



Projected Gradient Methods with Momentum

Matteo Lapucci¹ · Giampaolo Liuzzi² · Stefano Lucidi² ·
Marco Sciandrone² · Diego Scuppa²

Received: 23 January 2026 / Accepted: 19 June 2026 / Published online: 3 July 2026
© The Author(s) 2026

Abstract

We focus on the optimization problem with smooth, possibly nonconvex objectives and a convex constraint set for which the Euclidean projection operation is practically available. Focusing on this setting, we carry out a general convergence and complexity analysis for algorithmic frameworks. Consequently, we discuss theoretically sound strategies to integrate momentum information within classical projected gradient-type algorithms. One of these approaches is then developed in detail, up to the definition of a tailored algorithm with both theoretical guarantees and reasonable per-iteration cost. The proposed method is finally shown to outperform the standard (spectral) projected gradient method in two different experimental benchmarks, indicating that the addition of momentum terms is as beneficial in the constrained setting as it is in the unconstrained scenario.

Keywords Nonlinear optimization · Convex constraints · Projected gradient method · Momentum · Complexity

Mathematics Subject Classification 90C29 · 90C26 · 90C60

1 Introduction

In this work, we consider the constrained nonlinear optimization problem

$$\min f(x) \quad \text{s.t.} \quad x \in \mathcal{S}, \quad (1)$$

Communicated by Shiqian Ma.

✉ Diego Scuppa
diego.scuppa@uniroma1.it

¹ Dipartimento di Ingegneria dell'Informazione, Università di Firenze, Via di Santa Marta 3, 50139 Firenze, Italy

² Dipartimento di Ingegneria Informatica, Automatica e Gestionale, Sapienza Università di Roma, Via Ariosto 25, 00185 Roma, Italy

where $\mathcal{S} \subset \mathbb{R}^n$ is a closed, convex set and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a smooth, possibly nonconvex function.

A rather classical approach (analyzed in the 60s already [27]) to deal with problems of this form is by means of iterative descent schemes relying on the Euclidean projection operator $\mathcal{P}_{\mathcal{S}} : \mathbb{R}^n \rightarrow \mathcal{S}$, provided the latter is available at a low computational cost for the feasible set $\mathcal{S}$. In this scenario, two conceptually different types of *projected gradient (PG) methods* are, in fact, known from the literature, both collapsing to the standard gradient descent method if $\mathcal{S} = \mathbb{R}^n$. The first one [4, Sec. 2.3] is defined by the following update rule

$$x_{k+1} = \mathcal{P}_{\mathcal{S}}[x_k - \alpha_k \nabla f(x_k)] \quad (2)$$

where the stepsize α_k is determined by a curvilinear line search along the projection arc. The second PG approach is, instead, defined by the following iteration [21, Sec. 20.4]

$$x_{k+1} = x_k + \alpha_k \hat{d}_k, \quad (3)$$

where $\hat{d}_k$ is the *projected gradient direction*, i.e.,

$$\hat{d}_k = \mathcal{P}_{\mathcal{S}}[x_k - \eta_k \nabla f(x_k)] - x_k, \quad (4)$$

$\eta_k > 0$ is an algorithmic parameter, and the stepsize α_k is determined by an Armijo-type line search.

Taking inspiration from the Barzilai-Borwein spectral gradient method [3, 30], the selection of parameter η_k using the inverse Rayleigh quotient was proposed in [6] to be used in combination with a nonmonotone line search for the selection of α_k [20]. In the vast literature on this family of techniques, other relevant developments of PG methods deal with inexact projections [8, 16–18] and the use of gradient scaling [1, 9, 10].

Here, we are mainly interested in studying acceleration techniques for projected gradient algorithms using a *momentum*-based strategy, i.e., through the addition of a term that includes information from the past iteration within the update rule. In the unconstrained case, i.e., when $\mathcal{S} = \mathbb{R}^n$, gradient methods with momentum are defined according to iterations of the form

$$x_{k+1} = x_k - \alpha_k \nabla f(x_k) + \beta_k (x_k - x_{k-1}), \quad (5)$$

where $\alpha_k > 0$ is the stepsize, and $\beta_k > 0$ is the momentum weight. Relevant algorithms in this class are the famous Polyak's heavy-ball method [29] and the broad family of conjugate gradient methods [21, Ch. 12]. While the latter is widely considered among the best performing options for nonlinear unconstrained problems, the former has particularly received renewed interest in the context of stochastic gradient algorithms for finite-sum problems in machine learning [11, 33]. In both cases, deterministic and stochastic, the momentum term helps to prevent oscillations and yields faster convergence.

Very recently, some gradient methods with momentum have been proposed [24, 25, 31] sharing the innovative strategy of simultaneously determining α_k and β_k through a

two-dimensional search, instead of using predefined values chosen according to some rule. These methods share strong theoretical convergence guarantees and remarkable computational efficiency. Despite the flourishing activity of the literature on momentum methods for unconstrained optimization, extensions to the constrained case have barely been studied.

The natural extension of the update rule (5) to the constrained case, in a projection-based fashion, is to modify (2) as follows

$$x_{k+1} = \mathcal{P}_S[x_k - \alpha_k \nabla f(x_k) + \beta_k(x_k - x_{k-1})].$$

In [32], the above scheme was, in fact, proposed for tackling problem (1), assuming that the objective function f is possibly nonsmooth but convex, and that the feasible set S is bounded. For suitable, predefined sequences of α_k and β_k , the proposed method attains an optimal $\mathcal{O}(1/k)$ complexity bound.

In this work, we follow a completely different and new path for defining a projection-based extension to the constrained setting of momentum methods. In particular, we take inspiration from the line search-based projected gradient method (3), defining a feasible search direction

$$d_k = \alpha_k \hat{d}_k + \beta_k (\mathcal{P}_S[x_k + (x_k - x_{k-1})] - x_k),$$

where $\hat{d}_k$ comes from (4); if $S = \mathbb{R}^n$, the update collapses back to (5); here, the scalars α_k and β_k need not be fixed in advance but are rather determined by solving – in closed form – a constrained two-dimensional quadratic problem. Then, the iteration of the proposed method takes the form

$$x_{k+1} = x_k + \mu_k d_k, \quad (6)$$

where $\mu_k > 0$ is set by an Armijo-type line search; the backtracking procedure can be started with the unit stepsize which, by the construction of d_k , is optimal according to some objective function model in a subspace and is thus a computationally convenient guess. For the proposed class of algorithms, we prove asymptotic convergence results and a worst-case complexity bound of $\mathcal{O}(\varepsilon^{-2})$ in the nonconvex setting. In fact, these properties are proven for a more general class of constrained optimization algorithms and, to the best of our knowledge, are novel in the literature.

The results of preliminary computational experiments highlight the benefits of adding the momentum term compared to using the standard projected gradient direction; in particular, we show through numerical experiments on both bound-constrained and ball-constrained problems that a significant speed-up can be attained for the state-of-the-art spectral projected gradient solver from [6] if the proposed technique to introduce the momentum term is adopted.

The rest of the paper is organized as follows: we begin with the formalization of (approximate) stationarity conditions for problem (1) in Sect. 2; we then identify a general line search-based algorithmic framework in Sect. 3, for which we state and prove conditions guaranteeing global convergence and complexity bounds. The specific case of projected gradient methods with momentum is addressed in Sect. 4,

where the soundness of direction choice is proved considering both n -dimensional and two-dimensional subproblems. In Sect. 5, we present in detail the proposed method, which is then numerically compared to the standard spectral projected gradient in Sect. 6. Final remarks are given in Sect. 7.

2 Stationarity Conditions

Before turning to the introduction of algorithmic frameworks to deal with problem (1), we need to state some basic assumptions and introduce the stationarity concepts we will be using to characterize solutions of the problem.

We consider problem (1), stating the following set of assumptions on the objective function f .

The rendering is not good. Start the itemize at a new line, as we have done (e.g., with

Assumption 2.1 - f is bounded below over $\mathcal{S}$, i.e., there exists $f^* \in \mathbb{R}$ such that $f(x) \geq f^*$ for all $x \in \mathcal{S}$;
- f is L -smooth over an open convex set $\mathcal{D}$ containing $\mathcal{S}$, i.e.,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \quad \forall x, y \in \mathcal{D} \supseteq \mathcal{S}.$$

Now, the standard optimality condition for problem (1) can be given according to the following definition.

Definition 2.1 A point $\bar{x} \in \mathcal{S}$ is a stationary point for problem (1) if

$$\nabla f(\bar{x})^T (x - \bar{x}) \geq 0 \quad \forall x \in \mathcal{S}. \quad (7)$$

At this point, we shall also formally recall the Euclidean projector operator over $\mathcal{S}$, which will play a crucial role throughout the paper, defined as

$$\mathcal{P}_{\mathcal{S}}[x] = \arg \min_{y \in \mathcal{S}} \|x - y\|,$$

and the following well-known result that allows us to deal with stationarity in a practically more convenient way.

Proposition 2.1 Given $x \in \mathcal{S}$ and $\eta > 0$, let $\hat{x} = \mathcal{P}_{\mathcal{S}}[x - \eta \nabla f(x)]$ and $\hat{d} = \hat{x} - x$. Then:

i) the following inequality holds

$$\nabla f(x)^T \hat{d} \leq -\frac{1}{\eta} \|\hat{d}\|^2 = -\frac{1}{\eta} \|\mathcal{P}_{\mathcal{S}}[x - \eta \nabla f(x)] - x\|^2;$$

ii) point x is stationary for (1) if and only if $\|\hat{d}\| = \|\mathcal{P}_{\mathcal{S}}[x - \eta \nabla f(x)] - x\| = 0$.

Proof From [21, Prop. 20.2(i)], for every $y \in \mathbb{R}^n$ and $z \in \mathcal{S}$,

$$(y - \mathcal{P}_{\mathcal{S}}[y])^T (z - \mathcal{P}_{\mathcal{S}}[y]) \leq 0. \quad (8)$$

Then, for every $x \in \mathcal{S}$ and $\eta > 0$, taking $y = x - \eta \nabla f(x)$ and $z = x$ in (8),

$$((x - \eta \nabla f(x)) - \mathcal{P}_{\mathcal{S}}[x - \eta \nabla f(x)])^T (x - \mathcal{P}_{\mathcal{S}}[x - \eta \nabla f(x)]) \leq 0.$$

Then, point i) follows. Point ii) immediately follows from [21, Prop. 20.3]. $\square$

The main takeaway from point ii) of the above proposition is that we are allowed to select any arbitrary value of $\eta > 0$ to define the (continuous) function $\phi_{\eta} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ as

$$\phi_{\eta}(x) = \|\mathcal{P}_{\mathcal{S}}[x - \eta \nabla f(x)] - x\|$$

and use $\phi_{\eta}(\bar{x}) = 0$ as a stationarity condition for a point $\bar{x} \in \mathcal{S}$. Given $x_k \in \mathcal{S}$, we are also allowed to call x_k ε -stationary for problem (1) if

$$\phi_{\eta}(x_k) = \|\mathcal{P}_{\mathcal{S}}[x_k - \eta \nabla f(x_k)] - x_k\| \leq \varepsilon.$$

However, the ε -stationarity property is not independent of the choice of η , as we can deduce from the following lemma.

Lemma 2.1 Let $x \in \mathcal{S}$, $z \in \mathbb{R}^n$. Define functions $p, h : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{\geq 0}$ as

$$p(t) = \|\mathcal{P}_{\mathcal{S}}[x + tz] - x\| \quad \text{and} \quad h(t) = \frac{1}{t} \|\mathcal{P}_{\mathcal{S}}[x + tz] - x\| = \frac{p(t)}{t}.$$

Then, (a) $p(t)$ is monotonically nondecreasing and (b) $h(t)$ is monotonically nonincreasing.

Proof For part (b): see [4, Lemma 2.3.1]. For part (a), let $0 < t_1 < t_2$. We aim to show that $p(t_1) \leq p(t_2)$, i.e., $\|\mathcal{P}_{\mathcal{S}}[x + t_1 z] - x\|^2 \leq \|\mathcal{P}_{\mathcal{S}}[x + t_2 z] - x\|^2$. First, we prove that

$$z^T (\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z]) \geq 0. \quad (9)$$

Indeed, by [21, Prop. 20.2], we know that:

$$((x + t_1 z) - \mathcal{P}_{\mathcal{S}}[x + t_1 z])^T (\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z]) \leq 0 \quad (10)$$

and

$$(\mathcal{P}_{\mathcal{S}}[x + t_2 z] - (x + t_2 z))^T (\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z]) \leq 0. \quad (11)$$

Summing up (10) and (11), we obtain:

$$\begin{aligned} 0 &\geq (\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z] + (t_1 - t_2)z)^T (\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z]) \\ &= \|\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z]\|^2 + (t_1 - t_2)z^T (\mathcal{P}_{\mathcal{S}}[x + t_2 z] - \mathcal{P}_{\mathcal{S}}[x + t_1 z]) \end{aligned}$$

$$\geq (t_1 - t_2)z^T (\mathcal{P}_S[x + t_2z] - \mathcal{P}_S[x + t_1z]).$$

Since $t_1 - t_2 < 0$, (9) must be satisfied. Hence,

$$\begin{aligned} & \|\mathcal{P}_S[x + t_2z] - x\|^2 \\ &= \|(\mathcal{P}_S[x + t_1z] - x) + (\mathcal{P}_S[x + t_2z] - \mathcal{P}_S[x + t_1z])\|^2 \\ &\geq \|\mathcal{P}_S[x + t_1z] - x\|^2 + 2(\mathcal{P}_S[x + t_1z] - x)^T (\mathcal{P}_S[x + t_2z] - \mathcal{P}_S[x + t_1z]) \\ &= \|\mathcal{P}_S[x + t_1z] - x\|^2 + 2t_1z^T (\mathcal{P}_S[x + t_2z] - \mathcal{P}_S[x + t_1z]) \\ &\quad - 2((x + t_1z) - \mathcal{P}_S[x + t_1z])^T (\mathcal{P}_S[x + t_2z] - \mathcal{P}_S[x + t_1z]) \\ &\geq \|\mathcal{P}_S[x + t_1z] - x\|^2, \end{aligned}$$

where the last inequality follows from (9) and (10). The lemma is proved. $\square$

From point (a) of Lemma 2.1, we realize that given x_k , the value of $\phi_\eta(x_k)$ is nondecreasing as the value of η increases; in other words, by changing the values of η , we can retrieve a family of stationarity measures that become milder as the value of η gets closer to 0. As a consequence, if a point x_k satisfies the ε -stationarity condition given some value of $\bar{\eta} > 0$, then it will also be ε -stationary for any choice of ϕ_η with $\eta \in (0, \bar{\eta}]$.

Interestingly, we would obtain an analogous family of stationarity measures, this time with increasing strictness as η gets smaller, if we considered the function $\frac{1}{\eta} \|\mathcal{P}_S[x_k - \eta \nabla f(x_k)] - x_k\|$ instead of ϕ_η . However, in the following, we will only be referring to function ϕ_η to treat stationarity.

3 A General Line Search-Based Algorithmic Framework for Constrained Nonconvex Optimization and its Complexity Bounds

In this section, we define a general line search-based algorithmic scheme for constrained optimization, and we state general complexity and convergence results that, to the best of our knowledge, are novel in the literature. These results will be later employed for the specific theoretical analysis of the projected gradient algorithms with momentum proposed in this work.

We begin by stating a second set of assumptions, this time concerning the search directions to be employed in iterative line search-based methods.

Assumption 3.1 Let $\{x_k\} \subseteq \mathcal{S}$ and $\{d_k\} \subseteq \mathbb{R}^n$ be the sequences of points and corresponding search directions produced by an iterative algorithm and let $\eta > 0$. Then, there exist constants $a, c_1, c_2 > 0$ such that, for all k , the following conditions hold:

$$x_k + \alpha d_k \in \mathcal{S} \quad \forall \alpha \in [0, a], \quad (12)$$

$$\nabla f(x_k)^T d_k \leq -c_1 \|d_k\|^2, \quad (13)$$

$$\nabla f(x_k)^T d_k \leq -c_2 \|\mathcal{P}_S[x_k - \eta \nabla f(x_k)] - x_k\|^2 = -c_2 \phi_\eta(x_k)^2. \quad (14)$$

The two latter conditions in Assumption 3.1 somehow extend the classical gradient-related condition in unconstrained optimization [12] to the constrained setting; the former deals instead with the feasibility of directions; in particular, it is just slightly stricter than d_k being feasible: by the convexity of $\mathcal{S}$, the requirement in fact simply collapses to $x_k + \alpha d_k \in \mathcal{S}$ for all k .

Remark 3.1 For any $\eta > 0$, Assumption 3.1 always holds with $a = 1$ and $c_1 = c_2 = 1/\eta$ if d_k is always chosen to be equal to $\hat{d}_k = \mathcal{P}_{\mathcal{S}}[x_k - \eta \nabla f(x_k)] - x_k = \hat{x}_k - x_k$. The first condition, in fact, follows from $x_k + \hat{d}_k = \hat{x}_k \in \mathcal{S}$ by the definition of the projection operator. The two latter conditions instead collapse to the same inequality, which holds for $c_1 = c_2 = \frac{1}{\eta}$ by Proposition 2.1. As a consequence of this observation, the following complexity results will be applicable to the classical projected gradient method.

We are now able to state a useful preliminary result.

Proposition 3.1 *Let $x_k \in \mathcal{S}$ and let $d_k \in \mathbb{R}^n$ be a pair of point-direction produced under Assumption 3.1. Then, for any $\gamma \in (0, 1)$, we have that*

$$f(x_k + \alpha d_k) \leq f(x_k) + \alpha \gamma \nabla f(x_k)^T d_k$$

for all α such that $0 < \alpha \leq \alpha_L = \min \left[a, \frac{(1-\gamma)2c_1}{L} \right]$.

Proof Let $\alpha \in (0, \alpha_L]$. By the Lipschitz-continuity of ∇f and recalling the first condition of Assumption 3.1, noting $\alpha \leq a$, we can write $x_k + \alpha d_k \in \mathcal{S}$. We can also employ the descent lemma and (13) from Assumption 3.1 to write

$$\begin{aligned} f(x_k + \alpha d_k) - f(x_k) &\leq \alpha \nabla f(x_k)^T d_k + \alpha^2 \frac{L}{2} \|d_k\|^2 \\ &= \alpha \gamma \nabla f(x_k)^T d_k + (1 - \gamma) \alpha \nabla f(x_k)^T d_k + \alpha^2 \frac{L}{2} \|d_k\|^2 \\ &\leq \alpha \gamma \nabla f(x_k)^T d_k + (1 - \gamma) \alpha \nabla f(x_k)^T d_k - \alpha^2 \frac{L}{2c_1} \nabla f(x_k)^T d_k \\ &= \alpha \gamma \nabla f(x_k)^T d_k + \alpha \left((1 - \gamma) - \alpha \frac{L}{2c_1} \right) \nabla f(x_k)^T d_k \\ &\leq \alpha \gamma \nabla f(x_k)^T d_k, \end{aligned}$$

where the last inequality follows noting that, by definition, $\alpha \leq \frac{(1-\gamma)2c_1}{L}$. $\square$

Next, we define the general constrained optimization algorithm based on the standard Armijo-type line search. The method is formalized by the pseudocode reported in Algorithm 1.

The next proposition describes the theoretical properties of the presented line search-based algorithm.

Proposition 3.2 *Let $\eta > 0$ and let $\{x_k\}$, $\{d_k\}$ be the sequences produced by Algorithm 1, with $\{d_k\}$ satisfying Assumption 3.1. For any $\varepsilon > 0$, we denote by:*

Algorithm 1 Armijo line search-based algorithm

```

1: Input:  $x_0 \in \mathcal{S}, a > 0, \gamma \in (0, 1), \delta \in (0, 1), \eta > 0$ .
2: for  $k = 0, 1, \dots$  do
3:   /* Compute the search direction */
4:   Compute  $d_k$  satisfying the conditions of Assumption 3.1
5:   /* Perform Armijo line search along  $d_k$  */
6:   Set  $\alpha = a$ 
7:   while  $f(x_k + \alpha d_k) > f(x_k) + \gamma \alpha \nabla f(x_k)^T d_k$  do
8:     Set  $\alpha = \delta \alpha$ 
9:   Set  $\alpha_k = \alpha$ ,
10:  Set  $x_{k+1} = x_k + \alpha_k d_k$ 

```

- k_ε : the first iteration where the point x_{k_ε} satisfies $\phi_\eta(x_{k_\varepsilon}) \leq \varepsilon$;
- ni_ε : the total number of iterations k where the points x_k satisfy $\phi_\eta(x_k) > \varepsilon$.

Then, in the worst case, we have for Algorithm 1

$$k_\varepsilon \leq \frac{(f(x_0) - f^*)}{\delta \alpha_L \gamma c_2} \varepsilon^{-2} = \mathcal{O}(\varepsilon^{-2}) \quad \text{and} \quad ni_\varepsilon \leq \frac{(f(x_0) - f^*)}{\delta \alpha_L \gamma c_2} \varepsilon^{-2} = \mathcal{O}(\varepsilon^{-2}),$$

where f^* is the lower bound of f over $\mathcal{S}$ and $\alpha_L = \min \left[a, \frac{(1-\gamma)2c_1}{L} \right]$.

Proof The instructions of the Armijo-type line search and Proposition 3.1 imply that, for every iteration k , the stepsize α_k satisfies $\alpha_k \geq \delta \alpha_L$.

From the acceptance condition of the Armijo-type line search and the third condition of Assumption 3.1, for all $k \geq 0$, we then get

$$\begin{aligned} f(x_k) - f(x_{k+1}) &\geq \alpha_k \gamma |\nabla f(x_k)^T d_k| \\ &\geq \delta \alpha_L \gamma |\nabla f(x_k)^T d_k| \geq \delta \alpha_L \gamma c_2 \phi_\eta(x_k)^2. \end{aligned} \quad (15)$$

Recalling the definition of k_ε , for $k = 0, \dots, k_\varepsilon - 1$, we have $\phi_\eta(x_k) > \varepsilon$ and thus $f(x_k) - f(x_{k+1}) \geq \delta \alpha_L \gamma c_2 \varepsilon^2$. By summing the inequalities, for $k = 0, \dots, k_\varepsilon - 1$, we get $f(x_0) - f(x_{k_\varepsilon}) \geq k_\varepsilon \delta \alpha_L \gamma c_2 \varepsilon^2$, from which we can write

$$f(x_0) - f^* \geq f(x_0) - f(x_{k_\varepsilon}) \geq k_\varepsilon \delta \alpha_L \gamma c_2 \varepsilon^2;$$

it then follows

$$k_\varepsilon \leq \frac{(f(x_0) - f^*)}{\delta \alpha_L \gamma c_2} \varepsilon^{-2}.$$

Moreover, by adding together all the relations (15) from the first iteration up to the k -th one, we have

$$f(x_0) - f^* \geq f(x_0) - f(x_{k+1}) \geq \delta \alpha_L \gamma c_2 \sum_{i=0}^k \phi_\eta(x_i)^2.$$

Taking the limit for $k \rightarrow \infty$, we have $f(x_0) - f^* \geq \delta\alpha_L\gamma c_2 \sum_{i=0}^{\infty} \phi_\eta(x_i)^2$ so that we get

$$f(x_0) - f^* \geq \delta\alpha_L\gamma c_2 \sum_{i=0}^{\infty} \phi_\eta(x_i)^2 \geq \delta\alpha_L\gamma c_2 \sum_{\phi_\eta(x_i) \geq \varepsilon} \phi_\eta(x_i)^2 \geq \delta\alpha_L\gamma c_2 n_{i_\varepsilon} \varepsilon^2,$$

which finally implies:

$$n_{i_\varepsilon} \leq \frac{(f(x_0) - f^*)}{\delta\alpha_L\gamma c_2} \varepsilon^{-2}.$$

□

Remark 3.2 The bounds on k_ε and n_{i_ε} stated in Proposition 3.2 are equal. This result makes sense: in the worst case, all n_{i_ε} iterations with non- ε -stationary iterates will have to be performed before obtaining x_{k_ε} .

Remark 3.3 If we assume that the construction of d_k requires a fixed number of (additional) gradient evaluations and a fixed – possibly zero – number of projections, then we can immediately derive from the result of Proposition 3.2 an $\mathcal{O}(\varepsilon^{-2})$ bound for the number of gradient evaluations and projections required for obtaining x_{k_ε} .

By standard reasoning similar to that from [12, Lemma 2.2.1], we can also note that the number of function evaluations nf_{α_k} needed to satisfy the Armijo criterion at iteration k is bounded according to

$$nf_{\alpha_k} \leq \left\lceil \log_\delta \left(\frac{\alpha_L}{a} \right) \right\rceil + 1.$$

The bound is constant w.r.t. k . If we then assume that the number of function evaluations required to define d_k is also fixed, we can finally derive an $\mathcal{O}(\varepsilon^{-2})$ bound for the number of function evaluations required to reach x_{k_ε} .

The result from Proposition 3.2 also allows us to immediately derive a general global asymptotic convergence result.

Corollary 3.1 *Let $\{x_k\}$, $\{d_k\}$ be the sequences produced by Algorithm 1, and assume S is a compact set. Then $\{x_k\}$ admits accumulation points, each one being stationary for problem (1).*

Proof By the instructions of the algorithm, $\{x_k\} \subseteq S$; by the compactness of S , we immediately conclude that accumulation points exist for $\{x_k\}$. Now, let $\bar{x}$ be any such accumulation point, i.e., there exists $K \subseteq \{0, 1, \dots\}$ such that $x_k \rightarrow \bar{x}$ for $k \in K$, $k \rightarrow \infty$. Recalling from the proof of Proposition 3.2 that

$$\sum_{k=0}^{\infty} \phi_\eta(x_k)^2 \leq \frac{f(x_0) - f^*}{\delta\alpha_L\gamma c_2} < +\infty,$$

we have that $\lim_{k \rightarrow \infty} \phi_\eta(x_k) = 0$. By the continuity of ϕ_η , we then get $0 = \lim_{k \in K, k \rightarrow \infty} \phi_\eta(x_k) = \phi_\eta(\bar{x})$, which concludes the proof. □

4 Projected Gradient Directions with Momentum

To define a projected gradient method with momentum, we need to feed Algorithm 1 with directions of the form

$$d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k, \quad (16)$$

where $\alpha_k, \beta_k \geq 0$ and, for $k \geq 1$,

$$\hat{d}_k = \mathcal{P}_S[x_k - \eta_k \nabla f(x_k)] - x_k, \quad \hat{s}_k = \mathcal{P}_S[x_k + (x_k - x_{k-1})] - x_k. \quad (17)$$

In every iteration, the parameter η_k is allowed to vary as long as it satisfies

$$\eta_{\min} \leq \eta_k \leq \eta_{\max}, \quad (18)$$

where $0 < \eta_{\min} \leq \eta_{\max} < +\infty$ are two fixed scalars. To determine α_k and β_k , we can solve a two-dimensional quadratic programming problem defined according to two possible strategies, inspired by [24]: the first involves a $n \times n$ matrix B_k , whereas the second involves a 2×2 matrix H_k .

We can thus consider a quadratic subproblem with the additional constraint on direction structure:

$$\begin{aligned} \min_{d, \alpha, \beta} \quad & \nabla f(x_k)^T d + \frac{1}{2} d^T B_k d \\ \text{s.t.} \quad & x_k + d \in S, \quad d = \alpha \hat{d}_k + \beta \hat{s}_k. \end{aligned} \quad (19)$$

Now, the above problem may be too complex to deal with if the goal is to conveniently exploit it within an outer algorithmic framework. We thus aim to define a two-dimensional subproblem whose feasible set is contained in that of (19), thus possibly excluding optimal solutions, but whose solution can be found in closed form for any convex set S . First, we need to state the following result.

Proposition 4.1 *For every pair of values α, β satisfying $\alpha + \beta \leq 1$ and $\alpha \geq 0, \beta \geq 0$, the direction $d = \alpha \hat{d}_k + \beta \hat{s}_k$ is such that $x_k + d \in S$.*

Proof Let (α, β) be scalars satisfying the assumptions of the proposition. The thesis straightforwardly holds for both $(\alpha, \beta) = (1, 0)$ and $(\alpha, \beta) = (0, 1)$ by the definition of Euclidean projection. Let us now consider a generic choice of values $\alpha, \beta > 0$. We can write

$$x_k + \alpha \hat{d}_k + \beta \hat{s}_k = \alpha(x_k + \hat{d}_k) + \beta(x_k + \hat{s}_k) + x_k(1 - \alpha - \beta).$$

Let $c = 1 - \alpha - \beta$. Since $x_k + \hat{d}_k \in S$, $x_k + \hat{s}_k \in S$, $x_k \in S$, and $\alpha + \beta + c = 1$, we have that $x_k + \alpha \hat{d}_k + \beta \hat{s}_k$ is a convex combination of points in S and thus it belongs to S . $\square$

On the basis of the above result, we can then replace problem (19) with the following two-dimensional problem

$$\begin{aligned} \min_{d, \alpha, \beta} \quad & \nabla f(x_k)^T d + \frac{1}{2} d^T B_k d \\ \text{s.t.} \quad & d = \alpha \hat{d}_k + \beta \hat{s}_k, \quad \alpha + \beta \leq 1, \quad \alpha \geq 0, \quad \beta \geq 0, \end{aligned} \quad (20)$$

i.e., setting $P_k = [\hat{d}_k \ \hat{s}_k]$ and $g_k = \nabla f(x_k)$,

$$\min_{\alpha, \beta} \quad \frac{1}{2} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}^T P_k^T B_k P_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + g_k^T P_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad \text{s.t.} \quad \alpha + \beta \leq 1, \quad \alpha \geq 0, \quad \beta \geq 0.$$

We could further consider replacing the 2×2 matrix $P_k^T B_k P_k$ with a generic 2×2 matrix H_k , thus obtaining the problem

$$\min_{\alpha, \beta} \quad \frac{1}{2} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}^T H_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + g_k^T P_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad \text{s.t.} \quad \alpha + \beta \leq 1, \quad \alpha \geq 0, \quad \beta \geq 0. \quad (21)$$

As shown in the Appendix, the solution of (21) can indeed be determined in closed form; thus, we can efficiently rely on this problem for the construction of d_k in the algorithmic framework.

In the next two propositions, we state conditions on B_k and H_k respectively, ensuring that the obtained values α_k and β_k are such that a sequence of directions $\{d_k\}$ produced according to (16) satisfies Assumption 3.1.

Proposition 4.2 *Let $\{B_k\} \subseteq \mathbb{R}^{n \times n}$ be a sequence of positive definite symmetric matrices, assuming that $0 < v_1 \leq v_2$ exist such that $v_1 \leq \lambda_{\min}(B_k) \leq \lambda_{\max}(B_k) \leq v_2$ for all k . Let $\{\eta_k\}$ be a sequence of values such that, for all k , condition (18) holds, assuming that $\eta_{\max} < \frac{2}{v_1}$. Furthermore, let $d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k$ be the solution, for each k , of problem (20) induced by matrix B_k and vectors $\hat{d}_k, \hat{s}_k$ defined according to (17). Then, for any $\eta \in [\eta_{\min}, \eta_{\max}]$, the sequence of directions $\{d_k\}$ satisfies Assumption 3.1.*

Proof Proposition 4.1 ensures that d_k satisfies condition (12) of Assumption 3.1 with $a = 1$.

Let us now consider the vector $\hat{d}_k = \mathcal{P}_S[x_k - \eta_k \nabla f(x_k)] - x_k$. As already observed, $\hat{d}_k$ is a feasible point for problem (20) obtained setting $\alpha = 1, \beta = 0$. Recalling that d_k is optimal for the subproblem at iteration k and noting that the direction $\tilde{d}_k = \frac{v_1}{v_2} \hat{d}_k$ is feasible, we can then write

$$\begin{aligned} \nabla f(x_k)^T d_k &\leq \nabla f(x_k)^T \hat{d}_k + \frac{1}{2} \hat{d}_k^T B_k \hat{d}_k \\ &\leq \nabla f(x_k)^T \tilde{d}_k + \frac{1}{2} \tilde{d}_k^T B_k \tilde{d}_k \\ &\leq -\frac{v_1}{v_2} \left(\frac{1}{\eta_k} \|\hat{d}_k\|^2 - \frac{1}{2} \lambda_{\max}(B_k) \frac{v_1}{v_2} \|\hat{d}_k\|^2 \right) \end{aligned}$$

$$\begin{aligned}
&\leq -\frac{v_1}{v_2} \left(\frac{1}{\eta_k} - \frac{1}{2} v_1 \right) \|\hat{d}_k\|^2 \\
&\leq -\frac{v_1}{v_2} \left(\frac{1}{\eta_{\max}} - \frac{1}{2} v_1 \right) \|\mathcal{P}_S[x_k - \eta_k \nabla f(x_k)] - x_k\|^2 \\
&= -\frac{v_1}{v_2} \eta_k^2 \left(\frac{1}{\eta_{\max}} - \frac{1}{2} v_1 \right) \frac{1}{\eta_k^2} \|\mathcal{P}_S[x_k - \eta_k \nabla f(x_k)] - x_k\|^2 \\
&\leq -\frac{v_1}{v_2} \eta_{\min}^2 \left(\frac{1}{\eta_{\max}} - \frac{1}{2} v_1 \right) \frac{1}{\eta_{\max}^2} \|\mathcal{P}_S[x_k - \eta_{\max} \nabla f(x_k)] - x_k\|^2 \\
&\leq -\frac{v_1}{v_2} \eta_{\min}^2 \left(\frac{1}{\eta_{\max}} - \frac{1}{2} v_1 \right) \frac{1}{\eta_{\max}^2} \|\mathcal{P}_S[x_k - \eta \nabla f(x_k)] - x_k\|^2,
\end{aligned}$$

where the third inequality follows from Proposition 2.1, and the fifth one comes from $\eta_k \leq \eta_{\max}$. For the sixth one, we used that $\eta_k \geq \eta_{\min}$, and we exploited the inequality

$$\frac{1}{\eta_k} \|\mathcal{P}_S[x_k - \eta_k \nabla f(x_k)] - x_k\| \geq \frac{1}{\eta_{\max}} \|\mathcal{P}_S[x_k - \eta_{\max} \nabla f(x_k)] - x_k\|,$$

which follows from $\eta_k \leq \eta_{\max}$ and $h(t) = \frac{1}{t} \|\mathcal{P}_S[x - t \nabla f(x)] - x\|$ being non-increasing, according to Lemma 2.1. The last one finally follows from function $p(t) = \|\mathcal{P}_S[x - t \nabla f(x)] - x\|$ being nondecreasing by Lemma 2.1, which implies

$$\|\mathcal{P}_S[x_k - \eta_{\max} \nabla f(x_k)] - x_k\| \geq \|\mathcal{P}_S[x_k - \eta \nabla f(x_k)] - x_k\|,$$

as $\eta \leq \eta_{\max}$. Point (14) of Assumption 3.1 therefore follows for any $\eta \in [\eta_{\min}, \eta_{\max}]$ by defining

$$c_2 = \frac{v_1}{v_2} \frac{\eta_{\min}^2}{\eta_{\max}^2} \left(\frac{1}{\eta_{\max}} - \frac{1}{2} v_1 \right) > 0.$$

From the above chain of inequalities, we also get $\nabla f(x_k)^T d_k + \frac{1}{2} d_k^T B_k d_k \leq 0$, from which we obtain

$$\nabla f(x_k)^T d_k \leq -\frac{1}{2} d_k^T B_k d_k \leq -\frac{1}{2} v_1 \|d_k\|^2,$$

and hence, condition (13) of Assumption 3.1 also holds, setting $c_1 = v_1/2$. $\square$

We might want to elaborate on the condition $\eta_{\max} < 2/v_1$: large values of v_1 indicate that the matrix will induce a very strong regularization term within the subproblem, leading to a very small search direction; in an unconstrained scenario, this corresponds to perturbing the gradient direction with an almost null preconditioner. It is therefore not surprising that with large v_1 we can prove the property for the sequence of directions only if we consider correspondingly small values of η , i.e., taking into account less stringent measures of stationarity and of descent. A similar observation will analogously hold for the result in the following proposition concerning the case of the 2×2 matrix.

Proposition 4.3 Let $\{\hat{H}_k\} \subseteq \mathbb{R}^{2 \times 2}$ be a sequence of positive definite symmetric matrices, assuming that $0 < \hat{v}_1 \leq \hat{v}_2$ exist such that

$$\hat{v}_1 \leq \lambda_{\min}(\hat{H}_k) \quad \text{and} \quad (\hat{H}_k)_{11} \leq \hat{v}_2, \quad (22)$$

for all k . Let $\{\eta_k\}$ be a sequence of values such that, for all k , condition (18) holds, assuming that $\eta_{\max} < \frac{2}{\hat{v}_1}$. For all k , let $H_k \in \mathbb{R}^{2 \times 2}$ be the symmetric matrix in problem (21), defined as

$$H_k = D_k \hat{H}_k D_k \quad \text{where} \quad D_k = \begin{bmatrix} \|\hat{d}_k\| & 0 \\ 0 & \|\hat{s}_k\| \end{bmatrix},$$

and let $d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k$, where $[\alpha_k \ \beta_k]^T$ is the solution of problem (21). Then, for any $\eta \in [\eta_{\min}, \eta_{\max}]$, the sequence of directions $\{d_k\}$ satisfies Assumption 3.1.

Proof As already observed, any vector d belonging to the feasible set of (21) is such that $x_k + d \in \mathcal{S}$, so that the condition (12) of Assumption 3.1 is satisfied by d_k for all k with $a = 1$. The objective of problem (21) at iteration k can be rewritten as

$$\min_{\alpha, \beta} \frac{1}{2} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}^T D_k \hat{H}_k D_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + \begin{bmatrix} g_k^T \hat{d}_k \\ g_k^T \hat{s}_k \end{bmatrix}^T \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

Note that $(\hat{\alpha}, \hat{\beta})^T = (\frac{\hat{v}_1}{\hat{v}_2}, 0)^T$ is a feasible point of the two-dimensional problem. Then we can consider the corresponding direction $\tilde{d}_k = \hat{\alpha} \hat{d}_k + \hat{\beta} \hat{s}_k$ with $(\hat{\alpha}, \hat{\beta})^T = (\frac{\hat{v}_1}{\hat{v}_2}, 0)^T$. If $(\alpha_k, \beta_k)^T$ is the optimal solution of (21) and $d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k$, recalling Proposition 2.1, assumption (18) and $\eta_{\max} < 2/\hat{v}_1$, and by similar derivations as in the proof of Proposition 4.2, we can write:

$$\begin{aligned} \nabla f(x_k)^T d_k &\leq \nabla f(x_k)^T d_k + \frac{1}{2} \begin{bmatrix} \alpha_k \\ \beta_k \end{bmatrix}^T D_k \hat{H}_k D_k \begin{bmatrix} \alpha_k \\ \beta_k \end{bmatrix} \\ &\leq \nabla f(x_k)^T \frac{\hat{v}_1}{\hat{v}_2} \hat{d}_k + \frac{1}{2} (\hat{H}_k)_{11} \frac{\hat{v}_1^2}{\hat{v}_2^2} \|\hat{d}_k\|^2 \\ &\leq -\frac{\hat{v}_1}{\hat{v}_2} \frac{1}{\eta_k} \|\hat{d}_k\|^2 + \frac{1}{2} \frac{\hat{v}_1^2}{\hat{v}_2} \|\hat{d}_k\|^2 \\ &\leq -\frac{\hat{v}_1}{\hat{v}_2} \left(\frac{1}{\eta_k} - \frac{1}{2} \hat{v}_1 \right) \|\mathcal{P}_{\mathcal{S}}[x_k - \eta_k \nabla f(x_k)] - x_k\|^2 \\ &\leq -\frac{\hat{v}_1}{\hat{v}_2} \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right) \|\mathcal{P}_{\mathcal{S}}[x_k - \eta_k \nabla f(x_k)] - x_k\|^2 \\ &= -\frac{\hat{v}_1}{\hat{v}_2} \eta_k^2 \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right) \frac{1}{\eta_k^2} \|\mathcal{P}_{\mathcal{S}}[x_k - \eta_k \nabla f(x_k)] - x_k\|^2 \\ &\leq -\frac{\hat{v}_1}{\hat{v}_2} \eta_{\min}^2 \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right) \frac{1}{\eta_{\max}^2} \|\mathcal{P}_{\mathcal{S}}[x_k - \eta_{\max} \nabla f(x_k)] - x_k\|^2 \end{aligned}$$

$$\begin{aligned}
&\leq -\frac{\hat{v}_1}{\hat{v}_2} \eta_{\min}^2 \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right) \frac{1}{\eta_{\max}^2} \|\mathcal{P}_S[x_k - \eta \nabla f(x_k)] - x_k\|^2 \\
&\leq -c_2 \phi_\eta(x_k)^2,
\end{aligned}$$

where

$$c_2 = \frac{\hat{v}_1}{\hat{v}_2} \frac{\eta_{\min}^2}{\eta_{\max}^2} \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right) > 0,$$

so that condition (14) of Assumption 3.1 is proved.

Then, from the above chain of inequalities and by using (22), we have

$$\begin{aligned}
\nabla f(x_k)^T d_k &\leq -\frac{1}{2} \begin{bmatrix} \alpha_k \\ \beta_k \end{bmatrix}^T D_k \hat{H}_k D_k \begin{bmatrix} \alpha_k \\ \beta_k \end{bmatrix} \leq -\frac{1}{2} \hat{v}_1 \left((\alpha_k \|\hat{d}_k\|)^2 + (\beta_k \|\hat{s}_k\|)^2 \right) \\
&\leq -\frac{1}{4} \hat{v}_1 \|\alpha_k \hat{d}_k + \beta_k \hat{s}_k\|^2 = -\frac{1}{4} \hat{v}_1 \|d_k\|^2.
\end{aligned}$$

Therefore, condition (13) of the Assumption holds by setting $c_1 = \frac{\hat{v}_1}{4}$. $\square$

Remark 4.1 Following the approach presented in [6], it is possible to choose the parameter η_k by computing the inverse Rayleigh quotient. Specifically, given two constants $0 < \eta_{\min} \leq \eta_{\max} < +\infty$, we can choose η_k as:

$$\eta_k = \min \left\{ \eta_{\max}, \max \left\{ \eta_{\min}, \frac{s_k^T s_k}{s_k^T y_k} \right\} \right\},$$

where $y_k = \nabla f(x_k) - \nabla f(x_{k-1})$ and $s_k = x_k - x_{k-1}$, with the choice of $\eta_k = \eta_{\max}$ if $s_k^T y_k \leq 0$.

In the next section, we present a specific projected gradient method with momentum where the search direction is induced by a tailored 2×2 matrix H_k .

5 A Convergent Projected Gradient Method with Momentum

The analysis and the remarks from the preceding section drive us towards the design of a convergent projected gradient method with momentum. Summarizing for convenience the main insights discussed thus far, we will have to exploit search directions of the form $d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k$, where $\alpha_k, \beta_k \geq 0$, and, for $k \geq 1$,

$$\hat{d}_k = \mathcal{P}_S[x_k - \eta_k \nabla f(x_k)] - x_k, \quad \hat{s}_k = \mathcal{P}_S[x_k + (x_k - x_{k-1})] - x_k,$$

with $\eta_{\min} \leq \eta_k \leq \eta_{\max}$, in such a way that Assumption 3.1 holds.

While the analysis provides us with two sound options to obtain suitable values of α_k and β_k , we will focus here only on the latter, based on the two-dimensional subproblems, which appears to be more convenient from the numerical standpoint.

The key computational element of the algorithmic framework lies, therefore, in the definition of the 2×2 matrix H_k in subproblem (21), which we also rewrite below for convenience:

$$\min_{\alpha, \beta} \frac{1}{2} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}^T H_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + g_k^T P_k \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad \text{s.t.} \quad \alpha + \beta \leq 1, \quad \alpha \geq 0, \quad \beta \geq 0.$$

The solution $(\alpha^* \ \beta^*)^T$ of the above subproblem provides the scalars $\alpha_k = \alpha^*$ and $\beta_k = \beta^*$ defining the search direction d_k given by (16). In the above problem, all parameters are fixed given x_k , except for the three elements defining the symmetric matrix H_k , which thus leave three degrees of freedom.

To identify a meaningful matrix H_k to employ in the subproblem, we interpret the objective function as the quadratic Taylor polynomial approximation of the bivariate function of variables α and β , $f(x_k + \alpha \hat{d}_k + \beta \hat{s}_k)$:

$$\varphi_k(\alpha, \beta; H_k) = f(x_k) + \begin{bmatrix} g_k^T \hat{d}_k \\ g_k^T \hat{s}_k \end{bmatrix}^T \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}^T \begin{bmatrix} H_{11} & H_{12} \\ H_{12} & H_{22} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

The three elements defining H_k can then be determined by imposing the interpolation conditions on three points (α_1, β_1) , (α_2, β_2) , and (α_3, β_3) different from $(0, 0)$:

$$\begin{aligned} \varphi_k(\alpha_1, \beta_1; H_k) &= f(x_k + \alpha_1 \hat{d}_k + \beta_1 \hat{s}_k), \\ \varphi_k(\alpha_2, \beta_2; H_k) &= f(x_k + \alpha_2 \hat{d}_k + \beta_2 \hat{s}_k), \\ \varphi_k(\alpha_3, \beta_3; H_k) &= f(x_k + \alpha_3 \hat{d}_k + \beta_3 \hat{s}_k), \end{aligned} \quad (23)$$

and thus by simply solving a 3×3 linear system, which can be handled in closed form. The construction of a reliable quadratic matrix incurs the cost of a small number of additional function evaluations per iteration (at most 3).

Before formally presenting the algorithm, we summarize its main steps at each iteration:

- a quadratic subproblem of the form (21) is defined, where $H_k \in \mathbb{R}^{2 \times 2}$ is a symmetric matrix obtained by the aforementioned interpolation technique;
- once a solution $[\alpha_k \ \beta_k]^T$ of (21) has been computed, the corresponding search direction of the form (16) undergoes a check to ensure the gradient-related property (with fixed ϕ_η induced by a constant value η) of the whole sequence $\{d_k\}$;
- if the test is satisfied, then a standard Armijo-type line search can be performed along d_k ; otherwise, a suitable modification to H_k is made that guarantees the conditions of Proposition 4.3 to hold: steps (a) (with the modified $\tilde{H}_k$) and (b) (without the check on d_k) are then repeated, as well as the Armijo-type line search along the obtained d_k .

Algorithm 2 describes the instructions of the proposed procedure rigorously and in detail. In the following lemma we prove that the modified matrix $\tilde{H}_k$ computed as in lines 29–33 satisfies the assumptions of Proposition 4.3.

Algorithm 2 Projected Gradient Method with Momentum (PGMM)

```

1: Input:  $x_0 \in \mathcal{S}$ ,  $\gamma \in (0, 1)$ ,  $\delta \in (0, 1)$ ,  $0 < \eta_{\min} \leq \eta \leq \eta_{\max}$ ,  $\bar{c}_1, \bar{c}_2 > 0$ ,  $0 < \hat{v}_1 \leq \hat{v}_2$ .
2: Set  $k = 0$ 
3: Let  $x_{-1} = x_0$ 
4: for  $k = 0, 1, \dots$  do
5:   Set  $s_k = x_k - x_{k-1}$ ,  $y_k = \nabla f(x_k) - \nabla f(x_{k-1})$ 
6:   /* Define spectral parameter */
7:   if  $s_k^T y_k > 0$  then
8:     Set  $\eta_k = \min \left\{ \eta_{\max}, \max \left\{ \eta_{\min}, \frac{s_k^T s_k}{s_k^T y_k} \right\} \right\}$ 
9:   else
10:    Set  $\eta_k = \eta_{\max}$ 
11:    /* Define base directions */
12:    Set  $\hat{d}_k = \mathcal{P}_{\mathcal{S}}[x_k - \eta_k \nabla f(x_k)] - x_k$ ,  $\hat{s}_k = \mathcal{P}_{\mathcal{S}}[x_k + s_k] - x_k$ 
13:    /* Stop if  $x_k$  is a stationary point */
14:    if  $\hat{d}_k = 0$  then
15:      return  $x_k$ 
16:    if  $\hat{s}_k \neq 0$  then
17:      /* Common nondegenerate case */
18:      /* Define  $2 \times 2$  quadratic matrix by interpolation */
19:      Compute  $H_k$  imposing conditions (23) at three points  $(\alpha_1, \beta_1)$ ,  $(\alpha_2, \beta_2)$ ,  $(\alpha_3, \beta_3)$ 
20:      /* Compute the search direction */
21:      Set  $\alpha_k, \beta_k \in \arg \min_{\alpha, \beta \geq 0, \alpha + \beta \leq 1} \varphi_k(\alpha, \beta; H_k)$  obtained using Algorithm 3
22:      Set  $d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k$ 
23:      /* Gradient-related check */
24:      Set  $gr\_dir\_found = \text{False}$ 
25:      if  $\nabla f(x_k)^T d_k \leq -\bar{c}_1 \|d_k\|^2$  and  $\nabla f(x_k)^T d_k \leq -\bar{c}_2 \|\mathcal{P}_{\mathcal{S}}[x_k - \eta \nabla f(x_k)] - x_k\|^2$  then
26:        Set  $gr\_dir\_found = \text{True}$ 
27:      if  $gr\_dir\_found = \text{False}$  then
28:        /* Modify  $2 \times 2$  matrix */
29:        Set  $(\tilde{H}_k)_{11} = \max \{ \min \{ (H_k)_{11}, \hat{v}_2 \|\hat{d}_k\|^2 \}, \hat{v}_1 \|\hat{d}_k\|^2 \}$ 
30:        Set  $(\tilde{H}_k)_{22} = \max \{ (H_k)_{22}, \hat{v}_1 \|\hat{s}_k\|^2 \}$ 
31:        Set  $\tilde{r}_k = \sqrt{((\tilde{H}_k)_{11} - \hat{v}_1 \|\hat{d}_k\|^2)((\tilde{H}_k)_{22} - \hat{v}_1 \|\hat{s}_k\|^2)}$ 
32:        Set  $(\tilde{H}_k)_{12} = \max \{ \min \{ (H_k)_{12}, \tilde{r}_k \}, -\tilde{r}_k \}$ 
33:        Set  $\tilde{H}_k = \begin{pmatrix} (\tilde{H}_k)_{11} & (\tilde{H}_k)_{12} \\ (\tilde{H}_k)_{12} & (\tilde{H}_k)_{22} \end{pmatrix}$ 
34:        /* Recompute the search direction */
35:        Set  $\alpha_k, \beta_k \in \arg \min_{\alpha, \beta \geq 0, \alpha + \beta \leq 1} \varphi_k(\alpha, \beta; \tilde{H}_k)$  obtained using Algorithm 3
36:        Set  $d_k = \alpha_k \hat{d}_k + \beta_k \hat{s}_k$ 
37:      else
38:        /* No momentum, do simple projected gradient iteration */
39:         $d_k = \hat{d}_k$ 
40:        /* Perform Armijo line search along  $d_k$  */
41:        Set  $\mu = 1$ 
42:        while  $f(x_k + \mu d_k) > f(x_k) + \gamma \mu \nabla f(x_k)^T d_k$  do
43:          Set  $\mu = \delta \mu$ 
44:        Set  $\mu_k = \mu$ ,
45:        Set  $x_{k+1} = x_k + \mu_k d_k$ 

```

Lemma 5.1 Let $\|\hat{d}_k\| \neq 0$, $\|\hat{s}_k\| \neq 0$, and let $\tilde{H}_k$ be the matrix computed in lines 29–33 of Algorithm 2. Then, $\tilde{H}_k = D_k \hat{H}_k D_k$, where $D_k = \begin{bmatrix} \|\hat{d}_k\| & 0 \\ 0 & \|\hat{s}_k\| \end{bmatrix}$ and $\hat{H}_k$ is a 2×2 symmetric matrix such as $(\hat{H}_k)_{11} \leq \hat{v}_2$ and $\lambda_{\min}(\hat{H}_k) \geq \hat{v}_1$.

Proof We can explicitly compute the components of $\hat{H}_k$ as:

$$(\hat{H}_k)_{11} = \frac{(\tilde{H}_k)_{11}}{\|\hat{d}_k\|^2} = \max \left\{ \min \left\{ \frac{(H_k)_{11}}{\|\hat{d}_k\|^2}, \hat{v}_2 \right\}, \hat{v}_1 \right\}, \quad (24)$$

$$(\hat{H}_k)_{22} = \frac{(\tilde{H}_k)_{22}}{\|\hat{s}_k\|^2} = \max \left\{ \frac{(H_k)_{22}}{\|\hat{s}_k\|^2}, \hat{v}_1 \right\}, \quad (25)$$

$$(\hat{H}_k)_{12} = \frac{(\tilde{H}_k)_{12}}{\|\hat{d}_k\| \|\hat{s}_k\|} = \max \left\{ \min \left\{ \frac{(H_k)_{12}}{\|\hat{d}_k\| \|\hat{s}_k\|}, r_k \right\}, -r_k \right\}, \quad (26)$$

where

$$\begin{aligned} r_k &= \frac{\tilde{r}_k}{\|\hat{d}_k\| \|\hat{s}_k\|} = \sqrt{\left(\frac{(\tilde{H}_k)_{11}}{\|\hat{d}_k\|^2} - \hat{v}_1 \right) \left(\frac{(\tilde{H}_k)_{22}}{\|\hat{s}_k\|^2} - \hat{v}_1 \right)} \\ &= \sqrt{((\hat{H}_k)_{11} - \hat{v}_1)((\hat{H}_k)_{22} - \hat{v}_1)}. \end{aligned} \quad (27)$$

By (24), $(\hat{H}_k)_{11} \leq \hat{v}_2$. Moreover,

$$\begin{aligned} &\sqrt{((\hat{H}_k)_{11} - (\hat{H}_k)_{22})^2 + 4(\hat{H}_k)_{12}^2} \\ &\leq \sqrt{((\hat{H}_k)_{11} - (\hat{H}_k)_{22})^2 + 4((\hat{H}_k)_{11} - \hat{v}_1)((\hat{H}_k)_{22} - \hat{v}_1)} \\ &= \sqrt{((\hat{H}_k)_{11} - (\hat{H}_k)_{22})^2 + 4(\hat{H}_k)_{11}(\hat{H}_k)_{22} - 4\hat{v}_1((\hat{H}_k)_{11} + (\hat{H}_k)_{22}) + 4\hat{v}_1^2} \\ &= \sqrt{((\hat{H}_k)_{11} + (\hat{H}_k)_{22})^2 - 4\hat{v}_1((\hat{H}_k)_{11} + (\hat{H}_k)_{22}) + 4\hat{v}_1^2} \\ &= \sqrt{((\hat{H}_k)_{11} + (\hat{H}_k)_{22} - 2\hat{v}_1)^2} = (\hat{H}_k)_{11} + (\hat{H}_k)_{22} - 2\hat{v}_1, \end{aligned}$$

where in the first inequality we employed that, by (26)–(27),

$$(\hat{H}_k)_{12}^2 \leq ((\hat{H}_k)_{11} - \hat{v}_1)((\hat{H}_k)_{22} - \hat{v}_1),$$

and in the last equality that, by (24)–(25), $(\hat{H}_k)_{11} \geq \hat{v}_1$ and $(\hat{H}_k)_{22} \geq \hat{v}_1$. As a consequence of the above chain of inequalities, the minimum eigenvalue of $\hat{H}_k$ is bounded by:

$$\begin{aligned} \lambda_{\min}(\hat{H}_k) &= \frac{(\hat{H}_k)_{11} + (\hat{H}_k)_{22} - \sqrt{((\hat{H}_k)_{11} - (\hat{H}_k)_{22})^2 + 4(\hat{H}_k)_{12}^2}}{2} \\ &\geq \frac{(\hat{H}_k)_{11} + (\hat{H}_k)_{22} - ((\hat{H}_k)_{11} + (\hat{H}_k)_{22} - 2\hat{v}_1)}{2} = \hat{v}_1. \end{aligned}$$

□

Recalling the results proven in this paper, we are finally able to state the main convergence and complexity results for Algorithm 2.

Proposition 5.1 *Let Assumption 2.1 hold and let $\{x_k\}$ be the sequence produced by Algorithm 2. Then, if $\eta_{\max} < \frac{2}{\hat{v}_1}$, for any $\eta \in [\eta_{\min}, \eta_{\max}]$ and $\varepsilon > 0$ and recalling the notation of Proposition 3.2, we have $k_\varepsilon = \mathcal{O}(\varepsilon^{-2})$ and $ni_\varepsilon = \mathcal{O}(\varepsilon^{-2})$. The asymptotic bound of $\mathcal{O}(\frac{1}{\varepsilon^2})$ also holds for the number of function evaluations, gradient evaluations and projections carried out by the algorithm in the first k_ε iterations. Moreover, if $\mathcal{S}$ is compact, either the algorithm stops in a finite number $\bar{k}$ of iterations returning a stationary point $x_{\bar{k}}$, or the sequence $\{x_k\}$ has accumulation points, each one being stationary for problem (1).*

Proof We focus on the nontrivial case for which $\hat{d}_k \neq 0$ for all k ; otherwise, by definition of stationarity, we have finite termination with a stationary point. The proof for the complexity results directly follows from Proposition 3.2, noting that the algorithm carries out an Armijo line search and, for any choice of $\eta_{\min} \leq \eta \leq \eta_{\max}$, the sequence $\{d_k\}$ satisfies Assumption 3.1. In particular, condition (12) always holds with $a = 1$. As for conditions (13) and (14), we can note that:

- for each k such that $\hat{s}_k \neq 0$ and `gr_dir_found=True`, the safeguard at step 25 ensures that d_k satisfies the inequalities for the values $c_1 = \bar{c}_1$ and $c_2 = \bar{c}_2$ provided as inputs to the algorithm;
- if $\hat{s}_k \neq 0$ but the check at step 25 fails, the perturbation of matrix H_k carried out at steps 29–33 guarantees that the assumption of Proposition 4.3 holds (see Lemma 5.1); d_k is then constructed in such a way that the two inequalities hold for constants $c_1 = \frac{\hat{v}_1}{4}$ and $c_2 = \frac{\hat{v}_1}{\hat{v}_2} \frac{\eta_{\min}^2}{\eta_{\max}^2} \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right)$, retrieved in the proof of Proposition 4.3;
- if $\hat{s}_k = 0$, d_k is equal to $\hat{d}_k$ and, by Proposition 2.1 (see also Remark 3.1) and the fact that $\eta_k \leq \eta_{\max}$,

$$\nabla f(x_k)^T \hat{d}_k \leq -\frac{1}{\eta_k} \|\hat{d}_k\|^2 \leq -\frac{1}{\eta_{\max}} \|\hat{d}_k\|^2.$$

Therefore, the first inequality is satisfied with $c_1 = \frac{1}{\eta_{\max}}$. Moreover, reasoning similarly to the proof of Propositions 4.2–4.3,

$$-\frac{1}{\eta_{\max}} \|\hat{d}_k\|^2 \leq -\frac{\eta_{\min}^2}{\eta_{\max}^3} \|\mathcal{P}_{\mathcal{S}}[x_k - \eta \nabla f(x_k)] - x_k\|^2.$$

Hence, the second inequality is satisfied with $c_2 = \frac{\eta_{\min}^2}{\eta_{\max}^3}$.

Hence, the entire sequence satisfies (13) and (14) for all k with

$$c_1 = \min \left\{ \bar{c}_1, \frac{\hat{v}_1}{4}, \frac{1}{\eta_{\max}} \right\}, \quad c_2 = \min \left\{ \bar{c}_2, \frac{\hat{v}_1}{\hat{v}_2} \frac{\eta_{\min}^2}{\eta_{\max}^2} \left(\frac{1}{\eta_{\max}} - \frac{1}{2} \hat{v}_1 \right), \frac{\eta_{\min}^2}{\eta_{\max}^3} \right\}.$$

The asymptotic convergence result then follows from Corollary 3.1. $\square$

6 Numerical Results

In this section, we report the results of computational experiments carried out with the aim of assessing the potential of the proposed approach compared to baseline and state-of-the-art methods. The experimental analysis was divided into two parts, considering different classes of problems, which are discussed separately in the following.

Both scenarios, however, share the same general experimental setup: we compare our solver (PGMM) with the Spectral Projected Gradient Method (SPG) [6, 7], whose official MATLAB implementation is available at <https://www.ime.usp.br/~egbirgin/tango/index.php>. The SPG is widely regarded as the state-of-the-art projection based descent method, but it arguably also represents a baseline for our study, as it relies on simple gradient directions without momentum or preconditioning terms. For a fair comparison, the proposed Algorithm 2 was implemented within the SPG library code, introducing only the modifications required to define PGMM. The algorithm code, together with that concerning the experiments, is available at <https://github.com/DiegoScuppa/PGMM>. All experiments were carried out in a MATLAB R2025a environment.

In SPG, a non-monotone line search [20] is employed, considering the last M iterations, combined with a safeguarded quadratic interpolation strategy; the related parameters were set to the default values $M = 10$, $\gamma = 10^{-4}$, $\delta = 0.5$, $\sigma_{\min} = 0.1$, and $\sigma_{\max} = 0.9$. The same safeguarded quadratic interpolation strategy was maintained for use within the PGMM; however, we did not employ the non-monotone condition for the line search: contrary to SPG, which benefits from the non-monotone approach, we found the latter to be somewhat detrimental for the specific case of the proposed method in preliminary experiments. For the first iteration in both algorithms, we set $\eta_0 = 1/\|\mathcal{P}_S[x_0 - \nabla f(x_0)]\|_\infty$ and $d_0 = \hat{d}_0$. As for the computation of the matrix H_k , interpolation was performed using the points $(\alpha_1, \beta_1) = (0, \frac{1}{2})$, $(\alpha_2, \beta_2) = (\frac{1}{2}, 0)$, and $(\alpha_3, \beta_3) = (\frac{1}{2}, \frac{1}{2})$.

Concerning the stopping criterion, both algorithms were stopped as soon as condition $\|\mathcal{P}_S[x_k - \nabla f(x_k)] - x_k\|_\infty \leq 10^{-5}$ was achieved. Additional stopping conditions included a maximum number of 10^5 iterations and a numerical control on $\|x_k - x_{k-1}\|^2$, terminating the algorithm if it dropped below 10^{-15} .

6.1 Ball-Constrained Problems

As a first set of experiments, we considered nonlinear optimization problems with ℓ_1 -ball constraints. In particular, we focused on a small benchmark of machine learning problems consisting of constrained logistic regression (LR) tasks [22], where solution sparsity and generalization properties are enhanced by setting an upper bound on the ℓ_1 -norm of the weights rather than by setting an ℓ_1 -regularization term [26, 28].

The ℓ_1 -ball of radius R is defined as $B_1(R) = \{x \in \mathbb{R}^n : \|x\|_1 \leq R\}$, and the exact projection onto $B_1(R)$ is readily obtainable via fast procedures [13].

The structure of the resulting problem is, therefore, particularly well suited as a test case for the algorithm considered in this work. We employed a benchmark of ten LR training problems to evaluate the performance improvement achieved by Algorithm 2

Table 1 Logistic Regression problems: n denotes the number of variables (including the intercept), m the number of samples

Problem	n	m	Problem	n	m
a1a	124	1605	phishing	69	11055
a9a	124	32561	sonar	61	208
gisette	5001	6000	splice	61	1000
madelon	501	2000	w1a	301	2477
mushrooms	113	8124	w8a	301	49749

over SPG. The problems were built using the datasets reported in Table 1 and all available at <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>. The objective function is the (convex) negative log-likelihood function

$$f(w) = \frac{1}{m} \sum_{i=1}^m \log(1 + \exp(-y_i(Xw)_i)), \quad (28)$$

where $X \in \mathbb{R}^{m \times n}$ is the data matrix, $y \in \{-1, 1\}^m$ is the labels vector, and $w \in \mathbb{R}^n$ denotes the model parameters.

For each dataset, we computed an approximate unconstrained minimizer $\bar{w}$ of (28) using the L-BFGS solver provided by the MATLAB routine `fminunc`. We then performed experiments with ℓ_1 -ball constraints of radius $R = \frac{1}{2}\|\bar{w}\|_1$, so that the unconstrained minimizer would be cut out of the feasible set. For each problem, ten runs with different initial points were performed, for a total of 100 instances. CPU times, iteration counts, and function evaluations were recorded.

The comparison between PGMM and SPG is shown in Figure 1 in the form of performance profiles [15]. Detailed tables of results are available in the GitHub repository. All instances were successfully solved by both algorithms, except for five instances of the `madelon` problem where SPG failed. PGMM achieved the best CPU time on all problems except `splice` and `w8a`; nevertheless, all instances were solved by PGMM in a time not exceeding 1.5 times that of SPG. In terms of iterations, PGMM consistently outperformed SPG, while both methods exhibited comparable costs in terms of function evaluations. The above considerations suggest that projections and gradient evaluations are likely to account for a large part of the cost of the algorithms.

6.2 Bound-Constrained Problems

As a second benchmark, we considered box-constrained optimization problems, where the feasible set is $\mathcal{S} = \{x \in \mathbb{R}^n : l \leq x \leq u\}$, for given vectors l and u , whose components may also take infinite values. The projection onto $\mathcal{S}$ is trivial; indeed, $\mathcal{P}_{\mathcal{S}}[x] = \min\{u, \max\{l, x\}\}$.

We considered the bound-constrained problems from the CUTEst collection [19] (available at <https://github.com/ralna/CUTEst>), following the selection analyzed in [6, Tables 2–5]. A total of 50 problems were considered, using their default dimensions, bounds, and initial points. For each problem, ten runs starting from the same initial




Fig. 1 Performance profiles for the comparison between SPG and PGMM on the 100 LR instances with ℓ_1 -ball constraints




Fig. 2 Performance profiles for the comparison between SPG and PGMM on the 50 CUTEst problems with bound constraints

point were performed to mitigate variability in execution times. Average CPU times, iteration counts, and function evaluations were collected.

The results are summarized in Figure 2 using performance profiles. Tables containing the average results are available in the GitHub repository.

SPG failed on 9 problems, while PGMM failed on 8. Both solvers failed on BQPGAUSS, EXPQUAD, NCVXBQP2, NCVXBQP3, ODNAMUR, and QRTQUAD. Additionally, SPG failed on CHEBYQAD, QR3DLS, and SCON1LS, whereas PGMM failed on EXPLIN and EXPLIN2. For all failures except ODNAMUR, termination was due to the numerical control on successive iterates; for ODNAMUR, the maximum number of iterations was reached.

PGMM achieved the best CPU time on 36 problems, while SPG was the fastest on 8. Moreover, all problems solved by PGMM were completed in a time not exceeding 1.7 times that of SPG. PGMM also showed superior performance in terms of iterations, whereas SPG generally required fewer function evaluations. This behavior is expected since PGMM evaluates the objective function at three additional points per iteration to perform interpolation.


Overall, the numerical experiments on both problem classes highlight the superiority of PGMM over SPG. The proposed method can be interpreted as an enhancement of the classical SPG algorithm, where the introduction of a momentum term leads to a significant improvement in performance.

7 Conclusions

In this paper, we derived a rigorous convergence and complexity analysis for line search-based iterative descent algorithms for nonconvex optimization over convex sets. These novel results guided us towards the definition of convergent projected gradient-type methods exploiting momentum direction terms within their core mechanisms. From this general class of convergent methods, we drew a specific procedure that, with limited computational cost, minimizes at each iteration a two-dimensional quadratic model of the objective to construct the search direction, practically leading to a substantial speedup in convergence compared to the baseline gradient projection method. Experiments on a large benchmark of problems corroborated this claim.

Future work might concern the analysis of the proposed methods in scenarios where exact projection is not computationally reasonable. Also, combinations of the projection-based approach with tailored strategies for particular constraint sets (e.g., active-set strategies for linear constraints [2, 5, 14]) could further improve the performance of the method. Finally, adaptations of the approach to work with classes of nonconvex constraints (see, e.g., [23]) will also be of interest.

Fig. 3 Feasible set of the two-dimensional problem in α and β



Appendix: Closed Form Solution for the Problem of Minimizing a Two-Dimensional Quadratic Function Over a Polytope

We consider the two-dimensional optimization problem

$$\min_{\alpha, \beta} \varphi(\alpha, \beta) = \frac{1}{2} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}^T \begin{bmatrix} t & u \\ u & w \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + \begin{bmatrix} y \\ h \end{bmatrix}^T \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad \text{s.t.} \quad \alpha \geq 0, \beta \geq 0, \alpha + \beta \leq 1. \quad (29)$$

The feasible set of the problem is $S = \{(\alpha, \beta) \in \mathbb{R}^2 \mid \alpha \geq 0, \beta \geq 0, \alpha + \beta \leq 1\}$, which is depicted in Figure 3.

As S is compact, there exists at least one optimal solution to problem (29). We can split the set S into the interior part and the boundary: $S = \mathring{S} \cup \partial S$. In particular, the boundary could be further split into $\partial S = A_1 \cup A_2 \cup A_3 \cup \{(0, 0)\} \cup \{(0, 1)\} \cup \{(1, 0)\}$, where $A_1 = \{(\alpha, 0) \mid 0 < \alpha < 1\}$, $A_2 = \{(0, \beta) \mid 0 < \beta < 1\}$, $A_3 = \{(\alpha, 1 - \alpha) \mid 0 < \alpha < 1\}$. Hence, we can first search for a minimizer of φ in $\mathring{S}$ and then a (candidate) minimizer of φ in each subset of ∂S . First of all, we note that if the 2×2 quadratic matrix is not positive definite, then a global minimizer of φ on S can necessarily be found in ∂S . In this case, we directly skip to the analysis of ∂S . Otherwise, we compute the (unique) unconstrained minimizer of φ : if it belongs to S , then it is clearly the unique minimizer of φ on S . Otherwise, we analyze the behavior of φ in ∂S : we compute the value of φ in the three vertices of the simplex and the three minimizers of the (mono-dimensional) functions along the three edges. A sketch of an algorithm is presented as Algorithm 3.

Algorithm 3 Min of a 2-dim quadratic function over a polytope

Require: $t, u, w, y, h \in \mathbb{R}$.

```

1: /* Min on  $\hat{S}$  */
2:  $\det = tw - u^2$ 
3: if  $t > 0$  and  $\det > 0$  then
4:   /* The matrix is positive definite */
5:    $(\alpha^*, \beta^*) = (-wy + uh, uy - th) / \det$ 
6:   if  $(\alpha^*, \beta^*) \in S$  then
7:     return  $(\alpha^*, \beta^*)$ 
8:   /* Stop as it is the min on  $S$ . Otherwise: investigate  $\partial S$  */
9: /* Min on vertices */
10:  $f^* = 0$ 
11:  $(\alpha^*, \beta^*) = (0, 0)$ 
12: if  $\frac{t}{2} + y < f^*$  then
13:    $f^* = \frac{t}{2} + y$ 
14:    $(\alpha^*, \beta^*) = (1, 0)$ 
15: if  $\frac{w}{2} + h < f^*$  then
16:    $f^* = \frac{w}{2} + h$ 
17:    $(\alpha^*, \beta^*) = (0, 1)$ 
18: /* Min on  $A_1$ :  $\min_{0 < \alpha < 1} \varphi(\alpha, 0) = \frac{1}{2}t\alpha^2 + y\alpha$  */
19: if  $t > 0$  and  $0 < -\frac{y}{t} < 1$  and  $-\frac{y^2}{2t} < f^*$  then
20:    $f^* = -\frac{y^2}{2t}$ 
21:    $(\alpha^*, \beta^*) = (-\frac{y}{t}, 0)$ 
22: /* Min on  $A_2$ :  $\min_{0 < \beta < 1} \varphi(0, \beta) = \frac{1}{2}w\beta^2 + h\beta$  */
23: if  $w > 0$  and  $0 < -\frac{h}{w} < 1$  and  $-\frac{h^2}{2w} < f^*$  then
24:    $f^* = -\frac{h^2}{2w}$ 
25:    $(\alpha^*, \beta^*) = (0, -\frac{h}{w})$ 
26: /* Min on  $A_3$ :  $\min_{0 < \alpha < 1} \varphi(\alpha, 1 - \alpha) = \frac{1}{2}(t - 2u + w)\alpha^2 + (u - w + y - h)\alpha + (\frac{w}{2} + h)$  */
27:  $\bar{\alpha} = -\frac{u - w + y - h}{t - 2u + w}$ 
28: if  $t - 2u + w > 0$  and  $0 < \bar{\alpha} < 1$  and  $\varphi(\bar{\alpha}, 1 - \bar{\alpha}) < f^*$  then
29:    $f^* = \varphi(\bar{\alpha}, 1 - \bar{\alpha})$ 
30:    $(\alpha^*, \beta^*) = (\bar{\alpha}, 1 - \bar{\alpha})$ 
31: return  $(\alpha^*, \beta^*)$ 

```

Acknowledgements The authors are thankful to the editors and anonymous referees of this manuscript for the constructive comments that helped to improve the quality of the work.

Funding Open access funding provided by Università degli Studi di Roma La Sapienza within the CRUI-CARE Agreement. No funding was received for conducting this study.

Data Availability Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Code Availability Statement The code developed for the experimental part of this paper is publicly available at <https://github.com/DiegoScuppa/PGMM>.

Declarations

Competing interests The authors have no competing interests to declare that are relevant to the content of this article.

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. Andreani, R., Birgin, E.G., Martínez, J.M., Yuan, J.: Spectral projected gradient and variable metric methods for optimization with linear inequalities. *IMA J. Numer. Anal.* **25**(2), 221–252 (2005)
2. Andretta, M., Birgin, E.G., Martínez, J.M.: Partial spectral projected gradient method with active-set strategy for linearly constrained optimization. *Numerical Algorithms* **53**(1), 23–52 (2010)
3. Barzilai, J., Borwein, J.M.: Two-point step size gradient methods. *IMA J. Numer. Anal.* **8**, 141–148 (1988). <https://doi.org/10.1093/imanum/8.1.141>
4. Bertsekas, D.P.: *Nonlinear Programming*. Athena Scientific (1999)
5. Birgin, E.G., Mario Martínez, J.: Large-scale active-set box-constrained optimization method with spectral projected gradients. *Comput. Optim. Appl.* **23**(1), 101–125 (2002)
6. Birgin, E.G., Martínez, J.M., Raydan, M.: Nonmonotone spectral projected gradient methods on convex sets. *SIAM J. Optim.* **10**(4), 1196–1211 (2000). <https://doi.org/10.1137/S1052623497330963>
7. Birgin, E.G., Martínez, J.M., Raydan, M.: Algorithm 813: Spg-software for convex-constrained optimization. *ACM Trans. Math. Softw.* **27**(3), 340–349 (2001). <https://doi.org/10.1145/502800.502803>
8. Birgin, E.G., Martínez, J.M., Raydan, M.: Inexact spectral projected gradient methods on convex sets. *IMA J. Numer. Anal.* **23**(4), 539–559 (2003)
9. Bonettini, S., Prato, M.: New convergence results for the scaled gradient projection method. *Inverse Problems* **31**(9), 095,008 (2015)
10. Bonettini, S., Zanella, R., Zanni, L.: A scaled gradient projection method for constrained image deblurring. *Inverse Problems* **25**(1), 015,002 (2008)
11. Bottou, L., Curtis, F.E., Nocedal, J.: Optimization methods for large-scale machine learning. *SIAM Rev.* **60**(2), 223–311 (2018)
12. Cartis, C., Gould, N.I., Toint, P.L.: *Evaluation Complexity of Algorithms for Nonconvex Optimization: Theory*. SIAM, Computation and Perspectives (2022)
13. Condat, L.: Fast projection onto the simplex and the ℓ_1 ball. *Math. Program.* **158**(1), 575–585 (2016)
14. Cristofari, A., De Santis, M., Lucidi, S., Rinaldi, F.: Minimization over the ℓ_1 -ball using an active-set non-monotone projected gradient. *Comput. Optim. Appl.* **83**(2), 693–721 (2022)
15. Dolan, E.D., Moré, J.J.: Benchmarking optimization software with performance profiles. *Math. Program.* **91**(2), 201–213 (2002)
16. Ferreira, O.P., Lemes, M., Prudente, L.F.: On the inexact scaled gradient projection method. *Comput. Optim. Appl.* **81**(1), 91–125 (2022)
17. Golbabaee, M., Davies, M.E.: Inexact gradient projection and fast data driven compressed sensing. *IEEE Trans. Inf. Theory* **64**(10), 6707–6721 (2018)
18. Gonçalves, D.S., Gonçalves, M.L., Menezes, T.C.: Inexact variable metric method for convex-constrained optimization problems. *Optimization* **71**(1), 145–163 (2022)
19. Gould, N., Orban, D., Toint, P.: CUTEst: a Constrained and Unconstrained Testing Environment with safe threads for mathematical optimization. *Computational Optimization and Applications* **60** (2014). <https://doi.org/10.1007/s10589-014-9687-3>
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). <https://doi.org/10.1137/0723046>
21. Grippo, L., Sciandrone, M.: *Introduction to methods for nonlinear optimization*. Unitext. Springer, Cham (2023)
22. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference and Prediction*, (2009). (2nd edn.)

23. Jia, X., Kanzow, C., Mehlitz, P., Wachsmuth, G.: An augmented Lagrangian method for optimization problems with structured geometric constraints. *Math. Program.* **199**(1), 1365–1415 (2023)
24. Lapucci, M., Liuzzi, G., Lucidi, S., Pucci, D., Sciandrone, M.: A globally convergent gradient method with momentum. *Comput. Optim. Appl.* , 1–26 (2025)
25. Lee, C.P., Wang, P.W., Lin, C.J.: Limited-memory common-directions method for large-scale optimization: convergence, parallelization, and distributed optimization. *Math. Program. Comput.* **14**(3), 543–591 (2022)
26. Lee, S.I., Lee, H., Abbeel, P., Ng, A.Y.: Efficient L1 regularized logistic regression. In: *AAAI*, vol. 6, pp. 401–408 (2006)
27. Levitin, E.S., Polyak, B.T.: Constrained minimization methods. *USSR Comput. Math. Math. Phys.* **6**(5), 1–50 (1966)
28. Liu, J., Chen, J., Ye, J.: Large-scale sparse logistic regression. In: *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 547–556. (2009)
29. Polyak, B.: Some methods of speeding up the convergence of iteration methods. *USSR Comput. Math. Math. Phys.* **4**(5), 1–17 (1964)
30. Raydan, M.: The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem. *SIAM J. Optim.* **7**(1), 26–33 (1997). <https://doi.org/10.1137/S1052623494266365>
31. Tang, T., Toh, K.C., Xiao, N., Ye, Y.: A Riemannian dimension-reduced second-order method with application in sensor network localization. *SIAM J. Sci. Comput.* **46**(3), A2025–A2046 (2024). <https://doi.org/10.1137/23M1567229>
32. Tao, W., Wu, G.W., Tao, Q.: Momentum acceleration in the individual convergence of nonsmooth convex optimization with constraints. *IEEE Transactions on Neural Networks and Learning Systems* **33**(3), 1107–1118 (2022)
33. Wright, S.J., Recht, B.: *Optimization for Data Analysis*, Cambridge University Press (2022)

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