



Linear support vector regression with linear constraints

Quentin Klopfenstein¹ · Samuel Vaïter²

Received: 17 March 2020 / Revised: 2 April 2021 / Accepted: 31 May 2021 /
Published online: 24 June 2021

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2021

Abstract

This paper studies the addition of linear constraints to the Support Vector Regression when the kernel is linear. Adding those constraints into the problem allows to add prior knowledge on the estimator obtained, such as finding positive vector, probability vector or monotone data. We prove that the related optimization problem stays a semi-definite quadratic problem. We also propose a generalization of the Sequential Minimal Optimization algorithm for solving the optimization problem with linear constraints and prove its convergence. We show that an efficient generalization of this iterative algorithm with closed-form updates can be used to obtain the solution of the underlying optimization problem. Then, practical performances of this estimator are shown on simulated and real datasets with different settings: non negative regression, regression onto the simplex for biomedical data and isotonic regression for weather forecast. These experiments show the usefulness of this estimator in comparison to more classical approaches.

Keywords Support vector machine · Support vector regression · Sequential minimal optimization · Coordinate descent · Constrained linear regression

1 Introduction

Regression analysis seeks to find a relation between explanatory variables and an outcome variable. The most known linear regression estimator is the Ordinary Least Squares (OLS) which is the best linear unbiased estimator if the noise is *i.i.d* and Gaussian. Additional constraints on the OLS estimator can be added to ensure that it follows specific properties. For example, penalized regression such as the Ridge regression improves the efficiency of the estimation by introducing a bias and reducing the variance of the estimator in presence of collinearity (Hoerl and Kennard 1970). Another family of constrained regression is defined by the addition of hard constraints on the estimator, such as positivity constraints

Editor: Bart Baesens.

✉ Samuel Vaïter
samuel.vaïter@math.cnrs.fr

¹ Institut Mathématique de Bourgogne, Université de Bourgogne, Dijon, France

² Centre National de la Recherche Scientifique (CNRS) and the Institut Mathématiques de Bourgogne, Université de Bourgogne, Dijon, France

(Lawson and Hanson 1995) which can improve the estimation performance. Some applications in biology where the goal is to estimate cell proportions inside a tumor or the estimation of temperature in weather forecast justify the use of hard constraints on the estimator. The constrained OLS has been well studied, but the literature around the Support Vector Regression with additional constraints is scarcer.

Support Vector Machine. The Support Vector Machine (SVM) (Boser et al. 1992) is a class of supervised learning algorithms that have been widely used in the past 20 years for classification tasks and regression. These algorithms rely on two main ideas: the first one is the maximum margin hyperplane which consists in finding the hyperplane that maximises the distance between the vectors that are to be classified and the hyperplane. The second idea is the kernel method that allows the SVM to be used to solve non-linear problems. The technique is to map the vectors in a higher dimensional space which is done by using a positive definite kernel, then a maximum margin hyperplane is computed in this space which gives a linear classifier in the high dimensional space. In general, it leads to a non-linear classifier in the original input space.

From SVM to Support Vector Regression. Different implementations of the algorithms have been proposed such as C-SVM, ν -SVM (Schölkopf et al. 1999), Least-Squares SVM (Suykens and Vandewalle 1999), Linear Programming SVM (Friel and Harrison 1998) among others. Each of these versions have their strengths and weaknesses depending on which application they are used. They differ in terms of constraints considered for the hyperplane (C-SVM and Least-Squares SVM), in terms of norm considered on the parameters (C-SVM and Linear Programming SVM) and in terms of optimization problem formulation (C-SVM and ν -SVM). Overall, these algorithms are a great tool for classification tasks and they have been used in many different applications like facial recognition (Jia and Martinez 2009), image classification (Chapelle et al. 1999), cancer type classification (Haussler et al. 2000), text categorization (Joachims 1998) to only cite a few examples. Even though, SVM was first developed for classification, an adaptation for regression estimation was proposed in Drucker et al. (1997) under the name Support Vector Regression (SVR). In this case, the idea of maximum margin hyperplane is slightly changed into finding a tube around the regressors. The size of the tube is controlled by a hyperparameter chosen by the user: ϵ . This is equivalent to using an ϵ -insensitive loss function, $|y - f(x)|_{\epsilon} = \max\{0, |y - f(x)| - \epsilon\}$ which only penalizes the error above the chosen ϵ level. As for the classification version of the algorithm, a ν -SVR method exists. In this version, the hyperparameter ϵ is computed automatically but a new hyperparameter ν has to be chosen by the user which controls asymptotically the proportions of support vectors (Schölkopf et al. 1999). SVR has proven to be a great tool in the field of function estimation for many different applications: predicting times series in stock trades (Van Gestel et al. 2001), travel-time prediction (Chun-Hsin et al. 2004) and for estimating the amount of cells present inside a tumor (Newman et al. 2015).

Incorporating priors. In this last example of application, the authors used SVR to estimate a vector of proportions, however the classical SVR estimator does not take into account the information known about the space in which the estimator lives. Adding this prior information on the estimator may lead to better estimation performance. Incorporating information in the estimation process is a wide field of studies in statistical learning (we refer to Fig. 2 in Lauer and Bloch (2008) for a quick overview in the context of SVM). A growing interest in prior knowledge incorporated as regularization terms has emerged in the last decades. Lasso (Tibshirani 1996), Ridge (Hoerl and Kennard 1970), elastic-net (Zou and Hastie 2005) regression are examples of regularized problem where a prior information is used to fix an ill-posed problem or an overdetermined problem. The

ℓ_1 regularization of the Lasso will force the estimator to be sparse and bring statistical guarantees of the Lasso estimator in high dimensional settings. Another common way to add prior knowledge on the estimator is to add constraints known a-priori on this estimator. The most common examples are the ones that constrain the estimator to live in a subspace such as Non Negative Least Squares Regression (NNLS) (Lawson and Hanson 1995), isotonic regression (Barlow and Brunk 1972). These examples belong to a more general type of constraints: linear constraints. Other types of constraints exist like constraints on the derivative of the function that is to be estimated, smoothness of the function for example. Adding those constraints on the Least Squares estimator has been widely studied (Barlow and Brunk 1972; Liew 1976; Bro and De Jong 1997) and similar work has been done for the Lasso estimator (Gaines et al. 2018). Concerning the SVR, inequality and equality constraints added as prior knowledge were studied in Lauer and Bloch (2008). In this paper, the authors described a method for adding linear constraints on the Linear Programming SVR (Friel and Harrison 1998). This implementation of the algorithm considers the ℓ_1 norm of the parameters in the optimization problem instead of the classical ℓ_2 norm which leads to a linear programming optimization problem to solve instead of a quadratic programming problem. They also described a method for using information about the derivative of the function that is estimated.

Sequential Minimal Optimization. One of the main challenges of adding these constraints is that it often increases the difficulty of solving the optimization problem related to the estimator. For example, the Least Squares optimization problem has a closed-form solution whereas the NNLS uses sophisticated algorithms (Bro and De Jong 1997) to approach the solution. SVM and SVR algorithms were extensively studied and used in practise because very efficient algorithms were developed to solve the underlying optimization problems. One of them is called Sequential Minimal Optimization (SMO) (Platt 1998) and is based on a well known optimization technique called coordinate descent. The idea of the coordinate descent is to break the optimization problem into sub-problems selecting a subset of coordinates at each step and minimizing the function only via this chosen subset. The development of parallel algorithms have increased the interest in these coordinate descent methods which show to be very efficient for large scale problems. One of the key settings for the coordinate descent is the choice of the coordinate at each step, the choice's strategy will affect the efficiency of the algorithm. There exists three families of strategies for coordinate descent: cyclic (Tseng 2001), random (Nesterov 2012) and greedy. The SMO algorithm is a variant of a greedy coordinate descent (Wright 2015) and is the algorithm implemented in LibSVM (Chang and Lin 2011). It is very efficient to solve SVM/SVR optimization problems. In the context of linear kernel, other algorithm are used such as dual coordinate descent (Hsieh et al. 2008) or trust region newton methods (Lin et al. 2007).

Priors and SMO. In one of the application of SVR cited above, information a-priori about the estimator is not used in the estimation process and is only used in a post-processing step. This application comes from the cancer research field, where regression algorithms have been used to estimate the proportions of cell populations that are present inside a tumor (see Mohammadi et al. (2017) for a survey). Several estimators have been proposed in the biostatistics literature, most of them based on constrained least squares (Abbas et al. 2009; Qiao et al. 2012; Gong et al. 2011) but the gold standard is the estimator based on the Support Vector Regression (Newman et al. 2015). Our work is motivated by incorporating the fact that the estimator for this application belongs to the simplex: $\mathcal{S} = \{x \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1, x_i \geq 0\}$ in the SVR problem. We believe that for this application, it will lead to better estimation performance. From an optimization

point of view, our motivation is to find an efficient algorithm that is able to solve the SVR optimization problem where generic linear constraints are added to the problem as prior knowledge, including simplex prior as described. This work follows the one from Lauer and Bloch (2008) except that in our case, we keep the ℓ_2 norm on the parameters in the optimization problem which is the most common version of the SVR optimization problem and we only focus on inequality and equality constraints as prior knowledge.

Contributions. In this paper, we study a linear SVR with linear constraints optimization problem. We show that the dual of this new problem shares similar properties with the classical ν -SVR optimization problem (Proposition 2). We also prove that adding linear constraints to the SVR optimization problem does not change the nature of its dual problem, in the fact that the problem stays a semi-definite positive quadratic function subject to linear constraints. We propose a generalized SMO algorithm that allows the resolution of the new optimization problem. We show that the updates in the SMO algorithm keep a closed-form (Definition 5) and prove the convergence of the algorithm to a solution of the problem (Theorem 1). We illustrate on synthetic and real datasets the usefulness of our new regression estimator under different regression settings: non-negative regression, simplex regression and isotonic regression.

Outline. The article proceeds as follows: we introduce the optimization problem coming from the classical SVR and describe the modifications brought by adding linear constraints in Sect. 2. We then present the SMO algorithm, its generalization for solving constrained SVR and present our result on the convergence of the algorithm in Sect. 3. In Sect. 4, we use synthetic and real datasets on different regression settings to illustrate the practical performance of the new estimator.

Notations. We write $\|\cdot\|$ (resp. $\langle \cdot, \cdot \rangle$) for the euclidean norm (resp. inner product) on vectors. We use the notation $X_{:,i}$ (resp. $X_{i,:}$) to denote the vector corresponding to the i^{th} column of the matrix X (resp. i^{th} row of the matrix X). Throughout this paper, the design matrix will be $X \in \mathbb{R}^{n \times p}$ and $y \in \mathbb{R}^n$ will be the response vector. X^T will be used for the transposed matrix of X . The vector \mathbf{e} denotes the vector with only ones on each of its coordinates and e_j denotes the canonical vector with a one at the j^{th} coordinate. $\nabla_{x_i} f$ is the partial derivative $\frac{\partial f}{\partial x_i}$.

2 Constrained support vector regression

We study the underlying optimization problem that arises when adding linear constraints to the SVR estimator. One way to solve the optimization problem is to perform coordinate descent algorithms variants in its dual. In this section, we derive the dual optimization problem of the constrained SVR in order to apply coordinate descent variants to solve the constrained version, as it will be shown in Sect. 3. We describe the similarities and the differences between the classical problem and the constrained one.

2.1 Previous work : ν -support vector regression

The ν -SVR estimator (Schölkopf et al. 1999) is obtained solving the following quadratic optimization problem:

$$\begin{aligned} \min_{\beta, \beta_0, \xi_i, \xi_i^*, \epsilon} \quad & \frac{1}{2} \|\beta\|^2 + C \left(\nu \epsilon + \frac{1}{n} \sum_{i=1}^n (\xi_i + \xi_i^*) \right) \\ \text{subject to} \quad & y_i - \beta^T X_i - \beta_0 \leq \epsilon + \xi_i \\ & \beta^T X_i + \beta_0 - y_i \leq \epsilon + \xi_i^* \\ & \xi_i, \xi_i^* \geq 0, \epsilon \geq 0. \end{aligned} \tag{SVR-P}$$

By solving Problem (SVR-P), we seek a linear function $f(x) = \beta^T x + \beta_0$ where $\beta \in \mathbb{R}^p$ and $\beta_0 \in \mathbb{R}$, that is at most ϵ deviating from the response vector coefficient y_i . This function does not always exist which is why slack variables $\xi \in \mathbb{R}^n$ and $\xi^* \in \mathbb{R}^n$ are introduced in the optimization problem to allow some observations to break the condition given before. C and ν are two hyperparameters. $C \in \mathbb{R}$ controls the tolerated error and $\nu \in [0, 1]$ controls the number of observations that will lay inside the tube of size 2ϵ given by the two first constraints in (SVR-P). It can be seen as an ϵ -insensitive loss function where a linear penalization is put on the observations that lay outside the tube and the observations that lay inside the tube are not penalized (see Smola and Schölkopf (2004) for more details).

The different algorithms proposed to solve the primal Problem (SVR-P) often use its dual problem like in Platt (1998) and Hsieh et al. (2008). The dual problem is also a quadratic optimization problem with linear constraints but its structure allows an efficient resolution as we will see in more details in Sect. 3. The dual problem of (SVR-P) is the following optimization problem:

$$\begin{aligned} \min_{\alpha, \alpha^*} \quad & \frac{1}{2} (\alpha - \alpha^*)^T Q (\alpha - \alpha^*) + y^T (\alpha - \alpha^*) \\ \text{subject to} \quad & 0 \leq \alpha_i, \alpha_i^* \leq \frac{C}{n} \\ & \mathbf{e}^T (\alpha + \alpha^*) \leq C\nu \\ & \mathbf{e}^T (\alpha - \alpha^*) = 0, \end{aligned} \tag{SVR-D}$$

where $Q = XX^T \in \mathbb{R}^{n \times n}$. The equation link between (SVR-P) and (SVR-D) is given by the following formula: $\beta = -\sum_{i=1}^n (\alpha_i - \alpha_i^*) X_i$.

2.2 The constrained optimization problem

We propose a constrained version of Problem (SVR-P) that allows the addition of prior knowledge on the linear function f that we seek to estimate. The constrained estimator is obtained solving the optimization problem:

$$\begin{aligned}
 & \min_{\beta, \beta_0, \xi_i, \xi_i^*, \epsilon} \quad \frac{1}{2} \|\beta\|^2 + C \left(\nu \epsilon + \frac{1}{n} \sum_{i=1}^n (\xi_i + \xi_i^*) \right) \\
 & \text{subject to } \beta^T X_i + \beta_0 - y_i \leq \epsilon + \xi_i \\
 & \quad y_i - \beta^T X_i - \beta_0 \leq \epsilon + \xi_i^* \tag{LSVR-P} \\
 & \quad \xi_i, \xi_i^* \geq 0, \epsilon \geq 0 \\
 & \quad A\beta \leq b \\
 & \quad \Gamma\beta = d,
 \end{aligned}$$

where $A \in \mathbb{R}^{k_1 \times p}$, $\Gamma \in \mathbb{R}^{k_2 \times p}$, $\beta \in \mathbb{R}^p$, $\xi, \xi^* \in \mathbb{R}^n$ and $\beta_0, \epsilon, \in \mathbb{R}$.

The algorithm that we propose in Sect. 3 also uses the structure of the dual problem of (LSVR-P). The next proposition introduces the dual problem and some of its properties.

Proposition 1 *If the set $\{\beta \in \mathbb{R}^n, A\beta \leq b, \Gamma\beta = d\}$ is not empty then,*

1. *Strong duality holds for (LSVR-P).*
2. *The dual problem of (LSVR-P) is*

$$\begin{aligned}
 & \min_{\alpha, \alpha^*, \gamma, \mu} \quad \frac{1}{2} [(\alpha - \alpha^*)^T Q(\alpha - \alpha^*) + \gamma^T A A^T \gamma \\
 & \quad + \mu^T \Gamma \Gamma^T \mu + 2 \sum_{i=1}^n (\alpha_i - \alpha_i^*) \gamma^T A X_i \\
 & \quad - 2 \sum_{i=1}^n (\alpha_i - \alpha_i^*) \mu^T \Gamma X_i - 2 \gamma^T A \Gamma^T \mu] \\
 & \quad + \gamma^T (\alpha - \alpha^*) + \gamma^T b - \mu^T d \tag{LSVR-D} \\
 & \text{subject to } 0 \leq \alpha_i^{(*)} \leq \frac{C}{n} \\
 & \quad \mathbf{e}^T (\alpha + \alpha^*) \leq C\nu \\
 & \quad \mathbf{e}^T (\alpha - \alpha^*) = 0 \\
 & \quad \gamma_j \geq 0,
 \end{aligned}$$

3. *The equation link between primal and dual is*

$$\beta = - \sum_{i=1}^n (\alpha_i - \alpha_i^*) X_i - A^T \gamma + \Gamma^T \mu.$$

The proof of the first statement of the proposition is given in the discussion below whereas the proofs for the two other statements come from a straightforward derivation of the Lagrangian dual of (LSVR-P). We have that $\alpha, \alpha^* \in \mathbb{R}^n, \gamma \in \mathbb{R}^{k_1}$ are the vector of Lagrange multipliers associated the the inequality constraint $A\beta \leq b$ which explains the non-negative constraints on its coefficients. $\mu \in \mathbb{R}^{k_2}$ are the Lagrange multipliers associated to the equality constraint $\Gamma\beta = d$ which also explains that there are no constraints in the dual problem on μ . The objective function f which we will write in the stacked form as:

$$f(\theta) = \theta^T \bar{Q} \theta + l^T \theta,$$

where

$$\theta = \begin{bmatrix} \alpha \\ \alpha^* \\ \gamma \\ \mu \end{bmatrix}, \quad l = \begin{bmatrix} y \\ -y \\ b \\ -d \end{bmatrix} \in \mathbb{R}^{2n+k_1+k_2},$$

$$\bar{Q} = \begin{bmatrix} Q & -Q & XA^T & -X\Gamma^T \\ -Q & Q & -XA^T & X\Gamma^T \\ AX^T & -AX^T & AA^T & -A\Gamma^T \\ -\Gamma X^T & \Gamma X^T & -\Gamma A^T & \Gamma\Gamma^T \end{bmatrix}$$

is a square matrix of size $2n + k_1 + k_2$.

An important observation is that this objective function is always convex. The matrix \bar{Q} is the product of the matrix $\begin{bmatrix} X \\ -X \\ A \\ -\Gamma \end{bmatrix}$ and its transpose matrix. It means that \bar{Q} is a Gramian matrix and it is positive semi-definite which implies that f is convex. The problem (LSVR-D) is then a quadratic programming optimization problem which meets Slater's condition if there exists a θ that belongs to the feasible domain which we will denote by \mathcal{F} . If there is such a θ we have strong duality holding between problem (LSVR-P) and (LSVR-D). The only condition we need to have on A and Γ is that they define a non-empty polyhedron in order to be able to solve the optimization problem.

Our second observation on problem (LSVR-D) is that the inequality constraints $e^T(\alpha + \alpha^*) \leq Cv$ is replaced by an equality constraint in the same way that it was suggested in Chang and Lin (2002) for the classical problem (SVR-D).

Proposition 2 *If ϵ is strictly positive, all optimal solutions of (LSVR-D) satisfy*

1. $\alpha_i \alpha_i^* = 0, \forall i \leq n,$
2. $e^T(\alpha + \alpha^*) = Cv.$

The proof is given in Appendix A. This observation will be important for the algorithm that we propose in Sect. 3.

3 Generalized sequential minimal optimization

In this section we propose a generalization of the SMO algorithm (Platt 1998) to solve problem (LSVR-D) and present our main result on the convergence of the proposed algorithm to the solution of (LSVR-D). The SMO algorithm is a variant of greedy coordinate descent taking into consideration non-separable constraints, which in our case are the two equality constraints. We start by describing the previous algorithm that solve (SVR-D).

3.1 Previous work: sequential minimal optimization

In this subsection, we define

$$f(\alpha, \alpha^*) = \frac{1}{2}(\alpha - \alpha^*)^T Q(\alpha - \alpha^*) + y^T(\alpha - \alpha^*)$$

and we note $\nabla f \in \mathbb{R}^{2n}$ its gradient. From Keerthi and Gilbert (2002), we rewrite the Karush-Kuhn-Tucker (KKT) conditions in the following way:

$$\min_{i \in I_{\text{up}}} \nabla_{\alpha_i} f \geq \max_{j \in I_{\text{low}}} \nabla_{\alpha_j} f, \quad (1)$$

where $I_{\text{up}}(\alpha) = \{i \in \{1, \dots, n\} : \alpha_i < \frac{C}{n}\}$ and $I_{\text{low}}(\alpha) = \{i \in \{1, \dots, n\} : \alpha_i > 0\}$.

The same condition is written for the α^* variables replacing α_i by α_i^* above. These conditions leads to an important definition for the rest of this paper.

Definition 1 We will say that (i, j) is a violating pair of variables if one of these two conditions is satisfied:

$$i \in I_{\text{up}}(\alpha), j \in I_{\text{low}}(\alpha) \text{ and } \nabla_{\alpha_i} f < \nabla_{\alpha_j} f,$$

$$i \in I_{\text{low}}(\alpha), j \in I_{\text{up}}(\alpha) \text{ and } \nabla_{\alpha_i} f > \nabla_{\alpha_j} f.$$

Because the algorithm SMO does not provide in general an exact solution in a finite number of steps there is a need to relax the optimality conditions which gives a new definition.

Definition 2 We will say that (i, j) is a τ -violating pair of variables if one of these two conditions is satisfied:

$$i \in I_{\text{up}}(\alpha), j \in I_{\text{low}}(\alpha) \text{ and } \nabla_{\alpha_i} f < \nabla_{\alpha_j} f - \tau,$$

$$i \in I_{\text{low}}(\alpha), j \in I_{\text{up}}(\alpha) \text{ and } \nabla_{\alpha_i} f > \nabla_{\alpha_j} f + \tau.$$

The SMO algorithm will then choose at each iteration a pair of violating variables in the α block or in the α^* block. Once the choice is done, a subproblem of size two is solved, considering that only the two selected variables are to be minimized in problem (SVR-D). The outline of the algorithm is presented in Algorithm 1.

The choice of the violating pair of variables presented in Keerthi et al. (2001) was to always work with the most violating pairs of variables, which means the variables that leads to the largest gap compared to the optimality conditions given in (1). This choice is what makes a link with greedy coordinate descent, however greedy here is related to the

largest gap with the optimality score and is not related to the largest decrease in the objective function.

The resolution of the subproblem of size two has a closed-form. The idea is to use the two equality constraints to go from a problem of size two to a problem of size one. Then, the goal is to minimize a quadratic function of one variable under box constraints which is done easily. We will give more details of the resolution of these subproblems in Sect. 3.3 for our proposed algorithm.

The proof of convergence of SMO algorithm was given in Keerthi and Gilbert (2002) without convergence rate. The proof relies on showing that the sequence defined by the algorithm $f(\alpha^k, (\alpha^*)^k)$ is a decreasing sequence and that there cannot be the same violating pair of variables infinitely many times. The linear convergence rate was proved later by She (2017) as well as the identification of the support vectors in finite time.

Algorithm 1 SMO algorithm (Platt, 1998)

Require: $\tau > 0$

Initializing $\alpha^0 \in \mathbb{R}^n$, $(\alpha^*)^0 \in \mathbb{R}^n$ in \mathcal{F} and set $k = 0$

while $\Delta > \tau$ **do**

$i \leftarrow \operatorname{argmin}_{i \in I_{\text{up}}} \nabla_{\alpha_i} f$ $j \leftarrow \operatorname{argmax}_{i \in I_{\text{low}}} \nabla_{\alpha_j} f$

$i^* \leftarrow \operatorname{argmin}_{i \in I_{\text{up}}^*} \nabla_{\alpha_i^*} f$ $j^* \leftarrow \operatorname{argmax}_{i \in I_{\text{low}}^*} \nabla_{\alpha_j^*} f$

$\Delta_1 \leftarrow \nabla_{\alpha_j} f - \nabla_{\alpha_i} f$

$\Delta_2 \leftarrow \nabla_{\alpha_j^*} f - \nabla_{\alpha_i^*} f$

$\Delta \leftarrow \max(\Delta_1, \Delta_2)$

▷ Select the maximal violating pair

if $\Delta = \Delta_1$ **then**

$\alpha^{k+1} \leftarrow$ Solution of subproblem for α_i and α_j

$(\alpha^*)^{k+1} \leftarrow (\alpha^*)^k$

else

$(\alpha^*)^{k+1} \leftarrow$ Solution of subproblem for α_{i^*} and α_{j^*}

$\alpha^{k+1} \leftarrow \alpha^k$

$k \leftarrow k + 1$

return $\alpha^k, (\alpha^*)^k$

3.2 Optimality conditions for the constrained SVR

In this subsection we define f as the objective function of Problem (LSVR-D) and $\nabla f \in \mathbb{R}^{2n+k_1+k_2}$ its gradient. We will now give the KKT conditions of (LSVR-D) for the different blocks. The results derive from classical Lagrangian calculation.

The α block **Case 1-** $\alpha_i = 0$ then $\beta_i = 0$ and $\lambda_i \geq 0$, we have that

$$\nabla_{\alpha_i} f - \sigma - \delta \geq 0.$$

Case 2- $\alpha_i = \frac{C}{n}$ then $\lambda_i = 0$ and $\beta_i \geq 0$, we have that

$$\nabla_{\alpha_i} f - \sigma - \delta \leq 0.$$

Case 3- $0 < \alpha_i < \frac{C}{n}$ then $\beta_i = 0$, we have that

$$\nabla_{\alpha_i} f - \sigma - \delta = 0.$$

We then consider the sets of indices:

$$I_{\text{up}}(\alpha) = \left\{ i \in \{1, \dots, n\} : \alpha_i < \frac{C}{n} \right\},$$

$$I_{\text{low}}(\alpha) = \{ i \in \{1, \dots, n\} : \alpha_i > 0 \}.$$

The optimality conditions are satisfied if and only if

$$\min_{i \in I_{\text{up}}} \nabla_{\alpha_i} f \geq \max_{j \in I_{\text{low}}} \nabla_{\alpha_j} f.$$

The α^* block In this block, the conditions are very similar to the ones given for the block α , the only difference here is that we will have two new sets of indices:

$$I_{\text{up}}^*(\alpha^*) = \left\{ i \in \{1, \dots, n\} : \alpha_i^* < \frac{C}{n} \right\},$$

$$I_{\text{low}}^*(\alpha) = \{ i \in \{1, \dots, n\} : \alpha_i^* > 0 \}.$$

which gives the following optimality condition:

$$\min_{i \in I_{\text{up}}^*} \nabla_{\alpha_i^*} f \geq \max_{j \in I_{\text{low}}^*} \nabla_{\alpha_j^*} f.$$

The γ block

$$\nabla_{\gamma_j} L = \nabla_{\gamma_j} f - \eta_j = 0$$

$$\eta_j \gamma_j = 0$$

$$\eta_j \geq 0.$$

We will consider different possibilities of value for γ_j .

Case 1- $\gamma_j = 0$ then $\nabla_{\gamma_j} f \geq 0$.

Case 2- $\gamma_j > 0$ then $\nabla_{\gamma_j} f = 0$.

Definition 3 We will say that j is a τ -violating variable for the block γ if $\nabla_{\gamma_j} f + \tau < 0$.

The μ block

$$\nabla_{\mu_j} L = \nabla_{\mu_j} f = 0.$$

Definition 4 We will say that j is a τ -violating variable for the block μ if $|\nabla_{\mu_j} f| > \tau$.

From these conditions on each block, we build an optimization strategy that follows the idea of the SMO described in Sect. 3.1. For each block of variables, we compute what we call a *violating optimality score* based on the optimality conditions given above. Once the scores are computed for each block, we select the block which has the largest score and solve an optimization subproblem in the block selected. If the block α or the block α^* is selected, we will update a pair of variables by solving a minimization problem of size two. However if the block γ or the block μ is selected, we will update only one variable at a time. This is justified by the fact that the variables α and α^* have non-separable equality

constraints linking them together. The rest of this section will be dedicated to the presentation of our algorithm and to giving some interesting properties such as a closed-form for updates on each of the blocks and a convergence theorem.

Algorithm 2 Generalized SMO algorithm

Require: $\tau > 0$

Initializing $\alpha^0 \in \mathbb{R}^n$, $(\alpha^*)^0 \in \mathbb{R}^n$, $\gamma^0 \in \mathbb{R}^{k_1}$ and $\mu^0 \in \mathbb{R}^{k_2}$ in \mathcal{F} and set $k = 0$

while $\Delta > \tau$ **do**

$$i \leftarrow \underset{i \in I_{\text{up}}}{\operatorname{argmin}} \nabla_{\alpha_i} f \qquad j \leftarrow \underset{i \in I_{\text{low}}}{\operatorname{argmax}} \nabla_{\alpha_j} f$$

$$i^* \leftarrow \underset{i \in I_{\text{up}}^*}{\operatorname{argmin}} \nabla_{\alpha_i^*} f \qquad j^* \leftarrow \underset{i \in I_{\text{low}}^*}{\operatorname{argmax}} \nabla_{\alpha_j^*} f$$

$$\Delta_1 \leftarrow \nabla_{\alpha_j} f - \nabla_{\alpha_i} f \qquad \Delta_2 \leftarrow \nabla_{\alpha_j^*} f - \nabla_{\alpha_i^*} f$$

$$\Delta_3 \leftarrow - \min_{j \in \{1, \dots, k_1\}} \nabla_{\gamma_j} f \qquad \Delta_4 \leftarrow \max_{j \in \{1, \dots, k_2\}} |\nabla_{\mu_j} f|$$

$$\Delta \leftarrow \max(\Delta_1, \Delta_2, \Delta_3, \Delta_4) \qquad \triangleright \text{Select the maximal violating variables}$$

if $\Delta = \Delta_1$ **then**

$$\alpha^{k+1} \leftarrow \text{Solution of subproblem for } \alpha_i \text{ and } \alpha_j$$

else if $\Delta = \Delta_2$ **then**

$$(\alpha^*)^{k+1} \leftarrow \text{Solution of subproblem for } \alpha_{i^*} \text{ and } \alpha_{j^*}$$

else if $\Delta = \Delta_3$ **then**

$$u = \underset{i \in \{1, \dots, k_1\}}{\operatorname{argmin}} \nabla_{\gamma_i} f$$

$$\gamma^{k+1} \leftarrow \text{Solution of subproblem for } \gamma_u$$

else

$$u = \underset{i \in \{1, \dots, k_2\}}{\operatorname{argmax}} |\nabla_{\mu_i} f|$$

$$\mu^{k+1} \leftarrow \text{Solution of subproblem for } \mu_u$$

$$k \leftarrow k + 1$$

$$\text{return } \alpha^k, (\alpha^*)^k, \gamma^k, \mu^k$$

3.3 Updates rules and convergence

The first definition describes the closed-form updates for the different blocks of variables.

Definition 5 The update between iterate k and iterate $k + 1$ of the generalized SMO algorithm has the following form:

1. if the block α is selected and (i, j) is the most violating pair of variable then the update will be as follows:

$$\alpha_i^{k+1} = \alpha_i^k + t^*,$$

$$\alpha_j^{k+1} = \alpha_j^k - t^*,$$

where $t^* = \min(\max(I_1, -\frac{(\nabla_{\alpha_i} f - \nabla_{\alpha_j} f)}{(Q_{ii} - 2Q_{ij} + Q_{jj})}), I_2)$ with $I_1 = \max(-\alpha_i^k, \alpha_j^k - \frac{c}{n})$ and $I_2 = \min(\alpha_j^k, \frac{c}{n} - \alpha_i^k)$.

- if the block α^* is selected and (i^*, j^*) is the most violating pair of variable then the update will be as follows:

$$\begin{aligned} (\alpha_i^*)^{k+1} &= (\alpha_i^*)^k + t^*, \\ (\alpha_j^*)^{k+1} &= (\alpha_j^*)^k - t^*, \end{aligned}$$

where $t^* = \min(\max(I_1, -\frac{(\nabla_{\alpha_i^*} f - \nabla_{\alpha_j^*} f)}{(Q_{ii} - 2Q_{ij} + Q_{jj})}), I_2)$ with $I_1 = \max(-(\alpha_i^*)^k, (\alpha_j^*)^k - \frac{C}{n})$ and $I_2 = \min((\alpha_j^*)^k, \frac{C}{n} - (\alpha_i^*)^k)$.

- if the block γ is selected and i is the index of the most violating variable in this block then the update will be as follows:

$$\gamma_i^{k+1} = \max\left(-\frac{\nabla_{\gamma_i} f}{(AA^T)_{ii}} + \gamma_i^k, 0\right).$$

- if the block μ is selected and i is the index of the most violating variable in this block then the update will be as follows:

$$\mu_i^{k+1} = -\frac{\nabla_{\mu_i} f}{(\Gamma\Gamma^T)_{ii}} + \mu_i^k.$$

This choice of updates comes from solving the optimization problem (LSVR-D) considering that only one or two variables are updated at each step. One of the key elements of the algorithm is to make sure that at each step the iterate belongs to \mathcal{F} . Let us suppose that the block α is selected as the block in which the update will happen and let (i, j) be the most violating pair of variables. The update is the resolution of a subproblem of size 2, considering that only α_i and α_j are the variables, the rest remains constant. The two equality constraints in (LSVR-D), $\sum_{i=1}^n \alpha_i - \alpha_i^* = 0$ and $\sum_{i=1}^n \alpha_i + \alpha_i^* = Cv$, lead to the two following equalities: $\alpha_i^{k+1} + \alpha_j^{k+1} = \alpha_i^k + \alpha_j^k$. The later yields to using a parameter t for the update of the variables leading to $\alpha_i^{k+1} = \alpha_i^k + t$ and $\alpha_j^{k+1} = \alpha_j^k - t$.

Updating the variable in the block α this way will force the iterates of Algorithm 1 to meet the two equalities constraints at each step. We find t by solving the dual problem (LSVR-D) considering that we minimize only over t . Let the vector $u \in \mathbb{R}^{2n+p+k_1+k_2}$ which contains only zeros except at the i^{th} coordinate where it is equal to t and at j^{th} coordinate where it is equal to $-t$. Therefore, we find t by minimizing the following optimization problem:

$$\begin{aligned} \min_{t \in \mathbb{R}} \quad & \psi(t) = \frac{1}{2} [(\theta^k + u)^T \bar{Q}(\theta^k + u)] + l^T(\theta^k + u) \\ \text{subject to} \quad & 0 \leq \alpha_i^{k+1}, \alpha_j^{k+1} \leq \frac{C}{n}. \end{aligned}$$

First we minimize the objective function without the constraints and since it is a quadratic function of one variable we just clip the solution of unconstrained problem to have the solution of the constrained problem. We will use the term “clipped update” or “clipping” when the update is projected unto the constraints space and is not the result of the unconstrained optimization problem. As we only consider size one problem for the updates, it will mean that the update will be a bound of an interval. We will use the notation K as a term containing the terms that do not depend on t . We write that

$$\begin{aligned} \psi(t) &= \frac{1}{2}u^T \bar{Q}u + u^T \bar{Q}\theta^k + l^T u + K \\ &= \frac{1}{2}t^2(\bar{Q}_{ii} + \bar{Q}_{jj} - 2\bar{Q}_{ij}) + u^T \nabla f(\theta^k) + K \\ &= \frac{1}{2}t^2(\bar{Q}_{ii} + \bar{Q}_{jj} - 2\bar{Q}_{ij}) + t(\nabla_{\alpha_i} f(\theta^k) \\ &\quad - \nabla_{\alpha_j} f(\theta^k)) + K. \end{aligned}$$

It follows that the unconstrained minimum of $\psi(t)$ is $t_q = \frac{-(\nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k))}{(\bar{Q}_{ii} + \bar{Q}_{jj} - 2\bar{Q}_{ij})}$.

Taking the constraints into account we have that $0 \leq \alpha_i^k + t \leq \frac{C}{n}$ and $0 \leq \alpha_j^k - t \leq \frac{C}{n}$, which yields to $t^* = \min(\max(I_1, t_q), I_2)$ with $I_1 = \max(-\alpha_i, \alpha_j - \frac{nC}{n})$ and $I_2 = \min(\alpha_j, \frac{nC}{n} - \alpha_i)$. The definition of the updates for the block α^* relies on the same discussion.

Let us now make an observation that will explain the definition of the updates for the blocks γ and μ . Let i be the index of the variable that will be updated. Solving the problem; $\theta_i^{k+1} = \underset{\theta_i}{\operatorname{argmin}} \frac{1}{2}\theta^T \bar{Q}\theta + l^T \theta$, leads to the following solution $\theta_i^{k+1} = \frac{-\nabla_{\theta_i} f(\theta^k)}{\bar{Q}_{ii}} + \theta_i^k$.

Let us recall that the update for the block γ has to keep the coefficient of γ positive to stay in \mathcal{F} hence we have to perform the following clipped update with $i \in \{2n + p + 1, \dots, 2n + p + k_1\}$:

$$\theta_i^{k+1} = \max\left(\frac{-\nabla_{\gamma_i} f(\theta^k)}{\bar{Q}_{ii}} + \theta_i^k, 0\right).$$

Then noticing that $\bar{Q}_{ii} = AA_{ii}^T$ for this block, we obtain the update for the block γ .

There are no constraints on the variables in the block μ , so the update comes from the fact that $\bar{Q}_{ii} = \Gamma \Gamma_{ii}^T$ for $i \in \{2n + p + k_1 + 1, \dots, 2n + p + k_1 + k_2\}$ which corresponds to the indices of the block μ .

From these updates we have to make sure that

$$Q_{ii} + Q_{jj} - 2Q_{ij} \neq 0,$$

let us recall that $Q_{ij} = \langle X_i, X_j \rangle$ which means that

$$Q_{ii} + Q_{jj} - 2Q_{ij} = \|X_i - X_j\|^2.$$

This quantity is zero only when $X_i = X_j$: coordinate wise. It would mean that the same row appears two times in the design matrix which does not bring any new information for the regression and can be avoided easily. $(AA^T)_{ii} = \langle A_i, A_i \rangle$ is zero if and only if $A_i = 0$ which means that a row of the matrix A is zero, so there is no constraint on any variable of the optimization problem which will never happen. It is the same discussion for $(\Gamma \Gamma^T)_{ii}$.

The next proposition makes sure that once a variable (resp. pair of variables) is updated, it cannot be a violating variable (resp. pair of variables) at the next step. This proposition makes sure, for the two blocks α and α^* , that the update t^* cannot be 0.

Proposition 3 *If (i, j) (resp. i) was the pair of most violating variable (resp. the most violating variable) in the block α or α^* (resp. block γ or μ) at iteration k then at iteration $k + 1$, (i, j) (resp. i) cannot be violating the optimality conditions.*

The proof of this proposition is left in the Appendix B.

Finally, we show that the algorithm converges to a solution of (LSVR-D) and since strong duality holds it allows us to have a solution of (LSVR-P).

Theorem 1 *For any given $\tau > 0$ the sequence of iterates $\{\theta^k\}$, defined by the generalized SMO algorithm, converges to an optimal solution of the optimization problem (LSVR-D)*

The proof of this theorem relies on the same idea as the one proposed in Lopez and Dorronsoro (2012) for the classical SMO algorithm and is given in Appendix C. We show that it can be extended to our algorithm with some new observations. The general idea of the proof is to see that the distance between the primal vector generated by the SMO-algorithm and the optimal solution of the primal is controlled by the following expression $\frac{1}{2} \|\beta^k - \beta^{\text{opt}}\| \leq f(\theta^k) - f(\theta^{\text{opt}})$, where β^k is the k^{th} primal iterate obtained via the relationship primal-dual and θ^k and where β^{opt} is a solution of (LSVR-P). From this observation, we show that we can find a subsequence of the SMO-algorithm θ^{k_j} that converges to some $\bar{\theta}$, solution of the dual problem. Using the continuity of the objective function of the dual problem, we have that $f(\theta^{k_j}) \rightarrow f(\bar{\theta})$. Finally, we show that the sequence $\{f(\theta^k)\}$ is decreasing and bounded which implies its convergence and from the convergence monotone theorem we know that $f(\theta^k)$ converges to $f(\bar{\theta})$ since one of its subsequences converges. This proves that $\|\beta^k - \beta^{\text{opt}}\| \rightarrow 0$ and finishes the proof. The convergence rate for the SMO algorithm is difficult to obtain considering the greedy choice of the blocks and the greedy choice inside the blocks. A proof for the classical SMO exists but with uniformly at random choice of the block (She 2017). Convergence rate for greedy algorithms in optimization can be found in Nutini et al. (2015) for example but the assumption that the constraints must be separable is a major issue for our case. The study of this convergence rate is out of scope of this paper.

4 Numerical experiments

The code for the different regression settings is available on a GitHub repository¹, each setting is wrapped up in a package and is fully compatible with scikit learn (Pedregosa et al. 2011) BaseEstimator class.

In order to compare the estimators, we worked with the Mean Absolute Error; $\text{MAE} = \frac{1}{p} \sum_{i=1}^p |\beta_i^* - \hat{\beta}_i|$ and the Root Mean Squared Error; $\text{RMSE} = \sqrt{\frac{1}{p} \|\beta^* - \hat{\beta}\|^2}$, where β^* are the ground truth coefficients and $\hat{\beta}$ are the estimated coefficients. We also used the Signal-To-Noise Ratio (SNR) to control the level noise simulated in the data defined as $\text{SNR} = 10 \log_{10} \left(\frac{\mathbb{E}(X\beta(X\beta)^T)}{\text{Var}(e)} \right)$.

4.1 Non negative regression

First, the constraints are set to force the coefficient of β to be positive and we compare our constrained-SVR estimator with the NNLS (Lawson and Hanson 1995) estimator which is the result of the following optimization problem

¹ <https://github.com/Klopfel/LSVR>

$$\begin{aligned} \min_{\beta} \quad & \frac{1}{2} \|y - X\beta\|^2 \\ \text{subject to} \quad & \beta_i \geq 0. \end{aligned} \tag{NNLS}$$

In this special case of non-negative regression, $A = -I_p$, $b = 0$, $\Gamma = 0$, $d = 0$, the constrained-SVR optimization problem which we will call Non-Negative SVR (NNSVR) then becomes

$$\begin{aligned} \min_{\beta, \beta_0, \xi_i, \xi_i^*, \epsilon} \quad & \frac{1}{2} \|\beta\|^2 + C \left(\nu \epsilon + \frac{1}{n} \sum_{i=1}^n (\xi_i + \xi_i^*) \right) \\ \text{subject to} \quad & \beta^T X_i + \beta_0 - y_i \leq \epsilon + \xi_i \\ & y_i - \beta^T X_i - \beta_0 \leq \epsilon + \xi_i^* \\ & \xi_i, \xi_i^* \geq 0, \epsilon \geq 0 \\ & \beta_i \geq 0. \end{aligned} \tag{NNSVR}$$

Synthetic data. We generated the design matrix X from a Gaussian distribution $\mathcal{N}(0, 1)$ with 500 samples and 50 features. The true coefficients to be found β^* were generated taking the exponential of a Gaussian distribution $\mathcal{N}(0, 2)$ in order to have positive coefficients. Y was simply computed as the product between X and β^* . We wanted to test the robustness of our estimator compared to NNLS and variant of SVR estimators. To do so, we simulated noise in the data using different types of distributions, we tested Gaussian noise and Laplacian noise under different levels of noise. For this experiment, the noise distributions were generated to have an SNR equals to 10 and 20, for each type of noise we performed 50 repetitions. The noise was only added in the matrix Y , the design matrix X was left noiseless. The choice of the two hyperparameters C and ν was done using 5-folds cross validation on a grid of possible pairs. The values of C were taken evenly spaced in the \log_{10} base between $[-3, 3]$, we considered 10 different values. The values of ν were taken evenly spaced in the linear space between $[0.05, 1.0]$ and we also considered 10 possible values.

We compared different estimators: NNLS, NNSVR, the Projected-SVR (P-SVR) which is simply the projection of the classical SVR estimator onto the positive orthant and also the classical SVR estimator without constraints. The results of this experiment are in Table 1. We see that for a low Gaussian noise level (SNR = 20) the NNLS has a lower RMSE and lower MAE. However, we see that the differences between the four compared methods are small. When the level of noise increases (SNR = 10), the NNSVR estimator is the one with the lowest RMSE and MAE. The NNLS estimator performs poorly in the presence of high level of noise in comparison to the SVR based estimator. When a Laplacian noise is added to the data, the NNSVR is the estimator that has the lowest RMSE and MAE for low level of noise SNR = 20 and high level of noise SNR = 10.

4.2 Regression unto the simplex

In this subsection, we study the performance of our proposed estimator on simplex constraints Simplex Support Vector Regression (SSVR). In this case, $A = -I_p$, $b = 0$, $\Gamma = \mathbf{e}$ and $d = 1$. The optimization problem that we seek to solve is:

Table 1 Results for the support vector regression (SVR), projected support regression (P-SVR), non-negative support vector regression (NNSVR) and non-negative least squares (NNLS) for simulated data with $n = 500$ and $p = 50$

| Distribution | Estimator | RMSE | MAE |
|-----------------|-----------------------------------|----------------------|-----------------------|
| Gaussian noise | SVR | 2.238 (0.081) | 29.288 (2.452) |
| | P-SVR | 2.178 (0.087) | 27.248 (2.545) |
| | NNSVR | 2.174 (0.089) | 27.224 (2.480) |
| | SNR = 20 ($\sigma = 773.1$) | NNLS | 2.120 (0.114) |
| Gaussian noise | SVR | 2.732 (0.099) | 44.764 (4.230) |
| | P-SVR | 2.584 (0.154) | 39.687 (5.963) |
| | NNSVR | 2.536 (0.105) | 37.740 (3.866) |
| | SNR = 10 ($\sigma = 2444.9$) | NNLS | 3.478 (0.208) |
| Laplacian noise | SVR | 2.086 (0.109) | 25.538 (3.181) |
| | P-SVR | 2.039 (0.109) | 23.978 (3.059) |
| | NNSVR | 2.035 (0.115) | 23.827 (3.146) |
| | SNR = 20 ($b = 546.7$) | NNLS | 2.115 (0.103) |
| Laplacian noise | SVR | 2.665 (0.148) | 42.245 (5.777) |
| | P-SVR | 2.526 (0.198) | 37.745 (7.271) |
| | NNSVR | 2.480 (0.157) | 35.786 (5.761) |
| | SNR = 10 ($b = 1728.8$) | NNLS | 3.463 (0.230) |

The mean (standard deviation) of the root mean squared error (RMSE) and the mean absolute error (MAE) over 50 repetitions are reported. Different noise distribution (Gaussian and Laplacian) and different Signal to Noise Ratio (SNR) values were tested

$$\min_{\beta, \beta_0, \xi_i, \xi_i^*, \epsilon} \frac{1}{2} \|\beta\|^2 + C \left(\nu \epsilon + \frac{1}{n} \sum_{i=1}^n (\xi_i + \xi_i^*) \right)$$

$$\text{subject to } \beta^T X_i + \beta_0 - y_i \leq \epsilon + \xi_i$$

$$y_i - \beta^T X_i - \beta_0 \leq \epsilon + \xi_i^* \tag{SSVR}$$

$$\xi_i, \xi_i^* \geq 0, \epsilon \geq 0$$

$$\beta_i \geq 0$$

$$\sum_i \beta_i = 1.$$

Synthetic data. We first tested on simulated data generated by the function `make_regression` of `scikit-learn`. Once the design matrix X and the response vector y were generated using this function, we had access to the ground truth that we will write β^* . This function was not designed to generate data with a β^* that belongs to the simplex so we first projected β^* unto the simplex and then recomputed y multiplying the design matrix by the new projected β_S^* . We added a centered Gaussian noise in the data with the standard deviation chosen such as the signal-to-noise ratio (SNR) was equal to a defined number, we used the following formula for a given SNR $\sigma = \sqrt{\frac{\text{Var}(y)}{10^{\text{SNR}/10}}}$, where σ is the standard deviation used to simulate the noise in the data. The choice of the two hyperparameters C and ν was done using 5-folds cross validation on a grid of possible pairs. The values of C were taken

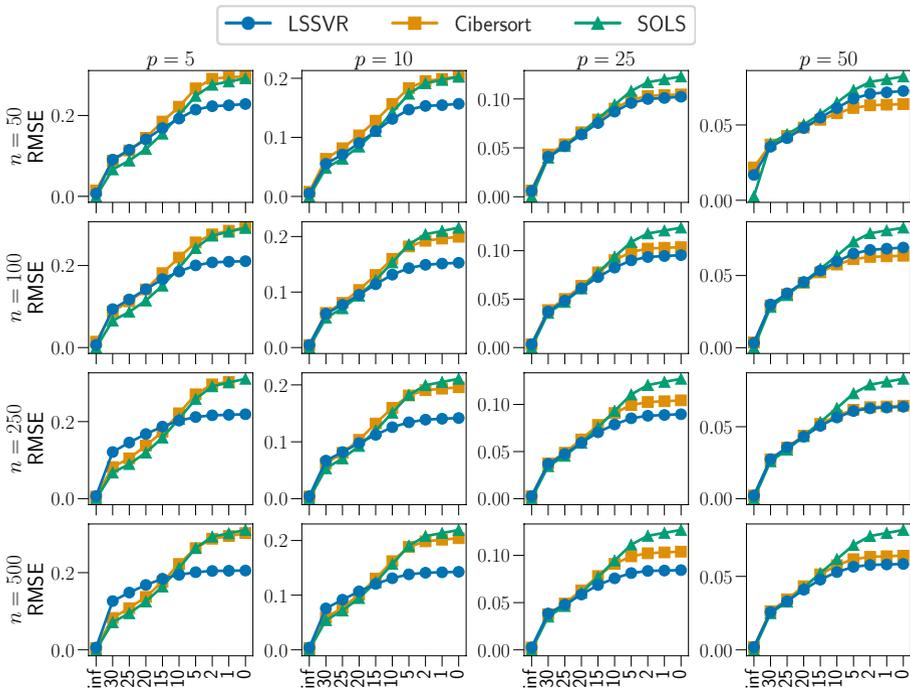


Fig. 1 The root mean squared error (RMSE) as a function of the signal to noise ratio (SNR) is presented. Different dimensions for the design matrix X and the response vector y were considered. n represents the number of rows of X and p the number of columns. For each plot, the blue line represents the RMSE for the Linear Simplex SVR (LSSVR) estimator, the green one the Simplex Ordinary Least Squares (SOLS) estimator and the orange one the Cibersort estimator. Each point of the curve is the mean RMSE of 50 repetitions. The noise in the data has a Gaussian distribution

evenly spaced in the \log_{10} base between $[-3, 3]$, we considered 10 different values. The values of v were taken evenly spaced in the linear space between $[0.05, 1.0]$ and we also considered 10 possible values. We tested different size for the matrix $X \in \mathbb{R}^{n \times p}$ to check the potential effects of the dimensions on the quality of the estimation and we did 50 repetitions for each point of the curves. The measure that was used to compare the different estimators is the RMSE between the true β and the estimated $\hat{\beta}$.

We compared the RMSE of our estimator to the Simplex Ordinary Least Squares (SOLS) which is the result of the following optimization problem:

$$\begin{aligned} \min_{\beta} \quad & \frac{1}{2} \|y - X\beta\|^2 \\ \text{subject to} \quad & \beta_i \geq 0, \\ & \sum_{i=1}^p \beta_i = 1, \end{aligned} \tag{SOLS}$$

and to the estimator proposed in the biostatistics literature that is called Cibersort. This estimator is simply the result of using the classical SVR and project the obtained estimator unto the simplex. The RMSE curves as a function of the SNR are presented in Fig. 1. We

observe that the SSVR is generally the estimator with the lowest RMSE, this observation becomes clearer as the level of noise increases in the data. We notice that when there is a low level of noise and when n is not too large in comparison to p , the three compared estimator perform equally. However, there is a setting when n is large in comparison to p (in this experiment for $n = 250$ or 500 and $p = 5$) where the SSVR estimator has a higher RMSE than the Cibersort and SOLS estimator until a certain level of noise ($\text{SNR} < 15$). Overall, this simulation shows that there is a significant improvement in the estimation performance of the SSVR mainly when there is noise in the data.

Real dataset. In the cancer research field, regression algorithms have been used to estimate the proportions of cell populations that are present inside a tumor. Indeed, a tumor is composed of different types of cells such as cancer cells, immune cells, healthy cells among others. Having access to the information of the proportions of these cells could be a key to understanding the interactions between the cells and the cancer treatment called immunotherapy (Couzin-Frankel 2013). The modelization done is that the RNA extracted from the tumor is seen as a mixed signal composed of different pure signals coming from the different types of cells. This signal can be unmixed knowing the different pure RNA signal of the different types of cells. In other words, y will be the RNA signal coming from a tumor and X will be the design matrix composed of the RNA signal from the isolated cells. The number of rows represent the number of genes that we have access to and the number of columns of X is the number of cell populations that we would like to quantify. The hypothesis is that there is a linear relationship between X and y . As said above, we want to estimate proportions which means that the estimator has to belong to the probability simplex $\mathcal{S} = \{x : x_i \geq 0, \sum_i x_i = 1\}$.

Several estimators have been proposed in the biostatistics literature most of them based on constrained least squares (Qiao et al. 2012; Gong et al. 2011; Abbas et al. 2009) but the gold standard is the estimator based on the SVR.

We compared the three same estimators on a real biological dataset where the real quantities of cells to obtain were known. The dataset can be found on the GEO website under the accession code GSE11103². For this example $n = 584$ and $p = 4$ and we have access to 12 different samples that are our repetitions. Following the same idea than previous benchmark performed in this field of application, we increased the level of noise in the data and compared the RMSE of the different estimators. Gaussian and Laplacian distributions of noise were added to the data. The choice of the two hyperparameters C and ν was done using 5-folds cross validation on a grid of possible pairs. The values of C were taken evenly spaced in the \log_{10} base between $[-5, -3]$, we considered 10 different values. The interval of C is different than the simulated data because of the difference in the range value of the dataset. The values of ν were taken evenly spaced in the linear space between $[0.05, 1.0]$ and we also considered 10 possible values.

We see that when there is no noise in the data ($\text{SNR} = \infty$) both Cibersort and SSVR estimator perform equally. The SOLS estimator already has a higher RMSE than the two others estimator probably due to the noise already present in the data. As the level of noise increases, the SSVR estimator remains the estimator with the lowest RMSE in both Gaussian and Laplacian noise settings (Fig. 2).

² The dataset can be downloaded from the <https://www.ncbi.nlm.nih.gov/geo/Gene Expression Omnibus> website under the accession code GSE11103.

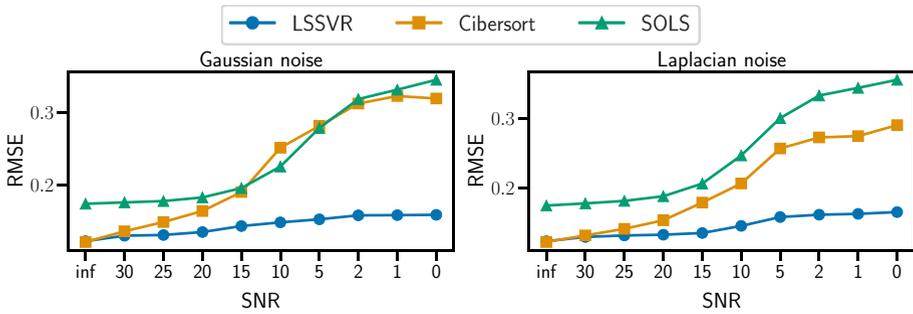


Fig. 2 The Root Mean Squared Error (RMSE) as a function of the Signal to Noise Ratio (SNR) is presented on a real dataset where noise was manually added. Two different noise distribution were tested: Gaussian and laplacian. Each point of the curve is the mean RMSE of 12 different response vectors and we repeated the process four times for each level of noise. This would be equivalent to having 48 different repetitions

4.3 Isotonic regression

In this subsection, we will consider constraints that impose an order on the variables. This type of regression is usually called isotonic regression. Such constraints appear when prior knowledge are known on a certain order on the variables. This partial order on the variables can also be seen as an acyclic directed graph. More formally, we note $G = (V, E)$ a directed acyclic graph where V is the set of vertices and E is the set of nodes. On this graph, we define a partial order on the vertices. We will say for $u, v \in V$ that $u \leq v$ if and only if there is a path joining u and v in G . This type of constraints seems natural in different applications such as biology, medicine, and weather forecast.

The most simple example of this type of constraints might be the monotonic regression where we force the variables to be in a increasing or decreasing order. It means that with our former notations that we would impose that $\beta_1 \leq \beta_2 \leq \dots \leq \beta_p$ on the estimator. This type of constraints can be coded in a finite difference matrix (or more generally any incidence matrix of a graph) where $A = (a_{i,j}) \in \mathbb{R}^{(p-1) \times p}$ defined by $a_{i,i} = 1, a_{i,i+1} = -1, a_{i,j} = 0$ otherwise, and $\Gamma = 0, b = 0, d = 0$ forming linear constraints as in the scope of this paper. The Isotonic Support Vector Regression (ISVR) optimization problem is written as follows:

$$\begin{aligned}
 & \min_{\beta, \beta_0, \xi_i, \xi_i^*, \epsilon} \quad \frac{1}{2} \|\beta\|^2 + C \left(\nu \epsilon + \frac{1}{n} \sum_{i=1}^n (\xi_i + \xi_i^*) \right) \\
 & \text{subject to } \beta^T X_i + \beta_0 - y_i \leq \epsilon + \xi_i \\
 & \quad \quad \quad y_i - \beta^T X_i - \beta_0 \leq \epsilon + \xi_i^* \\
 & \quad \quad \quad \xi_i, \xi_i^* \geq 0, \epsilon \geq 0 \\
 & \quad \quad \quad \beta_1 \leq \beta_2 \leq \dots \leq \beta_n.
 \end{aligned} \tag{ISVR}$$

We compare our proposed ISVR estimator with the classical least squares isotonic regression (IR) (Barlow and Brunk 1972) which is the solution of the following problem:

Table 2 Results for the isotonic support vector regression (ISVR), and the isotonic regression (IR) for simulated data with $p = 50$

| Distribution | Estimator | RMSE | MAE |
|-----------------------------|-----------|---------------------|---------------------|
| Gaussian noise SNR = 20 | ISVR | 0.212 (0.02) | 0.254 (0.06) |
| | IR | 0.203 (0.02) | 0.229 (0.04) |
| Gaussian noise SNR = 10 | ISVR | 0.284 (0.04) | 0.446 (0.12) |
| | IR | 0.311 (0.04) | 0.534 (0.12) |
| Laplacian noise SNR = 20 | ISVR | 0.202 (0.03) | 0.223 (0.05) |
| | IR | 0.203 (0.02) | 0.221 (0.04) |
| Laplacian noise SNR = 10 | ISVR | 0.276 (0.05) | 0.414 (0.11) |
| | IR | 0.312 (0.05) | 0.513 (0.13) |

The mean (standard deviation) of the root mean squared error (RMSE) and the mean absolute error (MAE) over 50 repetitions are reported. Different noise distribution (Gaussian and Laplacian) and different signal to noise ratio (SNR) values were tested

$$\min_{\beta} \quad \frac{1}{2} \|\beta - y\|^2 \tag{IR}$$

subject to $\beta_1 \leq \beta_2 \leq \dots \leq \beta_n$.

Synthetic dataset. We first generated data from a Gaussian distribution ($\mu = 0$, $\sigma = 1$) that we sorted and then added noise in the data following the same process as described in Sect. 4.2 with different SNR values (10 and 20). We tested Gaussian noise and Laplacian noise. We compared the estimation quality of both methods using MAE and RMSE. In this experiment, the design matrix X is the identity matrix. We performed grid search selection via cross validation for the hyperparameters C and ν . C had 5 different possible values taken on the logscale from 0 to 3, and ν had 5 different values taken between 0.05 and 1 on the linear scale. The dimension of the generated Gaussian vector was 50 and we did 50 repetitions. We present in Table 2 the results of the experiment, the value inside a cell is the mean RMSE or MAE over the 50 repetitions and the value between brackets is the standard deviation over the repetitions. Under a low level of Gaussian noise or Laplacian noise, both methods are close in term of RMSE and MAE with a little advantage for the classical isotonic regression estimator. When the level of noise is important (SNR = 10), our proposed ISVR has the lowest RMSE and MAE for the two noise distribution tested.

Real dataset. Isotonic types of constraints can be found in different applications such as biology, ranking and weather forecast for example. Focusing on global warming type of data, researchers have studied the anomaly of the average temperature over a year in comparison to the years 1961-1990. These temperature anomalies have a monotonous trend and keep increasing since 1850 until 2015. Isotonic regression estimator was used on this dataset³ in Gaines et al. (2018) and we compared our proposed ISVR estimator for anomaly prediction. The hyperparameter for the ISVR were set manually for this simulation. Figure 3 shows the result for the two estimators. The classical isotonic regression estimator perform better than our proposed estimator globally which is confirmed by the RMSE and MAE values of $\text{RMSE}_{\text{IR}} = 0.0067$

³ This dataset can be downloaded from the Carbon Dioxide Information Analysis Center <https://cdiac.ess-dive.lbl.gov/trends/temp/jonescru/jones.html> at the Oak Ridge National Laboratory.

Fig. 3 Global warming dataset. Annual temperature anomalies relative to 1961–1990 average, with estimated trend using Isotonic Support Vector Regression (ISVR) and the classical Isotonic Regression (IR) estimator

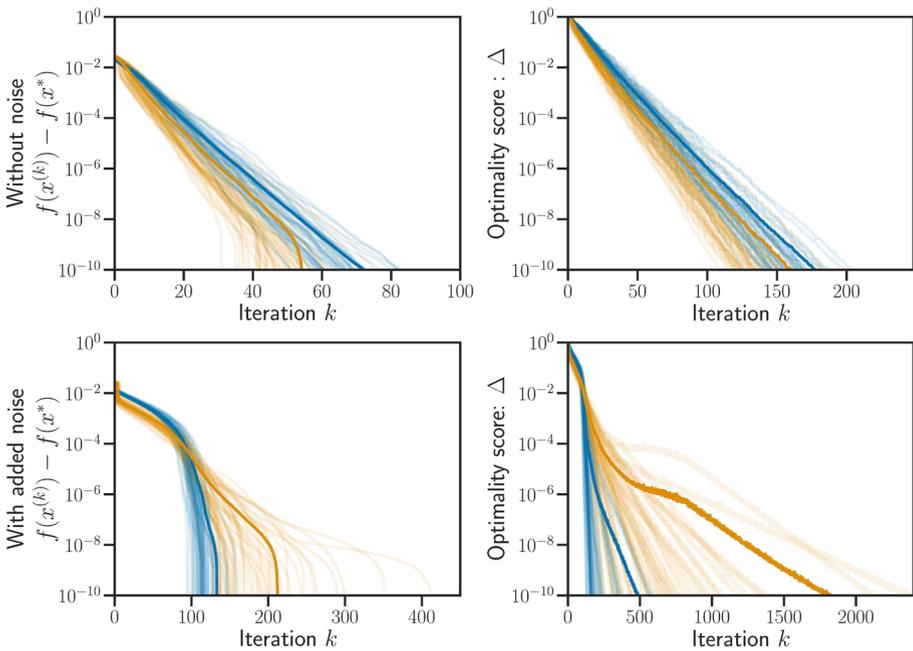
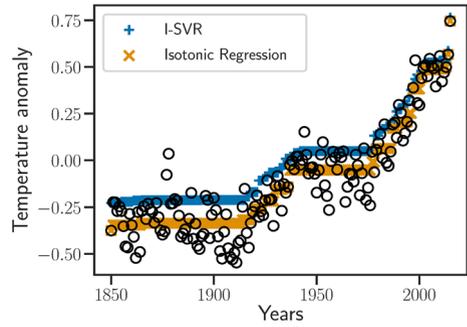


Fig. 4 Plots of 50 trajectories of the dual objective function value (first column) and the optimality score (second column) in function of the number of iterations for the classical SMO algorithm in blue and the proposed generalized SMO in red. Two settings were used, one without noise and another one with additive Gaussian noise

against $RMSE_{ISVR} = 0.083$ and $MAE_{IR} = 0.083$ against $MAE_{ISVR} = 0.116$. However, we notice that in the portions where there is a significant change like between 1910–1940 and 1980–2005, the IR estimation looks like a step function whereas the ISVR estimation follows an increasing trend without these piecewise constant portions. The bias induced by the use of constraints can be overcome with refitting methods such as Deledalle et al. (2017).

4.4 Performance of the GSMO versus SMO

We compared the efficiency of the SMO algorithm to solve the classical SVR optimization problem and the SSVR optimization problem in Fig. 4. To do so, we used the same data simulation process described earlier in this subsection and set the number of rows of the matrix X , $n = 200$ and the number of columns $p = 25$. Two different settings were considered here, one without any noise in the data and another one with Gaussian noise added such that the SNR would be equal to 30. The transparent trajectories represent the decrease of the objective function or the optimality score Δ for the classical SMO in blue and for the generalized SMO in red for the 50 repetitions considered. The average trajectory is represented in dense color. The first row of figures are the results for the noiseless setting and the second row for the setting with noise. When there is not noise in the data, the generalized SMO decreases faster than the classical SMO. It is important to remind that the true vector here belongs to the simplex so without any noise it is not surprising that our proposed algorithm goes faster than the classical SMO. However, when noise is adding to the data, it takes more iterations for the generalized SMO to find the solution of the optimization problem. Figure 4 illustrates the convergence towards a minimum the GSMO algorithm as stated in Theorem 1.

5 Conclusion

In this paper, we studied the optimization problem related to SVR with linear constraints. We showed that for this optimization problem, strong duality holds and that the dual problem is convex. We presented a generalized SMO algorithm that solve the dual problem and we proved its convergence to a solution. This algorithm uses a coordinate descent strategy where a closed-form of the updates were defined. The proposed algorithm is easy to implement and shows good performance in practise. We demonstrated the good performance of our proposed estimator on different regression settings. In presence of high level of noise, our estimator has shown to be robust and has better estimation performance in comparison to Least Squares based estimators or projected SVR estimators.

This work leaves several open questions for future works. The question of the convergence rate of the algorithm is very natural and will have to be address in the future. Another natural question rises about the possibility to extend our method on non-linear function estimation with linear constraints. From our point of view, it is a very challenging question because the dual optimization problem of the linearly constrained SVR loses its only dependence on the inner product between the columns of X , crossed terms appear in the objective function which makes it difficult to use the kernel trick as it would naturally be used for classical SVR.

Appendix A proof of Proposition 2

Proof To prove part 1., let's recall that α_i, α_i^* are the lagrange multipliers associated to the optimization problem (LSVR-P) constraints:

$$\begin{aligned} \beta^T X_i: + \beta_0 - y_i &\leq \epsilon + \xi_i \\ y_i - \beta^T X_i: - \beta_0 &\leq \epsilon + \xi_i^* \end{aligned} \tag{2}$$

The complementary optimality conditions leads to

$$\begin{aligned} \alpha_i(\beta^T X_i: + \beta_0 - y_i - \epsilon - \xi_i) &= 0 \\ \alpha_i^*(y_i - \beta^T X_i: - \beta_0 - \epsilon - \xi_i^*) &= 0. \end{aligned}$$

Let us now suppose that $\alpha_i > 0$ and $\alpha_i^* > 0$ which implies that

$$\begin{aligned} \beta^T X_i: + \beta_0 - y_i - \epsilon - \xi_i &= 0 \\ y_i - \beta^T X_i: - \beta_0 - \epsilon - \xi_i^* &= 0. \end{aligned}$$

It follows that $-2\epsilon = \xi_i + \xi_i^*$ and $\xi_i, \xi_i^* \geq 0$ which implies $\xi_i = \xi_i^* = \epsilon = 0$. This goes against our condition $\epsilon > 0$.

To prove part 2., we need to remind the optimality conditions deriving from (LSVR-P); $(Cv - \sum_{i=1}^l \alpha_i + \alpha_i^*)\epsilon = 0$. Thus, if $\epsilon > 0$ we have that $\sum_{i=1}^n \alpha_i + \alpha_i^* = Cv$. □

Appendix B proof of Proposition 3

We start by giving a lemma that will be useful to prove the proposition for the blocks α and α^* .

Lemma 1 *If the update between iteration k and $k + 1$ happens in the block α (or α^*) and that (i, j) is the most violating pair of variables then*

$$\begin{aligned} \nabla_{\alpha'} f(\theta^{k+1}) - \nabla_{\alpha'} f(\theta^k) &= \nabla_{\alpha'} f(\theta^k) - \nabla_{\alpha'} f(\theta^k) \\ &\quad + t^*(Q_{ii} + Q_{jj} - 2Q_{ij}) \end{aligned} \tag{3}$$

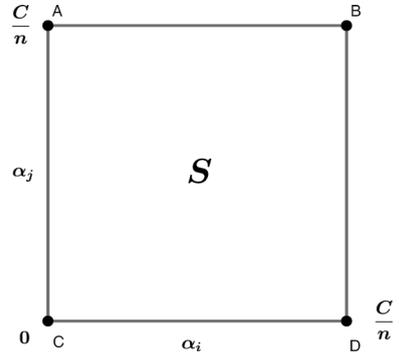
Proof Let us recall that the update in the block α (or α^*) has the following form ; $\alpha_i^{k+1} = \alpha_i^k + t^*$ and $\alpha_j^{k+1} = \alpha_j^k - t^*$, with t^* as defined in Definition 5. In a stacked form we have that

$$\begin{aligned} \nabla_{\alpha'} f(\theta^{k+1}) - \nabla_{\alpha'} f(\theta^k) &= (Q\theta^{k+1})_i + l_i - (Q\theta^{k+1})_j - l_j \\ &= \sum_{s=1}^{2n+k_1+k_2} Q_{is} \theta_s^{k+1} + l_i - \sum_{s=1}^{2n+k_1+k_2} Q_{js} \theta_s^{k+1} - l_j \\ &= \nabla_{\alpha'} f(\theta^k) - \nabla_{\alpha'} f(\theta^k) + t^*(Q_{ii} - Q_{ij}) + t^*(Q_{jj} - Q_{ij}) \\ &= \nabla_{\alpha'} f(\theta^k) - \nabla_{\alpha'} f(\theta^k) + t^*(Q_{ii} + Q_{jj} - 2Q_{ij}). \end{aligned}$$

□

This lemma is helpful for the proof the blocks γ and μ .

Fig. 5 Possible updates for the block α or α^*



Lemma 2 *If θ_i is the updated variable at iteration k , then it holds that:*
 $\nabla_{if}(\theta^{k+1}) = \bar{Q}_{ii}(\theta_i^{k+1} - \theta_i^k) + \nabla_{if}(\theta^k).$

Proof The proof is straightforward,

$$\begin{aligned} \nabla_{if}(\theta^{k+1}) &= (\bar{Q}\theta_i^{k+1})_i + l_i \\ &= \sum_{s=1}^{2n+k_1+k_2} \bar{Q}_{is}\theta_s^{k+1} + l_i \\ &= \sum_{s \neq i}^{2n+k_1+k_2} \bar{Q}_{is}\theta_s^{k+1} + l_i + \bar{Q}_{ii}\theta_i^{k+1} \\ &= \nabla_{if}(\theta^k) + \bar{Q}_{ii}\theta_i^{k+1} - \bar{Q}_{ii}\theta_i^k \\ &= \bar{Q}_{ii}(\theta_i^{k+1} - \theta_i^k) + \nabla_{if}(\theta^k). \end{aligned}$$

□

Let us now give the proof of Proposition 3.

Proof Let us consider that the update between iteration k and $k + 1$ takes place in the block α . We will define (i, j) as the most violating pair of variables as defined in Sect. 3. From the discussion in Sect. 3.3, we know that minimizing the objective function of (LSVR-D) considering that only the parameter t is a variable leads to minimizing the following function:

$$\psi(t) = \frac{1}{2}t^2(\bar{Q}_{ii} + \bar{Q}_{jj} - 2\bar{Q}_{ij}) + t(\nabla_{\alpha_i}f(\theta^k) - \nabla_{\alpha_j}f(\theta^k)) + K \tag{4}$$

We recall that t is the parameter that will be used for the update of α_i and α_j and K is a constant term. We also have the following result from Lemma 1:

$$\begin{aligned} \nabla_{\alpha_i}f(\theta^{k+1}) - \nabla_{\alpha_j}f(\theta^{k+1}) &= \nabla_{\alpha_i}f(\theta^k) - \nabla_{\alpha_j}f(\theta^k) \\ &\quad + (\bar{Q}_{ii} + \bar{Q}_{jj} - 2\bar{Q}_{ij})t^* \end{aligned} \tag{5}$$

The minimization update takes place in the square $S = [0, \frac{C}{n}] \times [0, \frac{C}{n}]$ illustrated in Fig. 5.

At points B and C of the square S, (i, j) cannot be a τ -violating pair of variables because they belong to the same set of indices I_{up} (or I_{low}). Everywhere else, violation can take place.

- On $]CA]$, $\alpha_i = 0$ and $\alpha_j > 0$ so $i \in I_{up}$ and $j \in I_{low}$ which means that by definition of τ -violating pair of variable

$$\nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k) < -\tau < 0$$

which means $t_q = \frac{-(\nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k))}{(\bar{Q}_{ii} + \bar{Q}_{jj} - 2\bar{Q}_{ij})} > 0$. Let us remind that :

$$\max\left(-\alpha_i, \alpha_j - \frac{C}{n}\right) \leq t^* \leq \min\left(\frac{C}{n} - \alpha_i, \alpha_j\right) \tag{6}$$

It means that on $]CA]$, (6) becomes : $0 \leq t^* \leq \alpha_j$ There are then two possibilities:

- if $t_q \geq \alpha_j$, it implies because of the constraints on t^* , that $t^* = \alpha_j$. The update becomes then $\alpha_i^{k+1} = \alpha_i^k + \alpha_j^k$ and $\alpha_j^{k+1} = 0$. Then j belongs to the set of indices I_{up} and i belongs to I_{low} . We deduce that $\nabla_{\alpha_i} f(\theta^{k+1}) - \nabla_{\alpha_j} f(\theta^{k+1}) \leq 0$ which proves that (i, j) is not a violating pair of variable anymore and that $\alpha^{k+1} \neq \alpha^k$
- Second possibility is that $t_q \leq \alpha_j$ then $t^* = t_q$, then $(\alpha_i^{k+1}, \alpha_j^{k+1})$ belongs to $\text{int}(S)$. We deduce that $\nabla_{\alpha_i} f(\theta^{k+1}) - \nabla_{\alpha_j} f(\theta^{k+1}) = 0$, (i, j) is not a τ -violating pair of variables anymore and $\alpha^{k+1} \neq \alpha^k$.

The same reasoning can be done on each segment of the edge of the square and also for points that are inside it. Moreover, it stays true for the block α^* and the proof is similar.

Let us now prove that when the update takes place at index i in the block γ then i is not violating variable at iteration $k + 1$. Then we need to show that $\nabla_{\gamma_i} f(\theta^{k+1}) \geq 0$. Let us start with the case where the update $\gamma_i^{k+1} = \frac{\nabla_{\gamma_i} f(\theta^k)}{\bar{Q}_{ii}} - \gamma_i^k$. Using (), we have that $\nabla_{\gamma_i} f(\theta^{k+1}) = 0$.

The second possible case is $\gamma_i^{k+1} = 0$ because $-\frac{\nabla_{\gamma_i} f(\theta^k)}{\bar{Q}_{ii}} + \gamma_i^k \leq 0$. If $\gamma_i^{k+1} = 0$ then $\nabla_{\gamma_i} f(\theta^{k+1}) = -\bar{Q}_{ii}\gamma_i^k + \nabla_{\gamma_i} f(\theta^k)$. \bar{Q}_{ii} is positive because it is a diagonal element of a Gram matrix $(A^T A)$ thus we get that $\nabla_{\gamma_i} f(\theta^{k+1}) \geq 0$, which proves that i is not a violating variable anymore.

The proof for the block μ relies on the same idea except that it is simpler because there is no clipped updates possible so $\nabla_{\mu_i} f(\theta^{k+1}) = 0$ if the updates takes place at μ_i which also proves that i is not a violating variable for this block of variables anymore. \square

Appendix C proof of Theorem 1

We begin the proof of the theorem by giving several preliminary results that will be helpful for giving the final proof. The first result gives a bound for controlling the distance of the primal iterates generated by the algorithm and the solution of (LSVR-P).

Lemma 3 For any SMO-LSSVR iterate

$$\beta^k = - \sum_{i=1}^n (\alpha_i^k - (\alpha_i^*)^k) X_{i:} - A^T \gamma^k + \Gamma^T \mu^k,$$

β^{opt} a solution of (LSVR-P) and θ^{opt} a solution of (LSVR-D), it holds that $\frac{1}{2} \|\beta^k - \beta^{\text{opt}}\| \leq f(\theta^k) - f(\theta^{\text{opt}})$.

Proof A first observation is that the relationship between the primal optimization problem and the dual leads to this equality

$$f(\theta^k) = \frac{1}{2} \|\beta^k\|^2 + l^T \theta^k. \tag{7}$$

Replacing β^k by $-\sum_{i=1}^n (\alpha_i^k - (\alpha_i^*)^k) X_{i:} - A^T \gamma^k + \Gamma^T \mu^k$ leads to (7). We have already seen that there is strong duality between both problems so the dual gap is zero at the solutions. Thus it means that for any primal optimal solution $(\beta^{\text{opt}}, \beta_0^{\text{opt}}, \xi^{\text{opt}}, \xi_i^{\text{opt}}, \epsilon^{\text{opt}})$ and any dual solution θ^{opt} , it holds true that

$$\begin{aligned} \frac{1}{2} \|\beta^{\text{opt}}\|^2 + C \left(v \epsilon^{\text{opt}} + \frac{1}{n} \sum_{i=1}^n \xi_i^{\text{opt}} + \xi_i^{\text{opt}} \right) &= -f(\theta^{\text{opt}}) \\ &= -\frac{1}{2} \|\beta^{\text{opt}}\|^2 - l^T \theta^{\text{opt}}. \end{aligned} \tag{8}$$

Using the equation link between primal and dual yields to

$$\begin{aligned} \langle \beta, \beta^{\text{opt}} \rangle &= \left\langle - \sum_{i=1}^n (\alpha_i - \alpha_i^*) X_{i:} - A^T \gamma + \Gamma^T \mu, \beta^{\text{opt}} \right\rangle \\ &= -\langle A^T \gamma, \beta^{\text{opt}} \rangle - \sum_{i=1}^n (\alpha_i - \alpha_i^*) \langle X_{i:}, \beta^{\text{opt}} \rangle \\ &\quad + \langle \Gamma^T \mu, \beta^{\text{opt}} \rangle. \end{aligned}$$

Since $\sum_{i=1}^n (\alpha_i - \alpha_i^*) = 0$, we have that

$$\begin{aligned} \langle \beta, \beta^{\text{opt}} \rangle &= -\langle A^T \gamma, \beta^{\text{opt}} \rangle - \sum_{i=1}^n (\alpha_i - \alpha_i^*) \langle X_{i:}, \beta^{\text{opt}} \rangle \\ &\quad + \langle \Gamma^T \mu, \beta^{\text{opt}} \rangle - \beta_0^{\text{opt}} \sum_{i=1}^n (\alpha_i - \alpha_i^*) \\ &= -\langle A^T \gamma, \beta^{\text{opt}} \rangle - \sum_{i=1}^n \alpha_i \langle X_{i:}, \beta^{\text{opt}} \rangle + \beta_0^{\text{opt}} \\ &\quad + \sum_{i=1}^n \alpha_i^* \langle X_{i:}, \beta^{\text{opt}} \rangle + \beta_0^{\text{opt}} + \langle \Gamma^T \mu, \beta^{\text{opt}} \rangle. \end{aligned}$$

Moreover, using the constraints of (LSVR-P) and the fact that $\alpha \geq 0$ and $\alpha^* \geq 0$ it holds that:

$$\begin{aligned}
 \langle \beta, \beta^{\text{opt}} \rangle &\geq -\langle A^T \gamma, \beta^{\text{opt}} \rangle + \sum_{i=1}^n \alpha_i (-y_i - \epsilon^{\text{opt}} - \xi_i^{\text{opt}}) \\
 &\quad + \sum_{i=1}^n \alpha_i^* (y_i - \epsilon^{\text{opt}} - (\xi_i^*)^{\text{opt}}) + \langle \Gamma^T \mu, \beta^{\text{opt}} \rangle \\
 &= -\langle A^T \gamma, \beta^{\text{opt}} \rangle - \sum_{i=1}^n (\alpha_i - \alpha_i^*) y_i - \epsilon^{\text{opt}} C_V \\
 &\quad - \sum_{i=1}^n \alpha_i \xi_i^{\text{opt}} + \alpha_i^* (\xi_i^*)^{\text{opt}} + \langle \Gamma^T \mu, \beta^{\text{opt}} \rangle.
 \end{aligned}$$

Finally we have

$$\begin{aligned}
 \frac{1}{2} \|\beta^k - \beta^{\text{opt}}\|^2 &= \frac{1}{2} \|\beta^k\|^2 - \langle \beta^k, \beta^{\text{opt}} \rangle + \frac{1}{2} \|\beta^{\text{opt}}\|^2 \\
 &\leq \frac{1}{2} \|\beta^k\|^2 + \langle A^T \gamma^k, \beta^{\text{opt}} \rangle + \epsilon^{\text{opt}} C_V \\
 &\quad + \sum_{i=1}^n (\alpha_i^k - (\alpha_i^*)^k) y_i + \sum_{i=1}^n \alpha_i^k \xi_i^{\text{opt}} \\
 &\quad + (\alpha_i^*)^k (\xi_i^*)^{\text{opt}} - \langle \Gamma^T \mu^k, \beta^{\text{opt}} \rangle + \frac{1}{2} \|\beta^{\text{opt}}\|^2
 \end{aligned}$$

Since β^{opt} satisfies the constraints of the primal optimization problem, it holds that $\langle \Gamma^T \mu, \beta^{\text{opt}} \rangle = \mu^T d$ and since $\gamma \geq 0$ we have $\langle A^T \gamma, \beta^{\text{opt}} \rangle \leq \gamma^T b$, thus

$$\begin{aligned}
 \frac{1}{2} \|\beta^k - \beta^{\text{opt}}\|^2 &\leq \frac{1}{2} \|\beta^k\|^2 + \gamma^T b + \sum_{i=1}^n (\alpha_i^k - (\alpha_i^*)^k) y_i \\
 &\quad + \epsilon^{\text{opt}} C_V + \sum_{i=1}^n \alpha_i^k \xi_i^{\text{opt}} \\
 &\quad + (\alpha_i^*)^k (\xi_i^*)^{\text{opt}} - \mu^T d + \frac{1}{2} \|\beta^{\text{opt}}\|^2.
 \end{aligned}$$

The linear term that we wrote l in the objective function of (LSVR-D) defines $l^T \theta = \sum_{i=1}^n (\alpha_i - \alpha_i^*) X_i + \gamma^T b - \mu^T d$ which in combination with the equality (7) gives

$$\begin{aligned}
 \frac{1}{2} \|\beta^k - \beta^{\text{opt}}\|^2 &\leq \frac{1}{2} f(\theta^k) + \epsilon^{\text{opt}} C_V \\
 &\quad + \sum_{i=1}^n \alpha_i^k \xi_i^{\text{opt}} + (\alpha_i^*)^k (\xi_i^*)^{\text{opt}} + \frac{1}{2} \|\beta^{\text{opt}}\|^2.
 \end{aligned}$$

Each $\alpha_i^k, (\alpha_i^*)^k$ is bounded by $\frac{C}{n}$ which yields to

$$\begin{aligned}
 \frac{1}{2} \|\beta^k - \beta^{\text{opt}}\|^2 &\leq f(\theta^k) + \epsilon^{\text{opt}} C_V \\
 &\quad + \frac{C}{n} \sum_{i=1}^n \xi_i^{\text{opt}} + (\xi_i^*)^{\text{opt}} + \frac{1}{2} \|\beta^{\text{opt}}\|^2.
 \end{aligned}$$

We recognize the objective function of the primal optimization problem and using that there is no dual gap at the optimum it follows that

$$\epsilon^{\text{opt}} C_V + \frac{C}{n} \sum_{i=1}^n \xi_i^{\text{opt}} + (\xi_i^*)^{\text{opt}} + \frac{1}{2} \|\beta^{\text{opt}}\|^2 = -f(\theta^{\text{opt}}),$$

which finishes the proof. □

Before the next statement, we need to give a definition that we will use in the next proofs.

Definition 6 Let (i, j) ($i \in I_{\text{low}}$ and $j \in I_{\text{up}}$) be the most violating pair of variables in the block α , (i^*, j^*) ($i^* \in I_{\text{low}}^*$ and $j^* \in I_{\text{up}}^*$) for the block α^* . Let s_1 be the index of the most violating variable in the block γ and s_2 in the block μ . We will call “optimality score” at iteration k the quantity $\Delta^k = \max(\Delta_1^k, \Delta_2^k, \Delta_3^k, \Delta_4^k)$, where $\Delta_1^k = \max(\nabla_{\alpha_\gamma} f(\theta^k) - \nabla_{\alpha_\gamma} f(\theta^k), 0)$, $\Delta_2^k = \max(\nabla_{\alpha_\mu} f(\theta^k) - \nabla_{\alpha_\mu} f(\theta^k), 0)$, $\Delta_3^k = \max(-\nabla_{\gamma_{s_1}} f(\theta^k), 0)$ and $\Delta_4^k = \max(|\nabla_{\mu_{s_2}} f(\theta^k)|, 0)$.

The next result states that the sequence $\{f(\theta^k)\}$ is a decreasing sequence. This result already states the convergence to a certain value \bar{f} because we know that the sequence is bounded by the existing global minimum of the function since f is convex.

Lemma 4 *The sequence generated by the Generalized SMO algorithm $\{f(\theta^k)\}$ is a decreasing sequence. This sequence converges to a value \bar{f} .*

Proof We first prove that $f(\theta^k) - f(\theta^{k+1}) \geq 0$ when minimization takes place in the block α . Let (i, j) be the indices of the variables selected to be optimized and let $u \in \mathbb{R}^{2n+k_1+k_2}$ be the vector with only zeros except at the i^{th} coordinate where it is equal to t^* as defined in Sect. 5 and at the j^{th} coordinate where it is equal to $-t^*$. We will also define $t_q = \frac{-(\nabla_{\alpha_\gamma} f(\theta^k) - \nabla_{\alpha_\mu} f(\theta^k))}{Q_{ii} + Q_{jj} - 2Q_{ij}}$, the unconstrained minimum for the update in α block. Let us compute

$$\begin{aligned} f(\theta^k) - f(\theta^{k+1}) &= \frac{1}{2} (\theta^k)^T \bar{Q} \theta^k + l^T \theta^k \\ &\quad - \frac{1}{2} (\theta^{k+1})^T \bar{Q} \theta^{k+1} + l^T \theta^{k+1} \\ &= \frac{1}{2} (\theta^k)^T \bar{Q} \theta^k + l^T \theta^k \\ &\quad - \frac{1}{2} (\theta^k + U)^T \bar{Q} (\theta^k + U) \\ &\quad + l^T (\theta^k + U) \\ &= -\frac{1}{2} u^T \bar{Q} u - u^T (Q \theta^k + l) \\ &= -\frac{1}{2} u^T \bar{Q} u - u^T (\nabla f(\theta^k)) \\ &= -\frac{(t^*)^2}{2} (Q_{ii} + Q_{jj} - 2Q_{ij}) \\ &\quad - t^* (\nabla_{\alpha_\gamma} f(\theta^k) - \nabla_{\alpha_\mu} f(\theta^k)). \end{aligned}$$

We first study the case when there is no clipping which means that $t^* = t_q$

1. No clipping. Replacing t^* by its expression leads to the following result:

$$\begin{aligned} f(\theta^k) - f(\theta^{k+1}) &= \frac{(\Delta_1^k)^2}{2(Q_{ii} + Q_{jj} - 2Q_{ij})} \\ &= \frac{(\Delta_1^k)^2}{2\|X_{i:} - X_{j:}\|^2} \geq 0. \end{aligned}$$

2. Clipping takes place because $t_q \leq t^* = \max(-\alpha_i, \alpha_j - \frac{C}{n})$

We notice that $t_q \leq \max(-\alpha_i, \alpha_j - \frac{C}{n}) \leq 0$ which implies that $i \in I_{low}$ and $j \in I_{up}$. In that case $\Delta_1^k = \nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k)$. Replacing t_q by its expression leads to

$$\begin{aligned} -(\nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k)) &\leq t^*(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ \frac{\Delta_1^k t^*}{2} &\leq \frac{-(t^*)^2}{2}(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ \frac{\Delta_1^k t^*}{2} - t^* \Delta_1^k &\leq \frac{-(t^*)^2}{2}(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ &\quad - t^*(\nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k)) \\ -\frac{1}{2} \Delta_1^k t^* &\leq \frac{-(t^*)^2}{2}(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ &\quad - t^*(\nabla_{\alpha_i} f(\theta^k) - \nabla_{\alpha_j} f(\theta^k)) \end{aligned}$$

Thus we have that if $t^* = -\alpha_i$, $f(\theta^k) - f(\theta^{k+1}) \geq \frac{1}{2} \Delta_1^k \alpha_i \geq 0$ and that if $t^* = \alpha_j - \frac{C}{n}$, $f(\theta^k) - f(\theta^{k+1}) \geq \frac{1}{2} \Delta_1^k (\frac{C}{n} - \alpha_j) \geq 0$. **3. Clipping takes place because** $t_q \geq t^* = \min(\frac{C}{n} - \alpha_i, \alpha_j)$.

This time $t_q \geq \min(\frac{C}{n} - \alpha_i, \alpha_j) \geq 0$ which also implies that $i \in I_{up}$ and $j \in I_{low}$ and that $\Delta_1^k = \nabla_{\alpha_j} f(\theta^k) - \nabla_{\alpha_i} f(\theta^k)$. The only difference here is that multiplying by $-t^*$ will imply a change in the inequality.

$$\begin{aligned} -(\nabla_j f(\theta^k) - \nabla_i f(\theta^k)) &\geq t^*(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ \frac{-\Delta_1^k t^*}{2} &\leq \frac{-(t^*)^2}{2}(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ \frac{-\Delta_1^k t^*}{2} + t^* \Delta_1^k &\leq \frac{-(t^*)^2}{2}(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ &\quad - t^*(\nabla_j f(\theta^k) - \nabla_i f(\theta^k)) \\ \frac{1}{2} \Delta_1^k t^* &\leq \frac{-(t^*)^2}{2}(Q_{ii} + Q_{jj} - 2Q_{ij}) \\ &\quad - t^*(\nabla_j f(\theta^k) - \nabla_i f(\theta^k)) \end{aligned}$$

Thus we have that if $t^* = \frac{C}{n} - \alpha_i$

$$f(\theta^k) - f(\theta^{k+1}) \geq \frac{1}{2} \Delta_1^k \left(\frac{C}{n} - \alpha_i \right) \geq 0,$$

and if $t^* = \alpha_j$,

$$f(\theta^k) - f(\theta^{k+1}) \geq \frac{1}{2} A_1^k \alpha_j \geq 0.$$

To prove that $f(\theta^k) - f(\theta^{k+1}) \geq 0$ when the update takes place in the block γ and μ we first need to observe that when only one variable is updated between iteration k and $k + 1$ it follows that

$$f(\theta^k) - f(\theta^{k+1}) = \frac{1}{2} \bar{Q}_{ii} (\theta_i^k - \theta_i^{k+1})^2.$$

Therefore, we now prove the result for the block γ . If the update is not a clipped update and i is the index of the updated variable, it holds that $\gamma_i^k - \gamma_i^{k+1} = \frac{\nabla_{\gamma_i} f(\theta^k)}{(AA^T)_{ii}}$, which gives the following bound

$$f(\theta^k) - f(\theta^{k+1}) = \frac{1}{2(AA^T)_{ii}} (\nabla_{\gamma_i} f(\theta^k))^2 \geq 0. \tag{9}$$

Moreover, if a clipped update takes place in this block, we know that it happens when $0 \leq \gamma_i^k \leq \frac{\nabla_{\gamma_i} f(\theta^k)}{(AA^T)_{ii}}$. It yields to the following bound

$$f(\theta^k) - f(\theta^{k+1}) = \frac{1}{2} (AA^T)_{ii} (\gamma_i^k)^2 \geq 0.$$

The result for the block μ is obtained using the same arguments except that there is no clipped updates. □

Lemma 5 *There exists a subsequence $\{\theta^{k_j}\}$ of iterations generated by the generalized SMO where clipping does not take place.*

Proof Let us suppose the contrary, which means that there exists an iteration K such that for all $k \geq K$ we only perform clipped updates. The number of variables N_B^k that belong to the boundary of its constraints (0 or $\frac{C}{n}$ for the blocks α or α^* and 0 for the block γ) is non-decreasing for all $k \geq K$ and it is bounded thus it must converge to another integer N^* .

This convergence implies that there exists k^* such that for all $k \geq k^*$, $N_B^k = N^*$ since N_B^k and N^* are integers. This observation allows us to conclude that for all $k \geq k^*$ clipped updates only take place in the blocks α or α^* since the updates in the block γ are made on only one variable and that the number of clipped variables has reached its maximum value. An update in the block γ would strictly increase the number of clipped variables which is not possible for all $k \geq k^*$ or the update would not change the value of θ and we showed before that this situation is not possible (Proposition 3).

For all $k \geq k^*$, we have that updates in the block α (resp. α^*) have this necessary scheme: α_i^k or α_j^k is equal to 0 or $\frac{C}{n}$ thus after the update, one of them will leave the boundary and the other one goes to it in order to keep the number of clipped variables equals to N^* . The different possibilities are then the following:

- if $\alpha_i^k = 0$ and $0 < \alpha_j^k \leq \frac{C}{l}$ the only possible update following the Definition 5 is

$$\begin{aligned} \alpha_i^{k+1} &= \alpha_i^k + \alpha_j^k = \alpha_j^k \\ \alpha_j^{k+1} &= \alpha_j^k - \alpha_j^k = 0. \end{aligned}$$

- if $\alpha_j^k = \frac{C}{l}$ and $0 \leq \alpha_i^k < \frac{C}{l}$ the only possible update following the Definition 5 is

$$\alpha_i^{k+1} = \alpha_i^k + \left(\frac{C}{l} - \alpha_i^k\right) = \frac{C}{l}$$

$$\alpha_j^{k+1} = \alpha_j^k - \left(\frac{C}{l} - \alpha_j^k\right) = \alpha_j^k.$$

It stays true for the block α^* and the discussion is similar. It is clear that from the description of the updates made above that there is only a finite number of ways to shuffle the values which means that there exists $k_1, k_2 \geq k^*$ such as $\theta^{k_1} = \theta^{k_2}$ and with $k_1 < k_2$. Therefore $f(\theta^{k_1}) = f(\theta^{k_2})$ which contradicts the decrease of the sequence $f(\theta^k)$ (Lemma 4). \square

Lemma 6 *Let $\{\theta^{k_j}\}$ be a subsequence generated by the Generalized SMO algorithm where clipping does not take place. We then have that $\Delta^{k_j} \rightarrow 0$.*

Proof We have that

$$f(\theta^{k_j}) - f(\theta^{k_{j+1}}) \geq \frac{(\nabla_{\alpha_p} f(\theta^{k_j}) - \nabla_{\alpha_q} f(\theta^{k_j}))^2}{2D^2} = \frac{(\Delta^{k_j})^2}{2D^2}$$

where $D = \max_{p,q} \|X_p - X_q\|$ when the update happens in the blocks α or α^* . When it happens in the block γ with no clipping we have the following inequality

$$f(\theta^{k_j}) - f(\theta^{k_{j+1}}) \geq \frac{(\nabla_{\gamma_i} f(\theta^{k_j}))^2}{2} = \frac{(-\Delta^{k_j})^2}{2} = \frac{(\Delta^{k_j})^2}{2}.$$

When the update takes place in the block μ , we have that

$$f(\theta^{k_j}) - f(\theta^{k_{j+1}}) \geq \frac{(\nabla_{\mu_i} f(\theta^{k_j}))^2}{2} = \frac{(\Delta^{k_j})^2}{2}.$$

We then define a sequence

$$u^{k_j} = \begin{cases} \frac{1}{2D^2}(\Delta^{k_j})^2 & \text{update in the blocks } \alpha \text{ or } \alpha^*. \\ \frac{1}{2}(\Delta^{k_j})^2 & \text{update in the blocks } \gamma \text{ or } \mu. \end{cases}$$

The sequence $\{u^{k_j}\} \rightarrow 0$ because of the bound given above and the fact that $f(\theta^{k_j}) - f(\theta^{k_{j+1}}) \rightarrow 0$ too (Lemma 4). This implies that $\Delta^{k_j} \rightarrow 0$ as well. \square

A consequence of the lemma above is that $\Delta_1^{k_j} \rightarrow 0, \Delta_2^{k_j} \rightarrow 0, \Delta_3^{k_j} \rightarrow 0$ and $\Delta_4^{k_j} \rightarrow 0$ because Δ^{k_j} is defined as the maximum of those four positive values.

Lemma 7 *Let $\{\theta^{k_j}\}$ be a subsequence generated by the generalized SMO algorithm where clipping does not take place. This subsequence is bounded.*

Proof To prove the statement, we will show that $\|\theta^{k_j} - \theta^{\text{opt}}\|^2$ is bounded where θ^{opt} belongs to the set of solution of (LSVR-D). Since each α_i and α_i^* is belongs to $[0, \frac{C}{n}]$, we have that

$$\begin{aligned} \|\theta^{k_j+1} - \theta^{\text{opt}}\|^2 &= \|\alpha^{k_j+1} - \alpha^{\text{opt}}\|^2 + \|(\alpha^*)^{k_j+1} - (\alpha^*)^{\text{opt}}\|^2 \\ &\quad + \|\gamma^{k_j+1} - \gamma^{\text{opt}}\|^2 + \|\mu^{k_j+1} - \mu^{\text{opt}}\|^2 \\ &\leq \frac{2C^2}{n} + \|\gamma^{k_j+1} - \gamma^{\text{opt}}\|^2 + \|\mu^{k_j+1} - \mu^{\text{opt}}\|^2. \end{aligned}$$

We will work on the bound for the quantity $\|\mu^{k_j+1} - \mu^{\text{opt}}\|^2$ first. If the update happens in the block μ at coordinate μ_j , we have the following

$$\begin{aligned} \|\mu^{k_j+1} - \mu^{\text{opt}}\|^2 &= \|\mu^{k_j} - e_j \frac{\nabla_{\mu_j} f(\theta^{k_j})}{(\Gamma \Gamma^T)_{jj}} - \mu^{\text{opt}}\|^2 \\ &= \|\mu^{k_j} - \mu^{\text{opt}}\|^2 - 2 \left\langle \mu^{k_j} - \mu^{\text{opt}}, e_j \frac{\nabla_{\mu_j} f(\theta^{k_j})}{(\Gamma \Gamma^T)_{jj}} \right\rangle \\ &\quad + \left\| e_j \frac{\nabla_{\mu_j} f(\theta^{k_j})}{(\Gamma \Gamma^T)_{jj}} \right\|^2 \\ &= \|\mu^{k_j} - \mu^{\text{opt}}\|^2 + \frac{\nabla_{\mu_j} f(\theta^{k_j})^2}{(\Gamma \Gamma^T)_{jj}^2} \\ &\quad - 2 \frac{\nabla_{\mu_j} f(\theta^{k_j})}{(\Gamma \Gamma^T)_{jj}} (\mu_j^{k_j} - \mu_j^{\text{opt}}). \end{aligned}$$

We then have that

$$\begin{aligned} -2 \frac{\nabla_{\mu_j} f(\theta^{k_j})}{(\Gamma \Gamma^T)_{jj}} (\mu_j^{k_j} - \mu_j^{\text{opt}}) &= 2(\mu_j^{k_j+1} - \mu_j^{k_j})(\mu_j^{k_j} - \mu_j^{\text{opt}}) \\ &= 2\langle \mu^{k_j+1} - \mu^{k_j}, \mu^{k_j} - \mu^{\text{opt}} \rangle \\ &\leq 2\|\mu^{k_j+1} - \mu^{k_j}\| \cdot \|\mu^{k_j} - \mu^{\text{opt}}\| \\ &\leq 2 \frac{|\nabla_{\mu_j} f(\theta^{k_j})|}{(\Gamma \Gamma^T)_{jj}} \|\mu^{k_j} - \mu^{\text{opt}}\| \\ &\leq 2 \frac{\Delta_4^{k_j}}{(\Gamma \Gamma^T)_{jj}} \|\mu^{k_j} - \mu^{\text{opt}}\| \end{aligned}$$

From 6, we have that $\Delta_4^{k_j} \rightarrow 0$ then it can be bounded by a constant M_0 . We know from (9) that

$$\frac{\nabla_{\mu_j} f(\theta^{k_j})^2}{(\Gamma \Gamma^T)_{jj}^2} = \frac{2}{(\Gamma \Gamma^T)_{jj}} (f(\theta^{k_j}) - f(\theta^{k_j+1})).$$

From Lemma 4, we know that $f(\theta^{k_j}) - f(\theta^{k_j+1}) \rightarrow 0$ then it can be bounded by a constant M_1 . Overall we have that

$$\begin{aligned} \|\mu^{k_j+1} - \mu^{\text{opt}}\|^2 &\leq \|\mu^{k_j} - \mu^{\text{opt}}\|^2 + 2\frac{M_0}{(\Gamma\Gamma^T)_{jj}}\|\mu^{k_j} - \mu^{\text{opt}}\| \\ &\quad + \frac{2}{(\Gamma\Gamma^T)_{jj}}M_1. \end{aligned}$$

By recursion we have

$$\begin{aligned} \|\mu^{k_j+1} - \mu^{\text{opt}}\|^2 &\leq \|\mu^0 - \mu^{\text{opt}}\|^2 + 2\frac{M_0}{(\Gamma\Gamma^T)_{jj}}\|\mu^0 - \mu^{\text{opt}}\| \\ &\quad + \frac{2}{(\Gamma\Gamma^T)_{jj}}M_1 \\ &< \infty. \end{aligned}$$

Since there is no clipped update on the subsequence $\{\theta^{k_j}\}$, the proof for the block γ is similar which proves that $\|\theta^{k_j} - \theta^{\text{opt}}\|$ is bounded. \square

Lemma 8 *Let $\{\theta^{k_j}\}$ be a subsequence generated by the generalized SMO algorithm where clipping does not take place. There exists a sub-subsequence that converges to $\bar{\theta}$, with $\bar{\theta}$ being a solution of (LSVR-D).*

Proof From Lemma 7, we have that $\{\theta^{k_j}\}$ is a bounded sequence, it means that we can extract a converging subsequence that we will write $\{\theta^{k_j}\}$ not to complicate the notations. Since \mathcal{F} is closed, $\bar{\theta}$ meets the constraints of the dual optimization problem and belongs to \mathcal{F} . We now want to prove that it belongs to the set of solution of (LSVR-D) by showing that $\bar{\Delta}_1(\bar{\theta}) \leq 0, \bar{\Delta}_2(\bar{\theta}) \leq 0, \bar{\Delta}_3(\bar{\theta}) \leq 0$ and $\bar{\Delta}_4(\bar{\theta}) \leq 0$. Let us make two observations that will be used for the following proof. The first one comes from the continuity of the gradient which implies that for all ϵ there exists K_1 such that for all $k_j \geq K_1, |\nabla_{\bar{i}}f(\theta^{k_j}) - \nabla_{\bar{i}}f(\bar{\theta})| < \epsilon$ for all i . The second observation is that it is possible too choose an ϵ small enough such that there exists K_2 such that for all $k_j \geq K_2$: if $\bar{\alpha}_i > 0$, we have $\alpha_i^{k_j} > 0$ and if $\bar{\alpha}_i < \frac{\epsilon}{n}$ we have $\alpha_i^{k_j} < \frac{\epsilon}{n}$. In other words, we say that all the indices in the set $I_{\text{low}}(\bar{\alpha})$ (resp. I_{up}) are also in $I_{\text{low}}(\alpha_n^{k_j})$ (resp. I_{up}). The same argument holds for indices in the block α^* .

Let us assume that $\bar{\Delta}_1 > 0$, it means that there exists at least one violating pair of variables that we will note (\bar{i}, \bar{j}) at $\bar{\theta}$. From the discussion above, we know that $\bar{i} \in I_{\text{low}}$ for all $k_j \geq K_2$ and that $\bar{j} \in I_{\text{up}}$ for all $k_j \geq K_2$. We then have that for all $\epsilon > 0$, there exists K_1 such as for all $k_j \geq \max(K_1, K_2)$,

$$\begin{aligned} \Delta_1^{k_j} &= \min_{i \in I_{\text{up}}} \nabla_{\bar{i}}f(\theta^{k_j}) - \max_{i \in I_{\text{low}}} \nabla_{\bar{i}}f(\theta^{k_j}) \\ &\geq \nabla_{\bar{j}}f(\theta^{k_j}) - \nabla_{\bar{j}}f(\theta^{k_j}) \\ &\geq (\nabla_{\bar{j}}f(\bar{\theta}) - \epsilon) - (\nabla_{\bar{j}}f(\bar{\theta}) + \epsilon) \\ &= \bar{\Delta}_1 - 2\epsilon. \end{aligned}$$

We choose $\epsilon = \frac{\bar{\Delta}_1}{2} - \epsilon'$ where $0 < \epsilon' < \frac{\bar{\Delta}_1}{2}$ which leads to

$$\Delta_1^{k_j} \geq \bar{\Delta}_1 - 2\epsilon' = 2\epsilon' > 0.$$

This inequality is true for all $k_j \geq \max(K_1, K_2)$ which contradicts the fact that $\Delta_1^{k_j} \rightarrow 0$. The proof is similar to show that $\bar{\Delta}_2 \leq 0$.

Let us now suppose that $\bar{\Delta}_3 > 0$ it means that there exists an index \bar{i} such that $\nabla_{\gamma_r} f(\bar{\theta}) < 0$. For all $\epsilon > 0$, there exists K_1, K_2 such as for all $k_j > \max(K_1, K_2)$

$$\begin{aligned} \Delta_3^{k_j} &= - \min_{i \in \{1, \dots, k_1\}} \nabla_{\gamma_i} f(\theta^{k_j}) \\ &\geq -\nabla_{\gamma_r} f(\theta^{k_j}) \\ &\geq -(\nabla_{\gamma_r} f(\bar{\theta}) + \epsilon) \\ &= \bar{\Delta}_3 - \epsilon. \end{aligned}$$

We choose $\epsilon = \bar{\Delta}_3 - \epsilon'$ where $0 < \epsilon' < \bar{\Delta}_3$ which leads to

$$\Delta_3^{k_j} \geq \bar{\Delta}_3 - \epsilon = \epsilon' > 0.$$

This inequality is true for all $k_j \geq \max(K_1, K_2)$ which contradicts the fact that $\Delta_3^{k_j} \rightarrow 0$.

Finally let's assume that $\Delta_4^{k_j} > 0$, it means that $|\nabla_{\mu_r} f(\theta^{k_j})| \neq 0$. Using the continuity of the gradient we write that for all $\epsilon > 0$ there exists K_1 such that for all $k_j \geq K_1$ we have $|\nabla_{\mu_r} f(\theta^{k_j}) - \nabla_{\mu_r} f(\bar{\theta})| < \epsilon$. Using triangle inequality we get that

$$\left| |\nabla_{\mu_r} f(\theta^{k_j})| - |\nabla_{\mu_r} f(\bar{\theta})| \right| \leq |\nabla_{\mu_r} f(\theta^{k_j}) - \nabla_{\mu_r} f(\bar{\theta})| < \epsilon.$$

Thus $-\epsilon \leq |\nabla_{\mu_r} f(\theta^{k_j})| - |\nabla_{\mu_r} f(\bar{\theta})| \leq \epsilon$, which means that

$$|\nabla_{\mu_r} f(\bar{\theta})| - \epsilon \leq |\nabla_{\mu_r} f(\theta^{k_j})|.$$

Then we have the following:

$$\begin{aligned} \Delta_4^{k_j} &= \max_{i \in \{1, \dots, k_2\}} |\nabla_{\mu_i} f(\theta^{k_j})| \\ &\geq |\nabla_{\mu_r} f(\theta^{k_j})| \\ &\geq |\nabla_{\mu_r} f(\bar{\theta})| - \epsilon \\ &= \bar{\Delta}_4 - \epsilon. \end{aligned}$$

We choose $\epsilon = \bar{\Delta}_4 - \epsilon'$ where $0 < \epsilon' < \bar{\Delta}_4$ which leads to

$$\Delta_4^{k_j} \geq \bar{\Delta}_4 - \epsilon = \epsilon' > 0.$$

This inequality is true for all $k_j \geq \max(K_1, K_2)$ which contradicts the fact that $\Delta_4^{k_j} \rightarrow 0$. □

Proof We are now able to give the proof of the Theorem 1. From Lemma 3, we have that $\frac{1}{2} \|\beta^k - \beta^{\text{opt}}\| \leq f(\theta^k) - f(\theta^{\text{opt}})$. Moreover, from Lemma 6 we know that there is a subsequence $\{\theta^{k_j}\}$ generated by the Generalized SMO algorithm where clipping does not take place and that converges to $\bar{\theta}$, with $\bar{\theta}$ a solution of (LSVR-D). The continuity of the objective function f allows us to say that $f(\theta^{k_j}) \rightarrow f(\bar{\theta})$. From Lemma 4, we know that $\{f(\theta^{k_j})\}$ is decreasing and bounded so the monotone convergence theorem implies that the whole sequence $f(\theta^k) \rightarrow f(\bar{\theta})$ and it follows that $\frac{1}{2} \|\beta^k - \beta^{\text{opt}}\| \rightarrow 0$ and finally that $\beta^k \rightarrow \beta^{\text{opt}}$. □

Acknowledgements This work was partly supported by ANR GraVa ANR-18-CE40-0005, Projet ANER RAGA G048CVCRB-2018ZZ and INSERM Plan cancer 18CP134-00.

References

- Abbas, A. R., Wolslegel, K., Seshasayee, D., Modrusan, Z., & Clark, H. F. (2009). Deconvolution of blood microarray data identifies cellular activation patterns in systemic lupus erythematosus. *PLoS ONE*, 4(7), 1–16.
- Barlow, R. E., & Brunk, H. D. (1972). The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337), 140–147.
- Boser, B. E., Guyon, I. M., & Vapnik, V. N. (1992). A training algorithm for optimal margin classifiers. In: Proceedings of the Fifth Annual Workshop on Computational Learning Theory, ACM, New York, NY, USA, COLT '92, pp 144–152.
- Bro, R., & De Jong, S. (1997). A fast non-negativity-constrained least squares algorithm. *Journal of Chemometrics*, 11(5), 393–401.
- Chang, C., & Lin, C. (2002). Training ν -support vector regression: theory and algorithms. *Neural Computation*, 14(8), 1959–1977.
- Chang, C. C., & Lin, C. J. (2011). LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2, 27:1-27:27.
- Chapelle, O., Haffner, P., & Vapnik, V. N. (1999). Support vector machines for histogram-based image classification. *IEEE Transactions on Neural Networks*, 10(5), 1055–1064.
- Chun-Hsin, Wu., Ho, Jan-Ming., & Lee, D. T. (2004). Travel-time prediction with support vector regression. *IEEE Transactions on Intelligent Transportation Systems*, 5(4), 276–281.
- Couzin-Frankel, J. (2013). Cancer immunotherapy. *Science*, 342(6165), 1432–1433.
- Deledalle, C. A., Papadakis, N., Salmon, J., & Vaiter, S. (2017). Clear: covariant least-square refitting with applications to image restoration. *SIAM Journal on Imaging Sciences*, 10(1), 243–284.
- Drucker, H., Burges, C. J. C., Kaufman, L., Smola, A. J., & Vapnik, V. (1997). Support vector regression machines. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems* (9th ed., pp. 155–161). Cambridge: MIT Press.
- Friel, T. T., & Harrison, R. (1998). Linear programming support vector machines for pattern classification and regression estimation: and the sr algorithm: improving speed and tightness of vc bounds in sv algorithms. Research report, Department of Automatic Control and Systems Engineering.
- Gaines, B. R., Kim, J., & Zhou, H. (2018). Algorithms for fitting the constrained lasso. *Journal of Computational and Graphical Statistics*, 27(4), 861–871.
- Gong, T., Hartmann, N., Kohane, I. S., Brinkmann, V., Staedtler, F., Letzkus, M., et al. (2011). Optimal deconvolution of transcriptional profiling data using quadratic programming with application to complex clinical blood samples. *PLoS ONE*, 6(11), 1–11.
- Haussler, D., Bednarski, D. W., Schummer, M., Cristianini, N., Duffy, N., & Furey, T. S. (2000). Support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics*, 16(10), 906–914.
- Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55–67.
- Hsieh, C. J., Chang, K. W., Lin, C. J., Keerthi, S. S., & Sundararajan, S. (2008). A dual coordinate descent method for large-scale linear svm. In: Proceedings of the 25th International Conference on Machine Learning, ACM, New York, NY, USA, ICML '08, pp 408–415.
- Jia, H., & Martinez, A. M. (2009). Support vector machines in face recognition with occlusions. In: Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp 136–141.
- Joachims, T. (1998). Text categorization with support vector machines: learning with many relevant features. In C. Nédellec & C. Rouveirol (Eds.), *Machine Learning: ECML-98* (pp. 137–142). Berlin, Heidelberg: Springer.
- Keerthi, S. S., & Gilbert, E. G. (2002). Convergence of a generalized smo algorithm for svm classifier design. *Machine Learning*, 46(1–3), 351–360.
- Keerthi, S. S., Shevade, S. K., Bhattacharyya, C., & Murthy, K. R. K. (2001). Improvements to platts smo algorithm for svm classifier design. *Neural Computation*, 13(3), 637–649.
- Lauer, F., & Bloch, G. (2008). Incorporating prior knowledge in support vector regression. *Machine Learning* 70.
- Lawson, C., & Hanson, R. (1995). Solving least squares problems. Society for Industrial and Applied Mathematics.

- Liew, C. K. (1976). Inequality constrained least-squares estimation. *Journal of the American Statistical Association*, 71(355), 746–751.
- Lin, C. J., Weng, R. C., & Keerthi, S. S. (2007). Trust region newton methods for large-scale logistic regression. In: Proceedings of the 24th international conference on Machine learning, ACM, pp 561–568.
- Lopez, J., & Dorronsoro, J. R. (2012). Simple proof of convergence of the smo algorithm for different svm variants. *IEEE Transactions on Neural Networks and Learning Systems*, 23(7), 1142–1147.
- Mohammadi, S., Zuckerman, N., Goldsmith, A., & Grama, A. (2017). A critical survey of deconvolution methods for separating cell types in complex tissues. *Proceedings of the IEEE*, 105(2), 340–366.
- Nesterov, Y. (2012). Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2), 341–362.
- Newman, A. M., Liu, C., Green, M. R., Gentles, A. J., Feng, W., Xu, Y., et al. (2015). Robust enumeration of cell subsets from tissue expression profiles. *Nature methods*, 12(5), 453–457.
- Nutini, J., Schmidt, M., Laradji, I. H., Friedlander, M., & Koepke, H. (2015). Coordinate descent converges faster with the gauss-southwell rule than random selection. In: Proceedings of the 32Nd International Conference on International Conference on Machine Learning, Vol 37, JMLR.org, ICML'15, pp 1632–1641.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Platt, J. (1998). Sequential minimal optimization: a fast algorithm for training support vector machines. Technical report, Microsoft Research.
- Qiao, W., Quon, G., Csaszar, E., Yu, M., Morris, Q., & Zandstra, P. W. (2012). Pert: A method for expression deconvolution of human blood samples from varied microenvironmental and developmental conditions. *PLOS Computational Biology*, 8(12), 1–14.
- Schölkopf, B., Bartlett, P., Smola, A., & Williamson, R. (1999). Shrinking the tube: a new support vector regression algorithm. In: Proceedings of the 1998 Conference on Advances in Neural Information Processing Systems II, MIT Press, Cambridge, MA, USA, pp 330–336.
- She, J. (2017). Linear convergence and support vector identification of sequential minimal optimization. In: Proceedings of the 10th NIPS Workshop on Optimization for Machine Learning, p 5.
- Smola, A. J., & Schölkopf, B. (2004). A tutorial on support vector regression. *Statistics and Computing*, 14(3), 199–222.
- Suykens, J., & Vandewalle, J. (1999). Least squares support vector machine classifiers. *Neural Processing Letters*, 9, 293–300.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3), 475–494.
- Van Gestel, T., Suykens, J. A. K., Baestaens, D., Lambrechts, A., Lanckriet, G., Vandaele, B., et al. (2001). Financial time series prediction using least squares support vector machines within the evidence framework. *IEEE Transactions on Neural Networks*, 12(4), 809–821.
- Wright, S. J. (2015). Coordinate descent algorithms. *Mathematical Programming*, 151(1), 3–34.
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320.

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