
GAP Safe Screening Rules for Sparse-Group Lasso

Eugene Ndiaye, Olivier Fercoq, Alexandre Gramfort, Joseph Salmon
LTCI, CNRS, Télécom ParisTech
Université Paris-Saclay
75013 Paris, France
first.last@telecom-paristech.fr

Abstract

For statistical learning in high dimension, sparse regularizations have proven useful to boost both computational and statistical efficiency. In some contexts, it is natural to handle more refined structures than pure sparsity, such as for instance group sparsity. Sparse-Group Lasso has recently been introduced in the context of linear regression to enforce sparsity both at the feature and at the group level. We propose the first (provably) safe screening rules for Sparse-Group Lasso, i.e., rules that allow to discard early in the solver features/groups that are inactive at optimal solution. Thanks to efficient dual gap computations relying on the geometric properties of ϵ -norm, safe screening rules for Sparse-Group Lasso lead to significant gains in term of computing time for our coordinate descent implementation.

1 Introduction

Sparsity is a critical property for the success of regression methods, especially in high dimension. Often, group (or block) sparsity is helpful when a known group structure needs to be enforced. This is for instance the case in multi-task learning [1] or multinomial logistic regression [7, Chapter 3]. In the multi-task setting, the group structure appears natural since one aims at jointly recovering signals whose supports are shared. In this context, sparsity and group sparsity are generally obtained by adding a regularization term to the data-fitting: ℓ_1 norm for sparsity and $\ell_{1,2}$ norm for group sparsity.

Along with recent works on hierarchical regularization [15, 20] have focused on a specific case: the Sparse-Group Lasso. This method is the solution of a (convex) optimization program with a regularization term that is a convex combination of the two aforementioned norms, enforcing sparsity and group sparsity at the same time.

With such advanced regularizations, the computational burden can be particularly heavy in high dimension. Yet, it can be significantly reduced if one can exploit the known sparsity of the solution in the optimization. Following the seminal paper on “safe screening rules” [11], many contributions have investigated such strategies [24, 23, 5]. These so called safe screening rules compute some tests on dual feasible points to eliminate primal variables whose coefficients are guaranteed to be zero in the exact solution. Still, the computation of a dual feasible point can be challenging when the regularization is more complex than ℓ_1 or $\ell_{1,2}$ norms. This is the case for the Sparse-Group Lasso as it is not straightforward to characterize if a dual point is feasible or not [23]. Here, we propose an efficient computation of the associated dual norm. It is all the more crucial since the naive implementation computes the Sparse-Group Lasso dual norm with a quadratic complexity w.r.t the groups dimensions.

We propose here efficient safe screening rules for the Sparse-Group Lasso that combine sequential rules (*i.e.*, rules that perform screening thanks to solutions obtained for a previously processed tuning parameter) and dynamic rules (*i.e.*, rules that perform screening as the algorithm proceeds) in a unified way. We elaborate on GAP safe rules, a strategy relying on dual gap computations introduced

for the Lasso [12] and to more general learning tasks in [18]. Note that alternative (unsafe) screening rules, for instance the “strong rules” [22], have been applied to the Lasso and its simple variants.

Our contributions are two fold here. First, we introduce the first safe screening rules for this problem, other alleged safe rules [23] for Sparse-Group Lasso were in fact not safe, as explained in detail in [18], and could lead to non-convergent implementation. Second, we link the Sparse-Group Lasso penalties to the ϵ -norm in [8]. This allows to provide a new algorithm to efficiently compute the required dual norms, adapting an algorithm introduced in [9]. We incorporate our proposed GAP Safe rules in a block coordinate descent algorithm and show its practical efficiency in climate prediction tasks. Another strategy leveraging dual gap computations and active sets has recently been proposed under the name *Blitz* [16]. It could naturally benefit from our fast dual norm evaluations in this context.

Notation For any integer $d \in \mathbb{N}$, we denote by $[d]$ the set $\{1, \dots, d\}$. The standard Euclidean norm is written $\|\cdot\|$, the ℓ_1 norm $\|\cdot\|_1$, the ℓ_∞ norm $\|\cdot\|_\infty$, and the transpose of a matrix Q is denoted by Q^\top . We also denote $(t)_+ = \max(0, t)$. Our observation vector is $y \in \mathbb{R}^n$ and the design matrix $X = [X_1, \dots, X_p] \in \mathbb{R}^{n \times p}$ has p features, stored column-wise. We consider problems where the vector of parameters $\beta = (\beta_1, \dots, \beta_p)^\top$ admits a natural group structure. A group of features is a subset $g \subset [p]$ and n_g is its cardinality. The set of groups is denoted by \mathcal{G} and we focus only on non-overlapping groups that form a partition of $[p]$. We denote by β_g the vector in \mathbb{R}^{n_g} which is the restriction of β to the indexes in g . We write $[\beta_g]_j$ the j -th coordinate of β_g . We also use the notation $X_g \in \mathbb{R}^{n \times n_g}$ for the sub-matrix of X assembled from the columns with indexes $j \in g$; similarly $[X_g]_j$ is the j -th column of $[X_g]$.

For any norm Ω , \mathcal{B}_Ω refers to the corresponding unit ball, and \mathcal{B} (resp. \mathcal{B}_∞) stands for the Euclidean (resp. ℓ_∞) unit ball. The soft-thresholding operator (at level $\tau \geq 0$), \mathcal{S}_τ , is defined for any $x \in \mathbb{R}^d$ by $[\mathcal{S}_\tau(x)]_j = \text{sign}(x_j)(|x_j| - \tau)_+$, while the group soft-thresholding (at level τ) is $\mathcal{S}_\tau^{\text{gp}}(x) = (1 - \tau/\|x\|)_+ x$. Denoting $\Pi_{\mathcal{C}}$ the projection on a closed convex set \mathcal{C} , this yields $\mathcal{S}_\tau = \text{Id} - \Pi_{\tau \mathcal{B}_\infty}$. The sub-differential of a convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ at x is defined by $\partial f(x) = \{z \in \mathbb{R}^d : \forall y \in \mathbb{R}^d, f(x) - f(y) \geq z^\top(x - y)\}$. We recall that the sub-differential $\partial \|\cdot\|_1$ of the ℓ_1 norm is $\text{sign}(\cdot)$,

defined element-wise by $\forall j \in [d], \text{sign}(x)_j = \begin{cases} \{\text{sign}(x_j)\}, & \text{if } x_j \neq 0, \\ [-1, 1], & \text{if } x_j = 0. \end{cases}$

Note that the sub-differential $\partial \|\cdot\|$ of the Euclidean norm is $\partial \|\cdot\|(x) = \begin{cases} \{x/\|x\|\}, & \text{if } x \neq 0, \\ \mathcal{B}, & \text{if } x = 0. \end{cases}$

For any norm Ω on \mathbb{R}^d , Ω^D is the dual norm of Ω , and is defined for any $x \in \mathbb{R}^d$ by $\Omega^D(x) = \max_{v \in \mathcal{B}_\Omega} v^\top x$, e.g., $\|\cdot\|_1^D = \|\cdot\|_\infty$ and $\|\cdot\|^D = \|\cdot\|$. We only focus on the Sparse-Group Lasso norm, so we assume that $\Omega = \Omega_{\tau, w}$, where $\Omega_{\tau, w}(\beta) := \tau \|\beta\|_1 + (1 - \tau) \sum_{g \in \mathcal{G}} w_g \|\beta_g\|$, for $\tau \in [0, 1], w = (w_g)_{g \in \mathcal{G}}$ with $w_g \geq 0$ for all $g \in \mathcal{G}$. The case where $w_g = 0$ for some $g \in \mathcal{G}$ together with $\tau = 0$ is excluded ($\Omega_{\tau, w}$ is not a norm in such a case).

2 Sparse-Group Lasso regression

For $\lambda > 0$ and $\tau \in [0, 1]$, the Sparse-Group Lasso estimator denoted by $\hat{\beta}^{(\lambda, \Omega)}$ is defined as a minimizer of the primal objective $P_{\lambda, \Omega}$ defined by:

$$\hat{\beta}^{(\lambda, \Omega)} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|^2 + \lambda \Omega(\beta) := P_{\lambda, \Omega}(\beta). \quad (1)$$

A dual formulation (see [6, Th. 3.3.5]) of (1) is given by

$$\hat{\theta}^{(\lambda, \Omega)} = \arg \max_{\theta \in \Delta_{X, \Omega}} \frac{1}{2} \|y\|^2 - \frac{\lambda^2}{2} \left\| \theta - \frac{y}{\lambda} \right\|^2 := D_\lambda(\theta), \quad (2)$$

where $\Delta_{X, \Omega} = \{\theta \in \mathbb{R}^n : \Omega^D(X^\top \theta) \leq 1\}$. The parameter $\lambda > 0$ controls the trade-off between data-fitting and sparsity, and τ controls the trade-off between features sparsity and group sparsity. In particular one recovers the Lasso [21] if $\tau = 1$, and the Group-Lasso [25] if $\tau = 0$.

For the primal problem, Fermat's rule (cf. Appendix for details) reads:

$$\lambda \hat{\theta}^{(\lambda, \Omega)} = y - X \hat{\beta}^{(\lambda, \Omega)} \quad \text{(link-equation)}, \quad (3)$$

$$X^\top \hat{\theta}^{(\lambda, \Omega)} \in \partial \Omega(\hat{\beta}^{(\lambda, \Omega)}) \quad \text{(sub-differential inclusion)}. \quad (4)$$

Remark 1 (Dual uniqueness). The dual solution $\hat{\theta}^{(\lambda, \Omega)}$ is unique, while the primal solution $\hat{\beta}^{(\lambda, \Omega)}$ might not be. Indeed, the dual formulation (2) is equivalent to $\hat{\theta}^{(\lambda, \Omega)} = \arg \min_{\theta \in \Delta_{X, \Omega}} \|\theta - y/\lambda\|$, so $\hat{\theta}^{(\lambda, \Omega)} = \Pi_{\Delta_{X, \Omega}}(y/\lambda)$ is the projection of y/λ over the dual feasible set $\Delta_{X, \Omega}$.

Remark 2 (Critical parameter: λ_{\max}). There is a critical value λ_{\max} such that 0 is a primal solution of (1) for all $\lambda \geq \lambda_{\max}$. Indeed, the Fermat's rule states $0 \in \arg \min_{\beta \in \mathbb{R}^p} \|y - X\beta\|^2/2 + \lambda \Omega(\beta) \iff 0 \in \{X^\top y\} + \lambda \partial \Omega(0) \iff \Omega^D(X^\top y) \leq \lambda$. Hence, the critical parameter is given by: $\lambda_{\max} := \Omega^D(X^\top y)$. Note that evaluating λ_{\max} highly relies on the ability to (efficiently) compute the dual norm Ω^D .

3 GAP safe rule for the Sparse-Group Lasso

The safe rule we propose here is an extension to the Sparse-Group Lasso of the GAP safe rules introduced for Lasso and Group-Lasso [12, 18]. For the Sparse-Group Lasso, the geometry of the dual feasible set $\Delta_{X, \Omega}$ is more complex (an illustration is given in Fig. 1). Hence, computing a dual feasible point is more intricate. As seen in Section 3.2, the computation of a dual feasible point strongly relies on the ability to evaluate the dual norm Ω^D . This crucial evaluation is discussed in Section 4. We first detail how GAP safe screening rules can be obtained for the Sparse-Group Lasso.

3.1 Description of the screening rules

Safe screening rules exploit the known sparsity of the solutions of problems such as (1). They discard inactive features/groups whose coefficients are guaranteed to be zero for optimal solutions. Then, a significant reduction in computing time can be obtained ignoring "irrelevant" features/groups. The Sparse-Group Lasso benefits from two levels of screening: the safe rules can detect both group-wise zeros in the vector $\hat{\beta}^{(\lambda, \Omega)}$ and coordinate-wise zeros in the remaining groups.

To obtain useful screening rules one needs a **safe region**, *i.e.*, a set containing the optimal dual solution $\hat{\theta}^{(\lambda, \Omega)}$. Following [11], when we choose a ball $\mathcal{B}(\theta_c, r)$ with radius r and centered at θ_c as a safe region, we call it a safe sphere. A safe sphere is all the more useful that r is small and θ_c close to $\hat{\theta}^{(\lambda, \Omega)}$. The safe rules for the Sparse-Group Lasso work as follows: for any group g in \mathcal{G} and any safe sphere $\mathcal{B}(\theta_c, r)$

$$\textbf{Group level safe screening rule:} \quad \max_{\theta \in \mathcal{B}(\theta_c, r)} \|\mathcal{S}_\tau(X_g^\top \theta)\| < (1 - \tau)w_g \Rightarrow \hat{\beta}_g^{(\lambda, \Omega)} = 0, \quad (5)$$

$$\textbf{Feature level safe screening rule:} \quad \forall j \in g, \max_{\theta \in \mathcal{B}(\theta_c, r)} |X_j^\top \theta| < \tau \Rightarrow \hat{\beta}_j^{(\lambda, \Omega)} = 0. \quad (6)$$

This means that provided one the last two test is true, the corresponding group or feature can be (safely) discarded. For screening variables, we rely on the following upper-bounds:

Proposition 1. For all group $g \in \mathcal{G}$ and $j \in g$,

$$\max_{\theta \in \mathcal{B}(\theta_c, r)} |X_j^\top \theta| \leq |X_j^\top \theta_c| + r \|X_j\|. \quad (7)$$

and

$$\max_{\theta \in \mathcal{B}(\theta_c, r)} \|\mathcal{S}_\tau(X_g^\top \theta)\| \leq \mathcal{T}_g := \begin{cases} \|\mathcal{S}_\tau(X_g^\top \theta_c)\| + r \|X_g\| & \text{if } \|X_g^\top \theta_c\|_\infty > \tau, \\ (\|X_g^\top \theta_c\|_\infty + r \|X_g\| - \tau)_+ & \text{otherwise.} \end{cases} \quad (8)$$

Assume now that one has found a safe sphere $\mathcal{B}(\theta_c, r)$ (their creation is deferred to Section 3.2), then the safe screening rules given by (5) and (6) read:

Theorem 1 (Safe rules for the Sparse-Group Lasso). Using \mathcal{T}_g defined in (8), we can state the following safe screening rules:

$$\textbf{Group level safe screening:} \quad \forall g \in \mathcal{G}, \quad \text{if } \mathcal{T}_g < (1 - \tau)w_g, \quad \text{then } \hat{\beta}_g^{(\lambda, \Omega)} = 0,$$

$$\textbf{Feature level safe screening:} \quad \forall g \in \mathcal{G}, \forall j \in g, \quad \text{if } |X_j^\top \theta_c| + r \|X_j\| < \tau, \quad \text{then } \hat{\beta}_j^{(\lambda, \Omega)} = 0.$$

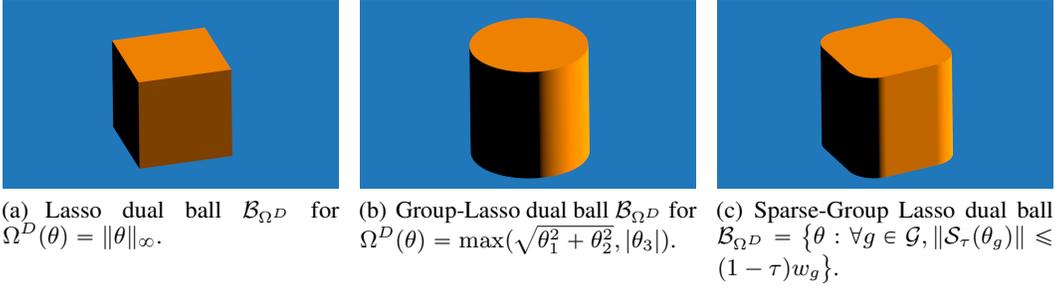


Figure 1: Lasso, Group-Lasso and Sparse-Group Lasso dual unit balls $\mathcal{B}_{\Omega^D} = \{\theta : \Omega^D(\theta) \leq 1\}$, for the case of $\mathcal{G} = \{\{1, 2\}, \{3\}\}$ (i.e., $g_1 = \{1, 2\}, g_2 = \{3\}$), $n = p = 3$, $w_{g_1} = w_{g_2} = 1$ and $\tau = 1/2$.

The screening rules can detect which coordinates or group of coordinates can be safely set to zero. This allows to remove the corresponding features from the design matrix X during the optimization process. While standard algorithms solve (1) scanning all variables, only active ones, i.e., non screened-out variables (using the terminology from Section 3.3) need to be considered with safe screening strategies. This leads to significant computational speed-ups, especially with a coordinate descent algorithm for which it is natural to ignore features (see Algorithm 2, in Appendix G).

3.2 GAP safe sphere

We now show how to compute the safe sphere radius and center using the duality gap.

3.2.1 Computation of the radius

With a dual feasible point $\theta \in \Delta_{X,\Omega}$ and a primal vector $\beta \in \mathbb{R}^p$ at hand, let us construct a safe sphere centered on θ , with radius obtained thanks to dual gap computations.

Theorem 2 (Safe radius). *For any $\theta \in \Delta_{X,\Omega}$ and $\beta \in \mathbb{R}^p$, one has $\hat{\theta}^{(\lambda,\Omega)} \in \mathcal{B}(\theta, r_{\lambda,\Omega}(\beta, \theta))$, for*

$$r_{\lambda,\Omega}(\beta, \theta) = \sqrt{\frac{2(P_{\lambda,\Omega}(\beta) - D_\lambda(\theta))}{\lambda^2}},$$

i.e., the aforementioned ball is a safe region for the Sparse-Group Lasso problem.

Proof. The result holds thanks to strong concavity of the dual objective, cf. Appendix C. \square

3.2.2 Computation of the center

In GAP safe screening rules, the screening test relies crucially on the ability to compute a vector that belongs to the dual feasible set $\Delta_{X,\Omega}$. The geometry of this set is illustrated in Figure 1. Following [5], we leverage the primal/dual link-equation (3) to construct a dual point based on a current approximation β of $\hat{\beta}^{(\lambda,\Omega)}$. When $\beta = \beta^{\lambda'}$ is obtained as an approximation for a previous value of $\lambda' \neq \lambda$ we call such a strategy **sequential screening**. When $\beta = \beta_k$ is the primal value at iteration k obtained by an iterative algorithm, we call this **dynamical screening**. Starting from a residual $\rho = y - X\beta$, one can create a dual feasible point by choosing¹:

$$\theta = \frac{\rho}{\max(\lambda, \Omega^D(X^\top \rho))}. \quad (9)$$

We refer to the sets $\mathcal{B}(\theta, r_{\lambda,\Omega}(\beta, \theta))$ as GAP safe spheres. Note that the generalization to any smooth data fitting term would be straightforward see [18].s

Remark 3. Recall that $\lambda \geq \lambda_{\max}$ yields $\hat{\beta}^{(\lambda,\Omega)} = 0$, in which case $\rho := y - X\hat{\beta}^{(\lambda,\Omega)} = y$ is the optimal residual and y/λ_{\max} is the dual solution. Thus, as for getting $\lambda_{\max} = \Omega^D(X^\top y)$, the scaling computation in (9) requires a dual norm evaluation.

¹We have used a simpler scaling w.r.t. [4] choice's (without noticing much difference in practice): $\theta = s\rho$ where $s = \min \left[\max \left(\frac{\rho^\top y}{\lambda \|\rho\|^2}, \frac{-1}{\Omega^D(X^\top \rho)} \right), \frac{1}{\Omega^D(X^\top \rho)} \right]$.

Algorithm 1 Computation of $\Lambda(x, \alpha, R)$.

Input:
 $x = (x_1, \dots, x_d)^\top \in \mathbb{R}^d$, $\alpha \in [0, 1]$, $R \geq 0$
Output: $\Lambda(x, \alpha, R)$

if $\alpha = 0$ and $R = 0$ **then**
 $\Lambda(x, \alpha, R) = \infty$
else if $\alpha = 0$ and $R \neq 0$ **then**
 $\Lambda(x, \alpha, R) = \|x\|/R$
else if $R = 0$ **then**
 $\Lambda(x, \alpha, R) = \|x\|_\infty/\alpha$
else
Get $I := \{i \in [d] : |x_i| > \frac{\alpha\|x\|_\infty}{\alpha+R}\}$
 $n_I := \text{Card}(I)$
Sort $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(n_I)}$

$S_0 = x_{(0)}$, $S_0^{(2)} = x_{(0)}^2$, $a_0 = 0$
for $k \in [n_I - 1]$ **do**
 $S_k = S_{k-1} + x_{(k)}$; $S_k^{(2)} = S_{k-1}^{(2)} + x_{(k)}^2$
 $a_{k+1} = \frac{S_k^{(2)}}{x_{(k+1)}^2} - 2\frac{S_k}{x_{(k+1)}} + k + 1$
if $\frac{R^2}{\alpha^2} \in [a_k, a_{k+1}[$ **then**
 $j_0 = k + 1$
break
if $\alpha^2 j_0 - R^2 = 0$ **then**
 $\Lambda(x, \alpha, R) = \frac{S_{j_0}^2}{2\alpha S_{j_0}}$
else
 $\Lambda(x, \alpha, R) = \frac{\alpha S_{j_0} - \sqrt{\alpha^2 S_{j_0}^2 - S_{j_0}^{(2)}(\alpha^2 j_0 - R^2)}}{\alpha^2 j_0 - R^2}$

3.3 Convergence of the active set

The next proposition states that the sequence of dual feasible points obtained from (9) converges to the dual solution $\hat{\theta}^{(\lambda, \Omega)}$ if $(\beta_k)_{k \in \mathbb{N}}$ converges to an optimal primal solution $\hat{\beta}^{(\lambda, \Omega)}$ (proof in Appendix). It guarantees that the GAP safe spheres $\mathcal{B}(\theta_k, r_{\lambda, \Omega}(\beta_k, \theta_k))$ are converging safe regions in the sense introduced by [12], since by strong duality $\lim_{k \rightarrow \infty} r_{\lambda, \Omega}(\beta_k, \theta_k) = 0$.

Proposition 2. *If $\lim_{k \rightarrow \infty} \beta_k = \hat{\beta}^{(\lambda, \Omega)}$, then $\lim_{k \rightarrow \infty} \theta_k = \hat{\theta}^{(\lambda, \Omega)}$.*

For any safe region \mathcal{R} , *i.e.*, a set containing $\hat{\theta}^{(\lambda, \Omega)}$, we define two levels of active sets, one for the group level and one for the feature level:

$$\mathcal{A}_{\text{gp}}(\mathcal{R}) := \{g \in \mathcal{G}, \max_{\theta \in \mathcal{R}} \|\mathcal{S}_\tau(X_g^\top \theta)\| \geq (1 - \tau)w_g\}, \quad \mathcal{A}_{\text{ft}}(\mathcal{R}) := \bigcup_{g \in \mathcal{A}_{\text{gp}}(\mathcal{R})} \{j \in g : \max_{\theta \in \mathcal{R}} |X_j^\top \theta| \geq \tau\}.$$

If one considers sequence of converging regions, then the next proposition (whose proof in Appendix) states that we can identify in finite time the optimal active sets defined as follows:

$$\mathcal{E}_{\text{gp}} := \left\{ g \in \mathcal{G} : \left\| \mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)}) \right\| = (1 - \tau)w_g \right\}, \quad \mathcal{E}_{\text{ft}} := \bigcup_{g \in \mathcal{E}_{\text{gp}}} \left\{ j \in g : |X_j^\top \hat{\theta}^{(\lambda, \Omega)}| \geq \tau \right\}.$$

Proposition 3. *Let $(\mathcal{R}_k)_{k \in \mathbb{N}}$ be a sequence of safe regions whose diameters converge to 0. Then, $\lim_{k \rightarrow \infty} \mathcal{A}_{\text{gp}}(\mathcal{R}_k) = \mathcal{E}_{\text{gp}}$ and $\lim_{k \rightarrow \infty} \mathcal{A}_{\text{ft}}(\mathcal{R}_k) = \mathcal{E}_{\text{ft}}$.*

4 Properties of the Sparse-Group Lasso

To apply our safe rule, we need to be able to evaluate the dual norm Ω^D efficiently. We describe such as step hereafter along with some useful properties of the norm Ω . Such evaluations are performed multiple times during the algorithm, motivating the derivation of an efficient algorithm, as presented in Algorithm 1.

4.1 Connections with ϵ -norms

Here, we establish a link between the Sparse-Group Lasso norm Ω and the ϵ -norm (denoted $\|\cdot\|_\epsilon$) introduced in [8]. For any $\epsilon \in [0, 1]$ and $x \in \mathbb{R}^d$, $\|x\|_\epsilon$ is defined as the unique nonnegative solution ν of the equation $\sum_{i=1}^d (|x_i| - (1 - \epsilon)\nu)_+^2 = (\epsilon\nu)^2$, ($\|x\|_0 := \|x\|_\infty$). Using soft-thresholding, this is equivalent to solve in ν the equation $\sum_{i=1}^d \mathcal{S}_{(1-\epsilon)\nu}(x_i)^2 = \|\mathcal{S}_{(1-\epsilon)\nu}(x)\|^2 = (\epsilon\nu)^2$. Moreover, the dual norm of the ϵ -norm is given by²: $\|y\|_\epsilon^D = \epsilon\|y\|^D + (1 - \epsilon)\|y\|_\infty^D = \epsilon\|y\| + (1 - \epsilon)\|y\|_1$. Now we can express the Sparse-Group Lasso norm Ω in term of the dual ϵ -norm and derive some basic properties.

²see [9, Eq. (42)] or Appendix

Proposition 4. For all groups g in \mathcal{G} , let us introduce $\epsilon_g := \frac{(1-\tau)w_g}{\tau+(1-\tau)w_g}$. Then, the Sparse-Group Lasso norm satisfies the following properties: for any β and ξ in \mathbb{R}^p

$$\Omega(\beta) = \sum_{g \in \mathcal{G}} (\tau + (1-\tau)w_g) \|\beta_g\|_{\epsilon_g}^D, \quad \text{and} \quad \Omega^D(\xi) = \max_{g \in \mathcal{G}} \frac{\|\xi_g\|_{\epsilon_g}}{\tau + (1-\tau)w_g}, \quad (10)$$

$$\mathcal{B}_{\Omega^D} = \{\xi \in \mathbb{R}^p : \forall g \in \mathcal{G}, \|\mathcal{S}_\tau(\xi_g)\| \leq (1-\tau)w_g\}. \quad (11)$$

The sub-differential at β reads $\partial\Omega(\beta) = \{z \in \mathbb{R}^p : \forall g \in \mathcal{G}, z_g \in \tau\partial\|\cdot\|_1(\beta_g) + (1-\tau)w_g\partial\|\cdot\|(\beta_g)\}$.

We obtain from the characterization of the unit dual ball (11) that for the Sparse-Group Lasso, any dual feasible point $\theta \in \Delta_{X,\Omega}$ verifies: $\forall g \in \mathcal{G}, X_g^\top \theta \in (1-\tau)w_g\mathcal{B} + \tau\mathcal{B}_\infty$.

From the dual norm formulation (10), a vector $\theta \in \mathbb{R}^n$ is feasible if and only if $\Omega^D(X^\top \theta) \leq 1$, i.e., $\forall g \in \mathcal{G}, \|X_g^\top \theta\|_{\epsilon_g} \leq \tau + (1-\tau)w_g$. Hence we deduce from (11) a new characterization of the dual feasible set: $\Delta_{X,\Omega} = \{\theta \in \mathbb{R}^n : \forall g \in \mathcal{G}, \|X_g^\top \theta\|_{\epsilon_g} \leq \tau + (1-\tau)w_g\}$.

4.2 Efficient computation of the dual norm

The following proposition shows how to compute the dual norm of the Sparse-Group Lasso (and the ϵ -norm). This is turned into an efficient procedure in Algorithm 1 (see the Appendix for details).

Proposition 5. For $\alpha \in [0, 1], R \geq 0$ and $x \in \mathbb{R}^d$, the equation $\sum_{i=1}^d \mathcal{S}_{\nu\alpha}(x_i)^2 = (\nu R)^2$ has a unique solution $\nu := \Lambda(x, \alpha, R) \in \mathbb{R}_+$, that can be computed in $O(d \log d)$ operations in the worst case. With $n_I = \text{Card}\{i \in [d] : |x_i| > \alpha\|x\|_\infty / (\alpha + R)\}$, the complexity of Algorithm 1 is $n_I + n_I \log(n_I)$, which is comparable to the ambient dimension d .

Thanks to Remark 2, we can explicit the critical parameter λ_{\max} for the Sparse-Group Lasso that is

$$\lambda_{\max} = \max_{g \in \mathcal{G}} \frac{\Lambda(X_g^\top y, 1 - \epsilon_g, \epsilon_g)}{\tau + (1-\tau)w_g} = \Omega^D(X^\top y), \quad (12)$$

and get a dual feasible point (9), since $\Omega^D(X^\top \rho) = \max_{g \in \mathcal{G}} \Lambda(X_g^\top \rho, 1 - \epsilon_g, \epsilon_g) / (\tau + (1-\tau)w_g)$.

5 Implementation

In this section we provide details on how to solve the Sparse-Group Lasso primal problem, and how we apply the GAP safe screening rules. We focus on the *block coordinate iterative soft-thresholding algorithm (ISTA-BC)*; see [19]. This algorithm requires a block-wise Lipschitz gradient condition on the data fitting term $f(\beta) = \|y - X\beta\|^2/2$. For our problem (1), one can show that for all group g in \mathcal{G} , $L_g = \|X_g\|_2^2$ (where $\|\cdot\|_2$ is the spectral norm of a matrix) is a suitable block-wise Lipschitz constant. We define the block coordinate descent algorithm according to the Majorization-Minimization principle: at each iteration l , we choose (e.g., cyclically) a group g and the next iterate β^{l+1} is defined such that $\beta_{g'}^{l+1} = \beta_{g'}^l$ if $g' \neq g$ and otherwise $\beta_g^{l+1} = \arg \min_{\beta_g \in \mathbb{R}^{n_g}} \|\beta_g - (\beta_g^l - \nabla_g f(\beta^l)/L_g)\|^2/2 + (\tau\|\beta_g\|_1 + (1-\tau)w_g\|\beta_g\|)\lambda/L_g$, where we denote for all g in \mathcal{G} , $\alpha_g := \lambda/L_g$. This can be simplified to $\beta_g^{l+1} = \mathcal{S}_{(1-\tau)w_g\alpha_g}^{\text{SP}}(\mathcal{S}_{\tau\alpha_g}(\beta_g^l - \nabla_g f(\beta^l)/L_g))$. The expensive computation of the dual gap is not performed at each pass over the data, but only every f^{ce} pass (in practice $f^{\text{ce}} = 10$ in all our experiments). A pseudo code is given in Appendix G.

6 Experiments

In this section we present our experiments and illustrate the numerical benefit of screening rules for the Sparse-Group Lasso.

6.1 Experimental settings and methods compared

We have run our ISTA-BC algorithm³ to obtain the Sparse-Group Lasso estimator for a non-increasing sequence of T regularization parameters $(\lambda_t)_{t \in [T-1]}$ defined as follows: $\lambda_t := \lambda_{\max} 10^{-\delta(t-1)/(T-1)}$.

³The source code can be found in https://github.com/EugeneNdiaye/GAPSAFE_SGL.

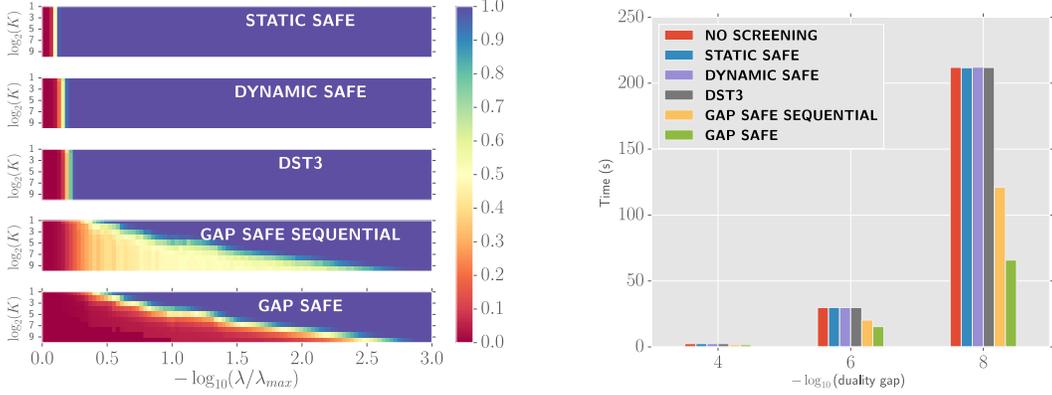


Figure 2: Experiments on a synthetic dataset ($\rho = 0.5, \gamma_1 = 10, \gamma_2 = 4, \tau = 0.2$). (a) Proportion of active variables, *i.e.*, variables not safely eliminated, as a function of parameters (λ_t) and the number of iterations K . More red, means more variables eliminated and better screening. (b) Time to reach convergence w.r.t the accuracy on the duality gap, using various screening strategies.

By default, we choose $\delta = 3$ and $T = 100$, following the standard practice when running cross-validation using sparse models (see R glmnet package [13]). The weights are always chosen as $w_g = \sqrt{n_g}$ (as in [20]).

We also provide a natural extension of the previous safe rules [11, 24, 5] to the Sparse-Group Lasso for comparisons (please refer to Appendix D for more details). The **static safe region** [11] is given by $\mathcal{B}(y/\lambda, \|y/\lambda_{\max} - y/\lambda\|)$. The corresponding **dynamic safe region** [5] is given by $\mathcal{B}(y/\lambda, \|\theta_k - y/\lambda\|)$, where $(\theta_k)_{k \in \mathbb{N}}$ is a sequence of dual feasible points obtained by dual scaling; *cf.* Equation (9). The **DST3**, is an improvement of the preceding safe region, see [24, 5], that we adapted to the Sparse-Group Lasso. The **GAP safe sequential** rules corresponds to using only GAP Safe spheres whose centers are the (last) dual point output by the solver for a former value of λ in the path. The **GAP safe** rules corresponds to performing our strategy both sequentially and dynamically. Presenting the sequential rule allows to measure the benefits due to sequential rules and to the dynamic rules.

We now demonstrate the efficiency of our method in both synthetic (Fig. (2)) and real datasets (Fig. 6.2). For comparison, we report computation times to reach convergence up to a certain tolerance on the duality gap for all the safe rules considered.

Synthetic dataset: We use a common framework [22, 23] based on the model $y = X\beta + 0.01\varepsilon$ where $\varepsilon \sim \mathcal{N}(0, \text{Id}_n)$, $X \in \mathbb{R}^{n \times p}$ follows a multivariate normal distribution such that $\forall(i, j) \in [p]^2$, $\text{corr}(X_i, X_j) = \rho^{|i-j|}$. We fix $n = 100$ and break randomly $p = 10000$ in 1000 groups of size 10 and select γ_1 groups to be active and the others are set to zero. In each selected groups, γ_2 coordinates are drawn with $[\beta_g]_j = \text{sign}(\xi) \times U$ for U is uniform in $[0.5, 10]$, ξ uniform in $[-1, 1]$.

Real dataset: NCEP/NCAR Reanalysis 1 [17] The dataset contains monthly means of climate data measurements spread across the globe in a grid of $2.5^\circ \times 2.5^\circ$ resolutions (longitude and latitude 144×73) from 1948/1/1 to 2015/10/31. Each grid point constitutes a group of 7 predictive variables (*Air Temperature, Precipitable water, Relative humidity, Pressure, Sea Level Pressure, Horizontal Wind Speed and Vertical Wind Speed*) whose concatenation across time constitutes our design matrix $X \in \mathbb{R}^{814 \times 73577}$. Such data have therefore a natural group structure.

In our experiments, we considered as target variable $y \in \mathbb{R}^{814}$, the values of *Air Temperature* in a neighborhood of Dakar. Seasonality and trend are first removed, as usually done in climate analysis for bias reduction in the regression estimates. Similar data has been used in [10], showing that the Sparse-Group Lasso estimator is well suited for prediction in climatology. Indeed, thanks to the sparsity structure, the estimates delineate via their support some predictive regions at the group level, as well as predictive features via coordinate-wise screening.

We choose τ in the set $\{0, 0.1, \dots, 0.9, 1\}$ by splitting in 50% the observations and run a training-test validation procedure. For each value of τ , we require a duality gap of 10^{-8} on the training part

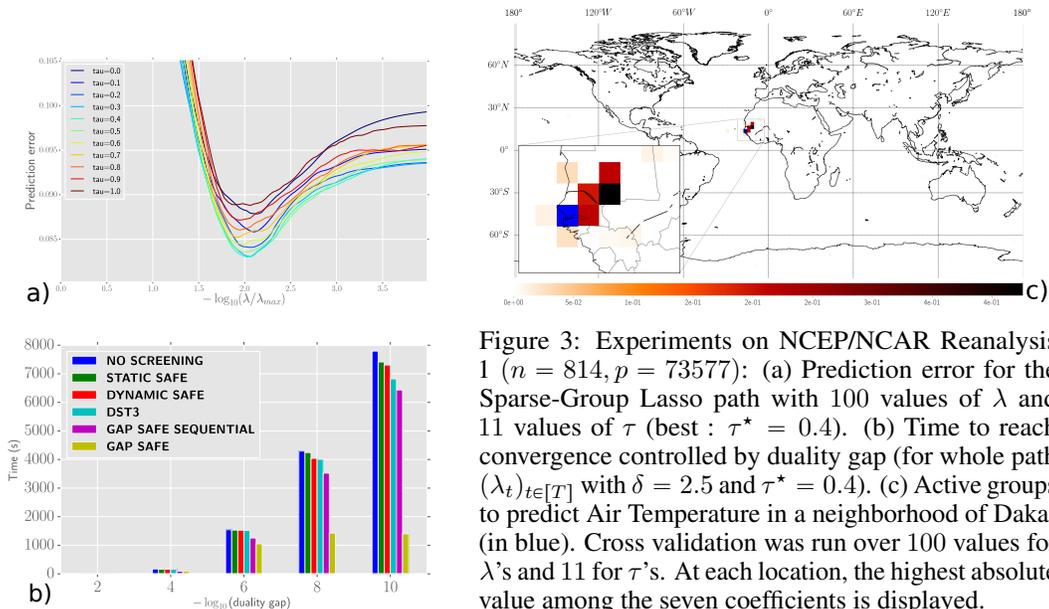


Figure 3: Experiments on NCEP/NCAR Reanalysis 1 ($n = 814, p = 73577$): (a) Prediction error for the Sparse-Group Lasso path with 100 values of λ and 11 values of τ (best : $\tau^* = 0.4$). (b) Time to reach convergence controlled by duality gap (for whole path $(\lambda_t)_{t \in [T]}$ with $\delta = 2.5$ and $\tau^* = 0.4$). (c) Active groups to predict Air Temperature in a neighborhood of Dakar (in blue). Cross validation was run over 100 values for λ 's and 11 for τ 's. At each location, the highest absolute value among the seven coefficients is displayed.

and pick the best one in term of prediction accuracy on the test part. The result is displayed in Figure 6.2.(a). We fixed $\delta = 2.5$ for the computational time benchmark in Figure 6.2.(b)

6.2 Performance of the screening rules

In all our experiments, we observe that our proposed GAP Safe rule outperforms the other rules in term of computation time. On Figure 2.(c), we can see that we need 65s to reach convergence whereas others rules need up to 212s at a precision of 10^{-8} . A similar performance is observed on the real dataset (Figure 6.2) where we obtain up to a 5x speed up over the other rules. The key reason behind this performance gain is the convergence of the GAP Safe regions toward the dual optimal point as well as the efficient strategy to compute the screening rule. As shown in the results presented on Figure 2, our method still manages to screen out variables when λ is small. It corresponds to low regularizations which lead to less sparse solutions but need to be explored during cross-validation.

In the climate experiments, the support map in Figure 6.2.(c) shows that the most important coefficients are distributed in the vicinity of the target region (in agreement with our intuition). Nevertheless, some active variables with small coefficients remain and cannot be screened out.

Note that we do not compare our method to the **TLFre** [23], since this sequential rule requires the exact knowledge of the dual optimal solution which is not available in practice. As a consequence, one may discard active variables which can prevent the algorithm from converging as shown in [18].

7 Conclusion

The recent GAP safe rules introduced have shown great improvements, for a wide range of regularized regression, in the reduction of computing time, especially in high dimension. To apply such GAP safe rules to the Sparse-Group Lasso, we have proposed a new description of the dual feasible set by establishing connections between the Sparse-Group Lasso norm and ϵ -norms. This geometrical connection has helped providing an efficient algorithm to compute the dual norm and dual feasible points, bottlenecks for applying the GAP Safe rules. Extending GAP safe rules on general hierarchical regularizations, is a possible direction for future research.

Acknowledgments: this work was supported by the ANR THALAMEEG ANR-14-NEUC-0002-01, the NIH R01 MH106174, by ERC Starting Grant SLAB ERC-YStG-676943 and by the Chair Machine Learning for Big Data at Télécom ParisTech.

References

- [1] A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
- [2] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Convex optimization with sparsity-inducing norms. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- [3] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, New York, 2011.
- [4] A. Bonnefoy, V. Emiya, L. Ralaivola, and R. Gribonval. A dynamic screening principle for the lasso. In *EUSIPCO*, 2014.
- [5] A. Bonnefoy, V. Emiya, L. Ralaivola, and R. Gribonval. Dynamic Screening: Accelerating First-Order Algorithms for the Lasso and Group-Lasso. *IEEE Trans. Signal Process.*, 63(19):20, 2015.
- [6] J. M. Borwein and A. S. Lewis. *Convex analysis and nonlinear optimization*. Springer, New York, second edition, 2006.
- [7] P. Bühlmann and S. van de Geer. *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg, 2011. Methods, theory and applications.
- [8] O. Burdakov. A new vector norm for nonlinear curve fitting and some other optimization problems. *33. Int. Wiss. Kolloq. Fortschr.reihe "Mathematische Optimierung | Theorie und Anwendungen"*, pages 15–17, 1988.
- [9] O. Burdakov and B. Merkulov. On a new norm for data fitting and optimization problems. *Linköping University, Linköping, Sweden, Tech. Rep. LiTH-MAT*, 2001.
- [10] S. Chatterjee, K. Steinhaeuser, A. Banerjee, S. Chatterjee, and A. Ganguly. Sparse group lasso: Consistency and climate applications. In *SIAM International Conference on Data Mining*, pages 47–58, 2012.
- [11] L. El Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination in sparse supervised learning. *J. Pacific Optim.*, 8(4):667–698, 2012.
- [12] O. Fercoq, A. Gramfort, and J. Salmon. Mind the duality gap: safer rules for the lasso. In *ICML*, pages 333–342, 2015.
- [13] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Ann. Appl. Stat.*, 1(2):302–332, 2007.
- [14] J.-B. Hiriart-Urruty. A note on the Legendre-Fenchel transform of convex composite functions. In *Nonsmooth Mechanics and Analysis*, pages 35–46. Springer, 2006.
- [15] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *J. Mach. Learn. Res.*, 12:2297–2334, 2011.
- [16] T. B. Johnson and C. Guestrin. Blitz: A principled meta-algorithm for scaling sparse optimization. In *ICML*, pages 1171–1179, 2015.
- [17] E. Kalnay, M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, et al. The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American meteorological Society*, 77(3):437–471, 1996.
- [18] E. Ndiaye, O. Fercoq, A. Gramfort, and J. Salmon. GAP safe screening rules for sparse multi-task and multi-class models. *NIPS*, pages 811–819, 2015.
- [19] Z. Qin, K. Scheinberg, and D. Goldfarb. Efficient block-coordinate descent algorithms for the group lasso. *Mathematical Programming Computation*, 5(2):143–169, 2013.
- [20] N. Simon, J. Friedman, T. Hastie, and R. Tibshirani. A sparse-group lasso. *J. Comput. Graph. Statist.*, 22(2):231–245, 2013.

- [21] R. Tibshirani. Regression shrinkage and selection via the lasso. *JRSSB*, 58(1):267–288, 1996.
- [22] R. Tibshirani, J. Bien, J. Friedman, T. Hastie, N. Simon, J. Taylor, and R. J. Tibshirani. Strong rules for discarding predictors in lasso-type problems. *JRSSB*, 74(2):245–266, 2012.
- [23] J. Wang and J. Ye. Two-layer feature reduction for sparse-group lasso via decomposition of convex sets. *arXiv preprint arXiv:1410.4210*, 2014.
- [24] Z. J. Xiang, H. Xu, and P. J. Ramadge. Learning sparse representations of high dimensional data on large scale dictionaries. In *NIPS*, pages 900–908, 2011.
- [25] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *JRSSB*, 68(1):49–67, 2006.
- [26] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *JRSSB*, 67(2):301–320, 2005.

A Convex optimization reminder

We first recall the necessary tools for building screening rules, namely the Fermat's first order optimality condition (also called Fermat's rule) and the characterization of the sub-differential of a norm by means of its dual norm.

Proposition 6 (Fermat's rule). ([3, Prop. 26.1]) For any convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$x^* \in \arg \min_{x \in \mathbb{R}^d} f(x) \iff 0 \in \partial f(x^*). \quad (13)$$

Proposition 7. ([2, Prop. 1.2]) The sub-differential of the norm Ω at x , denoted $\partial\Omega(x)$, is given by

$$\begin{cases} \{z \in \mathbb{R}^d : \Omega^D(z) \leq 1\} = \mathcal{B}_{\Omega^D} & \text{if } x = 0, \\ \{z \in \mathbb{R}^d : \Omega^D(z) = 1 \text{ and } z^\top x = \Omega(x)\} & \text{otherwise.} \end{cases} \quad (14)$$

B Additional convexity and optimization tools

In what follows we will use the dot product notation for any $x, x' \in \mathbb{R}^d$ we write $\langle x, x' \rangle = x^\top x'$.

We denote by ι_C the indicator function of a set C defined as

$$\iota_C : \mathbb{R}^d \rightarrow \mathbb{R}, \quad \iota_C(x) = \begin{cases} 0, & \text{if } x \in C, \\ +\infty, & \text{otherwise.} \end{cases} \quad (15)$$

We denote by $f^* : \mathbb{R}^d \rightarrow \mathbb{R}$ the Fenchel conjugate of f defined for any $z \in \mathbb{R}^d$ by $f^*(z) = \sup_{w \in \mathbb{R}^d} w^\top z - f(w)$.

Proposition 8. ([2, Prop. 1.4]) The Fenchel conjugate of the norm Ω is given by

$$\Omega^*(\xi) = \sup_{w \in \mathbb{R}^d} [\xi^\top w - \Omega(w)] = \iota_{\mathcal{B}_{\Omega^D}}(\xi). \quad (16)$$

C Proofs

We first remind the simple properties underlying the concept of safe screening rules [11] in our Sparse-Group Lasso context.

Proposition 0. The two levels of screening rules for the Sparse-Group Lasso are:
Feature level screening:

$$\forall j \in g, |X_j^\top \hat{\theta}^{(\lambda, \Omega)}| < \tau \implies \hat{\beta}_j^{(\lambda, \Omega)} = 0.$$

Group level screening:

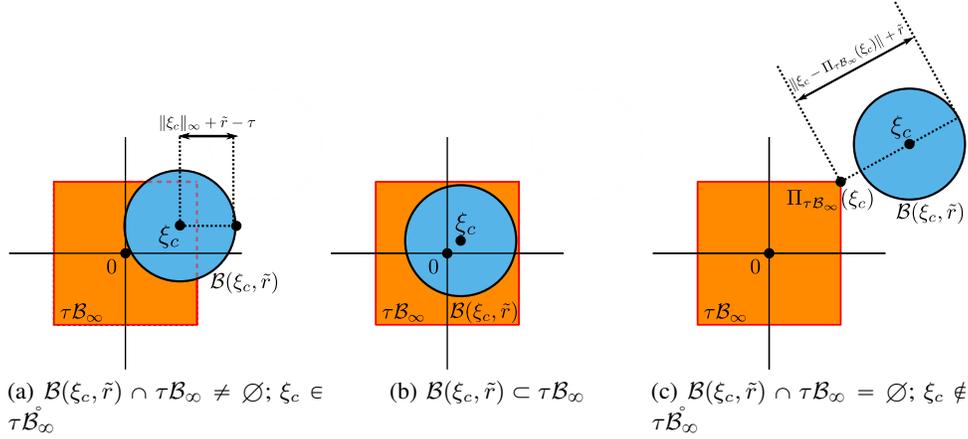
$$\forall g \in \mathcal{G}, \left\| \mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)}) \right\| < (1 - \tau)w_g \implies \hat{\beta}_g^{(\lambda, \Omega)} = 0.$$

Remark 4. The first rule is has a strict inequality, but can be relaxed to a non-strict one if $\tau \neq 1$.

Proof. Let us consider $\hat{\beta}_g^{(\lambda, \Omega)} \neq 0, g \in \mathcal{G}$. Then combining the **subdifferential inclusion** (4), the subdifferential of the ℓ_2 -norm and the decomposition of any dual feasible point from Proposition 4, we obtain :

$$\begin{aligned} X_g^\top \hat{\theta}^{(\lambda, \Omega)} &= \tau v_g + (1 - \tau)w_g \frac{\hat{\beta}_g^{(\lambda, \Omega)}}{\|\hat{\beta}_g^{(\lambda, \Omega)}\|} \text{ where } v \in \partial \|\cdot\|_1(\hat{\beta}_g^{(\lambda, \Omega)}), \\ X_g^\top \hat{\theta}^{(\lambda, \Omega)} &= \Pi_{\tau \mathcal{B}_{\infty}}(X_g^\top \hat{\theta}^{(\lambda, \Omega)}) + \mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)}). \end{aligned}$$

So we can deduce that $\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)}) \in (1 - \tau)w_g \left\{ \frac{\hat{\beta}_g^{(\lambda, \Omega)}}{\|\hat{\beta}_g^{(\lambda, \Omega)}\|} \right\}$. Since $\hat{\theta}^{(\lambda, \Omega)}$ is feasible then $\|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\| < (1 - \tau)w_g$ is equivalent to $\|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\| \neq (1 - \tau)w_g$ which implies, by



contrapositive, that $\hat{\beta}_g^{(\lambda, \Omega)} = 0$. Hence we obtain the group level safe rule. Furthermore, from the subdifferential of the ℓ_1 -norm, we have:

$$\forall j \in g, X_j^\top \hat{\theta}^{(\lambda, \Omega)} \in \begin{cases} (1 - \tau)w_g \left\{ \frac{\hat{\beta}_j^{(\lambda, \Omega)}}{\|\hat{\beta}^{(\lambda, \Omega)}\|} \right\} + \tau \left\{ \text{sign}(\hat{\beta}_j^{(\lambda, \Omega)}) \right\}, & \text{if } \hat{\beta}_j^{(\lambda, \Omega)} \neq 0, \\ [-\tau, \tau], & \text{if } \hat{\beta}_j^{(\lambda, \Omega)} = 0. \end{cases}$$

Hence, if $\hat{\beta}_j^{(\lambda, \Omega)} \neq 0$ then $X_j^\top \hat{\theta}^{(\lambda, \Omega)} = \text{sign}(\hat{\beta}_j^{(\lambda, \Omega)}) \left[(1 - \tau)w_g \frac{|\hat{\beta}_j^{(\lambda, \Omega)}|}{\|\hat{\beta}^{(\lambda, \Omega)}\|} + \tau \right]$ and so $|X_j^\top \hat{\theta}^{(\lambda, \Omega)}| \geq \tau$. By contrapositive, we obtain the feature level safe rule. \square

Proposition 1. For all group $g \in \mathcal{G}$ and $j \in g$,

$$\max_{\theta \in \mathcal{B}(\theta_c, r)} |X_j^\top \theta| \leq |X_j^\top \theta_c| + r \|X_j\|. \quad (17)$$

$\max_{\theta \in \mathcal{B}(\theta_c, r)} \|\mathcal{S}_\tau(X_g^\top \theta)\|$ is upper bounded by

$$\mathcal{T}_g = \begin{cases} \|\mathcal{S}_\tau(X_g^\top \theta_c)\| + r \|X_g\| & \text{if } \|X_g^\top \theta_c\|_\infty > \tau, \\ (\|X_g^\top \theta_c\|_\infty + r \|X_g\| - \tau)_+ & \text{otherwise.} \end{cases} \quad (18)$$

Proof. $|X_j^\top \theta| \leq |[X_g^\top (\theta - \theta_c)]_j| + |X_j^\top \theta_c| \leq r \|X_j\| + |X_j^\top \theta_c|$ as soon as $\theta \in \mathcal{B}(\theta_c, r)$.

Since $\theta \in \mathcal{B}(\theta_c, r)$ implies that $X_g^\top \theta \in \mathcal{B}(X_g^\top \theta_c, r \|X_g\|)$, we have $\max_{\theta \in \mathcal{B}(\theta_c, r)} \|\mathcal{S}_\tau(X_g^\top \theta)\| \leq \max_{\xi \in \mathcal{B}(\xi_c, \tilde{r})} \|\mathcal{S}_\tau(\xi)\|$ where $\xi_c = X_g^\top \theta_c$ and $\tilde{r} = r \|X_g\|$. From now, we just have to show how to compute $\max_{\xi \in \mathcal{B}(\xi_c, \tilde{r})} \|\mathcal{S}_\tau(\xi)\|$.

- In the case where $\xi_c \in \tau\mathring{\mathcal{B}}_\infty$, if $\|\xi_c\|_\infty + \tilde{r} \leq \tau$ (i.e., $\mathcal{B}(\xi_c, \tilde{r}) \subset \tau\mathcal{B}_\infty$), we have $\Pi_{\tau\mathcal{B}_\infty}(\xi) = \xi$ and thus, $\max_{\xi \in \mathcal{B}(\xi_c, \tilde{r})} \|\mathcal{S}_\tau(\xi)\| = \max_{\xi \in \mathcal{B}(\xi_c, \tilde{r})} \|\xi - \Pi_{\tau\mathcal{B}_\infty}(\xi)\| = 0$.
- Otherwise if $\xi_c \in \tau\mathring{\mathcal{B}}_\infty$ and $\|\xi_c\|_\infty + \tilde{r} > \tau$, for any vector $\xi \in \partial\mathcal{B}(\xi_c, \tilde{r}) \cap (\tau\mathcal{B}_\infty)^c$ and any vector $\tilde{\xi} \in \partial\tau\mathcal{B}_\infty \cap [\xi, \xi_c]$, $\|\xi - \Pi_{\tau\mathcal{B}_\infty}(\xi)\| \leq \|\xi - \tilde{\xi}\| = \tilde{r} - \|\tilde{\xi} - \xi_c\|$. Hence

$$\begin{aligned} \max_{\xi \in \mathcal{B}(\xi_c, \tilde{r})} \|\xi - \Pi_{\tau\mathcal{B}_\infty}(\xi)\| &\leq \max_{\substack{\xi \in \partial\mathcal{B}(\xi_c, \tilde{r}) \cap (\tau\mathcal{B}_\infty)^c \\ \tilde{\xi} \in \partial\tau\mathcal{B}_\infty \cap [\xi, \xi_c]}} \tilde{r} - \|\tilde{\xi} - \xi_c\| \\ &\leq \tilde{r} - \min_{\xi \in \partial\tau\mathcal{B}_\infty} \|\xi - \xi_c\| \\ &= \tilde{r} - \tau + \|\xi_c\|_\infty. \end{aligned}$$

This upper bound is attained. Indeed, $\max_{\theta \in \mathcal{B}(\theta_c, r)} \|\xi - \Pi_{\tau\mathcal{B}_\infty}(\xi)\| = \tilde{r} - \|\Pi_{\tau\mathcal{B}_\infty}(\hat{\xi}) - \xi_c\| = \tilde{r} - \tau + \|\xi_c\|_\infty$ where $\hat{\xi}$ is a vector in $\partial\mathcal{B}(\xi_c, \tilde{r})$ such that $\Pi_{\tau\mathcal{B}_\infty}(\hat{\xi}) = \xi_c + e_{j^*}(\tau - \|\xi_c\|_\infty)$ and $j^* \in \arg \max_{j \in [p]} |(\xi_c)_j|$.

- If $\xi_c \notin \tau\mathcal{B}_\infty$, since the projection operator on a convex set is a contraction, we have

$$\begin{aligned} \forall \xi \in \partial\mathcal{B}(\xi_c, \tilde{r}), \|\xi - \Pi_{\tau\mathcal{B}_\infty}(\xi)\| &\leq \|\xi - \Pi_{\tau\mathcal{B}_\infty}(\xi_c)\| \\ &\leq \|\xi_c - \Pi_{\tau\mathcal{B}_\infty}(\xi_c)\| + \|\xi - \xi_c\| \\ &= \|\xi_c - \Pi_{\tau\mathcal{B}_\infty}(\xi_c)\| + \tilde{r}. \end{aligned}$$

Moreover, it is straightforward to see that the vector $\tilde{\xi} := \tilde{\gamma}\xi_c + (1 - \tilde{\gamma})\Pi_{\tau\mathcal{B}_\infty}(\xi_c)$ where $\tilde{\gamma} = 1 + \frac{\tilde{r}}{\|\xi_c\| + \|\Pi_{\tau\mathcal{B}_\infty}(\xi_c)\|}$ belongs to $\partial\mathcal{B}(\xi_c, \tilde{r})$; it verifies $\Pi_{\tau\mathcal{B}_\infty}(\xi_c) = \Pi_{\tau\mathcal{B}_\infty}(\tilde{\xi})$ and it attains this bound. \square

Theorem 2 (Safe radius). *For any $\theta \in \Delta_{X,\Omega}$ and any $\beta \in \mathbb{R}^p$, one has $\hat{\theta}^{(\lambda,\Omega)} \in \mathcal{B}(\theta, r_{\lambda,\Omega}(\beta, \theta))$, for*

$$r_{\lambda,\Omega}(\beta, \theta) = \sqrt{\frac{2(P_{\lambda,\Omega}(\beta) - D_\lambda(\theta))}{\lambda^2}},$$

i.e., the aforementioned ball is a safe region for the Sparse-Group Lasso problem.

Proof. By weak duality, $\forall \beta \in \mathbb{R}^p$, $D_\lambda(\hat{\theta}^{(\lambda,\Omega)}) \leq P_{\lambda,\Omega}(\beta)$. Then, note that the dual objective function (1) is λ^2 -strongly concave. This implies:

$$\forall (\theta, \theta') \in \Delta_{X,\Omega} \times \Delta_{X,\Omega}, \quad D_\lambda(\theta) \leq D_\lambda(\theta') + \nabla D_\lambda(\theta')^\top (\theta - \theta') - \frac{\lambda^2}{2} \|\theta - \theta'\|^2.$$

Moreover, since $\hat{\theta}^{(\lambda,\Omega)}$ maximizes the concave function D_λ , the following inequality holds true:

$$\forall \theta \in \Delta_{X,\Omega}, \quad \nabla D_\lambda(\hat{\theta}^{(\lambda,\Omega)})^\top (\theta - \hat{\theta}^{(\lambda,\Omega)}) \leq 0.$$

Hence, we have for all $\theta \in \Delta_{X,\Omega}$ and $\beta \in \mathbb{R}^p$:

$$\begin{aligned} \frac{\lambda^2}{2} \|\theta - \hat{\theta}^{(\lambda,\Omega)}\|^2 &\leq D_\lambda(\hat{\theta}^{(\lambda,\Omega)}) - D_\lambda(\theta) \\ &\leq P_{\lambda,\Omega}(\beta) - D_\lambda(\theta). \end{aligned} \quad \square$$

Proposition 2. *If $\lim_{k \rightarrow \infty} \beta_k = \hat{\beta}^{(\lambda,\Omega)}$, then $\lim_{k \rightarrow \infty} \theta_k = \hat{\theta}^{(\lambda,\Omega)}$.*

Proof. Let $\alpha_k = \max(\lambda, \Omega_{\tau,w}^D(X^\top \rho_k))$ and recall that $\rho_k = y - X\beta_k$. We have :

$$\begin{aligned} \|\theta_k - \hat{\theta}^{(\lambda,\Omega)}\| &= \left\| \frac{1}{\alpha_k} (y - X\beta_k) - \frac{1}{\lambda} (y - X\hat{\beta}^{(\lambda,\Omega)}) \right\| \\ &= \left\| \left(\frac{1}{\alpha_k} - \frac{1}{\lambda} \right) (y - X\beta_k) - \frac{(X\hat{\beta}^{(\lambda,\Omega)} - X\beta_k)}{\lambda} \right\| \\ &\leq \left| \frac{1}{\alpha_k} - \frac{1}{\lambda} \right| \|y - X\beta_k\| + \left\| \frac{X\hat{\beta}^{(\lambda,\Omega)} - X\beta_k}{\lambda} \right\|. \end{aligned}$$

If $\beta_k \rightarrow \hat{\beta}^{(\lambda,\Omega)}$, then $\alpha_k \rightarrow \max(\lambda, \Omega_{\tau,w}^D(X^\top (y - X\hat{\beta}^{(\lambda,\Omega)}))) = \max(\lambda, \lambda \Omega_{\tau,w}^D(X^\top \hat{\theta}^{(\lambda,\Omega)})) = \lambda$ since $y - X\hat{\beta}^{(\lambda,\Omega)} = \lambda \hat{\theta}^{(\lambda,\Omega)}$ thanks to the **link-equation** (3) and since $\hat{\theta}^{(\lambda,\Omega)}$ is feasible i.e., $\Omega_{\tau,w}^D(X^\top \hat{\theta}^{(\lambda,\Omega)}) \leq 1$. Hence, both terms in the previous inequality converge to zero. \square

Proposition 3. *Let $(\mathcal{R}_k)_{k \in \mathbb{N}}$ be a sequence of safe regions whose diameters converge to 0. Then, $\lim_{k \rightarrow \infty} \mathcal{A}_{gp}(\mathcal{R}_k) = \mathcal{E}_{gp}$ and $\lim_{k \rightarrow \infty} \mathcal{A}_{ft}(\mathcal{R}_k) = \mathcal{E}_{ft}$.*

Proof. We proceed by double inclusion. First let us prove that $\exists k_0$ s.t. $\forall k \geq k_0, \mathcal{A}_{gp}(\mathcal{R}_k) \subset \mathcal{E}_{gp}$. Indeed, since the diameter of \mathcal{R}_k converges to zero, for any $\epsilon > 0$ there exist $k_0 \in \mathbb{N}, \forall k \geq k_0, \forall \theta \in \mathcal{R}_k, \|\theta - \hat{\theta}^{(\lambda,\Omega)}\| \leq \epsilon$. The triangle inequality implies that $\forall g \notin \mathcal{E}_{gp}, \|\mathcal{S}_\tau(X_g^\top \theta)\| \leq$

$\|\mathcal{S}_\tau(X_g^\top \theta) - \mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\| + \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\|$. Since the soft-thresholding operator is 1-Lipschitz, we have:

$$\|\mathcal{S}_\tau(X_g^\top \theta)\| \leq \|X_g(\theta - \hat{\theta}^{(\lambda, \Omega)})\| + \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\| \leq \epsilon \|X_g\| + \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\|,$$

as soon as $k \geq k_0$. Moreover, $\forall g \notin \mathcal{E}_{\text{gp}}$,

$$\|\mathcal{S}_\tau(X_g^\top \theta)\| \leq \max_{g \notin \mathcal{E}_{\text{gp}}} \|\mathcal{S}_\tau(X_g^\top \theta)\| \leq \epsilon \max_{g \notin \mathcal{E}_{\text{gp}}} \|X_g\| + \max_{g \notin \mathcal{E}_{\text{gp}}} \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\|.$$

It suffices to choose ϵ such that

$$\epsilon \max_{g \notin \mathcal{E}_{\text{gp}}} \|X_g\| + \max_{g \notin \mathcal{E}_{\text{gp}}} \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\| < (1 - \tau)w_g,$$

that is to say $\epsilon < \frac{(1-\tau)w_g - \max_{g \notin \mathcal{E}_{\text{gp}}} \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda, \Omega)})\|}{\max_{g \notin \mathcal{E}_{\text{gp}}} \|X_g\|}$, to remove the group g . For any $k \geq k_0$, $\mathcal{E}_{\text{gp}}^c = \{g \in \mathcal{G} : \|\mathcal{S}_\tau(X_g^\top \hat{\theta}^{(\lambda)})\| < (1 - \tau)w_g\} \subset \mathcal{A}_{\text{gp}}(\mathcal{R}_k)^c$, the set of variables removed by our screening rule. This proves the first inclusion.

Now we show that $\forall k \in \mathbb{N}$, $\mathcal{A}_{\text{gp}}(\mathcal{R}_k) \supset \mathcal{E}_{\text{gp}}$. Indeed, for all $g^* \in \mathcal{E}_{\text{gp}}$, $\|\mathcal{S}_\tau(X_{g^*}^\top \hat{\theta}^{(\lambda, \Omega)})\| = (1 - \tau)w_{g^*}$. Since for all k in \mathbb{N} , $\hat{\theta}^{(\lambda, \Omega)} \in \mathcal{R}_k$ then $\max_{\theta \in \mathcal{R}_k} \|\mathcal{S}_\tau(X_{g^*}^\top \theta)\| \geq \|\mathcal{S}_\tau(X_{g^*}^\top \hat{\theta}^{(\lambda, \Omega)})\| = (1 - \tau)w_{g^*}$ hence the second inclusion holds.

We have shown that $\forall k \geq k_0$, $\mathcal{A}_{\text{gp}}(\mathcal{R}_k) = \mathcal{E}_{\text{gp}}$ and so $\mathcal{A}_{\text{fit}}(\mathcal{R}_k) \subset \bigcup_{g \in \mathcal{E}_{\text{gp}}} \{j \in g : \max_{\theta \in \mathcal{R}_k} |X_j^\top \theta| \geq \tau\}$. Moreover, the same reasoning yields $\forall g \in \mathcal{G}$, $\{j \in g : \max_{\theta \in \mathcal{R}_k} |X_j^\top \theta| \geq \tau\} \subset \{j \in g : |X_j^\top \hat{\theta}^{(\lambda, \Omega)}| \geq \tau\}$. Hence $\forall k \geq k_0$, $\mathcal{A}_{\text{fit}}(\mathcal{R}_k) \subset \mathcal{A}_{\text{fit}}$. The reciprocal inclusion is straightforward. \square

Proposition 4. . For all group g in \mathcal{G} , let $\epsilon_g := \frac{(1-\tau)w_g}{\tau + (1-\tau)w_g}$ then the Sparse-Group Lasso norm satisfies the following properties: for any vectors β and ξ in \mathbb{R}^p

$$\Omega_{\tau, w}(\beta) = \sum_{g \in \mathcal{G}} (\tau + (1 - \tau)w_g) \|\beta_g\|_{\epsilon_g}^D, \quad (19)$$

$$\Omega_{\tau, w}^D(\xi) = \max_{g \in \mathcal{G}} \frac{\|\xi_g\|_{\epsilon_g}}{\tau + (1 - \tau)w_g}, \quad (20)$$

$$\mathcal{B}_{\Omega_{\tau, w}^D} = \{\xi \in \mathbb{R}^p : \forall g \in \mathcal{G}, \|\mathcal{S}_\tau(\xi_g)\| \leq (1 - \tau)w_g\}. \quad (21)$$

The subdifferential $\partial \Omega_{\tau, w}(\beta)$ of the norm $\Omega_{\tau, w}$ at β is given by

$$\left\{ x \in \mathbb{R}^p : \forall g \in \mathcal{G}, x_g \in \tau \partial \|\cdot\|_1(\beta_g) + (1 - \tau)w_g \partial \|\cdot\|(\beta_g) \right\}.$$

Proof.

$$\begin{aligned} \forall \beta \in \mathbb{R}^p, \Omega(\beta) &= \tau \|\beta\|_1 + (1 - \tau) \sum_{g \in \mathcal{G}} w_g \|\beta_g\| = \sum_{g \in \mathcal{G}} (\tau \|\beta_g\|_1 + (1 - \tau)w_g \|\beta_g\|) \\ &= \sum_{g \in \mathcal{G}} (\tau + (1 - \tau)w_g) \left[\frac{\tau}{\tau + (1 - \tau)w_g} \|\beta_g\|_1 + \frac{(1 - \tau)w_g}{\tau + (1 - \tau)w_g} \|\beta_g\| \right] \\ &= \sum_{g \in \mathcal{G}} (\tau + (1 - \tau)w_g) [(1 - \epsilon_g) \|\beta_g\|_1 + \epsilon_g \|\beta_g\|] \\ &= \sum_{g \in \mathcal{G}} (\tau + (1 - \tau)w_g) \|\beta_g\|_{\epsilon_g}^D. \end{aligned}$$

The definition of the dual norm reads $\Omega^D(\xi) = \max_{\beta: \Omega(\beta) \leq 1} \beta^\top \xi$, and solving this problem yields:

$$\begin{aligned}
\Omega^D(\xi) &= \sup_{\beta: \Omega(\beta) \leq 1} \langle \beta, \xi \rangle = \sup_{\beta} \inf_{\mu > 0} \langle \beta, \sum_{g \in \mathcal{G}} \xi_g \rangle - \mu \left(\sum_{g \in \mathcal{G}} \Omega_g(\beta_g) - 1 \right) \\
&= \inf_{\mu > 0} \left\{ \sum_{g \in \mathcal{G}} \sup_{\beta_g} [\langle \beta_g, \xi_g \rangle - \mu \Omega_g(\beta_g)] + \mu \right\} \\
&= \inf_{\mu > 0} \left\{ \sum_{g \in \mathcal{G}} \mu \Omega_g^* \left(\frac{\xi_g}{\mu} \right) + \mu \right\} = \inf_{\mu > 0} \left\{ \sum_{g \in \mathcal{G}} \iota_{\mathcal{B}_{\Omega_g^D}} \left(\frac{\xi_g}{\mu} \right) + \mu \right\} \\
&= \inf_{\mu > 0} \left\{ \max_{g \in \mathcal{G}} \iota_{\mathcal{B}_{\Omega_g^D}} \left(\frac{\xi_g}{\mu} \right) + \mu \right\} = \max_{g \in \mathcal{G}} \inf_{\mu > 0} \left\{ \Omega_g^* \left(\frac{\xi_g}{\mu} \right) + \mu \right\} \\
&= \max_{g \in \mathcal{G}} \inf_{\mu > 0} \sup_{\beta_g} \langle \beta_g, \frac{\xi_g}{\mu} \rangle - \Omega_g(\beta_g) + \mu \\
&= \max_{g \in \mathcal{G}} \inf_{\mu > 0} \sup_{u_g} \langle u_g, \xi_g \rangle - \mu (\Omega_g(u_g) - 1) \quad (\text{with } \mu u_g = \beta_g) \\
&= \max_{g \in \mathcal{G}} \sup_{u_g: \Omega_g(u_g) \leq 1} \langle u_g, \xi_g \rangle = \max_{g \in \mathcal{G}} \sup_{u_g} \langle u_g, \xi_g \rangle \quad \text{s.t. } (\tau + (1 - \tau)w_g) \|u_g\|_{\epsilon_g}^D \leq 1 \\
&= \max_{g \in \mathcal{G}} \sup_{u_g: \Omega_g(u_g) \leq 1} \langle u_g, \xi_g \rangle = \max_{g \in \mathcal{G}} \sup_{u'_g: \|u'_g\|_{\epsilon_g}^D \leq 1} \frac{u'^{\top}_g \xi_g}{\tau + (1 - \tau)w_g} = \max_{g \in \mathcal{G}} \frac{\|\xi_g\|_{\epsilon_g}}{\tau + (1 - \tau)w_g}.
\end{aligned}$$

We recall here the proof of [23] for the sake of completeness. First let us write $\Omega(\beta) = \Omega_1(\beta) + \Omega_2(\beta)$, where $\Omega_1(\beta) = \tau \|\beta\|_1$ and $\Omega_2(\beta) = (1 - \tau) \sum_{g \in \mathcal{G}} w_g \|\beta_g\|_2$. Since Ω_1 and Ω_2 are continuous everywhere, we have (see [14, Theorem 1]): $\Omega^*(\xi) = (\Omega_1 + \Omega_2)^*(\xi) = \min_{a+b=\xi} [\Omega_1^*(a) + \Omega_2^*(b)] = \min_a [\Omega_1^*(a) + \Omega_2^*(\xi - a)]$, which is also the inf-convolution (see [3, Chapter 12]) of these two norms. Using the Fenchel conjugate of the ℓ_1 norm ($\Omega_1^* = \iota_{\mathcal{B}_{\infty}}$) and of the ℓ_2 norm ($\Omega_2^* = \iota_{\mathcal{B}}$), we have

$$\Omega^*(\xi) = \sum_{g \in \mathcal{G}} \min_{a_g} \iota_{\mathcal{B}_{\infty}}(a_g) + \iota_{\mathcal{B}} \left(\frac{\xi_g - a_g}{(1 - \tau)w_g} \right) = \sum_{g \in \mathcal{G}} \iota_{\mathcal{B}} \left(\frac{\xi_g - \Pi_{\mathcal{B}_{\infty}}(\xi_g)}{(1 - \tau)w_g} \right).$$

Hence the indicator of the unit dual ball is $\iota_{\mathcal{B}_{\Omega^D}}(\xi) = \sum_{g \in \mathcal{G}} \iota_{(1 - \tau)w_g \mathcal{B}}(\xi_g - \Pi_{\mathcal{B}_{\infty}}(\xi_g))$ and using $\mathcal{S}_{\tau}(\xi_g) = \xi_g - \Pi_{\mathcal{B}_{\infty}}(\xi_g)$, we have:

$$\mathcal{B}_{\Omega^D} = \{\xi \in \mathbb{R}^p : \Omega^D(\xi) \leq 1\} = \{\xi \in \mathbb{R}^p : \forall g \in \mathcal{G}, \|\mathcal{S}_{\tau}(\xi_g)\| \leq (1 - \tau)w_g\}.$$

□

Proposition 5. . For $\alpha \in [0, 1]$, $R \geq 0$ and $x \in \mathbb{R}^d$, the equation $\sum_{j=1}^d \mathcal{S}_{\nu\alpha}(x_j)^2 = (\nu R)^2$ has a unique solution $\nu \in \mathbb{R}_+$, denoted by $\Lambda(x, \alpha, R)$ and that can be computed in $O(d \log d)$ operations in worst case. With $n_I = \text{Card}\{i \in [d] : |x_i| > \alpha \|x\|_{\infty} / (\alpha + R)\}$, the complexity of Algorithm 1 is $n_I + n_I \log(n_I)$, which is comparable to the ambient dimension d .

Proof. Dividing by ν^2 , which is positive as soon as $x \neq 0$, we get that $\sum_{j=1}^d \mathcal{S}_{\nu\alpha}(x_j)^2 = (\nu R)^2$ is equivalent to $\sum_{j=1}^d \mathcal{S}_{\alpha}(x_j/\nu)^2 = R^2$. Note that $\sum_{j=1}^d \mathcal{S}_{\alpha}(x_j/\nu)^2 = \sum_{j=1}^d \mathcal{S}_{\alpha}(|x_j|/\nu)^2$ so without loss of generality we assume $x \in \mathbb{R}_+^d$.

The case $\alpha = 0$ and $R = 0$ corresponds to the situation where all x_j are equal to zero or we impose ν equals to infinity. So we avoid this trivial case.

If $\alpha = 0$ and $R \neq 0$, $\nu = \|x\|/R$. Indeed,

$$\sum_{j=1}^d \mathcal{S}_0(x_j/\nu)^2 = R^2 \iff \sum_{j=1}^d (x_j/\nu)^2 = R^2 \iff \frac{\|x\|_2^2}{\nu^2} = R^2 \text{ hence the result.}$$

If $\alpha \neq 0$ and $R = 0$, we have :

$$\sum_{j=1}^d \mathcal{S}_\alpha \left(\frac{x_j}{\nu} \right)^2 = 0 \iff \forall j \in [d], \left(\frac{x_j}{\nu} - \alpha \right)_+ = 0 \iff \forall j \in [d], \frac{x_j}{\nu} \leq \alpha \iff \nu \geq \frac{\max_{j \in [d]} x_j}{\alpha}.$$

So we choose the smallest ν i.e., $\nu = \|x\|_\infty / \alpha$. In all the above cases, the computation is done in $O(d)$.

Otherwise $\alpha \neq 0$ and $R \neq 0$. The function $\nu \mapsto \sum_{j=1}^d \mathcal{S}_\alpha(x_j/\nu)^2$ is a non-increasing continuous function with limit $+\infty$ (resp. 0) when $\nu \rightarrow 0$ (resp. $\nu \rightarrow +\infty$). Hence, there is a unique solution to $\sum_{j=1}^d \mathcal{S}_\alpha(x_j/\nu)^2 = R^2$.

We denote by $x_{(1)}, \dots, x_{(d)}$ the coordinates of x ordered in decreasing order (with the convention $x_{(0)} = +\infty$ and $x_{(d+1)} = 0$). Note that $\sum_{j=1}^d \mathcal{S}_\alpha(x_j/\nu)^2 = \sum_{j=1}^d \mathcal{S}_\alpha(x_{(j)}/\nu)^2$. Then, there exists an index $j_0 \in [p]$ such that

$$R^2 \in \left[\sum_{j=0}^d \mathcal{S}_\alpha \left(\alpha \frac{x_{(j)}}{x_{(j_0)}} \right)^2, \sum_{j=0}^d \mathcal{S}_\alpha \left(\alpha \frac{x_{(j)}}{x_{(j_0+1)}} \right)^2 \right). \quad (22)$$

For such a j_0 , one can check that $\nu \in (x_{(j_0+1)}/\alpha, x_{(j_0)}/\alpha]$. The definition of the soft-thresholding operator yields

$$\mathcal{S}_\alpha(x_j/\nu)^2 = \begin{cases} (x_j/\nu - \alpha)^2 & \text{if } x_j \geq \nu\alpha, \\ 0 & \text{if } x_j < \nu\alpha. \end{cases} \quad (23)$$

It can be simplified thanks to $x_j \geq x_{(j_0)} \Rightarrow x_j \geq \nu\alpha$ and $x_j \leq x_{(j_0+1)} \Rightarrow x_j < \nu\alpha$. Hence, $R^2 = \sum_{j=1}^{j_0} (x_{(j)}/\nu - \alpha)^2 = \sum_{j=1}^{j_0} (x_{(j)}/\nu)^2 + \alpha^2 \sum_{j=1}^{j_0} 1 - 2\alpha \sum_{j=1}^{j_0} x_{(j)}/\nu$ so solving $\sum_{j=1}^p \mathcal{S}_\alpha(x_{(j)}/\nu)^2 = R^2$ is equivalent to solve on \mathbb{R}_+

$$(\alpha^2 j_0 - R^2)\nu^2 - \left(2\alpha \sum_{j=1}^{j_0} x_{(j)} \right) \nu + \sum_{j=1}^{j_0} x_{(j)}^2 = 0. \quad (24)$$

If $(\alpha^2 j_0 - R^2) = 0$, then $\nu = \sum_{j=1}^{j_0} x_{(j)}^2 / (2\alpha \sum_{j=1}^{j_0} x_{(j)})$. Otherwise ν is the unique solution lying in $(x_{(j_0+1)}/\alpha, x_{(j_0)}/\alpha]$ of the quadratic equation stated in Eq. (24).

In the worst case, to compute $\Lambda(x, \alpha, R)$, one needs to sort a vector of size d , what can be done in $O(d \log(d))$ operations, and finding j_0 thanks to (22) requires $O(d^2)$ if we apply a naive algorithm.

In the following, we show that one can easily reduce the complexity to $O(d \log(d))$ in worst case.

For all j in $[d]$, $\mathcal{S}_\alpha \left(\alpha \frac{x_j}{x_{(j_0)}} \right) = 0$ as soon as $x_j \leq x_{(j_0)}$. This implies that (22) is equivalent to

$$R^2 \in \left[\sum_{j=0}^{j_0-1} \mathcal{S}_\alpha \left(\alpha \frac{x_{(j)}}{x_{(j_0)}} \right)^2, \sum_{j=0}^{j_0} \mathcal{S}_\alpha \left(\alpha \frac{x_{(j)}}{x_{(j_0+1)}} \right)^2 \right). \quad (25)$$

Denoting $S_{j_0} := \sum_{j=1}^{j_0} x_{(j)}$ and $S_{j_0}^{(2)} := \sum_{j=1}^{j_0} x_{(j)}^2$, a direct calculation show that (25) can be rewritten as

$$R^2 \in \alpha^2 \left[\frac{S_{j_0-1}^{(2)}}{x_{(j_0)}^2} - 2 \frac{S_{j_0-1}}{x_{(j_0)}} + j_0, \frac{S_{j_0}^{(2)}}{x_{(j_0+1)}^2} - 2 \frac{S_{j_0}}{x_{(j_0+1)}} + j_0 + 1 \right). \quad (26)$$

Finally, solving $\sum_{j=1}^p \mathcal{S}_\alpha(x_{(j)}/\nu)^2 = R^2$ is equivalent to finding the solution of $(\alpha^2 j_0 - R^2)\nu^2 - (2\alpha S_{j_0})\nu + S_{j_0}^{(2)} = 0$ lying in $(x_{(j_0+1)}/\alpha, x_{(j_0)}/\alpha]$. Hence,

$$\Lambda(x, \alpha, R) = \frac{\alpha S_{j_0} - \sqrt{\alpha^2 S_{j_0}^2 - S_{j_0}^{(2)}(\alpha^2 j_0 - R^2)}}{\alpha^2 j_0 - R^2} =: \nu_1 \quad (27)$$

or

$$\Lambda(x, \alpha, R) = \frac{\alpha S_{j_0} + \sqrt{\alpha^2 S_{j_0}^2 - S_{j_0}^{(2)}(\alpha^2 j_0 - R^2)}}{\alpha^2 j_0 - R^2} =: \nu_2. \quad (28)$$

- If $\alpha^2 j_0 - R^2 < 0$, then $\nu_2 < 0$ and so it cannot be a solution since $\Lambda(x, \alpha, R)$ must be positive.
- Otherwise, we have

$$\nu_2 \geq \frac{\alpha S_{j_0}}{\alpha^2 j_0 - R^2} = \frac{1}{\alpha(j_0 - \frac{R^2}{\alpha^2})} \sum_{j=1}^{j_0} x_{(j)} > \frac{1}{\alpha j_0} \sum_{j=1}^{j_0} x_{(j)} \geq \frac{x_{(j_0)}}{\alpha},$$

where the second inequality results from the fact that $j_0 > j_0 - R^2/\alpha^2$. And again ν_2 cannot be a solution since $\Lambda(x, \alpha, R)$ belongs to $(x_{(j_0+1)}/\alpha, x_{(j_0)}/\alpha]$.

Hence, in all cases, the solution is given by ν_1 .

Moreover, we can significantly reduce the cost of the sort. In fact, for all ν , we have $\|\mathcal{S}_{\alpha\nu}(x)\| \geq \|\mathcal{S}_{\alpha\nu}(x)\|_\infty = \max_{j \in [d]} (|x_j| - \nu\alpha)_+$. Hence, $\|\mathcal{S}_{\alpha\nu}(x)\| - \nu R \geq \|x\|_\infty - \nu\alpha - \nu R > 0$ if and only if $\nu < \|x\|_\infty/(\alpha + R)$. Combining this with Equation (23), we take into account only the coordinates which have an absolute value greater than $\alpha\|x\|_\infty/(\alpha + R)$. Indeed, by contrapositive, if ν is a solution then $\nu \geq \|x\|_\infty/\alpha + R$ hence $x_j < \alpha\|x\|_\infty/\alpha + R \Rightarrow x_j < \nu\alpha \stackrel{(23)}{\Rightarrow} \mathcal{S}_\alpha(x_j/\nu) = 0$.

Finally, computing $\Lambda(x, \alpha, R)$ can be done by applying Algorithm 1. Note that this algorithm is similar to [9, Algorithm 4]. \square

D Notes about others methods

D.1 Extension of some previous methods to the Sparse-Group Lasso

D.1.1 Extension of [11]: *static safe region*

The *static safe region* can be obtained as in [11] using the ball $\mathcal{B}(y/\lambda, \|y/\lambda_{\max} - y/\lambda\|)$.

Indeed y/λ_{\max} is a dual feasible point. Hence the distance between y/λ and y/λ_{\max} is smaller than the distance between y/λ and $\hat{\theta}^{(\lambda, \Omega)}$ since the last point is the projection of y/λ over the (close and convex) dual feasible set $\Delta_{X, \Omega}$.

Extension of [4]: *dynamic safe region*

The *dynamic safe region* can be obtained as in [11] using the ball $\mathcal{B}(y/\lambda, \|\theta_k - y/\lambda\|)$, where the sequence $(\theta_k)_{k \in \mathbb{N}}$ converges to the dual optimal vector $\hat{\theta}^{(\lambda, \Omega)}$.

A sequence of dual points is required to construct such a ball, and following [11] we can consider the *dual scaling* procedure. We have chosen a simple procedure here: Let $\theta_k = \rho_k / \max(\lambda, \Omega_{\tau, w}^D(X^\top \rho_k))$, where $\rho_k := y - X\beta_k$, for a primal converging sequence β_k . Hence, one can use the safe sphere $\mathcal{B}(y/\lambda, \|\theta_k - y/\lambda\|)$ with the same reasoning as for the *static safe region*.

Hence, we can easily extend the corresponding screening rules to the Sparse-Group Lasso thanks to the formulation (6) and (5).

Extension of [4]: *DST3 safe region*

Now we show that the safe regions proposed in [24] for the Lasso and generalized in [4] to the Group-Lasso can be adapted to the Sparse-Group Lasso. For that, we define $g_\star = \arg \max_{g \in \mathcal{G}} \Omega^D(X^\top y)$,

$$\mathcal{V}_\star = \left\{ \theta \in \mathbb{R}^n : \|X_{g_\star}^\top \theta\|_{\epsilon_{g_\star}} \leq \tau + (1 - \tau)w_{g_\star} \right\} \text{ and } \mathcal{H}_\star = \left\{ \theta \in \mathbb{R}^n : \langle \theta, \eta \rangle = \tau + (1 - \tau)w_{g_\star} \right\}.$$

Where η is the vector normal to \mathcal{V}_\star at y/λ_{\max} and is given by $\eta := X_{g_\star}^\top \nabla \|\cdot\|_{\epsilon_{g_\star}}(X_{g_\star}^\top y/\lambda_{\max})$, where $\nabla \|\cdot\|_\epsilon(x) = \mathcal{S}_{(1-\epsilon)\|x\|_\epsilon}(x) / \|\mathcal{S}_{(1-\epsilon)\|x\|_\epsilon}(x)\|_\epsilon^D$, see Lemma 5 below. Let

$$\theta_c := \frac{y}{\lambda} - \left(\frac{\langle \frac{y}{\lambda}, \eta \rangle - (\tau + (1 - \tau)w_{g_\star})}{\|\eta\|^2} \right) \eta$$

be the projection of y/λ onto the hyperplane \mathcal{H}_\star and $r_{\theta_k} := \sqrt{\|y/\lambda - \theta_k\|^2 - \|y/\lambda - \theta_c\|^2}$ where θ_k is a dual feasible vector (which can be obtained by dual scaling). Then, the following proposition holds.

Proposition 9. *Let θ_c and r_{θ_k} defined as above. Then $\hat{\theta}^{(\lambda, \Omega)} \in \mathcal{B}(\theta_c, r_{\theta_k})$.*

Proof. We set $\mathcal{H}_\star^- := \{\theta \in \mathbb{R}^n : \langle \theta, \eta \rangle \leq \tau + (1 - \tau)w_{g_\star}\}$ the negative half-space induced by the hyperplane \mathcal{H}_\star . Since $\hat{\theta}^{(\lambda, \Omega)} \in \Delta_{X, \Omega} \subset \mathcal{V}_\star \subset \mathcal{H}_\star^-$ and $\mathcal{B}(\frac{y}{\lambda}, \|\frac{y}{\lambda} - \theta_k\|)$ is a safe region, then $\hat{\theta}^{(\lambda, \Omega)} \in \mathcal{H}_\star^- \cap \mathcal{B}(\frac{y}{\lambda}, \|\frac{y}{\lambda} - \theta_k\|)$.

Moreover, for any $\theta \in \mathcal{H}_\star^- \cap \mathcal{B}(\frac{y}{\lambda}, \|\frac{y}{\lambda} - \theta_k\|)$, we have:

$$\begin{aligned} \left\| \frac{y}{\lambda} - \theta_k \right\|^2 &\geq \left\| \frac{y}{\lambda} - \theta \right\|^2 = \left\| \left(\frac{y}{\lambda} - \theta_c \right) + (\theta_c - \theta) \right\|^2 = \\ &\left\| \frac{y}{\lambda} - \theta_c \right\|^2 + \|\theta_c - \theta\|^2 + 2 \left\langle \frac{y}{\lambda} - \theta_c, \theta_c - \theta \right\rangle. \end{aligned}$$

Since $\theta_c = \Pi_{\mathcal{H}_\star^-}(\frac{y}{\lambda})$ and \mathcal{H}_\star^- is convex, then $\langle \theta_c - \frac{y}{\lambda}, \theta_c - \theta \rangle \leq 0$. Thus

$$\left\| \frac{y}{\lambda} - \theta_k \right\|^2 \geq \left\| \frac{y}{\lambda} - \theta_c \right\|^2 + \|\theta_c - \theta\|^2, \text{ hence } \|\theta - \theta_c\| \leq \sqrt{\left\| \frac{y}{\lambda} - \theta_k \right\|^2 - \left\| \frac{y}{\lambda} - \theta_c \right\|^2} =: r_{\theta_k}.$$

Which show that $\mathcal{H}_\star^- \cap \mathcal{B}(\frac{y}{\lambda}, \|\frac{y}{\lambda} - \theta_k\|) \subset \mathcal{B}(\theta_c, r_{\theta_k})$. Hence the result. \square

E Sparse-Group Lasso plus Elastic Net

The Elastic-Net estimator ([26]) can be mixed with the Sparse-Group Lasso by considering

$$\arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|^2 + \lambda_1 \Omega(\beta) + \frac{\lambda_2}{2} \|\beta\|^2. \quad (29)$$

By setting $\tilde{X} = \begin{pmatrix} X \\ \sqrt{\lambda_2} \text{Id}_p \end{pmatrix} \in \mathbb{R}^{n+p, p}$ and $\tilde{y} = \begin{pmatrix} y \\ 0 \end{pmatrix} \in \mathbb{R}^{n+p}$, we can reformulate (29) as

$$\arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|\tilde{y} - \tilde{X}\beta\|^2 + \lambda_1 \Omega(\beta), \quad (30)$$

and we can adapt our GAP safe rule framework to this case.

F Properties of the ϵ -norm

We describe, for completeness, some properties of the ϵ -norm. The following material is inspired from [9].

Lemma 1. For all $\xi \in \mathbb{R}^d$, the ϵ -decomposition reads: $\xi = \xi^\epsilon + \xi^{1-\epsilon}$, where $\xi^\epsilon := \mathcal{S}_{(1-\epsilon)\|\xi\|_\epsilon}(\xi)$ and $\xi^{1-\epsilon} := \xi - \xi^\epsilon$. Moreover, $\|\xi^\epsilon\| = \epsilon \|\xi\|_\epsilon$ and $\|\xi^{1-\epsilon}\|_\infty = (1 - \epsilon) \|\xi\|_\epsilon$. Hence, the following decomposition holds for the ϵ -norm: $\|\xi\|_\epsilon = \|\xi^\epsilon\| + \|\xi^{1-\epsilon}\|_\infty$.

Proof. $\|\xi^\epsilon\| = \|\mathcal{S}_{(1-\epsilon)\|\xi\|_\epsilon}(\xi)\| = \epsilon \|\xi\|_\epsilon$ by definition of the ϵ -norm $\|\xi\|_\epsilon$. Moreover,

$$\xi^{1-\epsilon} = \sum_{i=1}^d [\xi_i - \text{sign}(\xi_i)(|\xi_i| - (1 - \epsilon) \|\xi\|_\epsilon)_+] = \sum_{i=1}^d \text{sign}(\xi_i) [|\xi_i| - (|\xi_i| - (1 - \epsilon) \|\xi\|_\epsilon)_+].$$

Thus, using the symbol $a \vee b$ to represent $\max(a, b)$, one has

$$\begin{aligned} \|\xi^{1-\epsilon}\|_\infty &= \max_{i \in [d]} |\text{sign}(\xi_i) [|\xi_i| - (|\xi_i| - (1 - \epsilon) \|\xi\|_\epsilon)_+]| \\ &= \max_{i \in [d]} \|\xi_i| - (|\xi_i| - (1 - \epsilon) \|\xi\|_\epsilon)_+ \vee \max_{i \in [d]} \|\xi_i| - (|\xi_i| - (1 - \epsilon) \|\xi\|_\epsilon)_+ \\ &= \max_{i \in [d]} |\xi_i| \vee (1 - \epsilon) \|\xi\|_\epsilon = (1 - \epsilon) \|\xi\|_\epsilon. \end{aligned}$$

\square

Lemma 2. Let us define $U(\|\xi\|_\epsilon) := \{u \in \mathbb{R}^d : \|u\| \leq \epsilon\|\xi\|_\epsilon\}$ and $V(\|\xi\|_\epsilon) := \{v \in \mathbb{R}^d : \|v\|_\infty \leq (1 - \epsilon)\|\xi\|_\epsilon\}$. Then

$$\xi^{(1-\epsilon)} = \arg \min_{\substack{u \in U(\|\xi\|_\epsilon) \\ \xi = u+v}} \|v\|_\infty \quad \text{and} \quad \xi^\epsilon = \arg \min_{\substack{v \in V(\|\xi\|_\epsilon) \\ \xi = u+v}} \|u\|.$$

Proof.

• **Existence and uniqueness of the solutions**

It is clear that

$$\arg \min_{\substack{u \in U(\|\xi\|_\epsilon) \\ \xi = u+v}} \|v\|_\infty = \arg \min_{\xi - U(\|\xi\|_\epsilon)} \|v\|_\infty,$$

and

$$\arg \min_{\substack{v \in V(\|\xi\|_\epsilon) \\ \xi = u+v}} \|u\| = \arg \min_{\xi - V(\|\xi\|_\epsilon)} \|u\|.$$

Thus, these two problems have unique solution because we minimize strict convex functions onto convex sets.

• **Uniqueness of the ϵ -decomposition**

From Lemma 1 we have $\xi = \xi^\epsilon + \xi^{1-\epsilon}$ where $\|\xi^\epsilon\| = \epsilon\|\xi\|_\epsilon$ and $\|\xi^{1-\epsilon}\|_\infty = (1 - \epsilon)\|\xi\|_\epsilon$. Hence $\xi^\epsilon \in U(\|\xi\|_\epsilon)$ and $\xi^{1-\epsilon} \in V(\|\xi\|_\epsilon)$. Now it suffices to show that this ϵ -decomposition is unique.

Suppose $\xi \neq 0$ (the uniqueness is trivial otherwise) and $v \in V(\|\xi\|_\epsilon)$. We show that for any $u \in \mathbb{R}^d$ such that $\xi = u + v$, $v \neq \xi^{1-\epsilon}$ implies $u \notin U(\|\xi\|_\epsilon)$.

$\|u\|^2 = \|\xi - v\|^2 = \|\xi^\epsilon + (\xi^{1-\epsilon} - v)\|^2 = \|\xi^\epsilon\|^2 + 2\langle \xi^\epsilon, \xi^{1-\epsilon} - v \rangle + \|\xi^{1-\epsilon} - v\|^2$,
hence $\|u\|^2 > \epsilon^2\|\xi\|_\epsilon^2 + 2\langle \xi^\epsilon, \xi^{1-\epsilon} - v \rangle$ because $\|\xi^\epsilon\| = \epsilon\|\xi\|_\epsilon$ and $\|\xi^{1-\epsilon} - v\| > 0$ ($v \neq \xi^{1-\epsilon}$).
Moreover,

$$\begin{aligned} \langle \xi^\epsilon, \xi^{1-\epsilon} - v \rangle &= \sum_{i=1}^d [\text{sign}(\xi_i)(|\xi_i| - (1 - \epsilon)\|\xi\|_\epsilon)_+] [\text{sign}(\xi_i)(|\xi_i| - (|\xi_i| - (1 - \epsilon)\|\xi\|_\epsilon)_+) - v_i] \\ &= \sum_{i=1}^d [(|\xi_i| - (1 - \epsilon)\|\xi\|_\epsilon)_+] [(|\xi_i| - (|\xi_i| - (1 - \epsilon)\|\xi\|_\epsilon)_+) - v_i \text{sign}(\xi_i)] \\ &\geq \sum_{\substack{i=1 \\ |\xi_i| > (1-\epsilon)\|\xi\|_\epsilon}}^d [|\xi_i| - (1 - \epsilon)\|\xi\|_\epsilon] [(1 - \epsilon)\|\xi\|_\epsilon - v_i \text{sign}(\xi_i)] \geq 0. \end{aligned}$$

The last inequality hold because $v \in V(\|\xi\|_\epsilon)$ i.e., $\forall i \in [d]$, $v_i \leq (1 - \epsilon)\|\xi\|_\epsilon$. Finally, $\|u\|^2 > \epsilon^2\|\xi\|_\epsilon^2$ hence the result. \square

Lemma 3. $\{\xi \in \mathbb{R}^d : \|\xi\|_\epsilon \leq \nu\} = \{u + v : u, v \in \mathbb{R}^d, \|u\| \leq \epsilon\nu, \|v\|_\infty \leq (1 - \epsilon)\nu\}$.

Proof. Thanks to Lemma 1, we have $\xi = \xi^\epsilon + \xi^{1-\epsilon}$, $\|\xi^\epsilon\| = \epsilon\|\xi\|_\epsilon$ and $\|\xi^{1-\epsilon}\|_\infty = (1 - \epsilon)\|\xi\|_\epsilon$. Hence, $\|\xi\|_\epsilon \leq \nu$ implies $\|\xi^\epsilon\| \leq \epsilon\nu$ and $\|\xi^{1-\epsilon}\|_\infty \leq (1 - \epsilon)\nu$.

Suppose $\xi = u + v$ such that $\|u\| \leq \epsilon\nu$ and $\|v\|_\infty \leq (1 - \epsilon)\nu$. From the ϵ -decomposition, we have $\|\xi\|_\epsilon = \|\xi^\epsilon\| + \|\xi^{1-\epsilon}\|_\infty$. Moreover, $\|\xi^\epsilon\| \leq \|u\|$ and $\|\xi^{1-\epsilon}\|_\infty \leq \|v\|_\infty$ thanks to Lemma 2. Hence $\|\xi\|_\epsilon \leq \|u\| + \|v\|_\infty \leq \epsilon\nu + (1 - \epsilon)\nu = \nu$. \square

Lemma 4 (Dual norm of the ϵ -norm). Let $\xi \in \mathbb{R}^d$, then $\|\xi\|_\epsilon^D = \epsilon\|\xi\| + (1 - \epsilon)\|\xi\|_1$.

Proof.

$$\begin{aligned} \|\xi\|_\epsilon^D &:= \max_{\|x\|_\epsilon \leq 1} \xi^\top x = \max_{\substack{\|u\| \leq \epsilon \\ \|v\|_\infty \leq 1-\epsilon}} \xi^\top (u + v) \quad \text{thanks to Lemma 3} \\ &= \epsilon \max_{\|u\| \leq 1} \xi^\top u + (1 - \epsilon) \max_{\|v\|_\infty \leq 1} \xi^\top v = \epsilon\|\xi\|^D + (1 - \epsilon)\|\xi\|_\infty^D. \quad \square \end{aligned}$$

Lemma 5. Let $\xi \in \mathbb{R}^d \setminus \{0\}$. Then $\nabla \|\cdot\|_\epsilon(\xi) = \frac{\xi^\epsilon}{\|\xi^\epsilon\|_\epsilon^D}$.

Proof. Let us define $h : \mathbb{R} \times \mathbb{R}^d \mapsto \mathbb{R}$ by $h(\nu, \xi) = \|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\| - \epsilon\nu$. Then we have

$$\begin{aligned} \frac{\partial h}{\partial \nu}(\nu, \xi) &= \frac{\mathcal{S}_{(1-\epsilon)\nu}(\xi)^\top}{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|} \frac{\partial \mathcal{S}_{(1-\epsilon)\nu}(\xi)}{\partial \nu} - \epsilon = -\frac{\mathcal{S}_{(1-\epsilon)\nu}(\xi)^\top}{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|} (1-\epsilon) \text{sign}(\xi) - \epsilon \\ &= -\frac{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|_1}{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|} (1-\epsilon) - \epsilon = -\frac{(1-\epsilon) \|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|_1 + \epsilon \|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|}{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|} \\ &= -\frac{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|_\epsilon^D}{\|\mathcal{S}_{(1-\epsilon)\nu}(\xi)\|} \quad (\text{thanks to Lemma 4}). \end{aligned}$$

By definition of the ϵ -norm, $h(\|\xi\|_\epsilon, \xi) = 0$. Since $\frac{\partial h}{\partial \nu}(\|\xi\|_\epsilon, \xi) = -\frac{\|\xi^\epsilon\|_\epsilon^D}{\epsilon\|\xi\|_\epsilon} \neq 0$, we obtain by applying the Implicit Function Theorem

$$\nabla \|\cdot\|_\epsilon(\xi) \times \frac{\partial h}{\partial \nu}(\|\xi\|_\epsilon, \xi) + \frac{\partial h}{\partial \xi}(\|\xi\|_\epsilon, \xi) = 0 \text{ hence } \nabla \|\cdot\|_\epsilon(\xi) = -\frac{\frac{\partial h}{\partial \xi}(\|\xi\|_\epsilon, \xi)}{\frac{\partial h}{\partial \nu}(\|\xi\|_\epsilon, \xi)}.$$

Moreover, $\frac{\partial h}{\partial \xi}(\|\xi\|_\epsilon, \xi) = \frac{\mathcal{S}_{(1-\epsilon)\|\xi\|_\epsilon}(\xi)}{\|\mathcal{S}_{(1-\epsilon)\|\xi\|_\epsilon}(\xi)\|} = \frac{\xi^\epsilon}{\|\xi^\epsilon\|} = \frac{\xi^\epsilon}{\epsilon\|\xi\|_\epsilon}$ hence the result: $\nabla \|\cdot\|_\epsilon(\xi) = \frac{\xi^\epsilon}{\|\xi^\epsilon\|_\epsilon^D}$. \square

G Implementation

Here we present the ISTA-BC we considered, and provide the GAP safe rules we have implemented. Note that the GAP safe rules we have used are both sequential and dynamical by nature and simply refer to as GAP safe.

Algorithm 2 ISTA-BC with GAP safe rules

Input : $X, y, \epsilon, K, f^{\text{ce}}, (\lambda_t)_{t \in [T-1]}$
 $\forall g \in \mathcal{G}$, compute $L_g = \|X_g\|_2^2$
Compute $\lambda_0 = \lambda_{\max}$ thanks to (12) and Algorithm 1
 $\beta^{\lambda_0} = 0$
for $t \in [T-1]$ **do**
 $\forall g \in \mathcal{G}, \alpha_g \leftarrow \lambda_t / L_g$
 $\beta \leftarrow \beta^{\lambda_{t-1}}$ // Get previous ϵ -solution
 for $k \in [K]$ **do**
 if $k \bmod f^{\text{ce}} = 1$ **then**
 Compute θ thanks to (9) and Algorithm 1.
 Set $\mathcal{R} = \mathcal{B}\left(\theta, \sqrt{\frac{2(P_{\lambda_t, \Omega}(\beta) - D_{\lambda_t}(\theta))}{\lambda_t^2}}\right)$ // Safe sphere
 if $P_{\lambda_t, \Omega}(\beta) - D_{\lambda_t}(\theta) \leq \epsilon$ // Stopping criterion
 then
 $\beta^{\lambda_t} \leftarrow \beta$
 break
 Update $\mathcal{A}_{\text{gp}}(\mathcal{R})$ and $\mathcal{A}_{\text{ft}}(\mathcal{R})$ thanks to Theorem 2
 for $g \in \mathcal{A}_{\text{gp}}(\mathcal{R})$ **do** // Loop over active groups
 for $j \in g \cap \mathcal{A}_{\text{ft}}(\mathcal{R})$ **do** // Loop over active features
 $\beta_j \leftarrow \mathcal{S}_{\tau\alpha_g}\left(\beta_j - \frac{\nabla_j f(\beta)}{L_g}\right)$ // Component-wise Soft-thresholding step
 $\beta_g \leftarrow \mathcal{S}_{(1-\tau)\omega_g\alpha_g}(\beta_g)$ // Block-wise Soft-thresholding step
 Output : $(\beta^{\lambda_t})_{t \in [T-1]}$
