

Dynamic Regret Bounds for Online Nonconvex Optimization

Julie Mulvaney-Kemp, SangWoo Park, Ming Jin, and Javad Lavaei

Abstract—Online optimization problems are well-understood in the convex case, where algorithmic performance is typically measured relative to the best fixed decision. In this paper, we shed light on online nonconvex optimization problems in which algorithms are evaluated against the optimal decision at each time using the more useful notion of dynamic regret. The focus is on loss functions which are arbitrarily nonconvex, but have global solutions that are slowly time-varying. We address this problem by first analyzing the region around the global solution at each time to define time-varying target sets, which contain the global solution and exhibit desirable properties under the projected gradient descent algorithm. All points in a target set satisfy the proximal Polyak-Łojasiewicz inequality, among other conditions. Then, we introduce two algorithms and prove that the dynamic regret for each algorithm is bounded by a function of the temporal variation in the optimal decision. The first algorithm assumes that the decision maker has some prior knowledge about the initial objective function. This algorithm ensures that decisions are within the target set at every time. The second algorithm makes no assumption about prior knowledge. It instead relies on random sampling and memory to find and then track the target sets over time. In this case, the landscape of the loss functions determines the likelihood that the dynamic regret will be small. Numerical experiments validate these theoretical results and highlight the impact of a single low-complexity problem early in the sequence.

I. INTRODUCTION

Nonconvex optimization is ubiquitous in real-world applications, such as the training of deep neural nets [1], matrix sensing/completion [2], [3], state estimation of dynamic systems [4], and the optimal power flow problem [5]. Moreover, most of these practical problems are solved sequentially over time with time-varying input data, leading to online (real-time) versions of the aforementioned examples [4], [6], [7].

In this paper, we study an online nonconvex optimization (ONO) problem whose loss (objective) function changes over discrete time periods, namely,

$$\underset{\mathbf{x} \in \mathbb{S}}{\text{minimize}} \quad f_t(\mathbf{x}) \quad (1)$$

where $t \in \mathbb{Z}_+$ denotes the time and $\mathbb{S} \subseteq \mathbb{R}^n$ is the time-invariant feasible region. At each time $t = 1, \dots, T$ in this ONO framework, the decision maker first chooses an action $\mathbf{x}_t \in \mathbb{S}$ while oblivious to the loss function $f_t : \mathbb{S} \rightarrow \mathbb{R}$. Once the action is played, it is evaluated against f_t , which may be chosen by an adversary in response to the action. Then, the decision maker is granted access to the loss function and its gradient.

Julie Mulvaney-Kemp, SangWoo Park and Javad Lavaei are with the Department of Industrial Engineering and Operations Research at the University of California, Berkeley. Emails: {julie_mulvaney-kemp, spark111, lavaei}@berkeley.edu. Ming Jin is with the Department of Electrical and Computer Engineering at Virginia Tech. Email: jinming@vt.edu. This work was supported by grants from ARO, ONR, AFOSR, and NSF.

The performance of a decision maker, or equivalently an algorithm, in online settings is typically evaluated by a metric called *regret* [8]. In this paper, we exclusively focus on the strictest version of regret, *dynamic regret*, which is defined as

$$\mathbf{Reg}_T^d(\mathbf{x}_1, \dots, \mathbf{x}_T) := \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t^*, \quad (2)$$

where f_t^* denotes the global optimal objective value of (1). Dynamic regret (also called *non-stationary regret*) compares the decision maker's actions to an optimal action at each time t . In comparison, *static regret* (also called *stationary regret* or simply *regret*) compares the decision maker's actions to the best fixed action in hindsight:

$$\mathbf{Reg}_T^s(\mathbf{x}_1, \dots, \mathbf{x}_T) = \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{y} \in \mathbb{S}} \sum_{t=1}^T f_t(\mathbf{y}). \quad (3)$$

In general, nonconvex optimization problems are NP-hard, and therefore commonly used local search algorithms, such as first-order and second-order descent algorithms, may converge to a spurious local minimum (i.e., a local minimum that is not globally optimal). As a result, dynamic regret can be arbitrarily high in a general setting due to the inability to efficiently find a near-optimal point \mathbf{x}_t . The existing works in the literature have derived regret bounds in terms of various quantities, such as the regularity of the comparator sequence [9], the temporal variation in the loss functions [10], the temporal variation in the gradient of the loss functions [11], and the temporal variation in the optimal decision (also called path length or path variation) [12], [13]. Details on many of these variation measures used to bound dynamic regret can be found in [14], where the online convex optimization problem is analyzed. The existing regret bounds for ONO either focus on static regret [15]–[18] or require the loss functions to be weakly pseudo-convex which is a restrictive condition that excludes spurious local minima [19]. In [20], the authors of this paper established probabilistic nonconvexity regret bounds for a variation of the ONO problem in which f_t is known to the decision maker at time t , future loss functions are unknown but fixed, the global minima are “sufficiently superior” to all local minima, and limit points of a continuous-time projected gradient algorithm can be found precisely. Finally, [21] and [22] explored how variability in the input data can help ONO solution trajectories escape non-global local solutions over time but they did not study dynamic regret and focused on asymptotic regret.

The main goal of this paper is to analyze how the quality of the obtained solutions evolves in ONO settings where the global solution changes slowly over time. To this end, we first develop mathematical tools for characterizing the landscape of constrained nonconvex optimization problems and analyze

the behavior of the projected gradient descent algorithm on such problems. There are many conditions in the literature that guarantee linear convergence of local search algorithms. In the unconstrained case when $\mathbb{S} = \mathbb{R}^n$, the Polyak-Łojasiewicz (PL) condition has been proven to be weaker than other common assumptions (such as strong convexity, essential strong convexity, weak convexity, and restricted secant inequality) that guarantee linear convergence [23]. Despite its favorable characteristics, requiring that a function satisfy the PL condition still significantly restricts the type of nonconvex functions that one can study. For instance, functions satisfying the PL condition cannot have local minima that are not globally optimal.

We leverage the generalization of the PL condition for constrained optimization, called the proximal-PL condition (originally proposed in [23]), to study dynamic regret minimization in a non-convex setting. The first contribution of this paper is to establish a target set for each time instance with the property that once the algorithm finds a point in the corresponding target set at a given time, the global minimizers of future problems can be found efficiently. These time-varying target sets are defined with respect to the proximal-PL condition and the global solution. We show several important properties of these sets, including linear convergence to the global minimizer and quadratic growth.

The design and regret analysis of two online algorithms constitute the second contribution of this paper. Specifically, we equip local search algorithms with memory and random exploration and establish dynamic regret bounds for each algorithm in terms of the path length of the optimal decision sequence, when the difference between consecutive points in this sequence is bounded appropriately. The first algorithm assumes that the decision maker has some prior knowledge about the initial function and can start at a point that is within its target set. This algorithm ensures bounded dynamic regret by producing decisions which track the time-varying target sets. The second algorithm obviates this initial condition assumption by using random exploration. In this case, dynamic regret depends on when the decision maker first finds a point within the corresponding target set, as after that time all decisions will track the time-varying target sets. Therefore, the relative volume of the time-varying target sets with respect to the entire feasible domain—a measure of how favorable the loss function landscape is—influences the likelihood that the dynamic regret will be small. In particular, a single low-complexity problem in the sequence can have a large influence on the outcomes.

The remainder of this paper is organized as follows. In Section II, we analyze the optimization problem for each fixed time step, focusing on a neighborhood of the global solution. In Section III, we introduce ONO algorithms, derive bounds on their dynamic regret, and support the analysis with empirical results. Finally, we conclude the paper in Section IV.

A. Notations

Let $\|\cdot\|$ indicate the ℓ^2 -norm of a vector and $|\cdot|$ represent the cardinality of a set. The symbols \mathbb{R}^n and \mathbb{Z}_+ denote the

space of n -dimensional real vectors and the set of positive integers, respectively. The globally optimal objective value of the optimization problem at time t is denoted by f_t^* . If there is a unique global optimum at time t , it will be denoted as \mathbf{x}_t^* , in which case $f_t(\mathbf{x}_t^*) = f_t^*$. The indicator function $\mathbb{I}_{\mathbb{S}}(\mathbf{x})$ returns zero if \mathbf{x} belongs to the set \mathbb{S} and infinity otherwise. We define the projection operator as follows:

$$\Pi_{\mathbb{S}}(\mathbf{x}) := \operatorname{argmin}_{\mathbf{y} \in \mathbb{S}} \|\mathbf{x} - \mathbf{y}\| \quad (4)$$

The tangent cone of a convex set \mathbb{S} at \mathbf{x} is denoted as $\mathbb{T}_{\mathbb{S}}(\mathbf{x})$. The sublevel set \mathcal{L}_t is defined as $\mathcal{L}_t(\alpha) := \{\mathbf{x} \in \mathbb{R}^n \mid f_t(\mathbf{x}) < \alpha\}$. Finally, $\mathbb{P}[\cdot]$ denotes the probability of the argument.

II. THEORETICAL RESULTS FOR A FIXED TIME STEP

A. Properties of the Problem Structure

Throughout this paper, we make the following assumptions on the problem structure:

- 1) The time-invariant feasible region $\mathbb{S} \subset \mathbb{R}^n$ is a compact, convex set known to the decision maker.
- 2) f_t is coercive and differentiable, but potentially nonconvex in \mathbf{x} with many local minima, for all $t \in \{1, 2, \dots, T\}$.
- 3) f_t has a unique global minimum \mathbf{x}_t^* over \mathbb{S} for all $t \in \{1, 2, \dots, T\}$.
- 4) The magnitude of the gradient is bounded above by a positive constant M_1 for all $t \in \{1, 2, \dots, T\}$. That is, $\sup_{\mathbf{x} \in \mathbb{S}, 1 \leq t \leq T} \|\nabla f_t(\mathbf{x})\| \leq M_1$.
- 5) The first derivative of f_t is L -Lipschitz continuous on \mathbb{S} for all $t \in \{1, 2, \dots, T\}$, implying the following inequality for some constant L :

$$f_t(\mathbf{y}) - f_t(\mathbf{x}) \leq \langle \nabla f_t(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2 \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{S}.$$

B. Proximal Polyak-Łojasiewicz Regions

In the context of unconstrained optimization problems, we say that a differentiable function f_t satisfies the Polyak-Łojasiewicz (PL) condition [24] if the following condition holds for some parameter $\mu > 0$:

$$\underbrace{\frac{1}{2} \|\nabla f_t(\mathbf{x})\|^2}_{\text{PL inequality}} \geq \mu(f_t(\mathbf{x}) - f_t^*) \quad \forall \mathbf{x} \in \mathbb{R}^n \quad (5)$$

If a function satisfies the PL condition and the magnitude of its gradient is small at some \mathbf{x} , then the function value at \mathbf{x} will be close to the global minimum. This is the reason why the PL condition is also referred to as the gradient domination condition [25]. For a general (unconstrained) nonconvex optimization problem, first-order methods such as gradient descent may not converge to a global minimizer. However, if a function f_t satisfies the Polyak-Łojasiewicz condition, then every stationary point is a global minimizer. Moreover, PL is one of the most general conditions under which gradient descent offers linear convergence to a global minimizer [23]. Note that, in general, functions satisfying the PL condition may not have a unique global minima.

The top plot in Figure 1 shows an example of a nonconvex function that satisfies the PL condition. On the other hand,




Fig. 1: The top figure shows the nonconvex function $f_1(x) = x^2 + 3 \sin^2(x)$, which satisfies the PL inequality with the parameter $\mu = 1/32$ for all $x \in \mathbb{R}$. The bottom figure shows an example of a nonconvex function that satisfies the PL inequality with the parameter $\mu = 1/32$ only for $x \in \{-55.9, -10.2\} \cup [-5.4, 5.4] \cup [10.2, 55.9]\}$. The function for the bottom figure is given below:

$$f_2(x) = \begin{cases} -\frac{1}{3\pi}x^3 - \frac{5}{2}x^2 - 6\pi x - \frac{17}{6}\pi^2, & \text{if } x < -\pi \\ x^2 + 3 \sin^2(x), & \text{if } -\pi \leq x \leq \pi \\ \frac{1}{3\pi}x^3 - \frac{5}{2}x^2 + 6\pi x - \frac{17}{6}\pi^2, & \text{if } \pi < x \end{cases}$$

the function in the bottom plot of Figure 1 manifests spurious local minima and therefore cannot satisfy the PL inequality for all \mathbf{x} for any $\mu > 0$. However, for a given μ , we can identify a subset of \mathbb{R} that satisfies the PL inequality. The idea of focusing on regions where the PL inequality is satisfied, rather than only considering functions satisfying the PL condition over the entire feasible region, leads to our definition of time-varying target sets in Section II-C.

Next, we return to considering constrained optimization problems. Constrained optimization can be cast in the framework of unconstrained optimization by appending the objective function with $\mathbb{I}_{\mathbb{S}}$, an indicator function of a convex set \mathbb{S} . This indicator function is non-smooth and convex. Subsequently, a natural generalization of gradient descent to the constrained case is the proximal gradient method, whose iteration is described by

$$\mathbf{x}_t^{k+1} = \underset{\mathbf{y}}{\operatorname{argmin}} \left[\langle \nabla f_t(\mathbf{x}_t^k), \mathbf{y} - \mathbf{x}_t^k \rangle + \frac{\|\mathbf{y} - \mathbf{x}_t^k\|^2}{2s} + \mathbb{I}_{\mathbb{S}}(\mathbf{y}) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^k) \right] \quad (6)$$

for every $k \in \mathbb{Z}_+$, where s is a positive constant. It can be shown that the above algorithm is equivalent to the projected gradient descent algorithm:

$$\mathbf{x}_t^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}_t^k - s \nabla f_t(\mathbf{x}_t^k)). \quad (7)$$

A matching generalization of the PL inequality, namely the proximal-PL inequality, was first proposed in [23].

Definition 1. (Proximal-PL inequality) For a function f_t , define the proximal-gradient with parameter $\beta > 0$ as

$$\mathcal{D}_t(\mathbf{x}, \beta) = -2\beta \min_{\mathbf{y}} \left[\langle \nabla f_t(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\beta \|\mathbf{y} - \mathbf{x}\|^2}{2} + \mathbb{I}_{\mathbb{S}}(\mathbf{y}) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}) \right] \quad (8)$$

We say that a point $\mathbf{x} \in \mathbb{S}$ satisfies the proximal-PL inequality with the parameters $\mu > 0$ and $\beta > 0$ if

$$\frac{1}{2} \mathcal{D}_t(\mathbf{x}, \beta) \geq \mu(f_t(\mathbf{x}) - f_t^*). \quad (9)$$

While [23] considers functions that satisfy the proximal-PL inequality at all points in \mathbb{S} , in this work we instead identify a subset of the entire space that satisfies the inequality. Hereby, we define the time-varying *proximal-PL region*, denoted $\mathcal{P}_t(\mu, \beta)$, as the set of all $\mathbf{x} \in \mathbb{S}$ satisfying the proximal-PL inequality with the parameters $\mu > 0$ and $\beta > 0$. That is,

$$\mathcal{P}_t(\mu, \beta) := \left\{ \mathbf{x} \in \mathbb{S} \mid \frac{1}{2} \mathcal{D}_t(\mathbf{x}, \beta) \geq \mu(f_t(\mathbf{x}) - f_t^*) \right\}. \quad (10)$$

Note that by virtue of the equivalence between equations (6) and (7), the proximal-gradient can also be expressed as follows:

$$\mathcal{D}_t(\mathbf{x}, \beta) = -2\beta \left[\langle \nabla f_t(\mathbf{x}), \Pi_{\mathbb{S}}(\mathbf{x} - \frac{1}{\beta} \nabla f_t(\mathbf{x})) - \mathbf{x} \rangle + \frac{\beta}{2} \|\Pi_{\mathbb{S}}(\mathbf{x} - \frac{1}{\beta} \nabla f_t(\mathbf{x})) - \mathbf{x}\|^2 \right]. \quad (11)$$

C. Time-Varying Regions of Attraction and Target Sets

A proximal-PL region can span over multiple regions of attraction associated with different local minima. Also, note that a region of attraction (RoA) is algorithm dependent. In this paper, we define RoAs with respect to the global minimizer under the projected gradient descent method and also under the projected gradient flow system, as a continuous version of the former.

Definition 2. (Region of Attraction) The region of attraction of a global minimizer \mathbf{x}_t^* of f_t based on the projected gradient descent method (discrete algorithm) with the step size s is defined as follows:

$$\mathcal{R}_t^D(s) := \{ \mathbf{x} \mid \lim_{k \rightarrow \infty} \mathbf{x}_t^k = \mathbf{x}_t^* \text{ where } \mathbf{x}_t^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}_t^k - s \nabla f_t(\mathbf{x}_t^k)) \text{ and } \mathbf{x}_t^0 = \mathbf{x} \}. \quad (12)$$

In addition, we define the region of attraction of a global minimizer based on the projected gradient flow system [26] (continuous algorithm) as follows:

$$\mathcal{R}_t^C := \{ \mathbf{x} \mid \lim_{\ell \rightarrow \infty} \mathbf{x}_t(\ell) = \mathbf{x}_t^*, \text{ where } \dot{\mathbf{x}}_t(\ell) = \Pi_{\mathbb{S}(\mathbf{x}_t)}[-\nabla f_t(\mathbf{x}_t(\ell))], \mathbf{x}_t(0) = \mathbf{x} \}. \quad (13)$$

Next, we define a region that we call the target set. In subsequent sections, we will show that if our proposed algorithm enters the target set at any point in time t , then it is possible to approach the global minimizer and track it henceforth.

Definition 3. (Target set) Let \mathcal{RP}_t^D and \mathcal{RP}_t^C denote the subsets of the discrete and continuous RoAs that are contained within the proximal-PL region at time t :

$$\mathcal{RP}_t^D(\mu, \beta, s) := \{\mathbf{x} \mid \mathbf{x}_t^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}_t^k - s\nabla f_t(\mathbf{x}_t^k)), \mathbf{x}_0 = \mathbf{x}, \lim_{k \rightarrow \infty} \mathbf{x}_t^k = \mathbf{x}_t^* \text{ and } \{\mathbf{x}_t^k\}_{k=0}^{\infty} \subset \mathcal{P}_t(\mu, \beta)\} \quad (14)$$

$$\mathcal{RP}_t^C(\mu, \beta) := \{\mathbf{x} \mid \dot{\mathbf{x}}_t = \Pi_{\mathbb{S}(\mathbf{x}_t)}(-\nabla f_t(\mathbf{x}_t)), \mathbf{x}_t(0) = \mathbf{x}, \lim_{\ell \rightarrow \infty} \mathbf{x}_t(\ell) = \mathbf{x}_t^* \text{ and } \mathbf{x}_t(\ell) \in \mathcal{P}_t(\mu, \beta) \forall \ell \geq 0\} \quad (15)$$

We define our target set for time t to be a subset of a sublevel set around the global minimizer that is contained within both \mathcal{RP}_t^D and \mathcal{RP}_t^C :

$$\mathcal{T}_t(\mu, \beta, s) := \mathcal{L}_t(\alpha_t) \cap \mathcal{RA}_t^D \cap \mathcal{RA}_t^C \quad (16)$$

where α_t is the largest α satisfying the following condition:

$$\left(\mathcal{L}_t(\alpha) \cap \mathcal{RA}_t^D \cap \mathcal{RA}_t^C\right) \subseteq \left(\mathcal{RP}_t^D \cap \mathcal{RP}_t^C\right) \quad (17)$$

In summary, all points in each target set are feasible, satisfy the proximal-PL inequality, and lead to the global solution under the continuous and discrete descent methods initialized at those points. Further, the target set is invariant for both of these methods. Note that the sets \mathcal{RA}_t^D , \mathcal{RA}_t^C , \mathcal{RP}_t^D and \mathcal{RP}_t^C depend on some or all of the parameters (μ, β, s) but this dependency has been omitted in order to simplify notation. As indicated by the subscript t , the target set varies over time.

One useful way to measure the size of a target set is with respect to the global solution.

Definition 4. (Reach) Define the reach of a target set as the maximum distance between the global minimum and any point in the target set:

$$\rho_t(\mu, \beta, s) := \max_{\mathbf{x} \in \mathcal{T}_t(\mu, \beta, s)} \|\mathbf{x}_t^* - \mathbf{x}\|. \quad (18)$$

D. Properties of Target Sets

In [23], the authors showed the linear convergence of the proximal-gradient algorithm when applied to functions satisfying the proximal-PL condition. In this paper, we show that initializing the proximal-gradient algorithm in the corresponding target set ensures linear convergence, regardless of whether the proximal-PL inequality is satisfied for all feasible points. Additionally, there is an open ball around the global solution whose intersection with the feasible set \mathbb{S} is also contained in the corresponding target set (see Appendix B).

Theorem 1. Given $\mu > 0, \beta \geq L$ and a fixed instance of t , consider the problem of minimizing f_t over \mathbb{S} (Problem (1)) via the projected gradient descent method (7) with the step size s . If $\mathbf{x}_t^0 \in \mathcal{T}_t(\mu, \beta, s)$, then the projected gradient descent method with $0 < s < \min(\frac{1}{\mu}, \frac{1}{\beta})$ converges linearly to the optimal value f_t^* , i.e.,

$$f_t(\mathbf{x}_t^N) - f_t^* \leq (1 - \mu s)^N [f_t(\mathbf{x}_t^0) - f_t^*], \quad (19)$$

where $N \in \{0, 1, 2, \dots\}$ indicates the number of iterations.

Proof. The proof is similar to that of Theorem 5 in [23]. Let $F_t(\mathbf{x}) := f_t(\mathbf{x}) + \mathbb{I}_{\mathbb{S}}(\mathbf{x})$. By using the Lipschitz continuity of the gradient of f_t , one can write:

$$F_t(\mathbf{x}_t^1) = f_t(\mathbf{x}_t^1) + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^1) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0)$$

$$\leq f_t(\mathbf{x}_t^0) + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) + \langle \nabla f_t(\mathbf{x}_t^0), \mathbf{x}_t^1 - \mathbf{x}_t^0 \rangle + \frac{L}{2} \|\mathbf{x}_t^1 - \mathbf{x}_t^0\|^2 + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^1) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0)$$

Then, noting that $\mathbf{x}_t^0 \in \mathcal{T}_t(\mu, \beta, s) \subset \mathbb{S}$ and $L \leq 1/s$, we obtain an upper bound of the form:

$$\begin{aligned} F_t(\mathbf{x}_t^1) &\leq f_t(\mathbf{x}_t^0) + \langle \nabla f_t(\mathbf{x}_t^0), \mathbf{x}_t^1 - \mathbf{x}_t^0 \rangle + \frac{1}{2s} \|\mathbf{x}_t^1 - \mathbf{x}_t^0\|^2 \\ &\quad + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^1) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) \\ &= f_t(\mathbf{x}_t^0) - \frac{s}{2} \mathcal{D}_t(\mathbf{x}_t^0, 1/s) \end{aligned}$$

where the equality follows from the definition of \mathbf{x}_t^{k+1} and the proximal-gradient. Finally, we upper bound the equation above by using the facts that \mathbf{x}_t^0 satisfies the proximal-PL inequality with parameters μ and β and that $\mathcal{D}_t(\mathbf{x}_t^0, 1/s) \geq \mathcal{D}_t(\mathbf{x}_t^0, \beta)$ since $\frac{1}{s} \geq \beta$ [23]:

$$F_t(\mathbf{x}_t^1) \leq f_t(\mathbf{x}_t^0) - \mu s [f_t(\mathbf{x}_t^0) - f_t^*]$$

Since \mathbf{x}_t^1 is feasible by the definition of projection, we have

$$f_t(\mathbf{x}_t^1) \leq f_t(\mathbf{x}_t^0) - \mu s [f_t(\mathbf{x}_t^0) - f_t^*],$$

which subsequently implies

$$f_t(\mathbf{x}_t^1) - f_t^* \leq (1 - \mu s) [f_t(\mathbf{x}_t^0) - f_t^*]. \quad (20)$$

Furthermore, by showing the decrease in objective value and directly following the definition of the target set, the above results also prove that $\mathbf{x}_t^1 \in \mathcal{T}_t(\mu, \beta, s)$. In other words, the target set is invariant under the projected gradient descent method, as was mentioned in Section II-C. Repeating the process for N steps, we have the final result:

$$f_t(\mathbf{x}_t^N) - f_t^* \leq (1 - \mu s)^N [f_t(\mathbf{x}_t^0) - f_t^*]$$

□

Theorem 1 also gives a lower bound on f_t^* after N iterations:

$$f_t^* \geq \frac{f_t(\mathbf{x}_t^N) - (1 - \mu s)^N f_t(\mathbf{x}_t^0)}{1 - (1 - \mu s)^N}, \quad \forall N \in \mathbb{Z}_+ \quad (21)$$

The next lemma establishes what we will refer to as the *robustness property* of a target set.

Lemma 1. (Robustness of a target set) Assume there exist parameters μ, β , and $s > 0$ such that $\alpha_t > 0$ for all $t \in \{1, \dots, T\}$ (where α_t is as defined in (17)). Then, the target set $\mathcal{T}_t(\mu, \beta, s)$ includes a feasible ball of radius at least r around the global solution for some $r > 0$. That is, $\exists \mu, \beta, s, r > 0$: $\mathcal{T}_t(\mu, \beta, s) \supseteq (\mathcal{B}(\mathbf{x}_t^*, r) \cap \mathbb{S})$ for all $t \in \{1, \dots, T\}$, where $\mathcal{B}(\mathbf{x}_t^*, r) := \{\mathbf{y} \mid \|\mathbf{x}_t^* - \mathbf{y}\| \leq r\}$.

Proof: See Appendix A.

For unconstrained problems, a function satisfying the PL condition implies that it also satisfies the quadratic growth condition [23]. Next, we prove a similar relationship between the proximal-PL inequality and quadratic growth.

Theorem 2. (Quadratic growth) The following inequality holds:

$$\sqrt{f_t(\mathbf{x}) - f_t^*} \geq \sqrt{\frac{\mu}{2}} \|\mathbf{x} - \mathbf{x}_t^*\| \quad \forall \mathbf{x} \in \mathcal{RP}_t^C(\mu, \beta). \quad (22)$$

Proof: See Appendix C.

While the proof of Theorem 2 relies on the continuous version of the projected gradient flow algorithm, this paper does not require implementing or solving this continuous dynamical system. The algorithms in Section III use the discrete-time projected gradient descent algorithm.

E. Visualization of a Proximal-PL Region and Target Set

To develop intuition about proximal-PL regions and target sets, it is beneficial to visualize these sets in an example. Consider the optimization problem

$$\begin{aligned} \min f(x_1, x_2) &= x_1^4 - 4x_1^3 + x_1^2 + 2x_1 + \frac{3}{2} \sin(2\pi x_1) \\ &\quad + x_2^4 - 4x_2^3 + x_2^2 + 2x_2 + \frac{3}{2} \sin(2\pi x_2) + 28.87 \\ \text{s.t.} \quad &-1 \leq x_1 \leq 3, \quad -1 \leq x_2 \leq 3 \end{aligned} \quad (23)$$

which is depicted in Figure 2a. This problem has the optimal value of 0 at $\mathbf{x}^* = (2.75, 2.75)$ and includes many spurious local solutions.

The proximal-PL region and target set for this problem with the parameters $\mu = 0.5$, $\beta = 250 > L$ and $s = \frac{1}{2\beta}$ are depicted in Figure 2b and Figure 2c, respectively. The proximal-PL region includes a neighborhood of the global solution, as well as points far from the global solution. However, many points in the feasible set do not satisfy the proximal-PL inequality, in particular those near local maxima or saddle points. Observe that the target set is a subset of $\mathcal{RA}^D \cap \mathcal{RA}^C$, $\mathcal{RP}^D \cap \mathcal{RP}^C$ and the proximal-PL region. The symmetry in Figure 2b and Figure 2c is a result of the symmetry in the loss function f .


III. ONLINE PROJECTED-GRADIENT DESCENT WITH RANDOM EXPLORATION

In this section, we leverage the results developed in Section II to study the ONO problem (1). We introduce and analyze two algorithms for different scenarios:


- 1) Scenario 1: An initial point in the target region around the global solution \mathbf{x}_1^* is known.
- 2) Scenario 2: No information about the loss functions or their minimizers is known in advance.

A. Scenario 1 - Known desirable initial point


Algorithm 1 provides a natural approach to solving the ONO problem (1) in the setting where a suitable initial point is known. At each time t , the decision maker performs S_t iterations of projected gradient descent on f_t , with the final iteration becoming the decision maker's action at $t + 1$. The assumption is that the decision maker has enough knowledge about the problem at $t = 1$ to select an initial point in the corresponding target set and that the change in the global optimum between time steps is upper-bounded based on parameters reflecting the functions' landscapes. The latter assumption restricts the adversary's choice of loss function and can be regarded as requiring the global solution sequence to have *steadiness*. This assumption is formalized next.



(a) Topology of the objective function f over the feasible set. Observe that this problem has many local minima.



(b) Points in the grey region satisfy the proximal-PL inequality for the function f over the set $[-1, 3] \times [-1, 3]$ with the parameters $\mu = 0.5$ and $\beta = 250$, while those points in the white regions do not. The unique optimal solution $\mathbf{x}^* = (2.75, 2.75)$ is identified by a red star.



(c) Illustration of the target set (yellow) and other sets critical to its definition. The red dashed circle demonstrates the robustness property established in Lemma 1. The length of the black dashed line is the reach of the target set.

Fig. 2: Visualization of the proximal-PL region and the target set for the optimization problem (23)

Assumption 1. (*Steadiness of global solution*) The change in global optimum between consecutive time steps is upper-bounded by $\bar{r} < r$, where r is as defined in Lemma 1. That is, for $t = 1, \dots, T - 1$,

$$\|\mathbf{x}_{t+1}^* - \mathbf{x}_t^*\| \leq \bar{r}. \quad (24)$$

where μ, β, s , and r collectively satisfy the robustness property in Lemma 1 and $\rho_t(\mu, \beta, s)$ is defined in (18). Furthermore, assume that S_t is large enough to satisfy the inequalities:

$$\sqrt{\frac{2M_1 \cdot \rho_t(\mu, \beta, s) \cdot (1 - \mu s)^{S_t}}{\mu}} \leq r - \bar{r} \quad (25a)$$

$$S_t > \frac{\log(\mu) - \log(2M_2)}{\log(1 - \mu s)}. \quad (25b)$$

The constant M_2 defined as

$$M_2 := \inf\{M \in \mathbb{R} : f_t(\mathbf{x}) - f_t^* \leq M \|\mathbf{x} - \mathbf{x}_t^*\|^2 \\ \forall t = 1, \dots, T, \forall \mathbf{x} \in \mathbb{S}\} \quad (26)$$

exists and is finite because the functions f_t are continuous over the bounded set \mathbb{S} .

Note that this assumption only limits the change in the global minimum; the overall landscape of the function can change arbitrarily. Under this assumption, we will establish a deterministic dynamic regret bound for Algorithm 1. To aid in establishing this bound, we first prove two lemmas:

- i) one showing the convergence in terms of the variables \mathbf{x} ,
- ii) another one proving that once the chosen action \mathbf{x}_t is within the target region at time t , all successive actions chosen by the algorithm will also lie within the target region of their respective time.

Lemma 2. Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 1. Under Assumption 1, if $\mathbf{x}_t \in \mathcal{T}_t(\mu, \beta, s)$ for some $t \in \{1, 2, \dots, T - 1\}$, then

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \gamma \|\mathbf{x}_t - \mathbf{x}_t^*\|. \quad (27)$$

where

$$\gamma = \max_{t=1, \dots, T} \sqrt{\frac{2M_2(1 - \mu s)^{S_t}}{\mu}} < 1 \quad (28)$$

Proof: From the convergence rate in Theorem 1 (specifically, equation (20)), we have

$$f_t(\mathbf{x}_{t+1}) - f_t^* \leq (1 - \mu s)^{S_t} [f_t(\mathbf{x}_t) - f_t^*]$$

Applying the quadratic growth inequality from Theorem 2 and taking the square root of all sides, we obtain

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \sqrt{\frac{2(f_t(\mathbf{x}_{t+1}) - f_t^*)}{\mu}} \\ \leq \sqrt{\frac{2(1 - \mu s)^{S_t} (f_t(\mathbf{x}_t) - f_t^*)}{\mu}}$$

Then, using the definition of M_2 , we arrive at

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \underbrace{\sqrt{\frac{2M_2(1 - \mu s)^{S_t}}{\mu}}}_{=\gamma_t} \|\mathbf{x}_t - \mathbf{x}_t^*\| \quad (29)$$

Algorithm 1 Online Projected Gradient Descent with Desirable Initialization

Require: $\mathbf{x}_1 \in \mathcal{T}_1(\mu, \beta, s)$, $0 < s < \min\{\frac{1}{\mu}, \frac{1}{\beta}\}$

- 1: **for** $t = 1, 2, \dots, T$ **do**
 - 2: Play \mathbf{x}_t
 - 3: Set $\mathbf{z}_0 = \mathbf{x}_t$
 - 4: **for** $i = 1, \dots, S_t$ **do**
 - 5: Observe $\nabla f_t(\mathbf{z}_{i-1})$
 - 6: Perform projected gradient descent update:
 $\mathbf{z}_i = \Pi_{\mathbb{S}}[\mathbf{z}_{i-1} - s \nabla f_t(\mathbf{z}_{i-1})]$
 - 7: **end for**
 - 8: Set $\mathbf{x}_{t+1} = \mathbf{z}_{S_t}$
 - 9: **end for**
-

Then $\gamma_t < 1$ since $S_t > \log(\mu/(2M_2))/\log(1 - \mu s)$. \square

The above lemma proves that given a sufficiently large S_t , we can make $\gamma \in [0, 1)$ arbitrarily close to zero, implying that the iterates can become arbitrarily close to the global minimizers at different times. The trade-off is between accuracy and computation time, which is driven by S_t . There is also an intuitive trade-off between the step size s and computation time: smaller step sizes require more algorithmic iterations.

Lemma 3. Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 1. Under Assumption 1, if $\mathbf{x}_t \in \mathcal{T}_t(\mu, \beta, s)$ for any $t \in \{1, 2, \dots, T - 1\}$, then $\mathbf{x}_{t+1} \in \mathcal{T}_{t+1}(\mu, \beta, s)$.

Proof: It is desirable to show that $\|\mathbf{x}_{t+1} - \mathbf{x}_{t+1}^*\| < r$, which ensures that $\mathbf{x}_{t+1} \in \mathcal{B}(\mathbf{x}_{t+1}^*, r)$. By Lemma 1 (the robustness property of target sets), we have $\mathcal{B}(\mathbf{x}_{t+1}^*, r) \subset \mathcal{T}_{t+1}(\mu, \beta, s)$. One can write:

$$\|\mathbf{x}_{t+1} - \mathbf{x}_{t+1}^*\| \leq \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| + \|\mathbf{x}_t^* - \mathbf{x}_{t+1}^*\| \\ \leq \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| + \bar{r} \\ \leq \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| + r - \sqrt{\frac{2M_1 \rho_t(\mu, \beta, s)(1 - \mu s)^{S_t}}{\mu}} \\ \leq \sqrt{\frac{2(f_t(\mathbf{x}_{t+1}) - f_t^*)}{\mu}} + r - \sqrt{\frac{2M_1 \rho_t(\mu, \beta, s)(1 - \mu s)^{S_t}}{\mu}} \\ \leq \sqrt{\frac{2(1 - \mu s)^{S_t} (f_t(\mathbf{x}_t) - f_t^*)}{\mu}} \\ \quad + r - \sqrt{\frac{2M_1 \rho_t(\mu, \beta, s)(1 - \mu s)^{S_t}}{\mu}} \\ \leq r + \sqrt{\frac{2M_1(1 - \mu s)^{S_t}}{\mu}} \left(\sqrt{\|\mathbf{x}_t - \mathbf{x}_t^*\|} - \sqrt{\rho_t(\mu, \beta, s)} \right) \\ \leq r$$

where the second and third inequalities use Assumption 1, the fourth inequality applies Theorem 2, the fifth inequality is due to Theorem 1, the sixth inequality applies the bounded gradient assumption from Section II-A, and the last inequality is due to (18). \square

Now, we present a dynamic regret bound for Algorithm 1.

Corollary 1. Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 1. Under Assumption 1, the dynamic regret satisfies the inequality

$$\text{Reg}_T^d(\mathbf{x}_1, \dots, \mathbf{x}_T) \leq \frac{M_1}{1-\gamma} \sum_{t=2}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| + \frac{\rho_1(\mu, \beta, s)}{(1-\gamma)/M_1} \quad (30)$$

Proof: This proof will follow the same line of reasoning as Theorem 1 and Corollary 1 of [14], where a similar result is proved for strongly convex functions. In the nonconvex setting considered in this paper, we will utilize Lemma 2 and Lemma 3 in our proof.

By the Intermediate Value Theorem, there exists $\mathbf{y} \in \{\mathbf{z} | \mathbf{z} = \omega \mathbf{x}_t + (1-\omega) \mathbf{x}_t^*, 0 \leq \omega \leq 1\}$ such that $f_t(\mathbf{x}) - f_t(\mathbf{x}_t^*) = \nabla f_t(\mathbf{y})^T (\mathbf{x}_t - \mathbf{x}_t^*)$. Therefore, by applying the bounded gradient assumption in Section II-A, we have

$$\text{Reg}_T^d(\mathbf{x}_1, \dots, \mathbf{x}_T) \leq M_1 \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^*\|. \quad (31)$$

Next we establish an upper bound on the summation in (31):

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^*\| &= \|\mathbf{x}_1 - \mathbf{x}_1^*\| + \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_t^*\| \\ &\leq \|\mathbf{x}_1 - \mathbf{x}_1^*\| + \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}^*\| + \sum_{t=2}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| \\ &\leq \|\mathbf{x}_1 - \mathbf{x}_1^*\| - \gamma \|\mathbf{x}_T - \mathbf{x}_T^*\| + \gamma \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^*\| \\ &\quad + \sum_{t=2}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| \\ \implies \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_t^*\| &\leq \frac{\|\mathbf{x}_1 - \mathbf{x}_1^*\| - \gamma \|\mathbf{x}_T - \mathbf{x}_T^*\|}{1-\gamma} \\ &\quad + \frac{1}{1-\gamma} \sum_{t=2}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| \quad (32) \\ &\leq \frac{\rho_1(\mu, \beta, s)}{1-\gamma} + \frac{1}{1-\gamma} \sum_{t=2}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| \quad (33) \end{aligned}$$

The first inequality invokes the triangle inequality. The second inequality applies Lemma 2 for each $t = 1, \dots, T$ and re-indexes the summation. This application of Lemma 2 is derived by recursively applying Lemma 3 to the requirement that $\mathbf{x}_1 \in \mathcal{T}_1(\mu, \beta, s)$. We rearrange terms to arrive at (32) and then apply the definition of the reach of the target set (18) to achieve the final inequality. Combining (33) with (31) completes the proof. \square

Observe that the dynamic regret is a function of the temporal variation in the optimal decision (also called path length or path variation), a common measure of variation discussed in the introduction. The path length is weighted by a function of γ that is large when γ is close to one and approximately one when γ is close to 0. Again, this trade-off between the strength of the dynamic regret bound and computation time is driven by S_t .

Algorithm 2 Online Projected Gradient Descent with Random Exploration

Require: $\mathbf{x}_1 \in \mathbb{S}$, $\mathcal{M}_1 = \emptyset$, $m = 0$, $0 < s < \min\{\frac{1}{\mu}, \frac{1}{\beta}\}$

- 1: **for** $t = 1, 2, \dots, T$ **do**
- 2: • Play \mathbf{x}_t
- 3: • Create $\mathcal{W}_t = \{\mathbf{w}_t^1, \dots, \mathbf{w}_t^q\}$ by uniformly sampling q random points from \mathbb{S}
- 4: • Set $\mathcal{Y}_t = \mathcal{W}_t \cup \mathcal{M}_t \cup \{\mathbf{x}_t\} := \{\mathbf{y}_t^1, \dots, \mathbf{y}_t^{q+m+1}\}$
- 5: **for** $k = 1, 2, \dots, q+m+1$ **do**
- 6: • Initialize $\mathbf{z}_0^k = \mathbf{y}_t^k$, $\mathbf{z}_*^k = \mathbf{y}_t^k$, $\underline{c}_k = \infty$, $\bar{b}_k = -\infty$
- 7: • Set $i = 1$
- 8: **while** $\underline{c}_k - \bar{b}_k > \epsilon$ or $i \leq S_t$ **do**
- 9: • Observe $\nabla f_t(\mathbf{z}_{i-1}^k)$
- 10: • Compute $\mathbf{z}_i^k = \Pi_{\mathbb{S}}[\mathbf{z}_{i-1}^k - s \nabla f_t(\mathbf{z}_{i-1}^k)]$
- 11: • Observe $c_i^k = f_t(\mathbf{z}_i^k)$
- 12: **if** $c_i^k < \underline{c}_k$ **then**
- 13: • $\mathbf{z}_*^k = \mathbf{z}_i^k$, $\underline{c}_k = c_i^k$
- 14: **end if**
- 15: • $b_i^k = (f_t(\mathbf{z}_i^k) - (1-\mu s)^i f_t(\mathbf{z}_0^k)) / (1 - (1-\mu s)^i)$
- 16: • Update $\bar{b}_k = \max\{\bar{b}_k, b_i^k\}$
- 17: • Update $i = i + 1$
- 18: **end while**
- 19: • Return $I_t^k = i$
- 20: **end for**
- 21: • Let $K = \text{argmin}_k \underline{c}_k$, and set $\mathbf{x}_{t+1} = \mathbf{z}_*^K$
- 22: • Store in memory all other points in $\{\mathbf{z}_*^k\}_{k=1}^{q+m+1}$ which could be in the proximal PL-region at time t :
 $\mathcal{M}_{t+1} = \{\mathbf{z}_*^k : \underline{c}_k \leq \underline{c}_K + \epsilon, k \in \{1, \dots, q+m+1\} \setminus K\}$
- 23: $m = |\mathcal{M}_{t+1}|$
- 24: **end for**

B. Scenario 2 - Blind initialization

The initialization scenario described in Scenario 1 – that a point in the target region is known at the initial time – is difficult to satisfy in practice. The reason is that the decision maker may have no information about how their adversary will design f_1 . In this case, it is advantageous to *explore* the landscape of f_t before selecting decision \mathbf{x}_{t+1} . Algorithm 2 explores by running the projected gradient descent algorithm from multiple initial points, which are sampled uniformly at random from \mathbb{S} .

The goal of exploration is to find a point in a time-varying target set. The decision maker cannot verify when this occurs, however, since they do not have knowledge of the landscape of the function. As a result, Algorithm 2 utilizes *memory* to make available at time $t+1$ points which may be in the target set at time t . Once a point in a time-varying target set is sampled, memory ensures that the decision maker has at least one initial point in the target sets for each future time step. This tracking guarantee is formalized in the following lemma.

Lemma 4. Consider sequences $\{\mathbf{x}_t\}_{t=1}^T$ and $\{\mathcal{Y}_t\}_{t=1}^T$ generated by Algorithm 2. Under Assumption 1, if $\mathbf{z}_0^k \in (\mathcal{T}_t(\mu, \beta, s) \cap \mathcal{Y}_t)$ for any $t \in \{1, 2, \dots, T-1\}$, then $\mathbf{z}_*^k \in (\mathcal{T}_{t+1}(\mu, \beta, s) \cap \mathcal{Y}_{t+1})$.

Proof: The number of iterations I_t^k is at least as large as S_t .

Therefore, applying the same logic as the proof of Lemma 3, we know that $\mathbf{z}_{I_t^k}^k \in \mathcal{T}_{t+1}(\mu, \beta, s)$. By Theorem 1, we have $\mathbf{z}_0^k \geq \mathbf{z}_1^k \geq \dots \geq \mathbf{z}_{S_t}^k \geq \mathbf{z}_{I_t^k}^k$ with $\mathbf{z}_i^k = \mathbf{z}_{i+1}^k$ only if $\mathbf{z}_i^k = \mathbf{x}_t^*$. This implies that $\mathbf{z}_*^k = \mathbf{z}_{I_t^k}^k$. It remains to show that $\mathbf{z}_*^k \in \mathcal{Y}_{t+1}$. If $\mathbf{z}_*^k = \mathbf{x}_{t+1}$, then $\mathbf{z}_*^k \in \mathcal{Y}_{t+1}$. Otherwise, since $\mathbf{z}_0^k \in \mathcal{T}_t(\mu, \beta, s)$, it holds that $\underline{c}_k \leq f_t^* + \epsilon \leq \underline{c}_K + \epsilon$, which implies $\mathbf{z}_*^k \in \mathcal{M}_{t+1} \subset \mathcal{Y}_{t+1}$. As a result, $\mathbf{z}_*^k \in \mathcal{Y}_{t+1}$, which completes the proof. \square

Since Algorithm 1 is a deterministic algorithm, the dynamic regret bound established in Corollary 1 is deterministic too. Algorithm 2 relies on sampling, and therefore its associated regret bound should be probabilistic. In the following culminating theorem, we provide an upper bound on the dynamic regret accrued using Algorithm 2 and a lower bound on the probability with which this bound holds.

Theorem 3. Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 2. Under Assumption 1, the dynamic regret satisfies the following probabilistic bound for all $\bar{T} \in \{1, \dots, T\}$:

$$\begin{aligned} \mathbb{P} \left[\mathbf{Reg}_T^d(\mathbf{x}_1, \dots, \mathbf{x}_T) \leq \mathbf{Reg}_{\bar{T}-1}^d(\mathbf{x}_1, \dots, \mathbf{x}_{\bar{T}-1}) \right. \\ \left. + \frac{M_1 \rho_{\bar{T}}(\mu, \beta, s)}{(1-\gamma)} + \frac{M_1}{1-\gamma} \sum_{t=\bar{T}+1}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| \right] \quad (34) \\ \geq 1 - \prod_{t=1}^{\bar{T}} \left(1 - \frac{\text{Vol}(\mathcal{T}_t(\mu, \beta, s))}{\text{Vol}(\mathbb{S})} \right)^q, \end{aligned}$$

where $\text{Vol}(\cdot)$ indicates the volume of the set and γ is defined in Lemma 2. This theorem relates the dynamic regret at time T to the dynamic regret at an earlier time \bar{T} , the variation within the optimal decision sequence after \bar{T} , and the relative sizes of the target sets through \bar{T} .

Proof: The probability that a point located in the time-varying target set $\mathcal{T}_t(\mu, \beta, s)$ appears in \mathcal{Y}_t by time \bar{T} is related to the volumes of $\mathcal{T}_t(\mu, \beta, s)$ and \mathbb{S} because, at each time step, q initial points are selected from \mathbb{S} uniformly at random. Hence,

$$\begin{aligned} \mathbb{P} [\mathcal{Y}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset \text{ for some } t \in \{1, \dots, \bar{T}\}] \quad (35) \\ \geq \mathbb{P} [\mathcal{W}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset \text{ for some } t \in \{1, \dots, \bar{T}\}] \\ = \mathbb{P} \left[\bigcup_{t=1}^{\bar{T}} \bigcup_{i=1}^q \mathbf{w}_t^i \in \mathcal{T}_t(\mu, \beta, s) \right] \\ = 1 - \mathbb{P} [\mathbf{w}_t^i \notin \mathcal{T}_t(\mu, \beta, s) \forall t=1, \dots, \bar{T}, \forall i=1, \dots, q] \\ = 1 - \prod_{t=1}^{\bar{T}} \left(1 - \frac{\text{Vol}(\mathcal{T}_t(\mu, \beta, s))}{\text{Vol}(\mathbb{S})} \right)^q \quad (36) \end{aligned}$$

Now, we will show that if $\mathcal{Y}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset$ for some $t \in \{1, \dots, \bar{T}\}$, then the dynamic regret is upper bounded by the expression in (34). Applying Lemma 4, Corollary 1 and Definition 4 yields that

$$\begin{aligned} \mathbf{Reg}_T^d(\mathbf{x}_1, \dots, \mathbf{x}_T) &= \sum_{t=1}^{\bar{T}-1} (f_t(\mathbf{x}_t) - f_t^*) + \sum_{t=\bar{T}}^T (f_t(\mathbf{x}_t) - f_t^*) \\ &= \mathbf{Reg}_{\bar{T}-1}^d(\mathbf{x}_1, \dots, \mathbf{x}_{\bar{T}-1}) + \sum_{t=\bar{T}}^T (f_t(\mathbf{x}_t) - f_t^*) \end{aligned}$$

$$\begin{aligned} \leq \mathbf{Reg}_{\bar{T}-1}^d(\mathbf{x}_1, \dots, \mathbf{x}_{\bar{T}-1}) + M_1 \frac{\|\mathbf{x}_{\bar{T}} - \mathbf{x}_{\bar{T}}^*\| - \gamma \|\mathbf{x}_T - \mathbf{x}_T^*\|}{(1-\gamma)} \\ + \frac{M_1}{1-\gamma} \sum_{t=\bar{T}+1}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| \\ \leq \mathbf{Reg}_{\bar{T}-1}^d(\mathbf{x}_1, \dots, \mathbf{x}_{\bar{T}-1}) + \frac{M_1 \rho_{\bar{T}}(\mu, \beta, s)}{(1-\gamma)} \\ + \frac{M_1}{1-\gamma} \sum_{t=\bar{T}+1}^T \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\|. \quad (37) \end{aligned}$$

This completes the proof. \square

Observe that the strength of this probabilistic bound depends on the landscape of loss functions around the global solution through the volume of the target sets. In particular, one can analyze the role that a ‘‘lower-complexity problem’’ at some time \bar{T} plays in determining the complexity of the entire online nonconvex optimization. As an extreme but important case, suppose that there is a time $\bar{T} \in \{1, \dots, T\}$ such that $f_{\bar{T}}$ is convex. Then the dynamic regret bound (34) holds with probability 1 since $\mathcal{T}_t(\mu, \beta, s) = \mathbb{S}$. In other words, the existence of a single convex problem, in between the sequence of nonconvex problems, is enough to break down the NP-hardness of solving nonconvex problems for all future times, under the steadiness of the global solution assumption. On the other hand, if the global solution is extremely ‘‘sharp’’ at all times, it is unrealistic to expect any algorithm with limited computation time to find the global solution. Thus, dynamic regret could be arbitrarily large in this case. Indeed, the target set of a sharp minima is small and therefore the probability of satisfying the dynamic regret bound in (34) is low, as expected.

Choices for the step size s , number of iterations S_t , and number of samples q , represent trade-offs between regret bound strength and computation time. As discussed in Section III-A, a smaller step size requires more algorithmic iterations to satisfy Assumption 1. Increasing the number of iterations may increase the time to execute the while loop (line 8). However, larger values of S_t improve the upper bound on dynamic regret in (34) by reducing γ . Increasing the number of random initial points improves the probability with which (34) holds, but also increases computation time.

C. Empirical study of Algorithm 2

The objective of this section is to support the results of Section III-B through numerical analysis. We will illustrate the performance of Algorithm 2 on online function sequences which satisfy the assumptions in Section II-A and Assumption 1 (steadiness of the global solution). To demonstrate the role that a single comparatively low-complexity problem can play in a sequence of nonconvex problems, we will consider two cases:

- A) ‘‘No low-complexity problem’’: In this case, $\{f_t : \mathbb{R}^2 \rightarrow \mathbb{R}\}_{t=1}^{40}$ each have many local minima over $\mathbb{S} = [-1, 3]$ and the target sets’ volumes represent between 2.47% and 4.14% of the feasible space. The geometry of f_1, \dots, f_6 , which are representative of the entire sequence, is shown in Figure 3.

TABLE I: Parameter and constant values

S	μ	β	s	ϵ	r
$[-1, 3]^2$	0.5	289	0.0031	0.1	0.29
$S_t(max)$	L	M_1	$\ \mathbf{x}_t^* - \mathbf{x}_{t-1}^*\ $		
7060	289	140	0.22		

B) “Lower-complexity problem at time 4”: In this case, $\{f_t : \mathbb{R}^2 \rightarrow \mathbb{R}\}_{t=1}^{40}$ is identical to Case A at every time period except $t = 4$. The target set corresponding to f_4 covers 20.7% of the feasible space. Meanwhile, \mathbf{x}_4^* is the same in both scenarios.

The parameter choices and key problem constants for these two online optimization problems are summarized in Table I.





Fig. 3: Contour plots of f_1, \dots, f_6 for Case A. The red star marks the unique global minimum of each function.


We conducted 500 trials of Algorithm 2 on Case A and Case B for 3 different sampling rates: $q = 1$, $q = 2$, and $q = 5$. Figure 4 plots the empirical probability that $\mathcal{Y}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset$ versus the theoretical lower bound provided in Theorem 3. (Note that, by Lemma 4, this is the same as the probability that $\mathcal{Y}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset$ for some $\bar{t} \in \{1, \dots, t\}$.) For the same value of q , the two cases are identical for $t = 1, 2, 3$ and diverge at $t = 4$ as a result of the “easy” problem in Case B. These results support Theorem 3. A gap between the lower bound and observed likelihood of initializing in the target region is expected, since the lower bound does not account for the possibility that \mathbf{x}_t or a memory point may be in the subsequent target set. The dynamic regret and optimality gap over time is shown in Figure 5. Regret accumulates quickly until the target set is found and then accumulates slowly as Algorithm 2 starts tracking the global solution.

IV. CONCLUSION

In this paper, we defined proximal-PL regions and target sets, characterized their properties, and used this new knowledge to propose and analyze algorithms for online nonconvex optimization problems. Linear convergence to the global minimizer and quadratic growth are the two key properties of the target sets that we established. Since dynamic regret can




(a) Case A




(b) Case B

Fig. 4: Empirical validation of Theorem 3 probability bound



(a) Case A



(b) Case B

Fig. 5: Regret resulting from Algorithm 2

be arbitrarily large when there are no restrictions on the loss functions, we constrain consecutive functions to have global solutions which are not too far apart, but do not limit the variation in the loss functions otherwise. In this setting, we propose two online algorithms. Algorithm 1 is relevant when the decision maker has a good initial point, and it provides a deterministic dynamic regret upper bound as a function of the path length of the optimal decision. Algorithm 2 utilizes exploration and memory to be relevant regardless of the initial point. It provides a probabilistic dynamic regret upper bound, which is also a function of the path length of the optimal decision. The strength of this probabilistic bound depends on the loss function landscapes. For example, the bound holds with probability 1 in the special case where one of the loss functions in the sequence is convex. Empirical studies support these bounds.

REFERENCES

- [1] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, “Visualizing the loss landscape of neural nets,” *Advances in Neural Information Processing Systems*, pp. 6389–6399, 2018.
- [2] E. J. Candes and B. Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
- [3] E. J. Candes and Y. Plan, “Matrix completion with noise,” *Proc. of the IEEE*, vol. 98, no. 6, pp. 925–936, 2010.
- [4] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, “Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations,” *IEEE Trans. on Automatic Control*, vol. 48, no. 2, 2003.
- [5] F. Zohrizadeh, C. Jozs, M. Jin, R. Madani, J. Lavaei, and S. Sojoudi, “Conic relaxations of power system optimization: Theory and algorithms,” *to appear in European Jnl. of Operational Research*, 2020.
- [6] M. S. Asif and J. Romberg, “Sparse recovery of streaming signals using 11-homotopy,” *IEEE Trans. on Signal Processing*, vol. 62, no. 16, pp. 4209–4223, 2014.
- [7] Y. Tang, K. Dvijotham, and S. Low, “Real-time optimal power flow,” *IEEE Trans. on Smart Grid*, vol. 8, no. 6, pp. 2963–2973, 2017.
- [8] E. Hazan, *Introduction to Online Convex Optimization*. Foundations and Trends in Optimization, 2015, vol. 2, no. 4.
- [9] E. C. Hall and R. M. Willett, “Online convex optimization in dynamic environments,” *IEEE Jnl. of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 647–662, 2015.
- [10] O. Besbes, Y. Gur, and A. Zeevi, “Non-stationary stochastic optimization,” *Operations Research*, vol. 63, no. 5, pp. 1227–1244, 2015.
- [11] A. Jadbabaie, S. S. Alexander Rakhlin, and K. Sridharan, “Online optimization: Competing with dynamic comparators,” *Artificial Intelligence and Statistics, Proc. of Machine Learning Research*, vol. 38, pp. 398–406, 2015.
- [12] M. Zinkevich, “Online convex programming and generalized infinitesimal gradient ascent,” *Proc. of the Int. Conf. on Machine Learning*, pp. 928–936, 2003.
- [13] T. Yang, L. Zhang, R. Jin, and J. Yi, “Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient,” in *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- [14] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, “Online optimization in dynamic environments: Improved regret rates for strongly convex problems,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 7195–7201.
- [15] L. Yang, L. Deng, M. H. Hajiesmaili, C. Tan, and W. S. Wong, “An optimal algorithm for online non-convex learning,” *Proc. of the ACM on Measurement and Analysis of Computing Syst.*, vol. 25, 2018.
- [16] O.-A. Maillard and R. Munos, “Online learning in adversarial lipschitz environments,” in *Joint european conference on machine learning and knowledge discovery in databases*. Springer, 2010, pp. 305–320.
- [17] E. Hazan and S. Kale, “Online submodular minimization,” *Journal of Machine Learning Research*, vol. 13, no. 10, 2012.
- [18] L. Zhang, T. Yang, R. Jin, and Z.-H. Zhou, “Online bandit learning for a special class of non-convex losses,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [19] X. Gao, X. Li, and S. Zhang, “Online learning with non-convex losses and non-stationary regret,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2018, pp. 235–243.
- [20] S. Park, J. Mulvaney-Kemp, M. Jin, and J. Lavaei, “Diminishing regret for online nonconvex optimization,” in *2021 American Control Conference (ACC)*, 2021, pp. 978–985.
- [21] Y. Ding, J. Lavaei, and M. Arcaç, “Escaping spurious local minimum trajectories in online time-varying nonconvex optimization,” *American Control Conference*, 2021.
- [22] S. Fattahi, C. Jozs, R. Mohammadi, J. Lavaei, and S. Sojoudi, “Absence of spurious local trajectories in time-varying optimization: A control-theoretic perspective,” 2020, *IEEE Conference on Control Technology and Applications (CCTA)*.
- [23] H. Karimi, J. Nutini, and M. Schmidt, “Linear convergence of gradient and proximal-gradient methods under the polyak-lojasiewicz condition,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2016, pp. 795–811.
- [24] B.T.Polyak, “Gradient methods for minimizing functionals,” *Zh. Vychisl. Mat. Mat. Fiz.*, vol. 3, no. 4, pp. 643–653, 1963.
- [25] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, “Global convergence of policy gradient methods for the linear quadratic regulator,” in *International Conference on Machine Learning*, 2018.
- [26] A. Hauswirth, S. Bolognani, and F. Dörfler, “Projected dynamical systems on irregular, non-euclidean domains for nonlinear optimization,” *SIAM Journal on Control and Optimization*, vol. 59, no. 1, pp. 635–668, 2021.
- [27] A. Nagurney and D. Zhang, *Projected dynamical systems and variational inequalities with applications*. Springer Science & Business Media, 2012, vol. 2.
- [28] J.-B. Hiriart-Urruty and C. Lemaréchal, *Convex analysis and minimization algorithms I: Fundamentals*. Springer science & business media, 2013, vol. 305.

APPENDIX

A. Proof of Lemma 1

Lemma 5 (see Appendix B) establishes the existence of an open set $\mathcal{S}_t^D \subset \mathcal{R}\mathcal{A}_t^D(s)$ such that $\mathbf{x}_t^* \in \mathcal{S}_t^D$. Then, by definition of an open set, $\mathcal{S}_t^D \supseteq (\mathcal{B}(\mathbf{x}_t^*, r_1) \cap \mathbb{S})$ for some $r_1 > 0$. It can be concluded from Proposition 8.5 and Lemma 8.3 in [26] that the projected gradient flow system described in (13) converges to the set of critical points of (1) and the sublevel sets of f_t are invariant under this system. Define \bar{l}_t as the second-lowest objective value among all critical points of (1). (f_t^* is the lowest objective value among all critical points.) Then \mathbf{x}_t^* is the only critical point in the sublevel set $\mathcal{L}_t(\bar{l})$, which implies $\mathcal{L}_t(\bar{l}) \subseteq \mathcal{R}\mathcal{A}_t^C$. By continuity of f_t and the unique global optimality of \mathbf{x}_t^* , we have $\mathcal{L}_t(\bar{l}) \supseteq (\mathcal{B}(\mathbf{x}_t^*, r_2) \cap \mathbb{S})$ for some $r_2 > 0$. Similarly, $\mathcal{L}_t(\alpha_t) \supseteq (\mathcal{B}(\mathbf{x}_t^*, r_3) \cap \mathbb{S})$ for some $r_3 > 0$. Therefore,

$$\begin{aligned} \mathcal{T}_t(\mu, \beta, s) &= \left(\mathcal{L}_t(\alpha_t) \cap \mathcal{R}\mathcal{A}_t^D(s) \cap \mathcal{R}\mathcal{A}_t^C \right) \\ &\supseteq \left(\mathcal{B}(\mathbf{x}_t^*, r_1) \cap \mathcal{B}(\mathbf{x}_t^*, r_2) \cap \mathcal{B}(\mathbf{x}_t^*, r_3) \cap \mathbb{S} \right) \\ &= \left(\mathcal{B}(\mathbf{x}_t^*, r) \cap \mathbb{S} \right) \end{aligned}$$

where $r = \min\{r_1, r_2, r_3\} > 0$. This completes the proof. \square

B. Capture property

Lemma 5. *Let f be a continuously differentiable function on a compact, convex set \mathbb{S} . Let $\{\mathbf{x}^k\}$ be a sequence of points in \mathbb{S} satisfying $f(\mathbf{x}^{k+1}) \leq f(\mathbf{x}^k)$ generated by the projected gradient descent method $\mathbf{x}^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}^k - s\nabla f(\mathbf{x}^k))$, which is convergent in the sense that every limit point of such sequences is a stationary point of $\min_{\mathbf{x} \in \mathbb{S}} f(\mathbf{x})$. Let \mathbf{x}^* be a local minimum of $\min_{\mathbf{x} \in \mathbb{S}} f(\mathbf{x})$, which is the only stationary*

point within some open set. Then there exists an open set \mathbb{B} containing \mathbf{x}^* such that if $\mathbf{x}^k \in \mathbb{B}$ for some $k \geq 0$, then $\mathbf{x}^k \in \mathbb{B}$ for all $k \geq \bar{k}$ and $\{\mathbf{x}^k\} \rightarrow \mathbf{x}^*$.

Proof: Let $\rho > 0$ be a constant such that

$$f(\mathbf{x}^*) < f(\mathbf{x}), \quad \forall \mathbf{x} \neq \mathbf{x}^* \text{ with } \|\mathbf{x} - \mathbf{x}^*\| \leq \rho.$$

For every $\delta \in [0, \rho]$, define

$$\phi(\delta) = \min_{\{\mathbf{x} \mid \delta \leq \|\mathbf{x} - \mathbf{x}^*\| \leq \rho\}} f(\mathbf{x}) - f(\mathbf{x}^*).$$

Note that $\phi(\delta)$ is a monotonically non-decreasing function of δ , and that $\phi(\delta) > 0$ for all $\delta \in (0, \rho]$. Given any $\epsilon \in (0, \rho]$, let $r \in (0, \epsilon]$ be such that

$$\|\mathbf{x} - \mathbf{x}^*\| < r \Rightarrow \|\mathbf{x} - \mathbf{x}^*\| + \frac{1}{\beta} \|\nabla f(\mathbf{x})\| < \epsilon.$$

Consider the open set

$$\mathbb{B} = \{\mathbf{x} \in \mathbb{S} \mid \|\mathbf{x} - \mathbf{x}^*\| < \epsilon, f(\mathbf{x}) < f(\mathbf{x}^*) + \phi(r)\}.$$

We claim that if $\mathbf{x}^k \in \mathbb{B}$ for some k , then $\mathbf{x}^{k+1} \in \mathbb{B}$. In order to prove the claim, assume that $\mathbf{x}^k \in \mathbb{B}$. Then,

$$\phi(\|\mathbf{x}^k - \mathbf{x}^*\|) \leq f(\mathbf{x}^k) - f(\mathbf{x}^*) < \phi(r),$$

where the first inequality is due to $\phi(\|\mathbf{x}^k - \mathbf{x}^*\|) = \min_{\{\mathbf{x} \mid \|\mathbf{x}^k - \mathbf{x}^*\| \leq \|\mathbf{x} - \mathbf{x}^*\| \leq \rho\}} f(\mathbf{x}) - f(\mathbf{x}^*) \leq f(\mathbf{x}^k) - f(\mathbf{x}^*)$ and the second inequality is due to the fact that $\mathbf{x}^k \in \mathbb{B}$. Since $\phi(\cdot)$ is monotonically non-decreasing, the above statement implies that $\|\mathbf{x}^k - \mathbf{x}^*\| < r$, which means that

$$\|\mathbf{x}^k - \mathbf{x}^*\| + \frac{1}{\beta} \|\nabla f(\mathbf{x}^k)\| < \epsilon$$

We also know that

$$\begin{aligned} \|\mathbf{x}^{k+1} - \mathbf{x}^*\| &= \|(\mathbf{x}^{k+1} - \mathbf{x}^k) + (\mathbf{x}^k - \mathbf{x}^*)\| \\ &\leq \|\mathbf{x}^{k+1} - \mathbf{x}^k\| + \|\mathbf{x}^k - \mathbf{x}^*\| \\ &= \|\Pi_{\mathbb{S}}(\mathbf{x}^k - \frac{1}{\beta} \nabla f(\mathbf{x}^k)) - \Pi_{\mathbb{S}}(\mathbf{x}^k)\| + \|\mathbf{x}^k - \mathbf{x}^*\| \\ &\leq \|(\mathbf{x}^k - \frac{1}{\beta} \nabla f(\mathbf{x}^k)) - \mathbf{x}^k\| + \|\mathbf{x}^k - \mathbf{x}^*\| \quad (39) \\ &= \frac{1}{\beta} \|\nabla f(\mathbf{x}^k)\| + \|\mathbf{x}^k - \mathbf{x}^*\| < \epsilon \end{aligned}$$

where equation (39) follows from the non-expansive property of the projection operator (when projected onto convex sets) and the final inequality follows from applying equation (B). Therefore by induction, this implies that if $\mathbf{x}^k \in \mathbb{B}$ for some k , we have $\mathbf{x}^k \in \mathbb{B}$ for all $k \geq \bar{k}$. Let $\bar{\mathbb{B}}$ be the closure of \mathbb{B} . Since $\bar{\mathbb{B}}$ is compact, the sequence $\{\mathbf{x}^k\}$ will have at least one limit point, which by assumption must be a stationary point of $\min_{\mathbf{x} \in \mathbb{S}} f(\mathbf{x})$. The only stationary point of $\min_{\mathbf{x} \in \mathbb{S}} f(\mathbf{x})$ within $\bar{\mathbb{B}}$ is \mathbf{x}^* since $\|\mathbf{x} - \mathbf{x}^*\| < \epsilon \leq \rho$ for all $\mathbf{x} \in \bar{\mathbb{B}}$. Hence, $\mathbf{x}^k \rightarrow \mathbf{x}^*$. \square

C. Proof of Theorem 2

Take the function f to be any f_t , $t \in \{1, \dots, T\}$. Define the function $g(\mathbf{x}) := \sqrt{f(\mathbf{x}) - f^*}$ and

$$\dot{\mathbf{x}}(\ell) = \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right), \quad \forall \ell \geq 0 \quad (40)$$

with $\mathbf{x}(0) = \mathbf{x}$. Then, by the fundamental theorem of calculus, we have

$$\begin{aligned} \sqrt{f(\mathbf{x}) - f^*} &= g(\mathbf{x}) - g(\mathbf{x}^*) \\ &= - \int_0^\infty \frac{d}{d\ell} g(\mathbf{x}(\ell)) d\ell = - \int_0^\infty \frac{\nabla f(\mathbf{x}(\ell))}{2g(\mathbf{x}(\ell))} \cdot \dot{\mathbf{x}}(\ell) d\ell \quad (41) \end{aligned}$$

The following lemma will be used to establish a lower bound on the term inside the integral.

Lemma 6. Consider the projected gradient flow (40) with $\mathbf{x}(0) \in \mathbb{S}$. There exists a unique solution $\mathbf{x}(\ell)$ to this projected dynamical system, and

$$\langle \nabla f(\mathbf{x}(\ell)), \dot{\mathbf{x}}(\ell) \rangle = \frac{-1}{2\beta} \cdot \lim_{h \rightarrow \infty} \mathcal{D}(\mathbf{x}(\ell), h) - \frac{\beta \|\dot{\mathbf{x}}(\ell)\|^2}{2}. \quad (42)$$

Proof: The existence and uniqueness of the solution of the projected dynamical system (40) is guaranteed by [27, Thm. 2.5] under the assumptions in Section II-A. Let $\{\mathbf{x}(\ell)\}_{\ell \geq 0}$ denote the unique solution to (40), and

$$\begin{aligned} \bar{\mathbf{x}}_\epsilon(\ell) &:= \Pi_{\mathbb{S}} \left(\mathbf{x}(\ell) - \frac{\epsilon}{\beta} \nabla f(\mathbf{x}(\ell)) \right) \\ &= \operatorname{argmin}_{\mathbf{y} \in \mathbb{S}} \left[\langle \mathbf{y} - \mathbf{x}(\ell), \nabla f(\mathbf{x}(\ell)) \rangle \epsilon + \frac{\beta \|\mathbf{y} - \mathbf{x}(\ell)\|^2}{2} \right] \end{aligned}$$

Then, it follows from [28, Sec III Prop. 5.3.5] that

$$\lim_{\epsilon \downarrow 0} \frac{\bar{\mathbf{x}}_\epsilon(\ell) - \mathbf{x}(\ell)}{\epsilon} = \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right),$$

where $\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))$ is the tangent cone of \mathbb{S} at $\mathbf{x}(\ell) \in \mathbb{S}$. By the definition of the proximal-gradient,

$$\begin{aligned} \lim_{h \rightarrow \infty} \mathcal{D}(\mathbf{x}(\ell), h) &= \lim_{\epsilon \downarrow 0} \mathcal{D}(\mathbf{x}(\ell), \beta/\epsilon) \\ &= \lim_{\epsilon \downarrow 0} \frac{-2\beta}{\epsilon} \cdot \min_{\mathbf{y} \in \mathbb{S}} \left[\langle \nabla f(\mathbf{x}(\ell)), \mathbf{y} - \mathbf{x}(\ell) \rangle + \frac{\beta}{2\epsilon} \|\mathbf{y} - \mathbf{x}(\ell)\|^2 \right] \\ &= \lim_{\epsilon \downarrow 0} \frac{-2\beta}{\epsilon^2} \cdot \min_{\mathbf{y} \in \mathbb{S}} \left[\langle \nabla f(\mathbf{x}(\ell)), \mathbf{y} - \mathbf{x}(\ell) \rangle \epsilon + \frac{\beta}{2} \|\mathbf{y} - \mathbf{x}(\ell)\|^2 \right] \\ &= \lim_{\epsilon \downarrow 0} \frac{-2\beta}{\epsilon^2} \left[\langle \nabla f(\mathbf{x}(\ell)), \bar{\mathbf{x}}_\epsilon(\ell) - \mathbf{x}(\ell) \rangle \epsilon + \frac{\beta}{2} \|\bar{\mathbf{x}}_\epsilon(\ell) - \mathbf{x}(\ell)\|^2 \right] \\ &= -2\beta \left[\langle \nabla f(\mathbf{x}(\ell)), \dot{\mathbf{x}}(\ell) \rangle + \frac{\beta}{2} \|\dot{\mathbf{x}}(\ell)\|^2 \right] \end{aligned}$$

where the last equation is due to the continuity of $\|\cdot\|^2$. Rearranging the above equation yields the desired result. \square

Returning to the proof of Theorem 2, next we establish a lower bound on the term inside the integral in (41).

$$\begin{aligned} &-\frac{\nabla f(\mathbf{x}(\ell))}{2g(\mathbf{x}(\ell))} \cdot \dot{\mathbf{x}}(\ell) \\ &= -\frac{1}{2g(\mathbf{x}(\ell))} \left\langle \nabla f(\mathbf{x}(\ell)), \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right) \right\rangle \\ &= \frac{1}{2g(\mathbf{x}(\ell))} \left(\frac{1}{2\beta} \lim_{h \rightarrow \infty} \mathcal{D}(\mathbf{x}(\ell), h) \right) \end{aligned}$$

$$\begin{aligned}
& + \frac{\beta}{2} \left\| \Pi_{\mathbb{T}_s(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right) \right\|^2 \\
& \geq \frac{1}{2g(\mathbf{x}(\ell))} \left(\frac{\mathcal{D}(\mathbf{x}(\ell), \beta)}{2\beta} + \frac{\beta}{2} \left\| \Pi_{\mathbb{T}_s(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right) \right\|^2 \right) \\
& \geq \frac{1}{2g(\mathbf{x}(\ell))} \left(\frac{\mu}{\beta} g(\mathbf{x}(\ell))^2 + \frac{\beta}{2} \left\| \Pi_{\mathbb{T}_s(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right) \right\|^2 \right) \\
& \geq \sqrt{\frac{\mu}{2}} \left\| \Pi_{\mathbb{T}_s(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right) \right\|
\end{aligned}$$

The first equality follows from the definition of the gradient flow system and the second equality is due to Lemma 6. The first inequality holds because of Lemma 1 of [23]. The second inequality applies the fact that $\mathbf{x}(t)$ satisfies the proximal-PL inequality with the parameters μ and β . The third inequality is the result of the arithmetic-geometric mean inequality. Finally, substituting this lower bound into (41) gives

$$\begin{aligned}
\sqrt{f(\mathbf{x}) - f^*} & \geq \sqrt{\frac{\mu}{2}} \left\| \int_0^\infty \Pi_{\mathbb{T}_s(\mathbf{x}(\ell))} \left(-\frac{1}{\beta} \nabla f(\mathbf{x}(\ell)) \right) d\ell \right\| \\
& = \sqrt{\frac{\mu}{2}} \left\| \int_0^\infty \dot{\mathbf{x}}(\ell) d\ell \right\| = \sqrt{\frac{\mu}{2}} \|\mathbf{x} - \mathbf{x}^*\|
\end{aligned}$$

This completes the proof. \square



Julie Mulvaney-Kemp is a Ph.D. candidate at the University of California, Berkeley in Industrial Engineering and Operations Research. She holds a B.S. in Systems Engineering from Washington University in St. Louis and an M.S. in Industrial Engineering and Operations Research from the University of California, Berkeley. In 2020 she received a “best of the best” conference paper award at the IEEE Power & Energy Society General Meeting and the Best Student Paper Award from the INFORMS Energy, Natural Resources and the Environment Section.



SangWoo Park received the B.S. degree in Environmental Engineering from Johns Hopkins University, Baltimore, MD, in 2016 and the M.S. degree in Industrial Engineering and Operations Research from the University of California, Berkeley, CA, in 2017. He is currently a Ph.D. candidate in the department of Industrial Engineering and Operations Research from the University of California, Berkeley, CA. He is the recipient of the Best Student Paper Award of 2020 American Control Conference.



Ming Jin is an assistant professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech. He was a postdoctoral researcher in the Department of Industrial Engineering and Operations Research at University of California, Berkeley. He received his doctoral degree from EECS department at University of California, Berkeley in 2017. He is the recipient of the Siebel scholarship, 2018 Best Paper Award of Building and Environment, and the 2016 Best Paper Award at the International Conference on Mobile Ubiquitous Computing, Systems,

Services and Technologies.



Javad Lavaei received the Ph.D. degree in control and dynamical systems from the California Institute of Technology, Pasadena, CA, in 2011. He is an Assistant Professor in the Department of Industrial Engineering and Operations Research, University of California, Berkeley, Berkeley, CA. Dr. Lavaei received multiple awards, including the NSF CAREER Award, Office of Naval Research Young Investigator Award, and the Donald P. Eckman Award.