
Screening Rules for Lasso with Non-Convex Sparse Regularizers

Alain Rakotomamonjy^{1,2} Gilles Gasso¹ Joseph Salmon³

Abstract

Leveraging on the convexity of the *Lasso* problem, screening rules help in accelerating solvers by discarding irrelevant variables, during the optimization process. However, because they provide better theoretical guarantees in identifying relevant variables, several non-convex regularizers for the *Lasso* have been proposed in the literature. This work is the first that introduces a screening rule strategy into a non-convex *Lasso* solver. The approach we propose is based on an iterative majorization-minimization (MM) strategy that includes a screening rule in the inner solver and a condition for propagating screened variables between iterations of MM. In addition to improve efficiency of solvers, we also provide guarantees that the inner solver is able to identify the zeros components of its critical point in finite time. Our experimental analysis illustrates the significant computational gain brought by the new screening rule compared to classical coordinate-descent or proximal gradient descent methods.

1. Introduction

Sparsity-inducing penalties are classical tools in statistical machine learning, especially in settings where the data available for learning is scarce and of high-dimension. In addition, when the solution of the learning problem is known to be sparse, using those penalties yield to models that can leverage this prior knowledge. The *Lasso* (Tibshirani, 1996) and the *Basis pursuit* (Chen et al., 2001; Chen & Donoho, 1994) where the first approaches that have employed ℓ_1 -norm penalty for inducing sparsity.

¹LITIS, Normandy University, University of Rouen Normandie, INSA Rouen Normandie, France ²Criteo AI Lab, Paris, France ³IMAG, Univ Montpellier, CNRS, Montpellier, France. Correspondence to: Alain Rakotomamonjy <alain.rakotomamonjy@univ-rouen.fr>.

While the *Lasso* has had a great impact on machine learning and signal processing communities given some success stories (Shevade & Keerthi, 2003; Donoho, 2006; Lustig et al., 2008; Ye & Liu, 2012) it also comes with some theoretical drawbacks (*e.g.*, biased estimates of large coefficient of the model). Hence, several authors have proposed non-convex penalties that approximate better the ℓ_0 - (pseudo)norm, the later being the original measure of sparsity though it leads to NP-hard learning problem. The most commonly used non-convex penalties are the *Smoothly Clipped Absolute Deviation* (SCAD) (Fan & Li, 2001), the *Log Sum penalty* (LSP) (Candès et al., 2008), the *capped- ℓ_1 penalty* (Zhang, 2010b), the *Minimax Concave Penalty* (MCP) (Zhang, 2010a). We refer the interested reader to (Soubies et al., 2017) for a discussion on the pros and cons of such non-convex formulations.

From an optimization point of view, the *Lasso* can benefit from a large palette of algorithms ranging from block-coordinate descent (Friedman et al., 2007; Fu, 1998) to (accelerated) proximal gradient approaches (Beck & Teboulle, 2009). In addition, by leveraging convex optimization theory, some of these algorithms can be further accelerated by combining them with sequential or dynamic safe screening rules (El Ghaoui et al., 2012; Bonnefoy et al., 2015; Fercoq et al., 2015), which allow to safely set some useless variables to 0 before terminating (or even sometimes before starting) the algorithm.

While learning sparse models with non-convex penalties are seemingly more challenging to address, block-coordinate descent algorithms (Breheny & Huang, 2011; Mazumder et al., 2011) or iterative shrinkage thresholding (Gong et al., 2013) algorithms can be extended to those learning problems. Another popular way of handling non-convexity is to consider the majorization-minimization (MM) (Hunter & Lange, 2004) principle which consists in iteratively minimizing a majorization of a (non-convex) objective function. When applied to non-convex sparsity enforcing penalty, the MM scheme leads to solving a sequence of weighted *Lasso* problems (Candès et al., 2008; Gasso et al., 2009; Mairal, 2013).

In this paper, we propose a screening strategy that can be applied when dealing with non-convex penalties. As far as we know, this work is the first attempt in that direction.

For that, we consider a MM framework and in this context, our contributions are : (i) the definition of a MM algorithm that produces a sequence of iterates known to converge towards a critical point of our non-convex learning problem, (ii) the proposition of a duality gap based screening rule for weighted *Lasso*, which is the core algorithm of our MM framework, (iii) the introduction of conditions allowing to propagate screened variables from one MM iteration to the next, (iv) we also empirically show that our screening strategy indeed improves the running time of MM algorithms with respect to block-coordinate descent or proximal gradient descent methods.

2. Global Non-Convex Framework

We introduce the problem we are interesting in, its first order optimality conditions and propose an MM approach for its resolution.

2.1. The optimization problem

We consider solving the problem of least-squares regression with a generic penalty of the form

$$\min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \sum_{j=1}^d r_\lambda(|w_j|) , \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^n$ is a target vector, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_d] \in \mathbb{R}^{n \times d}$ is the design matrix with column-wise features \mathbf{x}_j , \mathbf{w} is the coefficient vector of the model and the map $r_\lambda : \mathbb{R}_+ \mapsto \mathbb{R}_+$ is concave and differentiable on $[0, +\infty)$ with a regularization parameter $\lambda > 0$. In addition, we assume that $r_\lambda(|w|)$ is lower semi-continuous function. Note that most penalty functions such as SCAD, MCP or log sum (see their definitions in Table 1) admit such a property.

We consider tools such as Fréchet subdifferentials and limiting-subdifferentials (Kruger, 2003; Rockafellar & Wets, 2009; Mordukhovich et al., 2006) well suited for non-smooth and non-convex optimization, so that a vector \mathbf{w}^* belongs to the set of minimizers (not necessarily global) of Problem (1) if the following Fermat condition holds (Clarke, 1989; Kruger, 2003):

$$\mathbf{X}^\top (\mathbf{y} - \mathbf{X}\mathbf{w}^*) \in \sum_{j=1}^d \partial r_\lambda(|w_j^*|) , \quad (2)$$

with $\partial r_\lambda(\cdot)$ being the Fréchet subdifferential of r_λ , assuming it exists at \mathbf{w}^* . In particular this is the case for the MCP, log sum and SCAD penalties presented in Table 1. For an illustration, we present the optimality conditions for MCP and log sum.

Example 1. For the MCP penalty (see Table 1 for the definition and subdifferential), it is easy to show that the

$\partial r_\lambda(0) = [-\lambda, \lambda]$. Hence, the Fermat condition becomes

$$\begin{aligned} -\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\mathbf{w}^*) &= 0, & \text{if } |w_j^*| > \lambda\theta \\ |\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\mathbf{w}^*)| &\leq \lambda, & \text{if } w_j^* = 0 \\ -\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\mathbf{w}^*) + \lambda \operatorname{sign}(w_j^*) - \frac{w_j^*}{\theta} &= 0, & \text{otherwise} \end{aligned} \quad (3)$$

Example 2. For the log sum penalty, one can explicitly compute $\partial r_\lambda(0) = [-\frac{\lambda}{\theta}, \frac{\lambda}{\theta}]$ and leverage on smoothness of $r_\lambda(|w|)$ when $|w| > 0$ for computing $\partial r_\lambda(|w|)$. Then the above necessary condition can be translated as

$$\begin{aligned} -\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\mathbf{w}^*) + \lambda \frac{\operatorname{sign}(w_j^*)}{\theta + |w_j^*|} &= 0, & \text{if } w_j^* \neq 0 , \\ |\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\mathbf{w}^*)| &\leq \frac{\lambda}{\theta}, & \text{if } w_j^* = 0 \end{aligned} \quad (4)$$

As we can see, first order optimality conditions lead to simple equations and inclusion.

Remark 1. There exists a critical parameter λ_{\max} such that $\mathbf{0}$ is a critical point for the primal problem for all $\lambda \geq \lambda_{\max}$. This parameter depends on the subdifferential of r_λ at 0. For the MCP penalty we have $\lambda_{\max} \triangleq \max_j |\mathbf{x}_j^\top \mathbf{y}|$, and for the log sum penalty $\lambda_{\max} \triangleq \theta \max_j |\mathbf{x}_j^\top \mathbf{y}|$. From now on, we assume that $\lambda \leq \lambda_{\max}$ to avoid such irrelevant local solutions.

2.2. Majorization-Minimization approach

There exists several majorization-minimization algorithms for solving non-smooth and non-convex problems involving sparsity-inducing penalties (Gasso et al., 2009; Gong et al., 2013; Mairal, 2013).

In this work, we focus on MM algorithms provably convergent to a critical point of Problem 1 such as those described by Kang et al. (2015). Their main mechanism is to iteratively build a majorizing surrogate objective function that is easier to solve than the original learning problem. In the case of non-convex penalties that are either fully concave or that can be written as the sum of a convex and concave functions, the idea is to linearize the concave part, and the next iterate is obtained by optimizing the resulting surrogate function. Since $r_\lambda(|\cdot|)$ is a concave and differentiable function on $[0, +\infty)$, at any iterate k , we have :

$$r_\lambda(|w_j|) \leq r_\lambda(|w_j^k|) + r'_\lambda(|w_j^k|)(|w_j| - |w_j^k|) .$$

To take advantage of MM convergence properties (Kang et al., 2015), we also majorize the objective function and our algorithm boils down to the following iterative process

$$\begin{aligned} \mathbf{w}^{k+1} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} & \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \frac{1}{2\alpha} \|\mathbf{w} - \mathbf{w}^k\|_2^2 \\ & + \sum_{j=1}^d r'_\lambda(|w_j^k|) |w_j| , \end{aligned} \quad (5)$$

Table 1. Common non-convex penalties with their sub-differentials. Here $\lambda > 0$, $\theta > 0$ ($\theta > 1$ for MCP, $\theta > 2$ for SCAD).

| Penalty | $r_\lambda(w)$ | $\partial r_\lambda(w)$ |
|---------|---|---|
| Log sum | $\lambda \log(1 + w /\theta)$ | $\begin{cases} \left[\frac{-\lambda}{\theta}, \frac{\lambda}{\theta} \right] & \text{if } w = 0 \\ \left\{ \lambda \frac{\text{sign}(w)}{\theta + w } \right\} & \text{if } w > 0 \end{cases}$ |
| MCP | $\begin{cases} \lambda w - \frac{w^2}{2\theta} & \text{if } w \leq \lambda\theta \\ \theta\lambda^2/2 & \text{if } w > \lambda\theta \end{cases}$ | $\begin{cases} [-\lambda, \lambda] & \text{if } w = 0 \\ \left\{ \lambda \text{sign}(w) - \frac{w}{\theta} \right\} & \text{if } 0 < w \leq \lambda\theta \\ \{0\} & \text{if } w > \theta\lambda \end{cases}$ |
| SCAD | $\begin{cases} \lambda w & \text{if } w \leq \lambda \\ \frac{1}{2(\theta-1)}(-w^2 + 2\theta\lambda w - \lambda^2) & \text{if } \lambda < w \leq \lambda\theta \\ \frac{\lambda^2(1+\theta)}{2} & \text{if } w > \theta\lambda \end{cases}$ | $\begin{cases} [-\lambda, \lambda] & \text{if } w = 0 \\ \left\{ \lambda \text{sign}(w) \right\} & \text{if } 0 < w \leq \lambda \\ \left\{ \frac{1}{\theta-1}(-w + \theta\lambda \text{sign}(w)) \right\} & \text{if } 0 < w \leq \lambda\theta \\ \{0\} & \text{if } w > \theta\lambda \end{cases}$ |

where $\alpha > 0$ is some user-defined parameter controlling the proximal regularization strength. Note that as $\alpha \rightarrow \infty$, the above problem recovers the reweighted ℓ_1 iterations investigated by Candès et al. (2008); Gasso et al. (2009). Moreover, when using $\mathbf{w}^k = \mathbf{0}$ (e.g., when evaluating the first λ in a path-wise fashion, for $k = 0$) this recovers the Elastic-net penalty (Zou & Hastie, 2005).

3. Proximal Weighted Lasso: Coordinate Descent and Screening

As we have stated, the screening rule we have developed for regression with non-convex sparsity enforcing penalties is based on iterative minimization of Proximal Weighted Lasso problem. Hereafter, we briefly show that such sub-problems can be solved by iteratively optimizing coordinate wise. Then, we derive a duality gap based screening rule to screen coordinate-wise. In what follows, we assume that $\frac{0}{0} = 0^1$.

3.1. Primal and dual problems

Solving the following primal problem would prove useful in the application of our MM framework, as it encompasses Problem 5 as a special case (for some $\lambda_j = r'_\lambda(|w_j^k|)$):

$$\min_{\mathbf{w} \in \mathbb{R}^d} P_\Lambda(\mathbf{w}) \triangleq \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \frac{1}{2\alpha} \|\mathbf{w} - \mathbf{w}'\|_2^2 + \sum_{j=1}^d \lambda_j |w_j|. \quad (6)$$

with \mathbf{w}' being some pre-defined vector of \mathbb{R}^d , and $\Lambda = (\lambda_1 \dots \lambda_d)^\top$ with $\lambda_1 \geq 0, \dots, \lambda_d \geq 0$ some regularization parameters. In the sequel we denote Problem 6 as the Proximal Weighted Lasso (PWL) problem.

Problem (6) can be solved through proximal gradient descent algorithm (Beck & Teboulle, 2009) but in order to benefit from screening rules, coordinate descent algorithms are to be privileged. Friedman et al. (2010) have proposed a coordinate-wise update for the Elastic-net penalty and sim-

ilarly, it can be shown that for the above, the following update holds for any coordinate w_j and $\lambda_j \geq 0$:

$$w_j \leftarrow \frac{1}{\|\mathbf{x}_j\|_2^2 + \frac{1}{\alpha}} \text{sign}(t_j) \max(0, |t_j| - \lambda_j), \quad (7)$$

with $t_j = \mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\mathbf{w} + \mathbf{x}_j w_j) + \frac{1}{\alpha} w'_j$.

Typically, coordinate descent algorithms visit all the variables in a cyclic way (another popular choice is sampling uniformly at random among the coordinates) unless some screening rules prevent them from unnecessary updates.

Recent efficient screening methods rely on producing primal-dual approximate solutions and on defining tests based on these approximate solutions (Fercq et al., 2015; Shibagaki et al., 2016; Ndiaye et al., 2017). While our approach follows the same road, it does not result from a straightforward extension of the works of Fercq et al. (2015) and Shibagaki et al. (2016) due to the proximal regularizer.

The primal objective of Proximal Weighted Lasso (given in Equation (6)) is convex (and lower bounded) and admits at least one global solution. To derive our screening tests, we need to investigate the dual formulation associated to this problem, that reads:

$$\begin{aligned} \max_{\substack{\mathbf{s} \in \mathbb{R}^n \\ \mathbf{v} \in \mathbb{R}^d}} D(\mathbf{s}, \mathbf{v}) &\triangleq -\frac{1}{2} \|\mathbf{s}\|_2^2 - \frac{\alpha}{2} \|\mathbf{v}\|_2^2 + \mathbf{s}^\top \mathbf{y} - \mathbf{v}^\top \mathbf{w}' \quad (8) \\ \text{s.t.} \quad &|\mathbf{X}^\top \mathbf{s} - \mathbf{v}| \preceq \Lambda, \quad (9) \end{aligned}$$

with the inequality operator \preceq being applied in a coordinate-wise manner. As a side result of the dual derivation (see Appendix for the details), key conditions for screening any primal variable w_j are obtained as:

$$|\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| - \lambda_j < 0 \implies w_j^* = 0,$$

with \mathbf{s}^* and \mathbf{v}^* being solutions of the dual formulation from Equation (8).

¹This will be of interest for cases where $\lambda_j = 0$ in Problem 6.

3.2. Screening test

Screening test on the Proximal Weighted *Lasso* problem can thus be derived if we are able to provide an upper bound on $|\mathbf{x}_j^\top \mathbf{s}^* - v_j^*|$ that is guaranteed to be strictly smaller than λ_j . Suppose that we have an intermediate triplet of primal-dual solution $(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})$ with $\hat{\mathbf{s}}$ and $\hat{\mathbf{v}}$ being dual feasible², then we can derive the following bound

$$\begin{aligned} |\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| &= |\mathbf{x}_j^\top \hat{\mathbf{s}} - \hat{v}_j + \mathbf{x}_j^\top (\mathbf{s}^* - \hat{\mathbf{s}}) - (v_j^* - \hat{v}_j)| \\ &\leq |\mathbf{x}_j^\top \hat{\mathbf{s}} - \hat{v}_j| + \|\mathbf{x}_j\| \|\mathbf{s}^* - \hat{\mathbf{s}}\| + |v_j^* - \hat{v}_j|. \end{aligned}$$

Now we need an upper bound on the distance between the approximated and the optimal dual solution in order to make screening condition exploitable. By exploiting the property that the objective function $D(\mathbf{s}, \mathbf{v})$ of the dual problem given in Equation (8) is quadratic and strongly concave³, the following inequality holds

$$\begin{aligned} D(\hat{\mathbf{s}}, \hat{\mathbf{v}}) &\leq D(\mathbf{s}^*, \mathbf{v}^*) - \nabla_{\mathbf{s}} D(\mathbf{s}^*, \mathbf{v}^*)^\top (\hat{\mathbf{s}} - \mathbf{s}^*) \\ &\quad - \nabla_{\mathbf{v}} D(\mathbf{s}^*, \mathbf{v}^*)^\top (\hat{\mathbf{v}} - \mathbf{v}^*) \\ &\quad - \frac{1}{2} \|\hat{\mathbf{s}} - \mathbf{s}^*\|_2^2 - \frac{\alpha}{2} \|\hat{\mathbf{v}} - \mathbf{v}^*\|_2^2, \end{aligned}$$

with $\nabla_{\mathbf{s}} D = \left[\frac{\partial D}{\partial s_1}, \dots, \frac{\partial D}{\partial s_n} \right]^\top$, $\nabla_{\mathbf{v}} D = \left[\frac{\partial D}{\partial v_1}, \dots, \frac{\partial D}{\partial v_d} \right]^\top$. As the dual problem is a constrained optimization problem, the first-order optimality condition for $(\mathbf{s}^*, \mathbf{v}^*)$ reads $\nabla_{\mathbf{s}} D(\mathbf{s}^*, \mathbf{v}^*)^\top (\mathbf{s} - \mathbf{s}^*) + \nabla_{\mathbf{v}} D(\mathbf{s}^*, \mathbf{v}^*)^\top (\mathbf{v} - \mathbf{v}^*) \geq 0$, $\forall \mathbf{s} \in \mathbb{R}^n, \mathbf{v} \in \mathbb{R}^d$; thus we have

$$2(D(\mathbf{s}^*, \mathbf{v}^*) - D(\hat{\mathbf{s}}, \hat{\mathbf{v}})) \geq \|\hat{\mathbf{s}} - \mathbf{s}^*\|_2^2 + \alpha \|\hat{\mathbf{v}} - \mathbf{v}^*\|_2^2.$$

By strong duality, we have $P_\Lambda(\hat{\mathbf{w}}) \geq D(\mathbf{s}^*, \mathbf{v}^*)$, hence

$$2(P_\Lambda(\hat{\mathbf{w}}) - D(\hat{\mathbf{s}}, \hat{\mathbf{v}})) \geq \|\hat{\mathbf{s}} - \mathbf{s}^*\|_2^2 + \alpha \|\hat{\mathbf{v}} - \mathbf{v}^*\|_2^2.$$

We can now use the duality gap for bounding $\|\hat{\mathbf{s}} - \mathbf{s}^*\|$ and $|\hat{v}_j - v_j^*|$. Hence, given a primal-dual intermediate solution $(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})$, with duality gap $G_\Lambda(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}}) \triangleq P_\Lambda(\hat{\mathbf{w}}) - D(\hat{\mathbf{s}}, \hat{\mathbf{v}})$, the screening test for a variable j is

$$\underbrace{|\mathbf{x}_j^\top \hat{\mathbf{s}} - \hat{v}_j| + \sqrt{2G_\Lambda(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})}}_{T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})} \left(\|\mathbf{x}_j\| + \frac{1}{\alpha} \right) < \lambda_j, \quad (10)$$

and we can safely state that the j -th coordinate of \mathbf{w}^* is zero when this happens.

Finding approximate primal-dual solutions: In our case of interest, the Proximal Weighted *Lasso* problem is

²meaning $|\mathbf{X}^\top \hat{\mathbf{s}} - \hat{\mathbf{v}}| \preceq \Lambda$.

³see (Nesterov, 2004) for a precise definition of strong convexity/concavity.

solved in its primal form using a coordinate descent algorithm that optimizes one coordinate at a time. Hence, an approximate primal solution $\hat{\mathbf{w}}$ is easy to obtain by considering the current solution at a given iteration of the algorithm. From this primal solution, we show how to obtain a dual feasible solution $\hat{\mathbf{s}}$ that can be considered for the screening test.

One can check the following primal/dual link (for instance by deriving the first order conditions of the maximization in Equation 16 or equivalent formulation 21, see Appendix)

$$\mathbf{y} - \mathbf{X}\mathbf{w}^* = \mathbf{s}^* \quad \text{and} \quad \mathbf{w}^* - \mathbf{w}'^* = \alpha \mathbf{v}^*$$

with the constraints that $|\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| \leq \lambda_j, \forall j \in [d]$.

Hence, a good approximation of the dual solution can be obtained by scaling the residual vector $\mathbf{y} - \mathbf{X}\hat{\mathbf{w}}$ such that it becomes dual feasible. Indeed, the condition $|\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| \leq \lambda_j, \forall j \in [d]$ is guaranteed only at optimality. To avoid issues with dividing by vanishing λ_j 's, let us consider the set $\mathcal{S} = \{j \in [d] : \lambda_j > 0\}$ of associated indexes and assume that this set is non-empty. Then, one can define

$$j^\dagger = \arg \max_{j \in \mathcal{S}} \underbrace{\frac{1}{\lambda_j} |\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\hat{\mathbf{w}}) - \frac{1}{\alpha} (\hat{v}_j - v_j')|}_{\rho(j)}. \quad (11)$$

For all $j \in \mathcal{S}$, if $\rho(j^\dagger) \leq 1$ then $\hat{\mathbf{s}} \triangleq \mathbf{y} - \mathbf{X}\hat{\mathbf{w}}$ and $\mathbf{v} = \frac{1}{\alpha}(\hat{\mathbf{w}} - \mathbf{w}')$ are dual feasible, no scaling is needed. If $\rho(j^\dagger) > 1$, we define the approximate dual solution as $\hat{\mathbf{s}} = \frac{\mathbf{y} - \mathbf{X}\hat{\mathbf{w}}}{\rho(j^\dagger)}$ and $\hat{\mathbf{v}} = \frac{\hat{\mathbf{w}} - \mathbf{w}'}{\alpha \rho(j^\dagger)}$ which are dual feasible. Hence, in practice, we compute our screening test using the triplet $\{\hat{\mathbf{w}}, \frac{\mathbf{y} - \mathbf{X}\hat{\mathbf{w}}}{\max(1, \rho(j^\dagger))}, \frac{\hat{\mathbf{w}} - \mathbf{w}'}{\alpha \max(1, \rho(j^\dagger))}\}$.

Remark 2. As the above dual approximation is valid only for components of $\lambda_j > 0$, we have added a special treatment for components $j \in [d] \setminus \mathcal{S}$ setting $\hat{v}_j = \mathbf{x}_j^\top \hat{\mathbf{s}}$ for such indexes, guaranteeing dual feasibility. Note that the special case $\lambda_j = 0$ is not an innocuous case. For some non-convex penalties like MCP or SCAD, their gradient r'_λ is equal to 0 for large values and this is one of their key statistical property. Hence, the situation in which $\lambda_j = 0$ is very likely to occur in practice for these penalties.

Comparing with other screening tests: Our screening test exploits duality gap and strong concavity of the dual function, and the resulting test is similar to the one derived by Fercoq et al. (2015) or Raj et al. (2016). Indeed, it can be shown that Equation (10) boils down to be a test based on a sphere centered on $\hat{\mathbf{s}}$ with radius $\sqrt{2G_\Lambda(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})}$. The main difference relies on the threshold of the test λ_j , which differs on a coordinate basis instead of being uniform.

Similarly to the GAP test of Fercoq et al. (2015), our screening rule come with theoretical properties. It can be

Algorithm 1 Proximal Weighted *Lasso* (PWL)

Input: \mathbf{X} , \mathbf{y} , \mathbf{w}^0 , Λ , α , LstScreen

Output: \mathbf{w} , $\mathbf{r} = \mathbf{y} - \mathbf{X}\mathbf{w}$

```

1:  $k \leftarrow 0$ 
2: repeat
3:   for variable  $j \notin \text{LstScreen}$  do
4:      $w_j^{k+1} \leftarrow$  update coordinate  $w_j^k$  using Equation (7)
5:   end for
6:   compute duality gap,  $\mathbf{r} = \mathbf{y} - \mathbf{X}\mathbf{w}^{k+1}$ , approximate
   dual variables
7:   for variable  $j \notin \text{LstScreen}$  do
8:     update screening condition using Equation (10)
9:   end for
10:   $k \leftarrow k + 1$ 
11: until convergence
    
```

shown that if we have a converging sequence $\{\hat{\mathbf{w}}^k\}$ of primal coefficients $\lim_{k \rightarrow \infty} \hat{\mathbf{w}}^k = \mathbf{w}^*$, then $\hat{\mathbf{s}}^k$ and $\hat{\mathbf{v}}^k$ defined as above converge to \mathbf{s}^* and \mathbf{v}^* . Moreover, we can state a property showing the ability of our screening rule to remove irrelevant variables after a finite number of iterations.

Property 1. Define the *Equicorrelation set*⁴ of the Proximal Weighted Lasso as $\mathcal{E}^* \triangleq \{j \in [d] : |\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| = \lambda_j\}$ and $\mathcal{E}_k \triangleq \{j \in [d] : |\mathbf{x}_j^\top \hat{\mathbf{s}}^k - \hat{v}_j^k| \geq \lambda_j\}$ obtained at iteration k of an algorithm solving the Proximal Weighted Lasso. Then, there exists an iteration $k_0 \in \mathbb{N}$ s.t. $\forall k \geq k_0$, $\mathcal{E}_k \subset \mathcal{E}^*$.

Proof. Because $\hat{\mathbf{w}}^k$, $\hat{\mathbf{s}}^k$ and $\hat{\mathbf{v}}^k$ are convergent, owing to the strong duality, the duality gap also converges towards zero. Now, for any given ϵ , define k_0 such that $\forall k \geq k_0$, we have

$$\|\hat{\mathbf{s}}^k - \mathbf{s}^*\|_2 \leq \epsilon, \|\hat{\mathbf{v}}^k - \mathbf{v}^*\|_\infty \leq \epsilon \text{ and } \sqrt{2G_\Lambda} \leq \epsilon .$$

For $j \notin \mathcal{E}^*$, we have

$$\begin{aligned} |\mathbf{x}_j^\top (\hat{\mathbf{s}}^k - \mathbf{s}^*) - (\hat{v}_j^k - v_j^*)| &\leq |\mathbf{x}_j^\top (\hat{\mathbf{s}}^k - \mathbf{s}^*)| + |(\hat{v}_j^k - v_j^*)| \\ &\leq (\max_{j \notin \mathcal{E}^*} \|\mathbf{x}_j\| + 1)\epsilon . \end{aligned}$$

From the Triangle inequality we have:

$$\begin{aligned} |\mathbf{x}_j^\top \hat{\mathbf{s}}^k - \hat{v}_j^k| &\leq |\mathbf{x}_j^\top (\hat{\mathbf{s}}^k - \mathbf{s}^*) - (\hat{v}_j^k - v_j^*)| + |\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| \\ &\leq (\max_{j \notin \mathcal{E}^*} \|\mathbf{x}_j\| + 1)\epsilon + |\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| \end{aligned}$$

If we add $\sqrt{2G_\Lambda} \left(\|\mathbf{x}_j\| + \frac{1}{\alpha} \right)$ on both sides, we get

$$T_j^{(\lambda_j)} \leq \left(2 \max_{j \notin \mathcal{E}^*} \|\mathbf{x}_j\| + 1 + \frac{1}{\alpha} \right) \epsilon + |\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| .$$

⁴following a terminology introduced by Tibshirani (2013)

Now define the constant $C \triangleq \min_{j \notin \mathcal{E}^*} [\lambda_j - |\mathbf{x}_j^\top \mathbf{s}^* - v_j^*|]$ and because $j \notin \mathcal{E}^*$, $C > 0$. Hence, if we choose

$$\epsilon < \frac{C}{2 \max_{j \notin \mathcal{E}^*} \|\mathbf{x}_j\| + 1 + \frac{1}{\alpha}} ,$$

we have $T_j^{(\lambda_j)} < \lambda_j$ which means that the variable j has been screened out, hence $j \notin \mathcal{E}_k$. To conclude, $j \notin \mathcal{E}^*$ implies that $j \notin \mathcal{E}_k$, which also translates in $\mathcal{E}_k \subset \mathcal{E}^*$. \square

This property thus tells that all the zero variables of \mathbf{w}^* outside \mathcal{E}^* are correctly detected and screened out by our screening rule in a finite number of iterations of the algorithm. In real data scenarios, though \mathcal{E}^* can differ from the true support, it is generally equal to the true support⁵.

4. Screening Rules for Non-Convex Regularizers

Now that we have described the inner solver and its screening rule, we are going to analyze how this rule can be improved into a majorization-minimization (MM) approach.

4.1. Majorization-minimization approaches and screening

At first, let us discuss some properties of our MM algorithm. According to the first order condition of the Proximal Weighted *Lasso* related to the MM problem at iteration k , the following inequality holds for any $j \in [d]$

$$|\mathbf{x}_j^\top \mathbf{s}^{k,*} - v_j^{k,*}| \leq \lambda_j^k ,$$

where the superscript denotes the optimal solution at iteration k . Owing to the convergence properties of the MM algorithm Kang et al. (2015), we know that the sequence $\{\mathbf{w}^k\}$ converges towards a vector satisfying Equation (2). Owing to the continuity of $r_\lambda(\cdot)$, we deduce that the sequence $\{\lambda_j^k\}$ also converges towards a λ_j^* . Thus, by taking the limits of the above inequality, the following condition holds for $w_j^* = 0$:

$$|\mathbf{x}_j^\top \mathbf{s}^* - v_j^*| \leq \lambda_j^* ,$$

with $\lambda_j^* = r'_\lambda(|w_j^*|)$. This inequality basically tells us about the relation between vanishing primal component and the optimal dual variable at each iteration k . While this suggests that screening within each inner problem by defining λ_j in Equation (6) as $\lambda_j^k = r'_\lambda(|w_j^k|)$ should improve efficiency of the global MM solver, it does not tell whether screened variables at iteration k are going to be screened at the next iteration as the λ_j^k 's are also expected to vary between iterations.

⁵Tibshirani (2013) discusses this issue for the *Lasso* case

Remark 3. In general, the behavior of a λ_j across MM iterations strongly depends on an initial \mathbf{w}^0 and the optimal \mathbf{w}^* . For instance, if one variable w_j^0 is initialized at 0 and w_j^* is large, λ_j will tend to be decreasing across iterations. Conversely, λ_j will tend to increase.

4.2. Propagating screening conditions

In what follows, we derive conditions on allowing to propagate screened coefficients from one iteration to another in the MM framework.

Property 2. Consider a Proximal Weighted Lasso problem with weights $\{\lambda_j\}$ and its primal-dual approximate solutions $\hat{\mathbf{w}}$, $\hat{\mathbf{s}}$ and $\hat{\mathbf{v}}$ allowing to evaluate a screening test in Equation (10). Suppose that we have a new set of weight $\Lambda^\nu = \{\lambda_j^\nu\}_{j=1,\dots,d}$ defining a new Proximal Weighted Lasso problem. Given a primal-dual approximate solution defined by the triplet $(\hat{\mathbf{w}}^\nu, \hat{\mathbf{s}}^\nu, \hat{\mathbf{v}}^\nu)$ for the latter problem, a screening test for variable j reads

$$T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}}) + \|\mathbf{x}_j\| (a + \sqrt{2b}) + c + \frac{1}{\alpha} \sqrt{2b} < \lambda_j^\nu, \quad (12)$$

where $T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})$ is the screening test for j at λ_j , a , b and c are constants such that $\|\hat{\mathbf{s}}^\nu - \hat{\mathbf{s}}\|_2 \leq a$, $|G_\Lambda(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}}) - G_{\Lambda^\nu}(\hat{\mathbf{w}}^\nu, \hat{\mathbf{s}}^\nu, \hat{\mathbf{v}}^\nu)| \leq b$ and $|\hat{v}_j^\nu - \hat{v}_j| \leq c$.

Proof. The screening test for the Proximal Weighted Lasso problem for parameter Λ_j^ν can be written as

$$|\mathbf{x}_j^\top \hat{\mathbf{s}}^\nu - \hat{v}_j^\nu| + \sqrt{2G_{\Lambda^\nu}(\hat{\mathbf{w}}^\nu, \hat{\mathbf{s}}^\nu, \mathbf{v}^\nu)} (\|\mathbf{x}_j\| + \frac{1}{\alpha}) < \lambda_j^\nu.$$

Let us bound the terms in the left-hand side of this inequality. At first, we have:

$$\begin{aligned} |\mathbf{x}_j^\top \hat{\mathbf{s}}^\nu - \hat{v}_j^\nu| &\leq |\mathbf{x}_j^\top \hat{\mathbf{s}} - \hat{v}_j| + |\mathbf{x}_j^\top (\hat{\mathbf{s}}^\nu - \hat{\mathbf{s}})| + |\hat{v}_j^\nu - \hat{v}_j| \\ &\leq |\mathbf{x}_j^\top \hat{\mathbf{s}} - \hat{v}_j| + \|\mathbf{x}_j\| \|\hat{\mathbf{s}}^\nu - \hat{\mathbf{s}}\| + |\hat{v}_j^\nu - \hat{v}_j|, \end{aligned}$$

and

$$\begin{aligned} \sqrt{G_{\Lambda^\nu}} &\leq \sqrt{|G_{\Lambda^\nu} - G_\Lambda| + |G_\Lambda|} \\ &\leq \sqrt{|G_{\Lambda^\nu} - G_\Lambda|} + \sqrt{|G_\Lambda|}. \end{aligned} \quad (13)$$

where the last inequality holds by applying the norm property $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1$ to the 2-dimensional vector of component $[\sqrt{|G_{\Lambda^\nu} - G_\Lambda|}, \sqrt{|G_\Lambda|}]$, where we have drop the dependence on $(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})$ and $(\hat{\mathbf{w}}^\nu, \hat{\mathbf{s}}^\nu, \hat{\mathbf{v}}^\nu)$ for simplicity. By gathering the pieces together, the left-hand side of Equation (10) can be bounded by

$$\begin{aligned} |\mathbf{x}_j^\top \hat{\mathbf{s}} - \hat{v}_j| + \|\mathbf{x}_j\| \|\hat{\mathbf{s}}^\nu - \hat{\mathbf{s}}\| + \sqrt{2G_\Lambda} (\|\mathbf{x}_j\| + \frac{1}{\alpha}) \\ + \sqrt{2|G_{\Lambda^\nu} - G_\Lambda|} (\|\mathbf{x}_j\| + \frac{1}{\alpha}) + |\hat{v}_j^\nu - \hat{v}_j|. \end{aligned} \quad (14)$$

which leads us to the novel screening test

$$T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}}) + \|\mathbf{x}_j\| (a + \sqrt{2b}) + c + \frac{1}{\alpha} \sqrt{2b} \leq \lambda_j^\nu. \quad (15)$$

with $T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})$ defined in Equation (10). \square

Algorithm 2 MM algorithm for Lasso with penalty $r_\lambda(|w|)$

Input: $\mathbf{X}, \mathbf{y}, \lambda, \mathbf{w}^0, \alpha$

- 1: $k \leftarrow 0$
 - 2: **repeat**
 - 3: $\Lambda^k = \{\lambda_j^k\}_{j=1,\dots,d} \leftarrow \{r'_\lambda(|w_j^k|)\}_{j=1,\dots,d}$
 - 4: compute approximate dual \mathbf{s}^k and duality gap G_{Λ^k} given \mathbf{w}^k and Λ^k
 - 5: **if needed then**
 - 6: compute screening scores $T_j^{(\lambda_j)}(\mathbf{w}^k, \mathbf{s}^k, \mathbf{v}^k)$ according to Eq. (10)
 - 7: store reference duality gap and approximate dual
 - 8: **else**
 - 9: estimate screening scores according to Eq. (12)
 - 10: **end if**
 - 11: LstScreen \leftarrow updated screened variables list based on results of Line 6 or 9
 - 12: $\mathbf{w}^{k+1}, \mathbf{y} - \mathbf{X}\mathbf{w}^{k+1} \leftarrow \text{PWL}(\mathbf{X}, \mathbf{y}, \mathbf{w}^k, \Lambda^k, \alpha, \text{LstScreen})$
 - 13: $k \leftarrow k + 1$
 - 14: **until** convergence
-

In order to make this screening test tractable, we need at first an approximation $\hat{\mathbf{s}}^\nu$ of the dual solution, then an upper bound on the norm of $\|\hat{\mathbf{s}} - \hat{\mathbf{s}}^\nu\|$ and a bound on the difference in duality gap $|G_{\Lambda^\nu} - G_\Lambda|$. In practice, we will define $\hat{\mathbf{s}}^\nu = \frac{\mathbf{y} - \mathbf{X}\hat{\mathbf{w}}}{\rho^\nu(j^\top)}$ and then compute exactly $\|\hat{\mathbf{s}} - \hat{\mathbf{s}}^\nu\|$ and $|G_{\Lambda^\nu} - G_\Lambda|$. Interestingly, since we consider the primal solution $\hat{\mathbf{w}}$ as our approximate primal solution for the new problem, computing $\rho^\nu(j)$ only costs an element-wise division since $\mathbf{y} - \mathbf{X}\hat{\mathbf{w}}$ and $|\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\hat{\mathbf{w}})|$ have already been pre-computed at the previous MM iteration. Another interesting point to highlight is that given pre-computed screening values $\{T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})\}$, the screening test given in Equation (12) does not involve any additional dot product and thus is cheaper to compute than the exact rule in Equation (10).

4.3. Algorithm

When considering MM algorithms for handling non-convex penalties, we advocate the use of weighted Lasso with screening as a solver for Equation (5) and a screening variable propagation condition as given in Equation (12). This last property allows us to cheaply evaluating whether a variable can be screened before entering in the inner solver. However, Equation (12) also needs the screening score computed at some previous $\{\lambda_j\}$ and a trade-off has thus to be sought between relevance of the test when λ_j^ν is not too far from λ_j and the computational time needed for evaluating $T_j^{(\lambda_j)}$. In practice, we compute the exact score $T_j^{(\lambda_j)}(\hat{\mathbf{w}}, \hat{\mathbf{s}}, \hat{\mathbf{v}})$ every 10 iterations and apply Equation (12) for the rest of the iterations. This results in Algorithm 2 where the inner solvers, denoted as PWL and solved ac-




Figure 1. Comparing running time for coordinate descent (CD) and proximal gradient algorithms as well as screening-based CD and MM algorithms under different tolerances on the stopping condition and under different samples, features settings. For all experiments, we have 5 active variables. (left) $n = 50$, $d = 100$, $\sigma = 2$. (right) $n = 500$, $d = 5000$, $\sigma = 2$

ording to Algorithm 1, are warm-started with previous outputs and provided with a list of already screened variables. In practice, $\alpha > 0$ helps us guaranteeing theoretical convergence of the sequence $\{\mathbf{w}^k\}$ and we have set $\alpha = 10^9$ for all our experiments, leading to very small regularization allowing large deviations from \mathbf{w}^k .

5. Numerical Experiments

The goal of screening rules is to improve the efficiency of solvers by focusing only on variables that are non-vanishing. In this section, we report the computational gains we obtain owing to our screening strategy⁶.

5.1. Experimental set-up

Our goal is to compare the computational running time of different algorithms for computing a regularization path on problem (1). For most relevant non-convex regularizers, coordinatewise update (Breheny & Huang, 2011) and proximal operator (Gong et al., 2013) can be derived in a closed-form. For our experiments, we have used the log-sum penalty which has an hyperparameter θ . Hence, our regularization path involves 2 parameters (λ_t, θ_t) . The set of $\{\lambda_t\}_{t=0}^{N_t-1}$ has been defined as $\lambda_t \triangleq \lambda_{\max} 10^{-\frac{3t}{N_t-1}}$, N_t being dependent of the problems and $\theta \in \{0.01, 0.1, 1\}$.

Our baseline method should have been an MM algorithm, in which each subproblem as given in Equation (5) is solved using a coordinate descent algorithm without screening. However, due to its very poor running time, we have omitted its performance. Two other competi-

tors that directly address the non-convex learning problems have instead been investigated: the first one, denoted as **GIST**, uses a majorization of the loss function and iterative shrinkage-thresholding (Gong et al., 2013) while the second one, named **ncxCD** is a coordinate descent algorithm that directly handles non-convex penalties (Breheny & Huang, 2011; Mazumder et al., 2011). We have named **MM-screen** our method that screens within each Proximal Weighted *Lasso* and propagates screening scores while the genuine version drops the screening propagation (but screening inside inner solvers is kept), is denoted **MM-genuine**. All algorithms have been stopped when optimality conditions as described in Equation (4) are satisfied up to the same tolerance τ . Note that all algorithms have been implemented in Python/Numpy. Hence, this may give GIST a slight computational advantage over coordinate descent approaches that heavily benefit on loop efficiency of low level languages. For our MM approaches, we stop the inner solver of the Proximal Weighted *Lasso* when the duality gap is smaller than 10^{-4} and screening is computed every 5 iterations. In the outer iterations, we perform the screening every 10 iterations.

5.2. Toy problem

Our toy regression problem has been built as follows. The entries of the regression design matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ are drawn uniformly from a Gaussian distribution of zero mean and variance 4. For given n, d and a number p of active variables, the true coefficient vector \mathbf{w}^* is obtained as follows. The p non-zero positions are chosen randomly, and their values are drawn from a zero-mean unit variance Gaussian distribution, to which we added ± 0.1 according $\text{sign}(w_j^*)$. Finally, the target vector is obtained as $\mathbf{y} = \mathbf{X}\mathbf{w}^* + \mathbf{e}$ where \mathbf{e} is a random noise vector drawn

⁶Additional details are provided in the supplementary material and the Python code for the experiment is available at https://github.com/arakotom/screening_ncvx_penalty




Figure 2. Computational gain of propagating screening within MM iterations (left) *w.r.t.* number of features at a KKT tolerance of 10^{-4} . (middle) *w.r.t.* to noise level at a KKT tolerance of 10^{-4} . (right) *w.r.t.* number of features at different tolerance in a low-noise regime.




Figure 3. Running time for computing the regularization path on (left) the dense Leukemia dataset $n = 72$, $d = 7129$. (right) sparse dataset Newsgroup with $n = 961$ and $d = 21319$.

from a Gaussian distribution with zero-mean and standard deviation σ . For the toy problem, we have set $N_t = 50$.

Figure 1 presents the running time needed for the different algorithms to reach convergence under various settings. Indifferently to the settings, our screening rules help in reducing the computational time by a factor of about 5 compared to a plain non-convex coordinate descent. The gain compared to **GIST** is mostly notable for high-precision and noisy problems.

We have also analyzed the benefits brought by our screening propagation strategy. Figure 2 shows the gain in computation time when comparing **MM screening** and **MM genuine**. As we have kept the number of relevant features to 5, one can note that for a fixed amount of noise the gain is around 1.8 for a wide range of dimensions. Similarly, the gain is almost constant for a large range of noise, the best benefit occurring in low-noise settings. More interestingly, we have compared gains for increasing precision and for a low-noise situation $\sigma = 0.01$, a classical noise level in the screening literature (Ndiaye et al., 2016; Tibshirani et al., 2012). We can note from the most right panel of Figure 2 that the higher the precision targeted, the more one benefits from screening propagation. For a tolerance of $\tau = 10^{-8}$, the gain peaks around 6.5 whereas in most regimes of number of features, the gain is about 4. Even for a lower tolerance, the genuine screening approach is 4

times slower than the full approach we propose.

5.3. Real-world datasets

We have run comparisons on different real datasets. We have reported on Figure 3 (left) the case for the *leukemia* dataset (dense data with $n = 72$ and $d = 7129$). For the path computation, the number of λ_t 's has been fixed to $N_t = 20$. Note that the gain compared to **ncvxCD** varies from 3 to 5 depending on the tolerance and is about 4 compared to **GIST** at high tolerance. On Figure 3 (right), we compare the gain in running time brought by the screening propagation rule on the *newsgroups* dataset in which we have kept only 2 categories (*religion* and *graphics*) resulting in $n = 961$, $d = 21319$. The gain is similar to what we observed earlier, ranging in between 1.3 and 1.8.

6. Conclusion

We have presented the first screening rule strategy that handles sparsity-inducing non-convex regularizers. The approach we propose relies on a majorization-minimization framework in which each inner iteration solves a Proximal Weighted *Lasso* problem. We introduced a screening rule for this learning problem and a rule for propagating screened variables within MM iterations. Interestingly, our screening rule for the weighted *Lasso* identifies the screenable variables in a finite amount of time. We have carried out numerical experiments showing the benefits of the proposed approach *w.r.t.* methods directly handling the non-convexity of the regularizers and illustrating the situation in which our propagating-screening rule helps accelerating the solver. A planned extension of this work is to consider pruning variables using active set strategy (Johnson & Guestrin, 2015) while keeping screening rules (Massias et al., 2018).

Acknowledgments

This work benefited from the support of the project OAT-MIL ANR-17-CE23-0012 of the French National Research Agency (ANR), the Normandie Projet GRR-DAISI, European funding FEDER DAISI.

References

- Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- Bonnefoy, A., Emiya, V., Ralaivola, L., and Gribonval, R. Dynamic screening: Accelerating first-order algorithms for the lasso and group-lasso. *IEEE Trans. Signal Process.*, 63(19):5121–5132, 2015.
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- Breheny, P. and Huang, J. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Statist.*, 5(1):232, 2011.
- Candès, E. J., Wakin, M. B., and Boyd, S. P. Enhancing sparsity by reweighted l_1 minimization. *J. Fourier Anal. Applicat.*, 14(5-6):877–905, 2008.
- Chen, S. and Donoho, D. Basis pursuit. In IEEE (ed.), *Proceedings of 1994 28th Asilomar Conference on Signals, Systems and Computers*, 1994.
- Chen, S. S., Donoho, D. L., and Saunders, M. A. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- Clarke, F. H. *Method of Dynamic and Nonsmooth Optimization*. SIAM, 1989.
- Donoho, D. L. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006.
- El Ghaoui, L., Viallon, V., and Rabbani, T. Safe feature elimination in sparse supervised learning. *Journal of Pacific Optimization*, pp. 667–698, 4 2012.
- Fan, J. and Li, R. Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96(456):1348–1360, 2001.
- Fercoq, O., Gramfort, A., and Salmon, J. Mind the duality gap: safer rules for the lasso. In *ICML*, pp. 333–342, 2015.
- Friedman, J., Hastie, T., Höfling, H., Tibshirani, R., et al. Pathwise coordinate optimization. *Ann. Appl. Stat.*, 1(2):302–332, 2007.
- Friedman, J., Hastie, T., and Tibshirani, R. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1, 2010.
- Fu, W. J. Penalized regressions: the bridge versus the lasso. *J. Comput. Graph. Statist.*, 7(3):397–416, 1998.
- Gasso, G., Rakotomamonjy, A., and Canu, S. Recovering sparse signals with a certain family of nonconvex penalties and dc programming. *IEEE Trans. Signal Process.*, 57(12):4686–4698, 2009.
- Gong, P., Zhang, C., Lu, Z., Huang, J., and Ye, J. A general iterative shrinkage and thresholding algorithm for non-convex regularized optimization problems. In *ICML*, pp. 37–45, 2013.
- Hunter, D. R. and Lange, K. A tutorial on MM algorithms. *The American Statistician*, 58(1):30–37, 2004.
- Johnson, T. and Guestrin, C. Blitz: A principled meta-algorithm for scaling sparse optimization. In *ICML*, pp. 1171–1179, 2015.
- Kang, Y., Zhang, Z., and Li, W.-J. On the global convergence of majorization minimization algorithms for nonconvex optimization problems. *arXiv preprint arXiv:1504.07791*, 2015.
- Kruger, A. Y. On Fréchet subdifferentials. *Journal of Mathematical Sciences*, 116(3):3325–3358, 2003.
- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. Compressed sensing MRI. *IEEE Signal Processing Magazine*, 25(2):72–82, 2008.
- Mairal, J. Stochastic majorization-minimization algorithms for large-scale optimization. In *Advances in Neural Information Processing Systems*, pp. 2283–2291, 2013.
- Massias, M., Gramfort, A., and Salmon, J. Celer: a fast solver for the lasso with dual extrapolation. In *ICML*, volume 80, pp. 3321–3330, 2018.
- Mazumder, R., Friedman, J. H., and Hastie, T. Sparsenet: Coordinate descent with nonconvex penalties. *J. Amer. Statist. Assoc.*, 106(495):1125–1138, 2011.
- Mordukhovich, B. S., Nam, N. M., and Yen, N. Fréchet subdifferential calculus and optimality conditions in nondifferentiable programming. *Optimization*, 55(5-6):685–708, 2006.
- Ndiaye, E., Fercoq, O., Gramfort, A., and Salmon, J. Gap safe screening rules for sparse-group lasso. In *Advances in Neural Information Processing Systems*, pp. 388–396, 2016.

- Ndiaye, E., Fercoq, O., Gramfort, A., and Salmon, J. Gap safe screening rules for sparsity enforcing penalties. *Journal of Machine Learning Research*, 18(128): 1–33, 2017.
- Nesterov, Y. *Introductory lectures on convex optimization*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, Boston, MA, 2004.
- Raj, A., Olbrich, J., Gärtner, B., Schölkopf, B., and Jaggi, M. Screening rules for convex problems. *arXiv preprint arXiv:1609.07478*, 2016.
- Rockafellar, R. T. and Wets, R. J.-B. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- Shevade, S. K. and Keerthi, S. S. A simple and efficient algorithm for gene selection using sparse logistic regression. *Bioinformatics*, 19(17):2246–2253, 2003.
- Shibagaki, A., Karasuyama, M., Hatano, K., and Takeuchi, I. Simultaneous safe screening of features and samples in doubly sparse modeling. In *ICML*, pp. 1577–1586, 2016.
- Soubies, E., Blanc-Féraud, L., and Aubert, G. A unified view of exact continuous penalties for ℓ_2 - ℓ_0 minimization. *SIAM J. Optim.*, 27(3):2034–2060, 2017.
- Tibshirani, R. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, pp. 267–288, 1996.
- Tibshirani, R., Bien, J., Friedman, J., Hastie, T., Simon, N., Taylor, J., and Tibshirani, R. J. Strong rules for discarding predictors in lasso-type problems. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 74(2):245–266, 2012.
- Tibshirani, R. J. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7:1456–1490, 2013.
- Ye, J. and Liu, J. Sparse methods for biomedical data. *ACM Sigkdd Explorations Newsletter*, 14(1):4–15, 2012.
- Zhang, C.-H. Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38(2):894–942, 2010a.
- Zhang, T. Analysis of multi-stage convex relaxation for sparse regularization. *Journal of Machine Learning Research*, 11(Mar):1081–1107, 2010b.
- Zou, H. and Hastie, T. J. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320, 2005.