



Minimization over the ℓ_1 -ball using an active-set non-monotone projected gradient

Andrea Cristofari¹ · Marianna De Santis²  · Stefano Lucidi² · Francesco Rinaldi³

Received: 30 July 2021 / Accepted: 30 July 2022 / Published online: 27 August 2022
© The Author(s) 2022

Abstract

The ℓ_1 -ball is a nicely structured feasible set that is widely used in many fields (e.g., machine learning, statistics and signal analysis) to enforce some sparsity in the model solutions. In this paper, we devise an active-set strategy for efficiently dealing with minimization problems over the ℓ_1 -ball and embed it into a tailored algorithmic scheme that makes use of a non-monotone first-order approach to explore the given subspace at each iteration. We prove global convergence to stationary points. Finally, we report numerical experiments, on two different classes of instances, showing the effectiveness of the algorithm.

Keywords Active-set methods · ℓ_1 -ball · LASSO · Large-scale optimization

✉ Marianna De Santis
mdesantis@diag.uniroma1.it

Andrea Cristofari
andrea.cristofari@uniroma2.it

Stefano Lucidi
lucidi@diag.uniroma1.it

Francesco Rinaldi
rinaldi@math.unipd.it

¹ Dipartimento di Ingegneria Civile e Ingegneria Informatica, Università degli Studi di Roma “Tor Vergata”, Rome, Italy

² Dipartimento di Ingegneria Informatica, Automatica e Gestionale, Sapienza Università di Roma, Rome, Italy

³ Dipartimento di Matematica “Tullio Levi-Civita”, Università di Padova, Padua, Italy

1 Introduction

In this paper, we focus on the following problem:

$$\begin{aligned} \min \quad & \varphi(x) \\ & \|x\|_1 \leq \tau, \end{aligned} \quad (1)$$

where $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a function whose gradient is Lipschitz continuous with constant $L > 0$, $\|x\|_1$ denotes the ℓ_1 -norm of the vector x and τ is a suitably chosen positive parameter.

Problem (1) includes, as a special case, the so called *LASSO problem*, obtained when

$$\varphi(x) = \|Ax - b\|^2,$$

with A and b being a $m \times n$ matrix and a m -dimensional vector, respectively. Here and in the following, $\|\cdot\|$ denotes the Euclidean norm. Loosely speaking, in LASSO problems the ℓ_1 -norm constraint is able to induce sparsity in the final solution, and then these problems are widely used in statistics to build regression models with a small number of non-zero coefficients [17, 32].

Standard optimization algorithms (like, e.g., interior-point methods), besides being very expensive when the number of variables increases, do not properly exploit the main features and structure of the considered problem. This is the reason why, in the last decade, a number of first-order methods have been considered in the literature to deal with problem (1). Those methods can be divided into two main classes: projection-based approaches, like, e.g., gradient-projection methods [15, 31] and limited-memory projected quasi-Newton methods [30], which efficiently handle the problem by making use of tailored projection strategies [8, 16], and projection-free methods, like, e.g., Frank–Wolfe variants [5, 6, 25, 26], that embed a cheap linear minimization oracle.

As highlighted before, the main goal when using the ℓ_1 ball is to get very sparse solutions (i.e., solutions with many zero components). In this context, it hence makes sense to devise strategies that allow to quickly identify the set of zero components in the optimal solution. This would indeed guarantee a significant speed-up of the optimization process. A number of active-set strategies for structured feasible sets is available in the literature (see, e.g., [3, 4, 7, 9, 10, 13, 18, 19, 22–24, 28] and references therein), but none of those directly handles the ℓ_1 ball.

In this paper, inspired by the work carried out in [10], we propose a tailored active-set strategy for problem (1) and embed it into a first-order projection-based algorithm. At each iteration, the method first sets to zero the variables that are guessed to be zero at the final solution. This is done by means of the tailored active-set estimate, which aims at identifying the manifold where the solutions of problem (1) lie, while guaranteeing, thanks to a descent property, a reduction of the objective function at each iteration. Then, the remaining variables, i.e., those variables estimated to be non-zero at the final solution, are suitably modified by means of a non-monotone gradient-projection step.

The paper is organized as follows. In Sect. 2, we describe the active-set strategy and analyze the descent property connected to it. We then devise, in Sect. 3, our first-order optimization algorithm and carry out a global convergence analysis. We further report a numerical comparison with some well-known first order methods using two different classes of \mathcal{L}_τ -constrained problems (that is, LASSO and constrained sparse logistic regression) in Sect. 4. Finally, we draw some conclusions in Sect. 5.

2 The active-set estimate

Since the feasible set of problem (1) is convex and can be written as convex combination of the vectors $\pm\tau e_i, i = 1, \dots, n$, we can characterize the stationary points as follows.

Definition 1 A feasible point x^* of problem (1) is stationary if and only if

$$\begin{aligned} \nabla\varphi(x^*)^T(\tau e_i - x^*) &\geq 0, \quad i = 1, \dots, n, \\ \nabla\varphi(x^*)^T(-\tau e_i - x^*) &\geq 0, \quad i = 1, \dots, n. \end{aligned} \tag{2}$$

In the next proposition, we state some ‘‘complementarity-type’’ conditions for stationary points of problem (1).

Proposition 1 Let x^* be a stationary point of problem (1). Then

- (i) $x_i^* > 0 \Rightarrow \nabla\varphi(x^*)^T(\tau e_i - x^*) = 0$,
- (ii) $x_i^* < 0 \Rightarrow \nabla\varphi(x^*)^T(-\tau e_i - x^*) = 0$.

Proof If $|x_i^*| = \tau$, then $x^* = \tau \operatorname{sgn}(x_i^*) e_i$ and the result trivially holds. To prove point (i), now let $0 < x_i^* < \tau$. Taking into account (2), by contradiction we assume that

$$\nabla\varphi(x^*)^T(\tau e_i - x^*) > 0. \tag{3}$$

Let $d^+ \in \mathbb{R}^n$ be defined as follows:

$$d^+ = \frac{x_i^*}{\tau - x_i^*}(x^* - \tau e_i).$$

We have

$$\begin{aligned} \|x^* + d^+\|_1 &= \left(1 + \frac{x_i^*}{\tau - x_i^*}\right) \sum_{j \neq i} |x_j^*| + \left|x_i^* + \frac{x_i^*}{\tau - x_i^*}(x_i^* - \tau)\right| \\ &\leq \left(1 + \frac{x_i^*}{\tau - x_i^*}\right)(\tau - |x_i^*|) = \tau, \end{aligned} \tag{4}$$

so that d^+ is a feasible direction in x^* . Therefore, (3) and (4) imply that d^+ is a feasible descent direction for $\varphi(\cdot)$ in x^* . This contradicts the fact that x^* is a stationary point of problem (1) and point (i) is proved. To prove point (ii), we can use the same arguments as above, considering $-\tau < x_i^* < 0$ and, assuming by contradiction that $\nabla\varphi(x^*)^T(-\tau e_i - x^*) > 0$, we obtain that

$$d^- = \frac{|x_i^*|}{\tau + x_i^*}(x^* + \tau e_i)$$

is such that $\|x^* + d^-\|_1 \leq \tau$, that is, d^- is a feasible and descent direction for $\varphi(\cdot)$ in x^* , leading to a contradiction. \square

With little abuse of standard terminology, given a stationary point x^* we say that a variable x_i^* is *active* if $x_i^* = 0$, whereas a variable x_i^* is said to be *non-active* if $x_i^* \neq 0$. We can thus define the *active set* $\bar{A}_{\ell_1}(x^*)$ and the *non-active set* $\bar{N}_{\ell_1}(x^*)$ as follows:

$$\bar{A}_{\ell_1}(x^*) = \{i : x_i^* = 0\}, \quad \bar{N}_{\ell_1}(x^*) = \{1, \dots, n\} \setminus \bar{A}_{\ell_1}(x^*).$$

Now, we show how we estimate these sets starting from any feasible point x of problem (1). In order to obtain such an estimate we first need to suitably reformulate our problem (1) by introducing a dummy variable z . Let $\bar{\varphi}(x, z) : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ be the function defined as $\bar{\varphi}(x, z) = \varphi(x)$ for all (x, z) . Problem (1) can then be rewritten as

$$\begin{aligned} \min \quad & \bar{\varphi}(x, z) \\ & \|x\|_1 + z \leq \tau, \\ & z \geq 0. \end{aligned} \tag{5}$$

Every feasible point of problem (5) can be expressed as convex combination of $\{\pm\tau e_1, \dots, \pm\tau e_n, \tau e_{n+1}\} \subset \mathbb{R}^{n+1}$. Therefore, we can define the following matrix, where I denotes the $n \times n$ identity matrix:

$$\bar{M} = \tau \left[\begin{array}{ccc|ccc} & & & & & 0 \\ & & & & & \vdots \\ & & & & & 0 \\ I & & & -I & & \\ \hline 0 & \dots & 0 & 0 & \dots & 0 \\ & & & & & 1 \end{array} \right] \in \mathbb{R}^{(n+1) \times (2n+1)},$$

and we obtain the following reformulation of (1) as a minimization problem over the unit simplex:

$$\begin{aligned} \min \quad & f(y) = \bar{\varphi}(\bar{M}y) \\ & e^T y = 1, \\ & y \geq 0. \end{aligned} \tag{6}$$

Note that, given any feasible point x of problem (1), we can compute a feasible point y of problem (6) such that

$$\begin{aligned}
 y_i &= \frac{1}{\tau} \max\{0, x_i\}, \quad i = 1, \dots, n, \\
 y_{n+i} &= \frac{1}{\tau} \max\{0, -x_i\}, \quad i = 1, \dots, n, \\
 y_{2n+1} &= \frac{\tau - \|x\|_1}{\tau}.
 \end{aligned}
 \tag{7}$$

The rationale behind our approach is sketched in the three following points:

- (i) For any feasible point x of problem (1), by (7) we can compute a feasible point y of problem (6) such that

$$y_i = 0 \Leftrightarrow x_i \leq 0 \quad \text{and} \quad y_{n+i} = 0 \Leftrightarrow x_i \geq 0, \quad i = 1, \dots, n. \tag{8}$$

- (ii) According to (8), for every feasible point x of problem (1) we have that

$$x_i = 0 \Leftrightarrow y_i = y_{n+i} = 0, \quad i = 1, \dots, n. \tag{9}$$

Thus, it is natural to estimate a variable x_i as active at x^* if both y_i and y_{n+i} are estimated to be zero at the point corresponding to x^* in the y space. To estimate the zero variables among y_1, \dots, y_{2n+1} we use the active-set estimate described in [10], specifically devised for minimization problems over the unit simplex.

- (iii) Then, we are able to go back in the original x space to obtain an active-set estimate of problem (1) without explicitly considering the variables y_1, \dots, y_{2n+1} of the reformulated problem.

Remark 1 The introduction of the dummy variable z is needed in order to get a reformulation of problem (1) as a minimization problem over the unit simplex satisfying (8). Since every feasible point x of problem (1) can be expressed as a convex combination of the vertices of the polyhedron $\{x \in \mathbb{R}^n : \|x\|_1 \leq \tau\}$, a straightforward reformulation of problem (1) would then be the following:

$$\min\{\varphi(My) : e^T y = 1, y \geq 0\}, \tag{10}$$

with $M = \tau \begin{bmatrix} I & -I \end{bmatrix}$. However, this reformulation does not work for our purposes, as there exist feasible points x of problem (1) for which no y feasible for problem (10) satisfying (8) can be found. In particular, if x is in the interior of the ℓ_1 -ball (e.g., the origin), we cannot find any y feasible for problem (10) such that (8) holds.

Considering problem (6) and using the active-set estimate proposed in [10] for minimization problems over the unit simplex, given any feasible point y of problem (6) we define:

$$A(y) = \{i : y_i \leq \epsilon \nabla f(y)^T (e_i - y)\}, \tag{11}$$

$$N(y) = \{i : y_i > \epsilon \nabla f(y)^T (e_i - y)\}, \tag{12}$$

where ϵ is a positive parameter. $A(y)$ contains the indices of the variables that are estimated to be zero at a certain stationary point and $N(y)$ contains the indices of the variables that are estimated to be positive at the same stationary point (see [10] for details of how these formulas are obtained). As mentioned above, taking into account (9), we estimate a variable x_i as active for problem (1) if both y_i and y_{n+1} are estimated to be zero. Namely,

$$A_{\epsilon_1}(x) = \{i \in \{1, \dots, n\} : i \in A(y) \text{ and } (n + i) \in A(y)\}, \tag{13a}$$

$$N_{\epsilon_1}(x) = \{i \in \{1, \dots, n\} : i \in N(y) \text{ or } (n + i) \in N(y)\}. \tag{13b}$$

Now we show how $A_{\epsilon_1}(x)$ and $N_{\epsilon_1}(x)$ can be expressed without explicitly considering the variables y and the objective function $f(y)$ of the reformulated problem. This allows us to work in the original x space, avoiding to double the number of variables in practice.

To obtain the desired relations, first observe that

$$\nabla f(y) = \bar{M}^T \nabla \bar{\varphi}(x) = \tau \begin{bmatrix} \nabla \varphi(x) \\ -\nabla \varphi(x) \\ 0 \end{bmatrix}^T, \tag{14}$$

and

$$\nabla f(y)^T y = \nabla \bar{\varphi}(x)^T \bar{M} y = [\nabla \varphi(x)^T \ 0] \bar{M} y = \nabla \varphi(x)^T x.$$

Let us distinguish two cases:

(i) $x_i \geq 0$. Recalling (11)–(12), we have that $i \in A(y)$ if and only if

$$\begin{aligned} 0 \leq \frac{1}{\tau} x_i = y_i &\leq \epsilon \nabla f(y)^T (e_i - y) = \epsilon (\nabla_i f(y) - \nabla f(y)^T y) \\ &= \epsilon (\tau \nabla_i \varphi(x) - \nabla \varphi(x)^T x) = \epsilon \nabla \varphi(x)^T (\tau e_i - x) \end{aligned} \tag{15}$$

and $(n + i) \in A(y)$ if and only if

$$\begin{aligned} -\frac{1}{\tau} x_i \leq 0 = y_{n+i} &\leq \epsilon \nabla f(y)^T (e_{n+i} - y) = \epsilon (\nabla_{n+i} f(y) - \nabla f(y)^T y) \\ &= \epsilon (-\tau \nabla_i \varphi(x) - \nabla \varphi(x)^T x) = -\epsilon \nabla \varphi(x)^T (\tau e_i + x). \end{aligned} \tag{16}$$

(ii) $x_i < 0$. Similarly to the previous case, we have that $i \in A(y)$ if and only if

$$\begin{aligned} \frac{1}{\tau} x_i < 0 = y_i &\leq \epsilon \nabla f(y)^T (e_i - y) = \epsilon (\nabla_i f(y) - \nabla f(y)^T y) \\ &= \epsilon (\tau \nabla_i \varphi(x) - \nabla \varphi(x)^T x) = \epsilon \nabla \varphi(x)^T (\tau e_i - x) \end{aligned} \tag{17}$$

and $(n + i) \in A(y)$ if and only if

$$\begin{aligned}
 0 < -\frac{1}{\tau}x_i = y_{n+i} \leq \epsilon \nabla f(y)^T(e_{n+i} - y) &= \epsilon(\nabla_{n+i}f(y) - \nabla f(y)^T y) \\
 &= \epsilon(-\tau \nabla_i \varphi(x) - \nabla \varphi(x)^T x) = -\epsilon \nabla \varphi(x)^T(\tau e_i + x).
 \end{aligned} \tag{18}$$

From (15), (16), (17) and (18), we thus obtain

$$\begin{aligned}
 A_{\ell_1}(x) = \{i : \epsilon \tau \nabla \varphi(x)^T(\tau e_i + x) \leq 0 \leq x_i \leq \epsilon \tau \nabla \varphi(x)^T(\tau e_i - x) \text{ or} \\
 \epsilon \tau \nabla \varphi(x)^T(\tau e_i + x) \leq x_i \leq 0 \leq \epsilon \tau \nabla \varphi(x)^T(\tau e_i - x)\},
 \end{aligned} \tag{19}$$

$$N_{\ell_1}(x) = \{1, \dots, n\} \setminus A_{\ell_1}(x). \tag{20}$$

Let us highlight again that $A_{\ell_1}(x)$ and $N_{\ell_1}(x)$ do not depend on the variables y and on the objective function $f(y)$ of the reformulated problem, so no variable transformation is needed in practice to estimate the active set of problem (1). In the following, we prove that under specific assumptions, $\bar{A}_{\ell_1}(x^*)$ is detected by our active-set estimate, when evaluated in points sufficiently close to a stationary point x^* .

Proposition 2 *If x^* is a stationary point of problem (1), then there exists an open ball $\mathcal{B}(x^*, \rho)$ with center x^* and radius $\rho > 0$ such that, for all $x \in \mathcal{B}(x^*, \rho)$, we have*

$$A_{\ell_1}(x) \subseteq \bar{A}_{\ell_1}(x^*), \tag{21}$$

$$\bar{N}_{\ell_1}(x^*) \subseteq N_{\ell_1}(x). \tag{22}$$

Furthermore, if the following ‘‘strict-complementarity-type’’ assumption holds:

$$x_i^* = 0 \Rightarrow \nabla \varphi(x^*)^T(\tau e_i - x^*) > 0 \wedge \nabla \varphi(x^*)^T(\tau e_i + x^*) < 0, \tag{23}$$

then, for all $x \in \mathcal{B}(x^*, \rho)$, we have

$$A_{\ell_1}(x) = \bar{A}_{\ell_1}(x^*), \tag{24}$$

$$\bar{N}_{\ell_1}(x^*) = N_{\ell_1}(x). \tag{25}$$

Proof Let $i \in N_{\ell_1}(x^*)$, then $|x_i^*| > 0$. Proposition 1 implies that either

$$\nabla \varphi(x^*)^T(\tau e_i - x^*) = 0 \quad \text{if } x_i^* > 0,$$

or

$$\nabla \varphi(x^*)^T(-\tau e_i - x^*) = 0 \quad \text{if } x_i^* < 0.$$

Then, the continuity of $\nabla \varphi$ and the definition of $N_{\ell_1}(x)$ imply that there exists an open ball $\mathcal{B}(x^*, \rho)$ with center x^* and radius $\rho > 0$ such that, for all $x \in \mathcal{B}(x^*, \rho)$, we have that $i \in N_{\ell_1}(x)$. This proves (22) and, consequently, also (21). If (23) holds, the definition of $N_{\ell_1}(x)$ and the continuity of $\nabla \varphi$ ensures that $\bar{A}_{\ell_1}(x^*) \subseteq A_{\ell_1}(x)$ for all $x \in \mathcal{B}(x^*, \rho)$, implying that (24) and (25) hold. \square

2.1 Descent property

So far, we have obtained the active and non-active set estimates (19)–(20) passing through a variable transformation which allowed us to adapt the active and non-active set estimates proposed in [10] to our problem (1).

In [10], the active and non-active set estimates, designed for minimization problems over the unit simplex, guarantee a decrease in the objective function when setting (some of) the estimated active variables to zero and moving a suitable estimated non-active variable (in order to maintain feasibility).

In the following, we show that the same property holds for problem (1) using the active and non-active set estimates (19)–(20). To this aim, in the next proposition we first introduce the index set $J_{\ell_1}(x)$ and relate it with $N_{\ell_1}(x)$.

Proposition 3 *Let $x \in \mathbb{R}^n$ be a feasible non-stationary point of problem (1) and define*

$$J_{\ell_1}(x) = \left\{ j : j \in \operatorname{Argmax}_{i=1,\dots,n} \{ |\nabla_i \varphi(x)| \} \right\}.$$

Then, $J_{\ell_1}(x) \subseteq N_{\ell_1}(x)$.

Proof Let y be the point given by (7) and consider the reformulated problem (6). Let $A(y)$ and $N(y)$ be the index sets given in (11)–(12), that is, the active and non-active set estimates for problem (6), respectively.

From the expression of $\nabla f(y)$ given in (14), and exploiting the hypothesis that x is non-stationary (implying that $\nabla \varphi(x) \neq 0$), it follows that

$$\min_{i=1,\dots,2n+1} \{ \nabla_i f(y) \} < 0. \tag{26}$$

Since $\nabla_{2n+1} f(y) = 0$ (again from (14)), it follows that

$$(2n + 1) \notin \operatorname{Argmin}_{i=1,\dots,2n+1} \{ \nabla_i f(y) \}.$$

From Proposition 1 in [10], there exists $v \in \{1, \dots, 2n\}$ such that

$$v \in \operatorname{Argmin}_{i=1,\dots,2n} \{ \nabla_i f(y) \}, \tag{27}$$

$$v \in N(y). \tag{28}$$

In particular, we can rewrite (27) as

$$\nabla_v f(y) = \tau \min_{i=1,\dots,n} \{ \nabla_1 \varphi(x), \dots, \nabla_n \varphi(x), -\nabla_1 \varphi(x), \dots, -\nabla_n \varphi(x) \}.$$

Taking into account (26), we obtain

$$-\nabla_v f(y) \leq -\tau |\nabla_i \varphi(x)|, \quad \forall i = 1, \dots, n. \tag{29}$$

Now, let $j \in \{1, \dots, n\}$ be the following index:

$$j = \begin{cases} v, & \text{if } v \in \{1, \dots, n\}, \\ v - n, & \text{if } v \in \{n + 1, \dots, 2n\}. \end{cases} \tag{30}$$

Using again (14), we get $|\nabla_w f(y)| = |\nabla_j f(y)| = \tau |\nabla_j \varphi(x)|$. This, combined with (29), implies that

$$j \in \text{Argmax}_{i=1, \dots, n} \{|\nabla_i \varphi(x)|\}.$$

Finally, using (28) and (30), it follows that at least one index between j and $(n + j)$ belongs to $N(y)$. Therefore, from (13b) we have that $j \in N_{\ell_1}(x)$ and the assertion is proved. \square

Now, we need an assumption on the parameter ϵ appearing in (19)–(20). It will allow us to prove the subsequent proposition, stating that $\varphi(x)$ decreases if we set the variables in $A_{\ell_1}(x)$ to zero and suitably move a variable in $J_{\ell_1}(x)$.

Assumption 1 Assume that the parameter ϵ appearing in the estimates (19)–(20) satisfies the following conditions:

$$0 < \epsilon \leq \frac{1}{\tau^2 n L (2C + 1)},$$

where $C > 0$ is a given constant.

Proposition 4 Let Assumption 1 hold. Given a feasible non-stationary point x of problem (1), let $j \in J_{\ell_1}(x)$ and $I = \{1, \dots, n\} \setminus \{j\}$. Let $\hat{A}_{\ell_1}(x)$ be a set of indices such that $\hat{A}_{\ell_1}(x) \subseteq A_{\ell_1}(x)$. Let \tilde{x} be the feasible point defined as follows:

$$\tilde{x}_{\hat{A}_{\ell_1}(x)} = 0; \quad \tilde{x}_{I \setminus \hat{A}_{\ell_1}(x)} = x_{I \setminus \hat{A}_{\ell_1}(x)}; \quad \tilde{x}_j = x_j - \text{sgn}(\nabla_j \varphi(x)) \sum_{h \in \hat{A}_{\ell_1}(x)} |x_h|.$$

Then,

$$\varphi(\tilde{x}) - \varphi(x) \leq -CL \|\tilde{x} - x\|^2,$$

where $C > 0$ is the constant appearing in Assumption 1.

Proof Define

$$\hat{A}^+ = \hat{A}_{\ell_1}(x) \cap \{i : x_i \neq 0\}. \tag{31}$$

Since $\nabla \varphi$ is Lipschitz continuous with constant L , from known results (see, e.g., [29]) we can write

$$\begin{aligned} \varphi(\tilde{x}) &\leq \varphi(x) + \nabla \varphi(x)^T (\tilde{x} - x) + \frac{L}{2} \|\tilde{x} - x\|^2 \\ &= \varphi(x) + \nabla \varphi(x)^T (\tilde{x} - x) + \frac{L(2C + 1)}{2} \|\tilde{x} - x\|^2 - CL \|\tilde{x} - x\|^2 \end{aligned}$$

and then, in order to prove the proposition, what we have to show is that

$$\nabla\varphi(x)^T(\tilde{x} - x) + \frac{L(2C + 1)}{2} \|\tilde{x} - x\|^2 \leq 0. \tag{32}$$

From the definition of \tilde{x} , we have that

$$\begin{aligned} \|\tilde{x} - x\|^2 &= \sum_{i \in \hat{A}^+} x_i^2 + \left(\sum_{i \in \hat{A}^+} |x_i| \right)^2 \leq \sum_{i \in \hat{A}^+} x_i^2 + |\hat{A}^+| \sum_{i \in \hat{A}^+} x_i^2 \\ &= (|\hat{A}^+| + 1) \sum_{i \in \hat{A}^+} x_i^2. \end{aligned} \tag{33}$$

Furthermore,

$$\begin{aligned} \nabla\varphi(x)^T(\tilde{x} - x) &= - \sum_{i \in \hat{A}^+} \nabla_i\varphi(x)x_i - |\nabla_j\varphi(x)| \sum_{i \in \hat{A}^+} |x_i| \\ &= \sum_{i \in \hat{A}^+} |x_i|(-\nabla_i\varphi(x) \operatorname{sgn}(x_i) - |\nabla_j\varphi(x)|). \end{aligned} \tag{34}$$

Since $j \in J_{\ell_1}(x)$, from the definition of $J_{\ell_1}(x)$ it follows that $-\nabla_i\varphi(x) \geq -|\nabla_j\varphi(x)|$ for all $i \in \{1, \dots, n\}$. Therefore, we can write

$$\begin{aligned} \nabla\varphi(x)^T x &= \sum_{i=1}^n \nabla_i\varphi(x) \operatorname{sgn}(x_i) |x_i| \geq \sum_{i=1}^n -|\nabla_j\varphi(x)| |x_i| \\ &= -|\nabla_j\varphi(x)| \|x\|_1 \geq -|\nabla_j\varphi(x)| \tau. \end{aligned} \tag{35}$$

Using (19) and (35), for all $i \in \hat{A}^+$ we have that

$$\begin{aligned} x_i &\leq \epsilon\tau(\nabla_i\varphi(x)\tau - \nabla\varphi(x)^T x) \leq \epsilon\tau^2(\nabla_i\varphi(x) + |\nabla_j\varphi(x)|), \\ -x_i &\leq -\epsilon\tau(\nabla_i\varphi(x)\tau + \nabla\varphi(x)^T x) \leq \epsilon\tau^2(-\nabla_i\varphi(x) + |\nabla_j\varphi(x)|), \end{aligned}$$

and then,

$$|x_i| = \operatorname{sgn}(x_i) x_i \leq \epsilon\tau^2(\nabla_i\varphi(x) \operatorname{sgn}(x_i) + |\nabla_j\varphi(x)|), \quad \forall i \in \hat{A}^+.$$

Combining this inequality with (33), we obtain

$$\|\tilde{x} - x\|^2 \leq \epsilon\tau^2(|\hat{A}^+| + 1) \sum_{i \in \hat{A}^+} |x_i|(\nabla_i\varphi(x) \operatorname{sgn}(x_i) + |\nabla_j\varphi(x)|) \tag{36}$$

From (34) and (36), it follows that the left-hand side term of (32) is less than or equal to

$$\left(\epsilon\tau^2 \frac{L(2C + 1)}{2} (|\hat{A}^+| + 1) - 1 \right) \sum_{i \in \hat{A}^+} |x_i|(\nabla_i\varphi(x) \operatorname{sgn}(x_i) + |\nabla_j\varphi(x)|)$$

The desired result is hence obtained, since inequality (32) follows from the assumption we made on ϵ , using the fact that $|\hat{A}^+| \leq n - 1$ (as a consequence of Proposition 3) and $\sum_{i \in \hat{A}^+} |x_i|(|\nabla_i \varphi(x)| \operatorname{sgn}(x_i) + |\nabla_j \varphi(x)|) \geq 0$ (as a consequence of (36)). \square

We would like to highlight that the parameter ϵ depends on n by Assumption 1. However, from the proof of the above proposition, it is clear that n could be replaced by $|\hat{A}^+| + 1$, with \hat{A}^+ defined as in (31). Note that $|\hat{A}^+|$ might be much smaller than n .

3 The algorithm

Based on the active and non-active set estimates described above, we design a suitable active-set algorithm for solving problem (1), exploiting the property of our estimates and using an appropriate projected-gradient direction. At the beginning of each iteration k , we have a feasible point x^k and we compute $A_{\ell_1}(x^k)$ and $N_{\ell_1}(x^k)$, which, for ease of notation, we will refer to as $A_{\ell_1}^k$ and $N_{\ell_1}^k$, respectively. Then, we perform two main steps:

- First, we produce the point \tilde{x}^k as explained in Proposition 4, obtaining a decrease in the objective function (if $x^k \neq \tilde{x}^k$);
- Afterward, we move all the variables in $N_{\ell_1}^k$ by computing a projected-gradient direction d^k over the given non-active manifold and using a non-monotone Armijo line search. In particular, the reference value $\bar{\varphi}$ for the line search is defined as the maximum among the last n_m function evaluations, with n_m being a positive parameter.

In Algorithm 1, we report the scheme of the proposed algorithm, named Active-Set algorithm for minimization over the ℓ_1 -ball (AS- ℓ_1).

Algorithm 1 Active-Set algorithm for minimization over the ℓ_1 -ball (AS- ℓ_1)

- 1 Choose a feasible point x^0 and choose $\epsilon > 0$
 - 2 For $k = 0, 1, \dots$
 - 3 If x^k is a stationary point, then STOP
 - 4 Compute $A_{\ell_1}^k = A_{\ell_1}(x^k)$ and $N_{\ell_1}^k = N_{\ell_1}(x^k)$
 - 5 Compute $J_{\ell_1}^k = J_{\ell_1}(x^k)$, choose $j \in J_{\ell_1}^k$ and define $\tilde{N}_{\ell_1}^k = N_{\ell_1}^k \setminus \{j\}$
 - 6 Set $\tilde{x}_{A_{\ell_1}^k}^k = 0$, $\tilde{x}_{\tilde{N}_{\ell_1}^k}^k = x_{\tilde{N}_{\ell_1}^k}^k$ and $\tilde{x}_j^k = x_j^k - \operatorname{sgn}(\nabla_j \varphi(x^k)) \sum_{h \in A_{\ell_1}^k} |x_h^k|$
 - 7 Compute a projected-gradient type direction d^k such that $d_{A_{\ell_1}^k}^k = 0$
 - 8 Compute a stepsize $\alpha^k \in [0, 1]$ by Algorithm 2
 - 9 Set $x^{k+1} = \tilde{x}^k + \alpha^k d^k$
 - 10 End for
-

Algorithm 2 Non-monotone Armijo line search

- 0 Choose $\delta \in (0, 1)$, $n_m > 0$ and $\gamma \in (0, 1)$
 - 1 Update $\bar{\varphi}^k = \max_{0 \leq i \leq \min\{n_m, k\}} \varphi(\tilde{x}^{k-i})$
 - 2 If $\nabla\varphi(\tilde{x}^k)^T d^k < 0$ then
 - 3 Set $\alpha = 1$
 - 4 While $\varphi(\tilde{x}^k + \alpha d^k) > \bar{\varphi}^k + \gamma \alpha \nabla\varphi(\tilde{x}^k)^T d^k$
 - 5 Set $\alpha = \delta\alpha$
 - 6 End while
 - 7 Else
 - 8 Set $\alpha = 0$
 - 9 End if
 - 10 Set $\alpha^k = \alpha$
-

The search direction d^k at \tilde{x}^k (see line 7 of Algorithm 1) is made of two subvectors: $d^k_{A^k_{\ell_1}}$ and $d^k_{N^k_{\ell_1}}$. Since we do not want to move the variables in $A^k_{\ell_1}$, we simply set $d^k_{A^k_{\ell_1}} = 0$. For $d^k_{N^k_{\ell_1}}$, we compute a projected gradient direction in a properly defined manifold. In particular, let $\mathcal{B}_{N^k_{\ell_1}}$ be the set defined as

$$\mathcal{B}_{N^k_{\ell_1}} = \{x \in \mathbb{R}^n : \|x\|_1 \leq \tau, x_i = 0, \forall i \notin N^k_{\ell_1}\} \tag{37}$$

and let $P(\cdot)_{\mathcal{B}_{N^k_{\ell_1}}}$ denote the projection onto the $\mathcal{B}_{N^k_{\ell_1}}$. We also define

$$\hat{x}^k = P(\tilde{x}^k - m^k \nabla\varphi(\tilde{x}^k))_{\mathcal{B}_{N^k_{\ell_1}}}, \tag{38}$$

where $0 < \underline{m} \leq m^k \leq \bar{m} < \infty$ and with \underline{m} , \bar{m} being two constants. Then, $d^k_{N^k_{\ell_1}}$ is defined as

$$d^k_{N^k_{\ell_1}} = \hat{x}^k - \tilde{x}^k. \tag{39}$$

In the practical implementation of AS- ℓ_1 , we compute the coefficient m^k so that the resulting search direction is a spectral (or Barzilai–Borwein) gradient direction. This choice will be described in Sect. 4.

3.1 Global convergence analysis

In order to prove global convergence of AS- ℓ_1 to stationary points, we need some intermediate results. We first point out a property of our search directions, using standard results on projected directions.

Lemma 1 *Let Assumption 1 hold and let $\{x^k\}$ be the sequence of points produced by AS- ℓ_1 . At every iteration k , we have that*

$$\nabla\varphi(\tilde{x}^k)^T d^k \leq -\frac{1}{\bar{m}}\|d^k\|^2 \tag{40}$$

and $\{d^k\}$ is a bounded sequence.

Proof Using the properties of the projection, at very iteration k we have

$$(\tilde{x}^k - m^k \nabla\varphi(\tilde{x}^k) - \hat{x}^k)^T (x - \hat{x}^k) \leq 0, \quad \forall x \in \mathcal{B}_{N_{\ell_1}^k},$$

with $\mathcal{B}_{N_{\ell_1}^k}$ and \hat{x}^k being defined as in (37) and (38), respectively. Choosing $x = \tilde{x}^k$ in the above inequality and recalling the definition of d^k given in (39), we get

$$\nabla\varphi(\tilde{x}^k)^T d^k \leq -\frac{1}{m^k}\|d^k\|^2.$$

Since $m^k \leq \bar{m}$, for all k we obtain (40).

Furthermore, from the property of the projection we have that

$$\|d^k\| = \|P(\tilde{x}^k - m^k \nabla\varphi(\tilde{x}^k)) - \tilde{x}^k\| \leq m^k \|\nabla\varphi(\tilde{x}^k)\|.$$

Since $m^k \leq \bar{m}$ and $\{\nabla\varphi(\tilde{x}^k)\}$ is bounded, it follows that $\{d^k\}$ is bounded. □

We now prove that the sequence $\{\bar{\varphi}^k\}$ converges.

Lemma 2 *Let Assumption 1 hold and let $\{x^k\}$ be the sequence of points produced by AS- ℓ_1 . Then, the sequence $\{\bar{\varphi}^k\}$ is non-increasing and converges to a value $\bar{\varphi}$.*

Proof First note that the definition of $\bar{\varphi}^k$ ensures $\bar{\varphi}^k \leq \varphi(\tilde{x}^0)$ and hence $\varphi(\tilde{x}^k) \leq \varphi(\tilde{x}^0)$ for all k . Moreover, we have that

$$\bar{\varphi}^{k+1} = \max_{0 \leq i \leq \min\{n_m, k+1\}} \varphi(\tilde{x}^{k+1-i}) \leq \max\{\bar{\varphi}^k, \varphi(\tilde{x}^{k+1})\}.$$

Since $\varphi(\tilde{x}^{k+1}) \leq \bar{\varphi}^k$ by the definition of the line search, we derive $\bar{\varphi}^{k+1} \leq \bar{\varphi}^k$, which proves that the sequence $\{\bar{\varphi}^k\}$ is non-increasing. This sequence is bounded from below by the minimum of φ over the unit simplex and hence converges. □

The next intermediate result shows that the distance between $\{x^k\}$ and $\{\tilde{x}^k\}$ converges to zero and that the sequences $\{\varphi(x^k)\}$ and $\{\varphi(\tilde{x}^k)\}$ converge to the same point, using similar arguments as in [20].

Proposition 5 *Let Assumption 1 hold and let $\{x^k\}$ be the sequence of points produced by AS- ℓ_1 . Then,*

$$\lim_{k \rightarrow \infty} \|\tilde{x}^k - x^k\| = 0, \tag{41}$$

$$\lim_{k \rightarrow \infty} \varphi(\tilde{x}^k) = \lim_{k \rightarrow \infty} \varphi(x^k) = \bar{\varphi}. \tag{42}$$

Proof For each $k \in \mathbb{N}$, choose $l(k) \in \{k - \min(k, n_m), \dots, k\}$ such that $\bar{\varphi}^k = \varphi(\tilde{x}^{l(k)})$. From Proposition 4 we can write

$$\varphi(\tilde{x}^{l(k)}) \leq \varphi(x^{l(k)}) - CL\|\tilde{x}^{l(k)} - x^{l(k)}\|^2. \tag{43}$$

Furthermore, from the instructions of the line search and the fact that the sequence $\{\varphi(\tilde{x}^{l(k)})\}$ is non-increasing, for all $k \geq 1$ we have

$$\varphi(x^{l(k)}) \leq \varphi(\tilde{x}^{l(k-1)}) + \gamma \alpha^{l(k)-1} \nabla \varphi(\tilde{x}^{l(k-1)})^T d^{l(k)-1},$$

and then,

$$\varphi(\tilde{x}^{l(k)}) \leq \varphi(\tilde{x}^{l(k-1)}) + \gamma \alpha^{l(k)-1} \nabla \varphi(\tilde{x}^{l(k-1)})^T d^{l(k)-1} - CL\|\tilde{x}^{l(k)} - x^{l(k)}\|^2. \tag{44}$$

Since $\{\varphi(\tilde{x}^{l(k)})\}$ converges to $\bar{\varphi}$, we have that (43) and (44) imply

$$\begin{aligned} \lim_{k \rightarrow \infty} \|\tilde{x}^{l(k)} - x^{l(k)}\| &= 0, \\ \lim_{k \rightarrow \infty} \alpha^{l(k)-1} \nabla \varphi(\tilde{x}^{l(k-1)})^T d^{l(k)-1} &= 0. \end{aligned} \tag{45}$$

Furthermore, from Lemma 1 we have

$$\nabla \varphi(\tilde{x}^{l(k-1)})^T d^{l(k)-1} \leq -\frac{1}{m} \|d^{l(k)-1}\|^2,$$

and then the following limit holds:

$$\lim_{k \rightarrow \infty} \alpha^{l(k)-1} \|d^{l(k)-1}\| = 0. \tag{46}$$

Considering that $x^{l(k)} = \tilde{x}^{l(k)-1} + \alpha^{l(k)-1} d^{l(k)-1}$, (46) implies

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{l(k)-1} - x^{l(k)}\| = 0.$$

Furthermore, from the triangle inequality, we can write

$$\|\tilde{x}^{l(k)-1} - \tilde{x}^{l(k)}\| \leq \|\tilde{x}^{l(k)-1} - x^{l(k)}\| + \|x^{l(k)} - \tilde{x}^{l(k)}\|.$$

Then,

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{l(k)-1} - \tilde{x}^{l(k)}\| = 0 \tag{47}$$

and in particular, from the uniform continuity of φ over $\{x \in \mathbb{R}^n : \|x\|_1 \leq \tau\}$, we have

$$\lim_{k \rightarrow \infty} \varphi(\tilde{x}^{l(k)-1}) = \lim_{k \rightarrow \infty} \varphi(\tilde{x}^{l(k)}) = \bar{\varphi}. \tag{48}$$

Let

$$\hat{l}(k) = l(k + n_m + 2).$$

We show by induction that, for any given $j \geq 1$,

$$\lim_{k \rightarrow \infty} \|x^{\hat{l}(k)-(j-1)} - \tilde{x}^{\hat{l}(k)-(j-1)}\| = 0, \tag{49}$$

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{\hat{l}(k)-(j-1)} - \tilde{x}^{\hat{l}(k)-j}\| = 0, \tag{50}$$

$$\lim_{k \rightarrow \infty} \varphi(\tilde{x}^{\hat{l}(k)-j}) = \lim_{k \rightarrow \infty} \varphi(\tilde{x}^{\hat{l}(k)}). \tag{51}$$

If $j = 1$, since $\{\hat{l}(k)\} \subset \{l(k)\}$ we have that (49), (50) and (51) follow from (45), (47) and (48), respectively.

Assume now that (49), (50) and (51) hold for a given j . Then, reasoning as in the beginning of the proof, from the instructions of the line search and considering that $\{\varphi(\tilde{x}^{\hat{l}(k)})\}$ is non-increasing, we can write

$$\varphi(\tilde{x}^{\hat{l}(k)-j}) \leq \varphi(x^{\hat{l}(k)-j}) - CL\|\tilde{x}^{\hat{l}(k)-j} - x^{\hat{l}(k)-j}\|^2$$

and

$$\varphi(x^{\hat{l}(k)-j}) \leq \varphi(\tilde{x}^{\hat{l}(k)-(j+1)}) + \gamma\alpha^{\hat{l}(k)-(j+1)}\nabla\varphi(\tilde{x}^{\hat{l}(k)-(j+1)})^\top d^{\hat{l}(k)-(j+1)}.$$

Therefore we get

$$\begin{aligned} \varphi(\tilde{x}^{\hat{l}(k)-j}) &\leq \varphi(\tilde{x}^{\hat{l}(k)-(j+1)}) + \gamma\alpha^{\hat{l}(k)-(j+1)}\nabla\varphi(\tilde{x}^{\hat{l}(k)-(j+1)})^\top d^{\hat{l}(k)-(j+1)} + \\ &\quad - CL\|\tilde{x}^{\hat{l}(k)-j} - x^{\hat{l}(k)-j}\|^2, \end{aligned}$$

so that

$$\lim_{k \rightarrow \infty} \alpha^{\hat{l}(k)-(j+1)}\nabla\varphi(\tilde{x}^{\hat{l}(k)-(j+1)})^\top d^{\hat{l}(k)-(j+1)} = 0, \tag{52}$$

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{\hat{l}(k)-j} - x^{\hat{l}(k)-j}\| = 0. \tag{53}$$

The limit in (53) implies (49) for $j + 1$. The properties of the direction stated in Lemma 1, combined with (52), ensure that

$$\lim_{k \rightarrow \infty} \alpha^{\hat{l}(k)-(j+1)}\|d^{\hat{l}(k)-(j+1)}\| = 0. \tag{54}$$

Furthermore, since $x^{\hat{l}(k)-j} = \tilde{x}^{\hat{l}(k)-(j+1)} + \alpha^{\hat{l}(k)-(j+1)}d^{\hat{l}(k)-(j+1)}$, we have that (54) implies

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{\hat{l}(k)-(j+1)} - x^{\hat{l}(k)-j}\| = 0.$$

Using the triangle inequality, we can write

$$\|\tilde{x}^{\hat{l}(k)-(j+1)} - \tilde{x}^{\hat{l}(k)-j}\| \leq \|\tilde{x}^{\hat{l}(k)-(j+1)} - x^{\hat{l}(k)-j}\| + \|x^{\hat{l}(k)-j} - \tilde{x}^{\hat{l}(k)-j}\|.$$

Then,

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{\hat{l}(k)-(j+1)} - \tilde{x}^{\hat{l}(k)-j}\| = 0$$

and in particular, from the uniform continuity of φ over $\{x \in \mathbb{R}^n : \|x\|_1 \leq \tau\}$, we can write

$$\lim_{k \rightarrow \infty} \varphi(\tilde{x}^{\hat{l}(k)-(j+1)}) = \lim_{k \rightarrow \infty} \varphi(\tilde{x}^{\hat{l}(k)-j}) = \bar{\varphi}.$$

Thus we conclude that (50) and (51) hold for any given $j \geq 1$.

Recalling that

$$\begin{aligned} \hat{l}(k) - (k + 1) &= l(k + n_m + 2) - (k + 1) \leq n_m + 1, \\ \|\tilde{x}^{k+1} - \tilde{x}^{\hat{l}(k)}\| &\leq \sum_{j=k+1}^{\hat{l}(k)-1} \|\tilde{x}^{j+1} - \tilde{x}^j\|, \end{aligned}$$

we have that (50) implies

$$\lim_{k \rightarrow \infty} \|\tilde{x}^{k+1} - \tilde{x}^{\hat{l}(k)}\| = 0. \tag{55}$$

Furthermore, since

$$\|x^{k+1} - \tilde{x}^{\hat{l}(k)}\| \leq \|x^{k+1} - \tilde{x}^{k+1}\| + \|\tilde{x}^{k+1} - \tilde{x}^{\hat{l}(k)}\|,$$

from (55) and (49) we have

$$\lim_{k \rightarrow \infty} \|x^{k+1} - \tilde{x}^{\hat{l}(k)}\| = 0. \tag{56}$$

Since $\{\varphi(\tilde{x}^{\hat{l}(k)})\}$ has a limit, from the uniform continuity of φ over $\{x \in \mathbb{R}^n : \|x\|_1 \leq \tau\}$, (56) and (55) it follows that

$$\lim_{k \rightarrow \infty} \varphi(x^{k+1}) = \lim_{k \rightarrow \infty} \varphi(x^k) = \lim_{k \rightarrow \infty} \varphi(\tilde{x}^{\hat{l}(k)}) = \bar{\varphi}$$

and

$$\lim_{k \rightarrow \infty} \varphi(\tilde{x}^{k+1}) = \lim_{k \rightarrow \infty} \varphi(\tilde{x}^k) = \lim_{k \rightarrow \infty} \varphi(\tilde{x}^{\hat{l}(k)}) = \bar{\varphi},$$

proving (42). From the instructions of the algorithm and Proposition 4, we can write

$$\varphi(\tilde{x}^k) \leq \varphi(x^k) - CL\|\tilde{x}^k - x^k\|^2,$$

and then from (42) we have that (41) holds. □

The following proposition states that the directional derivative $\nabla\varphi(\tilde{x}^k)^T d^k$ tends to zero.

Proposition 6 *Let Assumption 1 hold and let $\{x^k\}$ be the sequence of points produced by AS- ℓ_1 . Then,*

$$\lim_{k \rightarrow \infty} \nabla \varphi(\tilde{x}^k)^T d^k = 0. \tag{57}$$

Proof To prove (57), assume by contradiction that it does not hold. Lemma 1 implies that the sequence $\{\nabla \varphi(\tilde{x}^k)^T d^k\}$ is bounded, so that there must exist an infinite set $K \subseteq \mathbb{N}$ such that

$$\nabla \varphi(\tilde{x}^k)^T d^k < 0, \quad \forall k \in K, \tag{58}$$

$$\lim_{k \rightarrow \infty, k \in K} \nabla \varphi(\tilde{x}^k)^T d^k = -\eta < 0, \tag{59}$$

for some real number $\eta > 0$. Taking into account (41) and the fact that the feasible set is compact, without loss of generality we can assume that both $\{x^k\}_K$ and $\{\tilde{x}^k\}_K$ converge to a feasible point x^* (passing into a further subsequence if necessary). Namely,

$$\lim_{k \rightarrow \infty, k \in K} x^k = \lim_{k \rightarrow \infty, k \in K} \tilde{x}^k = x^*. \tag{60}$$

Moreover, since the number of possible different choices of A^k and N^k is finite, without loss of generality we can also assume that

$$A^k = \hat{A}, \quad N^k = \hat{N}, \quad \forall k \in K,$$

and, using the fact that $\{d^k\}$ is a bounded sequence, that

$$\lim_{k \rightarrow \infty, k \in K} d^k = \bar{d} \tag{61}$$

(passing again into a further subsequence if necessary). From (59), (60), (61) and the continuity of $\nabla \varphi$, we can write

$$\nabla \varphi(x^*)^T \bar{d} = -\eta < 0. \tag{62}$$

Taking into account (58), from the instructions of AS- \mathcal{L}_1 we have that, at every iteration $k \in K$, a non-monotone Armijo line search is carried out (see line 2 in Algorithm 2) and a value $\alpha^k \in (0, 1]$ is computed such that

$$\varphi(x^{k+1}) \leq \varphi(\tilde{x}^{l(k)}) + \gamma \alpha^k \nabla \varphi(\tilde{x}^k)^T d^k,$$

or equivalently,

$$\varphi(\tilde{x}^{l(k)}) - \varphi(x^{k+1}) \geq \gamma \alpha^k |\nabla \varphi(\tilde{x}^k)^T d^k|.$$

From

(42), the left-hand side of the above inequality converges to zero for $k \rightarrow \infty$, hence

$$\lim_{k \rightarrow \infty, k \in K} \alpha^k |\nabla \varphi(\tilde{x}^k)^T d^k| = 0.$$

Using (59), we obtain that $\lim_{k \rightarrow \infty, k \in K} \alpha^k = 0$. It follows that there exists $\bar{k} \in K$ such that

$$\alpha^k < 1, \quad \forall k \geq \bar{k}, k \in K.$$

From the instructions of the line search procedure, this means that $\forall k \geq \bar{k}, k \in K$

$$\varphi\left(\tilde{x}^k + \frac{\alpha^k}{\delta} d^k\right) > \varphi(\tilde{x}^{l(k)}) + \gamma \frac{\alpha^k}{\delta} \nabla \varphi(\tilde{x}^k)^T d^k \geq \varphi(\tilde{x}^k) + \gamma \frac{\alpha^k}{\delta} \nabla \varphi(\tilde{x}^k)^T d^k. \quad (63)$$

Using the mean value theorem, $\xi^k \in (0, 1)$ exists such that

$$\varphi\left(\tilde{x}^k + \frac{\alpha^k}{\delta} d^k\right) = \varphi(\tilde{x}^{l(k)}) + \frac{\alpha^k}{\delta} \nabla \varphi\left(\tilde{x}^k + \xi^k \frac{\alpha^k}{\delta} d^k\right)^T d^k, \quad \forall k \geq \bar{k}, k \in K. \quad (64)$$

In view of (63) and (64), we can write

$$\nabla \varphi\left(\tilde{x}^k + \xi^k \frac{\alpha^k}{\delta} d^k\right)^T d^k > \gamma \nabla \varphi(\tilde{x}^k)^T d^k, \quad \forall k \geq \bar{k}, k \in K. \quad (65)$$

From (60), and exploiting the fact that $\{\xi^k\}_K, \{\alpha^k\}_K$ and $\{d^k\}_K$ are bounded sequences, we get

$$\lim_{k \rightarrow \infty, k \in K} \tilde{x}^k + \xi^k \frac{\alpha^k}{\delta} d^k = \lim_{k \rightarrow \infty, k \in K} \tilde{x}^k = x^*.$$

Therefore, taking the limits in (65) we obtain that $\nabla \varphi(x^*)^T \bar{d} \geq \gamma \nabla \varphi(x^*)^T \bar{d}$, or equivalently, $(1 - \gamma) \nabla \varphi(x^*)^T \bar{d} \geq 0$. Since $\gamma \in (0, 1)$, we get a contradiction with (62). \square

We are finally able to state the main convergence result.

Theorem 1 *Let Assumption 1 hold and let $\{x^k\}$ be the sequence of points produced by AS- ℓ_1 . Then, every limit point x^* of $\{x^k\}$ is a stationary point of problem (1).*

Proof From Definition 1, we can characterize stationarity using condition (2). In particular, we can define the following continuous functions $\Psi_i(x)$ to measure the stationarity violation at a feasible point x :

$$\Psi_i(x) = \max\{0, -\nabla \varphi(x)^T (\tau e_i - x), -\nabla \varphi(x)^T (-\tau e_i - x)\}, \quad i = 1, \dots, n,$$

so that a feasible point x is stationary if and only if $\Psi_i(x) = 0, i = 1, \dots, n$.

Now, let x^* be a limit point of $\{x^k\}$ and let $\{x^k\}_K, K \subseteq \mathbb{N}$, be a subsequence converging to x^* . Namely,

$$\lim_{k \rightarrow \infty, k \in K} x^k = x^*. \quad (66)$$

Note that x^* exists, as $\{x^k\}$ remains in the compact set $\{x \in \mathbb{R}^n \mid \|x\|_1 \leq \tau\}$. Since the number of possible different choices of A^k and N^k is finite, without loss of generality we can assume that

$$A^k = \hat{A}, \quad N^k = \hat{N}, \quad \forall k \in K$$

(passing into a further subsequence if necessary).

By contradiction, assume that x^* is non-stationary, that is, an index $\nu \in \{1, \dots, n\}$ exists such that

$$\Psi_\nu(x^*) > 0. \tag{67}$$

First, suppose that $\nu \in \hat{A}$. Then, from the expressions (19), we can write

$$0 \leq x_\nu^k \leq \epsilon \tau \nabla \varphi(x^k)^T (\tau e_\nu - x^k) \quad \text{or} \quad 0 \geq x_\nu^k \geq \epsilon \tau \nabla \varphi(x^k)^T (\tau e_\nu + x^k),$$

so that $\Psi_\nu(x^k) = 0$, for all $k \in \bar{K}$. Therefore, from (66), the continuity of $\nabla \varphi$ and the continuity of the functions Ψ_i , we get $\Psi_\nu(x^*) = 0$, contradicting (67).

Then, ν necessarily belongs to \hat{N} . Namely, x^* is non-stationary over $\mathcal{B}_{N_{\epsilon_1}^k}$, with $\mathcal{B}_{N_{\epsilon_1}^k}$ defined as in (37). This means that

$$x^* \neq P(x^* - \underline{m} \nabla \varphi(x^*))_{\mathcal{B}_{N_{\epsilon_1}^k}}. \tag{68}$$

Using Proposition 6 and Lemma 1, we have that $\lim_{k \rightarrow \infty, k \in K} \|d^k\| = 0$, that is, recalling the definition of d^k given in (38)–(39),

$$\lim_{k \rightarrow \infty, k \in K} \left\| \tilde{x}^k - P(\tilde{x}^k - m^k \nabla \varphi(\tilde{x}^k))_{\mathcal{B}_{N_{\epsilon_1}^k}} \right\| = 0.$$

From the properties of the projection we have that

$$\left\| \tilde{x}^k - P(\tilde{x}^k - m^k \nabla \varphi(\tilde{x}^k))_{\mathcal{B}_{N_{\epsilon_1}^k}} \right\| \geq \left\| \tilde{x}^k - P(\tilde{x}^k - \underline{m} \nabla \varphi(\tilde{x}^k))_{\mathcal{B}_{N_{\epsilon_1}^k}} \right\|,$$

so that the following holds

$$\lim_{k \rightarrow \infty, k \in K} \left\| \tilde{x}^k - P(\tilde{x}^k - \underline{m} \nabla \varphi(\tilde{x}^k))_{\mathcal{B}_{N_{\epsilon_1}^k}} \right\| = 0.$$

Using (66), the continuity of the projection and taking into account (41) in Proposition 6, we obtain

$$\left\| x^* - P(x^* - \underline{m} \nabla \varphi(x^*))_{\mathcal{B}_{N_{\epsilon_1}^k}} \right\| = 0.$$

This contradicts (68), leading to the desired result. □

4 Numerical results

In this section, we show the practical performances of AS- ℓ_1 on two classes of problems frequently arising in data science and machine learning that can be formulated as problem (1):

- LASSO problems [32], where

$$\varphi(x) = \|Ax - b\|^2, \tag{69}$$

- for given matrix $A \in \mathbb{R}^{m \times n}$ and vector $b \in \mathbb{R}^m$,
- ℓ_1 -constrained logistic regression problems, where

$$\varphi(x) = \sum_{i=1}^l \log(1 + \exp(-y_i x^T a_i)), \tag{70}$$

with given vectors a_i and scalars $y_i \in \{1, -1\}$, $i = 1, \dots, l$.

In our implementation of AS- ℓ_1 , we used a non-monotone line search with memory length $n_m = 10$ (see Algorithm 2) and a spectral (or Barzilai–Borwein) gradient direction for the variables in $N_{\ell_1}^k$. In particular, the coefficient m^k appearing in (38) was set to 1 for $k = 0$ and, for $k \geq 1$, we employed the following formula, adapting the strategy used in [2, 4, 11]:

$$m^k = \begin{cases} \max\{\underline{m}, m_a^k\}, & \text{if } 0 < m_a^k < \bar{m}, \\ \max\{\underline{m}, \min\{\bar{m}, m_b^k\}\}, & \text{if } m_a^k \geq \bar{m}, \\ \max\left\{\underline{m}, \min\left\{1, \frac{\|\nabla_{N_{\ell_1}^k} \varphi(\tilde{x}^k)\|}{\|\tilde{x}_{N_{\ell_1}^k}^k\|}\right\}\right\}, & \text{if } m_a^k \leq 0, \end{cases}$$

where $\underline{m} = 10^{-10}$, $\bar{m} = 10^{10}$, $m_a^k = \frac{(s^{k-1})^T y^{k-1}}{\|s^{k-1}\|^2}$, $m_b^k = \frac{\|y^{k-1}\|^2}{(s^{k-1})^T y^{k-1}}$, $s^{k-1} = \tilde{x}_{N_{\ell_1}^k}^k - \tilde{x}_{N_{\ell_1}^k}^{k-1}$ and $y^{k-1} = \nabla_{N_{\ell_1}^k} \varphi(\tilde{x}^k) - \nabla_{N_{\ell_1}^k} \varphi(\tilde{x}^{k-1})$.

The ϵ parameter appearing in the active-set estimate (19) should satisfy Assumption 1 to guarantee the descent property established in Proposition 4 and the convergence of the algorithm. Since the Lipschitz constant L is in general unknown, we approximate ϵ following the same strategy as in [9, 10, 12], where similar estimates are used. Starting from $\epsilon = 10^{-6}$, we update its value along the iterations, reducing it whenever the expected decrease in the objective, stated in Proposition 4, is not obtained.

In our experiments, we implemented AS- ℓ_1 in Matlab and compared AS- ℓ_1 with the two following first-order methods, implemented in Matlab as well:

- A spectral projected gradient method with non-monotone line search, which will be referred to as NM-SPG, downloaded from Mark Schmidt’s webpage <https://www.cs.ubc.ca/~schmidtm/Software/minConf.html>;

- The away-step Frank–Wolfe method with Armijo line search [5, 6], which will be referred to as AFW¹.

Finally, we report a comparison between AS- ℓ_1 and AS-PG, an active-set algorithm devised in [10], showing the benefit of explicitly handling the ℓ_1 -ball.

For every considered problem, we set the starting point equal to the origin and we first run AS- ℓ_1 , stopping when

$$\|x^k - P(x^k - \nabla\varphi(x^k))_{\ell_1}\| \leq 10^{-6},$$

where $P(\cdot)_{\ell_1}$ denotes the projection onto the ℓ_1 -ball. Then, the other methods were run with the same starting point and were stopped at the first iteration k such that

$$\varphi(x^k) \leq f^* + 10^{-6}(1 + |f^*|),$$

with f^* being the objective value found by AS- ℓ_1 . A time limit of 3600 s was also included in all the considered methods.

In NM-SPG, we used the default parameters (except for those concerning the stopping condition). Moreover, in AS- ℓ_1 and NM-SPG we employed the same projection algorithm [8], downloaded from Laurent Condat's webpage <https://lcondat.github.io/software.html>.

In all codes, we made use of the Matlab *sparse* operator to compute $\varphi(x)$ and $\nabla\varphi(x)$, in order to exploit the problem structure and save computational time. The experiments were run on an Intel Xeon(R) CPU E5-1650 v2 @ 3.50GHz with 12 cores and 64 Gb RAM.

The AS- ℓ_1 software is available at <https://github.com/acristofari/as-l1>.

4.1 Comparison on LASSO instances

We considered 10 artificial instances of LASSO problems, where the objective function $\varphi(x)$ takes the form of (69). Each instance was created by first generating a matrix $A \in \mathbb{R}^{m \times n}$ with elements randomly drawn from a uniform distribution on the interval $(0, 1)$, using $n = 2^{15}$ and $m = n/2$. Then, a vector x^* was generated with all zeros, except for $\text{round}(0.05m)$ components, which were randomly set to 1 or -1 . Finally, we set $b = Ax^* + 0.001v$, where v is a vector with elements randomly drawn from the standard normal distribution, and the ℓ_1 -sphere radius τ was set to $0.99\|x^*\|_1$.

The detailed comparison on the LASSO instances is reported in Table 1. For each instance and each algorithm, we report the final objective function value found, the CPU time needed to satisfy the stopping criterion and the percentage of zeros in the final solution, with a tolerance of 10^{-5} . In case an algorithm reached the time limit

¹ AFW was run by reformulating (1) as an optimization problem over the unit simplex, exploiting the fact that the feasible set is a convex combination of the vectors $\pm\tau e_i$, $i = 1, \dots, n$.

Table 1 Comparison on 10 LASSO instances. For each method, the first column (Obj) indicates the final objective value, the second column (CPU time) indicates the required time in seconds, where a star means that the time limit of 3600 s was reached, and the third column (%zeros) indicates the percentage of zeros in the final solution, with a tolerance of 10^{-5} . For each problem, the fastest algorithm is highlighted in bold

AS- ℓ_1			NM-SPG			AFW		
Obj	CPU time	%zeros	Obj	CPU time	%zeros	Obj	CPU time	%zeros
54.20	315.85	97.49	54.20	*	97.49	54.20	2762.90	97.49
52.32	366.90	97.50	5823.87	*	80.69	52.32	3046.89	97.50
53.95	449.67	97.50	1040.99	*	85.12	53.95	3023.94	97.50
54.04	292.83	97.50	2215.88	*	83.45	54.04	3050.17	97.50
52.98	330.65	97.50	841.57	*	85.21	52.98	2798.97	97.50
53.54	387.79	97.50	53.56	*	97.50	53.54	3006.38	97.50
52.71	806.80	97.50	3927.10	*	82.23	52.71	2837.90	97.50
53.58	580.89	97.50	4108.25	*	81.93	53.58	2768.45	97.50
52.61	402.03	97.50	1750.54	*	83.37	52.61	2924.38	97.50
53.36	535.10	97.50	53.89	*	97.49	53.36	2948.41	97.50

on an instance, we consider as final solution and final objective value those related to the last iteration performed. NM-SPG reached the time limit on all instances, being very far from f^* on 6 instances out of 10, with a difference of even two order of magnitude. AFW gets the same solutions as those obtained by AS- ℓ_1 , being however an order of magnitude slower than AS- ℓ_1 .

The same picture is given by Fig. 1, where we report the average optimization error $f(x^k) - f_{\text{best}}$ over the 10 instances, with f_{best} being the minimum objective value found by the algorithms. We can notice that AS- ℓ_1 clearly outperforms the other two methods.




Fig. 1 Average optimization error over LASSO instances (y axis) vs CPU time in seconds (x axis). The y axis is in logarithmic scale

4.2 Comparison on logistic regression instances

For the comparison among AS- ℓ_1 , NM-SPG and AFW on ℓ_1 -constrained logistic regression problems, where the objective function $\varphi(x)$ takes the form of (70), we considered 11 datasets for binary classification from the literature, with a number of samples l between 100 and 25,000, and a number of attributes n between 500 and 100,000. We report the complete list of datasets in Table 2.

For each dataset, we considered different values of the ℓ_1 -sphere radius τ , that is, $0.01n$, $0.03n$ and $0.05n$. The final results are shown in Table 3. As before, for each instance and each algorithm, we report the final objective function value found, the CPU time needed to satisfy the stopping criterion and the percentage of zeros in the final solution, with a tolerance of 10^{-5} . In case an algorithm reached the time limit on an instance, we consider as final solution and final objective value those related to the last iteration performed. Excluding the instance obtained from the Rev1_train.binary dataset with $\tau = 0.05n$, the three solvers get very similar solutions on all instances, with a difference of 0.02 at most in the final objective values. When considering $\tau = 0.01n$, AS- ℓ_1 is the fastest solver on 4 instances out of 11. Note that on the instance from the Farm-ads-vect dataset, AS- ℓ_1 is able to get the solution in a third of the CPU time needed by the other two solvers. On the other instances, the CPU time needed by AS- ℓ_1 is always comparable with the one needed by the fastest solver. Looking at the results for larger values of τ , we can notice that the instances get more difficult and in general less sparse. For $\tau = 0.03n$ and $\tau = 0.05n$, AS- ℓ_1 is the fastest solver on all the instances but two, those obtained from the Arcene and the Dorothea datasets, which are however addressed within 2 s. On other instances, such as those built from the Real-sim and the Rev1_train.binary datasets, AS- ℓ_1 is one or even two orders of magnitude faster with respect to NM-SPG and AFW.

In Fig. 2 we report the average optimization error $f(x^k) - f_{\text{best}}$ over the 11 instances, for each value of τ , with f_{best} being the minimum objective value found by the algorithms. We can notice that AFW is outperformed by the other two algorithms, which have similar performance when considering the average optimization

Table 2 Datasets used in the comparison on ℓ_1 -constrained logistic regression problems, where l is the number of instances and n is the number of attributes

Dataset	l	n	Reference
Arcene (training set)	100	10,000	[14, 21]
Dexter (training set)	300	19,999	[14, 21]
Dorothea (training set)	800	100,000	[14, 21]
Farm-ads-vect	4143	54,877	[14]
Gisette (training set)	6000	5000	[1, 14, 21]
Madelon (training set)	2000	500	[1, 14, 21]
Rcv1_train.binary (training set)	20,242	47,236	[1, 27]
Real-sim	72,309	20,958	[1]
Swarm (Aligned)	24,016	2400	[14]
Swarm (Flocking)	24,016	2400	[14]
Swarm (Grouped)	24,016	2400	[14]

Table 3 Comparison on ℓ_1 -constrained logistic regression problems with different values of the sphere radius τ

τ	Dataset	$\mathcal{AS} - \ell_1$			NM-SPG			AFW		
		Obj	CPU time	%zeros	Obj	CPU time	%zeros	Obj	CPU time	%zeros
0.01n	Arcene	0.20	8.31	99.31	0.20	22.34	99.30	0.20	12.14	99.32
	Dexter	3.88	1.09	99.41	3.88	1.30	99.41	3.88	11.81	99.41
	Dorothea	0.00	0.75	70.09	0.00	0.58	51.79	0.00	5.58	99.82
	Farm-ads-vect	146.38	576.62	98.39	146.38	1547.40	98.39	146.38	1699.47	98.41
	Gisette	0.00	9.70	8.20	0.00	18.42	8.22	0.00	354.81	78.60
	Madelon	1314.89	0.60 ara>	95.01	1314.89	0.41	95.01	1314.89	0.50	94.81
	Rcv1_train.binary	2344.19	11.15	98.86	2344.19	7.77	98.85	2344.19	218.41	98.87
	Real-sim	21693.00	4.33	99.47	21693.03	4.26	99.47	21693.03	104.79	99.47
	Swarm (Aligned)	2504.35	51.52	97.96	2504.35	34.05	97.96	2504.35	50.48	97.96
	Swarm (Flocking)	4823.39	56.99	97.79	4823.40	27.48	97.75	4823.40	29.68	97.79
	Swarm (Grouped)	5062.87	47.92	97.67	5062.88	43.53	97.67	5062.88	32.22	97.67
	0.03n	Arcene	0.00	1.91	81.27	0.00	1.80	83.80	0.00	1.26
Dexter		0.00	0.66	99.06	0.00	0.92	99.06	0.00	7.01	99.16
Dorothea		0.00	0.57	11.97	0.00	0.35	11.90	0.00	5.18	99.84
Farm-ads-vect		12.68	676.07	98.08	12.68	2191.57	98.08	12.69	*	98.10
Gisette		0.00	8.91	8.26	0.00	7.15	7.74	0.00	375.52	77.86
Madelon		1262.91	1.63	82.44	1262.91	2.38	82.24	1262.91	2.04	82.44
Rcv1_train.binary		965.13	78.43	97.03	965.13	620.44	97.00	965.13	3206.36	97.15
Real-sim		13174.84	9.40	98.01	13174.85	96.91	98.01	13174.85	785.63	98.04
Swarm (Aligned)		115.10	147.98	96.46	115.10	199.90	96.46	115.10	277.14	96.46
Swarm (Flocking)		1027.82	187.41	94.13	1027.82	130.88	94.13	1027.82	262.78	94.09
Swarm (Grouped)		966.16	199.21	94.29	966.16	475.48	94.25	966.16	281.29	94.29

Table 3 (continued)

τ	Dataset	AS- ℓ_1			NM-SPG			AFW		
		Obj	CPU time	%zeros	Obj	CPU time	%zeros	Obj	CPU time	%zeros
0.05n	Arcene	0.00	0.90	13.09	0.00	0.63	24.94	0.00	1.23	98.71
	Dexter	0.00	0.19	89.91	0.00	0.15	90.81	0.00	2.08	99.11
	Dorothea	0.00	0.59	11.90	0.00	0.32	11.90	0.00	4.82	99.86
	Farm-ads-vect	9.38	*	97.77	9.38	*	97.69	9.40	*	97.86
	Gisette	0.00	8.18	8.24	0.00	8.01	7.94	0.00	375.69	78.42
	Madelon	1225.19	2.25	72.26	1225.19	10.35	71.86	1225.19	3.61	72.26
	Rcv1_train.binary	470.53	275.16	96.06	470.54	2336.48	96.04	473.33	*	96.57
	Real-sim	10087.92	27.83	96.37	10087.93	161.01	96.36	10087.93	1771.75	96.43
	Swarm (Aligned)	6.80	296.74	95.34	6.80	429.97	95.34	6.80	510.29	95.34
	Swarm (Flocking)	237.26	285.46	92.54	237.26	413.54	92.54	237.26	532.85	92.54
	Swarm (Grouped)	210.80	260.47	92.71	210.80	937.63	92.71	210.80	663.67	92.71

For each method, the first column (Obj) indicates the final objective value, the second column (CPU time) indicates the required time in seconds, where a star means that the time limit of 3600 s was reached, and the third column (%zeros) indicates the percentage of zeros in the final solution, with a tolerance of 10^{-5} . For each problem, the fastest algorithm is highlighted in bold




Fig. 2 Average optimization error over ℓ_1 -constrained logistic regression instances (y axis) vs CPU time in seconds (x axis). The y axis is in logarithmic scale

error above 10^{-2} . When considering the average optimization error below 10^{-2} , we see that AS- ℓ_1 outperforms NM-SPG too.

4.3 Comparison with AS-PG

We now compare AS- ℓ_1 with AS-PG, an active set algorithm that uses a projected gradient direction, presented in [10]. As for AFW, AS-PG was run by reformulating (1) as an optimization problem over the unit simplex. In Figs. 3 and 4, we report the comparison over LASSO and logistic regression instances, respectively. The considered LASSO instances are obtained as before with $n = 2^{12}$ and $m = n/2$. As we can easily see by taking a look at the plots, the proposed approach guarantees better performances.




Fig. 3 Average optimization error over LASSO instances (y axis) vs CPU time in seconds (x axis). The y axis is in logarithmic scale




Fig. 4 Average optimization error over ℓ_1 -constrained logistic regression instances (y axis) vs CPU time in seconds (x axis). The y axis is in logarithmic scale

5 Conclusions

In this paper, we focused on minimization problems over the ℓ_1 -ball and described a tailored active-set algorithm. We developed a strategy to guess, along the iterations of the algorithm, which variables should be zero at a solution. A reduction in terms of objective function value is guaranteed by simply fixing to zero those variables estimated to be active. The active-set estimate is used in combination with a projected spectral gradient direction and a non-monotone Armijo line search. We analyzed in depth the global convergence of the proposed algorithm. The numerical results show the efficiency of the method on LASSO and sparse logistic regression instances, in comparison with two widely-used first-order methods.

Funding Open access funding provided by Università degli Studi di Roma La Sapienza within the CRUI-CARE Agreement.

Data availability The data analysed during the current study are available from the corresponding author on reasonable request.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. LIBSVM data: classification, regression, and multi-label. <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>
2. Andreani, R., Birgin, E.G., Martínez, J.M., Schuverdt, M.L.: Second-order negative-curvature methods for box-constrained and general constrained optimization. *Comput. Optim. Appl.* **45**(2), 209–236 (2010)
3. Bertsekas, D.P.: Projected Newton methods for optimization problems with simple constraints. *SIAM J. Control. Optim.* **20**(2), 221–246 (1982)

4. Birgin, E.G., Martínez, J.M.: Large-scale active-set box-constrained optimization method with spectral projected gradients. *Comput. Optim. Appl.* **23**(1), 101–125 (2002)
5. Bomze, I.M., Rinaldi, F., Bulò, S.R.: First-order methods for the impatient: support identification in finite time with convergent Frank–Wolfe variants. *SIAM J. Optim.* **29**(3), 2211–2226 (2019)
6. Bomze, I.M., Rinaldi, F., Zeffiro, D.: Active set complexity of the Away-step Frank–Wolfe Algorithm. arXiv preprint [arXiv:1912.11492](https://arxiv.org/abs/1912.11492) (2019)
7. Brás, C.P., Fischer, A., Júdice, J.J., Schönfeld, K., Seifert, S.: A block active set algorithm with spectral choice line search for the symmetric eigenvalue complementarity problem. *Appl. Math. Comput.* **294**, 36–48 (2017)
8. Condat, L.: Fast projection onto the simplex and the ℓ_1 ball. *Math. Program.* **158**(1), 575–585 (2016)
9. Cristofari, A., De Santis, M., Lucidi, S., Rinaldi, F.: A two-stage active-set algorithm for bound-constrained optimization. *J. Optim. Theory Appl.* **172**(2), 369–401 (2017)
10. Cristofari, A., De Santis, M., Lucidi, S., Rinaldi, F.: An active-set algorithmic framework for non-convex optimization problems over the simplex. *Comput. Optim. Appl.* **77**(1), 57–89 (2020)
11. Cristofari, A., Rinaldi, F., Tudisco, F.: Total variation based community detection using a nonlinear optimization approach. *SIAM J. Appl. Math.* **80**(3), 1392–1419 (2020)
12. De Santis, M., Lucidi, S., Rinaldi, F.: A fast active set block coordinate descent algorithm for ℓ_1 -regularized least squares. *SIAM J. Optim.* **26**(1), 781–809 (2016)
13. Di Serafino, D., Toraldo, G., Viola, M., Barlow, J.: A two-phase gradient method for quadratic programming problems with a single linear constraint and bounds on the variables. *SIAM J. Optim.* **28**(4), 2809–2838 (2018)
14. Dua, D., Graff, C.: UCI machine learning repository (2017). <http://archive.ics.uci.edu/ml>
15. Duchi, J., Gould, S., Koller, D.: Projected subgradient methods for learning sparse gaussians. arXiv preprint [arXiv:1206.3249](https://arxiv.org/abs/1206.3249) (2012)
16. Duchi, J., Shalev-Shwartz, S., Singer, Y., Chandra, T.: Efficient projections onto the l_1 -ball for learning in high dimensions. In: *Proceedings of the 25th International Conference on Machine Learning*, pp. 272–279 (2008)
17. Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.: Least angle regression. *Ann. Stat.* **32**(2), 407–499 (2004)
18. Facchinei, F., Fischer, A., Kanzow, C.: On the accurate identification of active constraints. *SIAM J. Optim.* **9**(1), 14–32 (1998)
19. Facchinei, F., Júdice, J., Soares, J.: An active set Newton algorithm for large-scale nonlinear programs with box constraints. *SIAM J. Optim.* **8**(1), 158–186 (1998)
20. Grippo, L., Lampariello, F., Lucidi, S.: A nonmonotone line search technique for Newton’s method. *SIAM J. Numer. Anal.* **23**(4), 707–716 (1986)
21. Guyon, I., Gunn, S.R., Ben-Hur, A., Dror, G.: Result analysis of the NIPS 2003 feature selection challenge. In: *NIPS*, vol. 4, pp. 545–552 (2004)
22. Hager, W.W., Zhang, H.: A new active set algorithm for box constrained optimization. *SIAM J. Optim.* **17**(2), 526–557 (2006)
23. Hager, W.W., Zhang, H.: An active set algorithm for nonlinear optimization with polyhedral constraints. *Sci. China Math.* **59**(8), 1525–1542 (2016)
24. Hager, W.W., Zhang, H.: Projection onto a polyhedron that exploits sparsity. *SIAM J. Optim.* **26**(3), 1773–1798 (2016)
25. Jaggi, M.: Revisiting Frank–Wolfe: projection-free sparse convex optimization. In: *International Conference on Machine Learning*, pp. 427–435. PMLR (2013)
26. Lacoste-Julien, S., Jaggi, M.: On the global linear convergence of Frank–Wolfe optimization variants. In: *NIPS 2015—Advances in Neural Information Processing Systems* (2015)
27. Lewis, D.D., Yang, Y., Russell-Rose, T., Li, F.: Rcv1: a new benchmark collection for text categorization research. *J. Mach. Learn. Res.* **5**, 361–397 (2004)
28. Moré, J.J., Toraldo, G.: Algorithms for bound constrained quadratic programming problems. *Numer. Math.* **55**(4), 377–400 (1989)
29. Nesterov, Y.: *Introductory lectures on convex optimization: a basic course*, vol. 87. Springer, Berlin (2013)
30. Schmidt, M., Berg, E., Friedlander, M., Murphy, K.: Optimizing costly functions with simple constraints: a limited-memory projected quasi-newton algorithm. In: *Artificial Intelligence and Statistics*, pp. 456–463. PMLR (2009)

31. Schmidt, M., Murphy, K., Fung, G., Rosales, R.: Structure learning in random fields for heart motion abnormality detection. In: 2008 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8. IEEE (2008)
32. Tibshirani, R.: Regression shrinkage and selection via the lasso. *J. R. Stat. Soc.: Ser. B (Methodol.)* **58**(1), 267–288 (1996)

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