

FIRST-ORDER OPTIMIZATION ALGORITHMS VIA DISCRETIZATION OF FINITE-TIME CONVERGENT FLOWS

Anonymous authors

Paper under double-blind review

ABSTRACT

In this paper the performance of several discretization algorithms for two first-order finite-time optimization flows. These flows are, namely, the rescaled-gradient flow (RGF) and the signed-gradient flow (SGF), and consist of non-Lipschitz or discontinuous dynamical systems that converge locally in finite time to the minima of gradient-dominated functions. We introduce three discretization methods for these first-order finite-time flows, and provide convergence guarantees. We then apply the proposed algorithms in training neural networks and empirically test their performances on three standard datasets, namely, CIFAR10, SVHN, and MNIST. Our results show that our schemes demonstrate faster convergences against standard optimization alternatives, while achieving equivalent or better accuracy.

1 INTRODUCTION

Consider the unconstrained minimization problem for a given cost function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. When f is sufficiently regular, the standard algorithm in continuous time (dynamical system) is given by

$$\dot{x} = F_{\text{GF}}(x) \triangleq -\nabla f(x) \quad (1)$$

with $\dot{x} \triangleq \frac{d}{dt}x(t)$, known as the *gradient flow* (GF). Generalizing GF, the *q-rescaled* GF (*q-RGF*) Wibisono et al. (2016) given by

$$\dot{x} = -c \frac{\nabla f(x)}{\|\nabla f(x)\|_2^{\frac{q-2}{q-1}}} \quad (2)$$

with $c > 0$ and $q \in (1, \infty]$ has an asymptotic convergence rate $f(x(t)) - f(x^*) = \mathcal{O}\left(\frac{1}{t^{q-1}}\right)$ under mild regularity, for $\|x(0) - x^*\| > 0$ small enough, where $x^* \in \mathbb{R}^n$ denotes a local minimizer of f . However, we recently proved Romero & Benosman (2020) that the *q-RGF*, as well as our proposed *q-signed* GF (*q-SGF*)

$$\dot{x} = -c \|\nabla f(x)\|_1^{\frac{1}{q-1}} \text{sign}(\nabla f(x)), \quad (3)$$

where $\text{sign}(\cdot)$ denotes the sign function, applied element-wise for (real-valued) vectors, are both finite-time convergent, provided that f is gradient dominated of order $p \in (1, q)$. In particular, if f is strongly convex, then *q-RGF* and *q-SGF* is finite-time convergent for any $q \in (2, \infty]$, since f must be gradient dominated of order $p = 2$.

CONTRIBUTION

In this paper, we explore three discretization schemes for the *q-RGF* (2) and *q-SGF* (3) and provide some convergence guarantees using results from hybrid dynamical control theory. In particular, we explore a forward-Euler discretization of RGF/SGF, followed by an explicit Runge-Kutta discretization, and finally a novel Nesterov-like discretization. We then test their performance on both synthetic and real-world data in the context of deep learning, namely, over the well-known datasets CIFAR10, SVHN, and MNIST.

RELATED WORK

Propelled by the work of Wang & Elia (2011) and Su et al. (2014), there has been a recent and significant research effort dedicated to analyzing optimization algorithms from the perspective of

dynamical systems and control theory, especially in continuous time Wibisono et al. (2016); Wilson (2018); Lessard et al. (2016); Fazlyab et al. (2017b); Scieur et al. (2017); França et al. (2018); Fazlyab et al. (2018); Fazlyab et al. (2018); Taylor et al. (2018); França et al. (2019a); Orvieto & Lucchi (2019); Muehlebach & Jordan (2019). A major focus within this initiative is in *acceleration*, both in terms of trying to gain new insight into more traditional optimization algorithms from this perspective, or even to exploit the interplay between continuous-time systems and their potential discretizations for novel algorithm design Muehlebach & Jordan (2019); Fazlyab et al. (2017a); Shi et al. (2018); Zhang et al. (2018); França et al. (2019b); Wilson et al. (2019). Many of these papers also focus on explicit mappings and matchings of convergence rates from the continuous-time domain into discrete time.

For older work connecting ordinary differential equations (ODEs) and their numerical analysis, with optimization algorithms, see Botsaris (1978a;b); Zghier (1981); Snyman (1982; 1983); Brockett (1988); Brown (1989). In Helmke & Moore (1994), the authors studied relationships between linear programming, ODEs, and general matrix theory. Further, Schropp (1995) and Schropp & Singer (2000) explored several aspects linking nonlinear dynamical systems to gradient-based optimization, including nonlinear constraints.

Tools from Lyapunov stability theory are often employed for this purpose, mainly because there already exists a rich body of work within the nonlinear systems and control theory community for this purpose. In particular, typically in previous works, one seeks asymptotically Lyapunov stable gradient-based systems with an equilibrium (stationary point) at an isolated extremum of the given cost function, thus certifying local convergence. Naturally, global asymptotic stability leads to global convergence, though such an analysis will typically require the cost function to be strongly convex everywhere.

For physical systems, a Lyapunov function can often be constructed from first principles via some physically meaningful measure of energy (*e.g.*, total energy = potential energy + kinetic energy). In optimization, the situation is somewhat similar in the sense that a suitable Lyapunov function may often be constructed by taking simple surrogates of the objective function as candidates. For instance, $V(x) \triangleq f(x) - f(x^*)$ can be a good initial candidate. Further, if f is continuously differentiable and x^* is an isolated stationary point, then another alternative is $V(x) \triangleq \|\nabla f(x)\|^2$. However, most fundamental and applied research conducted in systems and control regarding Lyapunov stability theory deals exclusively with *continuous-time* systems. Unfortunately, (dynamical) stability properties are generally not preserved for simple forward-Euler and sample-and-hold discretizations and control laws Stuart & Humphries (1998). Furthermore, practical implementations of optimization algorithms in modern digital computers demand discrete-time. Nonetheless, it has been extensively noted that a vast amount of general Lyapunov-based results appear to have a discrete-time equivalent.

In that sense, we aim here to start from the q -RGF, and q -SGF continuous flows, characterized by their Lyapunov-based finite-time convergence, and seek discretization schemes, which allow us to ‘shadow’ the solutions of these flows in discrete time, hoping to achieve an acceleration of the discrete methods inspired from the finite-time convergence characteristics of the continuous flows.

2 OPTIMIZATION ALGORITHMS AS DISCRETE-TIME SYSTEMS

Generalizing (1), (2), and (3), consider a continuous-time algorithm (dynamical system) modeled via an ordinary differential equation (ODE)

$$\dot{x} = F(x) \tag{4}$$

for $t \geq 0$, or, more generally, a differential inclusion

$$\dot{x}(t) \in \mathcal{F}(x(t)) \tag{5}$$

a.e. $t \geq 0$ (*e.g.* for the $q = \infty$ case), such that $x(t) \rightarrow x^*$ as $t \rightarrow t^*$. In the case of the q -RGF (2) and q -SGF (3) for f gradient dominated of order $p \in (1, q)$, we have finite-time convergence, and thus $t^* = t^*(x(0)) < \infty$.

Most of the popular numerical optimization schemes can be written in a state-space form (*i.e.*, recursively), as

$$X_{k+1} = F_d(k, X_k) \tag{6a}$$

$$x_k = G(X_k) \tag{6b}$$

for $k \in \mathbb{Z}_+ \triangleq \{0, 1, 2, \dots\}$ and a given $X_0 \in \mathbb{R}^m$ (typically $m \geq n$), where $F_d : \mathbb{Z}_+ \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $G : \mathbb{R}^m \rightarrow \mathbb{R}^n$.

Naturally, (6) can be seen as a discrete-time dynamical system constructed by discretizing (4) in time. In particular, we have $x_k \approx x(t_k)$, where $\{0 = t_0 < t_1 < t_2 < \dots\}$ denotes a time partition and $x(\cdot)$ a solution to (4) or (5) as appropriate. Therefore, we call X_k and x_k , respectively, the *state* and *output* at time step k . Whenever $F_d(k, X)$ does not depend on k , we will drop k and thus write $F_d(X)$ instead. Whenever $m = n$, we will denote take $G(X) \triangleq X$ and replace X and X_k by x and x_k , respectively.

Example 1. The standard gradient descent (GD) algorithm

$$x_{k+1} = x_k - \eta \nabla f(x_k) \quad (7)$$

with step size (learning rate) $\eta > 0$ can be readily written in the form (6) by taking $m = n$, $F_d(x) \triangleq x - \eta \nabla f(x)$, and $G(x) \triangleq x$.

- If the step sizes are adaptive, *i.e.* if we replace η by a sequence $\{\eta_k\}$ with $\eta_k > 0$, then we only need to replace $F_d(k, x) \triangleq x - \eta_k \nabla f(x)$, provided that $\{\eta_k\}$ is not computed using feedback from $\{x_k\}$ (*e.g.* through a line search method).
- If we do wish to use *feedback*¹ (and no memory past the most recent output and step size), then we can set $m = n + 1$, $G([x; \eta]) \triangleq x$, and $F_d([x; \eta]) \triangleq [F_d^{(1)}([x; \eta]); F_d^{(2)}([x; \eta])]$, where $F_d^{(1)}([x; \eta]) \triangleq x - \eta \nabla f(x)$, and $F_d^{(2)}$ is a user-defined function that dictates the updates in the step size. In particular, an open-loop (no feedback) adaptive step size $\{\eta_k\}$ may also be achieved under this scenario, provided that it is possible to write $\eta_{k+1} = F_d^{(2)}(\eta_k)$. If this is not possible (and still open-loop step size), then we may take $F_d^{(2)}(k, X) \triangleq \eta_{k+1}$, and of course add a k -argument in F_d .
- If we wish to use individual step sizes for each the n components of $\{x_k\}$, then it suffices to take η_k as an n -dimensional vector (thus $m = 2n$), and make appropriate changes in F_d and G .

In each of these cases, GD can be seen as a forward-Euler discretization of the GF (1), *i.e.*,

$$x_{k+1} = x_k + \Delta t_k F_{\text{GF}}(x_k) \quad (8)$$

with $F_{\text{GF}} = -\nabla f$ and adaptive time step $\Delta t_k \triangleq t_{k+1} - t_k$ chosen as $\Delta t_k = \eta_k$.

Example 2. The proximal point algorithm (PPA)

$$x_{k+1} = \arg \min_{x \in \mathbb{R}^n} \left\{ f(x) + \frac{1}{2\eta_k} \|x - x_k\|^2 \right\} \quad (9)$$

with step size $\eta_k > 0$ (open loop, for simplicity) can also be written in the form (6), by taking $m = n$, $F_d(k, x) \triangleq \arg \min_{x' \in \mathbb{R}^n} \{f(x') + \frac{1}{2\eta_k} \|x' - x\|^2\}$, and $G(x) \triangleq x$. Naturally, we need to assume sufficient regularity for $F_d(k, x)$ to exist and we must design a consistent way to choose $F_d(k, x)$ when multiple minimizers exist in the definition of $F_d(k, x)$. Alternatively, these conditions must be satisfied, at the very least, at every $(k, x) \in \{(0, x_0), (1, x_1), (2, x_2), \dots\}$ for a particular chosen initial $x_0 \in \mathbb{R}^n$.

Assuming sufficient regularity, we have $\nabla_x \{f(x) + \frac{1}{2\eta_k} \|x - x_k\|^2\}|_{x=x_{k+1}} = 0$, and thus

$$\nabla f(x_{k+1}) + \frac{1}{\eta_k} (x_{k+1} - x_k) = 0 \iff x_{k+1} = x_k + \Delta t_k F_{\text{GF}}(x_{k+1}) \quad (10)$$

with $\Delta t_k = \eta_k$, which is precisely the backward-Euler discretization of the GF (1).

¹Also known as *closed-loop* design in control-theoretic and reinforcement learning terminology, meaning that $\eta_k = \varphi(k, x_k)$ for some $\varphi : \mathbb{Z}_+ \times \mathbb{R}^n \rightarrow \mathbb{R}_+$ that does not depend on $\{X_0, X_1, X_2, \dots\}$. On the other hand, *open-loop* design can be seen as closed loop with $\varphi(k, \cdot)$ constant for each $k \in \mathbb{Z}_+$.

Example 3. The Nesterov accelerated gradient descent (N-AGD)

$$y_k = x_k + \beta_k(x_k - x_{k-1}) \quad (11a)$$

$$x_{k+1} = y_k - \eta_k \nabla f(y_k) \quad (11b)$$

with step size $\eta_k > 0$ and momentum coefficient $\beta_k > 0$ (both open loop, for simplicity), can be written in the form (6) by taking $m = 2n$,

$$F_d(k, [y; x]) \triangleq \begin{bmatrix} (1 + \beta_{k+1})(y - \eta_k \nabla f(y)) - \beta_{k+1}x \\ y - \eta_k \nabla f(y) \end{bmatrix} \quad (12)$$

and $G([y; x]) \triangleq x$ for $y, x \in \mathbb{R}^n$. In other words, $X_k = [y_k; x_k]$. Traditionally, $\beta_k = \frac{k-1}{k-2}$, but clearly, if we set $\eta_k = \eta > 0$ and $\beta_k = \beta \in (0, 1)$ (in practice, $\eta \approx 0$ and $\beta \approx 1$), then we can drop k from $F_d(k, [y; x])$.

There exist a few approaches in the literature on the interpretation of N-AGD (11b) as the discretization of a second-order continuous-time dynamical system, namely via a vanishing step size argument Su et al. (2014), or via symplectic Euler schemes of crafted Hamiltonian systems Muehlebach & Jordan (2019); França et al. (2019b).

3 PROPOSED ALGORITHMS VIA DISCRETIZATION

In this section, we propose three classes of optimization algorithms via discretization of the q -RGF (2) and q -SGF (3). But first, we review the necessary conditions to ensure finite-time convergence of these flows.

Given $q \in (1, \infty]$, let $F_{q\text{-RGF}}(x)$ and $F_{q\text{-SGF}}(x)$ be defined, respectively, by the RHS of (2) and (3). The hyperparameter $c > 0$ is not explicitly denoted in $F_{q\text{-RGF}}, F_{q\text{-SGF}}$. Next, borrowing terminology from Wilson et al. (2019), we say that f (assumed continuously differentiable) is μ -gradient dominated of order $p \in (1, \infty]$ (with $\mu > 0$) near the local minimizer x^* if

$$\frac{p-1}{p} \|\nabla f(x)\|^{\frac{p}{p-1}} \geq \mu^{\frac{1}{p-1}} (f(x) - f^*) \quad (13)$$

for every $x \in \mathbb{R}^n$ near $x = x^*$, where $f^* = f(x^*)$. When $\mu > 0$ is unknown or unimportant, but known to exist, we will omit it in the previous definition. It can be proved that continuously differentiable strongly convex functions are gradient dominated of order $p = 2$. Furthermore, if f is gradient dominated (of any order) w.r.t. x^* , then x^* is an isolated stationary point of f .

Remark 1. For strongly convex functions, gradient dominance of order $p = 2$ can be established. In fact, gradient dominance is usually defined exclusively for order $p = 2$, often referred to as the Polyak-Łojasiewicz (PL) inequality, which was introduced by Polyak (1963) to relax the (strong) convexity assumption commonly used to show convergence of the GD algorithm (7). The PL inequality can also be used to relax convexity assumptions of similar gradient and proximal-gradient methods Karimi et al. (2016); Attouch & Bolte (2009). Our adopted generalized notion of gradient dominance is strongly tied to the Łojasiewicz gradient inequality from real analytic geometry, established by Łojasiewicz (1963; 1965)² independently and simultaneously from Polyak (1963), and generalizing the PL inequality. More precisely, this inequality is typically written as

$$\|\nabla f(x)\| \geq C \cdot |f(x) - f^*|^\theta \quad (14)$$

for every $x \in \mathbb{R}^n$ in a small enough open neighborhood of the stationary point $x = x^*$, for some $C > 0$ and $\theta \in (\frac{1}{2}, 1]$. This inequality is guaranteed for analytic functions Łojasiewicz (1965). More precisely, when x^* is a local minimizer of f , the aforementioned relationship is explicitly given by

$$C = \left(\frac{p}{p-1} \right)^{\frac{p-1}{p}} \mu^{\frac{1}{p}}, \quad \theta = \frac{p-1}{p}. \quad (15)$$

Therefore, analytic functions are always gradient dominated. However, while analytic functions are always smooth, smoothness is not required to attain gradient dominance.

²For more modern treatments in English, see Łojasiewicz & Zullo (1999); Bolte et al. (2007)

We are now ready to state the finite-time convergence of the q -RGF (2) and q -SGF (3).

Theorem 1. *Romero & Benosman (2020)* Suppose that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and μ -gradient dominated of order $p \in (1, \infty)$ near a strict local minimizer $x^* \in \mathbb{R}^n$. Let $c > 0$ and $q \in (p, \infty]$. Then, any maximal solution $x(\cdot)$, in the sense of Filippov, to the q -RGF (2) or q -SGF (3) will converge in finite time to x^* , provided that $\|x(0) - x^*\| > 0$ is sufficiently small. More precisely, $\lim_{t \rightarrow t^*} x(t) = x^*$, where the convergence time $t^* < \infty$ may depend on which flow is used, but in both cases is upper bounded by

$$t^* \leq \frac{\|\nabla f(x_0)\|^{\frac{1}{\theta} - \frac{1}{\theta'}}}{cC^{\frac{1}{\theta}} \left(1 - \frac{\theta}{\theta'}\right)}, \quad (16)$$

where $x_0 = x(0)$, $C = \left(\frac{p}{p-1}\right)^{\frac{p-1}{p}} \mu^{\frac{1}{p}}$, $\theta = \frac{p-1}{p}$, and $\theta' = \frac{q-1}{q}$. In particular, given any compact and positively invariant subset $S \subset \mathcal{D}$, both flows converge in finite time with the aforementioned convergence time upper bound (which can be tightened by replacing \mathcal{D} with S) for any $x_0 \in S$. Furthermore, if $\mathcal{D} = \mathbb{R}^n$, then we have global finite-time convergence, i.e. finite-time convergence to any maximal solution (in the sense of Filippov) $x(\cdot)$ with arbitrary $x_0 \in \mathbb{R}^n$.

In essence, the analysis (introduced in Romero & Benosman (2020)) consists of leveraging the gradient dominance to show that the energy function $\mathcal{E}(t) \triangleq f(x(t)) - f^*$ satisfies the Lyapunov-like differential inequality $\dot{\mathcal{E}}(t) = \mathcal{O}(\mathcal{E}(t)^\alpha)$ for some $\alpha < 1$. The detailed proof is recalled in Appendix C for completeness.

3.1 MAIN RESULT 1: DISCRETIZATION ALGORITHMS AND THEIR CONVERGENCE ANALYSIS

3.1.1 FORWARD-EULER DISCRETIZATION

First, we propose a simple forward Euler discretization of the finite-time convergent flows

$$x_{k+1} = x_k + \eta F(x_k), \quad \eta > 0 \quad (17)$$

where $F \in \{F_{q\text{-RGF}}, F_{q\text{-SGF}}\}$. We show later, in Theorem 2, that this simple method leads, for small enough $\eta > 0$, to solutions that are ϵ -close to the finite-time flows.

3.1.2 EXPLICIT RUNGE-KUTTA DISCRETIZATION

We propose to use the following discretization

$$x_{k+1} = x_k + \eta \sum_{i=1}^K \alpha_i F(y_k^i), \quad y_k^1 = x_k, \quad \sum_{i=1}^K \alpha_i = 1, \quad (18a)$$

$$y_k^i = x_k + \eta \sum_{j=1}^{i-1} \beta_j F(y_k^j), \quad i > 1, \quad (18b)$$

for $\eta > 0$, $K \in \{1, 2, 3, \dots\}$, and $F \in \{F_{q\text{-RGF}}, F_{q\text{-SGF}}\}$. This method is well-known to be numerically stable under the consistency condition $\sum_{i=1}^K \alpha_i = 1$, Stuart & Humphries (1996). However, in our optimization framework, we want to be able to guarantee that the stable numerical solution of (18) remains close to the solution of the continuous flows. In other words, we seek arbitrarily small global error, also known as *shadowing*. This will be discussed in Theorem 2.

3.1.3 NESTEROV-LIKE DISCRETIZATION

First, we rewrite Nesterov's accelerated GD as

$$x_{k+1} = x_k + \beta y_k - \eta \nabla f(x_k + \beta y_k) \quad (19a)$$

$$y_{k+1} = x_{k+1} - x_k, \quad (19b)$$

where y_k now serves as a momentum term. We argue that Nesterov's acceleration can be interpreted as actually applying the discretization given by (19) to the GF (1), i.e., by seeing the term $-\eta \nabla f(x_k + \beta y_k)$ as a mapping applied to the GF flow (1) at $x_k + \beta y_k$, as $-\eta F_{GF}(x_k + \beta y_k)$.

Therefore, given any optimization flow represented by the continuous-time system $\dot{x} = F(x)$, locally convergent to a local minimizer $x^* \in \mathbb{R}^n$ of a cost function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we can replicate Nesterov’s acceleration of (1). More precisely, we obtain the algorithm

$$x_{k+1} = x_k + \eta F(x_k + \beta y_k) + \beta y_k \quad (20a)$$

$$y_{k+1} = x_{k+1} - x_k. \quad (20b)$$

Based on this idea, we propose two ‘Nesterov-like’ discrete optimization algorithms. The first one based on the q -RGF continuous flow, is defined as:

$$x_{k+1} = x_k + \eta \left(-c \frac{\nabla f(x_k + \beta y_k)}{\|\nabla f(x_k + \beta y_k)\|^{\frac{q-2}{q-1}}} \right) + \beta y_k \quad (21a)$$

$$y_{k+1} = x_{k+1} - x_k. \quad (21b)$$

The second algorithm is based on the q -SGF continuous flow, and is given by:

$$x_{k+1} = x_k + \eta \left(-c \|\nabla f(x_k + \beta y_k)\|_1^{\frac{1}{q-1}} \text{sign}(\nabla f(x_k + \beta y_k)) \right) + \beta y_k \quad (22a)$$

$$y_{k+1} = x_{k+1} - x_k. \quad (22b)$$

3.1.4 CONVERGENCE ANALYSIS

We present here some convergence results of the three proposed discretizations. The analysis summarized in Theorem 2 is based on tools from hybrid control theory, and is detailed in Appendix D³.

Theorem 2. *Suppose that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable, locally L_f -Lipschitz, and μ -gradient dominated of order $p \in (2, \infty)$ in a compact neighborhood S of a strict local minimizer $x^* \in \mathbb{R}^n$. Let $c > 0$ and $q \in (p, \infty]$. Then, for a given initial condition $x_0 \in S$ any maximal solution $x(t)$, $x(0) = x_0$, (in the sense of Filippov) to the q -RGF given by (2) or the q -SGF flow (3), there exists an arbitrarily small $\epsilon > 0$ such that the solution x_k of any of the discrete algorithms (17), (18), (21), or (22), with sufficiently small $\eta > 0$, are ϵ -close to $x(t)$, i.e., $\|x_k - x\| \leq \epsilon$, and s.t. the following convergence bound holds*

$$\|f(x_k) - f(x^*)\| \leq L_f \epsilon + [(f(x_0) - f(x^*))^{(1-\alpha)} - \tilde{c}(1-\alpha)\eta k]^{1/(1-\alpha)}, \quad L_f > 0, \quad k \leq k^*, \quad (23)$$

where $\alpha = \frac{\theta}{\theta'}$, $\theta = \frac{p-1}{p}$, $\theta' = \frac{q-1}{q}$, $\tilde{c} = c \left(\left(\frac{p}{p-1} \right)^{\frac{p-1}{p}} \mu^{\frac{1}{p}} \right)^{\frac{1}{\theta'}}$, and $k^* = \frac{(f(x_0) - f(x^*))^{(1-\alpha)}}{\tilde{c}(1-\alpha)\eta}$.

Theorem 2 thus shows that that ϵ -convergence of $x_k \rightarrow x^*$ can be achieved in a finite number of steps upper bounded by $k^* = \frac{(f(x_0) - f(x^*))^{(1-\alpha)}}{\tilde{c}(1-\alpha)\eta}$. This is a preliminary convergence result, which is meant to show the existence of discrete solutions obtained via the proposed discretization algorithms, which are ϵ -close to the continuous solutions of the finite-time flows. We also underline here that after x_k reaches an ϵ -neighborhood of x^* , then $x_{k+1} \approx x_k$, $\forall k > k^*$, since x^* is an equilibrium point of the continuous flows; see Definition 2 in Appendix B.

4 MAIN RESULT 2: NUMERICAL EXPERIMENTS

4.1 NUMERICAL TESTS ON AN ACADEMIC EXAMPLE

Let us show first on a simple numerical example that the acceleration in convergence, proven in continuous time for certain range of the hyperparameters, can translate to some convergence acceleration in discrete time, as shown in Theorem 2. We consider the Rosenbrock function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, given by $f(x_1, x_2) = (a - x_1)^2 + b(x_2 - x_1^2)^2$, with parameters $a, b \in \mathbb{R}$. This function admits exactly one stationary point $(x_1^*, x_2^*) = (a, a^2)$ for $b \geq 0$, and is locally strongly convex, hence locally satisfies gradient dominance of order $p = 2$, which allows us to select $q > 2$ in q -RGF

³Note that there might be several ways of approaching this proof. For instance, one could follow the general results on stochastic approximation of set-valued dynamical systems, using the notion of perturbed solutions to differential inclusions presented in Benaïm et al. (2005).

and q -SGF to achieve finite-time convergence in continuous-time. We report in Figure 1 the mean performance of all three discretizations for q -RGF and q -SGF⁴ with fixed step size⁵, for several values of q , for 10 random initial conditions in $[0, 2]$. We observe for all three discretizations that, as expected from the continuous flow analysis, for q close to 2, q -RGF behaves similar to GD in terms of convergence rate, whereas for $q > 2$ the finite-time convergence in continuous time seems to translate to some acceleration in this simple discretization method. Similarly for q -SGF, q closer to 2 translates to less accelerated algorithms, with a behavior similar to GD, whereas larger q values lead to accelerated convergence.

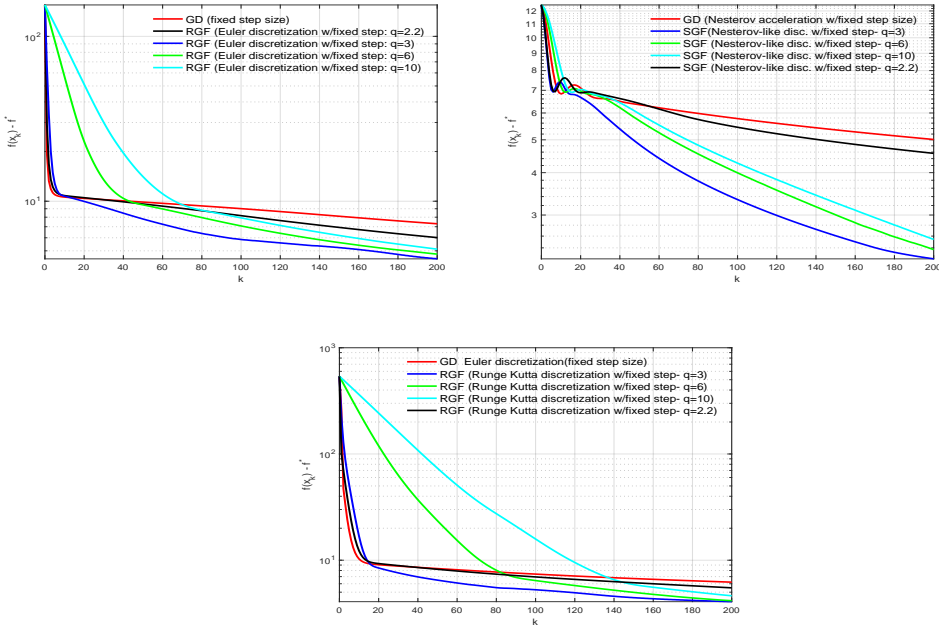


Figure 1: Example of the proposed discretization algorithms of finite-time q -RGF and q -SGF

4.2 NUMERICAL EXPERIMENTS ON REAL-WORLD DATA

We report here the results of our experiments using deep neural network (DNN) training on three well-known datasets, namely, CIFAR10, MNIST, and SVHN. We report results on CIFAR10, and SVHN in the sequel, while results on MNIST can be found in Appendix E. Note that, we use Pytorch platform to conduct all the tests reported in this paper, and do not use GPUs. We underline here that the DNNs are non-convex globally, however, one could assume at least local convexity, hence local gradient dominance of order $p = 2$, thus, we will select $q > 2$ in our experiments (see (Remark 2, Appendix E) for more explanations on the choice of q).

4.2.1 EXPERIMENT ON CIFAR10

In this experiment, we use the proposed algorithms to train a VGG16 CNN model with cross entropy loss, e.g., Simonyan & Zisserman (2015) on the CIFAR10 dataset. We divided the dataset into a training set spread in 50 batches of 1000 images each, and a test set of 10 batches with 1000 images each. We ran 20 epochs of training over all the training batches. Since Nesterov accelerated GD is one of the most efficient methods in DNN applications, to conduct fair comparisons we implemented our Nesterov-like discretization of q -RGF ($c = 1, q = 3, \eta = 0.04, \mu = 0.9$), and the Nesterov-like discretization of q -SGF ($c = 10^{-3}, q = 3, \eta = 0.04, \mu = 0.9$). We compare against the

⁴To avoid overloading the figures we had to choose one flow at a time, either q -RGF or q -SGF. More results can be found in Appendix E

⁵We did multiple iterations to find the best step size for each algorithm (best values were between 10^{-4} and 10^{-2} depending on the algorithm). Details of the step size for each test are given in Appendix E.

mainstream algorithms⁶, such as, Nesterov’s accelerated gradient descent (GD), Adam, Adaptive gradient (AdaGrad), per-dimension learning rate method for gradient descent (AdaDelta), and Root Mean Square Propagation (RMSprop)⁷. Note that, all algorithms have been implemented in their stochastic version⁸, i.e., using mini-batches implementation, with constant step size. For instance, in Figures 2, 3⁹, we see the training loss for the different optimization algorithms. We notice that the proposed Algorithms, RGF, and SGF, quickly separate from the GD, and RMS algorithms in terms of convergence speed, but also ends up with an overall better performance on the test set 84% vs. 83% for GD, and RMS. We also note that other algorithms, such as, AdaGrad and AdaDelta behave similarly to RGF in terms of convergence speed, but lack behind in terms of final performance 75% and 68%, respectively. Finally, in Figure 3, we notice that Adam is slower in terms of computation time w.r.t SGF, and RGF, with an average lag of 8 min, and 80 min, respectively. However, to be fair one, has to underline that Adam is an adaptive method based on the extra computations and memory overhead of lower order moments estimate and bias correction, Kingma & Ba (2015). Furthermore, to better compare the performance of these algorithms, we report in Figure 2 the loss on the test dataset over the learning iterations. We confirm that RGF and SGF performs better than GD, RMSprop, and Adam, while avoiding the overfitting observed with AdaGrad and AdaDelta.




Figure 2: Losses for several optimization algorithms- VGG-16- CIFAR10: Train loss (left), test loss (right)- We add an S before the name of an algorithm to denote its stochastic implementation.




Figure 3: Training loss vs. computation time for several optimization algorithms- VGG-16- CIFAR10

⁶We run several tests by trying to optimally tune the parameters of each algorithms on a validation set (tenth of training set), and we are reporting the best final accuracy performance we could obtain for each one. We have implemented all algorithms with the same total number of iterations, so that we can compare the convergence speed of all algorithms against a common iteration range. Details of the hyper-parameter values are given in Appendix.

⁷Original reference for each method can be found in: <https://pytorch.org/docs/stable/optim.html>

⁸We want to underline here that our first tests were done in the deterministic setting, however, to compare the proposed optimization methods against the best optimization algorithms available for DNNs training, we decided to also conduct comparison tests in the stochastic setting. Since the results remain qualitatively unchanged, we only report here the results due to the stochastic setting.

⁹To avoid overloading the figures we reported only the computation time plots of the three most competitive methods: RGF, SGF and Adam.

4.2.2 EXPERIMENTS ON SVHN DATASET

We test the proposed algorithms to train the same VGG16 CNN model with cross entropy loss on the SVHN dataset. We divided the dataset into a training set of 70 batches with 1000 images each, and a test set of 10 of 1000 images each, and ran 20 epochs of training over all the training batches. We tested the Nesterov-like discretization of q -RGF ($c = 1, q = 3, \eta = 0.04, \mu = 0.09$), and the Nesterov-like discretization of q -SGF ($c = 10^{-3}, q = 11, \eta = 0.04, \mu = 0.09$) against Nesterov’s accelerated gradient descent (GD), and Adam¹⁰. Note from Figures 4, 5 it is clear that RGF and SGF give a good performance in terms of convergence speed, and final test performance 93%. We can also observe in Figure 5 that SGF, RGF are faster than GD, and all three methods are faster (in average 41 min for GD, 58 min for SGF, 75 min for RGF) than Adam as expected since it is an adaptive scheme with more computation steps (see our discussion of Adam in Section 4.2.1). More




Figure 4: Losses for several optimization algorithms - CNN- SVHN: Train loss (left), test loss (right)




Figure 5: Training loss vs. computation time for several optimization algorithms- VGG-16- SVHN

numerical results on MNIST, and on SVHN using Euler and Runge-Kutta discretization, showing similar qualitative results, can be found in Appendix E.

5 CONCLUSION

We studied connections between optimization algorithms and continuous-time representations (dynamical systems) via discretization. We then reviewed two families of non-Lipschitz or discontinuous first-order optimization flows for continuous-time optimization, namely the q -RGF and q -SGF, whose distinguishing characteristic is their finite-time convergence. We then proposed three discretization methods for these flows, namely a forward-Euler discretization, followed by an explicit Runge-Kutta discretization, and finally a novel Nesterov-like discretization. Based on tools from hybrid systems control theory, we proved a convergence bound for these algorithms. Finally, we conducted numerical experiments on known deep neural net benchmarks, which showed that the proposed discrete algorithms can outperform some state of the art algorithms, when tested on large DNN models.

¹⁰We also tested Adaptive gradient (AdaGrad), per-dimension learning rate method for gradient descent (AdaDelta), and Root Mean Square Propagation (RMSprop). However, since their performance was not competitive we decided not to report the plots to avoid overloading the figures.

REFERENCES

- Hedy Attouch and Jerome Bolte. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Mathematical Programming B*, 116(1):5–16, 2009.
- Andrea Bacciotti and Francesca Ceragioli. Stability and stabilization of discontinuous systems and nonsmooth Lyapunov functions. *ESAIM: Control, Optimisation and Calculus of Variations*, 4: 361–376, 1999.
- Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44(1):328–348, 2005.
- Jérôme Bolte, Aris Daniilidis, and Adrian Lewis. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *Society for Industrial and Applied Mathematics*, 17:1205–1223, January 2007.
- C.A Botsaris. A class of methods for unconstrained minimization based on stable numerical integration techniques. *Journal of Mathematical Analysis and Applications*, 63(3):729–749, 1978a.
- C.A. Botsaris. Differential gradient methods. *Journal of Mathematical Analysis and Applications*, 63(1):177–198, 1978b.
- R.W. Brockett. Dynamical systems that sort lists, diagonalize matrices and solve linear programming problems. In *IEEE Conference on Decision and Control*, pp. 799–803, 1988.
- A.A. Brown. Some effective methods for unconstrained optimization based on the solution of systems of ordinary differential equations. *Journal of Optimization Theory and Applications*, 62(2):211–224, August 1989.
- Frank H. Clarke. Generalized gradients of Lipschitz functionals. *Advances in Mathematics*, 40(1): 52–67, 1981.
- Jorge Cortés. Finite-time convergent gradient flows with applications to network consensus. *Automatica*, 42(11):1993–2000, November 2006.
- Jorge Cortés. Discontinuous dynamical systems. *IEEE Control Systems Magazine*, 28(3):36–73, June 2008.
- Jorge Cortés and Francesco Bullo. Coordination and geometric optimization via distributed dynamical systems. *SIAM Journal on Control and Optimization*, 44(5):1543–1574, October 2005.
- M. Fazlyab, A. Koppel, V. M. Preciado, and A. Ribeiro. A variational approach to dual methods for constrained convex optimization. In *2017 American Control Conference (ACC)*, pp. 5269–5275, May 2017a.
- M. Fazlyab, A. Koppel, V. M. Preciado, and A. Ribeiro. A variational approach to dual methods for constrained convex optimization. In *2017 American Control Conference (ACC)*, pp. 5269–5275, May 2017b.
- M. Fazlyab, M. Morari, and V. M. Preciado. Design of first-order optimization algorithms via sum-of-squares programming. In *IEEE Conference on Decision and Control (CDC)*, pp. 4445–4452, December 2018.
- Mahyar Fazlyab, Alejandro Ribeiro, Manfred Morari, and Victor M. Preciado. Analysis of optimization algorithms via integral quadratic constraints: Nonstrongly convex problems. *SIAM J. Optim.*, 28(3):2654–2689, 2018.
- Aleksei Fedorovich Filippov and F. M. Arscott. *Differential equations with discontinuous righthand sides*. Kluwer Academic Publishers Group, Dordrecht, Netherlands, 1988.
- G. França, D. Robinson, and R. Vidal. ADMM and accelerated ADMM as continuous dynamical systems. July 2018.
- G. França, D.P. Robinson, and R. Vidal. A dynamical systems perspective on nonsmooth constrained optimization. *arXiv preprint 1808.04048*, 2019a.

- G. França, J. Sulam, D. Robinson, and R. Vidal. Conformal symplectic and relativistic optimization. *arXiv preprint 1903.04100*, 2019b.
- Uwe Helmke and John Barratt Moore. *Optimization and Dynamical Systems*. Springer-Verlag, 1994.
- Qing Hui, Wassim Haddad, and Sanjay Bhat. Semistability, finite-time stability, differential inclusions, and discontinuous dynamical systems having a continuum of equilibria. *IEEE Transactions on Automatic Control*, 54:2465–2470, November 2009.
- Hamed Karimi, Julie Nutini, , and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.
- Diederik P. Kingma and Jimmy Lei Ba. Adam: a method for stochastic optimization. *International Conference on Learning Representations*, pp. 1–15, May 2015.
- L. Lessard, B. Recht, , and A. Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM J. Optim*, 26(1):57–95, 2016.
- S. Łojasiewicz. A topological property of real analytic subsets (in French). *Les équations aux dérivées partielles*, pp. 87–89, 1963.
- S. Łojasiewicz. *Ensembles semi-analytiques*. Centre de Physique Theorique de l’Ecole Polytechnique, 1965. URL <https://perso.univ-rennes1.fr/michel.coste/Lojasiewicz.pdf>.
- Stanisław Łojasiewicz and Maria-Angeles Zurro. On the gradient inequality. *Bulletin of the Polish Academy of Sciences, Mathematics*, 47, January 1999.
- Michael Muehlebach and Michael Jordan. A dynamical systems perspective on Nesterov acceleration. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4656–4662, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- A. Orvieto and A. Lucchi. Shadowing properties of optimization algorithms. In *Neural Information Processing Systems*, December 2019.
- Bradley Paden and Shankar Sastry. A calculus for computing filippov’s differential inclusion with application to the variable structure control of robot manipulators. *IEEE Transactions on Circuits and Systems*, 34:73–82, February 1987.
- Boris Polyak. Gradient methods for the minimisation of functionals (in Russian). *USSR Computational Mathematics and Mathematical Physics*, 3:864–878, December 1963.
- O. Romero and M. Benosman. Finite-time convergence in continuous-time optimization. In *International Conference on Machine Learning*, Vienna, Austria, July 2020.
- R. G. Sanfelice and A. R. Teel. Dynamical properties of hybrid systems simulators. *Automatica*, 46: 239–248, 2010.
- Johannes Schropp. Using dynamical systems methods to solve minimization problems. *Applied Numerical Mathematics*, 18(1):321–335, 1995.
- Johannes Schropp and I Singer. A dynamical systems approach to constrained minimization. *Numerical Functional Analysis and Optimization*, 21:537–551, May 2000.
- D. Scieur, V. Roulet, F. Bach, , and A. d’Aspremont. Integration methods and optimization algorithms. In *Neural Information Processing Systems*, December 2017.
- Bin Shi, Simon Du, Michael Jordan, and Weijie Su. Understanding the acceleration phenomenon via high-resolution differential equations. *arXiv preprint 1810.08907*, October 2018.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, pp. 1–14, May 2015.

- J.A. Snyman. A new and dynamic method for unconstrained minimization. *Applied Mathematical Modelling*, 6(6):448–462, December 1982.
- J.A. Snyman. An improved version of the original leap-frog dynamic method for unconstrained minimization: LFOP1(b). *Applied Mathematical Modelling*, 7(3):216–218, June 1983.
- A. M. Stuart and A. R. Humphries. *Dynamical systems and numerical analysis*. Cambridge University Press, 1996.
- Andrew M. Stuart and A. R. Humphries. *Dynamical systems and numerical analysis*. Cambridge University Press, first edition, November 1998.
- W. Su, S. Boyd, and E. J. Candes. A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pp. 2510–2518. Curran Associates, Inc., 2014.
- Adrien Taylor, Bryan Van Scoy, and Laurent Lessard. Lyapunov functions for first-order methods: Tight automated convergence guarantees. In *International Conference on Machine Learning*, Stockholm, Sweden, July 2018.
- J. Wang and N. Elia. A control perspective for centralized and distributed convex optimization. In *IEEE Conference on Decision and Control and European Control Conference*, pp. 3800–3805, December 2011.
- Andre Wibisono, Ashia C. Wilson, and Michael I. Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.
- A. Wilson. *Lyapunov Arguments in Optimization*. PhD thesis, UC Berkeley, 2018.
- Ashia Wilson, Lester Mackey, and Andre Wibisono. Accelerating rescaled gradient descent: Fast optimization of smooth functions. In *Advances in Neural Information Processing Systems*. December 2019.
- Abbas K. Zghier. *The use of differential equations in optimization*. PhD thesis, Loughborough University, 1981.
- J. Zhang, A. Mokhtari, S. Sra, , and A. Jadbabaie. Direct runge-kutta discretization achieves acceleration. In *Neural Information Processing Systems*, December 2018.

A DISCONTINUOUS SYSTEMS AND DIFFERENTIAL INCLUSIONS

¹¹ Recall that for an initial value problem (IVP)

$$\dot{x}(t) = F(x(t)) \quad (24a)$$

$$x(0) = x_0 \quad (24b)$$

with $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the typical way to check for existence of solutions is by establishing continuity of F . Likewise, to establish uniqueness of the solution, we typically seek Lipschitz continuity. When F is discontinuous, we may understand (24a) as the Filippov differential inclusion

$$\dot{x}(t) \in \mathcal{K}[F](x(t)), \quad (25)$$

where $\mathcal{K}[F] : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ denotes the Filippov set-valued map given by

$$\mathcal{K}[F](x) \triangleq \bigcap_{\delta > 0} \bigcap_{\mu(S)=0} \overline{\text{co}} F(B_\delta(x) \setminus S), \quad (26)$$

where μ denotes the usual Lebesgue measure and $\overline{\text{co}}$ the convex closure, i.e. closure of the convex hull $\overline{\text{co}}$. For more details, see Paden & Sastry (1987). We can generalize (25) to the differential inclusion Bacciotti & Ceragioli (1999)

$$\dot{x}(t) \in \mathcal{F}(x(t)), \quad (27)$$

where $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is some set-valued map.

Definition 1 (Carathéodory/Filippov solutions). We say that $x : [0, \tau) \rightarrow \mathbb{R}^n$ with $0 < \tau \leq \infty$ is a *Carathéodory solution* to (27) if $x(\cdot)$ is absolutely continuous and (27) is satisfied a.e. in every compact subset of $[0, \tau)$. Furthermore, we say that $x(\cdot)$ is a *maximal Carathéodory solution* if no other Carathéodory solution $x'(\cdot)$ exists with $x = x'|_{[0, \tau)}$. If $\mathcal{F} = \mathcal{K}[F]$, then Carathéodory solutions are referred to as *Filippov solutions*.

For a comprehensive overview of discontinuous systems, including sufficient conditions for existence (Proposition 3) and uniqueness (Propositions 4 and 5) of Filippov solutions, see the work of Cortés (2008). In particular, it can be established that Filippov solutions to (24) exist, provided that the following assumption (Assumption 1) holds.

Assumption 1 (Existence of Filippov solutions). $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined almost everywhere (a.e.) and is Lebesgue-measurable in a non-empty open neighborhood $U \subset \mathbb{R}^n$ of $x_0 \in \mathbb{R}^n$. Further, F is locally essentially bounded in U , i.e., for every point $x \in U$, F is bounded a.e. in some bounded neighborhood of x .

More generally, Carathéodory solutions to (27) exist (now with arbitrary $x_0 \in \mathbb{R}^n$), provided that the following assumption (Assumption 2) holds.

Assumption 2 (Existence of Carathéodory solutions). $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ has nonempty, compact, and convex values, and is *upper semi-continuous*.

Filippov & Arscott (1988) proved that, for the Filippov set-valued map $\mathcal{F} = \mathcal{K}[F]$, Assumptions 1 and 2 are equivalent (with arbitrary $x_0 \in \mathbb{R}^n$ in Assumption 1).

Uniqueness of the solution requires further assumptions. Nevertheless, we can characterize the Filippov set-valued map in a similar manner to Clarke's generalized gradient, as seen in the following proposition.

Proposition 1 (Theorem 1 of Paden & Sastry (1987)). *Under Assumption 1, we have*

$$\mathcal{K}[F](x) = \left\{ \lim_{k \rightarrow \infty} F(x_k) : x_k \in \mathbb{R}^n \setminus (\mathcal{N}_F \cup S) \text{ s.t. } x_k \rightarrow x \right\} \quad (28)$$

for some (Lebesgue) zero-measure set $\mathcal{N}_F \subset \mathbb{R}^n$ and any other zero-measure set $S \subset \mathbb{R}^n$. In particular, if F is continuous at a fixed x , then $\mathcal{K}[F](x) = \{F(x)\}$.

¹¹The notions introduced here are solely needed for rigorously dealing with the singular discontinuity at the equilibrium point of the q -RGF and q -SGF flows. However, the reader can skip these definitions and still be able to intuitively follow the proofs of Theorems 1, and 2.

For instance, for the GF (1), we have $\mathcal{K}[-\nabla f](x) = \{-\nabla f(x)\}$ for every $x \in \mathbb{R}^n$, provided that f is continuously differentiable. Furthermore, if f is only locally Lipschitz continuous and regular (see Definition 3 of Appendix B), then $\mathcal{K}[-\nabla f](x) = -\partial f(x)$, where

$$\partial f(x) \triangleq \left\{ \lim_{k \rightarrow \infty} \nabla f(x_k) : x_k \in \mathbb{R}^n \setminus \mathcal{N}_f \text{ s.t. } x_k \rightarrow x \right\} \quad (29)$$

denotes Clarke's generalized gradient Clarke (1981) of f , with \mathcal{N}_f denoting the zero-measure set over which f is not differentiable (Rademacher's theorem). It can be established that ∂f coincides with the subgradient of f , provided that f is convex. Therefore, the GF (1) interpreted as Filippov differential inclusion may also be seen as a continuous-time variant of subgradient descent methods.

B FINITE-TIME STABILITY OF DIFFERENTIAL INCLUSIONS

We are now ready to focus on extending some notions from traditional Lipschitz continuous systems to differential inclusions.

Definition 2. We say that $x^* \in \mathbb{R}^n$ is an *equilibrium* of (27) if $x(t) \equiv x^*$ on some small enough non-degenerate interval is a Carathéodory solution to (27). In other words, if and only if $0 \in \mathcal{F}(x^*)$. We say that (27) is (*Lyapunov*) *stable* at $x^* \in \mathbb{R}^n$ if, for every $\varepsilon > 0$, there exists some $\delta > 0$ such that, for every maximal Carathéodory solution $x(\cdot)$ of (27), we have $\|x_0 - x^*\| < \delta \implies \|x(t) - x^*\| < \varepsilon$ for every $t \geq 0$ in the interval where $x(\cdot)$ is defined. Note that, under Assumption 2, if (27) is stable at x^* , then x^* is an equilibrium of (27) Bacciotti & Ceragioli (1999). Furthermore, we say that (27) is (*locally and strongly*) *asymptotically stable* at $x^* \in \mathbb{R}^n$ if it is stable at x^* and there exists some $\delta > 0$ such that, for every maximal Carathéodory solution $x : [0, \tau) \rightarrow \mathbb{R}^n$ of (27), if $\|x_0 - x^*\| < \delta$ then $x(t) \rightarrow x^*$ as $t \rightarrow \tau$. Finally, (27) is (*locally and strongly*) *finite-time stable* at x^* if it is asymptotically stable and there exists some $\delta > 0$ and $T : B_\delta(x^*) \rightarrow [0, \infty)$ such that, for every maximal Carathéodory solution $x(\cdot)$ of (27) with $x_0 \in B_\delta(x^*)$, we have $\lim_{t \rightarrow T(x_0)} x(t) = x^*$.

We will now construct a Lyapunov-based criterion adapted from the literature of finite-time stability of Lipschitz continuous systems.

Lemma 1. *Let $\mathcal{E}(\cdot)$ be an absolutely continuous function satisfying the differential inequality*

$$\dot{\mathcal{E}}(t) \leq -c\mathcal{E}(t)^\alpha \quad (30)$$

a.e. in $t \geq 0$, with $c, \mathcal{E}(0) > 0$ and $\alpha < 1$. Then, there exists some $t^ > 0$ such that $\mathcal{E}(t) > 0$ for $t \in [0, t^*)$ and $\mathcal{E}(t^*) = 0$. Furthermore, $t^* > 0$ can be bounded by*

$$t^* \leq \frac{\mathcal{E}(0)^{1-\alpha}}{c(1-\alpha)}, \quad (31)$$

with this bound tight whenever (30) holds with equality. In that case, but now with $\alpha \geq 1$, then $\mathcal{E}(t) > 0$ for every $t \geq 0$, with $\lim_{t \rightarrow \infty} \mathcal{E}(t) = 0$. This will be represented by $t^ = \infty$, with $\mathcal{E}(\infty) \triangleq \lim_{t \rightarrow \infty} \mathcal{E}(t)$.*

Proof. Suppose that $\mathcal{E}(t) > 0$ for every $t \in [0, T]$ with $T > 0$. Let t^* be the supremum of all such T 's, thus satisfying $\mathcal{E}(t) > 0$ for every $t \in [0, t^*)$. We will now investigate $\mathcal{E}(t^*)$. First, by continuity of \mathcal{E} , it follows that $\mathcal{E}(t^*) \geq 0$. Now, by rewriting

$$\dot{\mathcal{E}}(t) \leq -c\mathcal{E}(t)^\alpha \iff \frac{d}{dt} \left[\frac{\mathcal{E}(t)^{1-\alpha}}{1-\alpha} \right] \leq -c, \quad (32)$$

a.e. in $t \in [0, t^*)$, we can thus integrate to obtain

$$\frac{\mathcal{E}(t)^{1-\alpha}}{1-\alpha} - \frac{\mathcal{E}(0)^{1-\alpha}}{1-\alpha} \leq -ct, \quad (33)$$

everywhere in $t \in [0, t^*)$, which in turn leads to

$$\mathcal{E}(t) \leq [\mathcal{E}(0)^{1-\alpha} - c(1-\alpha)t]^{1/(1-\alpha)} \quad (34)$$

and

$$t \leq \frac{\mathcal{E}(0)^{1-\alpha} - \mathcal{E}(t)^{1-\alpha}}{c(1-\alpha)} \leq \frac{\mathcal{E}(0)^{1-\alpha}}{c(1-\alpha)}, \quad (35)$$

where the last inequality follows from $\mathcal{E}(t) > 0$ for every $t \in [0, t^*)$. Taking the supremum in (35) then leads to the upper bound (31). Finally, we conclude that $\mathcal{E}(t^*) = 0$, since $\mathcal{E}(t^*) > 0$ is impossible given that it would mean, due to continuity of \mathcal{E} , that there exists some $T > t^*$ such that $\mathcal{E}(t) > 0$ for every $t \in [0, T]$, thus contradicting the construction of t^* .

Finally, notice that if \mathcal{E} is such that (30) holds with equality, then (34) and the first inequality in (35) hold with equality as well. The tightness of the bound (31) thus follows immediately. Furthermore, notice that if $\alpha \geq 1$, and \mathcal{E} is a tight solution to the differential inequality (30), *i.e.* $\mathcal{E}(t) = [\mathcal{E}(0)^{1-\alpha} - c(1-\alpha)t]^{1/(1-\alpha)}$, then clearly $\mathcal{E}(t) > 0$ for every $t \geq 0$ and $\mathcal{E}(t) \rightarrow 0$ as $t \rightarrow \infty$. ■

Cortés & Bullo (2005) proposed (Proposition 2.8) a Lyapunov-based criterion to establish finite-time stability of discontinuous systems, which fundamentally coincides with our Lemma 1 for the particular choice of exponent $\alpha = 0$. Their proposition was, however, directly based on Theorem 2 of Paden & Sastry (1987). Later, Cortés (2006) proposed a second-order Lyapunov criterion, which, on the other hand, fundamentally translates to $\mathcal{E}(t) \triangleq V(x(t))$ being strongly convex. Finally, Hui et al. (2009) generalized Proposition 2.8 of Cortés & Bullo (2005) in their Corollary 3.1, to establish semistability. Indeed, these two results coincide for isolated equilibria.

We now present a novel result that generalizes the aforementioned first-order Lyapunov-based results, by exploiting our Lemma 1. More precisely, given a Lyapunov candidate function $V(\cdot)$, the objective is to set $\mathcal{E}(t) \triangleq V(x(t))$, and we aim to check that the conditions of Lemma 1 hold. To do this, and assuming V to be locally Lipschitz continuous, we first borrow and adapt from Bacciotti & Ceragioli (1999) the definition of *set-valued time derivative* of $V : \mathcal{D} \rightarrow \mathbb{R}$ w.r.t. the differential inclusion (27), given by

$$\dot{V}(x) \triangleq \{a \in \mathbb{R} : \exists v \in \mathcal{F}(x) \text{ s.t. } a = p \cdot v, \forall p \in \partial V(x)\}, \quad (36)$$

for each $x \in \mathcal{D}$. Notice that, under Assumption 2 for Filippov differential inclusions $\mathcal{F} = \mathcal{K}[F]$, the set-valued time derivative of V thus coincides with with the set-valued Lie derivative $\mathcal{L}_{\mathcal{F}}V(\cdot)$. Indeed, more generally \dot{V} could be seen as a set-valued Lie derivative $\mathcal{L}_{\mathcal{F}}V$ w.r.t. the set-valued map \mathcal{F} .

Definition 3. $V(\cdot)$ is said to be *regular* if every directional derivative, given by

$$V'(x; v) \triangleq \lim_{h \rightarrow 0} \frac{V(x + hv) - V(x)}{h}, \quad (37)$$

exists and is equal to

$$V^\circ(x; v) \triangleq \limsup_{x' \rightarrow x, h \rightarrow 0^+} \frac{V(x' + hv) - V(x')}{h}, \quad (38)$$

known as *Clarke's upper generalized derivative* Clarke (1981).

In practice, regularity is a fairly mild and easy to guarantee condition. For instance, it would suffice that V is convex or continuously differentiable to ensure that it is Lipschitz and regular.

Assumption 3. $V : \mathcal{D} \rightarrow \mathbb{R}$ is locally Lipschitz continuous and regular, with $\mathcal{D} \subseteq \mathbb{R}^n$ open.

Under Assumption 3, Clarke's generalized gradient

$$\partial V(x) \triangleq \{p \in \mathbb{R}^n : V^\circ(x; v) \geq p \cdot v, \forall v \in \mathbb{R}^n\} \quad (39)$$

is non-empty for every $x \in \mathcal{D}$, and is also given by

$$\partial V(x) = \left\{ \lim_{k \rightarrow \infty} \nabla V(x_k) : x_k \in \mathbb{R}^n \setminus \mathcal{N}_V \text{ s.t. } x_k \rightarrow x \right\}, \quad (40)$$

where \mathcal{N}_V denotes the set of points in $\mathcal{D} \subseteq \mathbb{R}^n$ where V is not differentiable (Rademacher's theorem) Clarke (1981).

Through the following lemma (Lemma 2), we can formally establish the correspondence between the set-valued time-derivative of V and the derivative of the energy function $\mathcal{E}(t) \triangleq V(x(t))$ associated with an arbitrary Carathéodory solution $x(\cdot)$ to the differential inclusion (27).

Lemma 2 (Lemma 1 of Bacciotti & Ceragioli (1999)). *Under Assumption 3, given any Carathéodory solution $x : [0, \tau) \rightarrow \mathbb{R}^n$ to (27), then $\mathcal{E}(t) \triangleq V(x(t))$ is absolutely continuous and $\dot{\mathcal{E}}(t) = \frac{d}{dt} V(x(t)) \in \dot{V}(x(t))$ a.e. in $t \in [0, \tau)$.*

We are now ready to state and prove our Lyapunov-based sufficient condition for finite-time stability of differential inclusions.

Theorem 3. *Suppose that Assumptions 2 and 3 hold for some set-valued map $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ and some function $V : \mathcal{D} \rightarrow \mathbb{R}$, where $\mathcal{D} \subseteq \mathbb{R}^n$ is an open and positively invariant neighborhood of a point $x^* \in \mathbb{R}^n$. Suppose that V is positive definite w.r.t. x^* and that there exist constants $c > 0$ and $\alpha < 1$ such that*

$$\sup \dot{V}(x) \leq -cV(x)^\alpha \quad (41)$$

a.e. in $x \in \mathcal{D}$. Then, (27) is finite-time stable at x^* , with settling time upper bounded by

$$t^* \leq \frac{V(x_0)^{1-\alpha}}{c(1-\alpha)}, \quad (42)$$

where $x(0) = x_0$. In particular, any Carathéodory solution $x(\cdot)$ with $x(0) = x_0 \in \mathcal{D}$ will converge in finite time to x^* under the upper bound (42). Furthermore, if $\mathcal{D} = \mathbb{R}^n$, then (27) is globally finite-time stable. Finally, if $\dot{V}(x)$ is a singleton a.e. in $x \in \mathcal{D}$ and (41) holds with equality, then the bound (42) is tight.

Proof. Note that, by Proposition 1 of Bacciotti & Ceragioli (1999), we know that (27) is Lyapunov stable at x^* . All that remains to be shown is local convergence towards x^* (which must be an equilibrium) in finite time. Indeed, given any maximal solution $x : [0, t^*) \rightarrow \mathbb{R}^n$ to (27) with $x(0) = x_0 \neq x^*$, we know by Lemma 2, that $\mathcal{E}(t) = V(x(t))$ is absolutely continuous with $\dot{\mathcal{E}}(t) \in \dot{V}(x(t))$ a.e. in $t \in [0, t^*)$. Therefore, we have

$$\dot{\mathcal{E}}(t) \leq \sup \dot{V}(x(t)) \leq -cV(x(t))^\alpha = -c\mathcal{E}(t)^\alpha \quad (43)$$

a.e. in $t \in [0, t^*)$. Since $\mathcal{E}(0) = V(x_0) > 0$, given that $x_0 \neq x^*$, the result then follows by invoking Lemma 1 and noting that $\mathcal{E}(t^*) = 0 \iff V(t^*, x(t^*)) = 0 \iff x(t^*) = x^*$. ■

Finite-time stability still follows without Assumption 2, provided that x^* is an equilibrium of (27). In practical terms, this means that trajectories starting arbitrarily close to x^* may not actually exist, but nevertheless there exists a neighborhood \mathcal{D} of x^* over which, any trajectory $x(\cdot)$ that indeed exists and starts at $x(0) = x_0 \in \mathcal{D}$ must converge in finite time to x^* , with settling time upper bounded by $T(x_0)$ (the bound still tight in the case that (41) holds with equality).

C PROOF OF THEOREM 1

Let us focus on the q -RGF (2) (the case of q -SGF (3) follows exactly the same steps) with the candidate Lyapunov function $V \triangleq f - f^*$. Clearly, V is Lipschitz continuous and regular (given that it is continuously differentiable). Furthermore, V is positive definite w.r.t. x^* .

Notice that, due to the dominated gradient assumption, x^* must be an isolated stationary point of f . To see this, notice that, if x^* were not an isolated stationary point, then there would have to exist some \tilde{x}^* sufficiently near x^* such that \tilde{x}^* is both a stationary point of f , and satisfies $f(\tilde{x}^*) > f^*$, since x^* is a strict local minimizer of f . But then, we would have

$$0 = \frac{p-1}{p} \|\nabla f(\tilde{x}^*)\|^{\frac{p}{p-1}} \geq \mu^{\frac{1}{p-1}} (f(\tilde{x}^*) - f^*) > 0, \quad (44)$$

and subsequently $0 > 0$, which is absurd.

Therefore, $F(x) \triangleq -c\nabla f(x) / \|\nabla f(x)\|^{\frac{q-2}{q-1}}$ is continuous for every $x \in \mathcal{D} \setminus \{x^*\}$, for some small enough open neighborhood \mathcal{D} of x^* . Let us assume that \mathcal{D} is positively invariant w.r.t. (2), which can be achieved, for instance, by replacing \mathcal{D} with its intersection with some small enough strict sublevel set of f . Notice that $\|F(x)\| = c\|\nabla f(x)\|^{\frac{1}{q-1}}$ with $q \in (p, \infty) \subset (1, \infty]$, i.e., $\frac{1}{q-1} \in [0, \infty)$. If $q = \infty$, which results in the normalized gradient flow $\dot{x} = -\frac{\nabla f(x)}{\|\nabla f(x)\|}$ proposed by Cortés (2006),

then $\|F(x)\| = c > 0$. We can thus show that $F(x)$ is discontinuous at $x = 0$ for $q = \infty$. On the other hand, if $q \in (p, \infty) \subset (1, \infty)$, then we have $\|F(x)\| \rightarrow 0$ as $x \rightarrow x^*$, and thus $F(x)$ is continuous (but not Lipschitz) at $x = x^*$. Regardless, we may freely focus exclusively on $\mathcal{D} \setminus \{x^*\}$ since $\{x^*\}$ is obviously a zero-measure set.

Let $\mathcal{F} \triangleq \mathcal{K}[F]$. We thus have, for each $x \in \mathcal{D} \setminus \{x^*\}$,

$$\sup \dot{V}(x) = \sup \{a \in \mathbb{R} : \exists v \in \mathcal{F}(x) \text{ s.t. } a = p \cdot v, \forall p \in \partial V(x)\} \quad (45a)$$

$$= \sup \{\nabla V(x) \cdot v : v \in \mathcal{F}(x)\} \quad (45b)$$

$$= \nabla V(x) \cdot F(x) \quad (45c)$$

$$= -c \|\nabla f(x)\|^{2-\frac{q-2}{q-1}} \quad (45d)$$

$$= -c \|\nabla f(x)\|^{\frac{1}{\theta'}} \quad (45e)$$

$$\leq -c [C(f(x) - f^*)^\theta]^{\frac{1}{\theta'}} \quad (45f)$$

$$= -c C^{\frac{1}{\theta'}} V(x)^{\frac{\theta}{\theta'}}. \quad (45g)$$

Since $\frac{\theta}{\theta'} < 1$, given that $s > 1 \mapsto \frac{s-1}{s}$ is strictly increasing, then the conditions of Theorem 3 are satisfied. In particular, we have finite-time stability at x^* with a settling time t^* upper bounded by

$$t^* \leq \frac{(f(x_0) - f^*)^{1-\frac{\theta}{\theta'}}}{c C^{\frac{1}{\theta'}} (1 - \frac{\theta}{\theta'})} \leq \frac{(\|\nabla f(x_0)\|/C)^{\frac{1}{\theta}(1-\frac{\theta}{\theta'})}}{c C^{\frac{1}{\theta'}} (1 - \frac{\theta}{\theta'})} = \frac{\|\nabla f(x_0)\|^{\frac{1}{\theta}-\frac{1}{\theta'}}}{c C^{\frac{1}{\theta}} (1 - \frac{\theta}{\theta'})} \quad (46)$$

for each $x_0 \in \mathcal{D}$, which completes the proof.

D PROOF OF THEOREM 2

To prove Theorem 2, we borrow some tools and results from hybrid control systems theory. Hybrid control systems are characterized by continuous flows with discrete jumps between the continuous flows. They are often modeled by differential inclusions added to discrete mappings to model the jumps between the differential inclusions. We see the case of the optimization flows proposed here as a simple case of a hybrid systems with one differential inclusion, with a possible jump or discontinuity at the optimum. Based on this, we will use the tools and results of Sanfelice & Teel (2010), which study how a certain class of hybrid systems behave after discretization with a certain class of discretization algorithms. In other words, Sanfelice & Teel (2010) quantifies, under some conditions, how close are the solutions of the discretized hybrid dynamical system to the solutions of the original hybrid system.

In this section we will denote the differential inclusion of the continuous optimization flow by $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, and its discretization in time by $\mathcal{F}_d : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$. We first recall a definition, which we will adapt from the general case of jumps between multiple differential inclusions (Definition 3.2, Sanfelice & Teel (2010)) to our case of one differential inclusion or flow.

Definition 4. ((T, ϵ) -closeness). Given $T > 0$, $\epsilon > 0$, $\eta > 0$, two solutions $x_t : [0, T] \rightarrow \mathbb{R}^n$, and $x_k : \{0, 1, 2, \dots\} \rightarrow \mathbb{R}^n$ are (T, ϵ) -close if:

- (a) for all $t \leq T$ there exists $k \in \{1, 2, \dots\}$ such that $|t - k\eta| < \epsilon$, and $\|x_t(t) - x_k(k)\| < \epsilon$,
- (b) for all $k \in \{1, 2, \dots\}$ there exists $t \leq T$ such that $|t - k\eta| < \epsilon$, and $\|x_t(t) - x_k(k)\| < \epsilon$.

Next, we will recall Theorem 5.2 in Sanfelice & Teel (2010), while adapting it to our special case of a hybrid system with one differential inclusion¹².

Theorem 4. (*Closeness of continuous and discrete solutions on compact domains*) Consider the differential inclusion

$$\dot{X}(t) \in \mathcal{F}(X(t)), \quad (47)$$

¹²A set-valued mapping $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is *outer semicontinuous* if for each sequence $\{x_i\}_{i=1}^\infty$ converging to a point $x \in \mathbb{R}^n$ and each sequence $y_i \in \mathcal{F}(x_i)$ converging to a point y , it holds that $y \in \mathcal{F}(x)$. It is *locally bounded* if, for each $x \in \mathbb{R}^n$, there exists compact sets $K, K' \subset \mathbb{R}^n$ such that $x \in K$ and $\mathcal{F}(K) \triangleq \cup_{x \in K} \mathcal{F}(x) \subset K'$. In what follows, we use the following notations: Given a set A , $\text{con}A$ denotes the convex hull, and \mathbb{B} denotes the closed unit ball in a Euclidean space.

for a given set-valued mapping $\mathcal{F} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$ assumed to be outer semicontinuous, locally bounded, nonempty, and with convex values for every $x \in \mathcal{C}$, for some closed set $\mathcal{C} \subseteq \mathbb{R}^m$. Consider the discrete-time system represented by the flow $\mathcal{F}_d : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, such that, for each compact set $K \subset \mathbb{R}^n$, there exists $\rho \in \mathcal{K}_\infty$, and $\eta^* > 0$ such that for each $x \in K$ and each $\eta \in (0, \eta^*]$,

$$\mathcal{F}_d(x) \subset x + \eta \operatorname{con}\mathcal{F}(x + \rho(\eta)\mathbb{B}) + \eta\rho(\eta)\mathbb{B}. \quad (48)$$

Then, for every compact set $K \subset \mathbb{R}^n$, every $\epsilon > 0$, and every time horizon $T \in \mathbb{R}_{>0}$ there exists $\eta^* > 0$ such that: for any $\eta \in (0, \eta^*]$ and any discrete solution x_k with $x_k(0) \in \bar{K} + \delta\mathbb{B}$, $\delta > 0$, there exists a continuous solution x_t with $x_t(0) \in K$ such that x_k and x_t are (T, ϵ) -close.

To prove Theorem 2 we will use the results of Theorem 4, where we will have to check that condition (48) is satisfied for the three proposed discretizations.

We are now ready to prove Theorem 2. First, note that outer semicontinuity follows from the upper semicontinuity and the closedness of the Filippov differential inclusion map. Furthermore, local boundedness follows from continuity everywhere outside stationary points, which are isolated.

Now, let us examine their discretization by the three proposed algorithms:

FORWARD-EULER DISCRETIZATION

The mapping \mathcal{F}_d in this case is a singleton, given by

$$\mathcal{F}_d(x) \triangleq x + \eta F(x), \quad (49)$$

where $\eta > 0$, which clearly satisfies condition (48).

RUNGE-KUTTA DISCRETIZATION

Once again, the mapping \mathcal{F}_d is singleton, this time given by

$$\mathcal{F}_d(x) \triangleq x + \eta \sum_{i=1}^K \alpha_i \mathcal{F}(y^i) \quad (50a)$$

$$y^i = x + \eta \sum_{j=1}^{i-1} \beta_j \mathcal{F}(y^j), \quad (50b)$$

where $\eta, \alpha_1, \dots, \alpha_K, \beta_1, \dots, \beta_{K-1} > 0$ are such that $\sum_{i=1}^K \alpha_i = 1$.

By equation (50b) one can establish a function $\rho \in \mathcal{K}_\infty$ such that for each $x_k \in K \subset \mathbb{R}^n$, for each $\eta > 0$, $y_k^i \in x_k + \eta\rho(\eta)\mathbb{B}$. Next, by equation (50a) together with the condition $\sum_{i=1}^K \alpha_i = 1$, one can write that for any $x_k \in K$ and $\eta > 0$, $\mathcal{F}_d(x_k) \subset x_k + \eta \operatorname{con}\mathcal{F}(y_k^i) \subset x_k + \eta \operatorname{con}\mathcal{F}(x_k + \eta\rho(\eta)\mathbb{B}) + \eta\rho(\eta)\mathbb{B}$.

NESTEROV-LIKE DISCRETIZATION

In this case the discrete-time flow \mathcal{F}_d is defined as

$$\mathcal{F}_d(x_k) = x_k + \eta \mathcal{F}(x_k + \mu y_k) + \mu y_k \quad (51a)$$

$$y_k = x_k - x_{k-1}. \quad (51b)$$

In this case, to take into account the integral effect of the Nesterov-like discretization, let us extend the continuous-time flow as

$$\dot{\tilde{x}}_t = \left(\frac{\dot{x}_t}{\frac{d}{dt} \int_t^{t+\eta} x_t(s) ds} \right) = \left(\frac{\mathcal{F}(x_t)}{\eta} \right), \quad (52)$$

we then compare the solution of the extended continuous-time system (52) with the extended discrete-time system

$$\tilde{x}_{k+1} = \begin{pmatrix} x_{k+1} \\ x_{k+1} - x_k \end{pmatrix} = \begin{pmatrix} x_k + \eta \mathcal{F}(x_k(1 + \mu) + x_{k-1}) + \mu(x_k - x_{k-1}) \\ \eta \mathcal{F}(x_k(1 + \mu) + x_{k-1}) + \mu(x_k - x_{k-1}) \end{pmatrix}, \quad (53)$$

which could be rearranged as

$$\mathcal{F}_d(\tilde{x}_k) = \tilde{x}_{k+1} = \tilde{x}_k + M\tilde{x}_k + \eta\tilde{\mathcal{F}}_t(\tilde{x}_k), \quad (54)$$

where $M = \begin{bmatrix} 0 & \mu \\ 0 & -1 + \mu \end{bmatrix}$, and $\tilde{\mathcal{F}}_t(\tilde{x}_k) = \begin{bmatrix} \mathcal{F}(\tilde{x}_k) \\ \mathcal{F}(\tilde{x}_k) \end{bmatrix}$, which shows that for any $\tilde{x}_k \in K \subset \mathbb{R}^{2n}$,

and $\eta > 0$, we have (using similar recursive reasoning as above) that $\mathcal{F}_d(\tilde{x}_k) \subset \tilde{x}_k + \eta \text{con} \tilde{\mathcal{F}}_t(\tilde{x}_k + \eta\rho(\eta)\mathbb{B}) + \eta\rho(\eta)\mathbb{B}$. Then, using Theorem 4 we conclude about the (T, ϵ) -closeness between the continuous-time solutions of the flows $\mathcal{F} : q$ -RGF (2), q -SGF (3), and the discrete-time solutions of their respective discretization by any of the three discretization methods. Furthermore, for the Nesterov-like discretization, we can also conclude about the (T, ϵ) -closeness of the integral of the continuous-time solutions $\tilde{x}_t(2)$ and its discretization $\tilde{x}_k(2)$.

Finally, using the Lyapunov function $V = f - f^*$ as defined in the proof of Theorem 1, together with inequalities (45g), (34), and a local Lipschitz bound on f , one can derive the convergence bound given by (23), as follows:

$$\begin{aligned} \|f(x_k) - f(x^*) - (f(x_t) - f(x^*))\| &= \|f(x_k) - f(x_t)\| \leq \tilde{\epsilon} = L_f\epsilon, \quad L_f > 0, \quad \epsilon > 0, \\ \|f(x_k) - f(x^*)\| - \|(f(x_t) - f(x^*))\| &\leq \|f(x_k) - f(x_t)\| \leq \tilde{\epsilon}, \\ \|f(x_k) - f(x^*)\| &\leq \tilde{\epsilon} + \|f(x_t) - f(x^*)\|, \\ \|f(x_k) - f(x^*)\| &\leq \tilde{\epsilon} + [(f(x_0) - f(x^*))^{(1-\alpha)} - \tilde{c}(1-\alpha)\eta k]^{1/(1-\alpha)}, \quad \text{for } k \leq k^*, \end{aligned}$$

where $\alpha = \frac{\theta}{\theta'}$, $\theta = \frac{p-1}{p}$, $\theta' = \frac{q-1}{q}$, $\tilde{c} = c \left(\left(\frac{p}{p-1} \right)^{\frac{p-1}{p}} \mu^{\frac{1}{p}} \right)^{\frac{1}{\theta'}}$, $k^* = \frac{(f(x_0) - f(x^*))^{(1-\alpha)}}{\tilde{c}(1-\alpha)\eta}$.

E ADDITIONAL DETAILS AND NUMERICAL RESULTS

In this section, we will expand upon the numerical results experiments discussed in the paper. In particular, we report more details on the hyper-parameters values used¹³ for the numerical tests, and report some results for the MNIST experiments.

E.1 HYPER PARAMETERS VALUES USED IN THE TESTS OF SECTION 4.1

- GD fixed step size: $\eta = 10^{-3}$
- RGF Euler disc. w/fixed step size: $q = 2.2$, $\eta = 10^{-3}$
- RGF Euler disc. w/fixed step size: $q = 3$, $\eta = 10^{-2}$
- RGF Euler disc. w/fixed step size: $q = 6$, $\eta = 10^{-2}$
- RGF Euler disc. w/fixed step size: $q = 10$, $\eta = 10^{-2}$
- GD Nesterov acceleration fixed step size: $\eta = 10^{-4}$; $\mu = 0.9$
- SGF Nesterov-like disc. w/fixed step size: $q = 2.2$, $\eta = 10^{-4}$, $\mu = 0.9$
- SGF Nesterov-like disc. w/fixed step size: $q = 3$, $\eta = 10^{-3}$, $\mu = 0.9$
- SGF Nesterov-like disc. w/fixed step size: $q = 6$, $\eta = 10^{-3}$, $\mu = 0.9$
- SGF Nesterov-like disc. w/fixed step size: $q = 10$, $\eta = 10^{-2}$, $\mu = 0.09$
- RGF Runge Kutta disc. w/fixed step size: $q = 2.2$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$
- RGF Runge Kutta disc. w/fixed step size: $q = 3$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$
- RGF Runge Kutta disc. w/fixed step size: $q = 6$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$
- RGF Runge Kutta disc. w/fixed step size: $q = 10$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$

¹³In all the tests, for q -RGF and q -SGF $c = 1$ unless otherwise stated.

E.2 HYPER PARAMETERS VALUES USED IN THE TESTS OF SECTION 4.2

Note that the description of the coefficients for each of the prior art methods can be found in: <https://pytorch.org/docs/stable/optim.html>.

- GD: $\eta = 4.10^{-2}$, $\mu = 0.9$, Nesterov=True
- RGF: $\eta = 4.10^{-2}$, $\mu = 0.9$
- SGF: $\eta = 4.10^{-3}$, $\mu = 0.9$
- ADAM: $\eta = 8.10^{-4}$ (remaining coefficients=nominal values)
- RMS: $\eta = 10^{-3}$ (remaining coefficients=nominal values)
- ADAGRAD: $\eta = 10^{-3}$ (remaining coefficients=nominal values)
- ADADELTA: $\eta = 4.10^{-2}$, $\rho = 0.9$, $\epsilon = 10^{-6}$, weight decay = 0

E.3 MORE TESTS ON THE ROSENBROCK FUNCTION

Due to space limitations, we have decided to report in the main paper only one test for q -RGF with Euler discretization, one test for q -SGF with Nesterov-like discretization, and one test for q -RGF with Runge Kutta discretization. For the sake of completeness we report here the remaining tests for each algorithm. One can observe similar qualitative behavior in Figure 6 as the one noticed in the results of Section 4.5.

The step-size and other hyper-parameters for each test are given below:

- GD fixed step size: $\eta = 10^{-3}$
- SGF Euler disc. w/fixed step size: $q = 2.1$, $\eta = 10^{-3}$
- SGF Euler disc. w/fixed step size: $q = 2.5$, $\eta = 10^{-3}$
- SGF Euler disc. w/fixed step size: $q = 2.8$, $\eta = 10^{-3}$
- SGF Euler disc. w/fixed step size: $q = 100$, $\eta = 10^{-2}$
- GD Nesterov acceleration fixed step size: $\eta = 10^{-4}$; $\mu = 0.9$
- RGF Nesterov-like disc. w/fixed step size: $q = 2.2$, $\eta = 10^{-4}$, $\mu = 0.9$
- RGF Nesterov-like disc. w/fixed step size: $q = 3$, $\eta = 10^{-3}$, $\mu = 0.9$
- RGF Nesterov-like disc. w/fixed step size: $q = 6$, $\eta = 10^{-3}$, $\mu = 0.9$
- RGF Nesterov-like disc. w/fixed step size: $q = 10$, $\eta = 10^{-3}$, $\mu = 0.9$
- SGF Runge Kutta disc. w/fixed step size: $q = 2.2$, $K = 2$, $\eta = 10^{-3}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$
- SGF Runge Kutta disc. w/fixed step size: $q = 3$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.9$, $\alpha_1 = \alpha_2 = 0.5$
- SGF Runge Kutta disc. w/fixed step size: $q = 6$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$
- SGF Runge Kutta disc. w/fixed step size: $q = 10$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 0.09$, $\alpha_1 = \alpha_2 = 0.5$

Remark 2. Choice of q : The settling time upper bound (15) decreases as $q \rightarrow \infty$, which appears to lead to faster convergence when discretized. On the other hand, the larger q is, the stiffer the ODE, so more prone to numerical instability, so q cannot be too large. Therefore, assuming p to be not too large, it appears that $q \in (p, p + \delta]$ works best, with $\delta > 0$ as small as needed to avoid numerical issues. For instance, if we know the cost function to be strongly convex (locally), then we search for q slightly larger than $p = 2$ at first, but continue to increase until performance deteriorates. If, on the other hand, we don't know the order $p > 1$, then it's currently unclear how to choose q . We will investigate this further in future work. Furthermore, there is evidence that gradient dominance does hold locally in many deep learning contexts (Zhou and Liang, 2017, <https://arxiv.org/abs/1710.06910>). Indeed, since convexity readily leads to gradient dominance of order $p = \infty$, it suffices that a slightly stronger form of it holds (but weaker than strong convexity), in order to have $p < \infty$, and thus for us to be able to choose $q > p$.

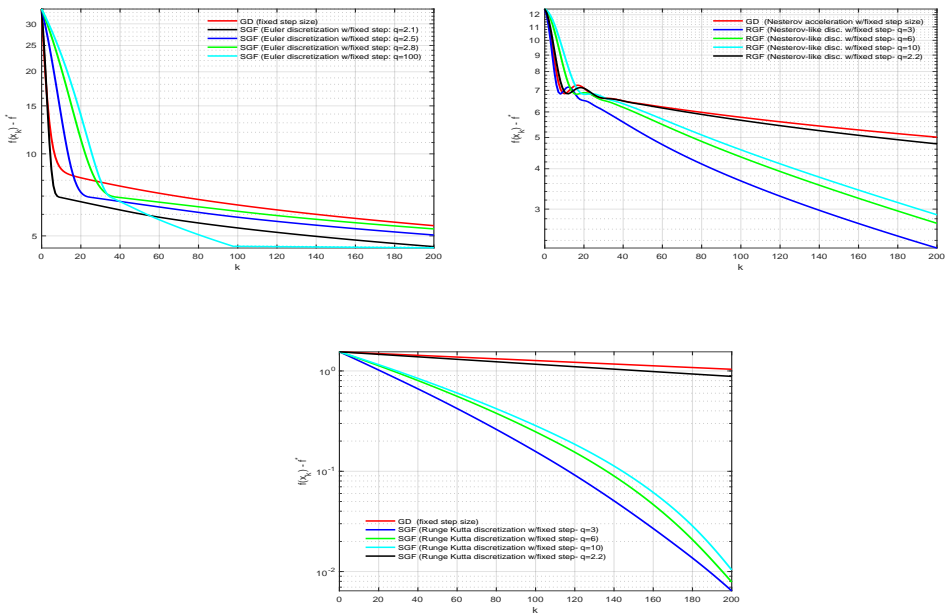


Figure 6: Example of the proposed discretization algorithms of finite-time q -RGF and q -SGF

E.4 EXPERIMENT 3: MNIST CNN MODEL

In this experiment we optimize the CNN network described using a Pytorch code sequence in **MODEL 1** with a negative log likelihood loss, on the MNIST dataset.

We use 10 epochs of training, with 60 batches of 1000 images for training, and 10 batches of 1000 images for testing. We tested Algorithm 1 RGF ($c = 1, q = 3, \eta = 0.06, \mu = 0.9$), and Algorithm 2 SGF ($c = 0.001, q = 2.1, \eta = 0.06, \mu = 0.9$) against Nesterov’s accelerated gradient descent (GD) ($\eta = 0.06, \mu = 0.9, \text{Nesterov}=\text{True}$), and Adam ($\eta = 0.004$, remaining coefficients=nominal values of Torch.Optim). We also tested other algorithms such as RMSprop, AdaGrad, and AdaDelta, but since their convergence performance, as well as, test performance were not very competitive w.r.t. GD, on this experiment, we decided not to report them here, to avoid overloading the graphs. In Figures 7, 8 we can see the training loss over the training iterations, where we see that GD, RGF and SGF perform better than Adam in terms of convergence speed (20 sec lead in average), and in terms of test performance 98% for Adam, 98% for GD, and 99% for both RGF and SGF. The RGF and SGF perform slightly better than GD in terms of convergence speed. The gain is relatively small (5 sec to 10 sec in average) which is expected in such a small DNN network (please refer to VGG16-CIFAR10, and VGG16-SVHN test results for larger computation-time gains). We also notice, in Figure 7 that all algorithms behave well in terms of avoiding overfitting the model.




Figure 7: Losses for several optimization algorithms- CNN- MNIST: Train loss (left), test loss (right)




Figure 8: Training loss vs. computation time for several optimization algorithms- MNIST

MODEL 1: MNIST-CNN

```

class Net(MNIST-CNN)
  def init(self):
    super(Net, self).init()
    self.conv1 = nn.Conv2d(1, 10, kernel.size=5)
    self.conv2 = nn.Conv2d(10, 20, kernel.size=5)
    self.conv2.drop = nn.Dropout2d()
    self.fc1 = nn.Linear(320, 50)
    self.fc2 = nn.Linear(50, 10)

  def forward(self, x):
    x = F.relu(F.max.pool2d(self.conv1(x), 2))
    x = F.relu(F.max.pool2d(self.conv2.drop(self.conv2(x)), 2))
    x = x.view(-1, 320)
    x = F.relu(self.fc1(x))
    x = F.dropout(x, training=self.training)
    x = self.fc2(x)
  return F.log.softmax(x)

```

E.5 EXPERIMENT 4: EULER DISCRETIZATION ON SVHN DATASET

In the main paper, due to space limitation, we decided to report the results of the Nesterov-like discretization only, which also seems like a fair comparison since our Nesterov-like discretization of q -RGF and q -SGF flows can be compared against Nesterov implementation of GD algorithm. However, we also wanted to test the performance of the simple Euler discretization of the proposed flows against the simple GD algorithm, to do so we run some extra tests on the SVHN dataset. These results are presented below.

We tested the proposed Euler algorithms to train the same VGG16 CNN model with cross entropy loss. We divided the dataset into a training set of 70 batches with 1000 images each, and a test set of 10 of 1000 images each, and ran 20 epochs of training over all the training batches. We tested the Euler discretization of q -RGF ($c = 1$, $q = 2.1$, $\eta = 0.04$, $\mu = 0.9$), and the Euler discretization of q -SGF ($c = 10^{-3}$, $q = 2.1$, $\eta = 0.04$, $\mu = 0.9$) against gradient descent (GD) and Adam (same optimal tuning as in Section 4.2). All algorithms have been implemented in their stochastic version.

In Figures 9, 10 we can see that both algorithms, Euler q -RGF and Euler q -SGF, converge faster (40 *min* lead in average) than SGD and Adam for these tests, and reach an overall better performance on the test-set.

E.6 EXPERIMENT 5: RUNGE-KUTTA DISCRETIZATION ON SVHN DATASET

Finally, for completeness, we also wanted to test the performance of the Runge-Kutta discretization of the proposed flows against SGD, to do so we run some extra tests on the SVHN dataset. These results are presented below.

We tested the proposed Runge-Kutta algorithms to train the same VGG16 CNN model with cross entropy loss. We divided the dataset into a training set of 70 batches with 1000 images each, and a test set of 10 of 1000 images each, and ran 20 epochs of training over all the training batches. We tested the Runge-Kutta discretization of q -RGF ($c = 1$, $q = 2.1$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 10^{-2}$, $\alpha_1 = \alpha_2 = 0.5$), and the Runge-Kutta discretization of q -SGF ($c = 10^{-3}$, $q = 2.1$, $K = 2$, $\eta = 10^{-2}$, $\beta_1 = 10^{-2}$, $\alpha_1 = \alpha_2 = 0.5$) against




Figure 9: Losses for several optimization algorithms- SVHN: Train loss (left), test loss (right)




Figure 10: Training loss vs. computation time for several optimization algorithms- VGG-16- SVHN

gradient descent (GD) and Adam (same optimal tuning as in Section 4.2). All algorithms have been implemented in their stochastic version.

In Figures 11, 12 we can see that both algorithms, Runge-Kutta q -RGF and Runge-Kutta q -SGF converge faster (40 *min* lead in average) than SGD and Adam for these tests.




Figure 11: Losses for several optimization algorithms- SVHN: Train loss (left), test loss (right)




Figure 12: Training loss vs. computation time for several optimization algorithms- VGG-16- SVHN