PROXIMAL GRADIENT METHODS BEYOND MONOTONY

Alberto De Marchi*

Abstract   We address composite optimization problems, which consist in minimizing the sum of a smooth and a merely lower semicontinuous function, without any convexity assumptions. Numerical solutions of these problems can be obtained by proximal gradient methods, which often rely on a line search procedure as globalization mechanism. We consider an adaptive nonmonotone proximal gradient scheme based on an averaged merit function and establish asymptotic convergence guarantees under weak assumptions, delivering results on par with the monotone strategy. Global worst-case rates for the iterates and a stationarity measure are also derived. Finally, a numerical example indicates the potential of nonmonotonicity and spectral approximations.

Keywords   Nonsmooth nonconvex optimization · Proximal algorithms · Gradient methods · Spectral stepsize · Nonmonotone linesearch.

AMS MSC   49J52 · 65K05 · 90C06.

1 INTRODUCTION

In this paper we consider the classical problem of minimizing the sum of a smooth function f and a nonsmooth function g. More precisely, we address optimization problems of the form

$$\text{(P)} \qquad \underset{x \in \mathbb{X}}{\text{minimize}} \; \varphi(x), \qquad \text{where} \quad \varphi := f + g,$$

$f \colon \mathbb{X} \to \mathbb{R}$ is differentiable, $g \colon \mathbb{X} \to \overline{\mathbb{R}} := \mathbb{R} \cup \{\infty\}$ admits an easily computable proximal mapping, and $\mathbb{X}$ denotes an Euclidean space, i.e., a real and finite-dimensional Hilbert space. The objective φ and its components f and g are allowed to be nonconvex.

The problem class (P) has been widely investigated, especially in the convex setting, and a number of splitting methods have been devised for its numerical solution [9, 23, 4, 28]. By exploiting its composite structure, proximal gradient (PG) methods offer a simple and versatile iterative technique for addressing (P). Introduced in [14], PG schemes are also known as forward-backward splittings, due to the close relation with the theory of monotone operators [3] and their equivalence in the convex setting [23, §3]. Combining a step in the negative gradient direction with a proximal point update, PG methods can cope with nonconvexity but require a mechanism to guarantee global convergence, that is, convergence from arbitrary initial points. Widely adopted since Armijo's work [1] in smooth unconstrained optimization, a (backtracking) line search procedure provides such globalization, if equipped with an appropriate criterion for validating (and therefore accepting) tentative iterates. Monotone strategies ensures that the objective value decreases *at each iteration*, but possibly at the cost of taking utterly small steps. In contrast, nonmonotone strategies adopt relaxed conditions for a tentative update to be accepted; this approach often proves beneficial in practice, as it can reduce conservatism and allow larger steps and faster convergence [15, 31, 28]. Similarly to watchdog techniques, nonmonotone line search procedures monitor progress and ensure that the objective value decreases, *at least eventually*, along the iterates generated by the algorithm.

*University of the Bundeswehr Munich, Department of Aerospace Engineering, Institute of Applied Mathematics and Scientific Computing, 85577 Neubiberg, Germany. alberto.demarchi@unibw.de, ORCID: 0000-0002-3545-6898.

We are interested in PG methods for (P) with provable well-definedness and convergence guarantees under minimal assumptions on the problem data. These properties have been recently established by Kanzow and Mehlitz [17] considering an adaptive *monotone* PG method. Therein, however, the same guarantees for a *nonmonotone* variant required stronger conditions, and more involved proofs. In particular, accumulation points are shown to be stationary for (P) under some continuity assumptions on φ and g [17, Ass. 4.1].

1.1 RELATED WORK

In their seminal work [15] in the context of unconstrained smooth optimization, Grippo, Lampariello and Lucidi (GLL) suggested a nonmonotone line search technique based on monitoring the *maximum* objective value attained by the latest iterates. Later Zhang and Hager [31] proposed a different flavor of nonmonotonicity, which takes a weighted *average* of the objective value at all iterates. Anticipating the notation of Algorithm 3.1 below, a certain iterate $x^{k+1} \in \operatorname{dom} g$ is deemed acceptable upon comparison of the objective $\varphi(x^{k+1})$ with a threshold based on the merit Φ_k, whose value is given by

$$(1.1) \qquad \max_{j=0,\ldots,\min\{M,k\}} \varphi(x^{k-j}) \qquad \text{and} \qquad (1-p)\Phi_{k-1} + p\varphi(x^k)$$

for the max and average flavors, respectively. Here, parameters $M \in \mathbb{N}$ and $p \in (0,1]$ control the level of nonmonotonicity allowed: for larger M, or smaller p, the merit Φ_k can attain larger values, hence imposing a weaker condition for the acceptance of x^{k+1}. Conversely, the values $M = 0$ and $p = 1$ result in monotone behavior, as the merit Φ_k coincides with the objective $\varphi(x^k)$. We notice that, in general, neither of these approaches is more conservative than the other. The GLL merit function has been adopted, e.g., in [24, 6] and [30] for smooth and composite optimization, respectively, under some convexity assumptions, before being considered in [17]. In the context of (fully nonconvex) composite optimization, a merit function based on averaging has been adopted, e.g., in [20] for safeguarding Nesterov-type accelerated PG methods as well as in [28] and [13, §4.3] for globalizing quasi-Newton-type schemes.

Nonmonotone techniques had a significant impact also on the development and practice of spectral gradient methods. Long after the classical steepest descent method was proposed by Cauchy [8] in 1847, Barzilai and Borwein [2] presented a simple and efficient tool for solving large-scale problems: a gradient method with step sizes motivated by Newton's method but not involving any Hessian matrix. The Barzilai-Borwein (BB) method uses the same search direction as for Cauchy's method, but its stepsize rule is different and, at nearly no extra cost, it often outperforms standard gradient methods [6, 7]. However, even when the objective function is strongly convex, it may not yield a monotone decrease in the function value and, in fact, it may not converge at all. Therefore, Raydan [24] suggested to control this behavior by combining the spectral stepsize with a nonmonotone line search as a safeguard for globalization. This, in particular, leads to global convergence for smooth unconstrained optimization problems and constrained on sets [6, 16].

The broad class (P) includes constrained optimization instances, namely problems involving the minimization of a smooth cost $f(x)$ subject to the constraint $x \in X$, with $X \subseteq \mathbb{X}$ some nonempty and closed set. In fact, it suffices to take $g \equiv \iota_X$ in (P), with $\iota_X \colon \mathbb{X} \to \overline{\mathbb{R}}$ denoting the indicator function of X, defined by $\iota_X(x) = 0$ if $x \in X$ and $\iota_X(x) = \infty$ otherwise. Therefore, if the projection operator associated to X (coinciding with the proximal mapping of ι_X) can be efficiently evaluated, PG methods are a viable option for dealing with such (potentially nonconvex) constrained problems. From the literature, one can observe the (temporal) progression from BB methods [2, 24] for smooth unconstrained problems, via spectral projected gradient (SPG) methods [6, 7] for smooth objectives over convex constraint sets and [16, §3] for nonconvex constraint sets, to [17] for nonconvex composite objectives, all using the nonmonotone GLL line search for globalization. Preferring a different kind of nonmonotonicity, the

methods we discuss in the following allow to tackle such broad class of problems, supported by the same theoretical guarantees and under weaker assumptions.

Finally, we shall remark that these methods often play the role of inner solvers in general purpose optimization packages, where subproblems of the form (P) have to be (approximately) solved. This is the case of SPG [6] within the augmented Lagrangian solver Algencan [5], [16, Alg. 3.1] within [16, Alg. 4.1], PANOC$^{(+)}$ [27, 13] within OpEn [26], ALS [11], ALPS [10], and IP-FB within an interior point method [12]. Thus, our findings have significant impact that possibly widens scope and aids practical performance of other methods and solvers too.

1.2 CONTRIBUTION

Intrigued by the theoretical gap emerged in [17] and motivated by the potential benefits of nonmonotone strategies in practice, we consider proximal gradient methods for addressing (P) in the fully nonconvex setting and focus on the role played by the line search procedure. Why does the nonmonotone strategy lead to the difficulties highlighted in [17, §4]? Is it possible to retain the simplicity and convergence guarantees of these (adaptive) methods while exploiting *nonmonotone* behaviour and minimal assumptions?

As we are going to reveal, nonmonotonicity is not responsible for the theoretical issues encountered, but it is its specific flavor that may be, in some sense, too weak. Instead of the GLL max-type condition adopted in [17], we study an adaptive PG method with a nonmonotone line search based on an *averaging* merit function [31, 28]. By establishing that every iteration yields some sufficient decrease for this merit function, we show that the convergence properties (and analysis) of the monotone setting naturally carry over to the nonmonotone one.

The key contributions of this paper are as follows.

- We present an adaptive proximal gradient method with nonmonotone line search and establish asymptotic convergence properties and guarantees under weak assumptions on the problem data (cf. Theorem 4.6). Although in general not more conservative than the GLL strategy, the averaging nonmonotonicity ensures sufficient decrease at every iteration (cf. Lemma 4.2), enabling a seamless extension from monotone to nonmonotone line search without additional assumptions (cf. Lemma 4.4).

- Under local Lipschitz continuity of ∇f, we provide global worst-case rate-of-convergence results. In particular, we show (at least) sublinear convergence of order $O(1/\sqrt{k})$ for the iterates $\{x^k\}_{k\in\mathbb{N}}$ (cf. Corollary 4.8) and, along converging subsequences, for the fixed-point residual as stationarity measure (cf. Theorem 4.9). Furthermore, assuming boundedness for the set of accumulation points of $\{x^k\}_{k\in\mathbb{N}}$, the result is extended to the whole sequence (cf. Corollary 4.11).

OUTLINE

In the following Section 2 we provide preliminary notions and optimality concepts. In Section 3 we detail and discuss the algorithm, whose convergence properties are investigated in Section 4 under appropriate conditions. We report on comparative numerical tests in Section 5 and conclude with some remarks in Section 6.

2   PRELIMINARIES

With $\mathbb{N}$, $\mathbb{R}$ and $\overline{\mathbb{R}} := \mathbb{R} \cup \{\infty\}$ we denote the natural numbers, real and extended-real line, respectively. Throughout the paper, the Euclidean space $\mathbb{X}$ will be equipped with the inner product $\langle \cdot, \cdot \rangle \colon \mathbb{X} \times \mathbb{X} \to \mathbb{R}$ and the associated norm $\|\cdot\|$. Given a nonempty set $E \subset \mathbb{X}$ and a point $p \in \mathbb{X}$, $\operatorname{dist}(p, E) := \inf_{x \in E} \|x - p\|$

denotes the distance of p from E. The sum of two sets $A, B \subseteq \mathbb{X}$ is meant in the sense of Minkowski, namely $A + B = B + A := \{a + b \mid a \in A, b \in B\}$. In case of a singleton $A = \{a\}$, we write $a + B$ as shorthand for $\{a\} + B$.

The continuous linear operator $f'(x)\colon \mathbb{X} \to \mathbb{R}$ denotes the derivative of f at $x \in \mathbb{X}$, and we will make use of $\nabla f(x) := f'(x)^*1$ where $f'(x)^*\colon \mathbb{R} \to \mathbb{X}$ is the adjoint of $f'(x)$. This way, ∇f is a mapping from $\mathbb{X}$ to $\mathbb{X}$ and for any $d \in \mathbb{X}$ it holds $f'(x)d = \langle \nabla f(x), d \rangle$.

The following concepts are standard in variational analysis [25, 21]. The effective *domain* of an extended-real-valued function $h\colon \mathbb{X} \to \overline{\mathbb{R}}$ is denoted by $\operatorname{dom} h := \{x \in \mathbb{X} \mid h(x) < \infty\}$. We say that h is *proper* if $\operatorname{dom} h \neq \emptyset$ and *lower semicontinuous* (lsc) if $h(\bar{x}) \leq \liminf_{x \to \bar{x}} h(x)$ for all $\bar{x} \in \mathbb{X}$. Following [25, Def. 8.3], we denote by $\hat{\partial} h\colon \mathbb{X} \rightrightarrows \mathbb{X}$ the *regular* (or *Fréchet*) *subdifferential* of h, defined for $\bar{x} \in \operatorname{dom} h$ by

$$\hat{\partial} h(\bar{x}) := \left\{ v \in \mathbb{X} \;\middle|\; \liminf_{\substack{x \to \bar{x} \\ x \neq \bar{x}}} \frac{h(x) - h(\bar{x}) - \langle v, x - \bar{x} \rangle}{\|x - \bar{x}\|} \geq 0 \right\}$$

and $\hat{\partial} h(\bar{x}) := \emptyset$ for $\bar{x} \notin \operatorname{dom} h$. Then, the *limiting* (or *Mordukhovich*) *subdifferential* of h is $\partial h\colon \mathbb{X} \rightrightarrows \mathbb{X}$, where $v \in \partial h(\bar{x})$ if and only if $\bar{x} \in \operatorname{dom} h$ and there exists a sequence $\{x^k, v^k\}_{k \in \mathbb{N}}$ such that $x^k \in \mathbb{X}$ and $v^k \in \hat{\partial} h(x^k)$ for all $k \in \mathbb{N}$ and $x^k \to \bar{x}$, $h(x^k) \to h(\bar{x})$, $v^k \to v$ as $k \to \infty$. Notice that, by construction, it is always $\hat{\partial} h(x) \subset \partial h(x)$ and, whenever h is convex, equality holds. The subdifferential of h at $\bar{x} \in \operatorname{dom} h$ satisfies $\partial (h + h_0)(\bar{x}) = \partial h(\bar{x}) + \nabla h_0(\bar{x})$ for any $h_0\colon \mathbb{X} \to \overline{\mathbb{R}}$ continuously differentaible around $\bar{x}$ [25, Ex. 8.8(c)], [21, Prop. 1.30(ii)].

The *proximal* mapping of $h\colon \mathbb{X} \to \overline{\mathbb{R}}$ with stepsize $\gamma > 0$ is defined by

$$(2.1) \qquad \operatorname{prox}_{\gamma h}(x) := \operatorname*{arg\,min}_{z \in \mathbb{X}} \left\{ h(z) + \tfrac{1}{2\gamma} \|z - x\|^2 \right\}.$$

We say that h is *prox-bounded* if it is proper and $h + \frac{1}{2\gamma}\|\cdot\|^2$ is bounded below on $\mathbb{X}$ for some $\gamma > 0$. The supremum of all such γ is the threshold γ_h of prox-boundedness for h. In particular, if h is bounded below by an affine function, then $\gamma_h = \infty$. When h is lsc, for any $\gamma \in (0, \gamma_h)$ the proximal mapping $\operatorname{prox}_{\gamma h}$ is locally bounded, nonempty- and compact-valued [25, Thm 1.25].

2.1 STATIONARITY CONCEPTS

Considering the minimization of a proper function $h\colon \mathbb{X} \to \overline{\mathbb{R}}$ and the definition of regular subdifferential, we have that the inclusion $0 \in \hat{\partial} h(x^*)$ is a necessary condition for $x^* \in \operatorname{dom} h$ to be a (local) minimizer of h [25, Thm. 10.1], [21, Prop. 1.30(i)]. We say a point $x^* \in \operatorname{dom} h$ is *M-stationary* (for h) if it satisfies (the potentially weaker condition) $0 \in \partial h(x^*)$. This denomination is due to the appearance of the limiting (or *Mordukhovich*) subdifferential. Moreover, a point $x^* \in \operatorname{dom} h$ is said to be ε-*M-stationary* for some tolerance $\varepsilon \geq 0$ (or simply *approximate M-stationary*) if $\operatorname{dist}(0, \partial h(x^*)) \leq \varepsilon$. When $\varepsilon = 0$, this notion reduces to M-stationarity, as it recovers the condition $0 \in \partial h(x^*)$ by closedness of $\partial h(x^*)$ [25, Thm. 8.6]. Notice that whenever $x^* \in \mathbb{X}$ is an (approximate) M-stationary point relative to h, it necessarily belongs to its domain $\operatorname{dom} h$, for otherwise the subdifferential $\partial h(x^*)$ would be empty.

Relative to (P), ε-M-stationarity of a point $x^* \in \operatorname{dom} \varphi$ means that it satisfies $\operatorname{dist}(0, \partial \varphi(x^*)) \leq \varepsilon$. Under mild assumptions (see our blanket Assumptions (A1) to (A3) given in Section 4), this condition turns into

$$(2.2) \qquad \operatorname{dist}(-\nabla f(x^*), \partial g(x^*)) \leq \varepsilon$$

because of the identity $\partial \varphi = \nabla f + \partial g$.

Finally, let us consider the proximal mapping $\text{prox}_{\gamma h}$ defined in (2.1) and, for any given $\gamma > 0$, the proximal-gradient mapping $x \mapsto \text{prox}_{\gamma g}(x - \gamma \nabla f(x))$ associated to (P). Then, for any $x \in \mathbb{X}$, the minimizing property of $\bar{x} \in \text{prox}_{\gamma g}(x - \gamma \nabla f(x))$ implies that $\bar{x} \in \text{dom}\, g$ and

$$(2.3) \qquad\qquad 0 \in \nabla f(x) + \frac{1}{\gamma}(\bar{x} - x) + \hat{\partial} g(\bar{x}),$$

owing to the relevant necessary conditions.

3 METHODS

In this section we detail and elaborate upon an adaptive nonmonotone proximal gradient scheme for solving (P), whose precise statement is given in Algorithm 3.1. An investigation of its well-definedness and convergence properties under some basic assumptions is postponed to the next Section 4.

Given a starting point x^0, the recurrence of proximal gradient updates executed at Step 5 form the core of Algorithm 3.1, whose *adaptivity* stems from the backtracking line search on the "proximal" stepsize γ_k at Step 9. Analogously to [17, Alg. 4.1], [13, Alg. 2], this line search on the stepsize γ_k obviates the need for a priori knowledge (or existence, at all) of a Lipschitz constant for ∇f and provides a globalization mechanism. Whenever the respective needed conditions at Step 9 are violated, the stepsize is reduced and the iteration restarted. With the particular choice $p_k \equiv 1$ for all $k \in \mathbb{N}$, the merit Φ_k monitoring progress coincides with the objective $\varphi(x^k)$ for all $k \in \mathbb{N}$, thus requiring the objective value to monotonically decrease along the iterates due to Step 9. Instead, the sequence $\{\varphi(x^k)\}_{k \in \mathbb{N}}$ can exhibit some degree of *nonmonotonicity* whenever $p_k < 1$, despite the monotone decrease of the monitoring sequence $\{\Phi_k\}_{k \in \mathbb{N}}$ generated by the *averaging* at Step 11.

Algorithm 3.1: Proximal Gradient Method for (P)

Data: starting point $x^0 \in \text{dom}\, g$, termination tolerance $\varepsilon > 0$
Result: ε-stationary point x^* for (P)

1 select $0 < \gamma_{\min} \leq \gamma_{\max} < \infty$, $0 < \alpha_{\min} \leq \alpha_{\max} < 1$, $0 < \beta_{\min} \leq \beta_{\max} < 1$, $p_{\min} \in (0, 1]$
2 set $\Phi_0 \leftarrow \varphi(x^0)$
3 **for** $k = 1, 2, \ldots$ **do**
4 $\quad$ choose $\gamma_k \in [\gamma_{\min}, \gamma_{\max}]$
5 $\quad$ compute $x^k \in \text{prox}_{\gamma_k g}(x^{k-1} - \gamma_k \nabla f(x^{k-1}))$
6 $\quad$ **if** $\|\frac{1}{\gamma_k}(x^k - x^{k-1}) - \nabla f(x^k) + \nabla f(x^{k-1})\| \leq \varepsilon$ **then**
7 $\quad\quad$ **return** $x^* \leftarrow x^k$
8 $\quad$ choose $\alpha_k \in [\alpha_{\min}, \alpha_{\max}]$ and $\beta_k \in [\beta_{\min}, \beta_{\max}]$
9 $\quad$ **if** $\varphi(x^k) > \Phi_{k-1} - \frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2$ **then**
10 $\quad\quad$ set $\gamma_k \leftarrow \beta_k \gamma_k$ and go back to Step 5
11 $\quad$ choose $p_k \in [p_{\min}, 1]$ and set $\Phi_k \leftarrow (1 - p_k)\Phi_{k-1} + p_k \varphi(x^k)$

The procedure relies on many algorithmic parameters γ_k, α_k, β_k and p_k, which can be selected anew (possibly several times) at every iteration k, at Steps 4, 8 and 11. These degrees of freedom provide significant flexibility in controlling the behavior of Algorithm 3.1 and the iterates it generates, and possibly in steering and speeding up the convergence.

It should be highlighted that it is not restrictive to require a feasible starting point for Algorithm 3.1, namely a point $x^0 \in \text{dom}\, g$. Conversely, it simplifies exposition and discussion of the convergence analysis in the following Section 4. That $x^0 \in \text{dom}\, g$ comes without loss of generality follows from

the observation that starting with some $x^0 \notin \mathrm{dom}\, g$ would result in a trivial iteration. More precisely, Step 5 would generate some $x^1 \in \mathrm{dom}\, g$ that is accepted by the line search at Step 9 without need for backtracking, since $\varphi(x^1) < \Phi_0 = \varphi(x^0) = \infty$. However, in order to avoid the infinite-valued merit Φ_0 spoiling the line search condition (due to Step 11, it would be $\Phi_k = \infty$ for all $k \in \mathbb{N}$), the procedure should be restarted with the newly set $x^0 \leftarrow x^1$.

3.1 TERMINATION CRITERIA

Step 6 of Algorithm 3.1 entails the termination criterion, which is designed to detect approximate stationary points for (P). Rearranging the necessary condition (2.3) associated to the proximal mapping evaluated at Step 5, we observe that the iterates satisfy

$$\frac{1}{\gamma_k}(x^{k-1} - x^k) - \nabla f(x^{k-1}) + \nabla f(x^k) \in \nabla f(x^k) + \hat{\partial} g(x^k) = \hat{\partial}\varphi(x^k)$$

for all $k \in \mathbb{N}$. This inclusion justifies the termination condition at Step 6 since

$$(3.1) \qquad \mathrm{dist}(0, \partial\varphi(x^k)) \leq \mathrm{dist}(0, \hat{\partial}\varphi(x^k)) \leq \left\| \frac{1}{\gamma_k}(x^{k-1} - x^k) - \nabla f(x^{k-1}) + \nabla f(x^k) \right\| \leq \varepsilon$$

encodes approximate stationarity of x^k for (P). Notice that the first inequality in (3.1) follows from the inclusion $\hat{\partial}\varphi(x^k) \subset \partial\varphi(x^k)$. As established in Theorem 4.7 below, for any given $\varepsilon > 0$ the condition at Step 6 is satisfied for some $k \in \mathbb{N}$ large enough, thus allowing the algorithm to return with the corresponding ε-M-stationary point $x^k \in \mathrm{dom}\,\varphi$.

Finally, we shall point out that the termination condition is checked within the backtracking loop, as opposed to after the line search procedure; analogously to [16, Alg. 3.1] and [17, Alg. 4.1], respectively. Although potentially costly, this allows to avoid a potential infinite loop at Step 9, thus making Algorithm 3.1 well-defined. If the problem data in (P) satisfy stronger regularity properties (e.g., ∇f is locally Lipschitz continuous, cf. Lemma 4.1), it may be possible to safely move the termination checking out of the backtracking procedure, thereby requiring fewer function evaluations.

4 CONVERGENCE ANALYSIS

In this section we analyze the properties of the iterates generated by Algorithm 3.1. The following blanket assumptions on (P) are considered throughout.

(A1) $f \colon \mathbb{X} \to \mathbb{R}$ is continuously differentiable;

(A2) $g \colon \mathbb{X} \to \overline{\mathbb{R}}$ is proper, lsc and prox-bounded;

(A3) $\inf \varphi \in \mathbb{R}$.

Notice that Assumption (A3) and properness and lsc of g in Assumption (A2) are relevant for the well-definedness of (P), as they imply that $\varphi := f + g$ is lsc and bounded from below on $\mathrm{dom}\,\varphi = \mathrm{dom}\, g \neq \emptyset$. Assumption (A1) and prox-boundedness of g in Assumption (A2) guarantee that the proximal mapping evaluations and the algorithm are well-defined, as we are about to show. Furthermore, since the iterates generated by Algorithm 3.1 satisfy $\{x^k\}_{k \in \mathbb{N}} \subseteq \mathrm{dom}\, g$, the smooth function f could be defined just on the domain of g.

The first basic result establishes that the algorithm is well-defined, namely that the inner (line search) loop requires finitely many backtrackings. In Lemma 4.1(i)–(ii) the finite termination of the line search procedure is due to the tentative update yielding sufficient decrease according to the condition at Step 9,

whereas in Lemma 4.1(iii) it follows from Step 6 detecting that the current iterate is ε-M-stationary for (P).

> **Lemma 4.1 (Well-definedness).** *Suppose that Assumptions (A1) to (A3) are satisfied. Consider the k-th iteration of Algorithm 3.1 and assume that at least one of the following conditions holds:*
>
>   (i) *x^{k-1} is not M-stationary for (P);*
>
>   (ii) *∇f is locally Lipschitz continuous (at least in a neighborhood of x^{k-1});*
>
>   (iii) *$\varepsilon > 0$.*
>
> *Then, the iteration terminates, and in particular Step 9 is passed in finitely many backtrackings.*

Proof. The line of proof builds upon [17, Lemma 3.1], relevant for the monotone counterpart of Algorithm 3.1, and the properties of the monitoring sequence $\{\Phi_k\}_{k\in\mathbb{N}}$ with respect to $\{\varphi(x^k)\}_{k\in\mathbb{N}}$.

(i) The claim follows immediately by induction since $\Phi_0 := \varphi(x^0)$ and, for $k \geq 1$, the conditions at Steps 9 and 11 imply

$$(4.1) \qquad \Phi_k := (1-p_k)\Phi_{k-1} + p_k\varphi(x^k) \geq (1-p_k)\left(\varphi(x^k) + \frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2\right) + p_k\varphi(x^k)$$

$$= \varphi(x^k) + (1-p_k)\frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2 \geq \varphi(x^k),$$

proving that the nonmonotone line search is not more conservative than the monotone one (which has $p_k = 1$ and $\Phi_k = \varphi(x^k)$ for all $k \in \mathbb{N}$). Thus, relying on [17, Lemma 3.1], the condition at Step 9 is violated in finitely many steps.

(ii) Since x^{k-1} remains untouched during the k-th iteration, γ_k is upper bounded by $\gamma_{\max}$, and $\operatorname{prox}_{\gamma_k g}: \mathbb{X} \rightrightarrows \mathbb{X}$ is locally bounded, we have that the (tentative) updates x^k at Step 5 remain in a bounded set [13, Lemma 4.1]; let L_k be the (local) Lipschitz constant of ∇f on this set. Therefore, for any given $\alpha_k \in (0,1)$ the quadratic upper bound

$$(4.2) \qquad f(x^k) \leq f(x^{k-1}) + \langle\nabla f(x^{k-1}), x^k - x^{k-1}\rangle + \frac{\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2$$

holds after finitely many backtrackings, at most when $\gamma_k \leq \alpha_k/L_k$, by [4, Lemma 5.7]. Noticing that the minimizing property of x^k at Step 5 gives

$$g(x^k) + \langle\nabla f(x^{k-1}), x^k - x^{k-1}\rangle + \frac{1}{2\gamma_k}\|x^k - x^{k-1}\|^2 \leq g(x^{k-1}),$$

adding terms with (4.2) leads to

$$\varphi(x^k) \leq \varphi(x^{k-1}) - \frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2 \leq \Phi_{k-1} - \frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2,$$

where the last inequality follows from the upper bound (4.1). This proves that the condition at Step 9 is violated and the inner loop terminates.

(iii) In view of Lemma 4.1(i), let us assume that x^{k-1} is M-stationary. Then, by Assumption (A1) and $\varepsilon > 0$, the termination condition at Step 6 is satisfied after finitely many attempts owing to the arguments of [17, Lemma 3.1]. □

We already mentioned that $\{\varphi(x^k)\}_{k\in\mathbb{N}}$ may possess a nonmonotone behavior, whereas the sequence $\{\Phi_k\}_{k\in\mathbb{N}}$ monitoring (and forcing) progress along the iterates is monotonically decreasing. As a corollary, the iterates of Algorithm 3.1 belong to the sublevel set of φ induced by the starting point $x^0 \in \operatorname{dom} g$. Furthermore, we establish that $\{\Phi_k\}_{k\in\mathbb{N}}$ exhibits some sufficient decrease at every iteration. This feature appears to provide the foundation for proving convergence under weak assumptions also in the nonmonotone setting, thus bridging the gap emerged in [17].

Lemma 4.2 (Descent behavior). *Suppose that Assumptions (A1) to (A3) are satisfied. The following hold for the sequence of iterates $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1:*

(i) *The sequence $\{\Phi_k\}_{k\in\mathbb{N}}$ is monotonically decreasing, lower bounded by $\inf \varphi$ and, for every $k \geq 1$, $k \in \mathbb{N}$, one has*

$$(4.3) \qquad \varphi(x^k) + (1-p_k)\delta_k \leq \Phi_k \leq \Phi_{k-1} - p_k\delta_k \qquad \text{where} \quad \delta_k := \frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2.$$

(ii) *Every iterate x^k remains in the sublevel set $\{x \in \mathbb{X} \mid \varphi(x) \leq \varphi(x^0)\} \subseteq \operatorname{dom} g$.*

Proof. Based on Lemma 4.1, we may assume the iterates of Algorithm 3.1 are well-defined.

(i) By the update rule at Step 11 and the line search condition at Step 9, we have that

$$\Phi_k := (1-p_k)\Phi_{k-1} + p_k\varphi(x^k)$$
$$\leq (1-p_k)\Phi_{k-1} + p_k\left(\Phi_{k-1} - \frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2\right)$$
$$= \Phi_{k-1} - p_k\frac{1-\alpha_k}{2\gamma_k}\|x^k - x^{k-1}\|^2.$$

The lower bound on Φ_k follows directly from Assumption (A3) and the inequalities in (4.1).

(ii) Owing to $x^0 \in \operatorname{dom} g$ and continuity of f, the sublevel set of φ induced by x^0 is a subset of $\operatorname{dom} g$. Then, since $p_k \in (0,1]$ and $\delta_k \geq 0$, it follows from $\Phi_0 := \varphi(x^0)$ and the inequality in (4.3) that $\varphi(x^k) \leq \Phi_k \leq \ldots \leq \Phi_0 = \varphi(x^0) < \infty$, whence the inclusion of x^k in the sublevel set. □

We next consider an asymptotic analysis of the algorithm. The following Lemma 4.3(iii) is the counterpart of [17, Prop 3.1] in the nonmonotone setting. Contrarily to [17, Prop 4.1], it is not more difficult to prove, thanks to the sufficient decrease on the averaging merit function; cf. (4.3).

Lemma 4.3. *Suppose that Assumptions (A1) to (A3) are satisfied. The following hold for the sequence of iterates $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1:*

(i) *$\{\varphi(x^k)\}_{k\in\mathbb{N}}$ and $\{\Phi_k\}_{k\in\mathbb{N}}$ converge to a finite value $\varphi_\star \geq \inf \varphi$; the latter from above.*

(ii) *$\sum_{k\in\mathbb{N}} \frac{1}{\gamma_k}\|x^k - x^{k-1}\|^2 < \infty$.*

(iii) *$\lim_{k\to\infty} \|x^k - x^{k-1}\| = 0$.*

Proof. We implicitly assume that Algorithm 3.1 generates an infinite sequence $\{x^k\}_{k\in\mathbb{N}}$.

(i) The assertion regarding $\{\Phi_k\}_{k\in\mathbb{N}}$ follows from (4.3) and Assumption (A3). Then, consider the update rule at Step 11, which gives $\varphi(x^k) = \frac{1}{p_k}(\Phi_k - \Phi_{k-1}) + \Phi_{k-1}$. Taking the limit for $k \to \infty$ and noticing that $p_k \geq p_{\min} > 0$ for all $k \in \mathbb{N}$, the convergence of $\{\Phi_k\}_{k\in\mathbb{N}}$ implies that of $\{\varphi(x^k)\}_{k\in\mathbb{N}}$ to the same finite value $\varphi_\star$.

(ii) A telescoping argument on (4.3), together with (i), yields

$$(4.4) \qquad \Phi_0 - \varphi_\star \geq \sum_{j=1}^k \Phi_{j-1} - \Phi_j \geq \sum_{j=1}^k p_j \delta_j \geq \sum_{j=1}^k p_{\min} \frac{1 - \alpha_{\max}}{2\gamma_j} \|x^j - x^{j-1}\|^2,$$

since $p_j \geq p_{\min}$ and $\alpha_j \leq \alpha_{\max}$ for all $j \geq 1$. The claimed finite sum follows from $p_{\min} > 0$, $\alpha_{\max} < 1$, and the independence of the finite upper bound $\Phi_0 - \varphi_\star$ on $k \in \mathbb{N}$.

(iii) Follows from (ii), since $\{\gamma_k\}_{k \in \mathbb{N}}$ is upper bounded by $\gamma_{\max}$. □

The following result is the nonmonotone counterpart of [17, Prop 3.2] and provides a key ingredient in relation to (the detection of) M-stationarity; see the termination condition at Step 6. In stark contrast with [17, Thm 4.1, Ass. 4.1], it does not require additional continuity assumptions on φ nor g, and its proof closely patterns that of [17, Prop 3.2], with minor modifications to account for the nonmonotone line search.

Lemma 4.4. *Suppose that Assumptions (A1) to (A3) are satisfied. Consider a sequence $\{x^k\}_{k \in \mathbb{N}}$ generated by Algorithm 3.1. Let $x^\star$ be an accumulation point of $\{x^k\}_{k \in \mathbb{N}}$ and $\{x^k\}_{k \in K}$ a subsequence such that $x^k \to_K x^\star$. Then, it is $\frac{1}{\gamma_k} \|x^k - x^{k-1}\| \to_K 0$.*

Proof. If the subsequence $\{\gamma_k\}_{k \in K}$ is bounded away from zero, the statement follows immediately from Lemma 4.3(iii). The remaining part of this proof therefore assumes, without loss of generality, that $\gamma_k \to_K 0$. Closely following the proof of [17, Prop. 3.2], let $\hat{\gamma}_k$ denote the attempted stepsize preceding the accepted stepsize γ_k; these values are well-defined as the inner loop terminates by Lemma 4.1. Owing to Step 10, it is $\gamma_k/\beta_{\max} \leq \hat{\gamma}_k \leq \gamma_k/\beta_{\min}$ for each $k \in \mathbb{N}$ and so, by $\gamma_k \to_K 0$, we also have $\hat{\gamma}_k \to_K 0$. However, the corresponding vector $\hat{x}^k$ does not satisfy the stepsize condition from Step 9 (for otherwise $\hat{\gamma}_k$ would have been accepted) with the associated parameter $\hat{\alpha}_k \in [\alpha_{\min}, \alpha_{\max}]$, namely it is

$$(4.5) \qquad \varphi(\hat{x}^k) > \Phi_{k-1} - \frac{1 - \hat{\alpha}_k}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 \geq \Phi_{k-1} - \frac{1 - \alpha_{\min}}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2.$$

On the other hand, since $\hat{x}^k$ solves the corresponding subproblem at Step 5, we have

$$(4.6) \qquad \langle \nabla f(x^{k-1}), \hat{x}^k - x^{k-1} \rangle + \frac{1}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 + g(\hat{x}^k) \leq g(x^{k-1}).$$

Then, using the Cauchy-Schwarz inequality and Lemma 4.2(ii), we obtain

$$g(\hat{x}^k) + \frac{1}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 \leq \|\nabla f(x^{k-1})\| \|\hat{x}^k - x^{k-1}\| + g(x^{k-1})$$

$$= \|\nabla f(x^{k-1})\| \|\hat{x}^k - x^{k-1}\| + \varphi(x^{k-1}) - f(x^{k-1})$$

$$\leq \|\nabla f(x^{k-1})\| \|\hat{x}^k - x^{k-1}\| + \Phi_0 - f(x^{k-1}).$$

By continuous differentiability of f and prox-boundedness of g, if $\{\|\hat{x}^k - x^{k-1}\|\}_{k \in K}$ is unbounded, then the left-hand side grows more rapidly than the right-hand side. Analogously, if $\{\|\hat{x}^k - x^{k-1}\|\}_{k \in K}$ remains bounded, but staying away from zero, at least on a subsequence, then the right-hand side is bounded but the left-hand side is unbounded, on the corresponding subsequence, as $\hat{\gamma}_k \to_K 0$. Therefore, it must be that $\hat{x}^k - x^{k-1} \to_K 0$, and thus $\hat{x}^k \to_K x^\star$ by Lemma 4.3(iii).

Now, by the mean-value theorem, there exists ξ^k on the line segment connecting x^{k-1} with $\hat{x}^k$ such that

$$\varphi(\hat{x}^k) - \varphi(x^{k-1}) = f(\hat{x}^k) + g(\hat{x}^k) - f(x^{k-1}) - g(x^{k-1})$$

$$= \langle \nabla f(\xi^k), \hat{x}^k - x^{k-1} \rangle + g(\hat{x}^k) - g(x^{k-1}).$$

Substituting $g(\hat{x}^k) - g(x^{k-1})$ from this expression into (4.6) yields

$$\langle \nabla f(x^{k-1}) - \nabla f(\xi^k), \hat{x}^k - x^{k-1} \rangle + \frac{1}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 + \varphi(\hat{x}^k) \le \varphi(x^{k-1}).$$

Therefore, exploiting (4.5) and Lemma 4.2(i), we obtain

$$\langle \nabla f(x^{k-1}) - \nabla f(\xi^k), \hat{x}^k - x^{k-1} \rangle + \frac{1}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 \le \varphi(x^{k-1}) - \varphi(\hat{x}^k)$$

$$\le \varphi(x^{k-1}) - \Phi_{k-1} + \frac{1 - \alpha_{\min}}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2$$

$$\le \frac{1 - \alpha_{\min}}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2.$$

Then, after rearranging, the Cauchy-Schwarz inequality yields

$$\frac{\alpha_{\min}}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 \le \|\nabla f(x^{k-1}) - \nabla f(\xi^k)\| \|\hat{x}^k - x^{k-1}\|.$$

Since $\hat{x}^k \ne x^{k-1}$ in view of (4.5) and Lemma 4.2(i), the previous inequality implies that

$$(4.7) \qquad\qquad \frac{\alpha_{\min}}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\| \le \|\nabla f(x^{k-1}) - \nabla f(\xi^k)\|.$$

Owing to $x^k \to_K x^\star$ and $\hat{x}^k \to_K x^\star$, it must be also $\xi^k \to_K x^\star$. Therefore, using the continuous differentiability of f and the fact that $\alpha_{\min} > 0$, it follows from (4.7) that $\frac{1}{\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\| \to_K 0$.

We now move toward the assertion. Exploiting the fact that x^k and $\hat{x}^k$ are solutions of the subproblems at Step 5 with stepsize γ_k and $\hat{\gamma}_k$, respectively, we obtain that

$$\langle \nabla f(x^{k-1}), x^k - \hat{x}^k \rangle + \frac{1}{2\gamma_k} \|x^k - x^{k-1}\|^2 + g(x^k) \le \frac{1}{2\gamma_k} \|\hat{x}^k - x^{k-1}\|^2 + g(\hat{x}^k)$$

and

$$\langle \nabla f(x^{k-1}), \hat{x}^k - x^k \rangle + \frac{1}{2\hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\|^2 + g(\hat{x}^k) \le \frac{1}{2\hat{\gamma}_k} \|x^k - x^{k-1}\|^2 + g(x^k).$$

Owing to $\gamma_k \le \beta_{\max} \hat{\gamma}_k < \hat{\gamma}_k$, adding these two inequalities yields the bound $\|x^k - x^{k-1}\| \le \|\hat{x}^k - x^{k-1}\|$. Therefore, by $\gamma_k \ge \beta_{\min} \hat{\gamma}_k > 0$, we have that

$$\frac{1}{\gamma_k} \|x^k - x^{k-1}\| \le \frac{1}{\beta_{\min} \hat{\gamma}_k} \|x^k - x^{k-1}\| \le \frac{1}{\beta_{\min} \hat{\gamma}_k} \|\hat{x}^k - x^{k-1}\| \to_K 0,$$

completing the proof. □

Patterning the proof above, local Lipschitz continuity of ∇f implies that the stepsize γ_k does not vanish along a convergent subsequence. This observation is stated explicitly in the following result, analogous to [17, Cor. 3.1]. A detailed proof is included in the Additional Material (p. 15).

> **Corollary 4.5.** *Suppose that Assumptions (A1) to (A3) are satisfied and consider a sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1. Let $x^\star$ be an accumulation point of $\{x^k\}_{k\in\mathbb{N}}$ and $\{x^k\}_{k\in K}$ a subsequence such that $x^k \to_K x^\star$. If ∇f is locally Lipschitz continuous (at least in a neighborhood of $x^\star$), then the subsequence $\{\gamma_k\}_{k\in K}$ is bounded away from zero.*

The following result provides the main global convergence guarantees for Algorithm 3.1. Notice that additional smoothness conditions are required on either f or g. These match those considered in [17, Thm 3.1] for the monotone variant.

Theorem 4.6. *Suppose that Assumptions (A1) to (A3) hold and at least one of the following conditions is satisfied for* (P):

(i) $g\colon \mathbb{X} \to \overline{\mathbb{R}}$ *is continuous relative to its domain;*

(ii) $\nabla f\colon \mathbb{X} \to \mathbb{X}$ *is locally Lipschitz continuous.*

Consider a sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1 and the finite value $\varphi_\star$ as in Lemma 4.3(i). Then, each accumulation point $x^\star$ of $\{x^k\}_{k\in\mathbb{N}}$ is M-stationary for (P) *and it is $\varphi(x^\star) = \varphi_\star$.*

Proof. Let $\{x^k\}_{k\in K}$ be a subsequence converging to $x^\star$. In view of Lemma 4.3(iii), it follows that also the subsequence $\{x^{k-1}\}_{k\in K}$ converges to $x^\star$. Furthermore, Lemma 4.4 yields $\frac{1}{\gamma_k}\|x^k - x^{k-1}\| \to_K 0$. Hence, if we can show the g-attentive convergence of $\{x^k\}_{k\in K}$ to $x^\star$, namely $x^k \to_K x^\star$ with $g(x^k) \to_K g(x^\star)$, the desired statement $0 \in \nabla f(x^\star) + \partial g(x^\star)$ is due to the optimality conditions (2.3) for Step 5 and is obtained by taking the limit $k \to_K \infty$ in

$$(4.8) \qquad 0 \in \nabla f(x^{k-1}) + \frac{1}{\gamma_k}(x^k - x^{k-1}) + \hat\partial g(x^k).$$

Moreover, by continuity of f and Lemma 4.3(i), it also implies that $\varphi(x^k) \to_K \varphi(x^\star) = \varphi_\star$.

(i) If g is continuous on its domain, the condition $g(x^k) \to_K g(x^\star)$ is readily obtained since all iterates x^k generated by Algorithm 3.1 belong to $\operatorname{dom} g$ and, owing to Lemma 4.2(ii) and lsc of g by Assumption (A2), $x^\star \in \operatorname{dom}\varphi = \operatorname{dom} g$.

(ii) It remains to consider the situation where g is merely lsc but ∇f is locally Lipschitz continuous. From $x^k \to_K x^\star$ and lsc of g, we find

$$g(x^\star) \le \liminf_{k\in K} g(x^k) \le \limsup_{k\in K} g(x^k).$$

Therefore, it suffices to show that $\limsup_{k\in K} g(x^k) \le g(x^\star)$. Since x^k solves the subproblem at Step 5 with stepsize γ_k, we obtain

$$\langle \nabla f(x^{k-1}), x^k - x^\star \rangle + \frac{1}{2\gamma_k}\|x^k - x^{k-1}\|^2 + g(x^k) \le \frac{1}{2\gamma_k}\|x^\star - x^{k-1}\|^2 + g(x^\star)$$

for each $k \in K$. We now take the upper limit for $k \to_K \infty$ on both sides. By Corollary 4.5, the stepsizes $\{\gamma_k\}_{k\in K}$ remain bounded away from zero, hence $\frac{1}{\gamma_k}\|x^\star - x^{k-1}\|^2 \to_K 0$. Using the continuity of ∇f, the convergences $x^k - x^{k-1} \to_K 0$ as well as $\frac{1}{\gamma_k}\|x^k - x^{k-1}\|^2 \to_K 0$ (see Lemma 4.3(ii)-(iii)), we obtain $\limsup_{k\in K} g(x^k) \le g(x^\star)$.

Altogether, we therefore get $x^k \to_K x^\star$ with $\varphi(x^k) \to_K \varphi(x^\star) = \varphi_\star$, completing the proof. □

We accompany our asymptotic characterization of Algorithm 3.1 with the finite termination property.

Theorem 4.7. *Suppose that Assumptions (A1) to (A3) hold. Let any $x^0 \in \operatorname{dom} g$ and $\varepsilon > 0$ be provided. If the sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1 admits an accumulation point, then the algorithm returns an ε-M-stationary point of* (P) *after finitely many iterations.*

Proof. Let $x^\star \in \mathbb{X}$ be an accumulation point of $\{x^k\}_{k\in\mathbb{N}}$ and let $\{x^k\}_{k\in K}$ be a subsequence such that $x^k \to_K x^\star$. Then, we have that $x^{k-1} \to_K x^\star$ by Lemma 4.3(iii) and $\frac{1}{\gamma_k}(x^k - x^{k-1}) \to_K 0$ by Lemma 4.4. Thus, owing to the continuity of ∇f by Assumption (A1), the stopping condition at Step 6 is satisfied for some $k \in K$ large enough. Finally, ε-M-stationarity of the returned point x^k follows from the bounds in (3.1). □

4.1 CONVERGENCE RATES

We now turn to analyzing the rate of asymptotic regularity for Algorithm 3.1 and, considering the iterates it generates, we recover the classical (worst-case) rate of (at least) $O(1/\sqrt{k})$ for nonconvex problems [4, Thm 10.15(c)]. Noticing that $\gamma_k \le \gamma_{\max}$ for all k and $\Phi_0 = \varphi(x^0) < \infty$, the result readily follows from the inequalities in (4.4). A detailed proof is included in the Additional Material (p. 15).

> **Corollary 4.8.** *Suppose that Assumptions (A1) to (A3) are satisfied and consider a sequence of iterates $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1. Then, for every $k \ge 1$ we have*
>
> $$\min_{1 \le j \le k} \|x^j - x^{j-1}\| \le \frac{1}{\sqrt{k}} \left(\frac{2\gamma_{\max}[\varphi(x^0) - \varphi_\star]}{p_{\min}(1-\alpha)} \right)^{1/2} < \infty$$
>
> *where $\varphi_\star$ is that of Lemma 4.3(i).*

We are also interested in characterizing the convergence rate of some stationarity measure for (P). In the nonconvex setting, we shall monitor the quantity $r_k := \frac{1}{\gamma_k}\|x^k - x^{k-1}\|$ as a measure of progress within Algorithm 3.1. Let us denote the proximal gradient mapping by $T_\gamma \colon \mathbb{X} \rightrightarrows \mathbb{X}$, $T_\gamma(x) := \operatorname{prox}_{\gamma g}(x - \gamma \nabla f(x))$. Since $x^k \in T_{\gamma_k}(x^{k-1})$ by Step 5, we have that $\operatorname{dist}(0, R_{\gamma_k}(x^{k-1})) \le r_k$, where $R_\gamma \colon \mathbb{X} \rightrightarrows \mathbb{X}$, $R_\gamma(x) := \frac{1}{\gamma}[x - T_\gamma(x)]$, denotes the fixed-point residual associated to (P) and provides a measure of stationarity thereof [4, 27, 28].

In order to obtain global complexity results, it seems indispensable the assumption of local Lipschitz continuity of ∇f. In fact, revealing that the stepsizes γ_k remain bounded away from zero (possibly only subsequentially or under some compactness assumptions), Corollary 4.5 plays a central role in obtaining suitable bounds for inferring rate-of-convergence results without additional assumptions.

We begin by stating in Theorem 4.9 that the global convergence rate of a stationarity measure for (P) is (at least) sublinear for convergent *subsequences*.

> **Theorem 4.9.** *Suppose Assumptions (A1) to (A3) are satisfied and ∇f is locally Lipschitz continuous. Consider a sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1. Let $x^\star$ be an accumulation point of $\{x^k\}_{k\in\mathbb{N}}$ and $\{x^k\}_{k\in K}$ a subsequence such that $x^k \to_K x^\star$. Then for every $k \ge 1$ we have*
>
> $$\min_{j \in K_k} \frac{\|x^j - x^{j-1}\|}{\gamma_j} \le \frac{1}{\sqrt{\#K_k}} \left(\frac{2[\varphi(x^0) - \varphi_\star]}{\gamma_\star p_{\min}(1-\alpha)} \right)^{1/2}$$
>
> *where $\varphi_\star$ is that of Lemma 4.3(i), $\gamma_\star > 0$ is a lower bound on $\{\gamma_k\}_{k\in K}$, and $\#K_k$ denotes the cardinality of the set of indices $K_k := \{j \in K \mid 1 \le j \le k\}$.*

Proof. Since ∇f is locally Lipschitz continuous and $x^k \to_K x^\star$, the subsequence $\{\gamma_k\}_{k\in K}$ is bounded away from zero by some $\gamma_\star > 0$, owing to Corollary 4.5. For all $k \in \mathbb{N}$, let $r_k := \|x^k - x^{k-1}\|/\gamma_k$. Extracting from (4.4), using that $\gamma_k \ge \gamma_\star > 0$ for all $k \in K$, and rearranging, we obtain

$$\frac{2[\Phi_0 - \varphi_\star]}{p_{\min}(1-\alpha)} \ge \sum_{j=1}^k \gamma_j r_j^2 \ge \sum_{j \in K_k} \gamma_j r_j^2 \ge \gamma_\star \sum_{j \in K_k} r_j^2 \ge \gamma_\star \#K_k \min_{j \in K_k} r_j^2,$$

proving the result, since $\Phi_0 = \varphi(x^0) < \infty$. $\qquad\square$

Now, if the iterates $\{x^k\}_{k\in\mathbb{N}}$ remain bounded, local Lipschitzness and compactness yield a uniform lower bound on the stepsizes $\{\gamma_k\}_{k\in\mathbb{N}}$. Then, a global sublinear rate of convergence for the stationarity measure readily follows from that of the iterates, given in Corollary 4.8. However, boundedness of

the iterates $\{x^k\}_{k\in\mathbb{N}}$ can be a strong assumption. In the following Lemma 4.10, we demonstrate that boundedness of the set of accumulation points is sufficient for the stepsizes to remain uniformly away from zero.

> **Lemma 4.10.** *Suppose Assumptions (A1) to (A3) are satisfied and ∇f is locally Lipschitz continuous. Consider a sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1. If the set of accumulation points of $\{x^k\}_{k\in\mathbb{N}}$ is bounded, then there exists $\overline{\gamma} > 0$ such that $\gamma_k \geq \overline{\gamma}$ for all $k \in \mathbb{N}$.*

Proof. Let $\omega \subset \mathbb{X}$ denote the set of accumulation points of $\{x^k\}_{k\in\mathbb{N}}$. Owing to Lemma 4.3(iii), if ω is bounded it is also compact. Now suppose, *ad absurdum*, that $\gamma_k \to 0$. Then, by compactness of ω, there exists a subsequence $\{x^k\}_{k\in K}$ such that $x^k \to_K x^\star$ for some $x^\star \in \omega$ and at the same time $\gamma_k \to_K 0$. On the contrary, due to the locally Lipschitz continuity assumption, Corollary 4.5 prevents $\{\gamma_k\}_{k\in K}$ from vanishing, thus proving the claim by contradiction. $\qquad\square$

Finally, patterning Theorem 4.9, the following result gives the corresponding global convergence rate with respect to the whole sequence. A detailed proof is included in the Additional Material (p. 16).

> **Corollary 4.11.** *Suppose Assumptions (A1) to (A3) are satisfied and ∇f is locally Lipschitz continuous. Consider a sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by Algorithm 3.1. If the set of accumulation points of $\{x^k\}_{k\in\mathbb{N}}$ is compact, then for every $k \geq 1$ we have*
>
> $$\min_{1\leq j\leq k} \frac{\|x^j - x^{j-1}\|}{\gamma_j} \leq \frac{1}{\sqrt{k}} \left(\frac{2[\varphi(x^0) - \varphi_\star]}{\gamma_\star p_{\min}(1-\alpha)} \right)^{1/2}$$
>
> *where $\varphi_\star$ is that of Lemma 4.3(i) and $\gamma_\star > 0$ is a lower bound for $\{\gamma_k\}_{k\in\mathbb{N}}$.*

Similar results can be obtained if the entire sequence $\{x^k\}_{k\in\mathbb{N}}$ converges to some accumulation point x^*. A situation leading to this circumstance is that of φ being locally strongly convex in a neighborhood of x^*. In this case, the convergence of $\{x^k\}_{k\in\mathbb{N}}$ to x^* directly follows from Lemma 4.3(iii) and Theorem 4.6, owing to [22, Lemma 4.10]. Another occurrence is that of φ satisfying some local error bound conditions [29, Thm 2], closely related to the Kurdyka-Łojasiewicz property [19].

5 SPECTRAL STEPSIZE AND NUMERICAL EXAMPLE

The general scheme we investigated comprises an arbitrary (yet bounded) choice of the proximal stepsize, enables nonmonotone behavior, and provides convergence guarantees under weak assumptions. These features lay the foundation for *spectral proximal gradient methods*, which can be constructed by considering the Barzilai-Borwein spectral estimate as a principled way for selecting the proximal stepsize [2]. In the spirit of quasi-Newton methods, a two-point approximation yields the spectral stepsize obtained in [2, Eq. 6]:

$$(5.1) \qquad \gamma_k^{\mathrm{BB}} := \frac{\langle \Delta x^k, \Delta x^k \rangle}{\langle \Delta x^k, \Delta g^k \rangle} \qquad \text{with } \Delta x^k := x^k - x^{k-1} \quad \text{and} \quad \Delta g^k := \nabla f(x^k) - \nabla f(x^{k-1}).$$

Associated to a local (isotropic, quadratic) surrogate of the smooth term f, the spectral choice is the "essential feature that puts efficiency in the projected gradient methodology" [6, §3]. When coupled with a linesearch for globalization, nonmonotonicity proved to be key in unleashing the performance of projected gradient methods with spectral approximations [24, 7]. We hope this behavior carries over to proximal methods.
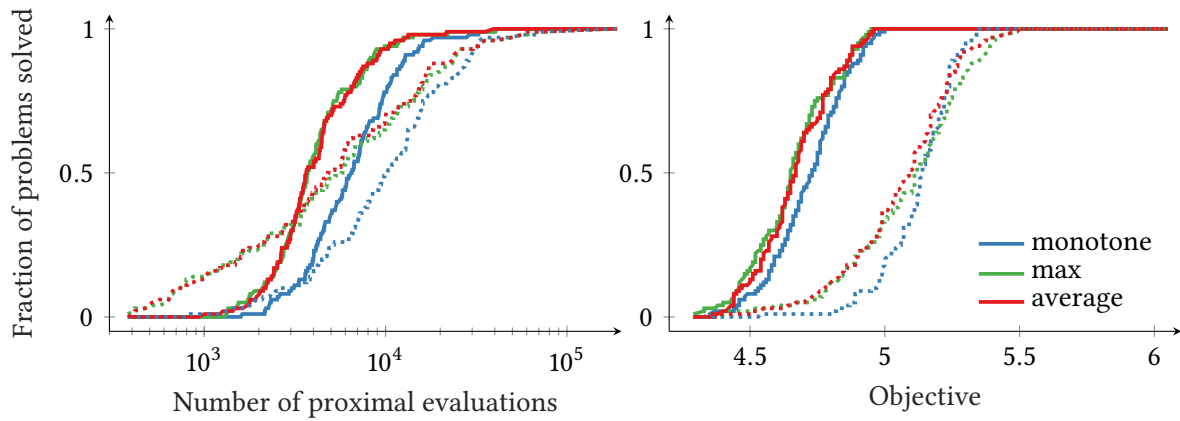
Figure 1: Comparing variants of Algorithm 3.1 on dictionary learning (5.2): combinations of plain (dotted) and spectral (solid) stepsizes with monotone and nonmonotone linesearch strategies.

We now compare the practical performance of several variants of Algorithm 3.1 on a numerical example. Two stepsize selection strategies are considered: "plain" simply inherits the previous value γ_{k-1}, whereas "spectral" uses the estimate $\gamma_{k-1}^{\mathrm{BB}}$ in (5.1); both choices are initialized with $\gamma_0 := 1$ and projected onto $[\gamma_{\min}, \gamma_{\max}]$ to comply with Step 4. These strategies are combined with three linesearch flavors: "monotone" ($p_k := 1$ for all $k \in \mathbb{N}$), nonmonotone with "average" merit ($p_k := 0.2$ for all $k \in \mathbb{N}$), and nonmonotone with "max" merit (memory $M = 5$) as in (1.1). Other parameters of Algorithm 3.1 are set to $\alpha_k := 0.999$, $\beta_k := 0.5$ for all $k \in \mathbb{N}$, $\gamma_{\min} := 10^{-12}$, $\gamma_{\max} := 10^{12}$, and tolerance $\varepsilon := 10^{-6}$.

DICTIONARY LEARNING

The challenge of dictionary learning is that of finding a collection of vectors that can sparsely yet accurately represent data signals. Given m signals $y_1, \ldots, y_m \in \mathbb{R}^n$, we wish to find ℓ dictionary atoms $d_1, \ldots, d_\ell \in \mathbb{R}^n$ such that each signal y_j can be approximated by a linear combination of these atoms via a sparse vector of coefficients $c_j \in \mathbb{R}^\ell$. Stacking the problem data in $Y \in \mathbb{R}^{n \times m}$, the unknown dictionary atoms in $D \in \mathbb{R}^{n \times \ell}$, and the unknown coefficients in $C \in \mathbb{R}^{\ell \times m}$, the problem can be expressed as

$$(5.2) \qquad \underset{D,C}{\text{minimize}} \quad \frac{1}{2}\|Y - DC\|_F^2 + \lambda\|C\|_0 \qquad\qquad \text{subject to} \quad \|d_j\|_2 = 1 \quad j = 1, \ldots, \ell,$$

where $\|\cdot\|_F$ and $\|\cdot\|_2$ denote the Frobenius and Euclidean norm, respectively. Parameter $\lambda \geq 0$ is introduced to tune the sparsity-promoting regularization $\|C\|_0$, which counts the number of nonzero entries in C. Without loss of generality, the dictionary atoms d_j are constrained to have unit norm; by scaling the coefficients in C, the objective remains unchanged. Relative to (P), function g encapsulates both the norm constraints on D and the cardinality cost on C, while f gives the quadratic loss.

We tested our algorithm on 100 problem instances with $n = 10$, $\ell = 20$ and $m = 30$, generated as follows. First, a dictionary $D \in \mathbb{R}^{n \times \ell}$ was constructed with entries sampled from a standard normal distribution, and each column was scaled to have unit norm. Then, a matrix $C \in \mathbb{R}^{\ell \times m}$ was randomly generated with $N = 3$ normally distributed nonzero coefficients per column. Finally, we set $Y := DC$, considered $\lambda := 10^{-2}$, and initialized the algorithm with $D^0 \in \mathbb{R}^{n \times \ell}$ and $C^0 \in \mathbb{R}^{\ell \times m}$ filled with normally distributed entries.

Numerical results are depicted in Figure 1, using profiles that represent the percentage of problems solved as a function of the available computational budget, considering either the evaluations of prox_g or the final value of the objective φ. Note that all solver variants find a solution, in the sense of (3.1), and terminate. Possibly inducing more erratic behaviors, spectral stepsize selection tends to reach lower

objective values, often with fewer evaluations compared to the plain counterpart. Conversely, for both plain and spectral stepsizes, nonmonotone strategies appear to perform better than monotone variants, especially in terms of evaluations (function, proximal, and iterations). Although the max and averaging nonmonotonicity flavors exhibit comparable capabilities, guarantees under weak assumptions are currently available for the latter only.

6 CONCLUSIONS

We studied an adaptive nonmonotone proximal gradient method for fully nonconvex composite optimization and its convergence properties under weak assumptions. Our findings indicate that, by attaining sufficient decrease at every iteration, an averaging merit function allows to establish convergence guarantees in the nonmonotone regime, matching those in the monotone case and thus bridging the gap emerged in the literature [17]. Furthermore, by exploiting these characterizations, we derived global rate-of-convergence results and bounds, for both iterates and a stationarity measure. Numerical comparisons demonstrated practical benefits of embracing nonmonotone linesearch strategies, whose averaging variant can rely on theoretical guarantees.

Future research work shall study the trade-offs and further compare (the performance of) spectral proximal gradient methods against Newton-type and variable metric approaches [18], which are based on more detailed (but potentially more intricate) models of the local landscape. Moreover, it remains an open question whether PANOC-type methods [27, 13] maintain or lose their convergence properties when the smooth term f is merely continuously differentiable.

ACKNOWLEDGMENTS

ADDITIONAL MATERIAL

Proof of Corollary 4.5. We may argue as in the proof of Lemma 4.4. By contradiction, let us assume, without loss of generality, that $\gamma_k \to_K 0$. For each $k \in K$, define $\hat{\gamma}_k$ and $\hat{x}^k$ as in that proof, and let $L_\star > 0$ denote the local Lipschitz constant of ∇f in a neighborhood of $x^\star$. Recall that $x^k \to_K x^\star$ and, from the proof of Lemma 4.4, that $\hat{x}^k \to_K x^\star$. Exploiting (4.7), we therefore obtain

$$\frac{\alpha}{2\hat{\gamma}_k}\|\hat{x}^k - x^{k-1}\| \leq \|\nabla f(x^{k-1}) - \nabla f(\xi^k)\| \leq L_\star\|x^{k-1} - \xi^k\| \leq L_\star\|\hat{x}^k - x^{k-1}\|$$

for all $k \in K$ sufficiently large, using the fact that ξ^k is on the line segment between x^{k-1} and $\hat{x}^k$. Since $\hat{x}^k \neq x^{k-1}$ by (4.5), this gives a contradiction as $\hat{\gamma}_k \to_K 0$. Hence, $\{\gamma_k\}_{k \in K}$ stays bounded away from zero. □

Proof of Corollary 4.8. Let us denote $\Delta_k := \|x^k - x^{k-1}\|$ for all $k \geq 1$. Owing to $\gamma_k \leq \gamma_{\max}$ for all k and $\Phi_0 = \varphi(x^0) < \infty$, from (4.4) we obtain that

$$\frac{2[\varphi(x^0) - \varphi_\star]}{p_{\min}(1-\alpha)} \geq \sum_{j=1}^{k}\frac{1}{\gamma_j}\Delta_j^2 \geq \frac{1}{\gamma_{\max}}\sum_{j=1}^{k}\Delta_j^2 \geq \frac{k}{\gamma_{\max}}\min_{1\leq j\leq k}\Delta_j^2.$$

Rearranging yields the claimed assertion.   □

Proof of Corollary 4.11. Since ∇f is locally Lipschitz continuous, it follows from Lemma 4.10 that $\{\gamma_k\}_{k\in\mathbb{N}}$ remains bounded away from zero if the set of accumulation points of $\{x^k\}_{k\in\mathbb{N}}$ is compact. Rearranging from (4.4), using that $\gamma_k \geq \gamma_\star > 0$ for all $k \in \mathbb{N}$, and rearranging, we obtain

$$\frac{2[\Phi_0 - \varphi_\star]}{p_{\min}(1-\alpha)} \geq \sum_{j=1}^{k} \frac{1}{\gamma_j} \|x^j - x^{j-1}\|^2 \geq \gamma_\star \sum_{j=1}^{k} \frac{\|x^j - x^{j-1}\|^2}{\gamma_j^2} \geq k\gamma_\star \min_{1 \leq j \leq k} \frac{\|x^j - x^{j-1}\|^2}{\gamma_j^2},$$

proving the result.   □

REFERENCES

[1] L. Armijo, Minimization of functions having Lipschitz continuous first partial derivatives, *Pacific Journal of Mathematics* 16 (1966), 1–3.

[2] J. Barzilai and J. M. Borwein, Two-point step size gradient methods, *IMA Journal of Numerical Analysis* 8 (1988), 141–148.

[3] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer International Publishing, 2017, doi:10.1007/978-3-319-48311-5.

[4] A. Beck, *First-Order Methods in Optimization*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017, doi:10.1137/1.9781611974997.

[5] E. G. Birgin and J. M. Martínez, *Practical Augmented Lagrangian Methods for Constrained Optimization*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2014.

[6] E. G. Birgin, J. M. Martínez, and M. Raydan, Nonmonotone spectral projected gradient methods on convex sets, *SIAM Journal on Optimization* 10 (2000), 1196–1211, doi:10.1137/s1052623497330963.

[7] E. G. Birgin, J. M. Martínez, and M. Raydan, Spectral projected gradient methods: Review and perspectives, *Journal of Statistical Software* 60 (2014), 1–21, doi:10.18637/jss.v060.i03.

[8] A. L. Cauchy, Méthode générale pour la résolution des systèms d'équations simultanées, *Comp. Rend. Sci. Paris* 25 (1847), 536–538.

[9] P. L. Combettes and J. C. Pesquet, Proximal splitting methods in signal processing, in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer, New York, 2011, 185–212.

[10] A. De Marchi, Implicit augmented Lagrangian and generalized optimization, 2023, arXiv:2302.00363.

[11] A. De Marchi, X. Jia, C. Kanzow, and P. Mehlitz, Constrained composite optimization and augmented Lagrangian methods, *Mathematical Programming* (2023), doi:10.1007/s10107-022-01922-4.

[12] A. De Marchi and A. Themelis, An interior proximal gradient method for nonconvex optimization, 2022, arXiv:2208.00799.

[13] A. De Marchi and A. Themelis, Proximal gradient algorithms under local Lipschitz gradient continuity, *Journal of Optimization Theory and Applications* 194 (2022), 771–794, doi:10.1007/s10957-022-02048-5.

[14] M. Fukushima and H. Mine, A generalized proximal point algorithm for certain non-convex minimization problems, *International Journal of Systems Science* 12 (1981), 989–1000, doi:10.1080/00207728108963798.

[15] L. Grippo, F. Lampariello, and S. Lucidi, A nonmonotone line search technique for Newton's method, *SIAM Journal on Numerical Analysis* 23 (1986), 707–716, doi:10.1137/0723046.

[16] X. Jia, C. Kanzow, P. Mehlitz, and G. Wachsmuth, An augmented Lagrangian method for optimization problems with structured geometric constraints, *Mathematical Programming* (2022), doi:10.1007/s10107-022-01870-z.

[17] C. Kanzow and P. Mehlitz, Convergence properties of monotone and nonmonotone proximal gradient methods revisited, *Journal of Optimization Theory and Applications* (2022), doi:10.1007/s10957-022-02101-3.

[18] J. D. Lee, Y. Sun, and M. A. Saunders, Proximal Newton-type methods for minimizing composite functions, *SIAM Journal on Optimization* 24 (2014), 1420–1443, doi:10.1137/130921428.

[19] G. Li and T. K. Pong, Calculus of the exponent of Kurdyka-Łojasiewicz inequality and its applications to linear convergence of first-order methods, *Foundations of Computational Mathematics* 18 (2018), 1199–1232, doi:10.1007/s10208-017-9366-8.

[20] H. Li and Z. Lin, Accelerated proximal gradient methods for nonconvex programming, in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (eds.), volume 28, Curran Associates, Inc., 2015.

[21] B. S. Mordukhovich, *Variational Analysis and Applications*, Springer, 2018, doi:10.1007/978-3-319-92775-6.

[22] J. J. Moré and D. C. Sorensen, Computing a trust region step, *SIAM Journal on Scientific and Statistical Computing* 4 (1983), 553–572, doi:10.1137/0904038.

[23] N. Parikh and S. Boyd, Proximal algorithms, *Foundations and Trends in Optimization* 1 (2014), 127–239, doi:10.1561/2400000003.

[24] M. Raydan, The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem, *SIAM Journal on Optimization* 7 (1997), 26–33, doi:10.1137/s1052623494266365.

[25] R. T. Rockafellar and R. J. Wets, *Variational Analysis*, volume 317, Springer, 1998.

[26] P. Sopasakis, E. Fresk, and P. Patrinos, OpEn: Code generation for embedded nonconvex optimization, *IFAC-PapersOnLine* 53 (2020), 6548–6554, doi:10.1016/j.ifacol.2020.12.071.

[27] L. Stella, A. Themelis, P. Sopasakis, and P. Patrinos, A simple and efficient algorithm for nonlinear model predictive control, in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, IEEE, 2017, 1939–1944, doi:10.1109/cdc.2017.8263933.

[28] A. Themelis, L. Stella, and P. Patrinos, Forward-backward envelope for the sum of two nonconvex functions: Further properties and nonmonotone linesearch algorithms, *SIAM Journal on Optimization* 28 (2018), 2274–2303, doi:10.1137/16m1080240.

[29] P. Tseng and S. Yun, A coordinate gradient descent method for nonsmooth separable minimization, *Mathematical Programming* 117 (2009), 387–423, doi:10.1007/s10107-007-0170-0.

[30] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, Sparse reconstruction by separable approximation, *IEEE Transactions on Signal Processing* 57 (2009), 2479–2493, doi:10.1109/tsp.2009.2016892.

[31] H. Zhang and W. W. Hager, A nonmonotone line search technique and its application to unconstrained optimization, *SIAM Journal on Optimization* 14 (2004), 1043–1056, doi:10.1137/s1052623403428208.