

Improved Analysis for Dynamic Regret of Strongly Convex and Smooth Functions

Peng Zhao
Lijun Zhang

*National Key Laboratory for Novel Software Technology
Nanjing University, Nanjing 210023, China*

ZHAOP@LAMDA.NJU.EDU.CN
ZHANGLJ@LAMDA.NJU.EDU.CN

Editors: A. Jadbabaie, J. Lygeros, G. J. Pappas, P. A. Parrilo, B. Recht, C. J. Tomlin, M. N. Zeilinger

Abstract

In this paper, we present an improved analysis for dynamic regret of strongly convex and smooth functions. Specifically, we investigate the Online Multiple Gradient Descent (OMGD) algorithm proposed by Zhang et al. (2017). The original analysis shows that the dynamic regret of OMGD is at most $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T\})$, where \mathcal{P}_T and \mathcal{S}_T are path-length and squared path-length that measures the cumulative movement of minimizers of the online functions. We demonstrate that by an improved analysis, the dynamic regret of OMGD can be improved to $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T\})$, where \mathcal{V}_T is the function variation of the online functions. Note that the quantities of $\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T$ essentially reflect different aspects of environmental non-stationarity—they are not comparable in general and are favored in different scenarios. Therefore, the dynamic regret presented in this paper actually achieves a *best-of-three-worlds* guarantee and is strictly tighter than previous results.

Keywords: Online Learning, Dynamic Regret, Strong Convexity, Smoothness, Gradient Descent

1. Introduction

In the development of online convex optimization, there are plenty of works devoted to designing online algorithms for minimizing static regret (Hazan, 2016), defined as

$$\text{S-Regret}_T = \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}), \quad (1)$$

which is the difference between the cumulative loss of the online algorithm and that of the best strategy in hindsight. When environment variables are stationary, minimizing static regret will lead to an algorithm that behaves well over the iterations. However, such a claim may not hold when environments are non-stationary and changing with time. To cope with non-stationary environments where the optimal decisions of online functions can be drifting over time, a more stringent measure—*dynamic regret*—is proposed and draws much attentions in recent years (Zinkevich, 2003; Hall and Willett, 2013; Besbes et al., 2015; Jadbabaie et al., 2015; Mokhtari et al., 2016; Yang et al., 2016; Zhang et al., 2017, 2018b; Baby and Wang, 2019; Zhang et al., 2020a,b), defined as

$$\text{D-Regret}_T = \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*), \quad (2)$$

where $\mathbf{x}_t^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x})$ is the minimizer of the online function. Dynamic regret enforces the player to compete with a time-varying comparator sequence, and thus is particularly favored in online

learning in non-stationary environments. The notion of dynamic regret is also referred to as tracking regret or shifting regret in the setting of prediction with expert advice (Herbster and Warmuth, 1998, 2001; Bousquet and Warmuth, 2002; Wei et al., 2016; Zheng et al., 2019).

It is known that in the worst case, a sub-linear dynamic regret is not attainable unless imposing certain regularities of the comparator sequence or the function sequence (Besbes et al., 2015; Jadbabaie et al., 2015). There are mainly three kinds of regularities used in the literature (Zinkevich, 2003; Besbes et al., 2015; Zhang et al., 2017).

- Path-length (Zinkevich, 2003): the variation of optimizers

$$\mathcal{P}_T = \sum_{t=2}^T \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2,$$

- Squared path-length (Zhang et al., 2017): the squared variation of optimizers

$$\mathcal{S}_T = \sum_{t=2}^T \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2^2,$$

- Function variation (Besbes et al., 2015): the variation over consecutive function values

$$\mathcal{V}_T = \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |f_{t-1}(\mathbf{x}) - f_t(\mathbf{x})|,$$

When the path-length \mathcal{P}_T is known in advance, dynamic regret of Online Gradient Descent (OGD) is at most $\mathcal{O}(\sqrt{T(1 + \mathcal{P}_T)})$ (Zinkevich, 2003; Yang et al., 2016) for convex functions. For strongly convex and smooth functions, Mokhtari et al. (2016) first show that an $\mathcal{O}(\mathcal{P}_T)$ dynamic regret is achievable; later, Zhang et al. (2017) propose Online Multiple Gradient Descent (OMGD) and prove an $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T\})$ dynamic regret. Yang et al. (2016) disclose that the $\mathcal{O}(\mathcal{P}_T)$ rate is also attainable for convex and smooth functions, provided that all the minimizers \mathbf{x}_t^* 's lie in the interior of the feasible set \mathcal{X} . Besides, Besbes et al. (2015) show that OGD with a restarting strategy attains an $\mathcal{O}(T^{2/3}\mathcal{V}_T^{1/3})$ dynamic regret when the function variation \mathcal{V}_T is available ahead of time. Later, Baby and Wang (2019) improve the rate to $\mathcal{O}(T^{1/3}\mathcal{V}_T^{2/3})$ for 1-dimensional square loss by trend filtering techniques. We finally remark that another strengthened form of dynamic regret is recently studied (Zhang et al., 2018a; Zhao et al., 2020a,b) which supports to compete with any sequence of changing comparators rather than the optimizers of online functions only.

In this paper, we focus on the dynamic regret measure (2) of online convex optimization for *strongly convex* and *smooth* functions. Specifically, we assume that the online functions f_1, \dots, f_T are λ -strongly convex and L -smooth, namely, for all $t = 1, \dots, T$,

- λ -strong convexity: for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, the following condition holds

$$f_t(\mathbf{y}) \geq f_t(\mathbf{x}) + \nabla f_t(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{\lambda}{2} \|\mathbf{y} - \mathbf{x}\|_2^2.$$

- L -smoothness: for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, the following condition holds

$$f_t(\mathbf{y}) \leq f_t(\mathbf{x}) + \nabla f_t(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|_2^2.$$

To minimize the dynamic regret of strongly convex and smooth functions, [Zhang et al. \(2017\)](#) propose an algorithm called Online Multiple Gradient Descent (OMGD) and prove that under certain mild assumptions, OMGD enjoys the following dynamic regret

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T\}). \quad (3)$$

In this paper, we present an improved analysis and demonstrate that the dynamic regret of OMGD can be also bounded by the function variation term \mathcal{V}_T . As a result, we actually show that the dynamic regret of OMGD is bounded by the minimization of path-length, squared path-length and function variation, namely,

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T\}). \quad (4)$$

As shown in the work of [Jadbabaie et al. \(2015\)](#), these three kinds of regularities $\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T$ are generally incomparable and are favored in different scenarios. Therefore, the dynamic regret bound (4) presented in this paper actually achieves a *best-of-three-worlds* guarantee, and is strictly tighter than the bound (3) of [Zhang et al. \(2017\)](#).

2. Algorithm: Online Multiple Gradient Descent

We first introduce the Online Gradient Descent (OGD) algorithm ([Zinkevich, 2003](#)), which starts from any $\mathbf{x}_1 \in \mathcal{X}$ and performs the following update at each iteration:

$$\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\mathbf{x}_t - \eta \nabla f_t(\mathbf{x}_t)],$$

where $\eta > 0$ is the step size, and $\Pi_{\mathcal{X}}[\cdot]$ denotes Euclidean projection onto the nearest point in \mathcal{X} .

To minimize the dynamic regret of strongly convex and smooth functions, [Zhang et al. \(2017\)](#) propose a variant of OGD, called Online Multiple Gradient Descent (OMGD). The algorithm performs gradient descent multiple times at each iteration. Specifically, at iteration t , given the current decision \mathbf{x}_t , OMGD will produce a sequence of $\mathbf{z}_t^1, \dots, \mathbf{z}_t^K, \mathbf{z}_t^{K+1}$, where K is the number of inner iterations, a constant *independent* from the time horizon T . The inner iterations start from $\mathbf{z}_t^1 = \mathbf{x}_t$ and then perform the following update procedure,

$$\mathbf{z}_t^{k+1} = \Pi_{\mathcal{X}}[\mathbf{z}_t^k - \eta \nabla f_t(\mathbf{z}_t^k)],$$

where $k = 1, \dots, K$ is the index of inner loop. The decision \mathbf{x}_{t+1} is set as the output of the inner iterations \mathbf{z}_t^{K+1} , i.e., $\mathbf{x}_{t+1} = \mathbf{z}_t^{K+1}$. The procedures of OMGD are summarized in Algorithm 1.

3. Dynamic Regret Analysis

In this section, we provide dynamic regret analysis for the OMGD algorithm. We first restate the (squared) path-length bounds of [Zhang et al. \(2017\)](#) in Section 3.1 and then prove the function variation bounds in Section 3.2. Finally, in Section 3.3 we present comparisons between various regularities and show the advantage of our results.

Before presenting the theoretical analysis, we state the following standard assumption adopted in the work of [Zhang et al. \(2017\)](#) and this work.

Algorithm 1 Online Multiple Gradient Descent (OMGD) (Zhang et al., 2017)

Input: number of inner iterations K and step size η

- 1: Let \mathbf{x}_1 be any point in \mathcal{X}
- 2: **for** $t = 1, \dots, T$ **do**
- 3: Submit \mathbf{x}_t and receive the loss $f_t : \mathcal{X} \mapsto \mathbb{R}$
- 4: $\mathbf{z}_t^1 = \mathbf{x}_t$
- 5: **for** $k = 1, \dots, K$ **do**
- 6:

$$\mathbf{z}_t^{k+1} = \Pi_{\mathcal{X}}[\mathbf{z}_t^k - \eta \nabla f_t(\mathbf{z}_t^k)]$$

- 7: **end for**
 - 8: $\mathbf{x}_{t+1} = \mathbf{z}_t^{K+1}$
 - 9: **end for**
-

Assumption 1. Suppose the following conditions hold for each online function $f_t : \mathcal{X} \mapsto \mathbb{R}$:

- The online function f_t is λ -strongly convex and L -smooth over \mathcal{X} ;
- The gradients are bounded by G , i.e., $\|\nabla f_t(\mathbf{x})\|_2 \leq G$ for any $\mathbf{x} \in \mathcal{X}$.

3.1. (Squared) Path-length Bounds

Zhang et al. (2017) prove that the dynamic regret of OMGD can be bounded by the path-length and squared path-length. We restate their results as follows.

Theorem 1 (Theorem 2 of Zhang et al. (2017)). *Under Assumption 1, by setting the step size $\eta \leq 1/L$ and the number of inner iterations $K = \lceil \frac{1/\eta + \lambda}{2\lambda} \ln 4 \rceil$ in Algorithm 1, for any constant $\alpha > 0$, we have*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \min \left\{ \begin{array}{l} 2G\mathcal{P}_T + 2G\|\mathbf{x}_1 - \mathbf{x}_1^*\|_2, \\ \frac{1}{2\alpha} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t^*)\|_2^2 + 2(L + \alpha)\mathcal{S}_T + (L + \alpha)\|\mathbf{x}_1 - \mathbf{x}_1^*\|_2^2. \end{array} \right. \quad (5)$$

Furthermore, suppose $\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t^*)\|_2^2 = \mathcal{O}(\mathcal{S}_T)$, we have

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T\}). \quad (6)$$

In particular, if \mathbf{x}_t^* belongs to the relative interior of \mathcal{X} (i.e., $\nabla f_t(\mathbf{x}_t^*) = 0$) for all $t \in [T]$, the dynamic regret bound in (5), as $\alpha \rightarrow 0$, implies

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \min \{ 2G\mathcal{P}_T + 2G\|\mathbf{x}_1 - \mathbf{x}_1^*\|_2, 2L\mathcal{S}_T + L\|\mathbf{x}_1 - \mathbf{x}_1^*\|_2^2 \}.$$

3.2. Function Variation Bounds

In this part, we show that by an improved analysis, the dynamic regret of OMGD can be further bounded by the function variation \mathcal{V}_T , as demonstrated in the following theorem.

Theorem 2. *Under Assumption 1, by setting the step size $\eta = 1/L$, and the number of inner iterations $K = \lceil \frac{4(L+\lambda)}{\lambda} \ln 4 \rceil$ in Algorithm 1, we have*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq 2\mathcal{V}_T + 2(f_1(\mathbf{x}_1) - f_T(\mathbf{x}_{T+1})). \quad (7)$$

Furthermore, suppose $\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t^*)\|_2^2 = \mathcal{O}(\mathcal{S}_T)$, from Theorem 1 and (7), we have

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T\}).$$

In particular, if \mathbf{x}_t^* belongs to the relative interior of \mathcal{X} (i.e., $\nabla f_t(\mathbf{x}_t^*) = 0$) for all $t \in [T]$, the dynamic regret bounds in (5) and (7), as $\alpha \rightarrow 0$, imply

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \min \{ & 2G\mathcal{P}_T + 2G\|\mathbf{x}_1 - \mathbf{x}_1^*\|_2, 2L\mathcal{S}_T + L\|\mathbf{x}_1 - \mathbf{x}_1^*\|_2^2, \\ & 2\mathcal{V}_T + 2(f_1(\mathbf{x}_1) - f_T(\mathbf{x}_{T+1})) \}. \end{aligned}$$

Remark 1. Notice that above settings of step size $\eta = 1/L$ and the number of inner iterations $K = \lceil \frac{4(L+\lambda)}{\lambda} \ln 4 \rceil$ also satisfy the condition of Theorem 1 (namely, $\eta \leq 1/L$ and $K \geq \lceil \frac{L+\lambda}{2\lambda} \ln 4 \rceil$ hold simultaneously), so the dynamic regret is also bounded by the path-length bounds in (6).

To prove Theorem 2, we introduce the following key lemma due to [Nesterov \(2013\)](#).

Lemma 1. *Assume that the function $f : \mathcal{X} \mapsto \mathbb{R}$ is λ -strongly convex and L -smooth, and denote by \mathbf{x}^* the optimizer, i.e., $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. Let*

$$\mathbf{v} = \Pi_{\mathcal{X}} \left[\mathbf{u} - \frac{1}{L} \nabla f(\mathbf{u}) \right]. \quad (8)$$

Then, we have

$$f(\mathbf{v}) - f(\mathbf{x}^*) \leq \gamma(f(\mathbf{u}) - f(\mathbf{x}^*)), \quad (9)$$

where

$$\gamma = \begin{cases} \frac{1}{2}, & \text{if } 3\lambda \geq 2L \\ 1 - \frac{\lambda}{4(L-\lambda)}, & \text{otherwise.} \end{cases}$$

Proof [of Lemma 1] Notice that the update procedure in (8) is equivalent to

$$\mathbf{v} = \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ f(\mathbf{u}) + \langle \nabla f(\mathbf{u}), \mathbf{x} - \mathbf{u} \rangle + \frac{L}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 \right\}. \quad (10)$$

Then, we have

$$\begin{aligned}
 f(\mathbf{v}) &\leq f(\mathbf{u}) + \langle \nabla f(\mathbf{u}), \mathbf{v} - \mathbf{u} \rangle + \frac{L}{2} \|\mathbf{v} - \mathbf{u}\|_2^2 && \text{(by } L\text{-smoothness)} \\
 &= \min_{\mathbf{x} \in \mathcal{X}} \left\{ f(\mathbf{u}) + \langle \nabla f(\mathbf{u}), \mathbf{x} - \mathbf{u} \rangle + \frac{L}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 \right\} && \text{(due to update procedure in (10))} \\
 &\leq \min_{\mathbf{x} \in \mathcal{X}} \left\{ f(\mathbf{x}) - \frac{\lambda}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 + \frac{L}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 \right\} && \text{(by } \lambda\text{-strong convexity)} \\
 &\leq \min_{\mathbf{x} = \alpha \mathbf{x}^* + (1-\alpha)\mathbf{u}, \alpha \in [0,1]} \left\{ f(\mathbf{x}) + \frac{L-\lambda}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 \right\} \\
 &= \min_{\alpha \in [0,1]} \left\{ f(\alpha \mathbf{x}^* + (1-\alpha)\mathbf{u}) + \frac{L-\lambda}{2} \|\alpha \mathbf{x}^* + (1-\alpha)\mathbf{u} - \mathbf{u}\|_2^2 \right\} \\
 &\leq \min_{\alpha \in [0,1]} \left\{ \alpha f(\mathbf{x}^*) + (1-\alpha)f(\mathbf{u}) + \frac{L-\lambda}{2} \alpha^2 \|\mathbf{x}^* - \mathbf{u}\|_2^2 \right\} \\
 &= \min_{\alpha \in [0,1]} \left\{ f(\mathbf{u}) - \alpha (f(\mathbf{u}) - f(\mathbf{x}^*)) + \frac{L-\lambda}{2} \alpha^2 \|\mathbf{x}^* - \mathbf{u}\|_2^2 \right\} \\
 &\leq \min_{\alpha \in [0,1]} \left\{ f(\mathbf{u}) - \alpha (f(\mathbf{u}) - f(\mathbf{x}^*)) + \frac{L-\lambda}{2} \frac{2}{\lambda} \alpha^2 (f(\mathbf{u}) - f(\mathbf{x}^*)) \right\} \\
 &= \min_{\alpha \in [0,1]} \left\{ f(\mathbf{u}) + \left(\frac{L-\lambda}{\lambda} \alpha^2 - \alpha \right) (f(\mathbf{u}) - f(\mathbf{x}^*)) \right\}.
 \end{aligned}$$

The last inequality is true because $f(\mathbf{x}) - f(\mathbf{x}^*) \geq \frac{\lambda}{2} \|\mathbf{x} - \mathbf{x}^*\|_2^2$ holds for all $\mathbf{x} \in \mathcal{X}$ due to the strong convexity (Nesterov, 2018, Theorem 2.1.8). Therefore, if $\frac{\lambda}{2(L-\lambda)} \geq 1$, then we set $\alpha = 1$ and obtain

$$f(\mathbf{v}) - f(\mathbf{x}^*) \leq \frac{L-\lambda}{\lambda} (f(\mathbf{u}) - f(\mathbf{x}^*)) \leq \frac{1}{2} (f(\mathbf{u}) - f(\mathbf{x}^*)).$$

Otherwise, we set $\alpha = \frac{\lambda}{2(L-\lambda)}$, and obtain

$$f(\mathbf{v}) - f(\mathbf{x}^*) \leq \left(1 - \frac{\lambda}{4(L-\lambda)} \right) (f(\mathbf{u}) - f(\mathbf{x}^*)).$$

This ends the proof of Lemma 1. ■

Based on Lemma 1, we now present the proof of Theorem 2.

Proof [of Theorem 2] From Lemma 1, we know that

$$f_t(\mathbf{x}_{t+1}) - f_t(\mathbf{x}_t^*) = f_t(\mathbf{z}_t^{K+1}) - f_t(\mathbf{x}_t^*) \stackrel{(9)}{\leq} \gamma^K (f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)) \leq \frac{1}{4} (f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)). \quad (11)$$

The last inequality holds due to the following facts. From the setting of the inner iteration $K = \lceil \frac{4(L+\lambda)}{\lambda} \ln 4 \rceil$, on one hand, it is clear that $\gamma^K \leq \frac{1}{4}$ holds when $\gamma = \frac{1}{2}$. On the other hand, when $\gamma = 1 - \frac{\lambda}{4(L-\lambda)}$, we have $\left(1 - \frac{\lambda}{4(L-\lambda)} \right)^K \leq \exp\left(-\frac{\lambda K}{4(L-\lambda)}\right) \leq \frac{1}{4}$.

Therefore, we can upper bound the dynamic regret as follows.

$$\begin{aligned}
 \sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \sum_{t=2}^T f_t(\mathbf{x}_t) - f_{t-1}(\mathbf{x}_t) + f_{t-1}(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) \\
 &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \mathcal{V}_T + \sum_{t=2}^T f_{t-1}(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) \\
 &= f_1(\mathbf{x}_1) - f_T(\mathbf{x}_{T+1}) + \mathcal{V}_T + \sum_{t=1}^{T-1} f_t(\mathbf{x}_{t+1}) - f_t(\mathbf{x}_t^*) \\
 &\stackrel{(11)}{\leq} f_1(\mathbf{x}_1) - f_T(\mathbf{x}_{T+1}) + \mathcal{V}_T + \frac{1}{4} \sum_{t=1}^{T-1} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*).
 \end{aligned} \tag{12}$$

Thus, by rearranging above terms, we prove the statement in Theorem 2. \blacksquare

3.3. Comparisons of Path-length and Function Variation

As demonstrated by [Jadbabaie et al. \(2015\)](#), the path-length and function variation are not comparable in general. Let us consider the following two instances.

Instance 1 (Online linear optimization over a d -dimensional simplex). Consider the online linear optimization problem, the feasible set \mathcal{X} is set as $\Delta_d = \{\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^d, x_i \geq 0, \sum_{i=1}^d x_i = 1\}$ and the online functions are $f_t(\mathbf{x}) = \langle \mathbf{w}_t, \mathbf{x} \rangle$, where

$$\mathbf{w}_t = \begin{cases} [\frac{1}{T}, 0, 0, \dots, 0]^\top & \text{when } t \text{ is odd} \\ [0, \frac{1}{T}, 0, \dots, 0]^\top & \text{when } t \text{ is even} \end{cases}$$

Then, $\mathbf{x}_t^* = [1, 0, 0, \dots, 0]^\top$ when t is odd and $\mathbf{x}_t^* = [0, 1, 0, \dots, 0]^\top$ when t is even. So we have

$$\mathcal{P}_T = \mathcal{S}_T = \Theta(T), \quad \mathcal{V}_T = \Theta(1).$$

Instance 2 (Online linear optimization over a 2-dimensional simplex, or prediction with two-expert advice). Let the feasible set \mathcal{X} be as $\Delta_2 = \{\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^2, x_i \geq 0, x_1 + x_2 = 1\}$, and the online functions be $f_t(\mathbf{x}) = \langle \mathbf{w}_t, \mathbf{x} \rangle$, where

$$\mathbf{w}_t = \begin{cases} [-\frac{1}{2}, 0]^\top & \text{when } t \text{ is odd} \\ [0, \frac{1}{2}]^\top & \text{when } t \text{ is even} \end{cases}$$

Then, we know that the optimal decision is fixed as the first expert, that is, $\mathbf{x}_t^* = [1, 0]^\top$. Thus,

$$\mathcal{P}_T = \mathcal{S}_T = 0, \quad \mathcal{V}_T = \Theta(T).$$

From above two instances, we conclude that function variation and (squared) path-length are not comparable in general. Our analysis shows that OMGD actually enjoys an $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T\})$ dynamic regret guarantee, which achieves *the best of three worlds* and thus strictly improves the previous result of $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T\})$ by [Zhang et al. \(2017\)](#).

4. Discussion

In this section, we discuss some aspects of our results.

4.1. Relationship between Squared Path-length and Function Variation

We further discuss the relationship between \mathcal{S}_T and \mathcal{V}_T providing that the online functions are λ -strongly convex. Denote by $\mathbf{x}_t^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x})$ and $\mathbf{x}_{t-1}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f_{t-1}(\mathbf{x})$. Indeed,

$$\begin{aligned} \|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\|_2^2 &\leq \frac{1}{\lambda} \left(f_{t-1}(\mathbf{x}_t^*) - f_{t-1}(\mathbf{x}_{t-1}^*) \right) \\ &= \frac{1}{\lambda} \left(f_{t-1}(\mathbf{x}_t^*) - f_t(\mathbf{x}_t^*) + f_t(\mathbf{x}_t^*) - f_{t-1}(\mathbf{x}_{t-1}^*) \right) \\ &\leq \frac{1}{\lambda} \left(f_{t-1}(\mathbf{x}_t^*) - f_t(\mathbf{x}_t^*) + f_t(\mathbf{x}_{t-1}^*) - f_{t-1}(\mathbf{x}_{t-1}^*) \right) \\ &\leq \frac{2}{\lambda} \sup_{\mathbf{x} \in \mathcal{X}} |f_t(\mathbf{x}) - f_{t-1}(\mathbf{x})|. \end{aligned}$$

Summing over all iterations yields $\mathcal{S}_T \leq 2\mathcal{V}_T/\lambda$. However, the right-hand side exhibits an explicit dependency on the strong convexity modulus λ , which could be very large when λ is small. Our analysis shows that the undesirable dependency can be eliminated by multiple gradient descent per round, such that the dynamic regret can be upper bounded by $\mathcal{V}_T + 2(f_1(\mathbf{x}_1) - f_T(\mathbf{x}_{T+1}))$ without the $1/\lambda$ factor, as demonstrated in Theorem 2.

4.2. Extensions to Non-strongly Convex Functions

When the online functions are not strongly convex, we discover that the following *greedy strategy* also enjoys nice dynamic regret guarantees. The greedy strategy picks one of the minimizers of the last online function as the current decision, namely,

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x}). \quad (13)$$

Since the online function f_t is not guaranteed to be strongly convex, it may have multiple minimizers. Denote by \mathcal{X}_t^* the set of all its minimizers, then the greedy strategy can choose an arbitrary one from \mathcal{X}_t^* as the current decision. In essence, this greedy strategy can be regarded as a version of OMGD with a sufficiently large number of inner iterations.

The following theorem demonstrates the greedy strategy enjoys an $\mathcal{O}(\min\{\bar{\mathcal{P}}_T, \bar{\mathcal{S}}_T, \mathcal{V}_T\})$ dynamic regret. Note that the definitions of (squared) path-length terms are slightly different from previous ones to handle the potential non-uniqueness of minimizers. The proof is in Appendix A.

Theorem 3. *Under Assumption 1 except for the strong convexity condition (i.e., it is allowed that $\lambda = 0$), suppose $\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t^*)\|_2^2 = \mathcal{O}(\bar{\mathcal{S}}_T)$, then the greedy strategy (13) satisfies*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \leq \mathcal{O}(\min\{\bar{\mathcal{P}}_T, \bar{\mathcal{S}}_T, \mathcal{V}_T\}),$$

where $\bar{\mathcal{P}}_T = \max_{\{\mathbf{x}_t^* \in \mathcal{X}_t^*\}_{t=1}^T} \sum_{t=2}^T \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2$ is the path-length, $\bar{\mathcal{S}}_T = \max_{\{\mathbf{x}_t^* \in \mathcal{X}_t^*\}_{t=1}^T} \sum_{t=2}^T \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2^2$ is the squared path-length, and $\mathcal{V}_T = \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |f_{t-1}(\mathbf{x}) - f_t(\mathbf{x})|$ is the function variation.

Note that the smoothness condition is only used in deriving the squared path-length bound, which is not necessary for the path-length bound and the function variation bound.

Remark 2. We further elucidate the connection between the greedy strategy and OMGD. Both methods decide the current decision \mathbf{x}_{t+1} based on the available online function f_t . The greedy strategy selects $\mathbf{x}_{t+1} = \mathbf{x}_t^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x})$ as the minimizer directly, while OMGD is a first-order method performing multiple gradient descent $\mathbf{z}_t^{k+1} = \Pi_{\mathcal{X}}[\mathbf{z}_t^k - \eta \nabla f_t(\mathbf{z}_t^k)]$ for $k = 1, \dots, K$ with $\mathbf{z}_t^1 = \mathbf{x}_t$ and $\mathbf{x}_{t+1} = \mathbf{z}_t^{K+1}$. Thus, OMGD can be regarded as a high-quality approximation of the greedy strategy via a *finite* number of inner gradient descent, providing that online functions are strongly convex. We note that the initialization of $\mathbf{z}_t^1 = \mathbf{x}_t$ is necessary for achieving a *constant* number of inner iterations (in our setting, $K = \lceil 4(L/\lambda + 1) \ln 4 \rceil$ as shown in Theorem 2).

Remark 3. It is worth noting that in the analysis of OMGD the strong convexity is only necessitated by the function-value decay lemma (Lemma 1) in our analysis. Thus, it would be possible to relax the strong convexity condition by some other properties that allow a similar function-value decay property exhibited in (9), with the purpose of accommodating broader class of problems.

5. Conclusion

In this paper, we investigate an existing online algorithm (Online Multiple Gradient Descent, OMGD) proposed by Zhang et al. (2017) for dynamic regret minimization of strongly convex and smooth functions. Under certain mild assumptions, OMGD was shown to attain $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T\})$ dynamic regret, where \mathcal{P}_T and \mathcal{S}_T are path-length and squared path-length, respectively. This paper contributes to an improved analysis and proves an $\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T\})$ dynamic regret without modifying the algorithm, where \mathcal{V}_T is the function variation. The key technique used to realize the improvement is a careful usage of function-value decay lemma. Since different regularities $\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T$ are generally not comparable and are favored in different scenarios, our presented dynamic regret achieves a best-of-three-worlds guarantee and is strictly tighter than previous results.

Many recent works consider control and dynamical systems in non-stationary environments (Goel and Wierman, 2019; Li et al., 2019; Shi et al., 2020; Zhao et al., 2021). As a matter of fact, online learning plays an important role in modern computational control theory, especially the online non-stochastic control problem introduced by Agarwal et al. (2019). We believe that the dynamic regret minimization results developed in this paper could be of interest for further development of online control in non-stationary environments.

Acknowledgment

This work was partially supported by the Open Research Projects of Zhejiang Lab (NO. 2021KB0AB02), and the Collaborative Innovation Center of Novel Software Technology and Industrialization. We are grateful for the anonymous reviewers for their valuable comments.

Appendix A. Proof of Theorem 3

Proof [of Theorem 3] We first consider the path-length bound. Note that the path-length dynamic regret for this greedy strategy was firstly proved by [Yang et al. \(2016\)](#). Below we restate their proof.

$$\begin{aligned}
 \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) &= f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \sum_{t=2}^T f_t(\mathbf{x}_{t-1}^*) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \\
 &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + G \sum_{t=2}^T \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2 \\
 &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + G\bar{\mathcal{P}}_T = \mathcal{O}(\bar{\mathcal{P}}_T).
 \end{aligned}$$

Next, we prove the squared path-length bound.

$$\begin{aligned}
 \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) &= f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \sum_{t=2}^T f_t(\mathbf{x}_{t-1}^*) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \\
 &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \sum_{t=2}^T \left(\nabla f_t(\mathbf{x}_t^*)^\top (\mathbf{x}_{t-1}^* - \mathbf{x}_t^*) + \frac{L}{2} \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2^2 \right) \\
 &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \frac{1}{2} \sum_{t=2}^T \|\nabla f_t(\mathbf{x}_t^*)\|_2^2 + \frac{L+1}{2} \sum_{t=2}^T \|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\|_2^2 \\
 &\leq \mathcal{O}(\bar{\mathcal{S}}_T),
 \end{aligned}$$

where the first inequality exploits L -smoothness condition and the last inequality makes use of the assumption that $\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t^*)\|_2^2 = \mathcal{O}(\bar{\mathcal{S}}_T)$.

We finally prove the function variation bound. Note that the function variation dynamic regret for this greedy strategy was firstly mentioned by [Jadbabaie et al. \(2015\)](#) in the paragraph under Eq. (5) of their paper. Following the derivation of (12), we have

$$\begin{aligned}
 \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) &= f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \sum_{t=2}^T f_t(\mathbf{x}_{t-1}^*) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \\
 &\leq f_1(\mathbf{x}_1) - f_T(\mathbf{x}_T^*) + \sum_{t=2}^T (f_t(\mathbf{x}_{t-1}^*) - f_{t-1}(\mathbf{x}_{t-1}^*)) \\
 &\leq f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*) + \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |f_t(\mathbf{x}) - f_{t-1}(\mathbf{x})| = \mathcal{O}(\mathcal{V}_T).
 \end{aligned}$$

Combining above three upper bounds of dynamic regret ends the proof. ■

References

- Naman Agarwal, Brian Bullins, Elad Hazan, Sham M. Kakade, and Karan Singh. Online control with adversarial disturbances. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 111–119, 2019.
- Dheeraj Baby and Yu-Xiang Wang. Online forecasting of total-variation-bounded sequences. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 11071–11081, 2019.
- Omar Besbes, Yonatan Gur, and Assaf J. Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, 2015.
- Olivier Bousquet and Manfred K. Warmuth. Tracking a small set of experts by mixing past posteriors. *Journal of Machine Learning Research*, 3:363–396, 2002.
- Gautam Goel and Adam Wierman. An online algorithm for smoothed regression and LQR control. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2504–2513, 2019.
- Eric C. Hall and Rebecca Willett. Dynamical models and tracking regret in online convex programming. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 579–587, 2013.
- Elad Hazan. Introduction to Online Convex Optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- Mark Herbster and Manfred K. Warmuth. Tracking the best expert. *Machine Learning*, 32(2): 151–178, 1998.
- Mark Herbster and Manfred K. Warmuth. Tracking the best linear predictor. *Journal of Machine Learning Research*, 1:281–309, 2001.
- Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 398–406, 2015.
- Yingying Li, Xin Chen, and Na Li. Online optimal control with linear dynamics and predictions: Algorithms and regret analysis. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 14858–14870, 2019.
- Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *Proceedings of the 55th IEEE Conference on Decision and Control (CDC)*, pages 7195–7201, 2016.
- Yurii Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- Yurii Nesterov. *Lectures on Convex Optimization*, volume 137. Springer, 2018.

- Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization with memory and competitive control. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.
- Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. Tracking the best expert in non-stationary stochastic environments. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 3972–3980, 2016.
- Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng 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 (ICML)*, pages 449–457, 2016.
- Lijun Zhang, Tianbao Yang, Jinfeng Yi, Rong Jin, and Zhi-Hua Zhou. Improved dynamic regret for non-degenerate functions. In *Advances in Neural Information Processing Systems 30 (NIPS)*, pages 732–741, 2017.
- Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 1330–1340, 2018a.
- Lijun Zhang, Tianbao Yang, Rong Jin, and Zhi-Hua Zhou. Dynamic regret of strongly adaptive methods. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 5877–5886, 2018b.
- Lijun Zhang, Shiyin Lu, and Tianbao Yang. Minimizing dynamic regret and adaptive regret simultaneously. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 309–319, 2020a.
- Yu-Jie Zhang, Peng Zhao, and Zhi-Hua Zhou. A simple online algorithm for competing with dynamic comparators. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 390–399, 2020b.
- Peng Zhao, Guanghui Wang, Lijun Zhang, and Zhi-Hua Zhou. Bandit convex optimization in non-stationary environments. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1508–1518, 2020a.
- Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 12510–12520, 2020b.
- Peng Zhao, Yu-Xiang Wang, and Zhi-Hua Zhou. Non-stationary online learning with memory and non-stochastic control. *ArXiv preprint*, arXiv:2102.03758, 2021.
- Kai Zheng, Haipeng Luo, Ilias Diakonikolas, and Liwei Wang. Equipping experts/bandits with long-term memory. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 5927–5937, 2019.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936, 2003.