

# Counter-examples in first-order optimization: a constructive approach

Baptiste Goujoud and Aymeric Dieuleveut and Adrien Taylor

**Abstract**—While many approaches were developed for obtaining worst-case complexity bounds for first-order optimization methods in the last years, there remain theoretical gaps in cases where no such bound can be found. In such cases, it is often unclear whether no such bound exists (e.g., because the algorithm might fail to systematically converge) or simply if the current techniques do not allow finding them. In this work, we propose an approach to automate the search for cyclic trajectories generated by first-order methods. This provides a constructive approach to show that no appropriate complexity bound exists, thereby complementing approaches providing sufficient conditions for convergence. Using this tool, we provide ranges of parameters for which the famous Polyak heavy-ball, Nesterov accelerated gradient, inexact gradient descent, and three-operator splitting algorithms fail to systematically converge, and show that it nicely complements existing tools searching for Lyapunov functions.

## I. INTRODUCTION

In the last years, first-order optimization methods (or algorithms) have attracted a lot of attention due to their practical success in many applications, including in machine learning (see, e.g., [5]). Theoretical foundations for those methods played a crucial role in this success, e.g., by enabling the development of momentum-type methods (see, e.g., [29], [27]). Formally, we consider the optimization problem

$$x_* \triangleq \arg \min_{x \in \mathbb{R}^d} f(x) \quad (\text{OPT})$$

for a function  $f$  belonging to a class of functions  $\mathcal{F}$  (e.g., the set of convex functions, or the set of strongly convex and smooth functions, etc.). Classical first-order optimization methods for solving this problem include *gradient descent* (GD), *Nesterov accelerated gradient method* (NAG) [27], and the *heavy-ball method* (HB) [29]. These families of algorithms are parametrized: for example, GD is parametrized by a step-size  $\gamma$  and HB is parametrized by both a step-size  $\gamma$  and a momentum parameter  $\beta$ . We generically denote by  $\mathcal{A}$  any such method, for a specific choice of its parameters. For a given class of function  $\mathcal{F}$  and an algorithm  $\mathcal{A}$ , we typically aim at answering the question

Does  $\mathcal{A}$  converge on every function of  $\mathcal{F}$  to their respective minimum?

B. Goujoud is PhD candidate at CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, Route de Saclay, 91120 Palaiseau, France. [baptiste.goujoud@gmail.com](mailto:baptiste.goujoud@gmail.com)

A. Dieuleveut is faculty at CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, Route de Saclay, 91120 Palaiseau, France. [aymeric.dieuleveut@polytechnique.edu](mailto:aymeric.dieuleveut@polytechnique.edu)

A. Taylor is faculty at INRIA Paris, 2 Rue Simone IFF, 75012 Paris. [adrien.taylor@inria.fr](mailto:adrien.taylor@inria.fr)

Common examples of function classes  $\mathcal{F}$  include the set  $\mathcal{F}_{\mu,L}$  of  $\mu$ -strongly convex and  $L$ -smooth functions, and the set  $\mathcal{Q}_{\mu,L}$  of  $\mu$ -strongly convex and  $L$ -smooth *quadratic* functions, for  $\mu, L \geq 0$ .

This type of analysis, requiring results to hold on every function of a given class  $\mathcal{F}$  is commonly referred to as *worst-case analysis* and is the most popular paradigm for the analysis of optimization algorithms, see, e.g., [27], [12], [6], [13], [7]. In this context, a very successful technique for proving worst-case convergence consists in looking for a decreasing sequence (called *Lyapunov sequence* [26], [23], [22]) of expressions  $V_t$  of the iterates  $x_t$ , i.e. such that

$$\forall f \in \mathcal{F}, \forall t, \forall x_t, \quad V_{t+1}((x_s)_{s \leq t+1}) \leq V_t((x_s)_{s \leq t}), \quad (1)$$

where some quantity of interest is upper-bounded by  $V_T((x_s)_{s \leq T})$  as  $T$  goes to infinity. For instance, when studying GD with step-size  $1/L$  on the class  $\mathcal{F}_{0,L}$  of  $L$ -smooth convex functions, we prove that  $\forall f \in \mathcal{F}_{0,L}$ ,  $(t+1)(f(x_{t+1}) - f(x_*)) + \frac{1}{2}\|x_{t+1} - x_*\|^2 \leq t(f(x_t) - f(x_*)) + \frac{1}{2}\|x_t - x_*\|^2$ . Therefore,  $V_t((x_s)_{s \leq t}) = t(f(x_t) - f(x_*)) + \frac{1}{2}\|x_t - x_*\|^2$  defines a decreasing sequence, and  $f(x_t) - f(x_*) \leq V_t((x_s)_{s \leq t})/t \leq V_0(x_0)/t$ , proving convergence of this method on this class of functions.

Due to the simplicity of the underlying proofs, the Lyapunov approach is particularly popular, e.g., for NAG [27], [4], and HB [16]. See [3], [13] for surveys on this topic.

**Necessary condition for worst-case convergence.** While finding a decreasing Lyapunov sequence guarantees convergence, not finding one does not guarantee anything: there may still exist a Lyapunov sequence, that the current analysis was not able to capture, or the method could converge without the existence of such Lyapunov sequence. Establishing that a method *provably does not admit a worst-case convergence analysis* is therefore critical for avoiding spending an indefinite amount of time and effort searching for a non-existent convergence guarantee. The existence of a *cycle* for the algorithm on a particular function means that it diverges on that function: in other words, the absence of cycle on all functions is a necessary condition for worst-case convergence. Moreover, a cycle can be observed after only a *finite* number of steps of the algorithm, while observing the divergence of a non periodic sequence is difficult or impossible. Overall, this makes the search for cycles a computationally practical way of proving divergence.

In order to discover cycles, we rely on computer-assisted worst-case analysis. *Performance estimation problems* (PEP [11], [33]) provide a systematic approach to obtain convergence guarantees, including the search for appropriate Lyapunov arguments. Some packages (especially `Pesto` [32])

and Pepit [18]) automate these tasks. We formulate cycle discovery as a minimization problem that can be cast in a PEP, and rely on the `Pepit` package to solve it.

**Examples:** We demonstrate the applicability of our method on several examples. In particular, the case of HB illustrates the potential of our methodology. In fact, the search for the step-size  $\gamma$  and momentum  $\beta$  parameters leading to the fastest worst-case convergence over  $\mathcal{F}_{\mu,L}$  is still an open problem, and the existence of parameters resulting in an accelerated rate remains a lingering question. Indeed, [25] exhibits a smooth and strongly convex function on which HB cycles, for parameters  $\gamma$  and  $\beta$  optimizing the worst-case guarantee on  $\mathcal{Q}_{\mu,L}$ . On the other hand, [16] obtains a worst-case convergence on  $\mathcal{F}_{\mu,L}$  for other parameters, but without acceleration. Recently, [34] proposes a very general procedure to find Lyapunov sequences and extended the region of parameters  $\gamma$  and  $\beta$  HB provably converges on, leveraging PEPs. However, outside this region of the parameter space, the question of the convergence of the HB method remains open in the absence of a proof of divergence. For this example, our approach demonstrates that a cycle exists for almost all parameters for which no Lyapunov is known.

**Summary of contributions:** This paper proposes a systematic approach to prove that no worst-case certificate of convergence can be obtained for a given algorithm  $\mathcal{A}$  on a class  $\mathcal{F}$ . To do so, we establish the existence of a function in  $\mathcal{F}$  over which  $\mathcal{A}$  cycles. We illustrate our approach by applying it to three famous first-order optimization algorithms, namely HB, NAG, inexact gradient descent with relatively bounded error. We further showcase the applicability of the approach to more general types of problems by studying the three-operator splitting method for monotone inclusions. For each method, we describe the set of parameters for which it is known to converge and the ones where we establish the existence of a cycle. In the first three examples, our approach enables to fill the gap: we show the existence of cycles for all parametrizations not known to result in convergence.

**Organization:** The rest of the paper is organized as follows. In Section II, we introduce the concept of a stationary algorithm and formally define a cycle. In Section III, we present our methodology to discover cycles, relying on PEP. Finally, in Section IV, we provide the numerical results.

## II. DEFINITIONS AND NOTATIONS

In this section, we consider a subclass of first-order methods, tailored for our analysis. It is chosen to ensure the periodicity of an algorithm that cycles once (see Proposition 3.1). The class reduces to “*p-stationary canonical linear iterative optimization algorithms*” (p-SCLI, see [2, Definition 1]) when the dependency to the previous iterates and gradients is linear which is a particular case of “*fixed-step first-order methods*” (FSFOM, see in [33, Definition 4]). Here, we consider *stationary first-order methods* (SFOM), whose iterates are defined as a fixed function of a given number of lastly observed iterates, as well as output of some oracles called on those iterates. Examples of such oracles include gradients, approximate gradients, function

evaluations, proximal step, exact line-search, Frank-Wolfe-type steps (see [31], [18] for lists of oracles that can be handled using PEPs). The oracles we use depend on the setting under consideration.

*Definition 2.1 (Stationary first-order method (SFOM)):*

A method  $\mathcal{A}$  is called order- $\ell$  *stationary first-order method* if there exists a deterministic first-order oracle  $\mathcal{O}^{(f)}$  and a function  $A$  such that the sequence generated on the function  $f$  verifies  $\forall t \geq \ell$ ,

$$x_t = A((x_{t-s}, \mathcal{O}^{(f)}(x_{t-s}))_{s \in \llbracket 1, \ell \rrbracket}). \quad (\text{SFOM})$$

For any given function of interest  $f$  and any initialization  $(x_t)_{t \in \llbracket 0, \ell-1 \rrbracket}$ , an order- $\ell$  (SFOM)  $\mathcal{A}$  iteratively generates a sequence  $(x_t)_{t \in \mathbb{N}}$  that we denote  $\mathcal{A}(f, (x_t)_{t \in \llbracket 0, \ell-1 \rrbracket})$ .

Definition 2.1 above is very similar to the definition of a general first-order method. However, the key assumption here is that the operation  $A$  does not depend on the iteration counter  $t$ : the algorithm is *stationary*. While this assumption is restrictive, many first-order methods are of the form (SFOM), including (but not limited to): GD, HB [29] and NAG [28] with constant step-sizes. On the other hand, any strategy involving decreasing step-size (e.g. for GD), or increasing momentum parameter (e.g. for NAG on  $\mathcal{F}_{0,L}$  as in [27]) are not in the scope of this definition. Note that the aforementioned examples use the first-order oracle  $\mathcal{O}^{(f)}(x) \triangleq (\nabla f(x), f(x))$ , although our methodology applies beyond this simple setting, as previously discussed. As an example, IV-D considers an algorithm relying on the resolvent (or proximal operation).

Stationarity is essential for being able to prove existence of a cyclical behavior in a finite number of steps. Next, we define a cyclic sequence.

*Definition 2.2 (Cycle):* For any positive integer  $K \geq 2$ , a sequence  $(x_t)_{t \geq 0}$  is said to be  $K$ -cyclic if  $\forall t \geq 0, x_t = x_{t+K}$ . A sequence  $x$  is said to be cyclic if there exists  $K \geq 2$  such that  $x$  is  $K$ -cyclic.

For any given order- $\ell$  (SFOM)  $\mathcal{A}$ , and any function class  $\mathcal{F}$ , we want to address the question

Does there exist a function  $f \in \mathcal{F}$  and an initialization  $(x_t)_{t \in \llbracket 0, \ell-1 \rrbracket}$  such that  $\mathcal{A}(f, (x_t)_{t \in \llbracket 0, \ell-1 \rrbracket})$  is cyclic?

*Example 2.3:* In [25, Equation 4.11], the authors answer positively to this question by providing a cycle of length 3, on the class  $\mathcal{F}_{\mu,L}$  with  $(\mu, L) = (1, 25)$ , and for  $\mathcal{A}$  the heavy-ball method with step-size  $\gamma = (\frac{2}{\sqrt{L} + \sqrt{\mu}})^2$  and momentum parameter  $\beta = (\frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}})^2$ . Those parameters are natural candidates, that correspond to the limit of the step-size and momentum in Chebychev acceleration [14], [24], [35], and result in an acceleration for quadratic functions.

In Section IV, we extend this result to more parameters.

## III. SEARCHING FOR CYCLES

In this section, we show how to find cyclic trajectories.

## A. Motivation

Finding diverging trajectories for an algorithm  $\mathcal{A}$  might be challenging. We thus focus on cycles, as they allow to focus on a finite sequences of iterates only. Indeed, for an SFOM, once we observe the cycle to be repeated once, we can easily extrapolate: this same cycle is repeated again and again. This statement is formalized in the following proposition.

*Proposition 3.1:* Let  $\mathcal{A}$  be an order- $\ell$  (SFOM), and  $(x_t)_{t \in \mathbb{N}}$  be any sequence generated by  $\mathcal{A}$ . Then the sequence  $(x_t)_{t \in \mathbb{N}}$  is cyclic if and only if there exists  $K \geq 2$  such that  $\forall t \in \llbracket 0, \ell - 1 \rrbracket, x_t = x_{t+K}$ .

The proof is provided in [17].

## B. Approach

We now present the approach used to search for cycles, based on performance estimation problems (PEPs) [11], [33]. We consider an algorithm  $\mathcal{A}$ , a function  $f$  and initial points  $(x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket}$ , and run  $\mathcal{A}$  on  $f$  starting on  $(x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket}$ . This generates the sequence  $x = \mathcal{A}(f, (x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket})$ . For any positive integer  $K$ , we then define the non-negative score

$$s_K(\mathcal{A}, f, (x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket}) = \sum_{t=0}^{\ell-1} \|x_t - x_{t+K}\|^2.$$

From Proposition 3.1, this score is identically zero if and only if  $\mathcal{A}$  cycles on  $f$  when starting from  $(x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket}$ . This suggests that one can search for cycles of length  $K$  by minimizing the score  $s_K(\mathcal{A}, f, (x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket})$  w.r.t. the function  $f$  and the initialization  $(x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket}$ .

Observe that fixed points of  $\mathcal{A}$ , that correspond to cycles of length 1, also cancel this score. Our goal is to search for cycles of length at least  $K \geq 2$ , that entail that the algorithm diverges for a particular function and initialization. As any convergent algorithm must admit the optimizer of  $f$  as fixed point, we have to exclude fixed points. To do so, we add the constraint that the two first iterates are far from each other. In most cases of interest, making this constraint can be done without loss of generality due to the homogeneity of the underlying problems. We arrive to the following formulation:

$$\left\{ \begin{array}{ll} \text{minimize} & \sum_{t=0}^{\ell-1} \|x_t - x_{t+K}\|^2 \\ \text{subject to} & \left\{ \begin{array}{l} x = \mathcal{A}(f, (x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket}) \\ \|x_1 - x_0\|^2 \geq 1. \end{array} \right. \end{array} \right. \quad (\mathcal{P})$$

As we see in the next sections, this problem can be used to answer the question of interest by testing the nullity of the solution of  $(\mathcal{P})$ .

As is,  $(\mathcal{P})$  looks intractable due to the minimization over the infinite-dimensional space  $\mathcal{F}$  and its non-convexity. This can be handled using the techniques proposed in [33], [31], developed for PEP. It consists in reformulating  $(\mathcal{P})$  into a semi-definite program (SDP) using interpolation / extension properties for the class  $\mathcal{F}$ , together with SDP lifting.

Indeed,  $(\mathcal{P})$  does not fully depend on  $f$ , but only on  $\mathcal{O}^{(f)}(x_t)$  where  $t \in \llbracket 0; K + \ell - 2 \rrbracket$ . By introducing the

variables  $\mathcal{O}_t \triangleq \mathcal{O}^{(f)}(x_t)$ , we can replace the constraint  $x = \mathcal{A}(f, (x_t)_{t \in \llbracket 0, \ell - 1 \rrbracket})$  of  $(\mathcal{P})$  by

$$\left\{ \begin{array}{l} x_\ell = A((x_{\ell-s}, \mathcal{O}_{\ell-s})_{s \in \llbracket 1, \ell \rrbracket}), \\ \vdots \\ x_{K+\ell-1} = A((x_{K+\ell-1-s}, \mathcal{O}_{K+\ell-1-s})_{s \in \llbracket 1, \ell \rrbracket}), \end{array} \right.$$

and minimize over the finite dimensional variables  $(\mathcal{O}_t)_{0 \leq t \leq K+\ell-2}$  instead of  $f$ , under the constraint that there exists a function  $f \in \mathcal{F}$  that interpolates those values, i.e. that verifies  $\mathcal{O}^{(f)}(x_t) = \mathcal{O}_t$  for all  $t \in \llbracket 0; K + \ell - 2 \rrbracket$ . For some classes  $\mathcal{F}$ , those interpolation property are equivalent to tractable inequalities, as in the following example.

*Example 3.2 (L-smooth convex functions):* If the oracles are only the gradients and the function values of the objective function  $f$ , denoting  $f_i \triangleq f(x_i)$  and  $g_i \triangleq \nabla f(x_i)$  (i.e.  $\mathcal{O}_i \triangleq (g_i, f_i)$ ), the interpolation conditions of  $\mathcal{F}_{0,L}$  are provided in [33] as

$$\forall i, j, f_i \geq f_j + \langle g_j, x_i - x_j \rangle + \frac{1}{2L} \|g_i - g_j\|^2. \quad (\text{IC})$$

This function class is considered in three of the four examples under consideration in the next section (HB, NAG, inexact GD). However, the methodology described in this paper applies to many other classes beyond  $\mathcal{F}_{\mu,L}$ , see for instance [18, Function classes]: an example of such a class is used in the fourth example of the following section. Each class considered must be described by its interpolation conditions, similar to (IC). Other examples of known interpolations conditions are provided in [31, Th. 3.4-3.6], [10, Cor.1&2], [21, Th.1] [19, Th. 2.6]. The key ingredient for a class to enter in the scope of this paper, is that its interpolation conditions are expressed as a degree 2 polynomial in  $x_t$  and  $\mathcal{O}_t$ , and that a given variable is not involved both in a monomial of degree 1 and one of degree 2, as in (IC).

Then, the SDP lifting part consists in introducing a Gram matrix  $G$  [33, Theorem 5] of vectors among  $x_t$  and  $\mathcal{O}_t$  that are involved in degree 2 monomials, so that those quadratic expressions of  $x_t$  and  $\mathcal{O}_t$  are then expressed linearly in term of  $G \succeq 0$ . Thereby, the problem can be cast a standard SDP.

In the case where the oracle is  $\mathcal{O}_t = (g_t, f_t)$ , and the class of interest  $\mathcal{F}$  is the class of  $L$ -smooth convex functions  $\mathcal{F}_{0,L}$ , the objective and all the constraints of  $(\mathcal{P})$  are written linearly in terms of  $(f_t)_t$  and quadratically in terms of  $(x_t, g_t)_t$ . Therefore, we define  $G$  as the Gram matrix of  $(x_t, g_t)_t$  and  $F$  as a vector storing all the values  $f_t$  leading to an SDP reformulation of the problem. See, e.g., [18, Section 2] for a detailed derivation on a simple example.

Setting  $u \triangleq (G, F)$ ,  $(\mathcal{P})$  is generally rewritten, under above mentioned key ingredients, as

$$\left\{ \begin{array}{ll} \text{minimize} & \langle u, v_{obj} \rangle \\ \text{subject to} & \left\{ \begin{array}{l} \langle u, v_1 \rangle \geq 0 \\ \vdots \\ \langle u, v_n \rangle \geq 0 \\ \langle u, v_{aff} \rangle \geq 1. \end{array} \right. \end{array} \right. \quad (\text{SDP-}\mathcal{P})$$

The objective is linear, as well as the first  $n$  constraints. The affine constraint  $\langle u, v_{aff} \rangle \geq 1$  enables to discard the trivial solution  $u = 0$  and corresponds in  $(\mathcal{P})$  to the constraint

$\|x_1 - x_0\|^2 \geq 1$ . Finally, the constraint  $u \in \mathcal{C}$  corresponds to the constraint  $G \succeq 0$ .  $\mathcal{C}$  is then a closed convex semi-cone.

By definition, if there exists a feasible vector  $u$  such that the objective of (SDP- $\mathcal{P}$ ) is zero, then it describes a cycle. Moreover (SDP- $\mathcal{P}$ ) is convex and efficiently solvable (due to the existence of a Slater point [33, Theorem 6]).

In the next sections, we numerically apply this methodology through the `Pepit` python package [18] which takes care about the tractable reformulations of ( $\mathcal{P}$ ) into (SDP- $\mathcal{P}$ ). (SDP- $\mathcal{P}$ ) is then solved using through a standard solver [1] to determine the infimum value of  $\langle u, v_{\text{obj}} \rangle$  over the feasible set of (SDP- $\mathcal{P}$ ). The next theorem allows to conclude about the existence of cycles.

*Theorem 3.3:* Assuming the infimum value of ( $\mathcal{P}$ ) to be 0, then there exists a cycle.

The proof is provided in [17]

#### IV. APPLICATION TO FOUR DIFFERENT (SFOM)S

In this section we illustrate this methodology on four examples: heavy-ball (HB), Nesterov accelerated gradient (NAG), gradient descent (GD) with inexact gradients, and three-operator splitting (TOS). For each, we apply the methodology proposed in Section III. The code is available in the public github repository <https://github.com/bgoujaud/cycles>. Since ingredients are the same as those of classical PEPs, we also use the python package `Pepit` [18]. We perform a grid search over the spaces of *parameters of interest*  $\Omega$ , described in the respective subsections. We compare the parameter region where Lyapunov functions can be obtained with the region in which we establish that the method cannot have a guaranteed worst-case convergence (due to the existence of cycles). More precisely, in Figures 1 to 4 below, green regions corresponds to parameter choices for which the methods converge (existence of a Lyapunov function, found using the code provided with [30]). Conversely, in the red regions, our methodology establishes that the method cycles on at least one function of  $\mathcal{F}$ . In short, the algorithms converge in the green regions and do not converge in the worst-case in the red ones.

Note that some parameters for which the algorithm  $\mathcal{A}$  admits a worst-case convergence guarantee could theoretically exist outside the green region: indeed, in [30], the authors do not guarantee that they necessarily find convergence. Similarly, parameters for which  $\mathcal{A}$  does *not* admit a worst-case convergence guarantee could theoretically exist outside of the red region: indeed ( $\mathcal{P}$ ) is defined for a fixed cycle length  $K$ , and we therefore run it several times with different values of  $K$ . Longer cycles are therefore not detected. Moreover, the non-existence of cycles does not necessarily imply that the algorithm always converges.

Interestingly, in practice, we observe on Figures 2 and 3 that the set  $\Omega$  of parameters is completely filled by the union of those 2 regions and that it is almost the case on Figure 1 (it may have been if we had searched for cycles of all lengths). As a consequence, we fully characterize for which tuning the algorithms admit a guaranteed worst-case convergence. On




Fig. 1. Heavy-ball (HB). Green area: set of parameters  $(\gamma, \beta) \in \Omega_{\text{HB}}$  for which a Lyapunov function exists; Red area: set of parameters  $(\gamma, \beta) \in \Omega_{\text{HB}}$  for which (HB) cycles on at least one function in  $\mathcal{F}_{0,L}$ .

the contrary, there remains a significant gap between the red and the green regions in our last example, see Figure 4.

##### A. Heavy-ball

The HB algorithm, as introduced by [29], corresponds to the following update, for a step-size parameter  $\gamma$  and a momentum parameter  $\beta$ :

$$x_{t+1} = x_t + \beta(x_t - x_{t-1}) - \gamma \nabla f(x_t). \quad (\text{HB})$$

Therefore (HB) is an order-2 (SFOM).

**Set  $\Omega_{\text{HB}}$  of parameters of interest:** HB converges on the set  $\mathcal{Q}_{0,L}$  if and only if the parameters  $\gamma, \beta$  verify  $0 \leq \gamma \leq 2(1 + \beta)/L \leq 4/L$  [29]. Note this condition enforces  $-1 \leq \beta \leq 1$ . Moreover, we restrict to  $\beta \geq 0$  as  $\beta < 0$  is not an interesting setting (slowing down convergence with respect to GD). Therefore, we limit our analysis to this set.

**Interpretation.** The red area in Figure 1 shows parameters where cycles of length  $K \in \llbracket 2; 25 \rrbracket$  are found by our methodology. The red color intensity indicates the length of the shortest cycle. A striking observation is that the space  $\Omega_{\text{HB}}$  is almost filled by the union of the red area and green one (where Lyapunov functions exist). Thereby, for almost all values of the parameters, we have a definitive answer on the existence of a certificate of convergence in the worst-case. That being said, there exists a small unfilled region in the top left corner (see the zoom on Figure 1) In this region, we do not know how HB behaves, and whether it accelerates. However, adding longer cycle length may enable to obtain cycles in that area. Indeed we considered only cycles of length  $K \leq 25$ , for computational reasons. Recently, [20] computed the analytical region of cycles of (HB) showing non-acceleration of the latter.

##### B. Nesterov accelerated gradient

NAG (also known as the *fast gradient method*) was introduced by [28] and corresponds to the following update, for a step-size parameter  $\gamma$  and a momentum parameter  $\beta$ :

$$\begin{cases} y_t &= x_t + \beta(x_t - x_{t-1}), \\ x_{t+1} &= y_t - \gamma \nabla f(y_t). \end{cases} \quad (\text{NAG})$$




Fig. 2. Nesterov Accelerated gradient (NAG). Green area: set of parameters  $(\gamma, \beta) \in \Omega_{\text{NAG}}$  for which a Lyapunov function exists; Red area: set of  $(\gamma, \beta) \in \Omega_{\text{NAG}}$  for which (NAG) cycles on at least one function in  $\mathcal{F}_{0,L}$ .

(NAG) is also written as follows  $y_{t+1} = (1 + \beta)(y_t - \gamma \nabla f(y_t)) - \beta(y_{t-1} - \gamma \nabla f(y_{t-1}))$ , and is therefore also an order-2 (SFOM).

**Set  $\Omega_{\text{NAG}}$  of parameters of interest:** As for HB, we consider the set of  $\beta$  and  $\gamma$  for which (NAG) converges on  $\mathcal{Q}_{0,L}$ . This corresponds to considering all  $\beta, \gamma$  verifying  $0 \leq \beta \leq 1$  and  $0 \leq \gamma \leq \frac{2}{L} \frac{1+\beta}{1+2\beta}$ .

**Interpretation:** (NAG) is known to converge, with an accelerated rate, on  $\mathcal{F}_{\mu,L}$ , for the tuning  $(\gamma, \beta) = (\frac{1}{L}, \frac{\sqrt{L}-\sqrt{\mu}}{\sqrt{L}+\sqrt{\mu}})$ , that optimizes the convergence rate on  $\mathcal{Q}_{\mu,L}$ . For this reason, (NAG) is considered to be more “robust” than HB.

Figure 2 shows that (NAG) admits a Lyapunov function for almost any parameters in  $\Omega_{\text{NAG}}$ . Moreover, our methodology does not detect any set of parameters at which a cycle of length  $K \in \llbracket 2; 25 \rrbracket$  exists, apart on the boundary  $\{(\gamma, \beta), \gamma = \frac{2}{L} \frac{1+\beta}{1+2\beta}\}$ . On the boundary, cycles of length 2 are observed, whose existences are theoretically verified on one-dimensional quadratic functions. This illustrates the robustness of our methodology when very few cycles exist.

### C. Inexact gradient method

Next, we consider the inexact gradient method, parameterized by  $\gamma$  and  $\varepsilon$  and the update

$$\begin{aligned} \text{Get } \mathcal{O}(x_t) &= d_t \text{ such that } \|d_t - \nabla f(x_t)\| \leq \varepsilon \|\nabla f(x_t)\|, \\ x_{t+1} &= x_t - \gamma d_t. \end{aligned} \quad (\text{IGD})$$

(IGD) is thus an (SFOM) of order 1.

**Set  $\Omega_{\text{IGD}}$  of parameters of interest:** Since the exact gradient method converges only for  $\gamma < \frac{2}{L}$ , we only consider such steps-sizes. Moreover,  $\varepsilon \geq 1$  allows  $d_t = 0$  and thereby does not make much sense. This motivates considering the set  $\Omega_{\text{IGD}} = \{(\gamma, \varepsilon) \in [0; \frac{2}{L}] \times [0, 1]\}$ .

**Interpretation:** We search for cycles of length  $K \in \llbracket 2; 25 \rrbracket$  and use color intensity to show the minimal cycle length. (IGD) is known to converge for any  $\gamma \leq \frac{2}{L(1+\varepsilon)}$  (see [9], [15]). Figure 3 shows that the complementary of this region of convergence is completely filled by parameters allowing

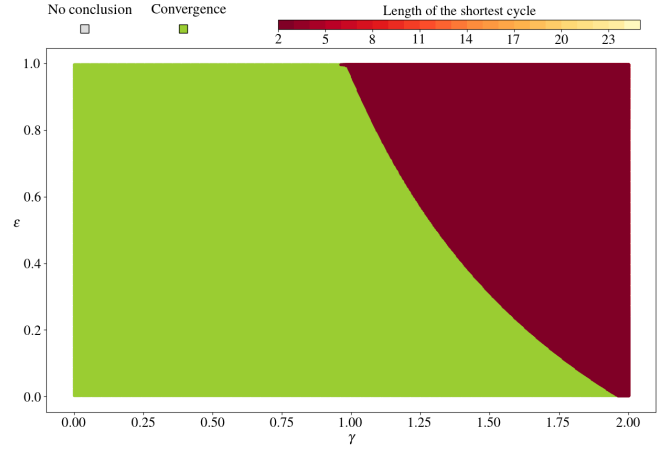


Fig. 3. Inexact GD (IGD). Green area: set of parameters  $(\gamma, \varepsilon) \in \Omega_{\text{IGD}}$  for which a Lyapunov function exists; Red area: set of  $(\gamma, \varepsilon) \in \Omega_{\text{IGD}}$  for which (IGD) cycles on at least one function in  $\mathcal{F}_{0,L}$ .

cycles, showing that no other parameters values than the known ones allow obtaining worst-case convergence of (IGD).

### D. Three-operator splitting

The three-operator splitting (TOS) method, introduced by [8], aims at solving the inclusion problem  $0 \in Ax + Bx + \partial f(x)$ , where  $A$  is a monotone operator,  $B$  is a co-coercive operator and  $\partial f$  denotes the differential of the smooth (strongly) convex function  $f$ . It corresponds to the following update, for a step-size parameter  $\gamma$ , a smoothing parameter  $\alpha$  and an update parameter  $\beta$ :

$$\begin{cases} x_{t+1} &= J_{\alpha B}(w_t), \\ y_{t+1} &= J_{\alpha A} \left( 2x_{t+1} - w_t - \frac{\gamma}{\beta} \nabla f(x_{t+1}) \right), \\ w_{t+1} &= w_t - \beta(x_{t+1} - y_{t+1}), \end{cases} \quad (\text{TOS})$$

where  $J_O$  denotes the resolvent of the operator  $O$ , i.e.  $J_O = (I + O)^{-1}$ . Note that (TOS) is therefore an order-1 (SFOM).




Fig. 4. Three-operator splitting (TOS). Green area: set of parameters  $(\gamma, \beta) \in \Omega_{\text{TOS}}$  for which a Lyapunov function exists; Red area: set of  $(\gamma, \beta) \in \Omega_{\text{TOS}}$  for which (TOS) cycles on at least a triplet of operators.

**Set  $\Omega_{\text{TOS}}$  of parameters of interest:** When considering  $A$ ,  $B$  and  $\nabla f$  to be linear symmetric and co-diagonalisable operators, the set of convergence of (TOS) is  $\Omega_{\text{TOS}} = \{(\gamma, \beta) \in [0, \frac{2}{L}] \times [0, 2]\}$ , which we therefore consider.

**Interpretation:** We search for cycles of length  $K \in \llbracket 2, 25 \rrbracket$  and use color intensity to show the minimal cycle length. Interestingly, there is a gap between the green region and the red ones. Unlike for (HB), it seems that increasing the length of the cycle does not help covering this gap and shows that some algorithms might have no Lyapunov function while not cycling. Understanding the behavior of (TOS) in the grey region is therefore still an open question.

## V. CONCLUSION

This work proposes a systematic approach for finding counter-examples to convergence of first-order methods, bringing a complementary tool to the existing systematic techniques for finding convergence guarantees (that include certifications through the existence of Lyapunov functions). Our approach is based on now classical tools and techniques used in the field of first-order optimization and a few existing packages [18], [32] allows for straightforward implementations of our methodology.

**ACKNOWLEDGMENTS.**The authors thank Margaux Zaffran for her feedbacks and fruitful discussions and her assistance in making plots. The work of B. Goujaud and A. Dieuleveut is partially supported by ANR-19-CHIA-0002-01/chaire SCAI, and Hi!Paris. A. Taylor acknowledges support from the European Research Council (grant SEQUOIA 724063). This work was partly funded by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

## REFERENCES

- [1] MOSEK ApS. *MOSEK Optimizer API for C 9.3.6*, 2019.
- [2] Yossi Arjevani, Shai Shalev-Shwartz, and Ohad Shamir. On lower and upper bounds in smooth and strongly convex optimization. *The Journal of Machine Learning Research*, 17(1):4303–4353, 2016.
- [3] Nikhil Bansal and Anupam Gupta. Potential-function proofs for gradient methods. *Theory of Computing*, 15(1):1–32, 2019.
- [4] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- [5] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- [6] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. and Trends in Machine Learning*, 8(3-4):231–357, 2015.
- [7] Antonin Chambolle and Thomas Pock. An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319, 2016.
- [8] Damek Davis and Wotao Yin. A three-operator splitting scheme and its optimization applications. *Set-valued and variational analysis*, 25:829–858, 2017.
- [9] Etienne De Klerk, Francois Glineur, and Adrien B Taylor. Worst-case convergence analysis of inexact gradient and newton methods through semidefinite programming performance estimation. *SIAM Journal on Optimization*, 30(3):2053–2082, 2020.
- [10] Radu-Alexandru Dragomir, Adrien B. Taylor, Alexandre d’Aspremont, and Jérôme Bolte. Optimal complexity and certification of Bregman first-order methods. *Math. Programming*, pages 1–43, 2021.
- [11] Yoel Drori and Marc Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. *Math. Programming*, 145(1):451–482, 2014.
- [12] Pavel Dvurechensky, Shimrit Shtern, and Mathias Staudigl. First-order methods for convex optimization. *EURO Journal on Computational Optimization*, 9, 2021.
- [13] Alexandre d’Aspremont, Damien Scieur, and Adrien Taylor. Acceleration methods. *Found. and Trends in Optimization*, 5(1-2), 2021.
- [14] Donald A. Flanders and George Shortley. Numerical determination of fundamental modes. *Journal of Applied Physics*, 21(12), 1950.
- [15] Oran Gannot. A frequency-domain analysis of inexact gradient methods. *Math. Programming*, 194(1-2):975–1016, 2022.
- [16] Euhanna Ghadimi, Hamid Reza Feyzmahdavian, and Mikael Johansson. Global convergence of the heavy-ball method for convex optimization. In *2015 European control conference (ECC)*, pages 310–315, 2015.
- [17] Baptiste Goujaud, Aymeric Dieuleveut, and Adrien Taylor. Counter-examples in first-order optimization: a constructive approach. *preprint arXiv:2303.10503*, 2023. (With appendices).
- [18] Baptiste Goujaud, Céline Moucer, François Glineur, Julien Hendrickx, Adrien Taylor, and Aymeric Dieuleveut. PEPit: computer-assisted worst-case analyses of first-order optimization methods in Python. *preprint arXiv:2201.04040*, 2022.
- [19] Baptiste Goujaud, Adrien Taylor, and Aymeric Dieuleveut. Optimal first-order methods for convex functions with a quadratic upper bound. *arXiv preprint arXiv:2205.15033*, 2022.
- [20] Baptiste Goujaud, Adrien Taylor, and Aymeric Dieuleveut. Provable non-accelerations of the heavy-ball method. *arXiv preprint arXiv:2307.11291*, 2023.
- [21] Charles Guille-Escuret, Baptiste Goujaud, Adam Ibrahim, and Ioannis Mitliagkas. Gradient descent is optimal under lower restricted secant inequality and upper error bound. In *NeurIPS*, 2022.
- [22] R. E. Kalman and J. E. Bertram. Control System Analysis and Design Via the “Second Method” of Lyapunov: II—Discrete-time systems. *Journal of Basic Engineering*, 82(2):394–400, 06 1960.
- [23] R. E. Kalman and J. E. Bertram. Control System Analysis and Design Via the “Second Method” of Lyapunov: I—Continuous-Time Systems. *Journal of Basic Engineering*, 82(2):371–393, 06 1960.
- [24] Cornelius Lanczos. Solution of systems of linear equations by. *Journal of research of the National Bureau of Standards*, 49(1):33, 1952.
- [25] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- [26] Aleksandr Mikhailovich Lyapunov and A.T. Fuller. The general problem of the stability of motion. *International journal of control*, 55(3):531–534, 1992. Original text in Russian, 1892.
- [27] Yurii Nesterov. A method of solving a convex programming problem with convergence rate  $O(1/k^2)$ . *Soviet Mathematics Doklady*, 27(2):372–376, 1983.
- [28] Yurii Nesterov. *Introductory Lectures on Convex Optimization*. Springer, 2003.
- [29] Boris T. Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4), 1963.
- [30] Adrien Taylor, Bryan Van Scoy, and Laurent Lessard. Lyapunov functions for first-order methods: Tight automated convergence guarantees. In *International Conference on Machine Learning*, 2018.
- [31] Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Exact worst-case performance of first-order methods for composite convex optimization. *SIAM Journal on Optimization*, 27(3):1283–1313, 2017.
- [32] Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Performance estimation toolbox (PESTO): automated worst-case analysis of first-order optimization methods. In *56th Annual Conference on Decision and Control (CDC)*, pages 1278–1283, 2017.
- [33] Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *Math. Programming*, 161(1-2):307–345, 2017.
- [34] Manu Upadhyaya, Sebastian Banert, Adrien B. Taylor, and Pontus Giselsson. Automated tight Lyapunov analysis for first-order methods. *preprint arXiv:2302.06713*, 2023.
- [35] David Young. On Richardson’s method for solving linear systems with positive definite matrices. *Journal of Mathematics and Physics*, 32(1-4):243–255, 1953.