



# A cost-sensitive constrained Lasso

Rafael Blanquero<sup>1,2</sup> · Emilio Carrizosa<sup>1,2</sup> · Pepa Ramírez-Cobo<sup>2,3</sup> ·  
M. Remedios Sillero-Denamiel<sup>1,2</sup>

Received: 16 April 2019 / Revised: 26 February 2020 / Accepted: 2 March 2020  
© Springer-Verlag GmbH Germany, part of Springer Nature 2020

## Abstract

The Lasso has become a benchmark data analysis procedure, and numerous variants have been proposed in the literature. Although the Lasso formulations are stated so that overall prediction error is optimized, no full control over the accuracy prediction on certain individuals of interest is allowed. In this work we propose a novel version of the Lasso in which quadratic performance constraints are added to Lasso-based objective functions, in such a way that threshold values are set to bound the prediction errors in the different groups of interest (not necessarily disjoint). As a result, a constrained sparse regression model is defined by a nonlinear optimization problem. This cost-sensitive constrained Lasso has a direct application in heterogeneous samples where data are collected from distinct sources, as it is standard in many biomedical contexts. Both theoretical properties and empirical studies concerning the new method are explored in this paper. In addition, two illustrations of the method on biomedical and sociological contexts are considered.

**Keywords** Performance constraints · Cost-sensitive learning · Sparse solutions · Sample average approximation · Heterogeneity · Lasso

**Mathematics Subject Classification** 62-07 Data analysis · 62H12 Estimation in multivariate analysis · 62J07 Ridge regression; shrinkage estimators (Lasso) · 62P99 Applications of statistics

## 1 Introduction

Let  $(Y, X)$  be a random vector, where  $X = (X_1, \dots, X_p)$  is a vector of  $p$  predictors and  $Y$  identifies the response variable. Given the observed response vector

---

✉ M. Remedios Sillero-Denamiel  
rsillero@us.es

<sup>1</sup> Departamento de Estadística e Investigación Operativa, Universidad de Sevilla, Seville, Spain

<sup>2</sup> Instituto de Matemáticas de la Universidad de Sevilla (IMUS), Seville, Spain

<sup>3</sup> Departamento de Estadística e Investigación Operativa, Universidad de Cádiz, Cádiz, Spain

$\mathbf{y} = (y_1, \dots, y_n)'$ ,  $n > p$ , and the related observed predictors,  $\mathbf{x}_j = (x_{1j}, \dots, x_{nj})'$ ,  $j = 1, \dots, p$ , the linear regression model predicts  $\mathbf{y}$  by

$$\hat{\mathbf{y}} = \hat{\beta}_0 + \hat{\beta}_1 \mathbf{x}_1 + \dots + \hat{\beta}_p \mathbf{x}_p.$$

Consider the well-known `prostate` database (Stamey et al. 1989), which consists of the measurements of  $p = 8$  predictors and one response variable (clinical measures) on  $n = 97$  men who were about to receive a radical prostatectomy. Further, assume that the dataset is divided into two groups: *Group 1*, corresponding to *young* individuals (aged less than 65) and *Group 2*, related to the population older than 65. If the goal is to minimize the overall mean squared error (MSE), the parameter vector  $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_p)$  can be estimated by a fitting procedure as ordinary least squares (OLS), yielding  $\hat{\boldsymbol{\beta}}^{ols}$ . The results obtained under the OLS are shown in the first two rows of Table 1, where 3/4 of the total set has been used to fit the model (training set) and the remaining samples for the assessment of the generalization error of the resulting model (test set). The overall MSE, the prediction errors over the two groups as well as the number of coefficients involved in the model are presented.

Once the model is fitted, there are two fundamental criteria for evaluating its performance: the accuracy of prediction and the identification of significant predictors, which provides a good interpretation of the solution. It is well-known that  $\hat{\boldsymbol{\beta}}^{ols}$  is not sparse, as can be observed from Table 1 where the eight predictor variables have been used by the model. This fact has entailed new penalization techniques as ridge regression (Hoerl and Kennard 1970), which is a continuous shrinkage method that achieves its better accuracy prediction through a bias-variance trade-off. Nevertheless, ridge regression is known not to be able to render a parsimonious solution. In contrast, best-subset selection (Garside 1965) achieves more sparse solutions, but it suffers from high variability and computational difficulties (Fan and Li 2001). To overcome those shortcomings, Tibshirani (1996) proposed the Lasso regularization technique, which achieves both estimation and selection of relevant predictors simultaneously by construction. Given  $\mathbf{X} = [\mathbf{1} \mid \mathbf{x}_1 \mid \dots \mid \mathbf{x}_p]$  the predictor matrix; then, the Lasso solution can be defined as

$$\hat{\boldsymbol{\beta}}^{Lasso}(\lambda) = \arg \min_{\boldsymbol{\beta}} \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 + \lambda \|(\beta_1, \dots, \beta_p)\|_1 \quad (1)$$

**Table 1** Results obtained using `prostate` dataset

Method		Overall MSE	Group 1 MSE	Group 2 MSE	Non-zero coefficients
OLS	Training set	0.344	0.355	0.333	8
	Test set	0.373	0.380	0.367	
Lasso	Training set	0.365	0.397	0.335	5
	Test set	0.408	0.414	0.403	
CSCLasso	Training set	0.355	0.357	0.352	6
	Test set	0.393	0.399	0.388	

where  $\lambda \geq 0$  is a tuning parameter and  $\|\cdot\|_1$  is the  $l_1$  norm. To visualize the effect of the penalty term in the Lasso formulation, consider the third and fourth rows in Table 1, which provide the results obtained under the Lasso for the `prostate` dataset. In this case, in comparison with OLS results, a sparser, and therefore a more interpretable solution, has been obtained at the expense of slightly worsening MSE values.

One of the advantages of the Lasso is that the entire path of solutions can be found thanks to the LARS algorithm (Efron et al. 2004). In addition, it is well-known that, under some conditions, the Lasso enjoys good theoretical and statistical properties (Donoho et al. 1995; Friedman et al. 2001; Bühlmann and Van-De Geer 2011). However, the Lasso presents some limitations; in particular, the literature related to the Lasso has not undertaken the problem of fully controlling the accuracy prediction on certain individuals of interest. In the previous `prostate` database, assume for instance that we are interested in fitting a sparse regression model to the dataset where, apart from obtaining a small overall mean squared error, also the prediction error for the *young* individuals should not exceed a given threshold. In this paper we propose a Lasso-based model that allows for such aim, namely the cost-sensitive constrained Lasso, denoted from now on as CSCLasso. The results obtained for the `prostate` database under the CSCLasso, whose definition and main properties shall be discussed in Sects. 2 and 3, are shown in the last two rows of Table 1. A threshold for the mean squared error over *Group 1* is set equal to 0.357, which represents an improvement of 10% over the prediction error of the Lasso (0.397). Note that in the training set the MSE satisfies the imposed constraint, as expected. Also note that the improvement in *Group 1* is at the expense of slightly increasing the prediction error over *Group 2*. In terms of sparsity, the CSCLasso model has needed an additional predictor variable comparing to Lasso in order to comply with the constraint.

As it will be seen in Sect. 2, the novel approach is set up by adding convex quadratic constraints to the Lasso formulation. Other approaches have considered constrained versions of the Lasso before, see for example (James et al. 2019; Gaines et al. 2018; Torres-Barrán et al. 2018; Hu et al. 2015 and references therein). In such works, equality or/and inequality linear constraints are considered for imposing prior knowledge and structure onto the coefficient estimates. In our approach instead, quadratic convex constraints are formulated and thus, our approach and results generalize those previously obtained in the literature.

Not only constrained versions of the Lasso can be found in the literature. Indeed, many different variants have been proposed. For example, in Zou (2006) adaptive weights for penalizing different coefficients in the  $l_1$  penalty are included as a way for fitting sparser models under more general conditions. Moreover, in the presence of highly correlated predictor variables (as is usual in microarray studies) or when predictors are structurally grouped (e.g. dummy variables), the Lasso sometimes does not perform well and, as a consequence, the *elastic net* (Zou and Hastie 2005) and the *group lasso* (Yuan and Lin 2006; Simon et al. 2011) were proposed. They combine  $l_2$  and  $l_1$  penalties to try to select (or remove) the correlated or structured predictor variables together. Other extension is to consider

$$\hat{\beta}^{Lasso}(\lambda) = \arg \min_{\beta} \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|^2 + \lambda \|\mathcal{A}\beta\|_1. \quad (2)$$

instead of (1), where  $\mathcal{A}$  is a fixed matrix (see Tibshirani and Taylor 2011). If  $\mathcal{A} = (0|I^{p \times p})$ , then the Lasso objective function is obtained; however, other forms of  $\mathcal{A}$  different from the identity can be found in the literature, see for example (Ollier and Viallon 2017). In fact, various choices of  $\mathcal{A}$  in (2) define problems that are already well-known in the literature as the *fused lasso* (Tibshirani et al. 2005). See Hastie et al. (2015) for an extensive review about Lasso problem and generalizations.

The motivation of this paper (controlling the performance measure on certain groups of interest) is not novel in the Data Analysis literature, and indeed, it has been exploited in classification contexts with the term *cost-sensitive learning* (Prati et al. 2015; He and Ma 2013). Many realworld problems, such as those found in medical diagnosis or credit card fraud detection, have asymmetric misclassification costs associated, since the consequences of wrong predictions across the classes may be very different. Therefore, for these problems, it is more important to achieve better classification rates for the individuals of interest (ill people, defaulting customers). See Carrizosa et al. (2008), Sun et al. (2009), Datta and Das (2015), Lee et al. (2017), Bradford et al. (1998) and Freitas et al. (2007) for more details and applications. Those methods are based on adapting the classifier construction or adding parameters, among others. As some examples, consider (Datta and Das 2015; Carrizosa et al. 2008) and Lee et al. (2017), which adapt the support vector machine (SVM) classifier. In Datta and Das (2015) the decision boundary shift is combined with unequal misclassification penalties. A biobjective problem is introduced in Carrizosa et al. (2008), which simultaneously minimizes the misclassification rates. In Lee et al. (2017), the authors propose a new weight adjustment factor that is applied to a weighted SVM.

This paper is structured as follows. In Sect. 2, the cost-sensitive constrained Lasso (CSCLasso) is introduced and some key issues are discussed. Section 3 considers theoretical properties of the CSCLasso, as the existence and uniqueness of solution, limit behaviour (in terms of the penalty parameter) and consistency. Section 4 presents a detailed numerical analysis with both simulated and real datasets, and finally, some conclusions and extensions are provided in Sect. 5. Technical proofs are relegated to the Appendix.

## 2 The cost-sensitive constrained Lasso: definition and key aspects

This section presents the cost-sensitive constrained Lasso, which, as will be seen, is defined through an optimization problem with constraints related to prediction errors for individuals of interest. In addition, some computational details, as well as different key aspects concerning the tuning parameters of our proposal, are presented.

### 2.1 Definition

The proposed CSCLasso is a novel variant of the Lasso where we shall demand that the prediction errors for the groups of interest are below certain threshold values,

$$\begin{aligned}
 \min_{\beta} \quad & \frac{1}{n_0} \|\mathbf{y}_0 - \mathbf{X}_0\beta\|^2 + \lambda \|\mathcal{A}\beta\|_1 \\
 \text{s.t.} \quad & \frac{1}{n_1} \|\mathbf{y}_1 - \mathbf{X}_1\beta\|^2 - f_1 \leq 0, \\
 & \vdots \\
 & \frac{1}{n_L} \|\mathbf{y}_L - \mathbf{X}_L\beta\|^2 - f_L \leq 0.
 \end{aligned} \tag{3}$$

In the previous formulation,  $(\mathbf{y}_0, \mathbf{X}_0)$  is the set of observations used to build the sparse model with overall minimum MSE, which can be the complete dataset  $(\mathbf{y}, \mathbf{X})$ , or a subset of smaller size. Additionally, let  $(\mathbf{y}_l, \mathbf{X}_l), l = 1, \dots, L$ , define groups of interest (not necessarily disjoint), where the MSE predictions are to be controlled. Then,  $n_l$  is the number of instances related to group  $l$ . Finally,  $f = (f_1, \dots, f_L)$  contains the different threshold values for the MSE on the different groups. The solution of optimization problem (3) will be denoted by  $\hat{\beta}^{CSCLasso}(\lambda)$ . From the formulation (3) it is natural to wonder whether running a Lasso on just the groups of interest is more advantageous. However, if a single Lasso is run on the groups of interest, dramatically bad predictions can be obtained when the resulting model is applied to new observations outside those groups, which is not the case for our approach. The same issue arises when a different Lasso model is built on each group of interest, but, in addition, new observations are not given with their group of origin. Contrary to what happens with our novel approach (3), the  $L$  predictions obtained through the  $L$  different estimated Lasso models may not be suitable to give a final prediction for such new samples.

The proposed method can be formulated as a Lasso with weighted quadratic penalties in the objective function associated with the different groups, but finding real meaning to their parameters (one per group) to be chosen is not an easy task (see Carrizosa and Romero-Morales (2001) and the references therein) and the full control over the accuracy prediction on certain individuals of interest would disappear. However, the parameters  $f = (f_1, \dots, f_L)$  involved in our model have a clear interpretation and, in addition, this formulation enables us to bound the prediction errors in the different groups of interest.

As an example, in the illustration of the method in Sect. 1 related to the prostate dataset, whereas the training set was used in the objective function with  $\mathcal{A} = (0|I^{8 \times 8})$ , the prediction error over the young population of the training set,  $(\mathbf{y}_1, \mathbf{X}_1)$ , is controlled through a performance constraint ( $f_1 = 0.357$ ). In a real application, once the  $L$  groups of interest are selected by the user, threshold values  $f_1, \dots, f_L$  have to be fixed. Note that these thresholds will depend directly on the dataset in question and the considered groups of interest. As a first option, they could be fixed by the user according to her demand, but therefore unfeasibility problems may appear when solving the CSCLasso problem (3). For that reason, in Sect. 2.3 two procedures for determining such threshold values so that (3) is feasible are given.

Next, some other aspects related to the formulation of the CSCLasso and its resolution will be discussed.

## 2.2 Computational details

The CSCLasso problem as defined by (3) is a non-differentiable convex optimization problem with quadratic and convex constraints. However, if we rewrite the non-differentiable term in (3) as

$$\mathbf{A}\boldsymbol{\beta} = \mathbf{u}^+ - \mathbf{u}^-,$$

where  $\mathbf{u}^+ = (u_1^+, \dots, u_p^+)$  and  $\mathbf{u}^- = (u_1^-, \dots, u_p^-)$  are new vectors of positive auxiliary variables, a differentiable version for the CSCLasso problem (3) is obtained in a straightforward manner as

$$\begin{aligned} \min_{\boldsymbol{\beta}, \mathbf{u}^+, \mathbf{u}^-} \quad & \frac{1}{n_0} \|\mathbf{y}_0 - \mathbf{X}_0\boldsymbol{\beta}\|^2 + \lambda \sum_{j=1}^p u_j^+ + \lambda \sum_{j=1}^p u_j^- \\ \text{s.t.} \quad & \frac{1}{n_1} \|\mathbf{y}_1 - \mathbf{X}_1\boldsymbol{\beta}\|^2 - f_1 \leq 0, \\ & \vdots \\ & \frac{1}{n_L} \|\mathbf{y}_L - \mathbf{X}_L\boldsymbol{\beta}\|^2 - f_L \leq 0, \\ & \mathbf{A}\boldsymbol{\beta} = \mathbf{u}^+ - \mathbf{u}^-, \\ & \mathbf{u}^+, \mathbf{u}^- \geq 0. \end{aligned}$$

This previous smooth formulation for the CSCLasso eases its resolution notably, since efficient solvers for quadratically constrained programming problems, such as Gurobi Optimization (2018), are available. In particular, the Gurobi R interface will be used in this work to obtain all numerical results.

Another remark concerning the formulation of the CSCLasso is that, instead of using the sum of squared deviations, least absolute deviations could have been considered. Then, (3) would be reduced to a regression problem under linear inequality constraints, as those described in James et al. (2019), Gaines et al. (2018) and Hu et al. (2015). Nevertheless, to cope the non-differentiability of the absolute value function, a huge number of constraints and new auxiliary variables, which would depend on  $n$ , should have been added. Consequently, these constrained approaches are likely to face severe numerical difficulties in practice for large datasets.

## 2.3 The choice of threshold values

As commented in Sect. 2.1, threshold values  $f_1, \dots, f_L$  could be fixed by the user. If the user is too demanding, imposing very low MSE threshold values for (some of) the different groups, the optimization problem may become unfeasible. Although a try-and-error procedure may be used, it would be very helpful to have strategies yielding feasible solutions. Here we propose two procedures for determining  $f_1, \dots, f_L$  in such a way that (3) is feasible.

First, we propose a choice of the threshold values so that they are close to the OLS results,

$$f_l = (1 + \tau)MSE_l(\hat{\beta}^{ols}), \quad l = 1, \dots, L, \tag{4}$$

where  $MSE_l(\beta) = \frac{1}{n_l} \|y_l - X_l \beta\|^2$ ,  $l = 1, \dots, L$  and  $\tau \geq 0$  is a small parameter whose meaning is the percentage of worsening with respect to the OLS prediction error. For the numerical example in Sect. 1, we could have imposed the threshold for the MSE over *Group 1* equal to 0.352, which is a 10% ( $\tau = 0.1$ ) more than  $MSE_1(\hat{\beta}^{ols}) = 0.319$ . The choice (4) deals with the heterogeneity coming from the variability of the different groups ( $MSE_l$  is different across groups). Nevertheless, when heterogeneity related to the importance of each group is also considered, the parameter  $\tau$  can be replaced in (4) by  $\tau_l$ ,  $l = 1, \dots, L$ .

Next, we shall compute the minimum value of  $\tau$ ,  $\tau_{min}$ , so as to (3) is feasible. That is, the minimum  $\tau$  so that there exists  $\beta^*$  satisfying

$$\left( \max_l \frac{MSE_l(\beta^*)}{MSE_l(\hat{\beta}^{ols})} \right) - 1 \leq \tau,$$

and, therefore,  $\tau_{min}$  will be given as

$$\tau_{min} = \left( \max_l \frac{MSE_l(\beta^*)}{MSE_l(\hat{\beta}^{ols})} \right) - 1.$$

Such  $\tau_{min}$  can be found as the optimal value of the following linear problem with convex and quadratic constraints

$$\begin{aligned} \min_{\beta, z} \quad & z \\ \text{s.t.} \quad & z \geq \frac{MSE_l(\beta)}{MSE_l(\hat{\beta}^{ols})} - 1, \quad \forall l = 1, \dots, L. \end{aligned} \tag{5}$$

The feasible version of the CSCLasso optimization problem can be formulated as

$$\begin{aligned} \min_{\beta} \quad & \frac{1}{n_0} \|y_0 - X_0 \beta\|^2 + \lambda \|A\beta\|_1 \\ \text{s.t.} \quad & \frac{1}{n_1} \|y_1 - X_1 \beta\|^2 - (1 + \tau)MSE_1(\hat{\beta}^{ols}) \leq 0, \\ & \vdots \\ & \frac{1}{n_L} \|y_L - X_L \beta\|^2 - (1 + \tau)MSE_L(\hat{\beta}^{ols}) \leq 0, \end{aligned} \tag{6}$$

where  $\tau \geq \tau_{min}$ .

Finally, note that if  $\tau$  is big enough, then solving (6) is equivalent to solve the unconstrained problem. Indeed, it is possible to find the value of  $\tau$ ,  $\tau_{max}(\lambda)$ , such that both the constrained and unconstrained problems are equivalent

$$\tau_{max}(\lambda) = \max_{l \in \{1, \dots, L\}} \frac{MSE_l(\hat{\beta}^{Lasso}(\lambda))}{MSE_l(\hat{\beta}^{ols})} - 1. \tag{7}$$

A second possible choice for the threshold values follows an analogous approach but, instead of considering the results of the OLS, we shall consider the mean squared error of the Lasso, as in the numerical example introduced in Sect. 1. For each  $l = 1, \dots, L$ ,

$$f_l = (1 - \gamma)MSE_l(\hat{\beta}^{Lasso}(\lambda)), \quad l = 1, \dots, L, \tag{8}$$

where  $\gamma \geq 0$  is related to the desired percentage of improvement over the Lasso solution ( $\gamma = 0.1$  in the numerical example of Sect. 1). In this case, we will compute the maximum value of  $\gamma$ ,  $\gamma_{max}$ , in such a way that (3) is feasible under (8), and the linear problem associated with  $\gamma_{max}$  is

$$\begin{aligned} \max_{\beta, z} \quad & z \\ \text{s.t.} \quad & 1 - \frac{MSE_l(\beta)}{MSE_l(\hat{\beta}^{Lasso}(\lambda))} \geq z, \quad \forall l = 1, \dots, L. \end{aligned} \tag{9}$$

Thus, another possible feasible version of the CSCLasso optimization problem can be formulated as

$$\begin{aligned} \min_{\beta} \quad & \frac{1}{n_0} \|y_0 - X_0\beta\|^2 + \lambda \|A\beta\|_1 \\ \text{s.t.} \quad & \frac{1}{n_1} \|y_1 - X_1\beta\|^2 - (1 - \gamma)MSE_1(\hat{\beta}^{Lasso}(\lambda)) \leq 0, \\ & \vdots \\ & \frac{1}{n_L} \|y_L - X_L\beta\|^2 - (1 - \gamma)MSE_L(\hat{\beta}^{Lasso}(\lambda)) \leq 0, \end{aligned} \tag{10}$$

where  $\gamma \leq \gamma_{max}$ .

Note that the two choices previously described for selecting the threshold values are not unique. Indeed, instead of using the MSE, another statistical measure as the R-squared can be considered. Further details about how they perform in numerical applications are described in Sects. 2.4 and 4.



### 2.4 The role of the tuning parameters

The CSCLasso, as defined by (6) or (10), is stated in terms of two tuning parameters,  $\lambda$  and  $\tau$  or  $\lambda$  and  $\gamma$ , respectively. The first one,  $\lambda$ , is related to the sparsity of the solution, and the second one is linked to the user’s demanding level, since the degree of requirement increases as  $\tau \rightarrow \tau_{min}$  (or  $\gamma \rightarrow \gamma_{max}$ ). In this section we investigate how the solution of the CSCLasso changes when  $\lambda$  and  $\tau$  jointly vary (analogous results are obtained if  $\lambda$  and  $\gamma$  are analyzed instead). With this purpose, consider again the experimental setting as in the example of Sect. 1 related to prostate dataset with  $\mathcal{A} = (0|I^{p \times p})$  (Lasso objective function), but in this occasion assume that the prediction errors of both groups (the young and the elderly people) shall be controlled.

The interval of variation of the parameter  $\lambda$  is set to  $I_\lambda = [0, 30]$ . Moreover, according to (5), the smallest value of  $\tau$  such that the CSCLasso optimization problem (6) is feasible is  $\tau_{min} = 0.055$ . On the other hand, following (7),  $\tau_{max} = \max_{\lambda \in [0, 30]} \tau_{max}(\lambda) = 2.355$ , although we will enlarge the interval of variation of  $\tau$  to also visualize the unconstrained solution; such interval will be finally set as  $I_\tau = [\tau_{min}, \tau_{max} + 2] = [0.055, 4.355]$ . Figure 1 represents, via a heat map, the solution for  $\hat{\beta}_1^{CSCLasso}(\lambda)$  for the different values of  $(\lambda, \tau)$  in a grid contained in  $I_\lambda \times I_\tau$ .

Some conclusions can be drawn from the figure. Consider first the cases where  $\tau$  and  $\lambda$  are big enough. Since, in this case,  $\tau \geq \tau_{max}$ , then, as commented at the end of the previous section, solving (6) is equivalent to solving the Lasso. Therefore,  $\hat{\beta}^{CSCLasso}(\lambda) = \hat{\beta}^{Lasso}(\lambda) = \mathbf{0}$  will be the optimal solution, provided that  $\lambda$  is big enough. Analogously, if  $\tau \geq \tau_{max}$  but  $\lambda$  is small, then  $\hat{\beta}^{CSCLasso}(\lambda) = \hat{\beta}^{Lasso}(\lambda)$ , which will be equal to zero or not depending on the importance of the variable. When  $\tau$  is small, the constraints are demanding and, even for large values of  $\lambda$ , it might happen that  $\hat{\beta}_1^{CSCLasso}(\lambda) \neq \hat{\beta}_1^{Lasso}(\lambda) = 0$ , as it is the case.

Figure 7 (see Appendix for further results) represents the analogous heat maps concerning  $\hat{\beta}_2^{CSCLasso}(\lambda), \dots, \hat{\beta}_8^{CSCLasso}(\lambda)$ . A similar discussion as with  $\hat{\beta}_1^{CSCLasso}(\lambda)$  is applicable to these figures. An interesting remark to be made concerns the importance of each variable: while variable 1 is the only one selected for the Lasso, the CSCLasso returns a less sparse solution in this case, since predictor variables 1, 2, 4 and 5 turn out to be significant. However, this is not the rule, since there are examples where the level of sparsity is higher for the CSCLasso, as will be shown in Sect. 4.

### 3 Theoretical properties

In this section we discuss some theoretical results concerning the CSCLasso model. In Sect. 3.1, the existence of a unique optimal solution to Problem (3) is proven for a fixed value of  $\lambda \geq 0$ . Section 3.2 deals with the limit behavior of the solution when  $\lambda$  approaches infinity. Finally, some consistency properties of the CSCLasso solution are derived in Sect. 3.3 from the Sample Average Approximation theory (see Shapiro et al. 2009).




Fig. 1 Heat map of  $\hat{\beta}_1^{CSCLasso}(\lambda)$  using prostate dataset

### 3.1 Existence and uniqueness of solution

For constrained versions of the Lasso in the literature, as in James et al. (2019), it is not possible to obtain the path of solutions and therefore approximations are made with the use of numerical algorithms. For the CSCLasso problem, a closed form solution of expression  $\hat{\beta}^{CSCLasso}(\lambda)$  is not available. However, an implicit characterization of the CSCLasso solution (with only one constraint) can be found as the following result states.

**Proposition 1** Consider the CSCLasso problem with one constraint,

$$\begin{aligned}
 \min_{\beta} \quad & \frac{1}{n_0} \|\mathbf{y}_0 - \mathbf{X}_0\beta\|^2 + \lambda \|\mathcal{A}\beta\|_1 \\
 \text{s.t.} \quad & \frac{1}{n_1} \|\mathbf{y}_1 - \mathbf{X}_1\beta\|^2 - (1 + \tau)MSE_1(\hat{\beta}^{ols}) \leq 0,
 \end{aligned} \tag{11}$$

where  $\mathcal{A} = (0|I^{p \times p})$  and assume that  $\mathbf{X}_0$  and  $\mathbf{X}_1$  are maximum rank matrices. Then

$$\begin{aligned}
 \hat{\beta}^{CSCLasso}(\lambda) = & \left( \frac{1}{n_0} \mathbf{X}'_0 \mathbf{X}_0 + \frac{1}{n_1} \eta(\lambda) \mathbf{X}'_1 \mathbf{X}_1 \right)^{-1} \left( \frac{1}{n_0} \mathbf{X}'_0 \mathbf{y}_0 + \frac{1}{n_1} \eta(\lambda) \mathbf{X}'_1 \mathbf{y}_1 \right) \\
 & - \frac{1}{2} \left( \frac{1}{n_0} \mathbf{X}'_0 \mathbf{X}_0 + \frac{1}{n_1} \eta(\lambda) \mathbf{X}'_1 \mathbf{X}_1 \right)^{-1} \mathbf{b}(\lambda)
 \end{aligned} \tag{12}$$

where  $\eta(\lambda)$  is the Lagrange multiplier of the constraint and the component  $s$ ,  $s = 0, 1, \dots, p$ , of the vector  $\mathbf{b}(\lambda)$  is given by

$$b_s(\lambda) = \begin{cases} \lambda, & \text{if } \hat{\beta}_s^{CSCLasso}(\lambda) > 0, \\ -\lambda, & \text{if } \hat{\beta}_s^{CSCLasso}(\lambda) < 0, \\ 0 & \text{else.} \end{cases}$$

From the previous proposition, it is clear that a closed form solution is hard to be obtained, even in the simplest scenario.

Nevertheless, given a fixed value of  $\lambda$ , the CSCLasso problem can be solved using quadratic programming via any of the standard solvers available in the literature. As an example, Fig. 2 depicts the path of solutions for the prostate example introduced in Sect. 1, for an assortment of values of  $\lambda$  in a grid (see Sect. 3.2 for details). Each line represents a component of  $\hat{\beta}^{CSCLasso}(\lambda)$ ,  $\hat{\beta}_j^{CSCLasso}(\lambda)$  with  $j = 1, \dots, 8$ . It can be observed from the figure that, contrary to what happens in the Lasso path of solutions (top panel of Fig. 2), the CSCLasso path of solutions is not piecewise linear (bottom panel of Fig. 2). Such non-linearity (due to the quadratic constraints) hinders the application of an iterative algorithm to obtain the path of solutions as those given in papers James et al. (2019) and Gaines et al. (2018).

Also note from Fig. 3 that as a consequence of the performance constraints, the solution is stabilized when  $\lambda$  increases, but does not shrink to 0, as with Lasso. This is detailed in Sect. 3.2.

Even without having the expression of the general solution of (3), we next prove that, under full rank assumptions, the solution is unique. First, in order to simplify the formulation of (3), henceforth its feasible set will be denoted by  $\mathbf{B}$ , which is convex and closed. This is also true (and the results which follow remain valid) if, on top of the performance constraints, one adds linear constraints modeling other aspects of interest (for example, the sign of a certain coefficient can be fixed to be positive or negative depending on the known relation between the corresponding predictors with the response variable). In the same vein, henceforth,  $(\mathbf{y}_0, \mathbf{X}_0) = (\mathbf{y}, \mathbf{X})$  is used to minimize the overall MSE. In this way, the CSCLasso problem (3) is rewritten as


$$\min_{\beta \in \mathbf{B}} \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|^2 + \lambda \|\mathcal{A}\beta\|_1. \tag{13}$$

The following result guarantees that the solution of problem (13) is unique.

**Theorem 1** Consider Problem (13) where  $\mathbf{X}$  is assumed to be a maximum rank matrix and its feasible region  $\mathbf{B}$  is a convex and closed set in  $\mathbb{R}^{p+1}$ . Then, Problem (13) has a unique optimal solution.

### 3.2 Asymptotic behaviour

One of the key points when dealing with Lasso-type problems is the choice of the regularization parameter  $\lambda$ . In the case of the Lasso, such choice is straightforward since the entire path of solutions is known to be piecewise linear, shrinking to 0.



**Fig. 2** Path of solutions under Lasso (top) and CSCLasso (bottom) for prostate dataset

In particular, it is known that there exists a value of  $\lambda$ ,  $\lambda^*$ , such that the solution  $\hat{\beta}^{Lasso}(\lambda) = \mathbf{0}$  is optimal for all values  $\lambda \geq \lambda^*$ .

The following result provides explicitly the value of such  $\lambda^*$ .




Fig. 3 Path of solutions under Lasso (top) and CSCLasso (bottom) for prostate dataset when  $\lambda$  increases

**Proposition 2** Consider the Lasso model (2). Define  $\lambda^*$  as the optimal value of the linear programming problem

$$\begin{aligned}
 \min_{z, \mathbf{t}} \quad & z \\
 \text{s.t.} \quad & \frac{2}{n} \mathbf{X}'\mathbf{y} = \mathcal{A}'\lambda\mathbf{t}, \\
 & -z \leq \lambda t_s \leq z, \quad s = 0, 1, \dots, p.
 \end{aligned}$$

Then,  $\hat{\beta}^{Lasso}(\lambda) = \mathbf{0}$  for all  $\lambda \geq \lambda^*$ . In particular, for  $\mathcal{A} = (0|I^{p \times p})$ ,

$$\lambda^* = \left\| \frac{2}{n} \mathbf{X}' \mathbf{y} \right\|_{\infty}$$

Some works dealing with (linear) constrained versions of the Lasso (as Gaines et al. 2018; James et al. 2019) have developed efficient algorithms to build the associated solutions path. Consider now the general problem (13). As commented in Sect. 3.1, the expression of  $\hat{\beta}^{CSCLasso}(\lambda)$  is not available in closed form and, consequently, the entire path cannot be computed. In this case, when  $\lambda$  tends to  $+\infty$ , the solution  $\hat{\beta}^{CSCLasso}(\lambda)$  stabilizes around  $\hat{\beta}^{CSCLasso}(+\infty) = \arg \min_{\beta \in B} \|\mathcal{A}\beta\|_1$ . This idea is also used in Gaines et al. (2018) and James et al. (2019), where, in order to find an initialization for the algorithms, the proposed constrained problems are solved by only considering the penalty term in the objective function. Such a limit solution is obtained by solving an optimization problem with a linear objective function and convex quadratic constraints, namely

$$\begin{aligned} \min_{\beta, u^+, u^-} \quad & \sum_{s=0}^p (u_s^+ + u_s^-) \\ \text{s.t.} \quad & \beta \in B \\ & \mathcal{A}\beta = u^+ - u^- \\ & u^+, u^- \geq \mathbf{0}. \end{aligned}$$

A grid search is carried out in the general CSCLasso problem (13) to obtain suitable values of  $\lambda$ . In order to fix the grid, we propose the following dynamic approach to find an approximate maximum value of  $\lambda$ ,  $\lambda^*$  (see Algorithm 1).

---

**Algorithm 1:** Dynamic approach for selecting  $\lambda^*$  in the CSCLasso

---

1. Fix  $\varepsilon > 0$  and  $c = (2^{-5})$ . Fix  $i = 1$  and compute  $\hat{\beta}^{CSCLasso}(c[i])$ .
  2. Compute  $\hat{\beta}^{CSCLasso}(+\infty) = \arg \min_{\beta \in B} \|\mathcal{A}\beta\|_1$ .
  3. While  $\|\hat{\beta}^{CSCLasso}(\lambda) - \hat{\beta}^{CSCLasso}(+\infty)\| > \varepsilon$ , repeat
    - a)  $i = i + 1$
    - b)  $c = (c, 2c[i - 1])$
    - c) compute  $\hat{\beta}^{CSCLasso}(c[i])$
  4.  $\lambda^* = c[i]$
-

Once the maximum value  $\lambda^*$  is found, then the grid ranges from 0 to  $\lambda^*$  with the desired step. Note that the previous algorithm already provides an initial grid of the form  $(2^{-5}, 2^{-4}, \dots, 2^0, \dots, \lambda^*)$ .

### 3.3 Consistency properties in the CSCLasso

The purpose of this section is to prove some results related to the consistency of both the solution and the objective value for CSCLasso problem (13). To do that, the theory of *Sample Average Approximation* (SAA) (Shapiro et al. 2009) will be applied. Consider the following stochastic programming problem

$$\min_{\beta \in \mathbf{B}} f(\beta) := E[F(\beta, (Y, X))], \tag{14}$$

where  $\mathbf{B}$  is a nonempty closed subset of  $\mathbb{R}^{p+1}$ ,  $(Y, X)$  is an absolutely continuous random vector whose probability distribution  $P$  is supported on a set  $\Xi \subset \mathbb{R}^{p+1}$  and  $F : \mathbf{B} \times \Xi \rightarrow \mathbb{R}$ . In Shapiro et al. (2009), under some conditions, the *true* problem (14) can be estimated by the SAA:

$$\min_{\beta \in \mathbf{B}} \hat{f}_n(\beta) := \frac{1}{n} \sum_{i=1}^n F(\beta, (y_i, x_i)), \tag{15}$$

where  $x_i = (1, x_{i1}, \dots, x_{ip})'$ , and  $\{(y_i, x_i)\}_{i=1, \dots, n}$  is a realization of the  $n$  random vectors  $\{(Y_i, X_i)\}_{i=1, \dots, n}$ , which are independent and identically distributed (i.i.d.) as the random vector  $(Y, X)$ . Note that the CSCLasso problem as in (13) takes the form of (15) as

$$\min_{\beta \in \mathbf{B}} \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \beta)^2 + \lambda \|A\beta\|_1 \tag{16}$$

and the *true* CSCLasso problem equivalent to (14) is

$$\min_{\beta \in \mathbf{B}} E[(Y - X' \beta)^2 + \lambda \|A\beta\|_1]. \tag{17}$$

Before proving the main result on the consistency of the CSCLasso, we first show the uniqueness of the solution of such a problem.

**Proposition 3** *The optimal solution of the true CSCLasso problem (17) is unique.*

Denote by  $v^{CSCLasso}(\lambda)$  and  $\beta^{CSCLasso}(\lambda)$ , respectively, the optimal value and the optimal solution of problem (17). Analogously, let  $\hat{v}^{CSCLasso}(\lambda)$  and  $\hat{\beta}^{CSCLasso}(\lambda)$  be the optimal value and the optimal solution, respectively, of the SAA CSCLasso problem (16). The following result shows the consistency of the SAA values to the *true* values.

**Theorem 2** Assume that  $E[\|X\|^2] < \infty$ ,  $E[Y^2] < \infty$ ,  $E[\|YX\|] < \infty$ . Then,  $\hat{v}^{CSCLasso}(\lambda)$  converges to  $v^{CSCLasso}(\lambda)$  and  $\hat{\beta}^{CSCLasso}(\lambda)$  converges to  $\beta^{CSCLasso}(\lambda)$  with probability one (w.p. 1).

Finally, note that the theoretical results that have been studied in this work are also applicable to other versions of constrained Lasso as long as the feasible set is convex and closed, as is the case with the above-mentioned works (James et al. 2019; Gaines et al. 2018; Hu et al. 2015).

## 4 Numerical experiments

In this section, the behaviour and performance of our approach is illustrated throughout an extensive empirical study. In particular, using both simulated and real datasets, the aim of the experiments shall be to improve the prediction errors of the Lasso in one or more groups of interest. Or in other words, threshold values shall be fixed as in (10). Since our proposal is a novel extension of the Lasso, we will also show the results under the Lasso, not only for those groups that are controlled (for which obviously, the Lasso performs worse) but also for the non-controlled groups. In this way the CSCLasso can be better inspected in comparison to the Lasso. Other aspects as the overall MSE and the percentage of non-zero coefficients in the regression model, among others, will be explored. Such measures will be estimated through median values using a fivefold cross validation approach. To this end, the dataset will be split at each fold into three sets: the so-called training, validation and test sets. The training set is used to fit the model, the validation set is used to estimate prediction error for model selection and the test set is used for assessment of the generalization error of the final chosen model.

### 4.1 A simulation study

The generation of the synthetic datasets in this section follows that of Ollier and Viallon (2017), where an overparameterized regression model is considered to cope with stratified data. A number of groups  $K = 20$  is set and two different sample sizes per group are considered,  $n_k = \{150, 500\}$ , for  $k = 1, \dots, K$ . The number of predictors  $p$  will be chosen from  $\{20, 100, 500\}$ . The matrix of predictor values  $\mathbf{X}$  is generated according to a multivariate normal distribution with zero mean and covariance matrix  $\Sigma$  being a Toeplitz matrix with element  $(i, j)$  equal to  $0.5^{|i-j|}$ . Regarding the response vector, a set of 20 predictors are randomly selected (with indexes included in a set  $P_0$ ), while the rest of predictors are noise (that is,  $\beta_j = 0$  for  $j \notin P_0$ ). The coefficients of the significant 20 predictors are chosen as follows. First, consider 10 random predictors out of the 20 selected. For such predictors, if the group  $k > 6$  then  $\beta_j = 1$  and, otherwise,  $\beta_j = 1 + K^{\frac{1}{2}}$ . For the other 10 predictors,  $\beta_j = 1$  if  $k \leq 6$  and  $\beta_j = 1 + K^{\frac{1}{2}}$  otherwise. In this way, the predictors behave differently depending on the group. Finally, the response vector for each group is generated according to the standard linear regression model with normal error.



Once the synthetic dataset is built and its response and predictor variables have been standardized, the CSCLasso with  $\mathcal{A} = (0|I^{p \times p})$  is run with constraints imposed over the first six groups. The choice of  $\lambda$  will change at each fold. A grid in  $\lambda$  is built as in Sect. 3.2, and, the value of  $\lambda$  which leads to the lower overall MSE in the validation set is selected. Table 2 shows the median prediction errors per group  $k$  ( $MSE_k$ ),  $k = 1, \dots, 6$ , obtained by the Lasso (rows in bold values) and the corresponding values obtained under the CSCLasso for different threshold values  $f$ .

In particular, the values of  $f$  have been set as improvement percentages over the Lasso values, where the improvement levels are 3%, 5%, 7%, 10%, 15% and 20% ( $\gamma$  equal to 0.03, 0.05, 0.07, 0.15 and 0.20, respectively). The results are obtained for different combinations  $(p, n_k)$ , where, as commented before,  $p$  is chosen from the set  $\{20, 100, 500\}$  and  $n_k$  from  $\{150, 500\}$ . For example, if  $n_k = 150$  and  $p = 20$ , the median of the mean squared error for the Lasso in *Group 1* was equal to 1.084. If the goal is to achieve an improvement of 3%,  $f$  must be chosen equal to 1.052. The median of the mean squared errors for the CSCLasso is equal to 0.933 in this case (results obtained on the test sample). It is important to remark that, for some levels of improvement, the CSCLasso problem is unfeasible due to the fact that  $\gamma_{max}$  for such datasets is smaller than the required  $\gamma$ , and, therefore, those cases are represented as empty spaces in the table. It must be noted that  $\gamma_{max}$  will also depend on each fold in the cross-validation because it is associated with the partition of the data, since the  $MSE_l$  in (9) depends on such partition ( $l = 1, \dots, L$ ). It is also worth mentioning that the constraints will always be satisfied on the training set but not necessarily on the test set, see, for example, the case  $k = 3$ ,  $n_k = 150$ ,  $p = 100$  with improvement level equal to 15%. This phenomenon is particularly common as  $p$  increases and  $n_k$  decreases (see for example the values corresponding to  $p = 500$  and  $n_k = 150$  in Table 2).

We next investigate how the improvement in the prediction errors of the groups of interest affects the prediction errors in the rest of the groups, the overall prediction error and the sparsity level. Figure 4 represents the percentage of non-zero (NZ) coefficients and the overall prediction error for different sample sizes, different levels of improvement and for  $p = 100$ . Lasso results are also included. For instance, when  $n_k = 300$  (black squares in Fig. 4), the NZ percentage for Lasso is 39.60 with an associated overall MSE of 0.734; whereas running the CSCLasso demanding a 3% of improvement over the first six groups, we achieve a NZ percentage of 38.61, and an overall MSE of 0.735. In general terms, it can be seen that the sparsity of the solution decreases with the improvement level: smaller squares, which represent smaller imposed improvement percentages, are on the left of bigger squares (which are associated with demanding percentage of improvement). Then, if the user is very demanding in predicting a specific group, this implies, in the majority of the cases, a less sparse solution. Notwithstanding, when no level of improvement is imposed (Lasso problem), the solution can be less sparse than in the CSCLasso, as in the case of  $n_k = 300$ . This also occurs when  $p = 500$  and  $n_k = 150$  (see the bottom graphic of Fig. 10 in Appendix: further results). Furthermore, the overall prediction error slightly worsens with the improvement level, due to the worsening of the predictions in the uncontrolled groups.

**Table 2** Median errors over test sets for synthetic datasets

$p$		$n_k = 150$						$n_k = 300$					
		Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
20	$MSE_k(Lasso)$	<b>1.084</b>	<b>0.904</b>	<b>0.768</b>	<b>0.925</b>	<b>0.944</b>	<b>0.674</b>	<b>0.801</b>	<b>0.770</b>	<b>0.966</b>	<b>0.937</b>	<b>0.776</b>	<b>0.938</b>
	<i>Improv. f</i>	1.052	0.877	0.745	0.898	0.916	0.653	0.777	0.747	0.937	0.909	0.753	0.910
	3% $MSE_k$	0.933	0.722	0.695	0.859	0.822	0.574	0.752	0.689	0.810	0.859	0.744	0.870
	<i>Improv. f</i>	1.030	0.859	0.730	0.879	0.897	0.640	0.761	0.732	0.918	0.890	0.737	0.891
	5% $MSE_k$	0.921	0.708	0.689	0.848	0.804	0.564	0.734	0.682	0.787	0.844	0.736	0.859
	<i>Improv. f</i>	-	-	-	-	-	-	0.745	0.717	0.899	0.871	0.722	0.873
	7%	-	-	-	-	-	-	0.717	0.677	0.768	0.834	0.721	0.847
	<i>Improv. f</i>	-	-	-	-	-	-	-	-	-	-	-	-
	10%	-	-	-	-	-	-	-	-	-	-	-	-
	<i>Improv. f</i>	-	-	-	-	-	-	-	-	-	-	-	-
100	$MSE_k(Lasso)$	<b>1.491</b>	<b>1.139</b>	<b>0.875</b>	<b>1.496</b>	<b>1.151</b>	<b>0.759</b>	<b>1.104</b>	<b>1.151</b>	<b>0.932</b>	<b>0.981</b>	<b>1.311</b>	<b>1.142</b>
	<i>Improv. f</i>	1.447	1.105	0.848	1.451	1.116	0.736	1.070	1.117	0.904	0.951	1.272	1.108
	3% $MSE_k$	1.234	0.948	0.781	1.322	0.945	0.708	1.044	1.145	0.911	0.963	1.306	1.139
	<i>Improv. f</i>	1.417	1.082	0.831	1.421	1.093	0.721	1.048	1.094	0.885	0.932	1.245	1.085
	5% $MSE_k$	1.226	0.941	0.777	1.317	0.937	0.706	1.025	1.123	0.897	0.947	1.311	1.142
	<i>Improv. f</i>	1.387	1.059	0.813	1.391	1.070	0.706	1.026	1.071	0.866	0.912	1.219	1.062
	7% $MSE_k$	1.219	0.933	0.773	1.311	0.929	0.704	1.005	1.111	0.883	0.934	1.309	1.141
	<i>Improv. f</i>	1.342	1.025	0.787	1.346	1.035	0.683	0.993	1.036	0.838	0.883	1.180	1.028
	10% $MSE_k$	1.207	0.921	0.767	1.303	0.917	0.702	0.978	1.094	0.864	0.915	1.311	1.131
	<i>Improv. f</i>	1.268	0.968	0.743	1.271	0.978	0.645	0.938	0.979	0.792	0.834	1.114	0.971

Table 2 continued

$p$		$n_k = 150$						$n_k = 300$					
		Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
500	15% $MSE_k$	1.128	0.903	0.765	1.250	0.919	0.695	0.936	1.061	0.833	0.882	1.274	1.087
	Improv. $f$	1.193	0.911	0.700	1.196	0.920	0.607	0.883	0.921	0.745	0.785	1.049	0.914
	20% $MSE_k$	1.155	0.909	0.771	1.251	0.900	0.700	0.895	1.036	0.803	0.855	1.227	1.046
	$MSE_k(Lasso)$	<b>1.306</b>	<b>1.133</b>	<b>1.318</b>	<b>1.261</b>	<b>1.246</b>	<b>1.473</b>	<b>1.151</b>	<b>1.204</b>	<b>1.171</b>	<b>1.148</b>	<b>1.278</b>	<b>1.068</b>
	Improv. $f$	1.267	1.099	1.279	1.223	1.208	1.429	1.116	1.168	1.136	1.114	1.240	1.035
	3% $MSE_k$	1.270	1.133	1.314	1.248	1.246	1.453	1.029	1.077	1.047	1.022	1.186	0.961
	Improv. $f$	1.241	1.076	1.252	1.198	1.183	1.400	1.093	1.144	1.113	1.091	1.214	1.014
	5% $MSE_k$	1.257	1.133	1.308	1.245	1.245	1.437	1.025	1.060	1.032	1.004	1.165	0.945
	Improv. $f$	1.215	1.053	1.226	1.173	1.158	1.370	1.070	1.119	1.089	1.068	1.189	0.993
	7% $MSE_k$	1.246	1.133	1.312	1.241	1.246	1.425	1.018	1.042	1.023	0.987	1.144	0.929
	Improv. $f$	1.176	1.019	1.186	1.135	1.121	1.326	1.036	1.083	1.054	1.033	1.150	0.961
	10% $MSE_k$	1.231	1.132	1.313	1.230	1.246	1.401	1.005	1.016	1.006	0.961	1.114	0.905
Improv. $f$	1.110	0.963	1.120	1.072	1.059	1.252	0.978	1.023	0.995	0.976	1.086	0.907	
15% $MSE_k$	1.207	1.114	1.271	1.193	1.241	1.376	0.971	0.977	1.039	0.923	1.060	0.866	
Improv. $f$	1.045	0.906	1.054	1.009	0.996	1.179	0.921	0.963	0.937	0.919	1.022	0.854	
20% $MSE_k$	1.184	1.081	1.266	1.150	1.230	1.354	0.952	0.972	0.995	0.908	1.057	0.857	

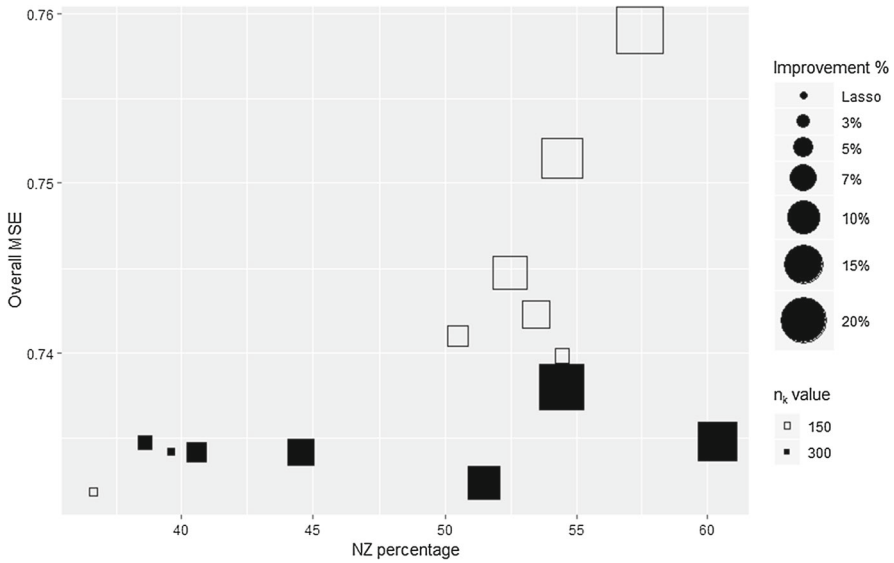


Fig. 4 Median overall MSE over the test sets and NZ percentage under the choice  $p = 100$


Figure 5 represents the prediction errors over the groups that are not controlled as well as the overall mean squared error, for different improvement levels,  $n_k = 150$  (top figure) and  $n_k = 300$  (bottom) when  $p = 100$ . In the figure, Lasso values are also shown. From the figure, it can be concluded that the Lasso performs better in the uncontrolled groups, since the prediction errors worsen under our proposal. However, the overall mean squared error remains almost constant (since the improved errors compensate the more deteriorated ones). Similar conclusions can be drawn under the choices  $p = 20$  and  $p = 500$  (see Figs. 8, 9 in Appendix: further results, respectively).

Next, we test how the solution of the CSCLasso behaves with respect to other global performance measures as the  $l_2$  distance, false positive and negative rates, which are defined not in terms of prediction errors, but on the correct fitting of the generator process (see Yu and Liu 2016). In particular, the  $l_2$  distance is defined as  $\|\hat{\beta} - \beta\|_2$ , where  $\beta$  is the vector of coefficients that generated the datasets (described at the beginning of this section), and  $\hat{\beta}(\lambda)$  are the estimators. In addition, the false positive rate (FPR) and false negative rate (FNR) are calculated as follows:

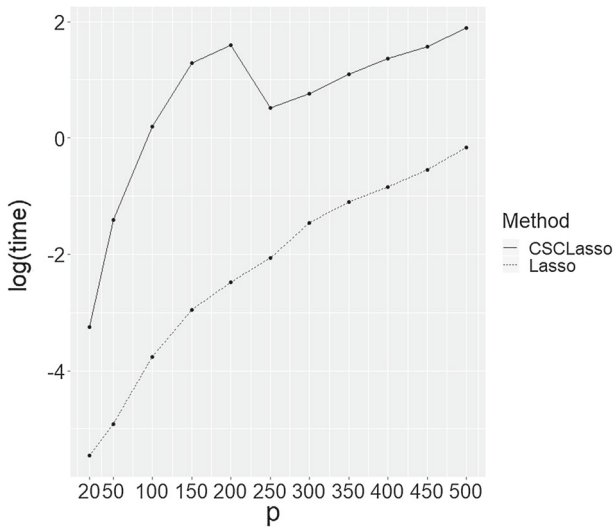
$$FPR = \frac{|j : \beta_j = 0 \ \& \ \hat{\beta}_j(\lambda) = 0|}{|j : \beta_j = 0|},$$

$$FNR = \frac{|j : \beta_j \neq 0 \ \& \ \hat{\beta}_j(\lambda) = 0|}{|j : \beta_j \neq 0|},$$

where  $j = 1, \dots, p$ . The median of these three measures as well as the median of the overall MSE (already shown in Figs. 4, 5 and 8, 9 and 10 in Appendix: further results), are presented in Table 3. For the choices where  $p = 20$ , the FPR values



**Fig. 5** Median  $MSE_k$  over the test sets for  $k = 7, \dots, 20$  under  $p = 100$  features and,  $n_k = 150$  (top) and  $n_k = 300$  (bottom). Each subgraph represents one group and the Y-axis shows the different percentages of improvement



**Fig. 6** A two-dimensional graph of the logarithm of the user times in seconds for  $n_k = 300$  as  $p$  increases

are not given since all the predictors have associated non-zero coefficients when the datasets were created. From this table it can be deduced that similar or even better results, comparing with those of the Lasso, are obtained in the majority of the cases across the four different measures.

A final remark concerning the computational cost of the CSCLasso when comparing with Lasso is as follows. The median user time required to solve the problem with the largest dataset considered in this study ( $n_k = 300$  and  $p = 500$ ) is 0.85 s when the Lasso is run on Intel(R) Core(TM) i7-7500U CPU at 2.70 GHz 2.90 GHz with 8.0 GB of RAM; whereas the CSCLasso requires 6.60 s. Nevertheless, to better understand how the computation time behaves depending on  $p$  value, a grid in this parameter has been inspected, while  $n_k$  is set to 300. Figure 6 depicts the logarithm of the user times in seconds obtained under Lasso and CSCLasso models when  $n_k = 300$  and  $p$  changes. Then, under a reasonable computational cost, the desired results are achieved. Further analyses regarding the computational times are shown in the Appendix (Fig. 11).

## 4.2 Leukemia dataset: a gene expression dataset

The real stratified dataset described in Kouno et al. (2013) is explored here. The data contain information related to myeloid monocytic leukemia cells undergoing differentiation to macrophages. In particular, the dataset is formed by expression levels of 45 transcription factors (response and predictor variables) measured at 8 distinct times (groups) of the differentiation process. As in Ollier and Viallon (2017), the aim is to predict the EGR2 transcription factor in terms of the other  $p = 44$  factors. The sample size per group is equal to 120. Similarly as in Sect. 4.1, the Lasso was run and the overall prediction errors, individual prediction errors per group and percentage of non-zero coefficients are recorded. The records in *Group 1* yield the best MSE

using Lasso model. Therefore, we may be interested in obtaining an even better fitting for such data. The CSCLasso problem is solved with threshold values smaller than the Lasso error, which turns out to be 0.370. Table 4 shows the obtained median results for an assortment of improvement levels, namely, 5%, 7%, 10% and 15% or, equivalently,  $\gamma$  is equal to 0.05, 0.05, 0.07 and 0.15, respectively. From that table, it can be seen how the prediction error of interest (corresponding to *Group 1*) decreases with the improvement level, as expected. Similarly as in Sect. 4.1, the overall mean squared error does not exhibit significant changes, while the prediction errors in the uncontrolled groups do not exhibit the same behaviour. Some of them slightly improve (*Group 6*), others slightly worsen (as the *Group 5* and *Group 8*) and others remain constant (as *Group 2*). Finally, in regards to the sparsity of the solution, for this dataset, less sparse solutions are obtained by CSCLasso in comparison with the Lasso ones.

### 4.3 Communities and crime dataset

In this section, a real dataset from the UCI Machine Learning Repository (Lichman 2013) will be analyzed. In particular, the so-called *Communities and Crime Unnormalized Data Set* shall be considered. The dataset is about communities within the United States and has already been inspected in the literature (see Redmond and Baveja 2002). This dataset combines crime information from the FBI databases (U.S. Department 1995) as well as socio-economic and law enforcement data from U.S. Department (1992) and U.S. Department (1992), respectively. The dataset is formed by  $p = 124$  predictors, 23 of which present missing values, and  $n = 2215$  instances, where the response variable measures the number of murders per 100K population. The predictor variables with missing values are not considered for the next experiments. As such, we finally consider  $p = 101$  predictors. Additionally, for each instance (community), the region from which it comes is known. Thus, if we were interested in obtaining a good prediction in a certain region, say Midwest, we could control these communities by including a performance constraint. Table 5 shows the median errors over the test set for *Group 1*, formed by the communities of Midwest, and over the rest of communities (*Group 2*). In terms of overall MSE and MSE over the two groups, similar conclusions as in Sect. 4.2 are drawn. Whereas different improvement levels are imposed, the MSE of interest ( $MSE_1$ ) is getting smaller but the overall prediction error is almost not affected by the constraint. Lastly, regarding the sparsity of the solution, an analogous behaviour as that observed in the case of simulated data is obtained: the solution becomes less sparse with the improvement level.

As previously commented, the groups of interest may overlap. As an illustration, assume that the interest is in controlling the prediction error in communities of Midwest or communities with a population density larger than or equal to the 75th percentile. Let *Group 1* denote the communities from Midwest, while *Group 2* represents the communities where the density of population is higher than the 75th percentile. For instance, if we aim to improve in a 7% the errors obtained by the Lasso model (equal to 0.513 and 0.442), then the CSCLasso results become 0.475 and 0.441, respectively.

**Table 3** Median performance measures over test sets for synthetic datasets

$p$	$n_k = 150$					$n_k = 300$				
	Overall MSE	$l_2$ distance	FPR	FNR	FNR	Overall MSE	$l_2$ distance	FPR	FNR	FNR
20	<b>Lasso</b>	<b>0.667</b>	<b>4.149</b>	–	<b>0.050</b>	<b>0.666</b>	<b>4.146</b>	–	<b>0.000</b>	<b>0.000</b>
	<i>Improv.</i> 3%	0.691	4.108	–	0.000	0.673	4.099	–	0.000	0.000
	<i>Improv.</i> 5%	0.695	4.104	–	0.000	0.683	4.089	–	0.000	0.000
	<i>Improv.</i> 7%	–	–	–	–	0.693	4.077	–	–	0.000
	<i>Improv.</i> 10%	–	–	–	–	–	–	–	–	–
	<i>Improv.</i> 15%	–	–	–	–	–	–	–	–	–
	<i>Improv.</i> 20%	–	–	–	–	–	–	–	–	–
	<b>Lasso</b>	<b>0.732</b>	<b>2.931</b>	<b>0.275</b>	<b>0.250</b>	<b>0.734</b>	<b>2.729</b>	<b>0.250</b>	<b>0.150</b>	<b>0.150</b>
100	<i>Improv.</i> 3%	0.740	2.857	0.388	0.100	0.735	2.721	0.163	0.150	0.150
	<i>Improv.</i> 5%	0.741	2.855	0.375	0.100	0.734	2.721	0.163	0.100	0.100
	<i>Improv.</i> 7%	0.742	2.853	0.400	0.100	0.734	2.718	0.163	0.100	0.100
	<i>Improv.</i> 10%	0.745	2.849	0.400	0.100	0.732	2.704	0.275	0.100	0.100
	<i>Improv.</i> 15%	0.751	2.844	0.425	0.100	0.735	2.678	0.288	0.050	0.050
	<i>Improv.</i> 20%	0.759	2.835	0.463	0.100	0.738	2.655	0.288	0.100	0.100
	<b>Lasso</b>	<b>0.772</b>	<b>2.778</b>	<b>0.135</b>	<b>0.300</b>	<b>0.744</b>	<b>2.750</b>	<b>0.065</b>	<b>0.300</b>	<b>0.300</b>
	<i>Improv.</i> 3%	0.772	2.779	0.040	0.350	0.744	2.658	0.063	0.250	0.250
500	<i>Improv.</i> 5%	0.772	2.775	0.050	0.300	0.745	2.653	0.075	0.250	0.250
	<i>Improv.</i> 7%	0.772	2.769	0.042	0.350	0.746	2.647	0.077	0.250	0.250
	<i>Improv.</i> 10%	0.772	2.760	0.056	0.350	0.748	2.632	0.104	0.250	0.250
	<i>Improv.</i> 15%	0.771	2.754	0.063	0.300	0.752	2.640	0.129	0.250	0.250
	<i>Improv.</i> 20%	0.774	2.735	0.069	0.350	0.760	2.607	0.148	0.200	0.200



**Table 4** Median errors over test set for gene expression dataset. Constraints imposed over Group 1

	$f_1$	Overall MSE	MSE <sub>1</sub>	MSE <sub>2</sub>	MSE <sub>3</sub>	MSE <sub>4</sub>	MSE <sub>5</sub>	MSE <sub>6</sub>	MSE <sub>7</sub>	MSE <sub>8</sub>	NZ
<i>Lasso</i>	–	<b>0.620</b>	<b>0.370</b>	<b>0.417</b>	<b>0.902</b>	<b>0.496</b>	<b>0.480</b>	<b>0.535</b>	<b>0.496</b>	<b>0.685</b>	<b>53.33</b>
<i>Improv. 5%</i>	<b>0.352</b>	0.623	<b>0.357</b>	0.417	0.918	0.500	0.555	0.512	0.497	0.719	73.33
<i>Improv. 7%</i>	<b>0.344</b>	0.626	<b>0.348</b>	0.418	0.915	0.504	0.565	0.514	0.497	0.722	66.67
<i>Improv. 10%</i>	<b>0.333</b>	0.630	<b>0.335</b>	0.418	0.911	0.510	0.574	0.513	0.498	0.728	66.67
<i>Improv. 15%</i>	<b>0.315</b>	0.636	<b>0.331</b>	0.415	0.903	0.523	0.591	0.512	0.502	0.737	73.33

**Table 5** Median errors over test set for communities and crime dataset. Constraints imposed over *Group 1*

	$f_1$	Overall MSE	$MSE_1$	$MSE_2$	NZ
<i>Lasso</i>	–	<b>0.488</b>	<b>0.433</b>	<b>0.453</b>	<b>21.57</b>
<i>Improv. 5%</i>	<b>0.411</b>	0.488	<b>0.422</b>	0.453	25.49
<i>Improv. 7%</i>	<b>0.403</b>	0.487	<b>0.420</b>	0.453	28.43
<i>Improv. 10%</i>	<b>0.390</b>	0.488	<b>0.416</b>	0.453	26.47
<i>Improv. 15%</i>	<b>0.368</b>	0.486	<b>0.403</b>	0.459	34.31

## 5 Conclusions and extensions

In this paper a new version of the Lasso regression model that strives to control the performance rates associated with individuals of interest is proposed. The method has a significant application in the context of heterogeneous data, where it is common that certain sources are more reliable than others, or simply the prediction on some groups of data are of higher interest, and thus a better fit is sought for some data. In order to control the individuals of interest, performance constraints are included in the regression model. This approach leads to a novel method (CSCLasso) which is not reported in the literature previously, up to our knowledge. Theoretical results concerning this novel methodology have been discussed and, in addition, the CSCLasso has been tested on six synthetic datasets with different properties, on a well-referenced real stratified biomedical dataset and on a real social sciences dataset. The numerical section shows that, with a low computational cost, the accuracy prediction errors for the groups of interest are controlled. This is done at the expense of reducing sparsity (if the regularization parameter is kept fixed) or the overall accuracy.

A number of extensions to this work are possible. Regarding the selection of the threshold values  $f_1, \dots, f_L$ , as commented in Sect. 2.4, a more flexible choice would be to consider a different fine-tuning parameter  $\tau$  (or  $\gamma$ ) for each group of interest  $l$ , say  $\tau_l$  ( $\gamma_l$ ). However, this generalization implies tuning (many) more parameters, making the model less usable by users. In addition, for the sake of dealing with strongly correlated predictors, it may be of interest to change the objective function by that of the elastic net (Zou and Hastie 2005). Another non-straightforward extension could be to address classification problems (via the logistic regression) instead of regression problems. In this case, we would not address a quadratic problem. Work in these issues is underway.

**Acknowledgements** Research partially supported by research grants and projects MTM2015-65915-R (Ministerio de Economía y Competitividad, Spain), FQM-329 and P18-FR-2369 (Junta de Andalucía, Spain), Fundación BBVA and EC H2020 MSCA RISE NeEDS Project (Grant agreement ID: 822214). In addition, we would like to thank the associated editor and two anonymous reviewers for carefully reading this work, and for their insightful comments, which have helped to improve the quality of this paper.

## Appendix: Proofs

### Proof of Proposition 1

Given  $\lambda \geq 0$ , consider problem (11). If  $\beta = \beta^+ - \beta^-$  with  $\beta^+ \geq 0$  and  $\beta^- \geq 0$  and  $\lambda = (0, \lambda, \dots, \lambda)'$ , a vector whose length is  $p + 1$ , then the differentiable version of that problem turns out to be

$$\begin{aligned} \min_{\beta^+, \beta^-} \quad & \frac{1}{n_0} \|y_0 - \mathbf{X}_0(\beta^+ - \beta^-)\|^2 + \lambda' \beta^+ + \lambda' \beta^- \\ \text{s.t.} \quad & \frac{1}{n_1} \|y_1 - \mathbf{X}_1(\beta^+ - \beta^-)\|^2 - (1 + \tau)MSE_1(\hat{\beta}^{ols}) \leq 0, \\ & \beta^+ \geq \mathbf{0} \Leftrightarrow -\beta^+ \leq \mathbf{0}, \\ & \beta^- \geq \mathbf{0} \Leftrightarrow -\beta^- \leq \mathbf{0}. \end{aligned}$$

From the Karush–Kuhn–Tucker conditions,

$$\begin{aligned} \mathcal{L}(\beta^+, \beta^-, \theta^+, \theta^-, \eta) &= \frac{1}{n_0} \|y_0 - \mathbf{X}_0(\beta^+ - \beta^-)\|^2 + \lambda' \beta^+ + \lambda' \beta^- - (\theta^+)' \beta^+ - (\theta^-)' \beta^- \\ &+ \eta \left( \frac{1}{n_1} \|y_1 - \mathbf{X}_1(\beta^+ - \beta^-)\|^2 - (1 + \tau)MSE_1(\hat{\beta}^{ols}) \right) \\ \frac{\partial}{\partial \beta^+} : & -\frac{2}{n_0} \mathbf{X}'_0(y_0 - \mathbf{X}_0(\beta^+ - \beta^-)) + \lambda - \theta^+ - \frac{2}{n_1} \eta \mathbf{X}'_1(y_1 - \mathbf{X}_1(\beta^+ - \beta^-)) = 0 \\ \frac{\partial}{\partial \beta^-} : & \frac{2}{n_0} \mathbf{X}'_0(y_0 - \mathbf{X}_0(\beta^+ - \beta^-)) + \lambda - \theta^- + \frac{2}{n_1} \eta \mathbf{X}'_1(y_1 - \mathbf{X}_1(\beta^+ - \beta^-)) = 0 \\ \theta^+, \theta^-, \eta &\geq 0 \\ (\theta^+)' \beta^+ &= 0 \\ (\theta^-)' \beta^- &= 0 \\ \eta \left( \frac{1}{n_1} \|y_1 - \mathbf{X}_1(\beta^+ - \beta^-)\|^2 - (1 + \tau)MSE_1(\hat{\beta}^{ols}) \right) &= 0 \end{aligned}$$

Thus,

$$\begin{aligned} - \text{if } \beta > 0 &\Rightarrow \beta^+ > 0, \beta^- = 0 \Rightarrow \theta^+ = 0 \Rightarrow -\frac{2}{n_0} \mathbf{X}'_0(y_0 - \mathbf{X}_0\beta) + \lambda - \frac{2}{n_1} \eta \mathbf{X}'_1(y_1 - \mathbf{X}_1\beta) = 0 \\ - \text{if } \beta < 0 &\Rightarrow \beta^+ = 0, \beta^- > 0 \Rightarrow \theta^- = 0 \Rightarrow \frac{2}{n_0} \mathbf{X}'_0(y_0 - \mathbf{X}_0\beta) + \lambda + \frac{2}{n_1} \eta \mathbf{X}'_1(y_1 - \mathbf{X}_1\beta) = 0 \end{aligned}$$

Therefore,

$$\frac{2}{n_0} \mathbf{X}'_0(y_0 - \mathbf{X}_0\beta) + \frac{2}{n_1} \eta(\lambda) \mathbf{X}'_1(y_1 - \mathbf{X}_1\beta) = \mathbf{b}(\lambda), \tag{18}$$

where  $\eta(\lambda)$  is the Lagrange multiplier associated with the first constraint and  $\mathbf{b}(\lambda)$  is a  $(p + 1)$ -dimensional vector whose  $s$ -th component,  $s = 0, 1, \dots, p$ , takes the following value

$$b_s(\lambda) = \begin{cases} \lambda, & \text{if } \beta_s > 0, \\ -\lambda, & \text{if } \beta_s < 0, \\ 0, & \text{else.} \end{cases}$$

Then, since  $\mathbf{X}_0$  and  $\mathbf{X}_1$  are maximum rank matrices, one obtains from (18) the following implicit expression for the solution  $\hat{\beta}^{CSCLasso}(\lambda)$  of Problem (11)

$$\begin{aligned} \hat{\beta}^{CSCLasso}(\lambda) &= \left( \frac{1}{n_0} \mathbf{X}'_0 \mathbf{X}_0 + \frac{1}{n_1} \eta(\lambda) \mathbf{X}'_1 \mathbf{X}_1 \right)^{-1} \left( \frac{1}{n_0} \mathbf{X}'_0 \mathbf{y}_0 + \frac{1}{n_1} \eta(\lambda) \mathbf{X}'_1 \mathbf{y}_1 \right) \\ &\quad - \frac{1}{2} \left( \frac{1}{n_0} \mathbf{X}'_0 \mathbf{X}_0 + \frac{1}{n_1} \eta(\lambda) \mathbf{X}'_1 \mathbf{X}_1 \right)^{-1} \mathbf{b}(\lambda). \end{aligned}$$

**Proof of Theorem 1**

Consider the function  $h : \beta \mapsto \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|^2 = \frac{1}{n} (\mathbf{y} - \mathbf{X}\beta)' (\mathbf{y} - \mathbf{X}\beta)$ , where  $\mathbf{X}$  is a maximum rank matrix by hypothesis. The matrix  $\frac{\mathbf{X}}{n}$  is of maximum rank and therefore the Hessian matrix  $H_h(\beta) = \frac{2}{n} \mathbf{X}'\mathbf{X}$  is positive definite, from where we conclude that  $h(\beta)$  is strictly convex, and hence,  $h(\beta) + \lambda \|\mathcal{A}\beta\|_1$  is also a strictly convex function.

We next prove that  $h(\beta)$  is a coercive function. Since  $\mathbf{X}'\mathbf{X}$  is positive definite, its eigenvalues are all positive. In particular, the smallest eigenvalue, say  $\gamma_r$ , will be nonzero. Moreover, using the spectral decomposition of a symmetric matrix,

$$\begin{aligned} \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|^2 &= \frac{1}{n} (\mathbf{y} - \mathbf{X}\beta)' (\mathbf{y} - \mathbf{X}\beta) = \frac{1}{n} \beta' \mathbf{X}' \mathbf{X} \beta - \frac{2}{n} \mathbf{y}' \mathbf{X} \beta + \frac{1}{n} \mathbf{y}' \mathbf{y} \\ &= \frac{1}{n} \beta' \mathbf{Q}' \mathbf{D} \mathbf{Q} \beta - \frac{2}{n} \mathbf{y}' \mathbf{X} \beta + \frac{1}{n} \mathbf{y}' \mathbf{y} \\ &\geq \frac{1}{n} \beta' \mathbf{Q}' \mathbf{D} \mathbf{Q} \beta - \left| \frac{2}{n} \mathbf{y}' \mathbf{X} \beta \right| + \frac{1}{n} \mathbf{y}' \mathbf{y} \geq \frac{\gamma_r}{n} \|\mathbf{Q}\beta\|^2 - \left\| \frac{2}{n} \mathbf{y}' \mathbf{X} \right\| \|\beta\| + \frac{1}{n} \mathbf{y}' \mathbf{y} \\ &= \frac{\gamma_r}{n} \|\beta\|^2 - \left\| \frac{2}{n} \mathbf{y}' \mathbf{X} \right\| \|\beta\| + \frac{1}{n} \mathbf{y}' \mathbf{y}, \end{aligned}$$

where, in the second-to-last step, the Cauchy-Schwarz inequality has been used. As  $\|\beta\| \rightarrow +\infty$ , then  $h(\beta) \rightarrow +\infty$  too, and thus  $h(\beta)$  is a coercive function.

Now we show that (13) has optimal solution. Let  $\beta^* \in \mathbf{B}$ . As  $h(\beta)$  is coercive, then there exists  $R > 0$  such that

$$\frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|^2 > \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta^*\|^2 + \lambda \|\mathcal{A}\beta^*\|_1,$$

for all  $\beta$  such that  $\|\beta\| > R$ . For that reason, the problem can be reduced to the feasible compact region  $B \cap \{\beta : \|\beta\| \leq R\}$ , which implies that the optimal solution is reached. Finally, the uniqueness of the solution follows from the fact that  $h(\beta) + \lambda\|\mathcal{A}\beta\|_1$  is strictly convex.

**Proof of Proposition 2**

Let us consider the optimization problem (2) and let  $\hat{\beta}^{Lasso}(\lambda)$  denotes its optimal solution. The necessary and sufficient optimality condition is:

$$\nabla \frac{1}{n} \|\mathbf{y} - \mathbf{X}\hat{\beta}^{Lasso}(\lambda)\|^2 + \lambda \partial \|\mathcal{A}\hat{\beta}^{Lasso}(\lambda)\|_1 \ni \mathbf{0}. \tag{19}$$

From the properties of subdifferential (see *Theorem 23.9* of Rockafellar (1972)) it follows that

$$\partial \|\mathcal{A}\hat{\beta}^{Lasso}(\lambda)\|_1 = \mathcal{A}' \partial \|\cdot\|_1|_{\mathcal{A}\hat{\beta}^{Lasso}(\lambda)},$$

which implies that (19) becomes

$$-\frac{2}{n} \mathbf{X}'(\mathbf{y} - \mathbf{X}\hat{\beta}^{Lasso}(\lambda)) + \lambda \mathcal{A}' \partial \|\cdot\|_1|_{\mathcal{A}\hat{\beta}^{Lasso}(\lambda)} \ni \mathbf{0}. \tag{20}$$

Consequently, the necessary and sufficient condition (20) in  $\hat{\beta}^{Lasso}(\lambda) = \mathbf{0}$  is

$$-\frac{2}{n} \mathbf{X}'\mathbf{y} + \lambda \{\mathcal{A}'\mathbf{t} : \|\mathbf{t}\|_\infty \leq 1\} \ni \mathbf{0},$$

since  $\partial \|\mathbf{0}\|_1$  is the unit ball of the  $\|\cdot\|_\infty$ . Equivalently,

$$\frac{2}{n} \mathbf{X}'\mathbf{y} \in \{\mathcal{A}'\lambda\mathbf{t} : \|\mathbf{t}\|_\infty \leq 1\}.$$

Therefore, the solution of the problem

$$\begin{aligned} \min_{\lambda, \mathbf{t}} \quad & \lambda \\ \text{s.t.} \quad & \frac{2}{n} \mathbf{X}'\mathbf{y} = \mathcal{A}'\lambda\mathbf{t}, \\ & \|\mathbf{t}\|_\infty \leq 1, \\ & \lambda \geq 0, \end{aligned} \tag{21}$$

will provide the minimum  $\lambda$  from which  $\hat{\beta}^{Lasso}(\lambda) = \mathbf{0}$  is the optimal solution. If  $\mathbf{q} = \lambda \mathbf{t}$ , then Problem (21) becomes

$$\begin{aligned} \min_{\lambda, \mathbf{q}} \quad & \lambda \\ \text{s.t.} \quad & \frac{2}{n} \mathbf{X}'\mathbf{y} = \mathcal{A}'\mathbf{q}, \\ & \|\mathbf{q}\|_{\infty} \leq \lambda. \end{aligned}$$

The constraint  $\|\mathbf{q}\|_{\infty}$  is equivalent to  $|q_s| \leq \lambda, s = 0, 1, \dots, p$  and the result follows.

**Proof of Proposition 3**

The proof follows very closely that of Theorem 1. First of all, it shall be proven that  $h : \beta \mapsto E[(Y - X'\beta)^2]$  is coercive. It is strictly convex on  $\beta$  since its Hessian matrix,  $2E[XX']$ , is positive definite due to  $X$  is an absolutely continuous  $p$ -dimensional random variable:

$$u'E[XX']u = E[u'XX'u] = E[(X'u)^2] > 0,$$

since  $P(X'u = 0) = 0$ . Moreover,  $\lambda\|\mathcal{A}\beta\|_1$  is a convex function on  $\beta$  and, therefore,  $E[(Y - X'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1]$  is also a strictly convex function on  $\beta$ .

On the one hand, the eigenvalues of the Hessian matrix are all positive and, in particular, the smallest eigenvalue, say  $\gamma_r$ , will be non-zero. On the other hand, using the spectral decomposition of a symmetric matrix,

$$\begin{aligned} E[(Y - X'\beta)^2] &= \beta'E[XX']\beta - 2E[XY]\beta + E[Y^2] = \beta'Q'DQ\beta - 2E[XY]\beta + E[Y^2] \\ &\geq \beta'Q'DQ\beta - |2E[XY]\beta| + E[Y^2] \geq \gamma_r\|Q\beta\|^2 - \|E[XY]\|\|\beta\| + E[Y^2] \\ &= \gamma_r\|\beta\|^2 - \|E[XY]\|\|\beta\| + E[Y^2], \end{aligned}$$

where, in the second-to-last step, the Cauchy-Schwarz inequality was used. As  $\|\beta\| \rightarrow +\infty$ , then  $E[(Y - X'\beta)^2] \rightarrow +\infty$ , that is, the quadratic function  $h(\beta) = E[(Y - X'\beta)^2]$  is coercive. The next step in the proof is to transform the original *true* problem (17) into an equivalent one with a feasible compact region  $\mathbf{B}^*$ . Given  $\beta^* \in \mathbf{B}$ , since  $h(\beta) = E[(Y - X'\beta)^2]$  is coercive, there exists  $R$  such that

$$E[(Y - X'\beta)^2] > E[(Y - X'\beta^*)^2] + \lambda\|\mathcal{A}\beta^*\|_1,$$

for all  $\beta$  with  $\|\beta\| > R$ . For that reason, the problem (17) can be reduced to the feasible compact region  $\mathbf{B}^* = \mathbf{B} \cap \{\beta : \|\beta\| \leq R\}$ , which implies that the optimal solution is reached.

Finally, the uniqueness of solution is a consequence of the strict convexity of the objective function.

**Proof of Theorem 2**

For the sake of simplicity,  $\beta^{CSCLasso}(\lambda)$  and  $\hat{\beta}^{CSCLasso}(\lambda)$  will be denoted henceforth by  $\beta$  and  $\hat{\beta}$ , respectively. In addition, let us consider the nonempty compact set  $C = B \cap \{\beta : \|\beta\| \leq R\}$ , where  $R$  is chosen according to the proof of Theorem 3.1.

Theorem 2 is a direct consequence of Theorem 5.3 in Shapiro et al. (2009) under some technical conditions, namely:

- C1. The expected value function  $E[(Y - X'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1]$  is finite valued and continuous on  $C$ .
- C2.  $\frac{1}{n} \sum_{i=1}^n ((y_i - x_i'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1)$  converges to  $E[(Y - X'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1]$  w.p. 1, as  $n \rightarrow \infty$ , uniformly in  $\beta \in C$ .

Let us denote  $F(\beta, (Y, X)) = (Y - X'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1$ . Then, the previous conditions C1 and C2 are consequences of Theorem 7.48 in Shapiro et al. (2009) provided that

- A1. for any  $\beta \in C$ , the function  $F(\cdot, (Y, X))$  is continuous at  $\beta$  for almost every  $(Y, X)$ ,
- A2. the function  $F(\beta, (Y, X))$ , with  $\beta \in C$ , is dominated by an integrable function,
- A3. the sample is i.i.d.

Given  $(Y, X)$ , the function  $(Y - X'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1$  is continuous at  $\beta$  for any  $\beta \in C$ , and therefore A1 is fulfilled. The sample is i.i.d. by hypothesis, and thus A3 holds too. Finally, in order to prove A2, it is necessary to find a measurable function  $g(Y, X) > 0$  such that  $E[g(Y, X)] < \infty$  and, for every  $\beta \in C$ ,  $|F(\beta, (Y, X))| \leq g(Y, X)$  w.p. 1. Using the Cauchy-Schwarz inequality, one has,

$$\begin{aligned} |F(\beta, (Y, X))| &= |(Y - X'\beta)^2 + \lambda\|\mathcal{A}\beta\|_1| \\ &= |Y^2 - 2YX'\beta + \beta'XX'\beta + \lambda\|\mathcal{A}\beta\|_1| \\ &\leq Y^2 + (X'\beta)^2 + 2|YX'\beta| + \lambda\|\mathcal{A}\beta\|_1 \\ &= Y^2 + |X'\beta|^2 + 2|YX'\beta| + \lambda\|\mathcal{A}\beta\|_1 \\ &\leq Y^2 + \|X\|^2\|\beta\|^2 + 2\|YX\|\|\beta\| + \lambda\|\mathcal{A}\beta\|_1. \end{aligned}$$

Let  $M_1$  and  $M_2$  be given by

$$M_1 = \max_{\beta \in C} \|\beta\| \qquad M_2 = \max_{\beta \in C} |\mathcal{A}\beta|$$

which are well defined due to the compactness of  $C$ . Therefore,  $g$  can be chosen as

$$g(Y, X) = Y^2 + M_1^2\|X\|^2 + 2M_1\|YX\| + \lambda M_2,$$

which is positive and, since  $E(\|X\|^2) < \infty$ ,  $E(Y^2) < \infty$ ,  $E(\|YX\|) < \infty$ , its expected value is finite. In consequence, A2 holds and the proof is concluded.

## Appendix: Further results

See Figs. 7, 8, 9 and 10.

To fully understand how the computation time behaves depending on  $n_k$  and  $p$  values, a grid in both parameters have been inspected. Figure 11 displays the logarithm of the user times in seconds obtained under Lasso and CSCLasso models when  $n_k$  and  $p$  change. The perspective drawn in the top left figure shows that Lasso model (bottom surface) is solved faster and in a smoother way. Besides, whereas smaller times are obtained for both methods when  $n_k$  and  $p$  are small, the biggest times are associated to  $n_k = 300$  and  $p = 500$ .

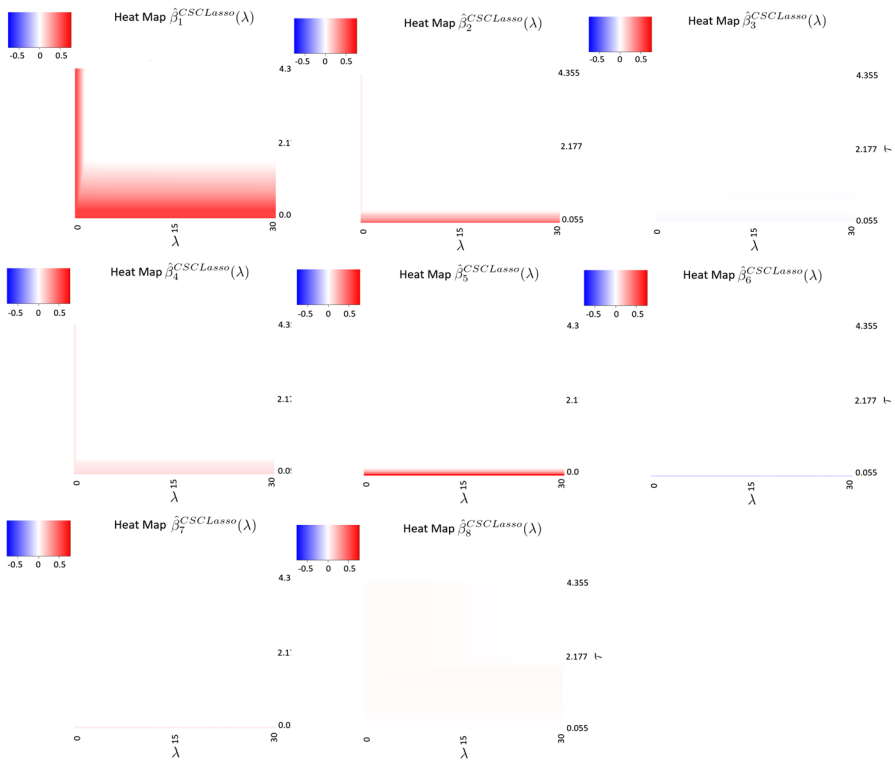

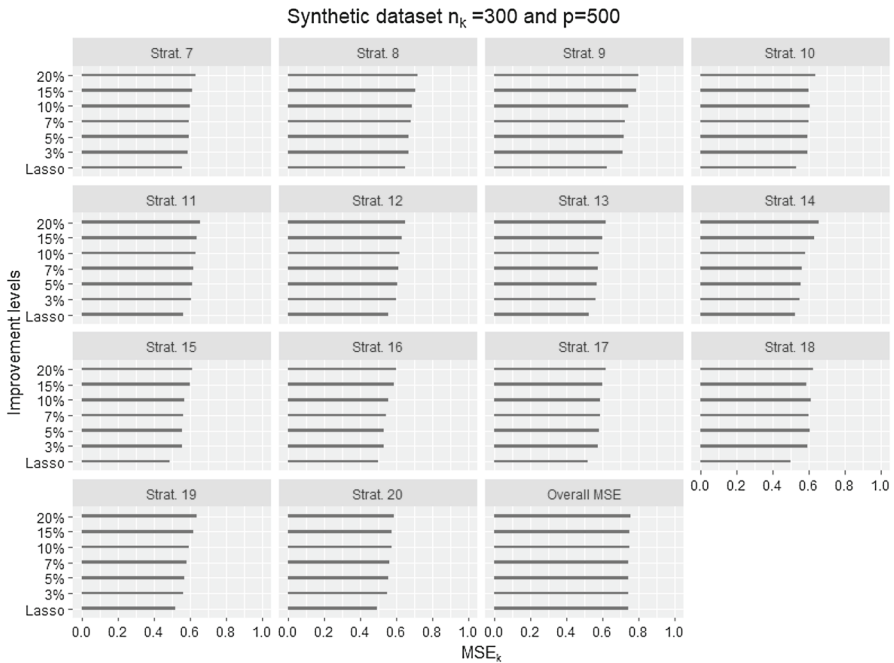
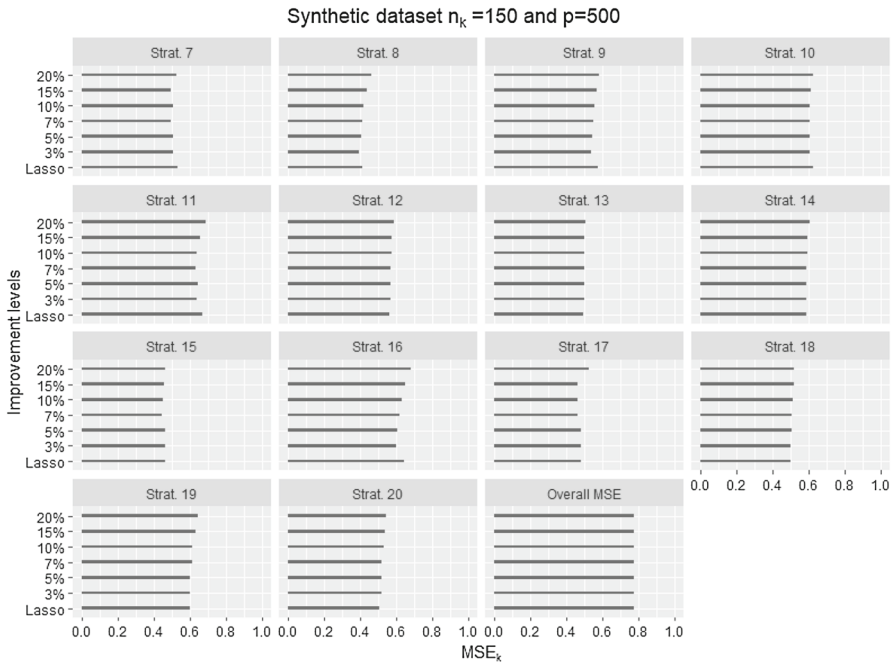


Fig. 7 Heat maps of  $\hat{\beta}^{CSCLasso}(\lambda) = (\hat{\beta}_1^{CSCLasso}(\lambda), \dots, \hat{\beta}_8^{CSCLasso}(\lambda))$  using prostate dataset

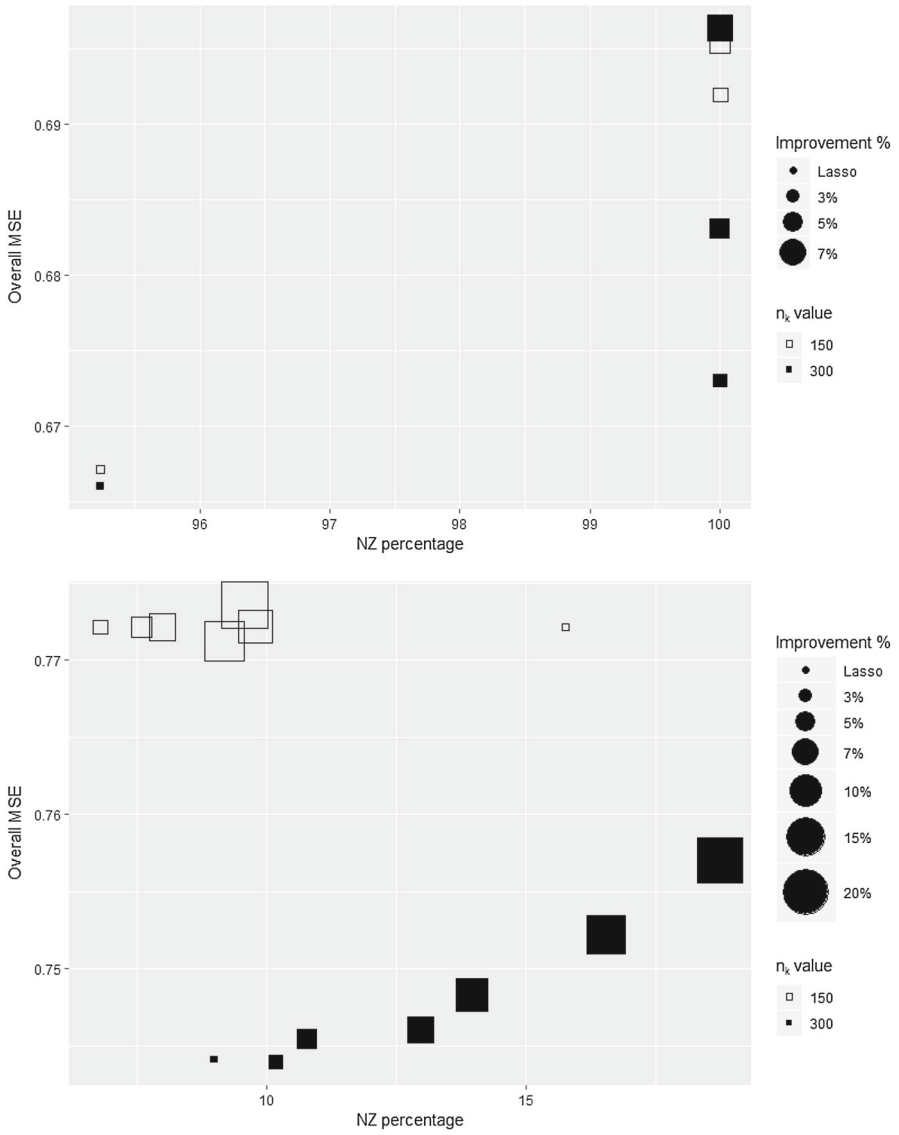




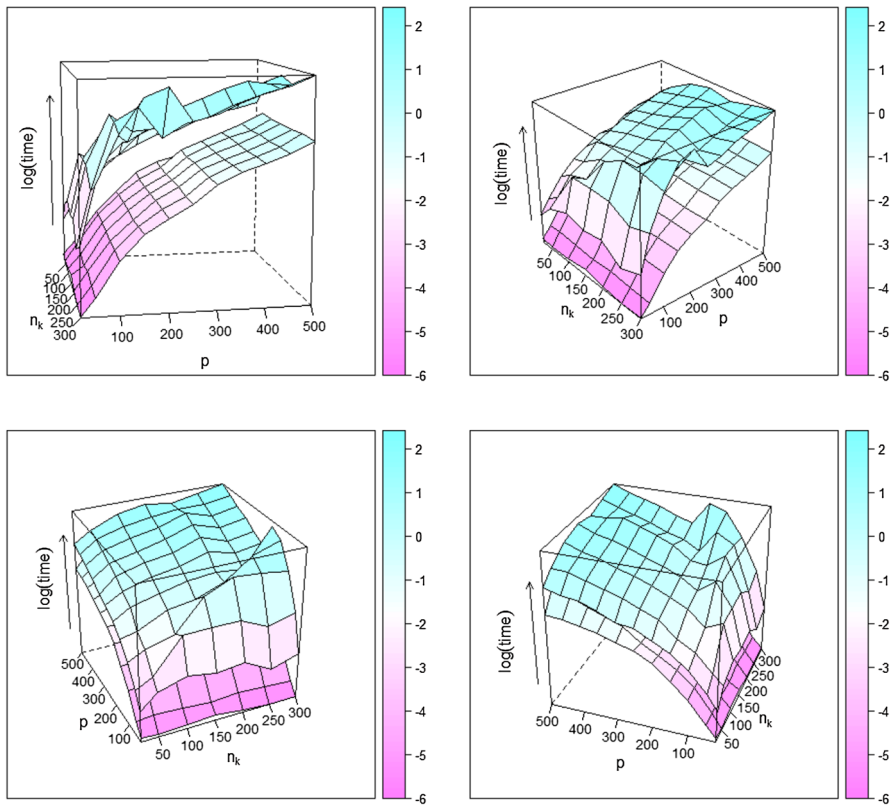
**Fig. 8** Median  $MSE_k$  over the test sets for  $k = 7, \dots, 20$  under  $p = 20$  features and the two  $n_k$  options. Each subgraph represents one group and the Y-axis shows the different percentages of improvement



**Fig. 9** Median  $MSE_k$  over the test sets for  $k = 7, \dots, 20$  under  $p = 500$  features and the two  $n_k$  options. Each subgraph represents one group and the Y-axis shows the different percentages of improvement



**Fig. 10** Median overall MSE over the test sets and NZ percentage under the choice  $p = 20$  (top) and  $p = 500$  (bottom)



**Fig. 11** Four perspectives of the logarithm of the user times in seconds for Lasso (bottom surface in the four graphics) and CSCLasso (top surfaces) models across a grid in  $n_k$  and  $p$

## References

- Bradford JP, Kunz C, Kohavi R, Brunk C, Brodley CE (1998) Pruning decision trees with misclassification costs. In: Nédellec C, Rouveiroi C (eds) *Machine learning: ECML-98*. Springer, Berlin, pp 131–136
- Bühlmann P, Van-De Geer S (2011) *Statistics for high-dimensional data*. Springer, Berlin
- Carrizosa E, Romero-Morales D (2001) Combining minsum and minmax: a goal programming approach. *Oper Res* 49(1):169–174
- Carrizosa E, Martín-Barragán B, Morales DR (2008) Multi-group support vector machines with measurement costs: a biobjective approach. *Discrete Appl Math* 156:950–966
- Datta S, Das S (2015) Near-Bayesian support vector machines for imbalanced data classification with equal or unequal misclassification costs. *Neural Netw* 70:39–52
- Donoho DL, Johnstone IM, Kerkyacharian G, Picard D (1995) Wavelet shrinkage: Asymptopia? *J R Stat Soc Ser B (Methodol)* 57(2):301–369
- Efron B, Hastie T, Johnstone I, Tibshirani R (2004) Least angle regression. *Ann Stat* 32(2):407–499
- Fan J, Li R (2001) Variable selection via nonconcave penalized likelihood and its oracle properties. *J Am Stat Assoc* 96(456):1348–1360
- Freitas A, Costa-Pereira A, Brazdil P (2007) Cost-sensitive decision trees applied to medical data. In: Song IY, Eder J, Nguyen TM (eds) *Data warehousing and knowledge discovery*. Springer, Berlin, pp 303–312
- Friedman J, Hastie T, Tibshirani R (2001) *The elements of statistical learning*. Springer, Heidelberg

- Gaines BR, Kim J, Zhou H (2018) Algorithms for fitting the constrained Lasso. *J Comput Graph Stat* 27(4):861–871
- Garside MJ (1965) The best sub-set in multiple regression analysis. *J R Stat Soc Ser C (Appl Stat)* 14(2–3):196–200
- Gurobi Optimization L (2018) Gurobi optimizer reference manual. <http://www.gurobi.com>
- Hastie T, Tibshirani R, Wainwright M (2015) *Statistical learning with sparsity*. Chapman and Hall/CRC, New York
- He H, Ma Y (2013) *Imbalanced learning: foundations, algorithms, and applications*. Wiley, Hoboken
- Hoerl AE, Kennard RW (1970) Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* 12(1):55–67
- Hu Q, Zeng P, Lin L (2015) The dual and degrees of freedom of linearly constrained generalized lasso. *Comput Stat Data Anal* 86:13–26
- James GM, Paulson C, Rusmevichientong P (2019) Penalized and constrained optimization: an application to high-dimensional website advertising. *J Am Stat Assoc* 1–31
- Kouno T, de Hoon M, Mar JC, Tomaru Y, Kawano M, Carninci P, Suzuki H, Hayashizaki Y, Shin JW (2013) Temporal dynamics and transcriptional control using single-cell gene expression analysis. *Genome Biol* 14(10):R118
- Lee W, Jun CH, Lee JS (2017) Instance categorization by support vector machines to adjust weights in adaboost for imbalanced data classification. *Inf Sci* 381(Supplement C):92–103
- Lichman M (2013) UCI machine learning repository. <http://archive.ics.uci.edu/ml>
- Ollier E, Viallon V (2017) Regression modelling on stratified data with the lasso. *Biometrika* 104(1):83–96
- Prati RC, Batista GEAPA, Silva DF (2015) Class imbalance revisited: a new experimental setup to assess the performance of treatment methods. *Knowl Inf Syst* 45(1):247–270
- Redmond M, Baveja A (2002) A data-driven software tool for enabling cooperative information sharing among police departments. *Eur J Oper Res* 141(3):660–678
- Rockafellar RT (1972) *Convex analysis*. Princeton University Press, Princeton
- Shapiro A, Dentcheva D, Ruszczyński A (2009) *Lectures on stochastic programming: modeling and theory*. SIAM, Philadelphia
- Simon N, Friedman J, Hastie T, Tibshirani R (2011) Regularization paths for Cox’s proportional hazards model via coordinate descent. *J Stat Softw* 39(5):1–13
- Stamey TA, Kabalin JN, McNeal JE, Johnstone IM, Freiha F, Redwine EA, Yang N (1989) Prostate specific antigen in the diagnosis and treatment of adenocarcinoma of the prostate: II. Radical prostatectomy treated patients. *J Urol* 141(5):1076–1083
- Sun Y, Wong AK, Kamel MS (2009) Classification of imbalanced data: a review. *Int J Pattern Recognit Artif Intell* 23:687–719
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B (Methodol)* 58(1):267–288
- Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K (2005) Sparsity and smoothness via the fused lasso. *J R Stat Soc Ser B (Stat Methodol)* 67(1):91–108
- Tibshirani RJ, Taylor J (2011) The solution path of the generalized Lasso. *Ann Stat* 39(3):1335–1371
- Torres-Barrán A, Alaíz CM, Dorronsoro JR (2018)  $\nu$ -SVM solutions of constrained Lasso and elastic net. *Neurocomputing* 275:1921–1931
- U.S. Department of Commerce, Bureau of the Census, Census of Population and Housing 1990 United States: Summary Tape File 1a & 3a (Computer Files), U.S. Department of Commerce, Bureau of the Census Producer, Washington, DC and Inter-university Consortium for Political and Social Research Ann Arbor, Michigan (1992)
- U.S. Department of Justice, Bureau of Justice Statistics, Law Enforcement Management and Administrative Statistics (Computer File) U.S. Department Of Commerce, Bureau of the Census Producer, Washington, DC and Inter-university Consortium for Political and Social Research Ann Arbor, Michigan (1992)
- U.S. Department of Justice, Federal Bureau of Investigation, Crime in the United States (Computer File) (1995)
- Yu G, Liu Y (2016) Sparse regression incorporating graphical structure among predictors. *J Am Stat Assoc* 111(514):707–720
- Yuan M, Lin Y (2006) Model selection and estimation in regression with grouped variables. *J R Stat Soc Ser B (Stat Methodol)* 68(1):49–67
- Zou H (2006) The adaptive lasso and its oracle properties. *J Am Stat Assoc* 101(476):1418–1429

Zou H, Hastie T (2005) Regularization and variable selection via the elastic net. *J R Stat Soc Ser B (Stat Methodol)* 67(2):301–320

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.