

Generalized Gradient Flows with Provable Fixed-Time Convergence and Fast Evasion of Non-Degenerate Saddle Points

Mayank Baranwal, Param Budhraja, Vishal Raj, and Ashish R. Hota, *Member, IEEE*

Abstract—Gradient-based first-order convex optimization algorithms find widespread applicability in a variety of domains, including machine learning tasks. Motivated by the recent advances in fixed-time stability theory of continuous-time dynamical systems, we introduce a generalized framework for designing accelerated optimization algorithms with strongest convergence guarantees that further extend to a subclass of non-convex functions. In particular, we introduce the *GenFlow* algorithm and its momentum variant that provably converge to the optimal solution of objective functions satisfying the Polyak-Łojasiewicz (PL) inequality in a fixed time. Moreover, for functions that admit non-degenerate saddle-points, we show that for the proposed GenFlow algorithm, the time required to evade these saddle-points is uniformly bounded for all initial conditions. Finally, for strongly convex-strongly concave minimax problems whose optimal solution is a saddle point, a similar scheme is shown to arrive at the optimal solution again in a fixed time. The superior convergence properties of our algorithm are validated experimentally on a variety of benchmark datasets.

Index Terms—Accelerated optimization, Fixed-time convergence, Minimax problem, Continuous-time optimization, Saddle point evasion

I. INTRODUCTION

Consider the unconstrained optimization problem

$$\min_{x \in \mathbb{R}^n} f(x), \quad (1)$$

where the function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a differentiable function. A large class of algorithms have been proposed to find an optimal solution of the above problem. Due to several computationally attractive properties, the class of gradient-based first-order optimization algorithms have seen wide applicability in machine learning tasks [1], [2]. In order to obtain faster convergence guarantees, accelerated gradient methods have been proposed and analyzed as well [3]–[5].

The above class of algorithms iteratively update the candidate solutions based on (local estimate of) the gradient of

Mayank Baranwal is with the Tata Consultancy Services Research, Mumbai, Maharashtra 400607, India (e-mail: baranwal.mayank@tcs.com).

Param Budhraja is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215, USA (e-mail: paramb@bu.edu).

Vishal Raj and Ashish R. Hota are with the Department of Electrical Engineering, Indian Institute of Technology, Kharagpur, West Bengal 721302, India (e-mails: vishalrajroy@iitkgp.ac.in, ahota@ee.iitkgp.ac.in).

the objective function. Thus, the evolution of the candidate solutions assumes the form of a discrete-time difference equation. Convergence of the iterates to the optimal solution are typically analyzed from a discrete-time viewpoint [1], [2].

More recently, inspired by the seminal paper [6], several works have started to view the evolution of the candidate solutions from the lens of a continuous-time dynamical system whose equilibrium point coincides with the optimal solution, and convergence to the equilibrium is tied to its stability property. For instance, the simple gradient descent algorithm for minimizing a convex function f can be mapped to a continuous-time dynamical system using the gradient flow $\dot{x}(t) = -\nabla f(x(t))$. Analysis of the resulting nonlinear dynamical system not only provides fresh insights into the convergence behavior of the underlying optimization algorithm [6], [7], but it also enables leveraging advanced results from the field of control theory [8], [9] to propose new algorithms with improved stability properties associated with the optimal solution leading to accelerated convergence guarantees and under more general assumptions [10]–[13].

In particular, the notion of finite-time stability [14] was examined in [15] and [12] in the context of continuous-time gradient-flows. In this work, we focus on the recently developed notion of fixed-time stability (FxTS) [16], which strengthens the notion of finite-time stability and provides fastest convergence guarantees. An equilibrium point is globally fixed-time stable if the dynamics reaches the equilibrium point in a finite amount of time which can be uniformly upper bounded for all initial conditions. Leveraging the above notion of stability, the authors in [13], [17] propose the fixed-time stable gradient flow (FxTS-GF) scheme, given by

$$\dot{x} = -c_1 \frac{\nabla f(x)}{\|\nabla f(x)\|^{\frac{p-2}{p-1}}} - c_2 \frac{\nabla f(x)}{\|\nabla f(x)\|^{\frac{q-2}{q-1}}},$$

with $c_1, c_2 > 0$, $p > 2$, $q \in (1, 2)$ under which the optimizer of f is shown to be a fixed-time stable equilibrium when f satisfies the PL-inequality [18]. In particular, a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies PL-inequality with modulus $\mu > 0$ if

$$\frac{1}{2} \|\nabla f(x)\|^2 \geq \mu (f(x) - f^*), \quad (2)$$

where $f^* := \min_{x \in \mathbb{R}^n} f(x)$ is the minimum value of the function. Note that all strongly-convex functions satisfy PL-inequality, however, the converse is not true in general. In fact, functions that satisfy the PL-inequality need not be convex. For

instance, the cost function in the logistic regression problem over a compact set is not strongly-convex, nevertheless, it satisfies the PL-inequality. Among non-convex problems of practical interest, recent studies [19], [20] demonstrate that both over-parameterized non-linear systems and neural network loss landscapes predominantly satisfy the PL condition on most (but not all) of the parameter space. This property ensures the existence of solutions and enables effective optimization through stochastic gradient descent (SGD). Notably, for functions satisfying PL-inequality, every stationary point is a global minimum.¹

The above scheme (FxTS-GF) is closely related to the class of first order algorithms with a re-scaled or normalized gradient descent [23]–[25] which have been shown to escape non-degenerate saddle points² more efficiently compared to the conventional gradient descent scheme; the latter may take exponentially long to escape certain saddle points.

The algorithms discussed above, namely (stochastic) gradient descent (SGD) [26], rescaled gradient descent [6], [24], [25] and fixed-time stable gradient flow [13] are characterized by constant learning rates across each dimension of the decision variable. As a result, the descent magnitude often gets dominated by the largest absolute component of the gradient vector, leading to larger descent steps in some dimensions and smaller steps in others. This poses a significant challenge in high-dimensional non-convex problems, such as the ones that arise in the context of training neural networks. In contrast, algorithms such as Adam [27] often exhibit improved convergence behavior due to element-wise normalization of the gradient vector resulting in uniform scaling of effective learning rates across each dimension [28], [29].

A. Statement of Contributions

In this work, we propose a novel continuous-time gradient flow, termed *GenFlow*, which includes element-wise normalization to adaptively scale the learning rates for each dimension. It generalizes element-wise scaling algorithms, such as Adam [27], to guarantee (accelerated) fixed-time convergence for a large class of convex as well as non-convex cost functions. We further provide bounds on escape time from non-degenerate saddle points. Our main contributions are summarized below.

- 1) **Fixed-Time and accelerated convergence for a subclass of non-convex problems:** Most existing accelerated gradient algorithms provide faster convergence guarantees only for a special class of convex functions satisfying strong convexity and/or Lipschitz-smoothness [3], [30]. In contrast, we prove that under the proposed GenFlow scheme and its momentum variant, the optimal solution is fixed-time stable for the class of

¹Although verifying the PL condition is challenging in the absence of a known globally optimal solution, Theorem 2 in [18] presents various implications that can potentially be utilized to verify the PL inequality. Similarly, in the context over-parameterized learning [21], [22], the existence of zero-loss solutions makes it tractable to verify PL condition in important special cases.

²A non-degenerate saddle point is a maximum in some directions and a minimum in others. See Section III for a formal definition.

cost functions satisfying the PL inequality (2), which also includes non-convex cost functions.

- 2) **Fixed-time evasion of non-degenerate saddle points:** In high-dimensional non-convex optimization problems, including problems that arise in the context of training neural networks, the proliferation of non-degenerate saddle points induces a significant challenge for optimization algorithms [31]. While the gradient descent algorithm can evade saddle points and converge to the local minima [32], the learning rate could be vanishingly slow and time to escape from the vicinity of a saddle point can be exponentially large [24], [31], [33]. Recent work has shown both empirically and theoretically that GD with normalization tends to evade saddle points more efficiently [23], [24], [34]. We show that under the proposed GenFlow scheme, the time required by the trajectory to escape any neighborhood of the saddle point can be upper bounded by a constant which is independent of the size of the neighborhood, which is in contrast with the results in [24], [31], [33].
- 3) **Accelerated convergence for minimax problems:** In addition to cost minimization problems, the class of minimax or *saddle point* problems arise in several settings including game theory [35], machine learning [36], [37] and statistics [38]. We extend the proposed GenFlow scheme to this class of problems and show that for a strongly convex-strongly concave objective function, a saddle point can be reached uniformly in a fixed time.
- 4) **Empirical performance:** The proposed GenFlow algorithm and its momentum variant are used for training neural networks and generative adversarial networks (GANs) and are shown to outperform state-of-the-art (SOTA) algorithms including Adam, root mean squared propagation (RMSProp), stochastic variance reduced gradient (SVRG) and fixed-time stable gradient flow (FxTS-GF).

Our proposed work introduces significant improvements and novel contributions compared to similar fixed-time convergent gradient flow schemes in the literature [13], [17]. We introduce adaptive gradient scaling through element-wise gradient normalization, which further accelerates convergence speed to optimal solutions when compared with FxTS-GF [13], [17] (see Fig. 2 for details). We also demonstrate, arguably for the first time, that GenFlow can escape non-degenerate saddle points in a fixed time. This class of non-convex functions was not addressed in [13], [17]. In addition, we develop a smooth momentum variant of GenFlow which is shown to converge to optimal solutions of functions satisfying the PL inequality in a fixed time. This is in contrast to [13] where the authors assume strong convexity and Lipschitz-smoothness on the cost function to design a momentum variant of a fixed-time convergent algorithm that is discontinuous and exhibits significant chattering near the optimal solution. We further extend the scope of GenFlow to solve minimax optimization problems in a fixed time, which commonly arise in two-player games, GANs, among others.

B. Technical Preliminaries

1) *Notation*: The set of all real numbers is denoted by \mathbb{R} , and the n -dimensional Euclidean space is denoted by \mathbb{R}^n . For a vector $x \in \mathbb{R}^n$, the notation x^\top is used for its transpose. The ℓ_p norm for $p \geq 1$ is denoted by $\|\cdot\|_p$, and when p is not specified $\|\cdot\|$ denotes the ℓ_2 norm. The set of functions $f : U \rightarrow V$, where $U \subseteq \mathbb{R}^n$ and $V \subseteq \mathbb{R}^m$, that are k -times continuously differentiable is denoted by $C^k(U, V)$. For a function $f \in C^1(\mathbb{R}^n, \mathbb{R})$, ∇f denotes its gradient. The i^{th} component of the gradient vector ∇f is denoted by $\nabla_i f$. The Hessian of $f \in C^2(\mathbb{R}^n, \mathbb{R})$ is denoted by $D^2 f$.

The set of all eigenvalues of a matrix $A \in \mathbb{R}^{n \times n}$ is called its spectrum, and is denoted by $\sigma(A)$. The minimum absolute value among all eigenvalues of A is denoted by $|\lambda|_{\min}(A)$, i.e., $|\lambda|_{\min}(A) = \min_{\lambda \in \sigma(A)} |\lambda|$. Similarly $|\lambda|_{\max}(A)$ denotes the maximum absolute value of all eigenvalues of A . The identity matrix of size $n \times n$ is denoted by I_n . For a square matrix A , we define $|A| := \sqrt{A^\top A}$.

The open ball of radius r around $x \in \mathbb{R}^n$ is denoted by $B_r(x) := \{y \in \mathbb{R}^n : \|y - x\|_2 < r\}$. For $n \geq 1$, the n -dimensional Lebesgue measure is denoted by \mathcal{L}^n .

2) *Background and Preliminary Results*: We here present a few preliminary results that enable us to establish fixed-time convergence guarantees. We first formally define fixed-time stability of an equilibrium point and then present a sufficient condition for it.

Definition 1 (Fixed-Time Stability [16]): Consider the autonomous differential equation:

$$\dot{x}(t) = g(x(t)), \quad \text{with} \quad g(x^*) = 0. \quad (3)$$

The equilibrium x^* of (3) is defined to be globally *fixed-time stable* (FxTS), if it is Lyapunov stable, and there exists a uniform settling time $\bar{T} < \infty$, such that for any initial condition $x(0) \in \mathbb{R}^n$, we have

$$x(t) = x^* \quad \text{for all} \quad t \geq \bar{T}.$$

Lemma 1 (Fixed-time stability [16]): Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ with $V \in C^1(\mathbb{R}^n, \mathbb{R})$ be a continuously differentiable Lyapunov function with $V(x^*) = 0$, $V(x) > 0$ for all $x \in \mathbb{R}^n \setminus \{x^*\}$, and

$$\dot{V}(x(t)) \leq -aV(x(t))^{\gamma_1} - bV(x(t))^{\gamma_2},$$

with $a, b > 0$, $\gamma_1 \in (0, 1)$ and $\gamma_2 > 1$, then the equilibrium x^* of (3) is globally FxTS with

$$\bar{T} \leq \frac{1}{a(1-\gamma_1)} + \frac{1}{b(\gamma_2-1)}.$$

Remark 1: Lemma 1 not only provides a sufficient condition for an equilibrium point to be fixed-time stable, but also provides an upper-bound on the uniform settling-time \bar{T} in terms of the parameters a, b, γ_1, γ_2 . Thus, given the total budget on the time required to arrive at the equilibrium point of (3), the parameters can potentially be tuned so that convergence is guaranteed within the prescribed budget.

Our main convergence theorem relies on the following power inequality lemma which is a direct consequence of Hölder's inequality [39].

Lemma 2: For a vector $z = (z_1, z_2, \dots, z_n) \in \mathbb{R}^n$, the following holds:

$$\left| \sum_{i=1}^n z_i \right|^p \leq \begin{cases} \sum_{i=1}^n |z_i|^p, & p \in (0, 1], \\ n^{p-1} \sum_{i=1}^n |z_i|^p, & p > 1. \end{cases} \quad (4)$$

II. FIXED-TIME CONVERGENCE OF GENFLOW

Motivated by the adaptive scaling properties of algorithms with element-wise normalization, and fixed-time stability properties of dynamical systems, we propose the following gradient flow scheme, termed as *GenFlow*, for the problem (1):

$$\dot{x}_i = -\frac{\nabla_i f(x)}{|\nabla_i f(x)|^{\frac{p-2}{p-1}}} - \frac{\nabla_i f(x)}{|\nabla_i f(x)|^{\frac{q-2}{q-1}}} \quad \text{for } i = 1, 2, \dots, n, \quad (5)$$

where $p > 2$ and $q \in (1, 2)$. Unlike the fixed-time convergent flow described in [13], the proposed GenFlow scheme adaptively scales the learning rate for each dimension through an element-wise normalization. We now establish convergence guarantees of GenFlow for minimization of functions that satisfy the PL-inequality (2).

Theorem 1 (GenFlow Fixed-Time Convergence): Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function that satisfies the PL inequality (2) and possesses a unique minimizer. Then, the GenFlow scheme defined by (5) converges to the optimal solution of f in a fixed time independent of the initialization.

Proof: We begin by considering the following candidate Lyapunov function:

$$V(x) := f(x) - f^*, \quad (6)$$

where $f^* = \min_{x \in \mathbb{R}^n} f(x)$. Let $x^* = \operatorname{argmin}_{x \in \mathbb{R}^n} f(x)$ and $f(x^*) = f^*$. Notice that the Lyapunov function is positive definite, i.e., $V(x) > 0$ for all $x \neq x^*$ with $V(x^*) = 0$. The choice of Lyapunov function is also motivated by the definition of PL-inequality. Taking the time-derivative of V along the trajectories of (5) yields:

$$\begin{aligned} \dot{V} &= \sum_{i=1}^n \nabla_i f(x) \dot{x}_i \\ &= -\sum_{i=1}^n |\nabla_i f(x)|^{2-\frac{p}{p-1}} - \sum_{i=1}^n |\nabla_i f(x)|^{2-\frac{q}{q-1}}, \\ &\leq -\left(\sum_{i=1}^n (\nabla_i f(x))^2 \right)^{\frac{p}{2(p-1)}} - n^{\frac{q-2}{2(q-1)}} \left(\sum_{i=1}^n (\nabla_i f(x))^2 \right)^{\frac{q}{2(q-1)}} \\ &\leq -(2\mu V)^{\frac{p}{2(p-1)}} - n^{\frac{q-2}{2(q-1)}} (2\mu V)^{\frac{q}{2(q-1)}}, \end{aligned}$$

where the second last inequality follows from Lemma 2 and the last inequality follows from the PL-inequality. Observe that for the given choices of p and q , we have $\frac{p}{2(p-1)} < 1$ and $\frac{q}{2(q-1)} > 1$, i.e., the sufficient conditions for fixed-time convergence stated in Lemma 1 are satisfied. Thus, the GenFlow scheme guarantees uniform convergence of $x(t)$ to x^* in a fixed time. ■

Remark 2: Note that the Lyapunov inequality in Lemma 1 is sufficient for uniform convergence of the state trajectory

$x(t)$ to x^* in a fixed time, rather than the weaker notion of convergence of $f(x(t))$ to f^* .³ In Theorem 1, we assume the presence of a unique optimizer x^* . However, it is important to note that when a function satisfies the PL inequality, every stationary point is a global minimum. Therefore, even without the requirement of a unique x^* , the convergence of $f(x(t))$ to f^* can still be guaranteed. In such cases, the trajectory of $x(t)$ will converge to one of the stationary points. Note that the uniform settling time can further be obtained in terms of the parameters μ, n, p and q as stated in Lemma 1.

Remark 3: In practice, the scheme (5) is numerically implemented by discretizing the continuous-time dynamics. When the gradient is close to 0, a small constant is often added to the denominators in order to avoid potential divisibility by zero. It is worth noting that this property is not unique to our algorithm, rather is present in any normalized gradient-based algorithm, such as Adam [27] and Adagrad [40], and the inclusion of a small positive constant in the denominator is a common practice to prevent potential division by zero; see, for instance, PyTorch implementation of Adagrad [41] and RMSprop [42]. Nevertheless, in most practical problems, the denominators are rarely if ever precisely zero, and the optimization process is typically terminated when the gradient norm reaches a very small value, such as 10^{-7} , in which case adding the small constant may not even be required.

A. Accelerated Convergence via Momentum: GenFlow(M)

Despite adaptive scaling for each dimension, pathological curvature conditions may halt optimizer's progress along certain directions. In such scenarios, optimizers can be further accelerated using momentum-like updates [43]. In light of this, we introduce the momentum variant of the GenFlow which we refer to as GenFlow(M):

$$\begin{aligned} \dot{v}_i &= \nabla_i f(x) - v_i \left(\frac{1}{|v_i|^{\frac{p-2}{p-1}}} + \frac{1}{|v_i|^{\frac{q-2}{q-1}}} \right), \\ \dot{x}_i &= -v_i - \nabla_i f(x) \left(\frac{1}{|\nabla_i f(x)|^{\frac{p-2}{p-1}}} + \frac{1}{|\nabla_i f(x)|^{\frac{q-2}{q-1}}} \right), \\ &\text{for } i = 1, 2, \dots, n, \end{aligned} \quad (7)$$

with $p > 2$ and $q \in (1, 2)$. Below, we show that the GenFlow(M) dynamics converges to the optimal solution uniformly in a fixed amount of time.

Theorem 2 (GenFlow(M) Fixed-Time Convergence): Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function that satisfies the PL inequality (2) and possesses a unique minimizer. Then, the GenFlow(M) scheme defined by (7) converges to the optimal solution of f in fixed time independent of the initialization.

Proof: Note that the equilibrium point of GenFlow(M) (7) is given by the tuple $(x, v) = (x^*, 0)$. Since f satisfies

³Interestingly, fixed-time convergence of $f(x(t))$ to f^* inherently entails the fixed-time convergence of $x(t)$ to x^* . This, however, does not extend to other notions of stability, such as with exponential stability, where the convergence rate of $f(x(t))$ to f^* does not automatically translate to the convergence rate of $x(t)$ to x^* .

PL-inequality, a natural choice for the candidate Lyapunov function is:

$$V(x, v) := (f(x) - f^*) + \frac{1}{2} \sum_{i=1}^n v_i^2. \quad (8)$$

Taking the time-derivative of V along the trajectories of (7) yields:

$$\begin{aligned} \dot{V} &= \sum_{i=1}^n \nabla_i f(x) \dot{x}_i + \sum_{i=1}^n v_i \dot{v}_i \\ &= - \sum_{i=1}^n (\nabla_i f(x))^2 \cdot \frac{p}{2(p-1)} - \sum_{i=1}^n (\nabla_i f(x))^2 \cdot \frac{q}{2(q-1)} \\ &\quad - \sum_{i=1}^n v_i^{2 \cdot \frac{p}{2(p-1)}} - \sum_{i=1}^n v_i^{2 \cdot \frac{q}{2(q-1)}} \\ &\leq - \left(\sum_{i=1}^n (\nabla_i f(x))^2 \right)^\alpha - \left(\sum_{i=1}^n v_i^2 \right)^\alpha \\ &\quad - n^{\frac{q-2}{2(q-1)}} \left(\left(\sum_{i=1}^n (\nabla_i f(x))^2 \right)^\beta + \left(\sum_{i=1}^n v_i^2 \right)^\beta \right), \end{aligned} \quad (9)$$

where the last inequality follows from Lemma 2 with $\alpha = \frac{p}{2(p-1)} < 1$ and $\beta = \frac{q}{2(q-1)} > 1$. Using the PL-inequality, (9) can further be written as:

$$\begin{aligned} \dot{V} &\leq - (2\mu(f(x) - f^*))^\alpha - \left(\sum_{i=1}^n v_i^2 \right)^\alpha \\ &\quad - \underbrace{n^{\frac{q-2}{2(q-1)}}}_{\kappa} \left((2\mu(f(x) - f^*))^\beta - \left(\sum_{i=1}^n v_i^2 \right)^\beta \right) \\ &\leq -2^\alpha \min(1, \mu^\alpha) \left((f(x) - f^*)^\alpha + \left(\frac{1}{2} \sum_{i=1}^n v_i^2 \right)^\alpha \right) \\ &\quad - \kappa^{2\beta} \min(1, \mu^\beta) \left((f(x) - f^*)^\beta + \left(\frac{1}{2} \sum_{i=1}^n v_i^2 \right)^\beta \right), \\ &\leq -2^\alpha \min(1, \mu^\alpha) V^\alpha - 2\kappa \min(1, \mu^\beta) V^\beta, \end{aligned} \quad (10)$$

where the last inequality is again obtained using Lemma 2. Recall that (10) satisfies the sufficient condition for fixed-time stability stated in Lemma 1, thus guaranteeing fixed-time convergence of $x(t)$ to x^* under GenFlow(M). ■

Under the proposed scheme, element-wise normalization facilitates adaptive gradient step-sizes along each coordinate, while momentum-based updates accumulate momentum in the directions of previous updates to overcome oscillations caused by noisy gradients and smoothly navigate flat regions where gradients approach zero. These two approaches mutually accelerate basic gradient flows, and their combination further enhances the optimization process which we emphasize in our simulation results as well.

Remark 4 (Note on algorithmic implementation): It is important to note that while continuous-time dynamical systems viewpoint enables rigorous convergence analysis of an optimization algorithm, in practice, it is inevitable that a discretized version is implemented on a digital computer. While we consider simple Euler-discretization of GenFlow in

our experiments, GenFlow(M) is discretized using step-sizes (β, η) as follows:

$$v_i(k+1) \approx \beta \left(\frac{v_i(k)}{|v_i(k)|^{\frac{p-2}{p-1}}} + \frac{v_i(k)}{|v_i(k)|^{\frac{q-2}{q-1}}} \right) + (1-\beta) \nabla_i f(x(k)),$$

$$x_i(k+1) = x_i(k) - \eta \left(v_i(k+1) + \frac{\nabla_i f(x(k))}{|\nabla_i f(x(k))|^{\frac{p-2}{p-1}}} + \frac{\nabla_i f(x(k))}{|\nabla_i f(x(k))|^{\frac{q-2}{q-1}}} \right).$$

Here $\beta \in (0, 1)$ can be understood as the momentum parameter, while η represents the learning rate or step-size.

The existing literature on discretization step-size for fixed-time convergent flows primarily consists of results of an existence nature rather than providing explicit bounds [44]. It is worth mentioning that even for optimization algorithms such as Nesterov's accelerated gradient method or other accelerated methods, the upper bound on the step-size is typically derived based on the Lipschitz constant of the gradient of the objective function [45]. However, in our work, we do not assume Lipschitz smoothness, which differentiates our approach. Deriving an upper bound on the discretization step-size remains an interesting direction for future research.

III. SADDLE POINT EVASION IN A FIXED TIME

As discussed earlier, classical gradient descent algorithm often slows down in the neighborhood of a saddle point as the gradient has small magnitude in this neighborhood. In this section, we show that due to the normalization term $|\nabla_i f(x)|^{-\frac{p-2}{p-1}}$, the magnitude of the vector field under the GenFlow (5) scheme does not decrease as rapidly as for GD. As a result, GenFlow can evade saddle points quickly. Specifically, we show that the time required to escape a neighborhood of the saddle point can be upper bounded by a constant independent of the size of the neighborhood. We impose the following assumption for the analysis in this section.

Assumption 1: The objective function f belongs to the class $C^2(\mathbb{R}^n, \mathbb{R})$.

For functions satisfying the above assumption, a non-degenerate saddle point x^* is a point such that $\nabla f(x^*) = 0$ and $D^2 f(x^*)$ is non-singular. A non-degenerate saddle point also has the property that there exists $\lambda \in \sigma(D^2(f(x^*)))$ such that $\lambda < 0$. Saddle points that satisfy these properties are referred to in the literature [34], [46] as strict (or rideable) saddle points. Non-degenerate saddle points also have a useful property that local improvements are always possible in their neighborhood.

Theorem 3 below provides an upper bound on the time spent by the trajectory under the GenFlow scheme in a neighborhood of a non-degenerate saddle point of the function f . The analysis is inspired by similar techniques in [24], and relies on the following key lemmas.

Lemma 3: Assume that function f satisfies Assumption 1 and let x^* denote its non-degenerate saddle point. Let $H := D^2(f(x^*))$. Define $\tilde{d}(x) := \sqrt{(x - x^*)^\top |H|(x - x^*)}$ for $x \in$

\mathbb{R}^n , where $|B| := \sqrt{B^\top B}$ for some square matrix B . Denote $\Lambda_{\max} := |\lambda|_{\max}(D^2 f(x^*))$ and $\Lambda_{\min} := |\lambda|_{\min}(D^2 f(x^*))$. Then, for $a \geq 0$ the following relationships hold:

$$\|x - x^*\| \leq \frac{a}{\sqrt{\Lambda_{\max}}} \implies \tilde{d}(x) \leq a, \quad \text{and} \quad (11)$$

$$\tilde{d}(x) \leq a \implies \|x - x^*\| \leq \frac{a}{\sqrt{\Lambda_{\min}}}. \quad (12)$$

Proof: The detailed proof is included in Appendix I. ■

In \mathbb{R}^n , the notion of distance is usually defined using the Euclidean norm. On the other hand, Lemma 3 introduces an alternative notion of distance $\tilde{d}(x)$, which is used later to prove Theorem 3. We also discuss on the relationship between the Euclidean distance and the alternative distance $\tilde{d}(x)$. The following lemma uses the Taylor's theorem to derive a few useful local approximations.

Lemma 4: Suppose the function f satisfies Assumption 1 and has a non-degenerate saddle point x^* . Then, there exist $k_1 > 0.5$, $k_2 < 1$, such that for some neighbourhood of x^* , the following inequalities hold:

$$|f(x) - f(x^*)| \leq k_1 \tilde{d}(x)^2, \quad (13)$$

$$\|\nabla f(x)\|_1 \geq k_2 \|H(x - x^*)\|_1. \quad (14)$$

Proof: The detailed proof is included in Appendix II. ■

Theorem 3 (Fast Evasion of Saddle-Points): Suppose the function f satisfies Assumption 1. Let x^* denote a non-degenerate saddle point of f . Let $x(t)$ be the solution of GenFlow (5) with initial condition x_0 such that the solution $x(t)$ does not converge to the saddle point x^* . Then, for all sufficiently small $r > 0$, the time spent in $B_r(x^*) \setminus \{x^*\}$ is upper bounded as

$$\mathcal{L}^1(\{t \geq 0 : x(t) \in B_r(x^*) \setminus \{x^*\}\}) \leq \min \left\{ n^{\frac{1}{p-1}} \frac{8k_1}{k_2^{\frac{p}{p-1}}} \frac{\Lambda_{\max}^{\frac{p-2}{2(p-1)}}}{\Lambda_{\min}^{\frac{p}{2(p-1)}}} r^{\frac{p-2}{p-1}}, n^{\frac{1}{q-1}} \frac{8k_1}{k_2^{\frac{q}{q-1}}} \frac{\Lambda_{\max}^{\frac{q-2}{2(q-1)}}}{\Lambda_{\min}^{\frac{q}{2(q-1)}}} r^{\frac{q-2}{q-1}} \right\}, \quad (15)$$

where $k_1 > 0.5$, $k_2 < 1$, $\Lambda_{\max} = |\lambda|_{\max}(D^2 f(x^*))$ and $\Lambda_{\min} = |\lambda|_{\min}(D^2 f(x^*))$.

Proof: The detailed proof is included in Appendix III. ■

Note that in equation (15) the exponent of the term $r^{\frac{p-2}{p-1}}$ is positive whereas the exponent of the term $r^{\frac{q-2}{q-1}}$ is negative. Hence as the radius r of the neighbourhood $B_r(x^*)$ increases the term $r^{\frac{p-2}{p-1}}$ increases whereas the term $r^{\frac{q-2}{q-1}}$ decreases. Thus, at some finite value both the terms must be the same. Let the constant \hat{r} be such that

$$n^{\frac{1}{p-1}} \frac{8k_1}{k_2^{\frac{p}{p-1}}} \frac{\Lambda_{\max}^{\frac{p-2}{2(p-1)}}}{\Lambda_{\min}^{\frac{p}{2(p-1)}}} \hat{r}^{\frac{p-2}{p-1}} = n^{\frac{1}{q-1}} \frac{8k_1}{k_2^{\frac{q}{q-1}}} \frac{\Lambda_{\max}^{\frac{q-2}{2(q-1)}}}{\Lambda_{\min}^{\frac{q}{2(q-1)}}} \hat{r}^{\frac{q-2}{q-1}}. \quad (16)$$

Then, for radius r satisfying the conditions of Theorem 3, it can be said that

$$\mathcal{L}^1(\{t \geq 0 : x(t) \in B_r(x^*) \setminus \{x^*\}\}) \leq n^{\frac{1}{p-1}} \frac{8k_1}{k_2^{\frac{p}{p-1}}} \frac{\Lambda_{\max}^{\frac{p-2}{2(p-1)}}}{\Lambda_{\min}^{\frac{p}{2(p-1)}}} \hat{r}^{\frac{p-2}{p-1}} = n^{\frac{1}{q-1}} \frac{8k_1}{k_2^{\frac{q}{q-1}}} \frac{\Lambda_{\max}^{\frac{q-2}{2(q-1)}}}{\Lambda_{\min}^{\frac{q}{2(q-1)}}} \hat{r}^{\frac{q-2}{q-1}},$$

where \hat{r} satisfies (16). Note that the time required to escape the neighbourhood $B_r(x^*)$ is upper bounded by a constant that is independent of the size of the neighbourhood r .

Remark 5: The above result holds true for initial conditions for which the trajectory does not converge to the saddle point. This is indeed true in most scenarios [32].

IV. FIXED-TIME CONVERGENT FLOWS FOR MINIMAX PROBLEMS

The GenFlow scheme can further be extended to obtain solutions of strongly convex-strongly concave minimax problem:

$$\min_{x \in \mathbb{R}^n} \max_{y \in \mathbb{R}^m} g(x, y). \quad (17)$$

The function $g(\cdot, \cdot)$ is strongly convex-strongly concave with moduli μ_1, μ_2 if it is strongly convex in its first argument with modulus μ_1 and strongly concave in the second argument with modulus μ_2 . We consider the following saddle point dynamics:

$$\begin{aligned} \dot{x}_i &= -\frac{\nabla_{x_i} g(x, y)}{\|\nabla G\|^{\frac{p-2}{p-1}}} - \frac{\nabla_{x_i} g(x, y)}{\|\nabla G\|^{\frac{q-2}{q-1}}}, & \text{for } i = 1, 2, \dots, n, \\ \dot{y}_j &= \frac{\nabla_{y_j} g(x, y)}{\|\nabla G\|^{\frac{p-2}{p-1}}} + \frac{\nabla_{y_j} g(x, y)}{\|\nabla G\|^{\frac{q-2}{q-1}}}, & \text{for } j = 1, 2, \dots, m, \end{aligned} \quad (18)$$

with $p > 2$, $q \in (1, 2)$, and the term in the denominator $\nabla G(x, y) := [\nabla_x g^\top(x, y) \quad -\nabla_y g^\top(x, y)]^\top$. For solving minimax optimization problems, achieving accelerated convergence in one variable does not automatically ensure accelerated convergence in the other variable. To guarantee joint accelerated convergence, it is essential to couple the dynamics of both x and y which is achieved through the norm of the total gradient vector ∇G . Along with the normalization, this coupling is also crucial for imparting fixed-time convergence as shown in the following theorem.

Theorem 4 (Fixed-Time Convergence to Saddle Point):

Suppose the function $g : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ in (17) is strongly convex-strongly concave with moduli μ_1 and μ_2 , respectively. Then, the saddle-point dynamics described in (18) converges to the unique optimal solution uniformly in a fixed amount of time, independent of the initialization.

Proof: Note that the saddle point is also characterized by the first-order stationarity condition, and since the function g is strongly convex-strongly concave in its arguments, the stationary point is indeed the saddle point (or the equilibrium of the saddle point dynamics (18)). The strong convexity-strong concavity translates to the following set of conditions on the second-order partial derivatives:

$$\nabla_{xx}^2 g(x, y) \succcurlyeq \mu_1 I_{n_1} \quad \text{and} \quad \nabla_{yy}^2 g(x, y) \preccurlyeq -\mu_2 I_{n_2}, \quad (19)$$

for every $x \in \mathbb{R}^n, y \in \mathbb{R}^m$. In light of the stationarity conditions and the strongly convex-strongly concave assumptions, a natural choice for the candidate Lyapunov function is:

$$V := \frac{1}{2} \|\nabla_x g(x, y)\|^2 + \frac{1}{2} \|\nabla_y g(x, y)\|^2. \quad (20)$$

Taking time-derivative of V along the trajectories of (18) yields:

$$\dot{V} = \nabla_x g(x, y)^\top \nabla_{xx}^2 g(x, y) \dot{x} + \nabla_y g(x, y)^\top \nabla_{yy}^2 g(x, y) \dot{y}$$

$$+ \underbrace{(\nabla_x g(x, y)^\top \nabla_{xy}^2 g(x, y) \dot{y} + \nabla_y g(x, y)^\top \nabla_{xy}^2 g(x, y) \dot{x})}_{=0}, \quad (21)$$

where the last term evaluates to zero following the saddle point dynamics. Substituting (18) into (21) results in:

$$\begin{aligned} \dot{V} &= \left(-\nabla_x g^\top \nabla_{xx}^2 g \nabla_x g + \nabla_y g^\top \nabla_{yy}^2 g \nabla_y g \right) \left[\frac{1}{\|\nabla G\|^{\frac{p-2}{p-1}}} \right. \\ &\quad \left. + \frac{1}{\|\nabla G\|^{\frac{q-2}{q-1}}} \right] \\ &\leq -\left(\mu_1 \|\nabla_x g(x, y)\|^2 + \mu_2 \|\nabla_y g(x, y)\|^2 \right) \left[\frac{1}{\|\nabla G\|^{\frac{p-2}{p-1}}} \right. \\ &\quad \left. + \frac{1}{\|\nabla G\|^{\frac{q-2}{q-1}}} \right] \end{aligned} \quad (22)$$

where the last inequality follows from (19). Recall that $\|\nabla G\|^2 = \|\nabla_x g(x, y)\|^2 + \|\nabla_y g(x, y)\|^2$, resulting in:

$$\begin{aligned} \dot{V} &\leq -\min(\mu_1, \mu_2) \left(\|\nabla G\|^{2 \cdot \frac{p}{2(p-1)}} + \|\nabla G\|^{2 \cdot \frac{q}{2(q-1)}} \right) \\ &= -\min(\mu_1, \mu_2) \left((2V)^{\frac{p}{2(p-1)}} + (2V)^{\frac{q}{2(q-1)}} \right), \end{aligned} \quad (23)$$

which satisfies the sufficient condition for fixed-time convergence stated in Lemma 1. ■

V. EXPERIMENTS

Having analyzed the theoretically superior convergence behavior of GenFlow and its momentum variant, we now focus our attention to empirical performance of the GenFlow scheme for a variety of tasks, involving (i) fast evasion of non-degenerate saddle points, (ii) training of deep neural networks for classification tasks, (iii) training of generative adversarial networks (GANs) [36]. The GenFlow and its momentum variant were implemented using PyTorch's inbuilt optimizer class, and can thus be easily integrated with any neural network training task. While GenFlow(M) is discretized as discussed in Remark 4, we use a simple forward-Euler discretization for GenFlow. Using similar techniques in [44, Theorem 2], it can be shown that under assumptions of strong-convexity and PL-inequality, the forward-Euler discretization results in (T, ϵ) -close discrete-time approximations of GenFlow and GenFlow(M), however, its detailed analysis is beyond the scope of the current study and left as part of the future work. In line with standard practice for optimization algorithms employing gradient normalization, we include a small numerical constant in the denominator term to prevent possible numerical issues related to division by zero. In the numerical implementation of GenFlow and GenFlow(M), we follow a similar approach by incorporating a small positive constant 10^{-7} to ensure numerical stability.

A. Fast Evasion of Non-Degenerate Saddle Points

To illustrate fast evasion of saddle points, we consider the test function, $f(x, y) = 0.5x^2 + 0.25y^4 - 0.5y^2$, described in [32]. Note that the function is strongly convex-strongly concave in its arguments. The function admits three isolated

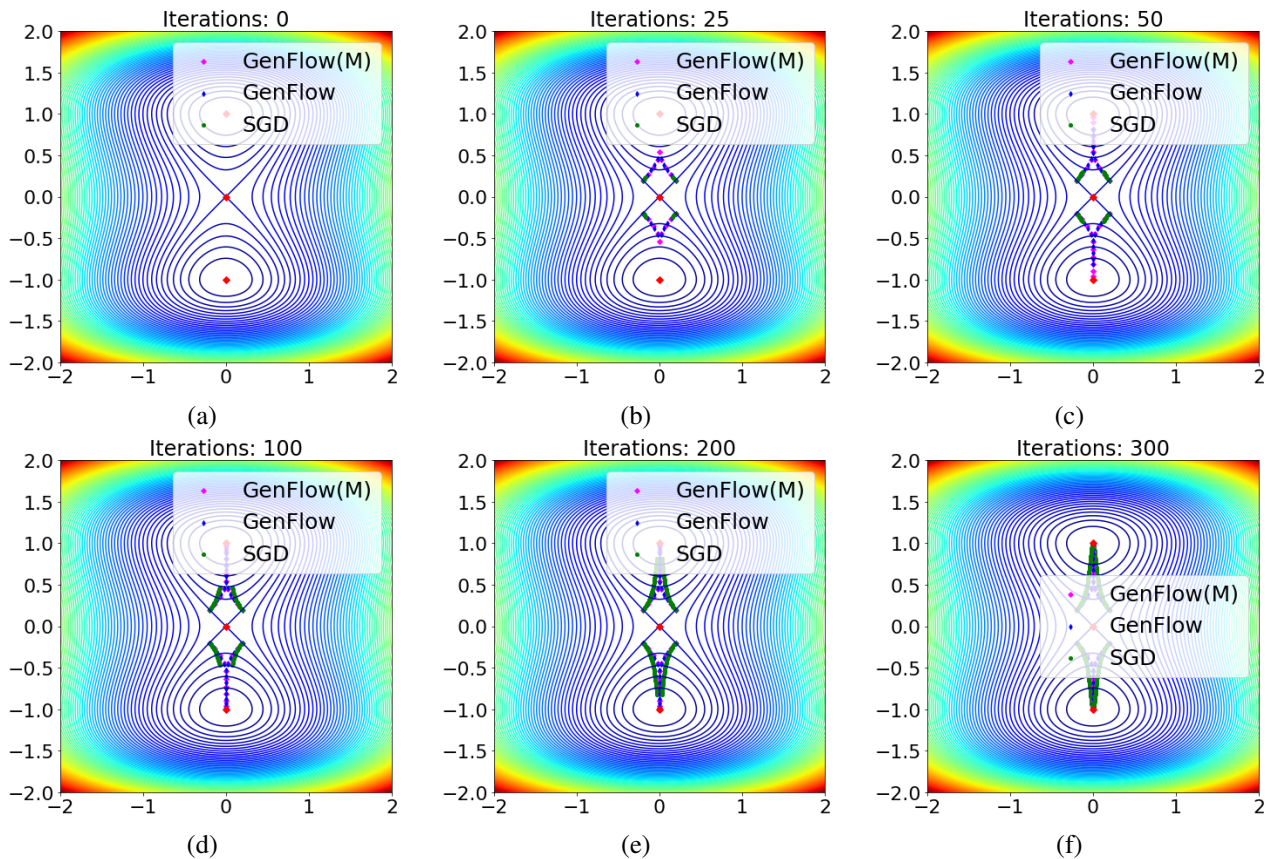


Fig. 1. Snapshots at different iterations for evasion of saddle points. GenFlow(M) is able to evade the saddle points and converge to optimum in less than 50 iterations (slightly better than GenFlow), while gradient descent requires more than 300 iterations for convergence.

critical points: $z_1 = (0, 0)$, $z_2 = (0, -1)$, $z_3 = (0, 1)$, with z_1 being a non-degenerate saddle point, while z_2 and z_3 are isolated local minima. It is well known that the vanilla gradient descent may take exponential time to evade non-degenerate saddle points [33]. We, thus, benchmark GenFlow and GenFlow(M) against the gradient descent for saddle point evasion of the aforementioned test function. It can be seen in Fig. 1 that while all the algorithms successfully escape the saddle point z_1 , the time of evasion is much smaller for GenFlow and GenFlow(M), as opposed to the gradient descent.

Interestingly, while the theoretical results on fixed-time saddle point evasion are specifically presented for GenFlow, similar arguments can be extended to analyze the saddle point evasion characteristics of GenFlow(M) as well. Notably, GenFlow(M) appears to be the fastest among the three methods due to the combined influence of element-wise normalization and accumulation of momentum.

B. Training of Deep Neural Networks

The accelerated convergence behavior of GenFlow for training of neural network classifiers is benchmarked on the two most widely datasets: MNIST (60k training images, 10k test images) [47]; and CIFAR10 (50k training images, 10k test images) [48].

1) *Baseline optimizers*: We benchmark GenFlow(M) against the most widely used optimizers belonging to different families of first-order optimization algorithms: (i) RMSprop [49], (ii) Stochastic Variance Reduced Gradient (SVRG) [50], (iii)

Adam [27], and (iv) FxTS-GF(M) [13].⁴ The hyperparameters for each algorithm are tuned for optimal performance. For instance, it was observed that a large learning rate for Adam or RMSprop significantly destabilizes the training, and hence a suitably tuned smaller learning rate was chosen for classification problems. In particular, we considered the following parameter choices to find the best performing hyperparameters: learning rates (Euler discretization step-sizes) in $\{10^{-4}, 2 \times 10^{-4}, 10^{-3}, 5 \times 10^{-3}, 10^{-2}, 5 \times 10^{-2}\}$ and momentum parameter (β) in $\{0.1, 0.2, \dots, 0.9\}$. Similarly, the exponents (p, q) utilized in gradient normalization of GenFlow and GenFlow(M), have well-defined ranges with the upper bound on convergence time decreasing as $p \rightarrow +\infty$ and $q \rightarrow 0^+$. In order to avoid any numerical instabilities upon discretization, these parameters are tuned gradually towards extreme values starting from $(p, q) = (2.1, 1.99)$.

2) *Evaluation details*: For the classification tasks, the optimization algorithms are evaluated on three criteria: (i) Training cross-entropy loss, (ii) Training accuracy, and (iii) Testing accuracy (for evaluating the quality of local minima in terms of generalization capability).

3) *Network details*: For classification on the MNIST dataset, we have considered a network with a single convolutional layer with ReLU activation (consisting of 32 filters of size 3×3), followed by a dense layer (with ReLU activation) of output size

⁴During our experiments, we observed that the performance of GenFlow is only marginally inferior to that of GenFlow(M). Therefore, in the interest of clearer visualization, we have not included the results from GenFlow for this case study.

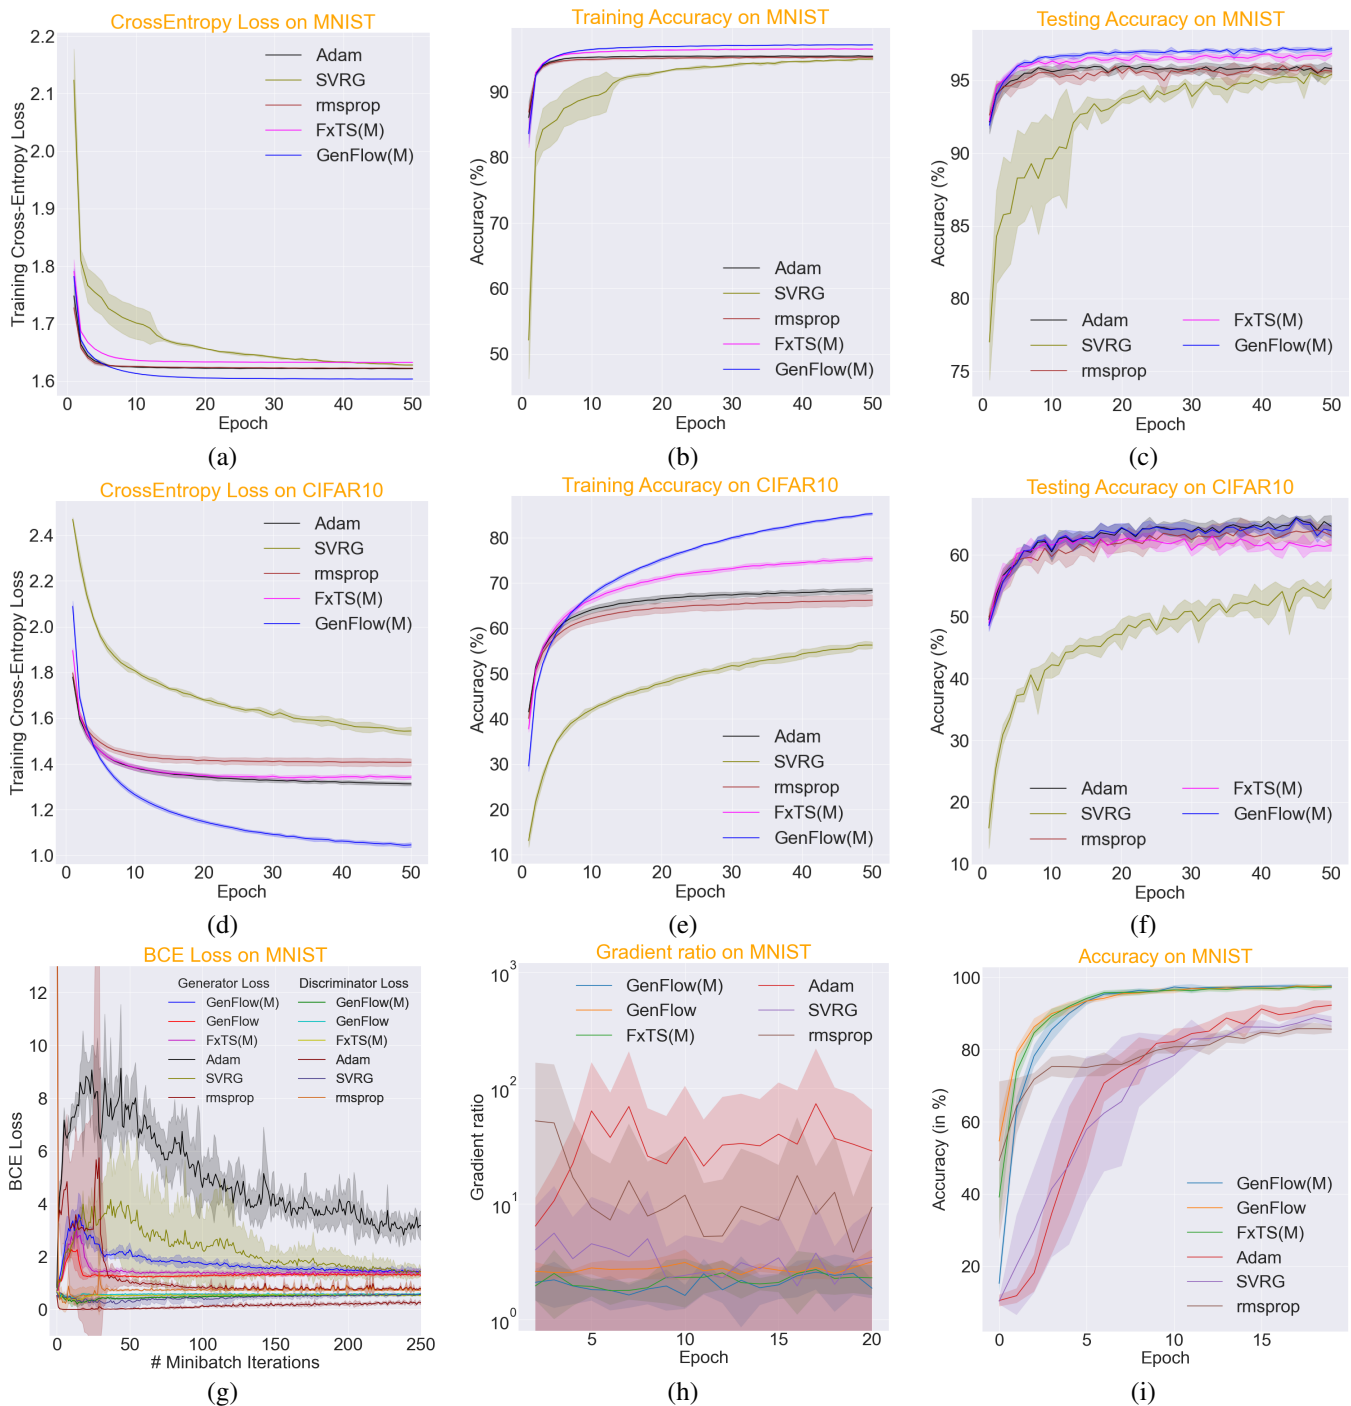


Fig. 2. Benchmarking GenFlow against SOTA optimizers for neural network based classification tasks and training stabilization of GANs: (a), (b), (c) indicate training and generalization performance of several optimizers on the MNIST dataset, while (d), (e), (f) analyze them on the CIFAR10 dataset. Items (g), (h), (i) depict training losses for GANs, gradient ratios for analyzing stabilization, and classification accuracy on images generated after each epoch, respectively.

128. The final linear layer transforms 128-dimensional input to a 10-dimensional output (corresponding to 10 classes) with SoftMax activation. Our architecture for CIFAR10 comprises of two convolutional layers, each followed by a max-pooling layer (with a 2×2 window). The convolutional layers comprise of 6 and 16 filters, each of size 5×5 with ReLU activation, respectively. This is followed by two dense layers with ReLU activation (of sizes 120 and 84, respectively) and an output

layer with SoftMax activation of size 10 (corresponding to 10 classes). We use ℓ_2 -regularized cross-entropy loss for the classification tasks.

4) **Results:** For classification tasks on both MNIST and CIFAR10 datasets, we tune the hyperparameters for optimal performance. In particular, we arrived at the following optimizer settings for classification task on MNIST: (i) Adam ($\text{lr}=10^{-3}$), (ii) SVRG ($\text{lr}=10^{-2}$), (iii) RMSprop ($\text{lr}=10^{-3}$),

(iv) FxTS(M) ($lr=10^{-2}$, momentum=0.7, $(p, q) = 20, 1.98$),
(v) GenFlow(M) ($lr=10^{-2}$, momentum=0.9, $(p, q) = 10, 1.98$).
Fig. 2a represents the training cross-entropy loss as a function of epochs, while Figs. 2b and 2c depict the corresponding training and test accuracy obtained using the aforementioned optimizers, and averaged across five random seeds. It can be seen that GenFlow(M) not only arrives at an optimal solution faster than other methods, the variance across the random seeds is significantly small. We attribute the lower variance of GenFlow(M) to its improved robustness against additive perturbations, a characteristic that can be demonstrated through a theoretical analysis similar to [13]. Both Adam and RMSprop could not be trained efficiently with higher learning rates, and the learning rates had to be reduced to 10^{-3} for optimal performance.

A similar behavior is observed on the CIFAR10 dataset (Figs. 2d-f). While the generalization performance of these optimizers on test dataset is largely very similar (barring the SVRG method), the performance on the training dataset in terms of ℓ_2 -regularized cross-entropy loss, as well as training accuracy, is far better for the proposed GenFlow(M). The optimal performances were obtained using following hyperparameters: (i) Adam ($lr=10^{-3}$), (ii) SVRG ($lr=10^{-3}$), (iii) RMSprop ($lr=10^{-3}$), (iv) FxTS(M) ($lr=5 \times 10^{-3}$, momentum=0.7, $(p, q) = 20, 1.98$), (v) GenFlow(M) ($lr=5 \times 10^{-3}$, momentum=0.5, $(p, q) = 10, 1.98$). Here, lr represents the learning rate of an optimization algorithm.

C. Training of GANs

GenFlow(M) and GenFlow are further evaluated by training GANs, modeled as minimax problems, on MNIST dataset. For this task, the optimization algorithms are evaluated on three criteria: (i) Accuracy of images generated by the Generator (calculated using a pre-trained model available at: github.com/aaron-xichen/pytorch-playground), (ii) Binary Cross Entropy (BCE) Loss of the Generator and the Discriminator, and (iii) Gradient Ratio, the ratio is equal to $\|grad_D^{(t)}\|_2 / \|grad_D^{(0)}\|_2 + \|grad_G^{(t)}\|_2 / \|grad_G^{(0)}\|_2$, where $grad_D^{(0)}$ (respectively $grad_G^{(0)}$) and $grad_D^{(t)}$ (respectively $grad_G^{(t)}$) are the initial and current gradients of Discriminator (respectively Generator) [51].

1) *Network details*: Conditional GAN (cGAN) with the following architecture is used: (i) Generator network consists of an input layer, two hidden layers and an output layer. Both the input and hidden layers comprise of 2D transposed batch normalized convolutional layer with ReLU activation, while the output layer is a simple 2D transposed convolutional layer with tanh activation. The convolutional filter parameters, represented by the tuple (input dimension, output dimension, kernel size, stride), for the four layers in order are (74,256,3,2), (256,128,4,1), (128,64,3,2), and (64,1,4,2). (ii) Discriminator network is a simpler architecture comprising of input, hidden and output layers. Here, we use 2D transposed batch normalized convolutional layer with leaky ReLU activation (with 0.2 negative slope) for the input and the hidden layers. The output layer is a 2D convolutional layer without any nonlinear

activation. The convolutional filter parameters for the three layers in order are (11,64,4,2), (64,128,4,2), and (128,1,4,2).

2) *Results*: The performance of GenFlow(M) and GenFlow is benchmarked against the FxTS-GF(M), Adam, SVRG and RMSprop optimizers. For each optimizer, best hyperparameters are selected based on Accuracy of images generated by the Generator for 20 epochs, averaged over 3 seeds. After finding the best hyperparameters, the plots in Fig. 2g-i are obtained for 5 seeds. The values of best hyperparameter tuple (generator step-size (lr), discriminator to generator step-size ratio, momentum (if applicable)) for GenFlow, GenFlow(M), FxTS-GF(M), Adam, SVRG and RMSprop are (0.01, 5, $p = 2.1$, $q = 1.9$, $\beta_1 = 0.9$, $\beta_2 = 1.8$), (0.005, 5, 0.3, $p = 2.1$, $q = 1.9$), (0.01, 5, 0.2, $p = 2.1$, $q = 1.9$), (0.0002, 10), (0.05, 5), and (0.005, 5, 0.5) respectively. From the Fig. 2i, we can see that GenFlow(M) and GenFlow outperforms other methods by achieving optimal solution faster and having low variance across seeds. Fig. 2h shows that GenFlow(M) and GenFlow has better training stability compared to other methods as the gradient ratio approximately stays near 1 [51]. Other details: (i) Mini-batch size = 128, (ii) Generator Gaussian noise dimension = 64 (iii) GenFlow hyperparameters β_1 and β_2 are multiplied with first and second terms in the RHS of (5) respectively. Images generated by the generator of each of the optimizers after 20 epochs are shown in Fig. 3. From the figure, it can be seen that the images generated by GenFlow(M) and GenFlow are qualitatively superior than the ones generated by Adam, RMSprop and SVRG optimizers. The difference is clearly visible when we compare digits, such as 4, 5, 6, and 7.

VI. CONCLUSION AND FUTURE WORK

In this work, we propose the *GenFlow* gradient-flow scheme and its momentum variant which were shown to converge to the optimum of cost functions satisfying PL-inequality in fixed time (i.e., the convergence time is finite and uniformly upper bounded irrespective of initial candidate solution). We further bound escape time from non-degenerate saddle points and show fixed-time convergence to the solution of minimax problems under suitable assumptions. Experimental results show superior convergence properties compared to state-of-the-art optimizers in training neural networks and GANs.

Future research directions involve analyzing consistent discretization strategies of the proposed GenFlow and GenFlow(M). In particular, we aim to explore the conditions under which the discretized GenFlow algorithm converges to an ϵ -neighborhood around the optimal solution within a uniformly bounded number of iterations. Of particular interest is the provision of upper bounds on the discretization step-size, considering additional assumptions of Lipschitz smoothness. Exploring faster momentum-based methods for solving saddle point problems is an area that has received relatively less attention. In our future work, we aim to extend the proposed saddle-point dynamics to incorporate a momentum variant, which has the potential to further accelerate convergence. Finally, we also aim to analyze the variance reduction properties and robustness of the proposed algorithms to noisy gradients in a follow up work.

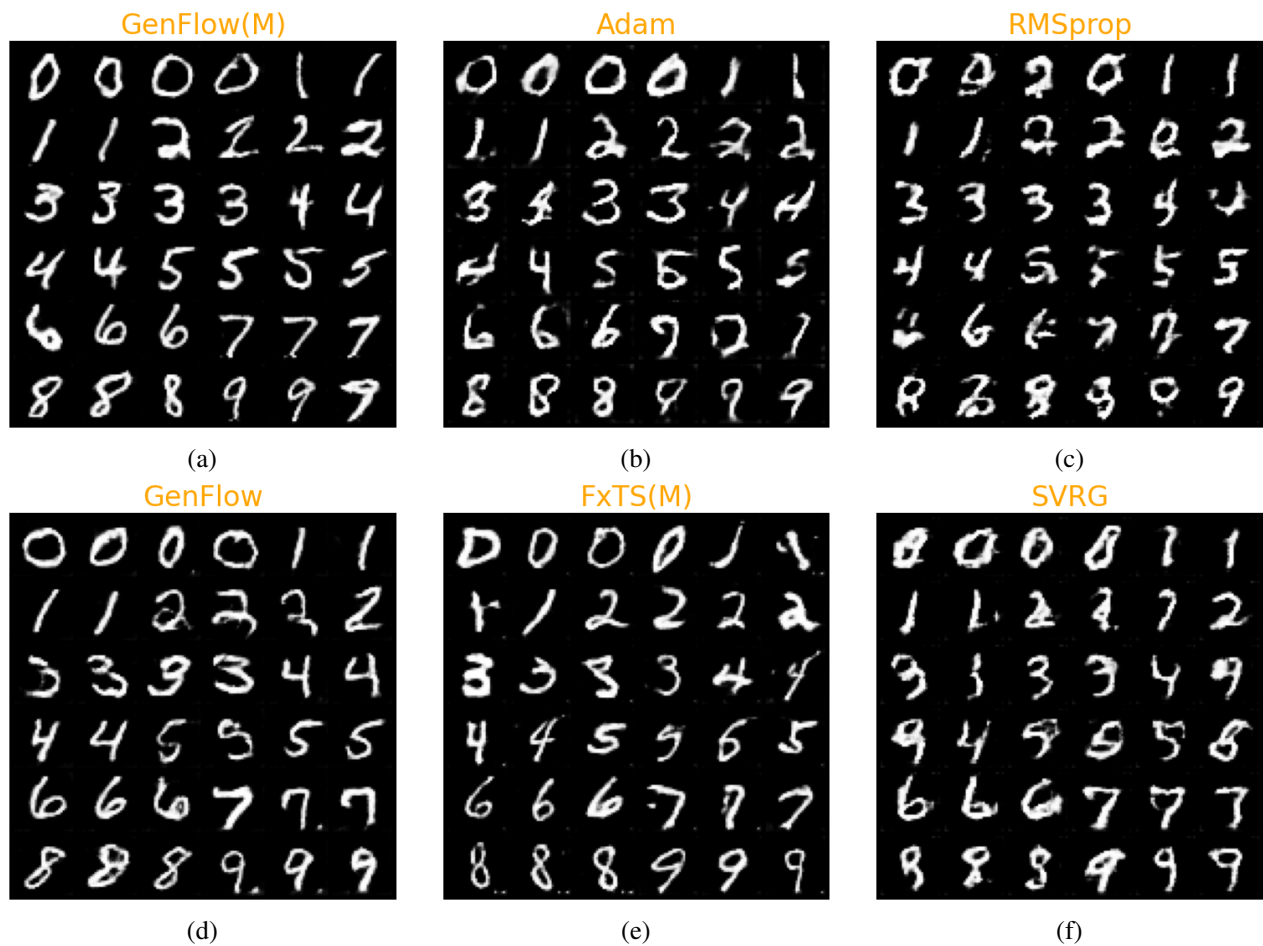


Fig. 3. Comparison of generator generated images after 20 epochs trained using different optimizers.

APPENDIX I PROOF OF LEMMA 3

Proof: The matrix H is a symmetric matrix. Thus its eigenvalue decomposition is given by $H = SDS^T$ where S is an orthogonal matrix and D is a diagonal matrix of the eigenvalues of H . Clearly $|D| = \sqrt{D^T D}$ is a diagonal matrix with modulus of the eigenvalues $\lambda \in \sigma(H)$ at the diagonal. Observe that $H^T H = SD^2 S^T$ and it can be easily proved that $|H| = \sqrt{H^T H} = S|D|S^T$. Now

$$\begin{aligned} \tilde{d}(x) &= \sqrt{(x-x^*)^T |H| (x-x^*)} \\ &= \sqrt{(x-x^*)^T S|D|S^T (x-x^*)} \\ &\leq \sqrt{(x-x^*)^T S\Lambda_{\max} I_n S^T (x-x^*)} \\ &= \sqrt{(x-x^*)^T \Lambda_{\max} (x-x^*)} \\ &= \sqrt{\Lambda_{\max}} \|x-x^*\|, \end{aligned} \quad (24)$$

where to prove the inequality the following fact has been used $\Lambda_{\max} I_n - |D| \succeq 0$. Thus, when $\|x-x^*\| \leq \frac{a}{\sqrt{\Lambda_{\max}}}$, then (11) holds. Similarly, we obtain

$$\begin{aligned} \tilde{d}(x) &\geq \sqrt{(x-x^*)^T S\Lambda_{\min} I_n S^T (x-x^*)} \\ &= \sqrt{\Lambda_{\min}} \|x-x^*\|, \end{aligned} \quad (25)$$

where the inequality is proved using $|D| - \Lambda_{\min} I_n \succeq 0$. If $\tilde{d}(x) \leq a$ then by (25) we obtain $\sqrt{\Lambda_{\min}} \|x-x^*\| \leq a$, i.e., (12) is established. This concludes the proof. \blacksquare

APPENDIX II PROOF OF LEMMA 4

Proof: For the function f , we apply Taylor's theorem at x^* to obtain

$$\begin{aligned} f(x) &= f(x^*) + \nabla f(x^*)^T (x-x^*) \\ &\quad + \frac{1}{2} (x-x^*)^T D^2(f(x^*)) (x-x^*) + R_2(x), \end{aligned}$$

where $R_2(x) \rightarrow 0$ as $\|x-x^*\| \rightarrow 0$. Using the fact $\nabla f(x^*) = 0$, we obtain

$$\begin{aligned} |f(x) - f(x^*)| &= \left| \frac{1}{2} (x-x^*)^T D^2(f(x^*)) (x-x^*) + R_2(x) \right| \\ &\leq \frac{1}{2} |(x-x^*)^T D^2(f(x^*)) (x-x^*)| + |R_2(x)| \\ &\leq \frac{1}{2} |(x-x^*)^T |H| (x-x^*)| + |R_2(x)| \\ &= \frac{1}{2} \tilde{d}(x)^2 + |R_2(x)|, \end{aligned}$$

where to obtain the second inequality, the relation $|H| - H \succeq 0$ is used. Since $\lim_{\|x-x^*\| \rightarrow 0} R_2(x) = 0$ for some

neighbourhood of x^* , for some $k_1 > 0.5$ it is true that $k_1 d(x)^2 \geq \frac{1}{2} \tilde{d}(x)^2 + |R_2(x)|$. Thus, the inequality (13) is proved.

We now invoke Taylor's theorem for $\nabla f(x)$ at x^* to obtain

$$\begin{aligned}\nabla f(x) &= \nabla f(x^*) + D^2(f(x^*))(x - x^*) + R_1(x) \\ &= H(x - x^*) + R_1(x),\end{aligned}$$

where the remainder term $R_1(x)$ satisfies $\lim_{\|x-x^*\|_1 \rightarrow 0} \|R_1(x)\|_1 = 0$. Reverse triangle inequality with respect to ℓ_1 norm yields

$$\begin{aligned}\|\nabla f(x)\|_1 &= \|H(x - x^*) + R_1(x)\|_1 \\ &\geq \|H(x - x^*)\|_1 - \|R_1(x)\|_1.\end{aligned}$$

Now, for some $k_2 < 1$ it can be proved that $\|H(x - x^*)\|_1 - \|R_1(x)\|_1 \geq k_2 \|H(x - x^*)\|_1$ holds for some neighbourhood of x^* which leads to (14), and concludes the proof. ■

APPENDIX III PROOF OF THEOREM 3

Proof: Without loss of generality, assume $x^* = 0$ and let $H := D^2(f(0))$. For $x \in \mathbb{R}^n$, define $\tilde{d}(x) := \sqrt{x^\top |H| x}$, where $|B| := \sqrt{B^\top B}$ for a square matrix B . Since the assumptions of Lemma 3-4 are satisfied, their results can be used here. This also means that for some neighbourhood of $0 \in \mathbb{R}^n$, the relations (13) and (14) hold true. Now

$$\begin{aligned}\|\nabla f(x(t))\|_1 &\geq k_2 \|Hx(t)\|_1 \geq k_2 \|Hx(t)\|_2 \\ &= k_2 \| |H|^{\frac{1}{2}} |H|^{\frac{1}{2}} x(t) \|_2 \\ &\geq k_2 \sqrt{\Lambda_{\min}} \| |H|^{\frac{1}{2}} x(t) \|_2 = k_2 \sqrt{\Lambda_{\min}} \tilde{d}(x(t)).\end{aligned}\tag{26}$$

To derive the second inequality above, the fact $\|v\|_1 \geq \|v\|_2$ for $v \in \mathbb{R}^n$ is used.

Using the chain rule, we obtain

$$\begin{aligned}\frac{d}{dt} f(x(t)) &= \sum_{i=1}^n \left(-|\nabla_i f(x(t))|^{\frac{p}{p-1}} - |\nabla_i f(x(t))|^{\frac{q}{q-1}} \right) \\ &= -\|\nabla f(x(t))\|_{\frac{p}{p-1}}^{\frac{p}{p-1}} - \|\nabla f(x(t))\|_{\frac{q}{q-1}}^{\frac{q}{q-1}} \\ &\leq -n^{\frac{-1}{p-1}} \|\nabla f(x(t))\|_1^{\frac{p}{p-1}} - n^{\frac{-1}{q-1}} \|\nabla f(x(t))\|_1^{\frac{q}{q-1}},\end{aligned}$$

where to derive the inequality, the relation $\|v\|_r \leq n^{\frac{1}{r}-\frac{1}{s}} \|v\|_s$ for $v \in \mathbb{R}^n$ and $s > r \geq 1$ has been used.

We now use the inequality (26) to obtain the following bound on the rate of change of $f(x(t))$:

$$\begin{aligned}-\frac{d}{dt} f(x(t)) &\geq \underbrace{\left(n^{\frac{-1}{p}} k_2 \sqrt{\Lambda_{\min}} \right)^{\frac{p}{p-1}}}_{k_3} \tilde{d}(x)^{\frac{p}{p-1}} \\ &\quad + \underbrace{\left(n^{\frac{-1}{q}} k_2 \sqrt{\Lambda_{\min}} \right)^{\frac{q}{q-1}}}_{k_4} \tilde{d}(x)^{\frac{q}{q-1}} \\ \implies -\frac{d}{dt} f(x(t)) &\geq k_3 \tilde{d}(x)^{\frac{p}{p-1}} + k_4 \tilde{d}(x)^{\frac{q}{q-1}}.\end{aligned}\tag{27}$$

Let the approximations (13) and (26) be true inside the ball $B_{\hat{r}}(0)$ for $\hat{r} > 0$. Suppose that $x(t) \in B_{\hat{r}}(0)$ for $t \in [t_1, t_2]$. Let $e(t) := \tilde{d}(x(t))$ and now integrate (27) to get

$$f(x(t_1)) - f(x(t_2)) \geq k_3 \int_{t_1}^{t_2} e(s)^{\frac{p}{p-1}} ds + k_4 \int_{t_1}^{t_2} e(s)^{\frac{q}{q-1}} ds.$$

Let $r := \kappa^{-\frac{1}{2}} \hat{r}$ where $\kappa = \frac{\Lambda_{\max}}{\Lambda_{\min}}$. For $\eta \leq \sqrt{\Lambda_{\max}} r$ it is true that $\tilde{d}(x) \leq \eta$ implies $x \in B_{\hat{r}}(0)$. Let $t_0 > 0$ be the first time where $e(t) \leq \eta$ and t_3 be the last time when $e(t) \leq \eta$, i.e., $t_3 = \sup\{t \in [0, \infty] : e(t) \leq \eta\}$. The time t_3 is not assumed to be finite at this point but later it is proved that the length of the interval $[t_0, t_3]$ is bounded which in turn implies that t_3 is finite. If $t_3 = \infty$, then in an abuse of notation let $f(x(\infty)) = \lim_{t \rightarrow \infty} f(x(t))$, where note that the limit exists since $f(x(t))$ is monotonically non-increasing by (27). Integrate over the interval $[t_0, t_3]$ to get the following

$$\begin{aligned}f(x(t_0)) - f(x(t_3)) &= \int_{t_0}^{t_3} -\frac{d}{ds} f(x(s)) ds \\ &\geq \int_{e(s) \leq \eta} -\frac{d}{ds} f(x(s)) ds \\ &\geq k_3 \int_{e(s) \leq \eta} e(s)^{\frac{p}{p-1}} ds \\ &\quad + k_4 \int_{e(s) \leq \eta} e(s)^{\frac{q}{q-1}} ds,\end{aligned}$$

where the fact $\frac{d}{dt} f(x(t)) \leq 0$ has been used to derive the first inequality and the second inequality has been derived using the inequality (27) over the intervals where $e(\cdot) \leq \eta$. Add and subtract $f(0)$ in the left hand side of the above inequality along with the bound (13) to get

$$2k_1 \eta^2 \geq k_3 \int_{e(s) \leq \eta} e(s)^{\frac{p}{p-1}} ds + k_4 \int_{e(s) \leq \eta} e(s)^{\frac{q}{q-1}} ds.$$

As $e(\cdot) \geq 0$, the following two inequalities can be deduced

$$\int_{e(s) \leq \eta} e(s)^{\frac{p}{p-1}} ds \leq \frac{2k_1}{k_3} \eta^2,\tag{28}$$

$$\int_{e(s) \leq \eta} e(s)^{\frac{q}{q-1}} ds \leq \frac{2k_1}{k_4} \eta^2.\tag{29}$$

Using Markov's inequality [52], we obtain

$$\begin{aligned}\mathcal{L}^1 \left(\left\{ s : \frac{\eta^{\frac{p}{p-1}}}{2} \leq e(s)^{\frac{p}{p-1}} \leq \eta^{\frac{p}{p-1}} \right\} \right) \\ \leq \frac{2}{\eta^{\frac{p}{p-1}}} \int_{e(s)^{\frac{p}{p-1}} \leq \eta^{\frac{p}{p-1}}} e(s)^{\frac{p}{p-1}} ds \\ = \frac{2}{\eta^{\frac{p}{p-1}}} \int_{e(s) \leq \eta} e(s)^{\frac{p}{p-1}} ds \leq \frac{4k_1}{k_3} \eta^{\frac{p-2}{p-1}},\end{aligned}$$

where to obtain the final inequality the relation (28) has been used. We now iteratively apply the Markov's inequality to obtain

$$\begin{aligned}\mathcal{L}^1 \left(\{s : 0 < e(s)^{\frac{p}{p-1}} \leq \eta^{\frac{p}{p-1}}\} \right) \\ = \sum_{i=0}^{\infty} \mathcal{L}^1 \left(\left\{ s : \frac{\eta^{\frac{p}{p-1}}}{2^{i+1}} \leq e(s)^{\frac{p}{p-1}} \leq \frac{\eta^{\frac{p}{p-1}}}{2^i} \right\} \right)\end{aligned}$$

$$\leq \sum_{i=0}^{\infty} \frac{4k_1}{k_3} \frac{\eta^{\frac{p-2}{p-1}}}{2^i} = \frac{8k_1}{k_3} \eta^{\frac{p-2}{p-1}}.$$

Since $\{s : 0 < e(s) \leq \eta\} = \{s : 0 < e(s)^{\frac{p}{p-1}} \leq \eta^{\frac{p}{p-1}}\}$ one gets

$$\mathcal{L}^1(\{s : 0 < e(s) \leq \eta\}) \leq \frac{8k_1}{k_3} \eta^{\frac{p-2}{p-1}}. \quad (30)$$

Observe that $\{s : 0 < \|x(s)\| \leq r\} \subset \{s : 0 < \tilde{d}(x(s)) \leq \sqrt{\Lambda_{\max}} r\}$. Substituting $\eta = \sqrt{\Lambda_{\max}} r$ in (30), we obtain

$$\mathcal{L}^1(\{s : 0 < \|x(s)\| \leq r\}) \leq n^{\frac{1}{p-1}} \frac{8k_1}{k_2^{\frac{p}{p-1}}} \frac{\Lambda_{\max}^{\frac{p-2}{2(p-1)}}}{\Lambda_{\min}^{\frac{p}{2(p-1)}}} r^{\frac{p-2}{p-1}}. \quad (31)$$

Instead of using (28), one can use (29) and follow similar steps to obtain the following set of inequalities:

$$\begin{aligned} \mathcal{L}^1\left(\{s : 0 < e(s)^{\frac{q}{q-1}} \leq \eta^{\frac{q}{q-1}}\}\right) &\leq \frac{8k_1}{k_4} \eta^{\frac{q-2}{q-1}}, \\ \mathcal{L}^1(\{s : 0 < e(s) \leq \eta\}) &\leq \frac{8k_1}{k_4} \eta^{\frac{q-2}{q-1}}, \\ \mathcal{L}^1(\{s : 0 < \|x(s)\| \leq r\}) &\leq n^{\frac{1}{q-1}} \frac{8k_1}{k_2^{\frac{q}{q-1}}} \frac{\Lambda_{\max}^{\frac{q-2}{2(q-1)}}}{\Lambda_{\min}^{\frac{q}{2(q-1)}}} r^{\frac{q-2}{q-1}}. \end{aligned} \quad (32)$$

Both (31) and (32) are upper bounds on the length of time spent in $B_r(x^*) \setminus \{x^*\}$ so it follows that both are satisfied if the smaller of the two bounds is satisfied. This concludes the proof. ■

REFERENCES

- [1] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *SIAM Review*, vol. 60, no. 2, pp. 223–311, 2018.
- [2] A. Beck, *First-order methods in optimization*. SIAM, 2017.
- [3] B. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Computational Mathematics and Mathematical Physics*, vol. 4, no. 5, pp. 1–17, 1964.
- [4] Y. Nesterov, "A method for unconstrained convex minimization problem with the rate of convergence $o(1/k^2)$," vol. 269. Doklady AN SSR, 1983, pp. 543–547.
- [5] H. Li, C. Fang, and Z. Lin, "Accelerated first-order optimization algorithms for machine learning," *Proceedings of the IEEE*, vol. 108, no. 11, pp. 2067–2082, 2020.
- [6] A. Wibisono, A. C. Wilson, and M. I. Jordan, "A variational perspective on accelerated methods in optimization," *Proceedings of the National Academy of Sciences*, vol. 113, no. 47, pp. E7351–E7358, 2016.
- [7] N. B. Kovachki and A. M. Stuart, "Continuous time analysis of momentum methods," *Journal of Machine Learning Research*, vol. 22, no. 17, pp. 1–40, 2021.
- [8] H. Khalil, *Nonlinear Systems*. Prentice Hall, 2002.
- [9] L. Lessard, B. Recht, and A. Packard, "Analysis and design of optimization algorithms via integral quadratic constraints," *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 57–95, 2016.
- [10] P. Xu, T. Wang, and Q. Gu, "Accelerated stochastic mirror descent: From continuous-time dynamics to discrete-time algorithms," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2018, pp. 1087–1096.
- [11] A. Orvieto and A. Lucchi, "Continuous-time models for stochastic optimization algorithms," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [12] O. Romero and M. Benosman, "Finite-time convergence in continuous-time optimization," in *International Conference on Machine Learning*. PMLR, 2020, pp. 8200–8209.
- [13] P. Budhraj, M. Baranwal, K. Garg, and A. Hota, "Breaking the convergence barrier: Optimization via fixed-time convergent flows," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 6, 2022, pp. 6115–6122.
- [14] S. P. Bhat and D. S. Bernstein, "Finite-time stability of continuous autonomous systems," *SIAM Journal on Control and Optimization*, vol. 38, no. 3, pp. 751–766, 2000.
- [15] J. Cortés, "Finite-time convergent gradient flows with applications to network consensus," *Automatica*, vol. 42, no. 11, pp. 1993–2000, 2006.
- [16] A. Polyakov, "Nonlinear feedback design for fixed-time stabilization of linear control systems," *IEEE Transactions on Automatic Control*, vol. 57, no. 8, pp. 2106–2110, 2011.
- [17] K. Garg and D. Panagou, "Fixed-time stable gradient flows: Applications to continuous-time optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 5, pp. 2002–2015, 2021.
- [18] H. Karimi, J. Nutini, and M. Schmidt, "Linear convergence of gradient and proximal-gradient methods under the polyak-lobasewicz condition," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2016, pp. 795–811.
- [19] S. Frei and Q. Gu, "Proxy convexity: A unified framework for the analysis of neural networks trained by gradient descent," *Advances in Neural Information Processing Systems*, vol. 34, pp. 7937–7949, 2021.
- [20] C. Liu, L. Zhu, and M. Belkin, "Loss landscapes and optimization in over-parameterized non-linear systems and neural networks," *Applied and Computational Harmonic Analysis*, vol. 59, pp. 85–116, 2022.
- [21] S. S. Du, X. Zhai, B. Póczos, and A. Singh, "Gradient descent provably optimizes over-parameterized neural networks," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=S1eK3i09YQ>
- [22] Z. Allen-Zhu, Y. Li, and Z. Song, "A convergence theory for deep learning via over-parameterization," in *International Conference on Machine Learning*. PMLR, 2019, pp. 242–252.
- [23] K. Y. Levy, "The power of normalization: Faster evasion of saddle points," *arXiv preprint arXiv:1611.04831*, 2016.
- [24] R. Murray, B. Swenson, and S. Kar, "Revisiting normalized gradient descent: Fast evasion of saddle points," *IEEE Transactions on Automatic Control*, vol. 64, no. 11, pp. 4818–4824, 2019.
- [25] A. C. Wilson, L. Mackey, and A. Wibisono, "Accelerating rescaled gradient descent: Fast optimization of smooth functions," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [26] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proceedings of COMPSTAT'2010*. Springer, 2010, pp. 177–186.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [28] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, no. 7, 2011.
- [29] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747*, 2016.
- [30] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2003, vol. 87.
- [31] Y. N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, "Identifying and attacking the saddle point problem in high-dimensional non-convex optimization," *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [32] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht, "Gradient descent only converges to minimizers," in *29th Annual Conference on Learning Theory*, ser. Proceedings of Machine Learning Research, vol. 49. PMLR, 2016, pp. 1246–1257.
- [33] S. S. Du, C. Jin, J. D. Lee, M. I. Jordan, B. Póczos, and A. Singh, "Gradient descent can take exponential time to escape saddle points," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, 2017, p. 1067–1077.
- [34] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points - online stochastic gradient for tensor decomposition," in *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, ser. JMLR Workshop and Conference Proceedings, vol. 40. JMLR.org, 2015, pp. 797–842.
- [35] D.-Z. Du and P. M. Pardalos, *Minimax and applications*. Springer Science & Business Media, 1995, vol. 4.
- [36] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [37] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *International Conference on Learning Representations*, 2018.
- [38] J. O. Berger, *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media, 2013.

- [39] Z. Zuo and L. Tie, "Distributed robust finite-time nonlinear consensus protocols for multi-agent systems," *International Journal of Systems Science*, vol. 47, no. 6, pp. 1366–1375, 2016.
- [40] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, no. 61, pp. 2121–2159, 2011.
- [41] "Pytorch Implementation of Adagrad," <https://pytorch.org/docs/stable/generated/torch.optim.Adagrad.html>, accessed: 2023-19-16.
- [42] "Pytorch Implementation of RMSprop," <https://pytorch.org/docs/stable/generated/torch.optim.RMSprop.html#torch.optim.RMSprop>, accessed: 2023-19-16.
- [43] M. Muehlebach and M. I. Jordan, "Optimization with momentum: Dynamical, control-theoretic, and symplectic perspectives," *Journal of Machine Learning Research*, vol. 22, no. 73, pp. 1–50, 2021.
- [44] K. Garg, M. Baranwal, R. Gupta, and M. Benosman, "Fixed-time stable proximal dynamical system for solving MVIPs," *IEEE Transactions on Automatic Control (Early Access)*, pp. 1–8, 2022.
- [45] J. J. Suh, G. Roh, and E. K. Ryu, "Continuous-time analysis of accelerated gradient methods via conservation laws in dilated coordinate systems," in *International Conference on Machine Learning*. PMLR, 2022, pp. 20 640–20 667.
- [46] J. Sun, Q. Qu, and J. Wright, "When are nonconvex problems not scary?" *arXiv preprint*, 2015.
- [47] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [48] A. Krizhevsky, "Learning multiple layers of features from tiny images," *Technical Report, University of Toronto*, 2009.
- [49] G. Hinton, N. Srivastava, and K. Swersky, "Neural networks for machine learning, Lecture 6a, Overview of mini-batch gradient descent," 2012. [Online]. Available: <https://www.cs.toronto.edu/~bonner/courses/2016s/csc321/lectures/lec6.pdf>
- [50] S. J. Reddi, A. Hefny, S. Sra, B. Póczos, and A. Smola, "Stochastic variance reduction for nonconvex optimization," in *International Conference on Machine Learning*. PMLR, 2016, pp. 314–323.
- [51] S. Jelassi, A. Mensch, G. Gidel, and Y. Li, "Adam is no better than normalized SGD: Dissecting how adaptivity improves GAN performance," 2022. [Online]. Available: <https://openreview.net/forum?id=D9SuLzhgK9>
- [52] H. Federer, *Geometric Measure Theory*. Springer, 2014.



Param Budhraja is a PhD student in the Department of Electrical and Computer Engineering at Boston University. He received his B.Tech. and M.Tech. (dual degree) from Indian Institute of Technology (IIT), Kharagpur in 2022 in Electrical Engineering. His research interests broadly lie in the areas of optimization and reinforcement learning.




Vishal Raj is a software engineer at Nutanix, Inc. He received his B.Tech. and M.Tech. (dual degree) in Electrical Engineering with a specialization in Instrumentation and Signal Processing Engineering from Indian Institute of Technology (IIT), Kharagpur in 2022. He also holds a minor in Computer Science and Engineering from Indian Institute of Technology (IIT), Kharagpur.



Mayank Baranwal is a senior scientist with the research division of Tata Consultancy Services in Mumbai. He also holds an Adjunct appointment with the Systems and Control group at the Indian Institute of Technology, Bombay. Prior to joining TCS, he was a postdoctoral scholar in the Department of Electrical and Computer Engineering at the University of Michigan, Ann Arbor. He received his Bachelors in Mechanical Engineering in 2011 from Indian Institute of Technology, Kanpur, and MS in Mechanical

Science and Engineering in 2014, MS in Mathematics in 2015 and PhD in Mechanical Science and Engineering in 2018, all from the University of Illinois at Urbana-Champaign. His research interests are in modeling, optimization, control and inference in network systems with applications to distributed optimization, bioinformatics and computational biology, transportation networks, control of microgrids and deep learning theory. Mayank is a recipient of the Institute Silver Medal in 2011 (from IIT Kanpur), the ME Outstanding Publication Award in 2017 (from the University of Illinois) and the Young Scientist Award (Tata Consultancy Services) in 2022.



Ashish R. Hota is an Assistant Professor in the Department of Electrical Engineering at Indian Institute of Technology (IIT), Kharagpur. He was a postdoctoral researcher at the Automatic Control Laboratory, ETH Zurich in 2018. He received his Ph.D. from Purdue University in 2017, and his B.Tech and M.Tech (dual degree) from Indian Institute of Technology (IIT), Kharagpur, in 2012, all in Electrical Engineering. From 2012 to 2014, he was a research assistant at the University of Waterloo, Canada. He received the Indian National Academy of Engineering (INAE) Young Associate Award in 2023, Outstanding Graduate Researcher Award from the College of Engineering, Purdue University in 2017, and the Institute Silver Medal from IIT Kharagpur in 2012. His research interests lie in the areas of (i) game theory and behavioral decision theory, (ii) stochastic optimization, control and learning and (iii) resilience and security of network systems.