arXiv:1506.03736v2 [stat.ML] 18 Nov 2015

# GAP Safe screening rules for sparse multi-task and multi-class models

Eugène Ndiaye [1], Olivier Fercoq [1], Alexandre Gramfort [1] and Joseph Salmon [1]

[1]LTCI, CNRS, Télécom ParisTech, Université Paris-Saclay,75013, Paris, France

November 17, 2015

**Abstract**

High dimensional regression benefits from sparsity promoting regularizations. Screening rules leverage the known sparsity of the solution by ignoring some variables in the optimization, hence speeding up solvers. When the procedure is proven not to discard features wrongly the rules are said to be *safe*. In this paper we derive new safe rules for generalized linear models regularized with $\ell_1$ and $\ell_1/\ell_2$ norms. The rules are based on duality gap computations and spherical safe regions whose diameters converge to zero. This allows to discard safely more variables, in particular for low regularization parameters. The GAP Safe rule can cope with any iterative solver and we illustrate its performance on coordinate descent for multi-task Lasso, binary and multinomial logistic regression, demonstrating significant speed ups on all tested datasets with respect to previous safe rules.

## 1 Introduction

The computational burden of solving high dimensional regularized regression problem has lead to a vast literature in the last couple of decades to accelerate the algorithmic solvers. With the increasing popularity of $\ell_1$-type regularization ranging from the Lasso [18] or group-Lasso [24] to regularized logistic regression and multi-task learning, many algorithmic methods have emerged to solve the associated optimization problems. Although for the simple $\ell_1$ regularized least square a specific algorithm (*e.g.,* the LARS [8]) can be considered, for more general formulations, penalties, and possibly larger dimension, coordinate descent has proved to be a surprisingly efficient strategy [12].

Our main objective in this work is to propose a technique that can speed-up any solver for such learning problems, and that is particularly well suited for coordinate descent method, thanks to active set strategies.

The *safe rules* introduced by [9] for generalized $\ell_1$ regularized problems, is a set of rules that allows to eliminate features whose associated coefficients are proved to be zero at the optimum. Relaxing the safe rule, one can obtain some more speed-up at the price of possible mistakes. Such heuristic strategies, called *strong rules* [19] reduce the computational cost using an active set strategy, but require difficult post-precessing to check for features possibly wrongly discarded. Another road to speed-up screening method has been the introduction of *sequential safe rules* [21, 23, 22]. The idea is to improve the screening thanks to the computations done for a previous regularization parameter. This scenario is particularly relevant in machine learning, where one computes solutions over a grid of regularization parameters, so as to select the best one (*e.g.,* to perform cross-validation). Nevertheless, such strategies suffer from the same problem as strong rules, since relevant features can be wrongly disregarded: sequential rules usually rely on theoretical quantities that are not known by the solver, but only approximated. Especially, for such rules to work one needs the exact dual optimal solution from the previous regularization parameter.

Recently, the introduction of *safe dynamic rules* [6, 5] has opened a promising venue by letting the screening to be done not only at the beginning of the algorithm, but all along the iterations. Following a

method introduced for the Lasso [11], we generalize this dynamical safe rule, called GAP Safe rules (because it relies on duality gap computation) to a large class of learning problems with the following benefits:

- a unified and flexible framework for a wider family of problems,
- easy to insert in existing solvers,
- proved to be safe,
- more efficient that previous safe rules,
- achieves fast true active set identification.

We introduce our general GAP Safe framework in Section 2. We then specialize it to important machine learning use cases in Section 3. In Section 4 we apply our GAP Safe rules to a multi-task Lasso problem, relevant for brain imaging with magnetoencephalography data, as well as to multinomial logistic regression regularized with $\ell_1/\ell_2$ norm for joint feature selection.

# 2 GAP Safe rules

## 2.1 Model and notations

We denote by $[d]$ the set $\{1, \dots, d\}$ for any integer $d \in \mathbb{N}$, and by $Q^\top$ the transpose of a matrix $Q$. Our observation matrix is $Y \in \mathbb{R}^{n \times q}$ where $n$ represents the number of samples, and $q$ the number of tasks or classes. The design matrix $X = [x^{(1)}, \dots, x^{(p)}] = [x_1, \dots, x_n]^\top \in \mathbb{R}^{n \times p}$ has $p$ explanatory variables (or features) column-wise, and $n$ observations row-wise. The standard $\ell_2$ norm is written $\|\cdot\|_2$, the $\ell_1$ norm $\|\cdot\|_1$, the $\ell_\infty$ norm $\|\cdot\|_\infty$. The $\ell_2$ unit ball is denoted by $\mathcal{B}_2$ (or simply $\mathcal{B}$) and we write $\mathcal{B}(c, r)$ the $\ell_2$ ball with center $c$ and radius $r$. For a matrix $\mathrm{B} \in \mathbb{R}^{p \times q}$, we denote by $\|\mathrm{B}\|_2^2 = \sum_{j=1}^p \sum_{k=1}^q \mathrm{B}_{j,k}^2$ the Frobenius norm, and by $\langle \cdot, \cdot \rangle$ the associated inner product.

We consider the general optimization problem of minimizing a separable function with a group-Lasso regularization. The parameter to recover is a matrix $\mathrm{B} \in \mathbb{R}^{p \times q}$, and for any $j$ in $\mathbb{R}^p$, $\mathrm{B}_{j,:}$ is the $j$-th row of B, while for any $k$ in $\mathbb{R}^q$, $\mathrm{B}_{:,k}$ is the $k$-th column. We would like to find

$$\widehat{\mathrm{B}}^{(\lambda)} \in \underset{\mathrm{B} \in \mathbb{R}^{p \times q}}{\arg\min} \underbrace{\sum_{i=1}^n f_i(x_i^\top \mathrm{B}) + \lambda \Omega(\mathrm{B})}_{P_\lambda(\mathrm{B})} \ , \tag{1}$$

where $f_i : \mathbb{R}^{1 \times q} \mapsto \mathbb{R}$ is a convex function with $1/\gamma$-Lipschitz gradient. So $F : \mathrm{B} \to \sum_{i=1}^n f_i(x_i^\top \mathrm{B})$ is also convex with Lipschitz gradient. The function $\Omega : \mathbb{R}^{p \times q} \mapsto \mathbb{R}_+$ is the $\ell_1/\ell_2$ norm $\Omega(\mathrm{B}) = \sum_{j=1}^p \|\mathrm{B}_{j,:}\|_2$ promoting a few lines of B to be non-zero at a time. The $\lambda$ parameter is a non-negative constant controlling the trade-off between data fitting and regularization.

Some elements of convex analysis used in the following are introduced here. For a convex function $f : \mathbb{R}^d \to [-\infty, +\infty]$ the Fenchel-Legendre transform[1] of $f$, is the function $f^* : \mathbb{R}^d \to [-\infty, +\infty]$ defined by $f^*(u) = \sup_{z \in \mathbb{R}^d} \langle z, u \rangle - f(z)$. The sub-differential of a function $f$ at a point $x$ is denoted by $\partial f(x)$. The dual norm of $\Omega$ is the $\ell_\infty/\ell_2$ norm and reads $\Omega_*(\mathrm{B}) = \max_{j \in [p]} \|B_{j,:}\|_2$.

**Remark 1.** For the ease of reading, all groups are weighted with equal strength, but extension of our results to non-equal weights as proposed in the original group-Lasso [24] paper would be straightforward.

## 2.2 Basic properties

First we recall the associated Fermat's condition and a dual formulation of the optimization problem:

[1] this is also often referred to as the (convex) conjugate of a function

**Theorem 1.** ***Fermat's condition*** *(see [3, Proposition 26.1] for a more general result) For any convex function $f : \mathbb{R}^n \to \mathbb{R}$:*

$$x \in \arg\min_{x \in \mathbb{R}^n} f(x) \Leftrightarrow 0 \in \partial f(x). \tag{2}$$

**Theorem 2** ([9])**.** *A dual formulation of* (1) *is given by*

$$\widehat{\Theta}^{(\lambda)} = \arg\max_{\Theta \in \Delta_X} \underbrace{-\sum_{i=1}^{n} f_i^*(-\lambda \Theta_{i,:})}_{D_\lambda(\Theta)}. \tag{3}$$

*where $\Delta_X = \{\Theta \in \mathbb{R}^{n \times q} : \forall j \in [p], \|{x^{(j)}}^\top \Theta\|_2 \leqslant 1\} = \{\Theta \in \mathbb{R}^{n \times q} : \Omega_*(X^\top \Theta) \leqslant 1\}$. The primal and dual solutions are linked by*

$$\forall i \in [n], \quad \widehat{\Theta}^{(\lambda)}_{i,:} = -\nabla f_i(x_i^\top \widehat{\mathrm{B}}^{(\lambda)})/\lambda. \tag{4}$$

*Furthermore, Fermat's condition reads:*

$$\forall j \in [p], \quad {x^{(j)}}^\top \widehat{\Theta}^{(\lambda)} \in \begin{cases} \left\{ \frac{\hat{\mathrm{B}}^{\lambda}_{j,:}}{\|\hat{\mathrm{B}}^{\lambda}_{j,:}\|_2} \right\}, & \text{if } \widehat{\mathrm{B}}^{(\lambda)}_{j,:} \neq 0, \\ \mathcal{B}_2, & \text{if } \widehat{\mathrm{B}}^{(\lambda)}_{j,:} = 0. \end{cases} \tag{5}$$

**Remark 2.** Contrarily to the primal, the dual problem has a unique solution under our assumption on $f_i$. Indeed, the dual function is strongly concave, hence strictly concave.

**Remark 3.** For any $\Theta \in \mathbb{R}^{n \times q}$ let us introduce $G(\Theta) = [\nabla f_1(\Theta_{1,:})^\top, \ldots, \nabla f_n(\Theta_{n,:})^\top] \in \mathbb{R}^{n \times q}$. Then the primal/dual link can be written $\widehat{\Theta}^{(\lambda)} = -G(X\widehat{\mathrm{B}}^{(\lambda)})/\lambda$ .

### 2.3 Critical parameter: $\lambda_{\max}$

For $\lambda$ large enough the solution of the primal problem is simply 0. Thanks to the Fermat's rule (2), 0 is optimal if and only if $-\nabla F(0)/\lambda \in \partial\Omega(0)$. Thanks to the property of the dual norm $\Omega_*$, this is equivalent to $\Omega_*(\nabla F(0)/\lambda) \leqslant 1$ where $\Omega_*$ is the dual norm of $\Omega$. Since $\nabla F(0) = X^\top G(0)$, 0 is a primal solution of $P_\lambda$ if and only if $\lambda \geqslant \lambda_{\max} := \max_{j \in [p]} \|{x^{(j)}}^\top G(0)\|_2 = \Omega_*(X^\top G(0))$.

This development shows that for $\lambda \geqslant \lambda_{\max}$, Problem (1) is trivial. So from now on, we will only focus on the case where $\lambda \leqslant \lambda_{\max}$.

### 2.4 Screening rules description

Safe screening rules rely on a simple consequence of the Fermat's condition:

$$\|{x^{(j)}}^\top \widehat{\Theta}^{(\lambda)}\|_2 < 1 \Rightarrow \widehat{\mathrm{B}}^{(\lambda)}_{j,:} = 0 \ . \tag{6}$$

Stated in such a way, this relation is useless because $\widehat{\Theta}^{(\lambda)}$ is unknown (unless $\lambda > \lambda_{\max}$). However, it is often possible to construct a set $\mathcal{R} \subset \mathbb{R}^{n \times q}$, called a *safe region*, containing it. Then, note that

$$\max_{\Theta \in \mathcal{R}} \|{x^{(j)}}^\top \Theta\|_2 < 1 \Rightarrow \widehat{\mathrm{B}}^{(\lambda)}_{j,:} = 0 \ . \tag{7}$$

The so called *safe screening rules* consist in removing the variable $j$ from the problem whenever the previous test is satisfied, since $\widehat{\mathrm{B}}^{(\lambda)}_{j,:}$ is then guaranteed to be zero. This property leads to considerable speed-up in practice especially with active sets strategies, see for instance [11] for the Lasso case. A natural goal is to find safe regions as narrow as possible: smaller safe regions can only increase the number of screened out variables. However, complex regions could lead to a computational burden limiting the benefit of screening. Hence, we focus on constructing $\mathcal{R}$ satisfying the trade-off:

- $\mathcal{R}$ is as small as possible and contains $\widehat{\Theta}^{(\lambda)}$.
- Computing $\max_{\Theta \in \mathcal{R}} \|{x^{(j)}}^\top \Theta\|_2$ is cheap.

## 2.5 Spheres as safe regions

Various shapes have been considered in practice for the set $\mathcal{R}$ such as balls (referred to as spheres) [9], domes [11] or more refined sets (see [23] for a survey). Here we consider the so-called "sphere regions" choosing a ball $\mathcal{R} = \mathcal{B}(c, r)$ as a safe region. One can easily obtain a control on $\max_{\Theta \in \mathcal{B}(c,r)} \|x^{(j)^\top}\Theta\|_2$ by extending the computation of the support function of a ball [11, Eq. (9)] to the matrix case: $\max_{\Theta \in \mathcal{B}(c,r)} \|x^{(j)^\top}\Theta\|_2 \leqslant \|x^{(j)^\top}c\|_2 + r\|x^{(j)}\|_2$ .

Note that here the center $c$ is a matrix in $\mathbb{R}^{p\times q}$. We can now state the safe sphere test:

$$\text{Sphere test:} \qquad \text{If} \quad \|x^{(j)^\top}c\|_2 + r\|x^{(j)}\|_2 < 1, \quad \text{then} \quad \widehat{\mathrm{B}}^{(\lambda)}_{j,:} = 0. \tag{8}$$

## 2.6 GAP Safe rule description

In this section we derive a GAP Safe screening rule extending the one introduced in [11]. For this, we rely on the strong convexity of the dual objective function and on weak duality.

**Finding a radius:** Remember that $\forall i \in [n], f_i$ is differentiable with a $1/\gamma$-Lipschitz gradient. As a consequence, $\forall i \in [n], f_i^*$ is $\gamma$-strongly convex [14, Theorem 4.2.2, p. 83] and so $D_\lambda$ is $\gamma\lambda^2$-strongly concave:

$$\forall (\Theta_1, \Theta_2) \in \mathbb{R}^{n\times q} \times \mathbb{R}^{n\times q}, \quad D_\lambda(\Theta_2) \leqslant D_\lambda(\Theta_1) + \langle \nabla D_\lambda(\Theta_1), \Theta_2 - \Theta_1\rangle - \frac{\gamma\lambda^2}{2}\|\Theta_1 - \Theta_2\|^2.$$

Specifying the previous inequality for $\Theta_1 = \widehat{\Theta}^{(\lambda)}, \Theta_2 = \Theta \in \Delta_X$, one has

$$D_\lambda(\Theta) \leqslant D_\lambda(\widehat{\Theta}^{(\lambda)}) + \langle \nabla D_\lambda(\widehat{\Theta}^{(\lambda)}), \Theta - \widehat{\Theta}^{(\lambda)}\rangle - \frac{\gamma\lambda^2}{2}\|\widehat{\Theta}^{(\lambda)} - \Theta\|^2.$$

By definition, $\widehat{\Theta}^{(\lambda)}$ maximizes $D_\lambda$ on $\Delta_X$, so we have: $\langle \nabla D_\lambda(\widehat{\Theta}^{(\lambda)}), \Theta - \widehat{\Theta}^{(\lambda)}\rangle \leqslant 0$. This implies

$$D_\lambda(\Theta) \leqslant D_\lambda(\widehat{\Theta}^{(\lambda)}) - \frac{\gamma\lambda^2}{2}\|\widehat{\Theta}^{(\lambda)} - \Theta\|^2.$$

By weak duality $\forall \mathrm{B} \in \mathbb{R}^{p\times q}, D_\lambda(\widehat{\Theta}^{(\lambda)}) \leqslant P_\lambda(\mathrm{B})$, so : $\forall \mathrm{B} \in \mathbb{R}^{p\times q}, \forall \Theta \in \Delta_X, D_\lambda(\Theta) \leqslant P_\lambda(\mathrm{B}) - \frac{\gamma\lambda^2}{2}\|\widehat{\Theta}^{(\lambda)} - \Theta\|^2$, and we deduce the following theorem:

**Theorem 3.**

$$\forall \mathrm{B} \in \mathbb{R}^{p\times q}, \forall \Theta \in \Delta_X, \quad \left\|\widehat{\Theta}^{(\lambda)} - \Theta\right\|_2 \leqslant \sqrt{\frac{2(P_\lambda(\mathrm{B}) - D_\lambda(\Theta))}{\gamma\lambda^2}} =: \hat{r}_\lambda(\mathrm{B}, \Theta). \tag{9}$$

Provided one knows a dual feasible point $\Theta \in \Delta_X$ and a $\mathrm{B} \in \mathbb{R}^{p\times q}$ , it is possible to construct a safe sphere with radius $\hat{r}_\lambda(\mathrm{B}, \Theta)$ centered on $\Theta$. We now only need to build a (relevant) dual point to center such a ball. Results from Section 2.3, ensure that $-G(0)/\lambda_{\max} \in \Delta_X$, but it leads to a static rule, a introduced in [9]. We need a dynamic center to improve the screening as the solver proceeds.

**Finding a center:** Remember that $\widehat{\Theta}^{(\lambda)} = -G(X\widehat{\mathrm{B}}^{(\lambda)})/\lambda$. Now assume that one has a converging algorithm for the primal problem, *i.e.,* $\mathrm{B}_k \to \widehat{\mathrm{B}}^{(\lambda)}$. Hence, a natural choice for creating a dual feasible point $\Theta_k$ is to choose it proportional to $-G(X\mathrm{B}_k)$, for instance by setting:

$$\Theta_k = \begin{cases} \frac{R_k}{\lambda}, & \text{if } \Omega_*(X^\top R_k) \leqslant \lambda, \\ \frac{R_k}{\Omega_*(X^\top R_k)}, & \text{otherwise.} \end{cases} \quad \text{where } R_k = -G(X\mathrm{B}_k) \ . \tag{10}$$

A refined method consists in solving the one dimensional problem: $\arg\max_{\Theta \in \Delta_X \cap \mathrm{Span}(R_k)} D_\lambda(\Theta)$. In the Lasso and group-Lasso case [5, 6, 11] such a step is simply a projection on the intersection of a line and the (polytope) dual set and can be computed efficiently. However for logistic regression the computation is more involved, so we have opted for the simpler solution in Equation (10). This still provides converging safe rules (see Proposition 1).

### Dynamic GAP Safe rule summarized

We can now state our dynamical GAP Safe rule at the $k$-th step of an iterative solver:

1. Compute $\mathrm{B}_k$, and then obtain $\Theta_k$ and $\hat{r}_\lambda(\mathrm{B}_k, \Theta_k)$ using (10).
2. If $\|x^{(j)^\top}\Theta_k\|_2 + \hat{r}_\lambda(\mathrm{B}_k, \Theta_k)\|x^{(j)}\|_2 < 1$, then set $\widehat{\mathrm{B}}^{(\lambda)}_{j,:} = 0$ and remove $x^{(j)}$ from $X$.

Dynamic safe screening rules are more efficient than existing methods in practice because they can increase the ability of screening as the algorithm proceeds. Since one has sharper and sharper dual regions available along the iterations, support identification is improved. Provided one relies on a primal converging algorithm, one can show that the dual sequence we propose is converging too.

The convergence of the primal is unaltered by our GAP Safe rule: screening out unnecessary coefficients of $\mathrm{B}_k$ can only decrease its distance with its original limits. Moreover, a practical consequence is that one can observe surprising situations where lowering the tolerance of the solver can reduce the computation time. This can happen for sequential setups.

**Proposition 1.** *Let* $\mathrm{B}_k$ *be the current estimate of* $\widehat{\mathrm{B}}^{(\lambda)}$ *and* $\Theta_k$ *defined in Eq.* (10) *be the current estimate of* $\widehat{\Theta}^{(\lambda)}$. *Then* $\lim_{k\to+\infty} \mathrm{B}_k = \widehat{\mathrm{B}}^{(\lambda)}$ *implies* $\lim_{k\to+\infty} \Theta_k = \widehat{\Theta}^{(\lambda)}$.

Note that if the primal sequence is converging to the optimal, our dual sequence is also converging. But we know that the radius of our safe sphere is $(2(P_\lambda(\mathrm{B}_k) - D_\lambda(\Theta_k))/(\gamma\lambda^2))^{1/2}$. By strong duality, this radius converges to 0, hence we have certified that our GAP Safe regions sequence $\mathcal{B}(\Theta_k, \hat{r}_\lambda(\mathrm{B}_k, \Theta_k))$ is a converging safe rules (in the sense introduced in [11, Definition 1]).

**Remark 4.** The active set obtained by our GAP Safe rule (*i.e.,* the indexes of non screened-out variables) converges to the equicorrelation set [20] $\mathcal{E}_\lambda := \{j \in p : \|x^{(j)^\top}\widehat{\Theta}^{(\lambda)}\|_2 = 1\}$, allowing us to early identify relevant features (see Proposition 2 in the supplementary material for more details).

# 3 Special cases of interest

We now specialize our results to relevant supervised learning problems, see also Table 1.

## 3.1 Lasso

In the Lasso case $q = 1$, the parameter is a vector: $\mathrm{B} = \beta \in \mathbb{R}^p$, $F(\beta) = 1/2\|y - X\beta\|_2^2 = \sum_{i=1}^n (y_i - x_i^\top\beta)^2$, meaning that $f_i(z) = (y_i - z)^2/2$ and $\Omega(\beta) = \|\beta\|_1$.

## 3.2 $\ell_1/\ell_2$ multi-task regression

In the multi-task Lasso, which is a special case of group-Lasso, we assume that the observation is $Y \in \mathbb{R}^{n\times q}$, $F(\mathrm{B}) = \frac{1}{2}\|Y - X\mathrm{B}\|_2^2 = \frac{1}{2}\sum_{i=1}^n \|Y_{i,:} - x_i^\top\mathrm{B}\|_2^2$ (*i.e.,* $f_i(z) = \|Y_{i,:} - z\|^2/2$) and $\Omega(\mathrm{B}) = \sum_{j=1}^p \|\mathrm{B}_{j,:}\|_2$. In signal processing, this model is also referred to as Multiple Measurement Vector (MMV) problem. It allows to jointly select the same features for multiple regression tasks [1, 2].

**Remark 5.** Our framework could encompass easily the case of non-overlapping groups with various size and weights presented in [6]. Since our aim is mostly for multi-task and multinomial applications, we have rather presented a matrix formulation.

### 3.3 $\ell_1$ regularized logistic regression

Here, we consider the formulation given in [7, Chapter 3] for the two classes logistic regression. In such a context, one observes for each $i \in [n]$ a class label $c_i \in \{1, 2\}$. This information can be recast as $y_i = \mathbb{1}_{\{c_i=1\}}$, and it is then customary to minimize (1) where

$$F(\beta) = \sum_{i=1}^{n} \left( -y_i x_i^\top \beta + \log\left(1 + \exp\left(x_i^\top \beta\right)\right)\right), \tag{11}$$

with $\mathrm{B} = \beta \in \mathbb{R}^p$ (*i.e.*, $q = 1$), $f_i(z) = -y_i z + \log(1 + \exp(z))$ and the penalty is simply the $\ell_1$ norm: $\Omega(\beta) = \|\beta\|_1$. Let us introduce Nh, the (binary) negative entropy function defined by [2]:

$$\mathrm{Nh}(x) = \begin{cases} x\log(x) + (1-x)\log(1-x), & \text{if } x \in [0,1] \ , \\ +\infty, & \text{otherwise } . \end{cases} \tag{12}$$

Then, one can easily check that $f_i^*(z_i) = \mathrm{Nh}(z_i + y_i)$ and $\gamma = 4$.

### 3.4 $\ell_1/\ell_2$ multinomial logistic regression

We adapt the formulation given in [7, Chapter 3] for the multinomial regression. In such a context, one observes for each $i \in [n]$ a class label $c_i \in \{1, \dots, q\}$. This information can be recast into a matrix $Y \in \mathbb{R}^{n\times q}$ filled by 0's and 1's: $Y_{i,k} = \mathbb{1}_{\{c_i=k\}}$. In the same spirit as the multi-task Lasso, a matrix $\mathrm{B} \in \mathbb{R}^{p\times q}$ is formed by $q$ vectors encoding the hyperplanes for the linear classification. The multinomial $\ell_1/\ell_2$ regularized regression reads:

$$F(\mathrm{B}) = \sum_{i=1}^{n} \left( \sum_{k=1}^{q} -Y_{i,k} x_i^\top \mathrm{B}_{:,k} + \log\left( \sum_{k=1}^{q} \exp\left(x_i^\top \mathrm{B}_{:,k}\right)\right)\right), \tag{13}$$

with $f_i(z) = \sum_{k=1}^{q} -Y_{i,k} z_k + \log\left(\sum_{k=1}^{q} \exp\left(z_k\right)\right)$ to recover the formulation as in (1). Let us introduce NH, the negative entropy function defined by (still with the convention $0\log(0) = 0$)

$$\mathrm{NH}(x) = \begin{cases} \sum_{i=1}^{q} x_i \log(x_i), & \text{if } x \in \Sigma_q = \{x \in \mathbb{R}_+^q : \sum_{i=1}^{q} x_i = 1\}, \\ +\infty, & \text{otherwise.} \end{cases} \tag{14}$$

Again, one can easily check that $f_i^*(z) = \mathrm{NH}(z + Y_{i,:})$ and $\gamma = 1$.

**Remark 6.** For multinomial logistic regression, $D_\lambda$ implicitly encodes the additional constraint $\Theta \in \mathrm{dom} D_\lambda = \{\Theta' \ : \ \forall i \in [n], -\lambda\Theta'_{i,:} + Y_{i,:} \in \Sigma_q\}$ where $\Sigma_q$ is the $q$ dimensional simplex, see (14). As 0 and $R_k/\lambda$ both belong to this set, any convex combination of them, such as $\Theta_k$ defined in (10), satisfies this additional constraint.

**Remark 7.** The intercept has been neglected in our models for simplicity. Our GAP Safe framework can also handle such a feature at the cost of more technical details (by adapting the results from [15] for instance). However, in practice, the intercept can be handled in the present formulation by adding a constant column to the design matrix $X$. The intercept is then regularized. However, if the constant is set high enough, regularization is small and experiments show that it has little to no impact for high-dimensional problems. This is the strategy used by the Liblinear package [10].

[2] with the convention $0\log(0) = 0$

| | Lasso | Multi-task regr. | Logistic regr. | Multinomial regr. |
|---|---|---|---|---|
| $f_i(z)$ | $\frac{(y_i - z)^2}{2}$ | $\frac{\|Y_{i,:} - z\|^2}{2}$ | $\log(1 + \mathrm{e}^z) - y_i z$ | $\log\big(\sum_{k=1}^{q} \mathrm{e}^{z_k}\big) - \sum_{k=1}^{q} Y_{i,k} z_k$ |
| $f_i^*(u)$ | $\frac{(y_i - u)^2 - y_i^2}{2}$ | $\frac{\|Y_{i,:} - u\|^2 - \|Y_{i,:}\|_2^2}{2}$ | $\mathrm{Nh}(u + y_i)$ | $\mathrm{NH}(u + Y_{i,:})$ |
| $\Omega(\mathrm{B})$ | $\|\beta\|_1$ | $\sum_{j=1}^{p} \|\mathrm{B}_{j,:}\|_2$ | $\|\beta\|_1$ | $\sum_{j=1}^{p} \|\mathrm{B}_{j,:}\|_2$ |
| $\lambda_{\max}$ | $\|X^\top y\|_\infty$ | $\Omega_*(X^\top Y)$ | $\|X^\top(\mathbf{1}_n/2 - y)\|_\infty$ | $\Omega_*(X^\top(\mathbf{1}_{n\times q}/q - Y))$ |
| $G(\Theta)$ | $\theta - y$ | $\Theta - Y$ | $\frac{\mathrm{e}^z}{1+\mathrm{e}^z} - y$ | $\mathrm{RowNorm}(\mathrm{e}^\Theta) - Y$ |
| $\gamma$ | 1 | 1 | 4 | 1 |

Table 1: Useful ingredients for computing GAP Safe rules. We have used lower case to indicate when the parameters are vectorial (*i.e.,* $q = 1$). The function RowNorm consists in normalizing a (non-negative) matrix row-wise, such that each row sums to one.

# 4 Experiments

In this section we present results obtained with the GAP Safe rule. Results are on high dimensional data, both dense and sparse. Implementation have been done in Python and Cython for low critical parts. They are based on the multi-task Lasso implementation of Scikit-Learn [17] and coordinate descent logistic regression solver in the Lightning software [4]. In all experiments, the coordinate descent algorithm used follows the pseudo code from [11] with a screening step every 10 iterations.

Note that we have not performed comparison with the sequential screening rule commonly acknowledge as the state-of-the-art "safe" screening rule (such as th EDDP+ [21]), since we can show that this kind of rule is not safe. Indeed, the stopping criterion is based on dual gap accuracy, and comparisons would be unfair since such methods sometimes do not converge to the prescribed accuracy. This is backed-up by a counter example given in the supplementary material. Nevertheless, modifications of such rules, inspired by our GAP Safe rules, can make them safe. However the obtained sequential rules are still outperformed by our dynamic strategies (see Figure 2 for an illustration).

## 4.1 $\ell_1/\ell_2$ multi-task regression

To demonstrate the benefit of the GAP Safe screening rule for a multi-task Lasso problem we used neuroimaging data. Electroencephalography (EEG) and magnetoencephalography (MEG) are brain imaging modalities that allow to identify active brain regions. The problem to solve is a multi-task regression problem with squared loss where every task corresponds to a time instant. Using a multi-task Lasso one can constrain the recovered sources to be identical during a short time interval [13]. This corresponds to a temporal stationary assumption. In this experiment we used a joint MEG/EEG data with 301 MEG and 59 EEG sensors leading to $n = 360$. The number of possible sources is $p = 22,494$ and the number of time instants $q = 20$. With a 1 kHz sampling rate it is equivalent to say that the sources stay the same for 20 ms.

Results are presented in Figure 1. The GAP Safe rule is compared with the dynamic safe rule from [6]. The experimental setup consists in estimating the solutions of the multi-task Lasso problem for 100 values of $\lambda$ on a logarithmic grid from $\lambda_{\max}$ to $\lambda_{\max}/10^3$. For the experiments on the left a fixed number of iterations from 2 to $2^{11}$ is allowed for each $\lambda$. The fraction of active variables is reported. Figure 1 illustrates that the GAP Safe rule screens out much more variables than the compared method, as well as the converging nature of our safe regions. Indeed, the more iterations performed the more the rule allows to screen variables. On

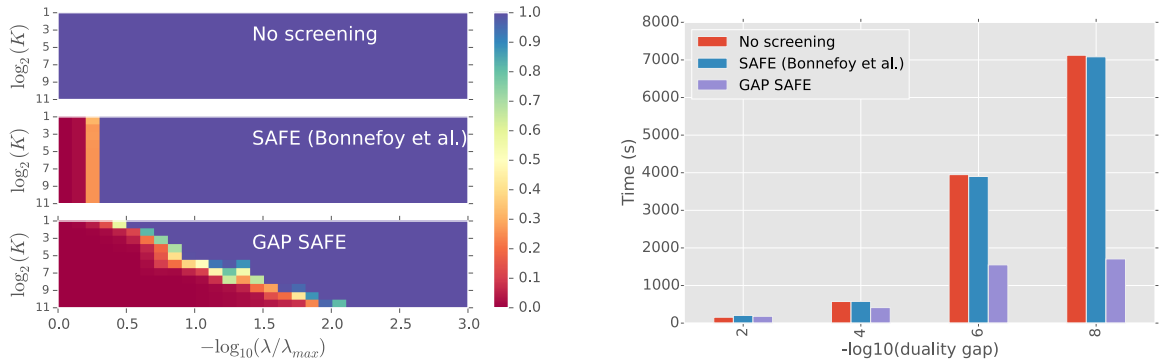


Figure 1: Experiments on MEG/EEG brain imaging dataset (dense data with $n = 360$, $p = 22494$ and $q = 20$). On the left: fraction of active variables as a function of $\lambda$ and the number of iterations $K$. The GAP Safe strategy has a much longer range of $\lambda$ with (red) small active sets. On the right: Computation time to reach convergence using different screening strategies.

the right, computation time confirms the effective speed-up. Our rule significantly improves the computation time for all duality gap tolerance from $10^{-2}$ to $10^{-8}$, especially when accurate estimates are required, *e.g.,* for feature selection.

## 4.2 $\ell_1$ binary logistic regression

Results on the Leukemia dataset are reported in Figure 2. We compare the dynamic strategy of GAP Safe to a sequential and non dynamic rule such as Slores [22]. We do not compare to the actual Slores rule as it requires the previous dual optimal solution, which is not available. Slores is indeed not a safe method (see Section B in the supplementary materials). Nevertheless one can observe that dynamic strategies outperform pure sequential one, see Section C in the supplementary material).

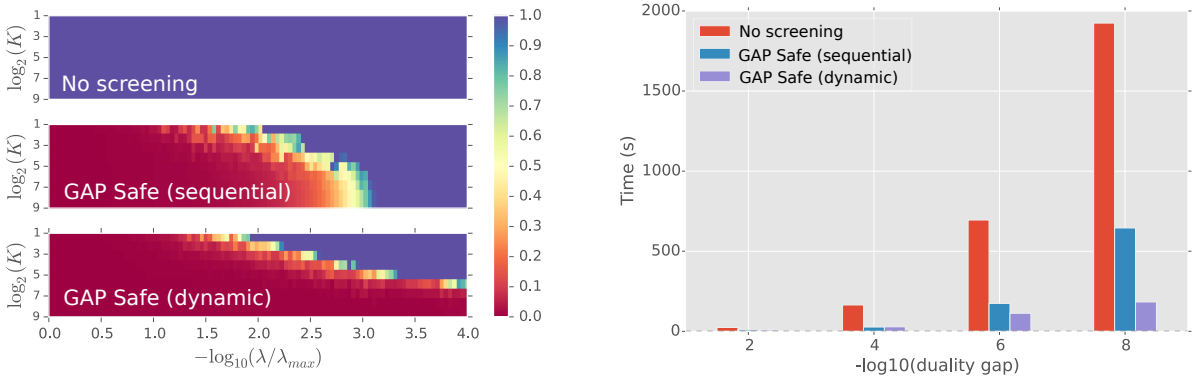


Figure 2: $\ell_1$ regularized binary logistic regression on the Leukemia dataset ($n = 72$ ; $m = 7{,}129$ ; $q = 1$). Simple sequential and full dynamic screening GAP Safe rules are compared. On the left: fraction of the variables that are active. Each line corresponds to a fixed number of iterations for which the algorithm is run. On the right: computation times needed to solve the logistic regression path to desired accuracy with 100 values of $\lambda$.

### 4.3 $\ell_1/\ell_2$ multinomial logistic regression

We also applied GAP Safe to an $\ell_1/\ell_2$ multinomial logistic regression problem on a sparse dataset. Data are bag of words features extracted from the News20 dataset (TF-IDF removing English stop words and words occurring only once or more than 95% of the time). One can observe on Figure 3 the dynamic screening and its benefit as more iterations are performed. GAP Safe leads to a significant speedup: to get a duality gap smaller than $10^{-2}$ on the 100 values of $\lambda$, we needed 1,353 s without screening and only 485 s when GAP Safe was activated.

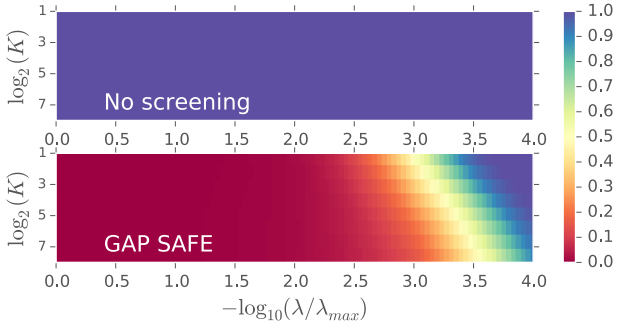


Figure 3: Fraction of the variables that are active for $\ell_1/\ell_2$ regularized multinomial logistic regression on 3 classes of the News20 dataset (sparse data with $n$ = 2,757 ; $m$ = 13,010 ; $q = 3$). Computation was run on the best 10% of the features using $\chi^2$ univariate feature selection [16]. Each line corresponds to a fixed number of iterations for which the algorithm is run.

## 5 Conclusion

This contribution detailed new safe rules for accelerating algorithms solving generalized linear models regularized with $\ell_1$ and $\ell_1/\ell_2$ norms. The rules proposed are safe, easy to implement, dynamic and converging, allowing to discard significantly more variables than alternative safe rules. The positive impact in terms of computation time was observed on all tested datasets and demonstrated here on a high dimensional regression task using brain imaging data as well as binary and multiclass classification problems on dense and sparse data. Extensions to other generalized linear model, *e.g.,* Poisson regression, are expected to reach the same conclusion. Future work could investigate optimal screening frequency, determining when the screening has correctly detected the support.

## Acknowledgment

We acknowledge the support from Chair Machine Learning for Big Data at Télécom ParisTech and from the Orange/Télécom ParisTech think tank phi-TAB. This work benefited from the support of the "FMJH Program Gaspard Monge in optimization and operation research", and from the support to this program from EDF.

## References

[1] A. Argyriou, T. Evgeniou, and M. Pontil. Multi-task feature learning. In *NIPS*, pages 41–48, 2006.

[2] A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

[3] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, New York, 2011.

[4] M. Blondel, K. Seki, and K. Uehara. Block coordinate descent algorithms for large-scale sparse multiclass classification. *Machine Learning*, 93(1):31–52, 2013.

[5] A. Bonnefoy, V. Emiya, L. Ralaivola, and R. Gribonval. A dynamic screening principle for the lasso. In *EUSIPCO*, 2014.

[6] 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.

[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] B. Efron, T. Hastie, I. M. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004. With discussion, and a rejoinder by the authors.

[9] L. El Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination in sparse supervised learning. *J. Pacific Optim.*, 8(4):667–698, 2012.

[10] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *J. Mach. Learn. Res.*, 9:1871–1874, 2008.

[11] O. Fercoq, A. Gramfort, and J. Salmon. Mind the duality gap: safer rules for the lasso. In *ICML*, pages 333–342, 2015.

[12] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.

[13] A. Gramfort, M. Kowalski, and M. Hämäläinen. Mixed-norm estimates for the M/EEG inverse problem using accelerated gradient methods. *Phys. Med. Biol.*, 57(7):1937–1961, 2012.

[14] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex analysis and minimization algorithms. II*, volume 306. Springer-Verlag, Berlin, 1993.

[15] K. Koh, S.-J. Kim, and S. Boyd. An interior-point method for large-scale l1-regularized logistic regression. *J. Mach. Learn. Res.*, 8(8):1519–1555, 2007.

[16] C. D. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, USA, 1999.

[17] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830, 2011.

[18] R. Tibshirani. Regression shrinkage and selection via the lasso. *JRSSB*, 58(1):267–288, 1996.

[19] 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.

[20] R. J. Tibshirani. The lasso problem and uniqueness. *Electron. J. Stat.*, 7:1456–1490, 2013.

[21] J. Wang, P. Wonka, and J. Ye. Lasso screening rules via dual polytope projection. *arXiv preprint arXiv:1211.3966*, 2012.

[22] J. Wang, J. Zhou, J. Liu, P. Wonka, and J. Ye. A safe screening rule for sparse logistic regression. In *NIPS*, pages 1053–1061, 2014.

[23] Z. J. Xiang, Y. Wang, and P. J. Ramadge. Screening tests for lasso problems. *arXiv preprint arXiv:1405.4897*, 2014.

[24] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *JRSSB*, 68(1):49–67, 2006.

# Supplementary Material

# A Proofs

## A.1 Proof of variable identification

**Proposition 2.** *There exists $k_0 \in \mathbb{N}$ such that for all $k \geqslant k_0$, an index $j \in [p]$ is screened out by the GAP Safe rule if and only if $j \in \mathcal{E}_\lambda := \{j \in p : \|x^{(j)\top} \widehat{\Theta}^{(\lambda)}\|_2 = 1\}$.*

*Proof.* For simplicity we use the notation $\mathcal{R}_k = \mathcal{B}(\Theta_k, \hat{r}_\lambda(\mathrm{B}_k, \Theta_k))$ for the safe region at step $k$. Define $\max_{j \notin \mathcal{E}_\lambda} |x^{(j)\top} \widehat{\Theta}^{(\lambda)}| = t < 1$. Fix $\epsilon > 0$ such that $\epsilon < (1-t)/(\max_{j \notin \mathcal{E}_\lambda} \|x^{(j)}\|)$. As $\Theta_k$ is converging to $\widehat{\Theta}^{(\lambda)}$, and $\lim_{k \to \infty} \hat{r}_\lambda(\mathrm{B}_k, \Theta_k) = 0$, there exists $k_0 \in \mathbb{N}$ such that $\forall k \geqslant k_0, \forall \Theta \in \mathcal{R}_k, \|\Theta - \widehat{\Theta}^{(\lambda)}\| \leqslant \epsilon$. Hence, for any $j \notin \mathcal{E}_\lambda$ and any $\Theta \in \mathcal{R}_k$, $|x^{(j)\top}(\Theta - \widehat{\Theta}^{(\lambda)})| \leqslant (\max_{j \notin \mathcal{E}_\lambda} \|x^{(j)}\|)\|\Theta - \widehat{\Theta}^{(\lambda)}\| \leqslant (\max_{j \notin \mathcal{E}_\lambda} \|x^{(j)}\|)\epsilon$. Using the triangle inequality, one gets

$$\begin{aligned} |x^{(j)\top}\Theta| \leqslant & (\max_{j \notin \mathcal{E}_\lambda} \|x^{(j)}\|)\epsilon + \max_{j \notin \mathcal{E}_\lambda} |x^{(j)\top} \widehat{\Theta}^{(\lambda)}| \\ \leqslant & (\max_{j \notin \mathcal{E}_\lambda} \|x^{(j)}\|)\epsilon + t < 1, \end{aligned}$$

provided that $\epsilon < (1-t)/(\max_{j \notin \mathcal{E}_\lambda} \|x^{(j)}\|)$. Hence, for all $k \geqslant k_0, j \notin \mathcal{E}_\lambda$ implies that $j$ is screened out by the GAP Safe rule thanks to the last inequality. For the reverse inclusion take $j \in \mathcal{E}_\lambda$, *i.e.,* $|x^{(j)\top} \widehat{\Theta}^{(\lambda)}| = 1$. Since by construction of our GAP Safe screening rule $\forall k \in \mathbb{N}, \widehat{\Theta}^{(\lambda)} \in \mathcal{R}_k$, then $j \in \{j' \in [p] : \max_{\Theta \in \mathcal{R}_k} |x^{(j')\top} \Theta| \geqslant 1\}$. This means that the variable $j$ can not be eliminated by our safe rule, and we have shown that in the limit we have exactly identified the equicorrelation set. □

## A.2 Proof that the GAP Safe rule is converging (Proposition 1)

*Proof.* We consider two cases.

First let us assume that $\theta_k = R_k/\Omega_*(X^\top G(X\mathrm{B}_k))$

$$\begin{aligned} \left\|\Theta_k - \widehat{\Theta}^{(\lambda)}\right\|_2 &= \left\| \frac{-G(X\mathrm{B}_k)}{\Omega_*(X^\top G(X\mathrm{B}_k))} + \frac{1}{\lambda} G(X\widehat{\mathrm{B}}^{(\lambda)}) \right\|_2 \\ &\leqslant \left\| \frac{G(X\mathrm{B}_k)}{\lambda} - \frac{G(X\mathrm{B}_k)}{\Omega_*(X^\top G(X\mathrm{B}_k))} \right\|_2 + \left\| \frac{G(X\widehat{\mathrm{B}}^{(\lambda)}) - G(X\mathrm{B}_k)}{\lambda} \right\|_2 \\ &\leqslant \left| \frac{1}{\lambda} - \frac{1}{\Omega_*(X^\top G(X\mathrm{B}_k))} \right| \|G(X\mathrm{B}_k)\|_2 + \left\| \frac{G(X\widehat{\mathrm{B}}^{(\lambda)}) - G(X\mathrm{B}_k)}{\lambda} \right\|_2 \end{aligned}$$

The second term converges to zero whenever $\mathrm{B}_k \to \widehat{\mathrm{B}}^{(\lambda)}$ since G is continuous (it is $\gamma$-Lipschitz). For the first term, note that $\Omega_*(X^\top G(X\mathrm{B}_k)) \to \Omega_*(X^\top G(X\widehat{\mathrm{B}}^{(\lambda)})) = \lambda \Omega_*(X^\top \widehat{\Theta}^{(\lambda)}) = \lambda$ (thanks to the primal/dual link, and that $\widehat{\Theta}^{(\lambda)}$ is dual feasible). Then, as $G$ is a Lipschitz function and all norms are equivalent in a finite dimension space, the right hand side converges to zero in the previous inequality, and the results stated follows.

In the second case $\Theta_k = R_k/\lambda$, so $\left\|\Theta_k - \widehat{\Theta}^{(\lambda)}\right\|_2 = \left\| \frac{-G(X\mathrm{B}_k) + G(X\widehat{\mathrm{B}}^{(\lambda)})}{\lambda} \right\|_2$ and the proof proceeds as in the first case.

□

# B EDPP is not safe

In the two last sections, we present a study on the EDDP method [21], a screening rule that relies on the dual optimal point obtained for the previous $\lambda$ in the path. Note that the same conclusion would hold true for generalization of the sequential approach given in [22], as well as for any other screening rule that needs exact dual solution at one step. To simplify the reading we use the vectorial (with no capital letters) notation used earlier. In the remainder we consider $\lambda_0 = \lambda_{\max}$ and a non-increasing sequence of $T-1$ tuning parameters $(\lambda_t)_{t\in[T-1]}$ in $(0, \lambda_{\max})$. In practice, we choose the common grid [7][2.12.1]): $\lambda_t = \lambda_0 10^{-\delta t/(T-1)}$. Wang *et al.* [21] proposed a sequential screening rule based on properties of the projection onto a convex set. Their rule is based on the exact knowledge of the true optimal solution for the previous parameter. Such a rule can be used to compute $\hat{\theta}^{(\lambda_1)}$ since $\hat{\theta}^{(\lambda_0)} = y/\lambda_0$ $(= y/\lambda_{\max})$ is known. However for $t > 1$, $\hat{\theta}^{(\lambda_t)}$ is only known approximately and the rules introduced in [21] are not safe anymore: some active groups may be wrongly disregarded if one does not use the exact value of $\hat{\theta}^{(\lambda_t)}$.

We first first recall the property they proved. Then, we give a counter-example that shows that the rule is indeed not safe. In Section C, we propose to modify their rule in order to make it safe in all cases.

Recall that in this case $q = 1$, the parameters are vectors: $\mathrm{B} = \beta \in \mathbb{R}^p$ and $\Theta = \theta \in \mathbb{R}^n$.

**Proposition 3** ([21, Theorem 19])**.** *Assume that $\lambda_{t-1} < \lambda_{\max}$, then the dual optimal solution of the group-Lasso with parameter $\lambda_t$, satisfies*

$$\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\big(\hat{\theta}^{(\lambda_{t-1})} + \frac{1}{2} v^{\perp}(\lambda_{t-1}, \lambda_t), \frac{1}{2} \left\| v^{\perp}(\lambda_{t-1}, \lambda_t) \right\|_2 \big) \tag{15}$$

*where*

$$v^{\perp}(\lambda_{t-1}, \lambda_t) = \frac{y}{\lambda_t} - \hat{\theta}^{(\lambda_{t-1})} - \alpha[\hat{\theta}^{(\lambda_{t-1})}](\frac{y}{\lambda_{t-1}} - \hat{\theta}^{(\lambda_{t-1})})$$

*and*

$$\begin{aligned} \alpha[\hat{\theta}^{(\lambda_{t-1})}] &:= \operatorname*{arg\,min}_{\alpha\in\mathbb{R}_+} \left\| \frac{y}{\lambda_t} - \hat{\theta}^{(\lambda_{t-1})} - \alpha(\frac{y}{\lambda_{t-1}} - \hat{\theta}^{(\lambda_{t-1})}) \right\|_2 \\ &= \frac{\langle \frac{y}{\lambda_{t-1}} - \hat{\theta}^{(\lambda_{t-1})}, \frac{y}{\lambda_t} - \hat{\theta}^{(\lambda_{t-1})} \rangle}{\| \frac{y}{\lambda_{t-1}} - \hat{\theta}^{(\lambda_{t-1})} \|_2^2}. \end{aligned} \tag{16}$$

Note that the rule proposed by [21] (as pointed out in [6]) relies on the exact knowledge of a dual optimal solution for a previously solved Lasso problem. This is impossible to obtain in practice and even if it is possible to find accurate solutions, the search for high accuracy may hinder the benefits of the screening when it was not actually needed. Using inaccurate solutions may lead to discarding variables that should have been active and so the screened optimization algorithm will not converge to a solution of the original problem.

We illustrate this issue on Figure 4. Knowing an approximation $\beta$ to the optimal primal point, returned by the optimization algorithm at the previous regularization parameter $\lambda_{t-1}$, we need to choose an approximation $\theta$ to the optimal dual point to run EDPP.

- If we choose to approximate the dual optimal point by $\theta = \frac{1}{\lambda_{t-1}}(y - X\beta)$ (blue curve with diamonds), then the result is catastrophic. Indeed, at $\lambda_1$, $\beta = 0$ is a valid $\epsilon$-solution for $\epsilon = 10^{-1.5}$ and the screening rule tries to perform a division by 0 when computing $\alpha[\theta]$.
- If we choose to approximate the dual optimal point by $\frac{1}{\max(\lambda_{t-1}, \|X^\top(y-X\beta)\|_\infty)}(y - X\beta)$, we have a better behavior (purple curve with triangles) but we may still have an algorithm which does not converge to an $\epsilon$-solution. Here, for the $13^{\text{th}}$ Lasso problem a variable is erroneously removed and the problem can only be solved to accuracy $0.03515 > 10^{-1.5} \approx 0.03162$. This may look like a small issue but when the stopping criterion is based on the duality gap, this causes the algorithm to continue until the maximum number of iterations is reached.

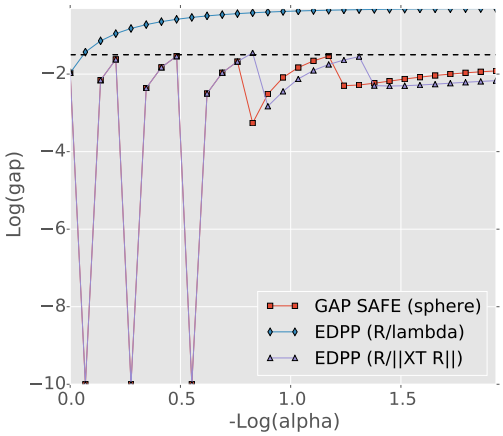


$$X = \begin{bmatrix} 1/\sqrt{2} & \sqrt{2}/\sqrt{3} \\ 0 & -1/\sqrt{6} \\ -1/\sqrt{2} & -1/\sqrt{6} \end{bmatrix}, \qquad y = \begin{bmatrix} 1/\sqrt{6} \\ 1/\sqrt{6} \\ -\sqrt{2}/\sqrt{3} \end{bmatrix}.$$

Figure 4: EDPP is not safe. We run GAP SAFE and two interpretations of EDPP (described in the main text) to solve the Lasso path on the dataset defined by $X$ and $y$ above with target accuracy $10^{-1.5}$. For each Lasso problem, we plot the final duality gap returned by the optimization solver.

# C Making EDDP screening rule safe

## C.1 The simpler screening rule

In the present paper, we give computable guarantees on the distance between the current dual feasible point and the solution of the problem. We show here how we can combine our result with Wang *et al.* 's in order to make their screening rule work even with approximate solutions to the previous Lasso problem.

For simplicity, we first consider the initial version of Wang *et al.* 's sphere test:

$$\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\big(\hat{\theta}^{(\lambda_{t-1})}, \big\|v^{\perp}(\lambda_{t-1}, \lambda_t)\big\|_2\big), \tag{17}$$

proved in [21, Theorem 7]. As we do not know $\hat{\theta}^{(\lambda_{t-1})}$, we cannot readily use this ball. However, we can modify it to make it a safe screening rules as follows:

**Proposition 4.** *Assume that $\lambda_{t-1} < \lambda_{\max}$, $\theta \in \Delta_X$ is a dual feasible point and $r_{\lambda_{t-1}} > 0$ is a radius satisfying $\hat{\theta}^{(\lambda_{t-1})} \in \mathcal{B}(\theta, r_{\lambda_{t-1}})$, then*

$$\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\Big(\theta, r_{\lambda_{t-1}}(1 + |1 - \alpha[\theta]|) + \Big\|\frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta)\Big\|_2\Big), \tag{18}$$

*where*

$$\alpha[\theta] := \operatorname*{arg\,min}_{\alpha \in \mathbb{R}_+} \Big\|\frac{y}{\lambda_t} - \theta - \alpha(\frac{y}{\lambda_{t-1}} - \theta)\Big\|_2 = \left(\frac{\langle \frac{y}{\lambda_{t-1}} - \theta, \frac{y}{\lambda_t} - \theta\rangle}{\|\frac{y}{\lambda_{t-1}} - \theta\|_2^2}\right)_+, \tag{19}$$

and for any $t \in \mathbb{R}$, $(t)_+ = \max(0, t)$.

*Proof.* Start first by noting that (17) implies

$$\hat{\theta}^{(\lambda_t)} \in \bigcup_{\theta' \in \mathcal{B}(\theta, r_{\lambda_{t-1}})} \mathcal{B}\Big(\theta', \min_{\alpha \in \mathbb{R}_+} \Big\|\frac{y}{\lambda_t} - \theta' - \alpha(\frac{y}{\lambda_{t-1}} - \theta')\Big\|_2\Big).$$

Let us denote

$$H = \max_{\theta' \in \mathcal{B}(\theta, r_{\lambda_{t-1}})} \min_{\alpha \in \mathbb{R}_+} \Big\|\frac{y}{\lambda_t} - \theta' - \alpha(\frac{y}{\lambda_{t-1}} - \theta')\Big\|_2,$$

then $\hat{\theta}^{(\lambda_t)} \in \mathcal{B}(\theta, r_{\lambda_{t-1}} + H)$. We now need to upper bound $H$. A simple choice is to take $\alpha$ to be $\alpha[\theta]$ defined in Eq. (19) The motivation for such a choice is because it is optimal when $r_{\lambda_{t-1}} = 0$. This provides the following bound on $H$:

$$\begin{aligned}
H &\leqslant \max_{\theta' \in \mathcal{B}(\theta, r_{\lambda_{t-1}})} \left\| \frac{y}{\lambda_t} - \theta' - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta') \right\|_2, \\
&= \left\| \frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta) + r_{\lambda_{t-1}}(\alpha[\theta] - 1) \frac{\frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta)}{\left\| \frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta) \right\|} \right\|_2, \\
&\leqslant r_{\lambda_{t-1}} |\alpha[\theta] - 1| + \left\| \frac{y}{\lambda_t} - \theta - \alpha[\theta].(\frac{y}{\lambda_{t-1}} - \theta) \right\|. \qquad (20)
\end{aligned}$$

Hence, after some simplifications:

$$\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\Big(\theta, r_{\lambda_{t-1}}(1 + |1 - \alpha[\theta]|) + \left\| \frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta) \right\|_2 \Big). \qquad \square$$

**Remark 8.** In the case that $\|y/\lambda_{t-1}\| \leqslant \|y/\lambda_{t-1} - \theta\| \leqslant 1$ then with the definition of $\alpha[\theta]$ and the Cauchy-Schwartz inequality one has that $1 + |\alpha[\theta] - 1| \leqslant \frac{\lambda_{t-1}}{\lambda_t}$. This means that the multiplicative ratio in front of $r_{\lambda_{t-1}}$ is $\lambda_{t-1}/\lambda_t$. In [11, Proposition 3], the bound obtained would only lead to the smaller ratio: $\sqrt{\lambda_{t-1}/\lambda_t}$.

**Remark 9.** From the proof of Theorem 7 in [21], it holds that for $\lambda < \lambda_{\max}$ then

$$\left\| \hat{\theta}^{(\lambda)} \right\| \leqslant \frac{\|y\|}{\lambda} \Leftrightarrow \hat{\theta}^{(\lambda)} \in B\left(0, \frac{\|y\|}{\lambda}\right). \qquad (21)$$

## C.2 The complete screening rule (EDDP+)

Let us now consider the EDDP+ screening rule [21] relying on the property (15): $\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\big(\hat{\theta}^{(\lambda_{t-1})} + \frac{1}{2} v^\perp(\lambda_{t-1}, \lambda_t), \frac{1}{2} \left\| v^\perp(\lambda_{t-1}, \lambda_t) \right\|_2\big)$. Using the same technique as for Proposition 4, we can strengthen our previous proposition with the following result.

**Proposition 5.** *Assume that $\lambda_{t-1} < \lambda_{\max}$, $\theta \in \Delta_X$ is a dual feasible point and $r_{\lambda_{t-1}} > 0$ is a radius satisfying $\hat{\theta}^{(\lambda_{t-1})} \in \mathcal{B}(\theta, r_{\lambda_{t-1}})$. Define $\alpha[\theta]$ as in* (19)*,*

$$\begin{aligned}
r_{\lambda_t} =& \frac{|1 - \alpha[\theta]| + 1 + \alpha[\theta]}{2} r_{\lambda_{t-1}} + \frac{1}{2} \left\| \frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta) \right\|_2 \\
&+ \frac{\left\| \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \right\|_2 r_{\lambda_{t-1}}}{2\|\frac{y}{\lambda_{t-1}} - \theta\|_2^2} \Big( 3 \left\| \frac{y}{\lambda_{t-1}} - \theta \right\|_2 + 2 r_{\lambda_{t-1}} \Big)
\end{aligned}$$

*and*

$$v^\perp(\theta, \lambda_{t-1}, \lambda_t) = \frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta). \qquad (22)$$

*Then $\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\Big(\theta + \frac{1}{2} v^\perp(\theta, \lambda_{t-1}, \lambda_t), r_{\lambda_t}\Big)$.*

*Proof.* As before, we do not know exactly $\hat{\theta}^{(\lambda_{t-1})}$ but we know that denoting

$$v^\perp(\theta', \lambda_{t-1}, \lambda_t) = \frac{y}{\lambda_t} - \theta' - \alpha[\theta'](\frac{y}{\lambda_{t-1}} - \theta') \qquad (23)$$

with

$$\alpha[\theta'] = \left( \frac{\langle \frac{y}{\lambda_{t-1}} - \theta', \frac{y}{\lambda_t} - \theta' \rangle}{\|\frac{y}{\lambda_{t-1}} - \theta'\|_2^2} \right)_+, \qquad (24)$$

we have

$$\hat{\theta}^{(\lambda_t)} \in \bigcup_{\theta' \in \mathcal{B}(\theta, r_{\lambda_{t-1}})} \mathcal{B}\Big(\theta' + \frac{1}{2} v^{\perp}(\theta', \lambda_{t-1}, \lambda_t), \frac{1}{2} \left\| v^{\perp}(\theta', \lambda_{t-1}, \lambda_t) \right\|_2 \Big).$$

Our goal is to find a ball centered at $\theta + \frac{1}{2} v^{\perp}(\theta, \lambda_{t-1}, \lambda_t)$ that contains all these balls, thus containing $\hat{\theta}^{(\lambda_t)}$. First, reminding (20)

$$\begin{aligned}
\left\| v^{\perp}(\theta', \lambda_{t-1}, \lambda_t) \right\|_2 &= \min_{\alpha \in \mathbb{R}_+} \left\| \frac{y}{\lambda_t} - \theta' - \alpha(\frac{y}{\lambda_{t-1}} - \theta') \right\|_2 \\
&\leqslant \max_{\theta' \in \mathcal{B}(\theta, r_{\lambda_{t-1}})} \min_{\alpha \in \mathbb{R}_+} \left\| \frac{y}{\lambda_t} - \theta' - \alpha(\frac{y}{\lambda_{t-1}} - \theta') \right\|_2 \\
&\leqslant r_{\lambda_{t-1}} \left| 1 - \alpha[\theta] \right| + \left\| \frac{y}{\lambda_t} - \theta - \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta) \right\|_2 .
\end{aligned}$$

We continue as

$$\begin{aligned}
&\theta' + \frac{1}{2} v^{\perp}(\theta', \lambda_{t-1}, \lambda_t) - \theta - \frac{1}{2} v^{\perp}(\theta, \lambda_{t-1}, \lambda_t) \\
&\qquad = (\theta' - \theta) + \frac{1}{2} \Big( \frac{y}{\lambda_t} - \theta' - \alpha[\theta'](\frac{y}{\lambda_{t-1}} - \theta') - \frac{y}{\lambda_t} + \theta + \alpha[\theta](\frac{y}{\lambda_{t-1}} - \theta) \Big) \\
&\qquad = \frac{1}{2} \Big( \theta' - \theta - (\alpha[\theta'] - \alpha[\theta])(\frac{y}{\lambda_{t-1}} - \theta') + \alpha[\theta](\theta' - \theta) \Big).
\end{aligned}$$

Taking the norm on both sides of the previous display,

$$\left\| \theta' + \frac{1}{2} v^{\perp}(\theta', \lambda_{t-1}, \lambda_t) - \theta - \frac{1}{2} v^{\perp}(\theta, \lambda_{t-1}, \lambda_t) \right\|_2 \leqslant \frac{1 + \alpha[\theta]}{2} \left\| \theta' - \theta \right\|_2 + \frac{|\alpha[\theta'] - \alpha[\theta]|}{2} \left\| \frac{y}{\lambda_{t-1}} - \theta' \right\|_2 .$$

Now, reminding that $x \mapsto (x)_+$ is a 1-Lipschitz function,

$$\begin{aligned}
\left| \alpha[\theta'] - \alpha[\theta] \right| &\leqslant \left| \frac{\langle \frac{y}{\lambda_{t-1}} - \theta', \frac{y}{\lambda_t} - \theta' \rangle}{\| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2} - \frac{\langle \frac{y}{\lambda_{t-1}} - \theta, \frac{y}{\lambda_t} - \theta \rangle}{\| \frac{y}{\lambda_{t-1}} - \theta \|_2^2} \right| \\
&= \left| \frac{\langle \frac{y}{\lambda_{t-1}} - \theta', \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \rangle}{\| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2} + 1 - \frac{\langle \frac{y}{\lambda_{t-1}} - \theta, \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \rangle}{\| \frac{y}{\lambda_{t-1}} - \theta \|_2^2} - 1 \right| \\
&= \left| \frac{\langle \| \frac{y}{\lambda_{t-1}} - \theta \|_2^2 (\frac{y}{\lambda_{t-1}} - \theta') - \| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2 (\frac{y}{\lambda_{t-1}} - \theta), \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \rangle}{\| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2 \| \frac{y}{\lambda_{t-1}} - \theta \|_2^2} \right| \\
&\leqslant \frac{\left\| \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \right\|_2}{\| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2 \| \frac{y}{\lambda_{t-1}} - \theta \|_2^2} \Big( \| \frac{y}{\lambda_{t-1}} - \theta' \|_2 \left| \| \frac{y}{\lambda_{t-1}} - \theta \|_2^2 - \| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2 \right| + \left\| \theta - \theta' \right\|_2 \| \frac{y}{\lambda_{t-1}} - \theta' \|_2^2 \Big) \\
&\leqslant \frac{\left\| \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \right\|_2}{\| \frac{y}{\lambda_{t-1}} - \theta' \|_2 \| \frac{y}{\lambda_{t-1}} - \theta \|_2^2} \Big( 2 \| \frac{y}{\lambda_{t-1}} - \frac{\theta' + \theta}{2} \|_2 \| \theta - \theta' \|_2 + \left\| \theta - \theta' \right\|_2 \| \frac{y}{\lambda_{t-1}} - \theta' \|_2 \Big) \\
&\leqslant \frac{\left\| \frac{y}{\lambda_t} - \frac{y}{\lambda_{t-1}} \right\|_2 \| \theta - \theta' \|_2}{\| \frac{y}{\lambda_{t-1}} - \theta' \|_2 \| \frac{y}{\lambda_{t-1}} - \theta \|_2^2} \Big( 2 \| \frac{y}{\lambda_{t-1}} - \theta \|_2 + \left\| \theta - \theta' \right\|_2 + \| \frac{y}{\lambda_{t-1}} - \theta \|_2 + \left\| \theta - \theta' \right\|_2 \Big). \qquad (25)
\end{aligned}$$

where the second inequality comes from the triangle inequality and the Cauchy-Schwartz Inequality, and the

third is obtained by factorizing the difference of squares. Plugging this in the former, we get:

$$\begin{aligned}
&\Big\|\theta' + \frac{1}{2}v^{\perp}(\theta',\lambda_{t-1},\lambda_t) - \theta - \frac{1}{2}v^{\perp}(\theta,\lambda_{t-1},\lambda_t)\Big\|_2 \\
&\qquad\leqslant \frac{1+\alpha[\theta]}{2}\left\|\theta'-\theta\right\|_2 + \frac{1}{2}\frac{\left\|\frac{y}{\lambda_t}-\frac{y}{\lambda_{t-1}}\right\|_2\|\theta-\theta'\|_2}{\|\frac{y}{\lambda_{t-1}}-\theta\|_2^2}\Big(3\left\|\frac{y}{\lambda_{t-1}}-\theta\right\|_2 + 2\left\|\theta-\theta'\right\|_2\Big).
\end{aligned}$$

One could check that there exists $\theta' \in \mathcal{B}(\theta, r_{\lambda_{t-1}})$ satisfying $\hat{\theta}^{(\lambda_t)} \in \mathcal{B}\big(\theta' + \frac{1}{2}v^{\perp}(\theta',\lambda_{t-1},\lambda_t), \frac{1}{2}\left\|v^{\perp}(\theta',\lambda_{t-1},\lambda_t)\right\|_2\big)$ and so combining the last inequality with (25)

$$\begin{aligned}
\Big\|\hat{\theta}^{(\lambda_t)} - \theta - \frac{1}{2}v^{\perp}(\theta,\lambda_{t-1},\lambda_t)\Big\|_2 &\leqslant \Big\|\hat{\theta}^{(\lambda_t)} - \theta' - \frac{1}{2}v^{\perp}(\theta',\lambda_{t-1},\lambda_t)\Big\|_2 \\
&\quad + \Big\|\theta' + \frac{1}{2}v^{\perp}(\theta',\lambda_{t-1},\lambda_t) - \theta - \frac{1}{2}v^{\perp}(\theta,\lambda_{t-1},\lambda_t)\Big\|_2 \\
&\leqslant \frac{|1-\alpha[\theta]|+1+\alpha[\theta]}{2}r_{\lambda_{t-1}} + \frac{1}{2}\left\|\frac{y}{\lambda_t}-\theta-\alpha[\theta](\frac{y}{\lambda_{t-1}}-\theta)\right\|_2 \\
&\quad + \frac{\left\|\frac{y}{\lambda_t}-\frac{y}{\lambda_{t-1}}\right\|_2 r_{\lambda_{t-1}}}{2\|\frac{y}{\lambda_{t-1}}-\theta\|_2^2}\Big(3\left\|\frac{y}{\lambda_{t-1}}-\theta\right\|_2 + 2r_{\lambda_{t-1}}\Big) \qquad \square
\end{aligned}$$