe an appropriate scheme. Finally, except for a few exceptions such as the `adalasso` function of the `parcor` (Krämer et al., 2009) and the `cvlasso` function of the `lassopack` package in STATA, proper implementations of the cross-validation for the adaptive lasso are lacking in publicly available adaptive lasso solvers. In particular, when using the popular `glmnet` R package (Friedman et al., 2010), it is tempting to implement the simple, and in this case improper, scheme, by simply using the `cv.glmnet` function to select the hyperparameter of the adaptive lasso after computing the initial estimates (step 1) on the full sample. This improper cross-validation scheme, was used in the illustrative example of Section 2.8.1 of Bühlmann and van De Geer (2011), and can also be found in the original version of the `adapt4pv` R package (Courtois et al., 2021). Consequently, many statistical analysts still seem to apply this improper cross-validation scheme for calibrating the adaptive lasso (Chang et al., 2020, Dessie et al., 2021, Pollard et al., 2021). As will be illustrated in Sections 2 and 3, this generally leads to sub-optimal performance in terms of both support recovery and prediction accuracy.

The rest of the article is organized as follows. In Section 2, we start with a brief overview of the principles of the adaptive lasso, focusing on linear regression models for simplicity. Then, we illustrate the flaw of the simple K -fold cross-validation scheme when used for calibrating the adaptive lasso.

A proper cross-validation scheme in this context is recalled and detailed in Section 3. In Section 4, we present results from a comprehensive simulation study where we empirically illustrate the suboptimality of the simple cross-validation scheme, in terms of support recovery and prediction accuracy. In Section 5, we present results from an application on single-cell data. Concluding remarks are given in Section 6.

2 The adaptive lasso under the linear regression model

2.1 Main notation and working model

As above, we will denote the sample size and the number of covariates by n and p , respectively. For simplicity, we will focus on linear regression models of the form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\xi}, \quad (1)$$

where $\mathbf{y} = (y_1, \dots, y_n)^T \in \mathbb{R}^n$ is the response vector, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}^{n \times p}$ is the design matrix, $\boldsymbol{\beta}^* = (\beta_1^*, \dots, \beta_p^*)^T \in \mathbb{R}^p$ is the p -dimensional vector of unknown parameters to be estimated, and $\boldsymbol{\xi} \in \mathbb{R}^n$ is some random noise. We further denote the support of $\boldsymbol{\beta}^*$ by $J = \{j : \beta_j^* \neq 0\}$.

For any positive integer $d \geq 1$, and any vector $\mathbf{u} \in \mathbb{R}^d$, we will denote the usual Euclidian norm (or L_2 -norm) by $\|\mathbf{u}\|_2$. We will let $\mathbf{0}_d$ and $\mathbf{1}_d$ be the vectors of size d with components all equal to 0 and 1 respectively, and \mathbf{I}_d be the $d \times d$ identity matrix. For any real matrix $\mathbf{M} = (M_1, \dots, M_d) \in \mathbb{R}^{n \times d}$, and any subset $E \subseteq \{1, \dots, d\}$, \mathbf{M}_E will denote the submatrix composed of the columns $(M_j)_{j \in E}$. We will further denote the cardinality of E by $|E|$.

Finally, for any sample $D_0 = \{y_i, \mathbf{x}_i\}_{i \in I_0}$, with I_0 a given set of integers, and any estimator $\hat{\boldsymbol{\beta}}$ of $\boldsymbol{\beta}^*$, we will denote by

$$\text{Pred.Error}(D_0, \hat{\boldsymbol{\beta}}) = \frac{1}{|I_0|} \sum_{i \in I_0} (y_i - \mathbf{x}_i^T \hat{\boldsymbol{\beta}})^2$$

the prediction error corresponding to $\hat{\boldsymbol{\beta}}$ evaluated on the sample D_0 .

2.2 The lasso and adaptive lasso

For any $\lambda \geq 0$, the lasso estimator $\hat{\boldsymbol{\beta}}_{\text{lasso}}(\lambda)$ (Tibshirani, 1996) is defined as any minimizer over $\boldsymbol{\beta} \in \mathbb{R}^p$ of the penalized criterion

$$\frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2}{n} + \lambda \sum_{j=1}^p |\beta_j|. \quad (2)$$

The tuning parameter λ controls the amount of regularization through the L_1 -norm $\|\boldsymbol{\beta}\|_1 = \sum_{j=1}^p |\beta_j|$. In practice, an appropriate value for this parameter is needed to guarantee good statistical performance for $\hat{\boldsymbol{\beta}}_{\text{lasso}}(\lambda)$,

with respect to both support recovery and prediction accuracy. As mentioned above, a popular strategy relies on K -fold cross-validation (Hastie et al., 2009) whose pseudo-code is recalled in Algorithm 3 in Appendix A. Let λ^{CV} be the value of λ selected by K -fold cross-validation, and let $\hat{\boldsymbol{\beta}}_{\text{lasso}}^{\text{CV}} = \hat{\boldsymbol{\beta}}_{\text{lasso}}(\lambda^{\text{CV}})$ denote one particular solution of (2) with λ set to λ^{CV} .

Now, denote by $\mathbf{w} = (w_1, \dots, w_p) \in \mathbb{R}_{\geq 0}^p$ a given vector of non-negative weights. For any $\lambda \geq 0$, the adaptive lasso estimator $\hat{\boldsymbol{\beta}}_{\text{ada}}(\lambda; \mathbf{w})$ is defined (Zou, 2006) as any minimizer over $\boldsymbol{\beta} \in \mathbb{R}^p$ of the criterion

$$\frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2}{n} + \lambda \sum_{j=1}^p w_j |\beta_j|. \quad (3)$$

The adaptive lasso reduces to the standard lasso for the particular choice of the weight vector $\mathbf{w} = \mathbf{1}_p$: any solution $\hat{\boldsymbol{\beta}}_{\text{lasso}}(\lambda)$ is also a solution $\hat{\boldsymbol{\beta}}_{\text{ada}}(\lambda; \mathbf{1}_p)$. In practice, weights are usually set to $w_j = 1/|\tilde{\beta}_j|$, or $w_j = 1/(|\tilde{\beta}_j| + \varepsilon)$, with $\tilde{\boldsymbol{\beta}} = (\tilde{\beta}_1, \dots, \tilde{\beta}_p)^T$ an initial estimator of $\boldsymbol{\beta}^*$, and ε some non-negative real number. A positive value for the ε parameter guarantees that every component of the weight vector is finite, so that every component of $\hat{\boldsymbol{\beta}}_{\text{ada}}(\lambda; \mathbf{w})$ has a chance to be non-zero. If the initial estimator is good enough, then $\tilde{\beta}_j$ is close to 0 for $j \notin J$, and less so for $j \in J$: if, in addition, ε is null or close enough to 0, weights w_j are large for $j \notin J$, and less so for $j \in J$. Then, components $j \notin J$ of the parameter vector are more heavily penalized than components $j \in J$. Several initial estimates can be used to derive the weight vector \mathbf{w} . When $p < n$, Zou (2006) suggests the use of $\mathbf{w}_{\text{OLS}} = 1/|\tilde{\boldsymbol{\beta}}_{\text{OLS}}|$, where $\tilde{\boldsymbol{\beta}}_{\text{OLS}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$ is the OLS estimator. Under mild conditions, and considering the fixed p case, Zou (2006) established the asymptotic oracular property for this ols-adaptive lasso: specifically, the ols-adaptive lasso is consistent in terms of variable selection (or sparsistent), and the distribution of $(\hat{\beta}_{\text{adap},j}(\lambda; \mathbf{w}_{\text{OLS}}))_{j \in J}$ is Gaussian with the same expectation and covariance matrix than that of $(\mathbf{X}_J^T \mathbf{X}_J)^{-1} \mathbf{X}_J^T \mathbf{y}$, the OLS estimator that would be obtained if J were known in advance. The ols-adaptive lasso inherits these good properties from the \sqrt{n} -consistency of the OLS estimate $\tilde{\boldsymbol{\beta}}_{\text{OLS}}$ in low-dimensional settings. However, the OLS estimate is less attractive when $p > n$. In these high-dimensional settings, Zhang et al. (2008) proposed the ridge-adaptive lasso, where ridge regression is used for the computation of the weights: $\mathbf{w}_{\text{ridge}} = 1/|\hat{\boldsymbol{\beta}}_{\text{ridge}}^{\text{CV}}|$. More precisely, $\hat{\boldsymbol{\beta}}_{\text{ridge}}^{\text{CV}}$ denotes the estimator produced by a ridge regression with tuning parameter selected by K -fold cross-validation, and we recall that the ridge regression is a penalized approach similar to the lasso, but with the L_1 -norm in the penalty replaced by the squared L_2 -norm $\|\boldsymbol{\beta}\|_2^2$ (Hoerl and Kennard, 1970). Alternatively, Bühlmann and Meier (2008) suggest the use of weights $\mathbf{w}_{1\text{-step}}$ derived from $\hat{\boldsymbol{\beta}}_{\text{lasso}}^{\text{CV}}$. This leads to what they refer to as the one-step lasso. Thanks to the so-called screening property of $\hat{\boldsymbol{\beta}}_{\text{lasso}}^{\text{CV}}$, the

one-step lasso was shown to be sparsistent under weaker irrepresentability conditions than those required for the lasso (Bühlmann and van De Geer, 2011). Other choices for the weights have been proposed in the literature: for example, Huang et al. (2008b) suggested the use of univariate OLS estimators.

2.3 Illustration of the defect of the simple K -fold cross-validation scheme for the calibration of the adaptive lasso

As recalled in the Introduction, the cross-validated prediction error $(1/K) \times \sum_k \text{Pred.Err}(D^{(k)}, \hat{\beta}_k)$ is widely used to calibrate penalized approaches, including the adaptive lasso. However, many statistical analysts still seem to implement too simple of a scheme for the calibration of the adaptive lasso (Chang et al., 2020, Courtois et al., 2021, Dessie et al., 2021, Pollard et al., 2021). More precisely, they consider the weight vector $\mathbf{w} \in \mathbb{R}_{\geq 0}^p$ derived from initial estimates computed on the full original sample as a “given” (fixed) weight vector. Then, they improperly use the simple cross-validation scheme to compare the predictive performance of the set of estimators $(\hat{\beta}_{\text{ada}}(\lambda_r; \mathbf{w}))_{1 \leq r \leq R}$, for any given sequence $\Lambda = (\lambda_1, \dots, \lambda_R)$ of candidate values for the hyperparameter, and finally select the optimal hyperparameter value, say, $\lambda^{\text{CV}}(\mathbf{w})$, and the corresponding optimal estimator $\hat{\beta}_{\text{ada}}(\lambda^{\text{CV}}(\mathbf{w}); \mathbf{w})$. See the pseudo-code given in Algorithm 3 in Appendix A for the detailed description of this simple, and in this case improper, cross-validation scheme.

Below, we present results from a first simulation study, whose main objectives are to illustrate (i) the good performance of the simple scheme when used for the calibration of the standard lasso, and (ii) its poor performance when applied for the calibration of adaptive lasso. Results from a more comprehensive simulation study will be presented in Section 4.

In this first synthetic example, we generate one sample $D = (y_i, \mathbf{x}_i)_{1 \leq i \leq n}$, made of $n = 1,000$ observations under the linear regression model (1) with $p = 1,000$. We first set $\beta^* = (\beta_1^*, \dots, \beta_p^*)^T$ with $\beta_j^* = 0$ for all $j \geq 11$, and $\beta_j^* = \iota_j 0.5$ for all $j \leq 10$, where ι_j is a $\{-1, 1\}$ -binary random variable, with $\mathbb{P}(\iota_j = 1) = 1/2$. Then, for each $i = 1, \dots, n$, we generate a Gaussian random noise $\xi_i \sim N(0, 1)$, a Gaussian vector of covariates $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p}) \sim N(\mathbf{0}_p, \mathbf{I}_p)$, and finally the outcome $y_i = \mathbf{x}_i^T \beta^* + \xi_i$. Similarly, we generate one independent test sample $\mathcal{D} = (y_i, \mathbf{x}_i)_{n+1 \leq i \leq n+N}$, made of $N = 10,000$ observations drawn under the same linear model. Then, for any given weight vector \mathbf{w} , the `glmnet` R package is used to compute the (adaptive) lasso estimator on an appropriate sequence $\Lambda = \Lambda(D, \mathbf{w}) = (\lambda_1(D, \mathbf{w}), \dots, \lambda_{100}(D, \mathbf{w}))$ of 100 decreasing values for the tuning parameter. The `glmnet` R package is used to compute the 10-fold cross-validated prediction error of $\hat{\beta}_{\text{ada}}(\lambda_r(D, \mathbf{w}); \mathbf{w})$, for $r = 1, \dots, 100$. For comparison, the true prediction error of each estimator is estimated on the independent

test sample \mathcal{D} .

Given the relatively high-dimensional setting of this first simulation study, the ols-adaptive lasso is not considered here. Figure 1 presents the results for the lasso ($\mathbf{w} = \mathbf{1}_p$), the one-step lasso ($\mathbf{w} = \mathbf{w}_{1\text{-step}} = 1/|\hat{\boldsymbol{\beta}}_{\text{lasso}}^{\text{CV}}|$), and the ridge-adaptive lasso ($\mathbf{w} = 1/|\hat{\boldsymbol{\beta}}_{\text{ridge}}^{\text{CV}}|$). For the latter two, we also report the cross-validated prediction error using a proper nested cross-validation scheme, which we will describe in more details in the next Section. For comparability, the x -axis corresponds to the tuning parameter sequence represented as a fraction of the data-specific maximal value $\lambda_1(D, \mathbf{w})$.

First consider the standard lasso (left panel). In this case, the cross-validated prediction error using the simple cross-validation scheme does a fairly good job in approximating the "true" prediction error (estimated on the test sample) on a wide range of λ -values. In particular, the λ -value at which the cross-validated prediction error is minimized (vertical dotted red line) is very close to that at which the prediction error on the test sample is minimized (vertical dotted blue line). Moreover, the prediction errors evaluated on the test sample at these two λ -values (horizontal dotted red and blue lines, respectively) are indistinguishable: in this example, the lasso estimate $\hat{\boldsymbol{\beta}}_{\text{lasso}}^{\text{CV}}$ selected by cross-validation is therefore nearly optimal with respect to prediction error. However, the simple cross-validation scheme does not perform that well for the two versions of the adaptive lasso. In particular, when using the simple scheme for the one-step lasso, the cross-validated prediction error constantly decreases as the tuning-parameter decreases, a behavior that we observed on many other simulation designs as well (results not shown). Then, the λ -value at which the cross-validated prediction error is minimized (vertical dotted red line) is very different from the minimizer of the prediction error evaluated on the test sample. In this example, the support of the one-step lasso estimator is too large if the tuning parameter is selected via the simple cross-validation scheme. This will also usually be the case in the more comprehensive simulation study presented in Section 4. Moreover, the prediction errors evaluated on the test sample at these two λ -values (horizontal dotted red and blue lines, respectively) differ substantially: the one-step-lasso estimator calibrated using the simple cross-validation scheme is far from optimal with respect to prediction error in this example. A similar albeit less pronounced behavior is observed in the case of the ridge-adaptive lasso, confirming that the simple K -fold cross-validation scheme is not valid for the calibration of the adaptive lasso. Conversely, these defects disappear when using a proper scheme whose principle will be recalled in details in the next Section (Ahrens et al., 2020, He et al., 2019, Krämer et al., 2009), which is based on nested cross-validation scheme (Krstajic et al., 2014, Varma and Simon, 2006). Using this proper scheme, the λ -value at which the corresponding cross-validated prediction error is minimized (vertical dotted green line) is close to the minimizer of

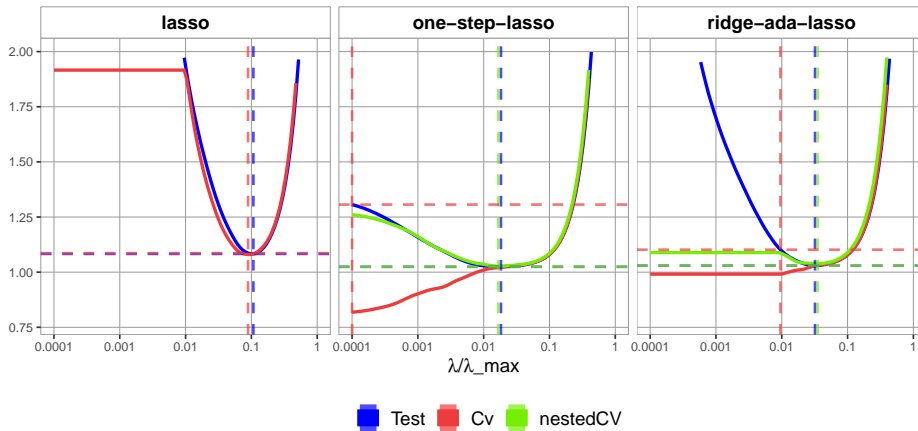


Figure 1: Comparison between the prediction error estimated on an independent test sample (in blue) and the prediction error estimated via the simple cross-validation scheme (in red), for the lasso (left), the one-step-lasso (middle), and the ridge-adaptive lasso (right). For the latter two, the cross-validated prediction error estimated using a proper nested cross-validation scheme is also presented (in green). Vertical dotted lines represent the value of the tuning parameter for which each particular curve is minimized (red: cross-validated prediction error using the simple scheme; green: cross-validated prediction error using the nested scheme; blue: prediction error estimated on an independent test sample). Horizontal dotted lines represent the value of the estimated true prediction error for these particular values of the tuning parameter.

the true prediction error estimated on the test sample for both the one-step lasso and ridge-adaptive lasso. Moreover, the prediction errors evaluated on the test sample at these two λ -values (horizontal dotted green line) are very close to the optimal prediction error: in this example, the one-step-lasso and the ridge-adaptive lasso estimators calibrated using the proper nested cross-validation scheme are nearly optimal with respect to prediction error.

3 Proper cross-validation schemes for the adaptive lasso

In this Section, we provide details on the rationale and principle of proper cross-validation schemes for the adaptive lasso, as previously introduced in Krämer et al. (2009), He et al. (2019) and Ahrens et al. (2020).

3.1 Additional notation

For any sample of observations $D_0 = \{y_i, \mathbf{x}_i\}_{i \in I_0}$, with I_0 a given set of integers, any vector of non-negative weights \mathbf{w} , and any non-negative λ value, we let $\text{Lasso}(D_0, \mathbf{w}, \lambda)$ denote one particular solution of the adaptive lasso (3), when computed on sample D_0 with weights \mathbf{w} and tuning parameter λ . For any positive integer $K \geq 2$, let $\text{sCVLasso}(D_0, \mathbf{w}, K)$ be an adaptive lasso estimator computed on D_0 with weights \mathbf{w} and the tuning parameter set to its optimal value according to the simple K -fold cross-validation scheme of Algorithm 3 in Appendix A; this scheme is improper unless the weights \mathbf{w} are independent from D_0 . On the other hand, $\text{nestedCVLasso}(D_0, \mathbf{w}, K)$ will denote an adaptive lasso estimator computed on D_0 with weights \mathbf{w} and the tuning parameter set to its optimal value according to an appropriate K -fold cross-validation scheme. This proper scheme will be described in details below. We use the notation nestedCVLasso to emphasize that this cross-validation scheme relies on two nested cross-validations in the case of the one-step lasso and the ridge-adaptive lasso; however, this is not the case for the ols-adaptive lasso, and our notation “ nestedCVLasso ” is a bit misleading in the case of the ols-adaptive-lasso. On the other hand, we will use the shorthand $\text{sCVRidge}(D_0, K)$ to denote the ridge estimator computed on D_0 with a tuning parameter set to its optimal value according to the simple K -fold cross-validation, and that will be used to compute the weights in the ridge-adaptive lasso.

3.2 Rationale and detailed presentation

We start by recalling the rationale for the defect of the simple K -fold cross-validation when used for the calibration of the adaptive lasso described and illustrated in Section 2.3. As mentioned in the Introduction, the principle of cross-validation is to mimic “independent” test samples. Then, for the cross-validation to be valid, at each step k of the K -fold cross-validation, the whole estimation procedure should be performed on the training sample $D \setminus D^{(k)}$, and should not use any information from the test sample $D^{(k)}$. However, the weights used in the adaptive lasso are derived from initial estimates computed on the entire original sample D . Therefore, when using the simple cross-validation scheme of Algorithm 3 in Appendix A for the calibration of the adaptive lasso, and considering the estimation of the adaptive lasso estimator as a whole, data from the test samples are used in the first step of the adaptive lasso, which makes the simple cross-validation scheme not valid in this particular framework.

A proper cross-validation scheme for the calibration of the adaptive lasso estimator is described in Algorithm 1 (Ahrens et al., 2020, He et al., 2019, Krämer et al., 2009). The main difference with the simple scheme lies in Step 2. Under the simple scheme, this step consists of Step 2-a, which re-

lies on $\text{sCVLasso}(D, \mathbf{w}, K)$, while it is replaced by Step 2-a' in the proper scheme, which relies on $\text{nestedCVLasso}(D, \mathbf{w}, K)$, detailed in Algorithm 2 in the particular case of the one-step lasso. The key difference between sCVLasso (see Algorithm 3 in Appendix A) and nestedCVLasso is highlighted in blue in Algorithm 2: at each step k of the proper cross-validation scheme, weights used for the adaptive lasso are first recomputed on the training sample $D \setminus D^{(k)}$, so that the whole estimation of the adaptive lasso estimator uses information from the training sample only, before computing the corresponding prediction error on the test sample $D^{(k)}$. In the case of the one-step lasso (or the ridge-adaptive lasso), this leads to a nested cross-validation scheme (Krstajic et al., 2014, Varma and Simon, 2006). More precisely, an appropriate sequence Λ of candidate λ values for the hyperparameter has first to be chosen, which typically depends on the weights computed on the whole original sample (we briefly get back to this point below). Then, at each step k of the “outer” K -fold cross-validation, one “inner” standard cross-validation is performed to compute the optimal lasso (or ridge) estimator on the training sample $D \setminus D^{(k)}$, from which the weights are derived, before the corresponding adaptive lasso estimator is computed for each of the λ values of the sequence, and their predictive performance is eventually evaluated on the test sample $D^{(k)}$. For each λ value of the sequence, the predictive performance is then averaged over the K folds. The optimal value for the hyperparameter is defined as the value that minimizes this averaged criterion. The optimal adaptive lasso estimator finally corresponds to the adaptive lasso computed on D , with weights \mathbf{w} also computed on D , and the hyperparameter set to this optimal value.

A first remark is that, in the case of the ols-adaptive lasso, the proper cross-validation scheme still involves the computation of weights on each training sample in the “outer” loop; however, because weights are derived from OLS estimation, no hyperparameter has to be selected for their computation, hence no inner loop is required. In other words, for the ols-adaptive lasso, the proper cross-validation scheme does not involve nested loops. A second important remark concerns the sequence of candidate values for the hyperparameter. Above, we propose to use the data-driven sequence $\Lambda := \Lambda(D, \mathbf{w}) = (\lambda_1, \dots, \lambda_R)$ computed from the entire original sample, with weights \mathbf{w} also computed on the entire original sample. It could be argued that, by doing so, the estimation process is not completely independent from the test samples. We tested alternative strategies for the selection of the sequence Λ in our simulation, including the data-driven sequence $\Lambda(D, \mathbf{1}_p)$ computed from the full original sample with weight vector set to $\mathbf{1}_p$ that is used in the `adalasso` function of the `parcor` R package (Krämer

et al., 2009), but the results were virtually unchanged.

Algorithm 1: Cross-validation for the adaptive lasso. Versions (i), (ii) and (iii) of step (1-a) correspond to the one-step lasso, the ols-adaptive Lasso, and the ridge-adaptive lasso, respectively. The simple (and improper) scheme corresponds to the algorithm ran with Step (2-a), while Step (2-a') (detailed in Algorithm 2) should be used instead.

Data: Sample: $D = \{y_i, \mathbf{x}_i\}_{i=1}^n$, Version of the adaptive lasso,
Number of folds: K , and parameter $\varepsilon \geq 0$

Result: $\hat{\beta}_{\text{ada}}^{\text{CV}}$

Step 1: Computation of the initial estimates and weights;

(1-a) *Initial estimates: either (i), (ii) or (iii) below, depending on the considered version of the adaptive lasso ;*

(i): $\tilde{\beta} = \text{sCVLasso}(D, \mathbf{1}_p, K)$ /* one-step lasso */

(ii): $\tilde{\beta} = \text{OLS}(D)$ /* ols-adaptive lasso */

(iii): $\tilde{\beta} = \text{sCVRidge}(D, K)$ /* ridge-adaptive lasso */

(1-b) *Weights: $\mathbf{w} = 1/(|\tilde{\beta}| + \varepsilon)$;*

Step 2: Computation of the final estimates;

(2-a) $\hat{\beta}_{\text{ada}}^{\text{CV}}(\mathbf{w}) = \text{sCVLasso}(D, \mathbf{w}, K)$;

(2-a') $\hat{\beta}_{\text{ada}}^{\text{CV}}(\mathbf{w}) = \text{nestedCVLasso}(D, \mathbf{w}, K)$;

Algorithm 2: Details on the nested K -fold cross-validation scheme (nestedCVLasso) for the calibration of the adaptive lasso, in the particular case of the one-step lasso.

Data: Sample: $D = \{y_i, \mathbf{x}_i\}_{i=1}^n$, Weights: \mathbf{w} , Number of folds: K ,
and parameter $\varepsilon \geq 0$

Result: $\hat{\beta}_{\text{ada}}^{n\text{CV}}(\mathbf{w}) = \hat{\beta}_{\text{ada}}(\lambda^{n\text{CV}}, \mathbf{w}, D)$

Computation of a sequence $\Lambda := \Lambda(D, \mathbf{w}) = (\lambda_1, \dots, \lambda_R)$;

Division of D into K folds: $D = \cup_{k=1}^K D^{(k)}$;

for $k \in \{1, \dots, K\}$ **do**

 /* Computation of the weights on $D \setminus D^{(k)}$ */

$\tilde{\beta}^{(k)} = \text{sCVLasso}(D \setminus D^{(k)}, \mathbf{1}_p, K)$;

$\mathbf{w}_k = 1/(|\tilde{\beta}^{(k)}| + \varepsilon)$;

for $r \in \{1, \dots, R\}$ **do**

$\hat{\beta}_{r,k} = \text{Lasso}(D \setminus D^{(k)}, \mathbf{w}_k, \lambda_r)$;

$E_{r,k} = \text{Pred.Error}(D^{(k)}, \hat{\beta}_{r,k})$;

end

end

$r^* = \text{argmin}_r \{ \sum_{k=1}^K E_{r,k} \}$;

$\lambda^{n\text{CV}} = \lambda_{r^*}$;

$\hat{\beta}_{\text{ada}}^{n\text{CV}}(\mathbf{w}) = \text{Lasso}(D, \mathbf{w}, \lambda^{n\text{CV}})$;

4 Simulation study

We now present results from a more comprehensive simulation study, which extends the simple one presented in Section 2.3. As before, we generate a sample $D = (y_i, \mathbf{x}_i)_{1 \leq i \leq n}$, made of $n = 500$ observations drawn under linear regression model (1). Here, we make the number of covariates p vary in $\{100, 500, 1000\}$, and set the number of relevant covariates to $p_0 = 10$. We randomly select the support J of β^* , with $|J| = p_0$. Then, we set $\beta_j^* = 0$ for all $j \notin J$, and $\beta_j^* = \iota_j \beta$, for all $j \in J$, where ι_j is a $\{-1, 1\}$ -binary variable, with $\mathbb{P}(\iota_j = 1) = 1/2$. As for the signal strength β , we make it vary in $\{1/4, 1/2, 1, 3/2\}$. Then, for each $i = 1, \dots, n$, we generate a Gaussian random noise $\xi_i \sim N(0, 1)$, a Gaussian vector of covariates $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p}) \sim N(\mathbf{0}_p, \mathbf{\Sigma})$, with $\mathbf{\Sigma}$ a symmetric $p \times p$ matrix, with element (k, j) equal to $\sigma_{k,j} = 0.3^{|k-j|}$, for $1 \leq k, j \leq p$. Finally the outcome for observation i is generated as $y_i = \mathbf{x}_i^T \beta^* + \xi_i$. Additionally, we generate one independent test sample $\mathcal{D} = (y_i, \mathbf{x}_i)_{n+1 \leq i \leq n+N}$, made of $N = 10,000$ observations drawn under the same linear model.

We consider the one-step lasso and the ridge-adaptive lasso as before. The ols-adaptive lasso is also considered in the low-dimensional scenario where $p = 100$. As in the previous study, weights were derived from initial estimates $\tilde{\beta}$ as $w_j = 1/|\tilde{\beta}_j|$; other choices for the parameter $\varepsilon \in \{0, 10^{-6}, 10^{-4}, 10^{-2}\}$ were tested, and led to very similar results (not shown). For each method, the optimal adaptive lasso estimator is computed following either the simple 10-fold cross-validation scheme (Algorithm 1 with Step 2-a; we will refer to these estimators as the one-step lasso CV, the ridge-adaptive lasso CV, and the ols-adaptive lasso CV respectively) or the 10-fold ‘‘nested’’ cross-validation scheme (Algorithm 1 with Step 2-a’; we will refer to these estimators as the one-step lasso nested CV, the ridge-adaptive lasso nested CV, and the ols-adaptive lasso nested CV, respectively; keep in mind that for the ols-adaptive lasso nested CV, the terminology is a bit misleading, as the cross-validation scheme does not involve nested loops). For comparison, results from the lasso estimator calibrated by standard 10-fold cross-validation (see Algorithm 3 in Appendix A) are presented; we will refer to this estimator as the lasso CV. We used functions from the `glmnet` R package to compute these different estimators. We evaluate the (signed) support accuracy, the precision, the recall and the prediction error corresponding to each estimator $\hat{\beta}$. More precisely, the signed support accuracy is defined as $\text{sACC} = \{\sum_{j=1}^p \mathbb{I}(\text{sign}(\beta_j^*) = \text{sign}(\hat{\beta}_j))\}/p$, where \mathbb{I} is the indicator function, and $\text{sign}(\cdot)$ the sign function, that is $\text{sign}(x) = +1$ if $x > 0$, $\text{sign}(x) = -1$ if $x < 0$, and $\text{sign}(x) = 0$ if $x = 0$. Precision and recall are defined as $\sum_{j=1}^p \mathbb{I}[\beta_j^* \neq 0, \hat{\beta}_j \neq 0] / \sum_{j=1}^p \mathbb{I}[\hat{\beta}_j \neq 0]$ and $\sum_{j=1}^p \mathbb{I}[\beta_j^* \neq 0, \hat{\beta}_j \neq 0] / p_0$, respectively. As for the prediction error, we simply compute $\text{Pred.Err}(\mathcal{D}, \hat{\beta})$, as before. These four criteria are averaged

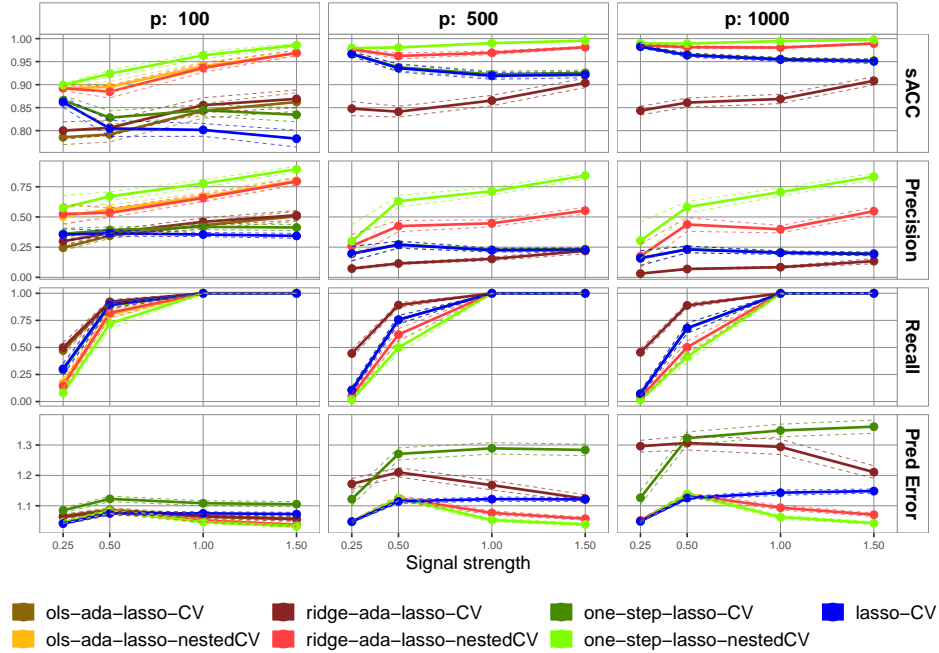


Figure 2: Results of the simulation study. Solid lines represent the averaged criteria, while dashed lines represent the corresponding 95% confidence intervals.

over the 100 replications we consider for each combination of values for the parameters (p, β) .

Figure 2 presents the results. First, focusing on the adaptive lasso estimators, our results illustrate that the proper cross-validation scheme yields better performance than the simple scheme, in terms of both prediction error and support accuracy. In terms of support recovery, the gain actually only concerns precision (or, equivalently, specificity). This observation confirms what we already observed in Section 2.3: the simple cross-validation scheme generally leads to the selection of too small of hyperparameters, hence poor precision/specificity. Conversely, recall (which is related to sensitivity) is usually better with the simple cross-validation scheme, but the balance is clearly in favour of the proper cross-validation scheme, as it usually yields substantially better overall support recovery. We also observe that, when correctly calibrated, the adaptive lasso usually outperforms the lasso in terms of support recovery (again, the gain comes from a better precision/specificity, as expected), and prediction error. However, the two methods achieve very similar performance for low signal strength.

Finally, Figure 3 presents results from an additional simulation study, using a more realistic scenario, where predictors are highly correlated. Specifically, our synthetic data are generated just as above, but this time with

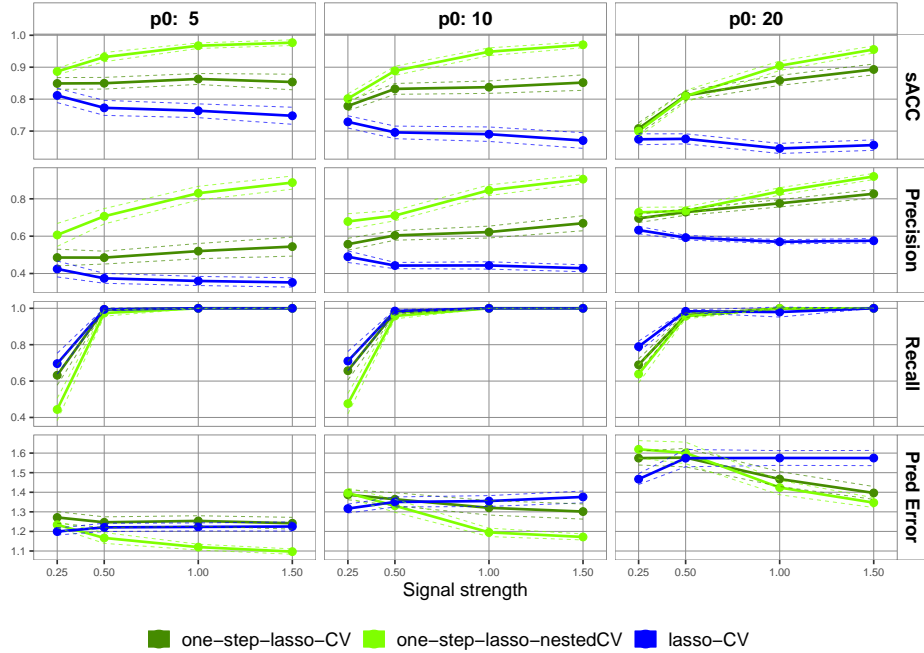


Figure 3: Results of the additional simulation study inspired by the single cell data. Solid lines represent the averaged criteria, while dashed lines represent the corresponding 95% confidence intervals.

$n = 100$ and $p = 44$, and with the variance-covariance matrix Σ of the predictors set to the empirical variance-covariance matrix of 44 of the transcription factors from the real data presented in Section 5. In this case, we make the number of non-zero parameters in vector β^* vary in $p_0 \in \{5, 10, 20\}$. We compute the one-step lasso, calibrated with either the simple cross-validation scheme or the nested one, as well as the standard lasso. Results from this additional simulation study are mostly consistent with those of the previous simulation study. They confirm that calibration with the nested cross-validation scheme generally leads to better performance of the one-step lasso, in terms of support recovery (in particular precision/specificity) and prediction error. They also highlight that for low signal strengths, the proper nested scheme does not always lead to improved prediction error, especially for large p_0 , compared to the simple (and improper) scheme. Moreover, for low signal strengths, the standard lasso performs a bit better than the one-step lasso in terms of prediction error in this setting.

5 Application to single-cell data

In this Section, we use single-cell data to illustrate the performance of the one-step lasso calibrated by 10-fold cross-validation, using either the simple or the nested scheme. For comparison, we also implement the lasso calibrated by the simple cross-validation scheme, which is appropriate in this case. We analyse data describing myelocytic leukemia cells undergoing differentiation to macrophage, described in Kouno et al. (2013); see also Ollier and Viallon (2017). Briefly, expression levels of 45 transcription factors are measured on $n = 120$ cells at 8 distinct time-points of this differentiation process. The main objective of Kouno et al. (2013) was to determine how associations among the 45 transcription factors vary over time. The authors focused on marginal associations and used univariate analyses, although graphical models, which describe conditional associations, might be better suited. Identification of the structure of a graphical model can be reduced to the identification and description of the neighbourhood of each covariate (Meinshausen et al., 2006). For illustration, we here focus on the first time-point, and consider, successively, each of the 45 transcription factors as the outcome of interest and the other 44 factors as the predictors. This leads to the construction of 45 linear regression models, which together can be used to construct a graphical model. Because it is out of the scope of the present work, we do not discuss, nor present, the estimated graphical models here. For each linear regression model, the prediction performance of each method (one-step lasso with the simple or nested cross-validation scheme and the standard lasso) is evaluated by 5-fold cross-validation repeated 100 times; see Figure 4. For 39 of the 45 considered models, the one-step lasso achieves a lower prediction error when calibrated using the proper nested cross-validation scheme. Moreover, even if the prediction error is generally better for the standard lasso compared to the one-step lasso (which may be explained by the low sample size in this application, and possibly low signal too), the one-step lasso outperforms the lasso for 13 models when using the proper nested cross-validation scheme for its calibration, while it never does when using the improper cross-validation scheme. Comparing the sizes of the supports for the parameter vectors produced by each method, the cumulative size of the supports for the 45 models is 464 for the lasso, 395 for the one-step lasso with improper cross-validation, and only 217 when proper cross-validation is used. Overall, these results confirm that the simple cross-validation scheme is generally not appropriate for the calibration of the one-step lasso, and that the nested cross-validation scheme usually yields sparser parameter estimates and lower prediction error.

6 Discussion-Conclusion

In this article, we described a defect of the simple K -fold cross-validation scheme when applied for the calibration of the adaptive lasso, with emphasis on the ols-adaptive and ridge-adaptive lasso, as well as the one-step lasso. We further described a refined cross-validation scheme, which is appropriate for the calibration of the adaptive lasso. Although well known by specialists, we believe the defect of the simple scheme has not been stressed enough in the literature, and many statistical analysts still overlook it when calibrating the adaptive lasso. In particular, although the general principle according to which the dataset used to train the model should not contain information from hold-out sample used to evaluate prediction performance is clearly violated when using the simple scheme for the calibration of the adaptive lasso, the specific implications of applying the cross-validation after constructing the weights on the full dataset are not as straightforward to appreciate as in other situations, i.e., when applying cross-validation after screening features on the full dataset.

We shall also stress that, although we focused on the K -fold cross-validation here, our description applies to other versions of cross-validation (Arlot and Celisse, 2010). Also, we focused on the adaptive lasso under linear regression models for simplicity, but similar defects are expected for the simple cross-validation scheme in the case of the adaptive generalized fused lasso (Viallon et al., 2016), the adaptive data shared lasso (Ballout et al., 2020, Gross and Tibshirani, 2016, Ollier and Viallon, 2017)), as well as the adaptive lasso under alternative regression models, such as generalized linear models and Cox proportional hazard models (He et al., 2019).

Disclaimer

Where authors are identified as personnel of the International Agency for Research on Cancer/World Health Organization, the authors alone are responsible for the views expressed in this article and they do not necessarily represent the decisions, policy or views of the International Agency for Research on Cancer /World Health Organization.

Declarations

Funding

Not applicable

Conflicts of interest/Competing interests

The authors declare no conflict of interest

Availability of data and material

Not applicable

Code availability

All our empirical results can be replicated using our R scripts available on <https://github.com/NadimBLT/AdapGlmnet-NestedCv>. In particular the `Adalasso.Nestedcv` function returns the optimal adaptive lasso estimator, with hyperparameter selected by the “nested” cross-validated scheme, in the case of the ols-adaptive lasso, the ridge-adaptive lasso and the one-step lasso.

References

- Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2020). lassopack: Model selection and prediction with regularized regression in stata. The Stata Journal, 20(1):176–235.
- Arlot, S. (2008). V -fold cross-validation improved: V -fold penalization. arXiv preprint arXiv:0802.0566v2.
- Arlot, S. (2019). Minimal penalties and the slope heuristics: a survey. arXiv preprint arXiv:1901.07277.
- Arlot, S. and Celisse, A. (2010). A survey of cross-validation procedures for model selection. Statistics surveys, 4:40–79.
- Ballout, N., Garcia, C., and Viallon, V. (2020). Sparse estimation for case-control studies with multiple disease subtypes. Biostatistics. kxz063.
- Ballout, N. and Viallon, V. (2019). Structure estimation of binary graphical models on stratified data: Application to the description of injury tables for victims of road accidents. Statistics in Medicine, 38(14):2680–2703.
- Bates, S., Hastie, T., and Tibshirani, R. (2021). Cross-validation: what does it estimate and how well does it do it? arXiv preprint stat.ME:2104.00673.
- Bühlmann, P. and Meier, L. (2008). Discussion: One-step sparse estimates in nonconcave penalized likelihood models. Ann. Statist., 36(4):1534–1541.
- Bühlmann, P. and van De Geer, S. (2011). Statistics for high-dimensional data: methods, theory and applications. Springer Science & Business Media.

- Chang, C., Sung, C.-Y., Hsiao, H., Chen, J., Chen, I.-H., Kuo, W.-T., Cheng, L.-F., Korla, P. K., Chung, M.-J., Wu, P.-J., et al. (2020). Hd-mac: a web-based interactive program for high-dimensional analysis of molecular alterations in cancer. Scientific reports, 10(1):1–10.
- Chen, J. and Chen, Z. (2008). Extended bayesian information criteria for model selection with large model spaces. Biometrika, 95(3):759–771.
- Chichignoud, M., Lederer, J., and Wainwright, M. J. (2016). A practical scheme and fast algorithm to tune the lasso with optimality guarantees. The Journal of Machine Learning Research, 17(1):8162–8181.
- Courtois, É., Tubert-Bitter, P., and Ahmed, I. (2021). New adaptive lasso approaches for variable selection in automated pharmacovigilance signal detection. Research Square preprint DOI: 10.21203/rs.3.rs-384204/v1.
- Dessie, E. Y., Tsai, J. J., Chang, J.-G., and Ng, K.-L. (2021). A novel mirna-based classification model of risks and stages for clear cell renal cell carcinoma patients. BMC bioinformatics, 22(10):1–16.
- Donoho, D. L. et al. (2000). High-dimensional data analysis: The curses and blessings of dimensionality. AMS math challenges lecture, 1(2000):32.
- Fan, J. and Li, R. (2006). Statistical challenges with high dimensionality: Feature selection in knowledge discovery. arXiv preprint math/0602133.
- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. Journal of statistical software, 33(1):1.
- Giacobino, C., Sardy, S., Diaz-Rodriguez, J., Hengartner, N., et al. (2017). Quantile universal threshold. Electronic Journal of Statistics, 11(2):4701–4722.
- Giraud, C. (2014). Introduction to high-dimensional statistics. Chapman and Hall/CRC.
- Gross, S. M. and Tibshirani, R. (2016). Data shared lasso: A novel tool to discover uplift. Computational statistics & data analysis, 101:226–235.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). The elements of statistical learning: data mining, inference, and prediction. Springer Science & Business Media.
- Hastie, T., Tibshirani, R., and Wainwright, M. (2015). Statistical learning with sparsity: the lasso and generalizations. Chapman and Hall/CRC.

- He, K., Wang, Y., Zhou, X., Xu, H., and Huang, C. (2019). An improved variable selection procedure for adaptive lasso in high-dimensional survival analysis. Lifetime data analysis, 25(3):569–585.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. Technometrics, 12(1):55–67.
- Huang, J., Horowitz, J. L., Ma, S., et al. (2008a). Asymptotic properties of bridge estimators in sparse high-dimensional regression models. The Annals of Statistics, 36(2):587–613.
- Huang, J., Ma, S., and Zhang, C.-H. (2008b). Adaptive lasso for sparse high-dimensional regression models. Statistica Sinica, pages 1603–1618.
- Jacob, L., Obozinski, G., and Vert, J.-P. (2009). Group lasso with overlap and graph lasso. In Proceedings of the 26th annual international conference on machine learning, pages 433–440.
- Kouno, T., de Hoon, M., Mar, J. C., Tomaru, Y., Kawano, M., Carninci, P., Suzuki, H., Hayashizaki, Y., and Shin, J. W. (2013). Temporal dynamics and transcriptional control using single-cell gene expression analysis. Genome biology, 14(10):1–12.
- Krämer, N., Schäfer, J., and Boulesteix, A.-L. (2009). Regularized estimation of large-scale gene association networks using graphical gaussian models. BMC bioinformatics, 10(1):1–24.
- Krstajic, D., Buturovic, L. J., Leahy, D. E., and Thomas, S. (2014). Cross-validation pitfalls when selecting and assessing regression and classification models. Journal of cheminformatics, 6(1):1–15.
- Meinshausen, N., Bühlmann, P., et al. (2006). High-dimensional graphs and variable selection with the lasso. The annals of statistics, 34(3):1436–1462.
- Ollier, E. and Viallon, V. (2017). Regression modeling on stratified data with the lasso. Biometrika, 104(1):84–96.
- Pollard, J. D., Haq, K. T., Lutz, K. J., Rogovoy, N. M., Paternostro, K. A., Soliman, E. Z., Maher, J., Lima, J. A., Musani, S. K., and Tereshchenko, L. G. (2021). Electrocardiogram machine learning for detection of cardiovascular disease in african americans: the jackson heart study. European Heart Journal-Digital Health, 2(1):137–151.
- Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. Journal of the Royal Statistical Society, Series B, 36:111–147.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Methodological), 58(1):267–288.

- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused lasso. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 67(1):91–108.
- Varma, S. and Simon, R. (2006). Bias in error estimation when using cross-validation for model selection. BMC bioinformatics, 7(1):1–8.
- Viallon, V., Lambert-Lacroix, S., Hoefling, H., and Picard, F. (2016). On the robustness of the generalized fused lasso to prior specifications. Statistics and Computing, 26(1):285–301.
- Zhang, J., Jeng, X. J., and Liu, H. (2008). Some two-step procedures for variable selection in high-dimensional linear regression. arXiv preprint arXiv:0810.1644.
- Zou, H. (2006). The adaptive lasso and its oracle properties. Journal of the American statistical association, 101(476):1418–1429.

A Pseudo-code of the simple K -fold cross-validation scheme for the calibration of the (adaptive) lasso

The pseudo-code detailed in Algorithm 3 below describes the function sCVLasso: when applied to the sample $D = \{y_i, \mathbf{x}_i\}_{i=1}^n$, with weights \mathbf{w} and number of folds K , it produces the adaptive Lasso estimator $\hat{\beta}_{\text{ada}}^{\text{CV}}(\mathbf{w}) = \hat{\beta}_{\text{ada}}(\lambda^{\text{CV}}, \mathbf{w})$, with hyperparameter λ^{CV} set to its value minimizing a K -fold cross-validated prediction error. Keep in mind that this simple scheme is not appropriate if weights \mathbf{w} are derived from the data, as is the case for the the ols-adaptive lasso, ridge-adaptive lasso and one-step lasso, estimators.

Algorithm 3: The simple K -fold cross-validation sCVLasso(D, \mathbf{w}, K) for the calibration the lasso. We recall that this scheme is not valid and should therefore not be used in the case of the adaptive lasso.

Data: Sample: $D = \{y_i, \mathbf{x}_i\}_{i=1}^n$, Weights: \mathbf{w} , Number of folds: K

Result: $\hat{\beta}_{\text{ada}}^{\text{CV}}(\mathbf{w})$

Computation of a sequence $\Lambda := \Lambda(D, \mathbf{w}) = (\lambda_1, \dots, \lambda_R)$;

Division of D into K folds: $D = \cup_{k=1}^K \{D^{(k)}\}$;

for $k \in \{1, \dots, K\}$ **do**

for $r \in \{1, \dots, R\}$ **do**

$\hat{\beta}_{r,k} = \text{Lasso}(D \setminus D^{(k)}, \mathbf{w}, \lambda_r)$;

$E_{r,k} = \text{Pred.Error}(D^{(k)}, \hat{\beta}_{r,k})$;

end

end

$r^* = \text{argmin}_r \{ \sum_{k=1}^K E_{r,k} \}$;

$\lambda^{\text{CV}} = \lambda_{r^*}$;

$\hat{\beta}_{\text{ada}}^{\text{CV}}(\mathbf{w}) = \hat{\beta}_{\text{ada}}(\lambda^{\text{CV}}, \mathbf{w})$;

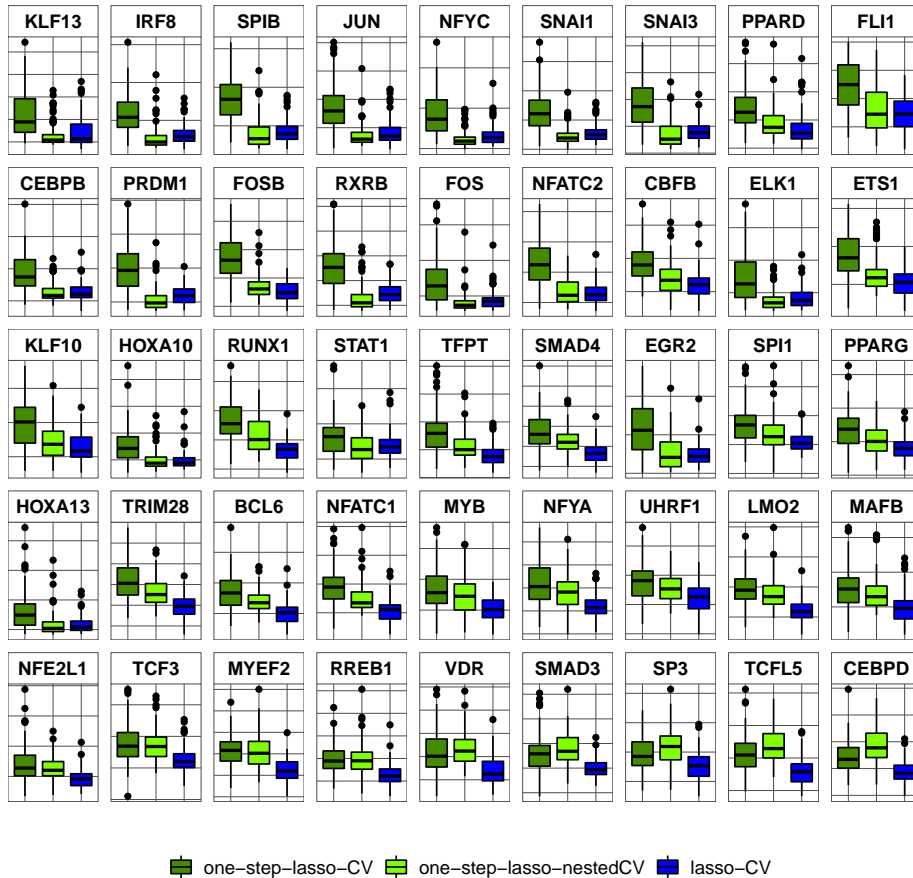


Figure 4: Prediction performance of the three methods in the 45 considered linear regression models (each model corresponds to the choice of one particular transcription factor as the outcome of interest), evaluated by 5-fold cross-validation repeated 100 times. Each box-plot represents the empirical distribution of the 5-fold cross-validated prediction error over the 100 repetitions.