On Online Optimization: Dynamic Regret Analysis of Strongly Convex and Smooth Problems

Ting-Jui Chang, Shahin Shahrampour
Wm Michael Barnes '64 Department of Industrial and Systems Engineering
Texas A&M University, College Station, TX, USA
{tingjui.chang, shahin}@tamu.edu

Abstract

The regret bound of *dynamic* online learning algorithms is often expressed in terms of the variation in the function sequence (V_T) and/or the path-length of the minimizer sequence after T rounds. For strongly convex and smooth functions, Zhang et al. (2017) establish the squared path-length of the minimizer sequence ($C^*_{2,T}$) as a lower bound on regret. They also show that online gradient descent (OGD) achieves this lower bound using multiple gradient queries per round. In this paper, we focus on unconstrained online optimization. We first show that a preconditioned variant of OGD achieves $O\big(\min\{C^*_T, C^*_{2,T}\}\big)$ with one gradient query per round (C^*_T refers to the normal path-length). We then propose online optimistic Newton (OON) method for the case when the first and second order information of the function sequence is *predictable*. The regret bound of OON is captured via the quartic path-length of the minimizer sequence ($C^*_{4,T}$), which can be much smaller than $C^*_{2,T}$. We finally show that by using multiple gradients for OGD, we can achieve an upper bound of $O(\min\{C^*_{2,T}, V_T\})$ on regret.

1 Introduction

Online optimization is modeled as a repeated game between a learner and an adversary (Hazan 2016). At the t-th round, $t \in [T] \triangleq \{1, \ldots, T\}$, the learner selects an action $\mathbf{x}_t$ from a convex set $\mathcal{X} \subseteq \mathrm{R}^n$ based on the information from previous rounds. Then, the adversary reveals a convex function $f_t : \mathcal{X} \to \mathrm{R}$ to the learner that incurs the loss $f_t(\mathbf{x}_t)$. The goal of online learning is to minimize the *regret*, which is the difference between the cumulative loss of the learner and that of a comparator sequence in hindsight. Depending on the comparator sequence, the regret can be either static or dynamic. The static regret is defined with respect to a fixed comparator as follows

$$\mathbf{Reg}^s_T \triangleq \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x}). \tag{1}$$

Static regret is well-studied in the literature of online optimization. Zinkevich (2003) shows that online gradient descent (OGD) provides a $O(\sqrt{T})$ upper bound on static regret for convex functions. Hazan, Agarwal, and Kale (2007) improve this bound to $O(\log T)$ for exp-concave functions as well as strongly convex functions. These bounds turn out to be optimal given their corresponding lower bounds (Hazan 2016). A more stringent benchmark for regret can be defined when the comparator sequence is *time-varying*, introducing the notion of *dynamic* regret (Besbes, Gur, and Zeevi 2015; Jadbabaie et al. 2015). In this case, the learner's performance is measured against the best sequence of actions (minimizers) at each round as follows

$$\mathbf{Reg}^d_T \triangleq \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{x}^*_t), \tag{2}$$

where $\mathbf{x}^*_t \triangleq \mathrm{argmin}_{\mathbf{x}\in\mathcal{X}} f_t(\mathbf{x})$. More generally, dynamic regret against an arbitrary comparator sequence $\{\mathbf{u}_t\}_{t=1}^T$ is defined as (Zinkevich 2003),

$$\mathbf{Reg}^d_T(\mathbf{u}_1, \ldots, \mathbf{u}_T) \triangleq \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}_t). \tag{3}$$

It is well-known that since the function sequence can fluctuate arbitrarily, the worst-case dynamic regret scales linearly with respect to T. However, when the environment varies slowly, there is hope to bound dynamic regret.

In the past few years, various studies on *dynamic online learning* have provided regret bounds in terms of the variation in the function sequence and/or the path-length of the minimizer sequence (Besbes, Gur, and Zeevi 2015; Hall and Willett 2015; Jadbabaie et al. 2015). The path-length of an arbitrary sequence $\{\mathbf{u}_t\}_{t=1}^T$ is defined as (Zinkevich 2003),

$$C_T(\mathbf{u}_1, \ldots, \mathbf{u}_T) \triangleq \sum_{t=2}^{T} \left\|\mathbf{u}_t - \mathbf{u}_{t-1}\right\|. \tag{4}$$

Zinkevich (2003) shows that applying OGD for convex functions results in an upper bound of $O(\sqrt{T}C_T)$ on dynamic regret. If the function sequence is assumed to be strongly convex and smooth, the upper bound can be further improved to $O(C^*_T)$ (Mokhtari et al. 2016), where

$$C^*_T \triangleq C_T(\mathbf{x}^*_1, \ldots, \mathbf{x}^*_T) = \sum_{t=2}^{T} \left\|\mathbf{x}^*_t - \mathbf{x}^*_{t-1}\right\|.$$

Let us now define a new variation measure $C^*_{p,T}$ (the path-length of order p) as follows

$$C^*_{p,T} \triangleq C_{p,T}(\mathbf{x}^*_1, \ldots, \mathbf{x}^*_T) = \sum_{t=2}^{T} \left\|\mathbf{x}^*_t - \mathbf{x}^*_{t-1}\right\|^p , \quad (5)$$

with the convention that $C^*_T = C^*_{1,T}$. It can be immediately seen that any bound of order $C^*_{p,T}$ implies a bound of order $C^*_{q,T}$ for $q < p$, as long as the minimizer sequence is assumed to be uniformly bounded (which is the case in the present work). Recently, Zhang et al. (2017) prove that by using multiple gradient queries in one round, the regret of OGD can be improved to $O(C^*_{2,T})$, which can be much smaller than C^*_T when the local variations are small. Zhang et al. (2017) also prove that the bound $O(C^*_{2,T})$ is optimal in the worst-case.

Besides the path-length, another commonly used regularity measure is V_T, the cumulative variation in the function sequence, defined as

$$V_T \triangleq \sum_{t=2}^{T} \sup_{\mathbf{x}\in\mathcal{X}} |f_t(\mathbf{x}) - f_{t-1}(\mathbf{x})|. \quad (6)$$

Besbes, Gur, and Zeevi (2015) show that the dynamic regret can be bounded by $O(T^{2/3}(V_T + 1)^{1/3})$ and $O(\sqrt{T(1 + V_T)})$ for convex functions and strongly convex functions, respectively. Note that in general C_T and V_T are not directly comparable, and Jadbabaie et al. (2015) provide problem environments where V_T and C_T are significantly different in terms of the order.

In this work, we focus on *unconstrained* online optimization for strongly convex and smooth functions and study dynamic regret in the sense of (2). Our contribution is threefold:

- We propose online preconditioned gradient descent (OPGD), where the gradient direction is re-scaled by a time-varying positive-definite matrix at each round. We show that OPGD achieves the dynamic regret bound of $O\big(\min\{C^*_T, C^*_{2,T}\}\big)$ with one gradient query in each round. Beside matching the lower bound in (Zhang et al. 2017), the result also entails that OGD and a regularized variant of online Newton method enjoy the same regret bound (Section 3.2).
- Inspired by optimistic mirror descent (Rakhlin and Sridharan 2013b), where predictions of the gradient sequence are used, we propose optimistic online Newton (OON) by incorporating predictions of the Hessian and gradient to online Newton method. We prove that in this case the dynamic regret bound can be further improved to $O(C^*_{4,T})$ if the predictions are accurate enough (Section 3.3). This is verified empirically in Section 6.
- We finally show that by applying multiple gradient descents, the dynamic regret is upper bounded by $O(\min\{V_T, C^*_{2,T}\})$. We also construct problem setups where V_T is much larger than $C^*_{2,T}$ in order, and vice versa (Section 4).

The proofs of our results are provided in the supplementary material.

2 Related Literature

In this section, we provide related literature on dynamic regret defined in (2) and (3). A summary of the results is tabulated in Tables 1 and 2. In particular, Table 1 summarizes results with one gradient query per round, whereas Table 2 exhibits those with multiple gradient queries per round.

As previously mentioned, Zinkevich (2003) shows that when the functions are convex, by applying OGD with a diminishing step size of $1/\sqrt{t}$, the dynamic regret defined in (3) can be bounded by $O(\sqrt{T}(1 + C_T))$. Zhang, Lu, and Zhou (2018) combine OGD with expert advice to improve the bound to $O(\sqrt{T(1 + C_T)})$. Focusing on regret in the sense of (2), Mokhtari et al. (2016) establish a regret bound of $O(C^*_T)$ for OGD under strong convexity and smoothness of the function sequence. Lesage-Landry, Taylor, and Shames (2020) show the same bound for online Newton method.

To further express the existing regret bounds, we need to define several other regularity measures. The first one is similar to the path-length (4) and is defined as

$$C'_T(\mathbf{u}_1, \ldots, \mathbf{u}_T) \triangleq \sum_{t=2}^{T} \|\mathbf{u}_t - \Phi_t(\mathbf{u}_{t-1})\| , \quad (7)$$

where $\Phi_t(\cdot)$ is a given dynamics (available to the learner). Hall and Willett (2015) propose a dynamic mirror descent algorithm that incorporates the dynamics $\{\Phi_t(\cdot)\}_{t=1}^T$ into online mirror descent and achieves a regret bound of $O(\sqrt{T}(1 + C'_T))$.

Besbes, Gur, and Zeevi (2015) propose a restarted OGD and analyze its performance for the case when only *noisy* gradients are available to the learner. They prove that the expected dynamic regret is bounded by $O(T^{2/3}(V_T + 1)^{1/3})$ and $O(\sqrt{T(1 + V_T)})$ for convex and strongly convex functions, respectively. The restarted OGD of (Besbes, Gur, and Zeevi 2015) is designed under the assumption that V_T (or an upper bound on V_T) is available to the learner from the outset.

Another measure is D_T, the variation in gradients, which is defined as

$$D_T \triangleq \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \mathbf{m}_t\|^2 , \quad (8)$$

where $\mathbf{m}_t$ is a predictable sequence computed by the learner before round t (Rakhlin and Sridharan 2013a,b). A special version of D_T with $\mathbf{m}_t = \nabla f_{t-1}$ is introduced by (Chiang et al. 2012), and the current definition is used by (Rakhlin and Sridharan 2013a,b) for studying optimistic mirror descent. Nevertheless, all of these works deal with static regret. Motivated by the fact that various regularity measures are not directly comparable, Jadbabaie et al. (2015) propose an adaptive version of optimistic mirror descent to bound dynamic regret (2). They establish a regret bound in terms of C_T, D_T, and V_T for convex functions with the assumption that the learner can accumulate each of these measures on-the-fly. When $V_T = 0$ or $C_T = 0$, their bound recovers that of (Rakhlin and Sridharan 2013b) on static regret.

Reference	Regret Definition	Setup	Regret Bound	
(2003)	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t)$	Convex	$O\big(\sqrt{T}(1 + C_T(\mathbf{u}_1, \dots, \mathbf{u}_T))\big)$	
(2015)	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t)$	Convex	$O\big(\sqrt{T}(1 + C'_T(\mathbf{u}_1, \dots, \mathbf{u}_T))\big)$	
(2015)	$\sum_t \mathbf{E}[f_t(\mathbf{x}_t)] - f_t(\mathbf{x}_t^*)$	Convex	$O\big(T^{2/3}(1 + V_T)^{1/3}\big)$	
(2015)	$\sum_t \mathbf{E}[f_t(\mathbf{x}_t)] - f_t(\mathbf{x}_t^*)$	Strongly Convex	$O\big(\sqrt{T(1 + V_T)}\big)$	
(2015)	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)$	Convex	$O\big(\sqrt{D_T + 1} + \min\big\{\sqrt{(D_T + 1)C_T^*}, [(D_T + 1)V_T T]^{\frac{1}{3}}\big\}\big)$	
(2016)	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)$	Strongly Convex and Smooth	$O\big(C_T^*\big)$	
(2018)	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t)$	Convex	$O\big(\sqrt{T(1 + C_T(\mathbf{u}_1, \dots, \mathbf{u}_T))}\big)$	
This work	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)$	Strongly Convex and Smooth	$O\big(\min\{C_T^*, C_{2,T}^*\}\big)$	OPGD Algorithm
This work	$\sum_t f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)$	Strongly Convex and Smooth	$O\big(C_{4,T}^* + D'_T\big)$	OON Algorithm

Table 1: Related works on dynamic online learning (single gradient query in each round). The bounds presented for this work are on unconstrained setup, whereas other works deal with constrained setup.

Reference	Regret Definition	Setup	Regret Bound
(2017)	$\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)$	Strongly Convex and Smooth	$O\big(\min\{C_T^*, C_{2,T}^*\}\big)$
This work	$\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*)$	Strongly Convex and Smooth	$O\big(\min\{V_T, C_{2,T}^*\}\big)$

Table 2: Related works on dynamic online learning (multiple gradient queries in each round). The bound presented for this work is on unconstrained setup, whereas other works deal with constrained setup.

Inspired by the notion of D_T, we introduce in Section 3.3 a new regularity D'_T defined as

$$D'_T \triangleq \sum_{t=2}^T \left\|\mathbf{M}_t^{-1}(\cdot)\mathbf{m}_t(\cdot) - (\nabla^2 f_t(\cdot))^{-1}\nabla f_t(\cdot)\right\|^2, \quad (9)$$

where $\mathbf{M}_t(\cdot)$ and $\mathbf{m}_t(\cdot)$ denote the predictions of $\nabla^2 f_t(\cdot)$ and $\nabla f_t(\cdot)$, respectively. We later show that by applying OON, the dynamic regret bound is $O(D'_T + C_{4,T}^*)$. This is an improvement over $O(C_{2,T}^*)$ only if D'_T is small, i.e., the predictions of the learner are accurate enough. Nevertheless, we also show that if the learner uses stale gradient/Hessian information in the form of $\mathbf{m}_t = \nabla f_{t-1}$ and $\mathbf{M}_t = \nabla^2 f_{t-1}$, the regret is still $O(C_{2,T}^*)$.

Other related works on dynamic regret include (Ravier, Calderbank, and Tarokh 2019; Yuan and Lamperski 2020). Ravier, Calderbank, and Tarokh (2019) assume the function sequence has a parametrizable structure and quantify the functional difference in terms of the variation in parameters. They propose an online gradient method combined with the prediction of the function parameters and show that the dynamic regret can be bounded in terms of C_T^* as well as the accumulation error in the parameters. Yuan and Lamperski (2020) analyze the trade-off between static and dynamic regret through studying the effect of forgetting factors for a class of online Newton algorithms. Ajalloeian, Simonetto, and Dall'Anese (2020) study online proximal-gradient method to track the minimizers of a composite convex function sequence. They provide, for strongly convex and convex functions, the regret bounds which take the approximation error of gradients and the proximal operator into consideration.

We note that Zhang et al. (2017) prove a $O(\min\{C_T^*, C_{2,T}^*\})$ regret bound with multiple gradient queries for OGD. We revisit the same algorithm (in an unconstrained setup) and establish a bound of $O(\min\{V_T, C_{2,T}^*\})$. The main benefit of the latter is that V_T and $C_{2,T}^*$ are not comparable, whereas $C_{2,T}^* = O(C_T^*)$ as long as the minimizer sequence is bounded.

Adaptive Regret: Beside the works related to the dynamic regret, the notion of adaptive regret (Hazan and Seshadhri 2007; Daniely, Gonen, and Shalev-Shwartz 2015; Zhang et al. 2018; Zhang, Liu, and Zhou 2019; Zhang, Lu, and Yang 2020) is also proposed to capture the dynamics in the environment. Adaptive regret characterizes a local version of static regret, where

$$\mathbf{Reg}_T^a([r, s]) \triangleq \sum_{t=r}^s f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{X}} \sum_{t=r}^s f_t(\mathbf{x}),$$

for each interval $[r, s] \subseteq [T]$. Zhang et al. (2018) draw a connection between strongly adaptive regret and dynamic regret and propose an adaptive algorithm which can bound the dynamic regret without prior knowledge of the functional variation. Zhang, Lu, and Yang (2020) propose a novel algorithm which can minimize the dynamic regret and the adaptive regret simultaneously.

3 Main Results

In this section, we present our main results. We first prove that OPGD achieves the optimal bound of $O\big(\min\{C_T^*, C_{2,T}^*\}\big)$ for dynamic regret (2), matching the lower bound of (Zhang et al. 2017). Then, we develop a variant of online Newton method (called OON), which employs predicted first and second order information in the update. The bound on the dynamic regret of OON can be improved to $O(C_{4,T}^*)$ if the predictions are accurate.

3.1 Preliminaries

Since our results are on strongly convex and smooth functions, we start by their formal definitions below. Through-out, we assume that the function sequence $\{f_t\}_{t=1}^T$ is differentiable.

Definition 1. *A function $f : \mathcal{X} \rightarrow \mathrm{R}$ is μ-strongly convex ($\mu > 0$) over the convex set $\mathcal{X}$ if*

$$f(\mathbf{x}) \geq f(\mathbf{y}) + \nabla f(\mathbf{y})^\top (\mathbf{x}-\mathbf{y}) + \frac{\mu}{2} \|\mathbf{x}-\mathbf{y}\|^2, \forall \mathbf{x}, \mathbf{y} \in \mathcal{X}.$$

Definition 2. *A function $f : \mathcal{X} \rightarrow \mathrm{R}$ is L-smooth ($L > 0$), when its gradient is Lipschitz continuous over the convex set $\mathcal{X}$, where*

$$f(\mathbf{x}) \leq f(\mathbf{y}) + \nabla f(\mathbf{y})^\top (\mathbf{x}-\mathbf{y}) + \frac{L}{2} \|\mathbf{x}-\mathbf{y}\|^2, \forall \mathbf{x}, \mathbf{y} \in \mathcal{X}.$$

3.2 Dynamic Regret Bound for Online Preconditioned Gradient Descent (OPGD)

When the functions are μ-strongly convex and L-smooth, Mokhtari et al. (2016) prove that the dynamic regret of OGD can be upper bounded by $O(C_T^*)$. Zhang et al. (2017) show that multiple gradient descents can help improving it to $O(\min\{C_T^*, C_{2,T}^*\})$.

We observe that in the unconstrained setup (which implies $\|\nabla f_t(\mathbf{x}_t^*)\| = 0$ for every $t \in [T]$), querying just one gradient in each round is enough to obtain $O(C_{2,T}^*)$. We analyze this observation for a slightly more general case where OGD is preconditioned, i.e., the gradient direction at each round is re-scaled according to a positive definite matrix. Our method, called OPGD, is summarized in Algorithm 1. We establish the regret bound in the following theorem, under an assumption that characterizes the relationship of the condition number of the matrices used for preconditioning to parameters μ and L.

Theorem 1. *Suppose that for any $t \in [T]$:*

1. *The function $f_t : \mathrm{R}^n \rightarrow \mathrm{R}$ is μ-strongly convex and L-smooth.*
2. *The preconditioning matrix $\mathbf{A}_t$ satisfies $\lambda' \cdot \mathbf{I} \preceq \mathbf{A}_t \preceq \lambda \cdot \mathbf{I}$.*
3. *The condition number satisfies $\frac{\lambda}{\lambda'} < 1 + \frac{\mu^2}{4L^2}$.*

If we set $\eta = \frac{\lambda' \mu}{2L^2}$, the dynamic regret for the sequence of actions $\mathbf{x}_t$ generated by OPGD is bounded as follows

$$\mathbf{Reg}_T^d \leq \min \begin{cases} (1) & \left(\frac{L^2}{\mu}\right)\left(\frac{4L^2\lambda - \mu^2\lambda'}{\mu^2\lambda' - 4L^2(\lambda-\lambda')}\right) \sum_{t=2}^{T+1} \left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\|^2 \\ & + \left(\frac{L^2\lambda}{\lambda'\mu} - \frac{\mu}{4}\right) \|\mathbf{x}_1 - \mathbf{x}_1^*\|^2 \\ (2) & \frac{LD}{2}\left[\frac{\|\mathbf{x}_1 - \mathbf{x}_1^*\| - \gamma\|\mathbf{x}_T - \mathbf{x}_T^*\|}{1-\gamma}\right. \\ & \left. + \frac{1}{1-\gamma} \sum_{t=2}^{T} \left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\|\right], \end{cases}$$

where $\gamma = \sqrt{(\frac{\lambda}{\eta} - \mu)/(\mu + \frac{\lambda'}{\eta})}$ and $D \triangleq \max_{t=1,\ldots,T} \|\mathbf{x}_t - \mathbf{x}_t^*\|$.

The theorem above shows that OPGD achieves $\min\{O(C_{1,T}^*), O(C_{2,T}^*)\}$ regret. An immediate corollary is that OGD also achieves the same rate if we set $\mathbf{A}_t = \mathbf{I}$, which implies $\lambda = \lambda' = 1$.

Algorithm 1 Online Preconditioned Gradient Descent (OPGD)

Require: Initial vector $\mathbf{x}_1 \in \mathrm{R}^n$, step size η, a sequence of positive definite matrices $\mathbf{A}_t$
for $t = 1, 2, \ldots, T$ **do**
  Play $\mathbf{x}_t$
  Observe the gradient of the current action $\nabla f_t(\mathbf{x}_t)$
  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{A}_t^{-1} \nabla f_t(\mathbf{x}_t)$
end for

Algorithm 2 Optimistic Online Newton (OON)

Require: Initial vector $\mathbf{x}_1 = \hat{\mathbf{x}}_0 \in \mathrm{R}^n$,
for $t = 1, 2, \ldots, T$ **do**
  Play $\mathbf{x}_t$
  Get the predicted second order information $\mathbf{M}_{t+1}$ and first order information $\mathbf{m}_{t+1}$ of function f_{t+1}
  $\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_{t-1} - \mathbf{H}_t^{-1}(\hat{\mathbf{x}}_{t-1}) \nabla f_t(\hat{\mathbf{x}}_{t-1})$
  $\mathbf{x}_{t+1} = \hat{\mathbf{x}}_t - \mathbf{M}_{t+1}^{-1}(\hat{\mathbf{x}}_t) \mathbf{m}_{t+1}(\hat{\mathbf{x}}_t)$
end for

Corollary 2. *Suppose that for any $t \in [T]$, the function $f_t : \mathrm{R}^n \rightarrow \mathrm{R}$ is μ-strongly convex and L-smooth, and let $\mathbf{A}_t = \mathbf{I}$. Then, OPGD amounts to OGD, and it achieves a regret bound of $O\big(\min\{C_T^*, C_{2,T}^*\}\big)$.*

Another corollary of Theorem 1 is on a regularized version of online Newton method as follows.

Corollary 3. *Suppose that for any $t \in [T]$, the function $f_t : \mathrm{R}^n \rightarrow \mathrm{R}$ is μ-strongly convex and L-smooth. Let $\mathbf{A}_t = \nabla^2 f_t(\mathbf{x}_t) + \zeta \cdot \mathbf{I}$ where $\zeta > \frac{(L-\mu)4L^2}{\mu^2} - \mu$. Then, OPGD corresponds to a regularized variant of online Newton method, and it achieves a regret bound of $O\big(\min\{C_T^*, C_{2,T}^*\}\big)$.*

Proof. We just need to verify the third condition in Theorem 1 for which we require

$$\frac{L+\zeta}{\mu+\zeta} < 1 + \frac{\mu^2}{4L^2} \Longleftrightarrow \zeta > \frac{(L-\mu)4L^2}{\mu^2} - \mu.$$

$\square$

Compared to the regret bound of (Lesage-Landry, Taylor, and Shames 2020) for online Newton method, Corollary 3 puts no constraint on the relative location of the starting point, i.e., the result is global (and not local). Moreover, the regret bound is tighter as $O(C_{2,T}^*)$ always implies $O(C_T^*)$ when the minimizer sequence is uniformly bounded.

3.3 Improved Dynamic Regret Bound for Optimistic Online Newton (OON)

The optimal bound on static regret for convex functions is $O(\sqrt{T})$ (Hazan 2016). However, Rakhlin and Sridharan (2013b) show that by using gradient predictions, the regret bound can be $O(\sqrt{D_T})$, which is tighter if the predicted gradients are close enough to the actual gradients. Essentially, by using the predicted sequence, the learner aims at taking advantage of the niceness of the adversarial sequence (if possible).

In this section, we extend this idea to the case that the learner can use first and second order information simultaneously. The resulting algorithm, called optimistic online Newton (OON), is summarized in Algorithm 2. The learner performs a Newton update based on the predicted information but corrects this update using the true information. The following theorem presents the regret bound of OON in a local sense, which has the potential to significantly outperform first-order methods. The local nature of the result is not surprising as for the classical Newton method, quadratic convergence guarantee is only local (see e.g., Theorem 1.2.5 of (Nesterov 1998)).

Theorem 4. *Suppose that for any $t \in [T]$:*

1. *The function $f_t : \mathrm{R}^n \to \mathrm{R}$ is μ-strongly convex and L-smooth.*
2. *$\exists L_H > 0$ such that $\|\mathbf{H}_t(\mathbf{x}) - \mathbf{H}_t(\mathbf{x}_t^*)\| \le L_H \|\mathbf{x} - \mathbf{x}_t^*\|$, where $\mathbf{H}_t(\mathbf{x}) = \nabla^2 f_t(\mathbf{x})$.*
3. *$\exists \mathbf{x}_1 \in \mathrm{R}^n$ such that $\|\mathbf{x}_1 - \mathbf{x}_1^*\| \le \frac{\mu}{L_H}$.*
4. *There exists a bound on local variations, where $\bar{c} \triangleq \max_{t\in\{2,\dots,T\}} \left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\| \le \frac{\mu}{2L_H}$.*

Then, the dynamic regret for the sequence of actions $\mathbf{x}_t$ by OON is bounded by

$$\begin{aligned}\mathbf{Reg}_T^d \le L\Bigg(& \frac{\|\hat{\mathbf{x}}_1 - \mathbf{x}_1^*\|^2 - \rho' \|\hat{\mathbf{x}}_T - \mathbf{x}_T^*\|^2}{1-\rho'} \\ & + \frac{\rho''}{1-\rho'} \sum_{t=2}^{T} \left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\|^4 \Bigg) \\ & + LD'_T + L\left\|\mathbf{H}_1^{-1}(\hat{\mathbf{x}}_0)\nabla f_1(\hat{\mathbf{x}}_0)\right\|^2,\end{aligned} \tag{10}$$

where $\rho' \triangleq \frac{1}{16}(1+c_1)^2(1+c_2)$, $\rho'' \triangleq (\frac{L_H}{2\mu})^2(1+\frac{1}{c_1})^2(1+\frac{1}{c_2})$, c_1 and c_2 are any positive constants such that $0 < \rho' < 1$, and $D'_T \triangleq \sum_{t=2}^{T} \left\|\mathbf{M}_t^{-1}(\hat{\mathbf{x}}_{t-1})\mathbf{m}_t(\hat{\mathbf{x}}_{t-1}) - \mathbf{H}_t^{-1}(\hat{\mathbf{x}}_{t-1})\nabla f_t(\hat{\mathbf{x}}_{t-1})\right\|^2$.

The theorem indicates that when the predicted information $(\mathbf{M}_t, \mathbf{m}_t)$ is close to $(\mathbf{H}_t, \nabla f_t)$, D'_T would be small and the dynamic regret bound would be close to $O(C_{4,T}^*)$. On the other hand, it can also be shown that if the learner uses stale information, i.e., $\mathbf{M}_t = \mathbf{H}_{t-1}$ and $\mathbf{m}_t = \nabla f_{t-1}$, the regret bound of (10) is $O(C_{2,T}^*)$, which matches the optimal worst-case bound.

Corollary 5. *Suppose that the assumptions of Theorem 4 hold. If for $t = 2, \dots, T$, $(\mathbf{M}_t, \mathbf{m}_t) = (\mathbf{H}_{t-1}, \nabla f_{t-1})$, the dynamic regret for the sequence of actions $\mathbf{x}_t$ by OON is bounded by $O(C_{2,T}^*)$.*

4 Dynamic Regret Bound for Online Multiple Gradient Descents (OMGD)

In this section, we revisit the OMGD algorithm developed by (Zhang et al. 2017), outlined in Algorithm 3. Zhang et al. (2017) prove that by applying multiple gradient descents, the dynamic regret bound is $O(\min\{C_T^*, C_{2,T}^*\})$, which basically translates to $O(C_{2,T}^*)$ as soon as the minimizers are uniformly bounded. We show that the regret bound of Zhang et al. (2017) can be made more comprehensive by including V_T, the variation in the function sequence. The new bound, which takes the form of $O(\min\{V_T, C_{2,T}^*\})$, is presented below.

Algorithm 3 Online Multiple Gradient Descent (OMGD) (Zhang et al. 2017)

Require: Initial vector $\mathbf{x}_1 \in \mathrm{R}^n$, step size η, function parameters μ and L.
for $t = 1, 2, \dots, T$ **do**
  Play $\mathbf{x}_t$
  Receive the information of f_t
  Let $\mathbf{z}_{t+1}^{(0)} = \mathbf{x}_t$ and $K_t = \lceil \frac{-2\log(t)}{\log(1-\frac{2\eta\mu L}{\mu+L})} \rceil$
  for $j = 1, 2, \dots, K_t$ **do**
    $\mathbf{z}_{t+1}^{(j)} = \mathbf{z}_{t+1}^{(j-1)} - \eta\nabla f_t(\mathbf{z}_{t+1}^{(j-1)})$
  end for
  $\mathbf{x}_{t+1} = \mathbf{z}_{t+1}^{(K_t)}$
end for

Theorem 6. *Suppose that for any $t \in [T]$, the function $f_t : \mathrm{R}^n \to \mathrm{R}$ is μ-strongly convex and L-smooth. For any $0 < \eta \le \frac{2}{\mu+L}$, the dynamic regret of OMGD is bounded as follows*

$$\mathbf{Reg}_T^d \le \min \begin{cases} (f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*)) + 2V_T + \frac{\pi^2 D^2 L}{12} \\ \frac{L}{2}\left[\, \|\mathbf{x}_1 - \mathbf{x}_1^*\|^2 + 2D^2(\frac{\pi^2}{6}) + 2C_{2,T}^* \right] \end{cases}, \tag{11}$$

where $D \triangleq \max_{t\in[T]} \|\mathbf{x}_t - \mathbf{x}_t^*\|$, and V_T is defined with respect to $\mathcal{X}$ being the convex hull of $\{\mathbf{x}_t, \mathbf{x}_t^*\}_{t=1}^T$.

4.1 Comparison of $C_{2,T}^*$ and V_T

We now show that V_T and $C_{2,T}^*$ are not directly comparable to each other. Therefore, having both of them present in the regret bound can only make the bound tighter. We construct two problem environments where $C_{2,T}^* \ll V_T$ and $V_T \ll C_{2,T}^*$, respectively.

Consider the following function sequence $f_t : \mathrm{R}^n \to \mathrm{R}$

$$f_t(\mathbf{x}) = \begin{cases} \|\mathbf{x} - \mathbf{x}^*\|^2, \text{if } t \text{ is odd} \\ \|\mathbf{x} - \mathbf{x}^*\|^2 + 1, \text{if } t \text{ is even} \end{cases}$$

For this function sequence, based on the regularity definitions (5) and (6), it is clear that

$$C_{2,T}^* = \sum_{t=2}^{T} \left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\|^2 = \sum_{t=2}^{T} \|\mathbf{x}^* - \mathbf{x}^*\|^2 = 0.$$

$$V_T = \sum_{t=2}^{T} \sup_{\mathbf{x}\in\mathcal{X}} |f_t(\mathbf{x}) - f_{t-1}(\mathbf{x})| = \Theta(T).$$

In this case, we see that V_T is much larger than $C_{2,T}^*$. On the other hand, consider another function sequence $f_t : \mathrm{R}^n \to \mathrm{R}$

$$f_t(\mathbf{x}) = \begin{cases} \frac{\|\mathbf{x}\|^2}{t}, \text{if } t \text{ is odd} \\ \frac{\|\mathbf{x}-\mathbf{y}\|^2}{t}, \text{if } t \text{ is even} \end{cases}$$

We have that

$$C_{2,T}^{*} = \sum_{t=2}^{T} \left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\|^2 = \sum_{t=2}^{T} \|\mathbf{y}\|^2 = \Theta(T).$$

$$\begin{aligned} V_T &= \sum_{t=2}^{T} \sup_{\mathbf{x}\in\mathcal{X}} |f_t(\mathbf{x}) - f_{t-1}(\mathbf{x})| \\ &= \sum_{t=2}^{T} \sup_{\mathbf{x}\in\mathcal{X}} \left|\frac{2\mathbf{x}^\top \mathbf{y} + \mathbf{y}^\top \mathbf{y}}{t}\right| \leq O(\log T). \end{aligned}$$

We can see that V_T is considerably smaller than $C_{2,T}^*$ in this scenario. With these two examples, it can be seen that if the regret bound only uses one regularity, it is possible that the resulting bound is not tight.

5 Discussion on Constrained Setup

In this section, we show that for OPGD and OMGD, if the function domain is constrained, similar theoretical results can be achieved. Since the domain set is constrained, a projection step needs to be added (see Algorithms 4 & 5). Note that $\Pi_{\mathcal{X}}(\mathbf{y}) = \text{argmin}_{\mathbf{x}\in\mathcal{X}} \|\mathbf{y}-\mathbf{x}\|$ and $\Pi_{\mathcal{X}}^{\mathbf{A}}(\mathbf{y}) = \text{argmin}_{\mathbf{x}\in\mathcal{X}} \|\mathbf{y}-\mathbf{x}\|_{\mathbf{A}}$, where $\|\mathbf{x}\|_{\mathbf{A}} = \sqrt{\mathbf{x}^\top \mathbf{A}\mathbf{x}}$ for a positive-definite matrix $\mathbf{A}$.

Algorithm 4 Online Preconditioned Gradient Descent (OPGD) for Constrained Setup

---

1: **Require:** Initial vector $\mathbf{x}_1 \in \mathcal{X}$, step size η, a sequence of positive definite matrices $\mathbf{A}_t$
2: **for** $t = 1, 2, \ldots, T$ **do**
3: Play $\mathbf{x}_t$
4: Observe the gradient of the current action $\nabla f_t(\mathbf{x}_t)$
5: $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}^{\mathbf{A}_t}\left(\mathbf{x}_t - \eta \mathbf{A}_t^{-1}\nabla f_t(\mathbf{x}_t)\right)$
6: **end for**

Theorem 7. *(The constrained version of Theorem 1) Suppose that for any $t \in [T]$:*

1. *The function $f_t : \mathcal{X} \to \mathrm{R}$ is μ-strongly convex and L-smooth over $\mathcal{X}$.*
2. *The preconditioning matrix $\mathbf{A}_t$ satisfies $\lambda'\cdot\mathbf{I} \preceq \mathbf{A}_t \preceq \lambda\cdot\mathbf{I}$.*
3. *The condition number satisfies $\frac{\lambda}{\lambda'} < 1 + \frac{\mu^2}{4L^2}$.*

If we set $\eta = \frac{\lambda'\mu}{2L^2}$, the dynamic regret for the sequence of actions $\mathbf{x}_t$ generated by Algorithm 4 is bounded as follows

$$\begin{aligned} \mathbf{Reg}_T^d \leq\ & \frac{L^2}{\mu}\left(\frac{4L^2\lambda - \mu^2\lambda'}{\mu^2\lambda' - 4L^2(\lambda-\lambda')}\right)\sum_{t=2}^{T+1}\left\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\right\|^2 \\ & + \left(\frac{L^2\lambda}{\lambda'\mu} - \frac{\mu}{4}\right)\|\mathbf{x}_1 - \mathbf{x}_1^*\|^2 + \frac{\mu D}{2L}\sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t^*)\| \\ & + \frac{\mu}{4L^2}\sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t^*)\|^2, \end{aligned}$$

where $D = \max_{\mathbf{x},\mathbf{y}\in\mathcal{X}} \|\mathbf{x}-\mathbf{y}\|$.

Theorem 8. *Suppose that for any $t \in [T]$, the function $f_t : \mathcal{X} \to \mathrm{R}$ is μ-strongly convex and L-smooth. For any $0 < \eta \leq \frac{1}{L}$, the dynamic regret of OMGD is bounded as follows*

$$\begin{aligned} &\mathbf{Reg}_T^d \\ &\leq \min\begin{cases} (1) & (f_1(\mathbf{x}_1) - f_1(\mathbf{x}_1^*)) + 2V_T + \frac{\pi^2 D^2 L}{12} + \\ & D\sum_{t=2}^{T}\left\|\nabla f_{t-1}(\mathbf{x}_{t-1}^*)\right\| \\ (2) & \frac{L}{2}\left[\|\mathbf{x}_1 - \mathbf{x}_1^*\|^2 + 2D^2(\frac{\pi^2}{6}) + 2C_{2,T}^*\right] + \\ & D\sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t^*)\| \end{cases}, \end{aligned} \tag{12}$$

where $D \triangleq \max_{\mathbf{x},\mathbf{y}\in\mathcal{X}} \|\mathbf{x}-\mathbf{y}\|$.

For Theorems 7 & 8, if $\mathbf{x}_t^*$ is inside the relative interior of $\mathcal{X}$ (i.e., $\nabla f_t(\mathbf{x}_t^*) = 0$) for any $t \in [T]$, the theoretical results of the unconstrained case can be recovered.

Algorithm 5 Online Multiple Gradient Descent (OMGD) (Zhang et al. 2017) for Constrained Setup

---

1: **Require:** Initial vector $\mathbf{x}_1 \in \mathcal{X}$, step size η, function parameters μ and L.
2: **for** $t = 1, 2, \ldots, T$ **do**
3: Play $\mathbf{x}_t$
4: Receive the information of f_t
5: Let $\mathbf{z}_{t+1}^{(0)} = \mathbf{x}_t$ and $K_t = \lceil \frac{-2\log(t)}{\log(1-\frac{2\mu}{1/\eta+\mu})} \rceil$
6: **for** $j = 1, 2, \ldots, K_t$ **do**
7: $\mathbf{z}_{t+1}^{(j)} = \Pi_{\mathcal{X}}\left(\mathbf{z}_{t+1}^{(j-1)} - \eta\nabla f_t(\mathbf{z}_{t+1}^{(j-1)})\right)$
8: **end for**
9: $\mathbf{x}_{t+1} = \mathbf{z}_{t+1}^{(K_t)}$
10: **end for**

---

6 Experimental Results

(Zhang et al. 2017) show that the regret bound $O(C_{2,T}^*)$ is optimal in the worst-case. In Theorem 4, we prove that if well-predicted function information is applied, the bound $O(C_{2,T}^*)$ is possibly improved to $O(C_{4,T}^*)$ with OON. We now provide the simulation results verifying this property. Consider a function sequence of the following form

$$f_t(\mathbf{x}) = \left((\mathbf{x}-\mathbf{x}_t^*)^\top \mathbf{Q}_t(\mathbf{x}-\mathbf{x}_t^*)\right)^2 + \frac{1}{2}(\mathbf{x}-\mathbf{x}_t^*)^\top \mathbf{Q}_t(\mathbf{x}-\mathbf{x}_t^*),$$

where $\mathbf{Q}_t$ is a positive definite matrix and $\alpha\mathbf{I} \preceq \mathbf{Q}_t \preceq \beta\mathbf{I}$ ($\alpha = 1$ and $\beta = 30$). Since OON operates in a local sense, we just need an L_H for which Assumption 2 in Theorem 4 holds, when executing step 5 in Algorithm 2.

The hessian of $f_t(\cdot)$ can be written as follows

$$\begin{aligned} \nabla^2 f_t(\mathbf{x}) =& 4(\mathbf{x}-\mathbf{x}_t^*)^\top \mathbf{Q}_t(\mathbf{x}-\mathbf{x}_t^*)\mathbf{Q}_t \\ & + 8\mathbf{Q}_t(\mathbf{x}-\mathbf{x}_t^*)(\mathbf{x}-\mathbf{x}_t^*)^\top \mathbf{Q}_t + \mathbf{Q}_t, \end{aligned}$$

which implies $f_t(\cdot)$ is α-strongly convex $\forall t$. We now discuss the choice of L_H to ensure Assumptions 2 in Theorem 4

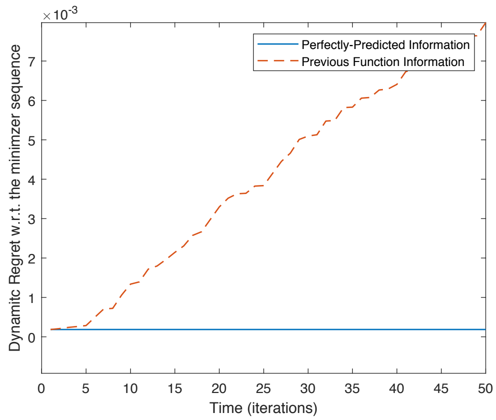

Figure 1: Dynamic regrets of OON using perfectly-predicted and stale function information.

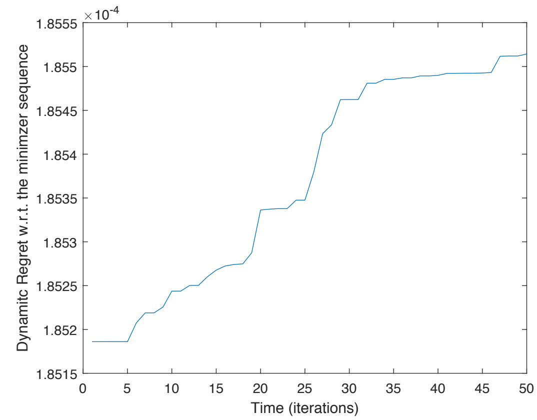

Figure 2: Dynamic regrets of OON using perfectly-predicted function information.

hold. Based on the hessian expression, we have

$$\begin{aligned}
&\|\mathbf{H}_t(\hat{\mathbf{x}}_{t-1}) - \mathbf{H}_t(\mathbf{x}_t^*)\| \\
\leq &\left\| 4(\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*)^\top \mathbf{Q}_t(\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*)\mathbf{Q}_t \right\| \\
+ &\left\| 8\mathbf{Q}_t(\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*)(\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*)^\top \mathbf{Q}_t \right\| \\
\leq &12\left\|\mathbf{Q}_t\right\|^2 \left\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*\right\|^2 \leq 12\beta^2 \left\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*\right\|^2 \\
\leq &12\beta^2 \frac{\alpha}{L_H} \left\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*\right\|,
\end{aligned}$$

where the last inequality comes from $\|\hat{\mathbf{x}}_{t-1} - \mathbf{x}_t^*\| \leq \frac{\alpha}{L_H}$ (see the supplementary material). Let $L_H = 12\beta^2 \frac{\alpha}{L_H}$, Assumption 2 holds for iterate $\hat{\mathbf{x}}_t$. Knowing that $L_H = \sqrt{12\beta^2\alpha}$, we have

$$\|\mathbf{x}_1 - \mathbf{x}_1^*\| = \|\hat{\mathbf{x}}_0 - \mathbf{x}_1^*\| \leq \frac{\alpha}{L_H} = \sqrt{\frac{\alpha}{12\beta^2}},$$

which defines the range to generate $\mathbf{x}_1$.

To let the function sequence evolve in an adversarial way, the optimal point of the next function is randomly selected from the sphere centered at the current optimal point with radius $\frac{\alpha}{L_H}$.

In this experiment, we compare the results of two strategies: one uses the perfectly-predicted information $((\mathbf{M}_t, \mathbf{m}_t) = (\mathbf{H}_t, \nabla f_t))$, and the other one uses the information of the previous function $((\mathbf{M}_t, \mathbf{m}_t) = (\mathbf{H}_{t-1}, \nabla f_{t-1}))$. Theoretically, the regret of the first strategy is upper bounded by $O(C_{4,T}^*)$ and the regret of the second one is bounded by $O(C_{2,T}^*)$ (see Corollary 5). In Fig. 1, we can see that the regret incurred using the stale function information grows much faster than the one using the perfectly-predicted information, which verifies the theoretical advantage of OON.

7 Concluding Remarks

In this paper, we revisited the dynamic regret for online optimization, where the function sequence is strongly convex and smooth. We first proposed online preconditioned gradient descent (OPGD), which achieves the optimal regret bound of $O\big(\min\{C_T^*, C_{2,T}^*\}\big)$ with one gradient query per round. Next, we developed optimistic online Newton (OON) method, which uses predictions of Hessians and gradients in the update process. We proved a (local) dynamic regret bound scaling as $O(D_T' + C_{4,T}^*)$, where D_T' measures the dissimilarity between the predicted information and the true information. If D_T' is small, this algorithm provides an improvement over the regret rate of $O(C_{2,T}^*)$. It is also intriguing to see if the idea of optimistic online learning can be extended to quasi Newton like methods and get improved bounds in terms of D_T when only predictable gradients are available. We further verified that a conservative learner, that uses stale (previous round) information, always incurs a regret that is no worse than $O(C_{2,T}^*)$, recovering the optimal worst-case. We finally revisited the online multiple gradient descent (OMGD) algorithm of (Zhang et al. 2017) and provided complementary analysis that shows the regret rate of $O(\min\{V_T, C_{2,T}^*\})$ for OMGD. The main benefit of the bound is that V_T and $C_{2,T}^*$ are not comparable, and including V_T in the regret bound can make it tighter.

Our results mainly provide the following insight about dynamic online learning. It is unlikely that the smoothness assumption alone can provide improved regret rates over convex setting. In classical optimization with *time-invariant* functions, smoothness does provide faster convergence rates for optimization algorithms (e.g., gradient descent). Whereas in dynamic online setting, since the adversary can change the function sequence drastically, the action $\mathbf{x}_t$ depending on the information of $\{f_s\}_{s=1}^{t-1}$ may suffer a large cost. Therefore, the assumption of strong convexity comes to rescue by translating the closeness of objective values to the closeness of actions. As a result, with smoothness and strong convexity together, dynamic regret bounds can be considerably better than the convex case. Given the potential improvement of OON over OGD, it would also be interesting to see whether we can leverage predicted higher-order information to further improve the dynamic regret bound.

Broader Impact

This paper should be of interest to the online learning and optimization community. We do not anticipate any future societal consequences as this work contributes to theoretical online optimization.

Acknowledgments

The authors gratefully acknowledge the support of NSF ECCS-1933878 Award.

References

Ajalloeian, A.; Simonetto, A.; and Dall'Anese, E. 2020. Inexact online proximal-gradient method for time-varying convex optimization. In *American Control Conference (ACC)*, 2850–2857.

Besbes, O.; Gur, Y.; and Zeevi, A. 2015. Non-stationary stochastic optimization. *Operations research* 63(5): 1227–1244.

Chiang, C.-K.; Yang, T.; Lee, C.-J.; Mahdavi, M.; Lu, C.-J.; Jin, R.; and Zhu, S. 2012. Online optimization with gradual variations. In *Conference on Learning Theory*, 6–1.

Daniely, A.; Gonen, A.; and Shalev-Shwartz, S. 2015. Strongly adaptive online learning. In *International Conference on Machine Learning*, 1405–1411.

Hall, E. C.; and Willett, R. M. 2015. Online convex optimization in dynamic environments. *IEEE Journal of Selected Topics in Signal Processing* 9(4): 647–662.

Hazan, E. 2016. Introduction to Online Convex Optimization. *Foundations and Trends in Optimization* 2(3-4): 157–325.

Hazan, E.; Agarwal, A.; and Kale, S. 2007. Logarithmic regret algorithms for online convex optimization. *Machine Learning* 69(2-3): 169–192.

Hazan, E.; and Seshadhri, C. 2007. Adaptive algorithms for online decision problems. In *Electronic colloquium on computational complexity (ECCC)*, volume 14.

Jadbabaie, A.; Rakhlin, A.; Shahrampour, S.; and Sridharan, K. 2015. Online optimization: Competing with dynamic comparators. In *Artificial Intelligence and Statistics*, 398–406.

Lesage-Landry, A.; Taylor, J. A.; and Shames, I. 2020. Second-order Online Nonconvex Optimization. *arXiv preprint arXiv:2001.10114* .

Mokhtari, A.; Shahrampour, S.; Jadbabaie, A.; and Ribeiro, A. 2016. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *IEEE 55th Conference on Decision and Control (CDC)*, 7195–7201.

Nesterov, Y. 1998. Introductory lectures on convex programming volume i: Basic course. *Lecture notes* 3(4): 5.

Rakhlin, A.; and Sridharan, K. 2013a. Online learning with predictable sequences. In *Conference on Learning Theory*, 993–1019. PMLR.

Rakhlin, S.; and Sridharan, K. 2013b. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, 3066–3074.

Ravier, R. J.; Calderbank, A. R.; and Tarokh, V. 2019. Prediction in online convex optimization for parametrizable objective functions. In *IEEE 58th Conference on Decision and Control (CDC)*, 2455–2460.

Yuan, J.; and Lamperski, A. 2020. Trading-off static and dynamic regret in online least-squares and beyond. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 6712–6719.

Zhang, L.; Liu, T.-Y.; and Zhou, Z.-H. 2019. Adaptive Regret of Convex and Smooth Functions. In *International Conference on Machine Learning*, 7414–7423.

Zhang, L.; Lu, S.; and Yang, T. 2020. Minimizing dynamic regret and adaptive regret simultaneously. In *International Conference on Artificial Intelligence and Statistics*, 309–319. PMLR.

Zhang, L.; Lu, S.; and Zhou, Z.-H. 2018. Adaptive online learning in dynamic environments. In *Advances in neural information processing systems*, 1323–1333.

Zhang, L.; Yang, T.; Yi, J.; Jin, R.; and Zhou, Z.-H. 2017. Improved dynamic regret for non-degenerate functions. In *Advances in Neural Information Processing Systems*, 732–741.

Zhang, L.; Yang, T.; Zhou, Z.-H.; et al. 2018. Dynamic Regret of Strongly Adaptive Methods. In *International Conference on Machine Learning*, 5882–5891.

Zinkevich, M. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, 928–936.