

Dynamic Regret of Online Mirror Descent for Relatively Smooth Convex Cost Functions

Nima Eshraghi, *Student member, IEEE*, and Ben Liang, *Fellow, IEEE*

Abstract—The performance of online convex optimization algorithms in a dynamic environment is often expressed in terms of the dynamic regret, which measures the decision maker’s performance against a sequence of time-varying comparators. In the analysis of the dynamic regret, prior works often assume Lipschitz continuity or uniform smoothness of the cost functions. However, there are many important cost functions in practice that do not satisfy these conditions. In such cases, prior analyses are not applicable and fail to guarantee the optimization performance. In this letter, we show that it is possible to bound the dynamic regret, even when neither Lipschitz continuity nor uniform smoothness is present. We adopt the notion of relative smoothness with respect to some user-defined regularization function, which is a much milder requirement on the cost functions. We first show that under relative smoothness, the dynamic regret has an upper bound based on the path length and functional variation. We then show that with an additional condition of relatively strong convexity, the dynamic regret can be bounded by the path length and gradient variation. These regret bounds provide performance guarantees to a wide variety of online optimization problems that arise in different application domains. Finally, we present numerical experiments that demonstrate the advantage of adopting a regularization function under which the cost functions are relatively smooth.

Index Terms—Optimization algorithms, online optimization, time-varying systems, first-order methods

I. INTRODUCTION

WE STUDY the problem of online convex optimization, which can be modeled by a sequential decision-making process over a finite number of rounds T . In every round $t \in \{1, 2, \dots, T\}$, the decision maker chooses a point x_t from a convex set \mathcal{X} based on the information from previous rounds. Then, the convex cost function $f_t(x) : \mathcal{X} \rightarrow \mathbb{R}$ is revealed to the decision maker, which suffers the corresponding cost $f_t(x_t)$. The goal of the decision maker is to minimize some regret, which is the difference between the cumulative cost of the decision maker and that of an optimal fixed point chosen in hindsight with knowledge of $f_t(x)$ for all rounds. Most early works consider the static regret:

$$\text{Reg}_T^s = \sum_{t=1}^T f_t(x_t) - \arg\min_{x \in \mathcal{X}} \sum_{t=1}^T f_t(x). \quad (1)$$

N. Eshraghi and B. Liang are with the Department of Electrical and Computer Engineering, University of Toronto, Ontario, Canada. (e-mail: {neshraghi, liang}@ece.utoronto.ca). This work was funded in part by the Natural Sciences and Engineering Research Council of Canada.

The benchmark variable in (1) is a *static* point that resides in the feasible set \mathcal{X} . In the literature of online learning and control, there are many algorithms that guarantee sublinear upper bounds on the static regret under various settings [1]–[4].

Such a static metric can accurately reflect the decision maker’s performance as long as the static benchmark performs consistently well over different rounds. However, this may not always hold in a dynamic environment, where the sequence of cost functions are time-varying. Thus, the *dynamic regret* has been proposed as a more stringent metric that measures the algorithm performance against a time-varying sequence. Previous studies often focus on the the sequence of minimizers in their regret analysis [5]–[9]. In this work, we target a more general form of the dynamic regret that allows comparison against an arbitrary comparator sequence $\{u_1, u_2, \dots, u_T\}$ [3]:

$$\text{Reg}_T^d = \sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(u_t). \quad (2)$$

Since the online cost functions can fluctuate arbitrarily, obtaining a sublinear upper bound on the dynamic regret may not be possible. Therefore, the dynamic regret is usually expressed in terms of certain *regularity measures* of the comparators or the cost function sequence. The regularity measures reflect how fast an environment evolves as time progresses. Previous works [3], [6]–[16] have utilized different regularity measures to bound the dynamic regret, namely the path length, gradient variation, and functional variation. The path length of an arbitrary sequence $\{u_1, u_2, \dots, u_T\}$ is defined as

$$C_T = \sum_{t=1}^T \|u_{t+1} - u_t\|, \quad (3)$$

which measures the variation in the comparator sequence, where $\|\cdot\|$ could be any norm. Several online learning algorithms provide an $O(\sqrt{T}(1 + C_T))$ upper bound on the dynamic regret of convex cost functions [3], [10], [11], which can be improved to $O(\sqrt{TC_T})$ when prior knowledge of C_T and T is available [12]. The path length has also been recently used in the study of online convex optimization with constraint violation [13], where upper bounds of $O(\sqrt{T(1 + C_T)})$ and $O(\sqrt{T})$ are derived on the dynamic regret and cumulative constraint violation, respectively. Furthermore, the dynamic regret can also be bounded by both the path length and its squared form using multiple mirror descent steps per online round [14]. In addition to the path length, gradient variation is another regularity measure often utilized in the online learning

literature [16]–[18]. It reflects how fast the gradient of the online functions changes. Here, we use a non-squared version of the gradient variation, which is defined as

$$G_T = \sum_{t=1}^T \sup_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|_*, \quad (4)$$

where $\|x\|_* = \sup_y \{x^T y \mid \|y\| \leq 1\}$ represents the dual norm. Another regularity measure related to the sequence of online functions is functional variation [9], [15], [16], defined as

$$V_T = \sum_{t=1}^T \sup_{x \in \mathcal{X}} |f_{t+1}(x) - f_t(x)|, \quad (5)$$

which collects the variation in the cost functions over time.

In the analysis of dynamic regret, prior works often impose assumptions such as Lipschitz continuity and/or uniform smoothness [3], [6]–[13], [15], [16]. However, some cost functions that arise in well-known applications do not satisfy these conditions, e.g., the Poisson inverse problem, the D-optimal design problem, and support vector machine training. To address this limitation, a generalization of these traditional assumptions is required. In this work, we remove the restrictive assumptions of Lipschitz continuity and uniform smoothness.

Instead, we adopt the notion of relative smoothness, which allows measuring the cost functions behavior relative to a user-specific function [19]. Relative smoothness generalizes the traditional form of smoothness used in earlier analyses and provides more flexibility in functions. Consequently, it is applicable to a broader class of cost functions. Such generalization has been proposed recently in the study of *offline* optimization methods based on mirror descent [19]–[22]. However, this is not applicable to many systems in practice since they are often time-varying, requiring an online solution. Performance analysis in the online setting is more challenging due to the appearance of some dynamic terms that require careful handling to reflect the speed of the changes in the problem environment. There is no prior work to study the feasibility of generalization to relative smoothness in the online setting.¹

As far as we are aware, this is the first study to bound the dynamic regret of online mirror descent for relatively smooth convex functions. First, we show that the dynamic regret has an upper bound of $O(1 + C_T + V_T)$. This compares favorably to the closest related works in [10], [11], which obtain the upper bound of $O(\sqrt{T}(1 + C_T))$ for online mirror descent under the Lipschitz continuity assumption. Besides removing the requirement of Lipschitz continuity, our bound further reduces the dependence on T and can be much smaller especially when the problem environment does not drift too fast. Second, we show that when the cost functions are in addition relatively strongly convex, the dynamic regret can be further tightened to $O(1 + C_T + \min(V_T, G_T))$. Thus, our results show that even when the cost functions are not Lipschitz continuous

¹In the online setting, the recent work of [23] studies the static regret without Lipschitz continuity. They introduce the notion of Riemann-Lipschitz continuity, to bound the static regret of online mirror descent under this assumption. However, as discussed above, the dynamic regret often is a more suitable measure of performance.

or uniformly smooth, it is still possible to guarantee performance in terms of the dynamic regret, by leveraging relative smoothness. Finally, our numerical experiments demonstrate substantial improvement in the performance of online mirror descent when the regularization function is chosen to provide relative smoothness.

II. DYNAMIC REGRET OF ONLINE MIRROR DESCENT

We consider the standard problem of online optimization with respect to a sequence of convex cost functions over a finite number of rounds, denoted by T . At the beginning of every round t , the decision maker submits a decision represented by x_t , which is taken from a convex and compact set \mathcal{X} . Then, the cost function of the current round $f_t(\cdot)$ is revealed, and the decision maker becomes aware that it has suffered the corresponding cost $f_t(x_t)$. The decision maker then updates its decision in the next round.

The mirror descent algorithm is a classical method for convex optimization problems. An appealing feature of mirror descent is the extension of the projection step using the Bregman divergence instead of the Euclidean distance, which makes the algorithm applicable to a wide range of problems. The Bregman divergence with respect to the function $r(\cdot)$ is defined as

$$D_r(x, y) = r(x) - r(y) - \langle \nabla r(y), x - y \rangle.$$

The Bregman divergence is a general way of measuring the distance through the lens of function $r(\cdot)$. An instructive example is the Bregman divergence associated with the squared ℓ_2 -norm, i.e., $r(x) = \frac{1}{2}\|x\|_2^2$. In this case, the Bregman divergence reduces to the Euclidean distance. In another example, the Bregman divergence corresponding to the choice of $r(x) = \sum_{i=1}^d x_i \log(x_i) - x_i$ on the d -dimensional probability simplex recovers the Kullback-Leibler divergence. Many other interesting examples are given in [22].

The online version of mirror descent [4] is given by

$$x_{t+1} = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle \nabla f_t(x_t), x \rangle + \frac{1}{\eta_t} D_r(x, x_t) \right\} \quad (6)$$

where η_t is the step size, and $D_r(x, y)$ is the Bregman divergence corresponding to the function $r(\cdot)$. In the context of mirror descent, $r(\cdot)$ is often called the *regularization function*. The term with the Bregman divergence helps to limit the changes in the sequence of decisions from one round to the next. In particular, the update in (6) suggests that the decision maker aims to stay close to the current decision x_t as measured by the Bregman divergence, while taking a step in a direction close to the negative gradient to reduce the current cost at round t . We make the following assumption, which is common in the study of online mirror descent [16], [24]:

Assumption 1. The Bregman divergence satisfies a Lipschitz condition of the form

$$|D_r(x, z) - D_r(y, z)| \leq \gamma \|x - y\|, \quad \forall x, y, z \in \mathcal{X}, \quad (7)$$

where γ is a positive constant.

Note that when the function $r(\cdot)$ is Lipschitz continuous on the feasibility domain, the Lipschitz condition on the Bregman divergence is automatically satisfied.

In this work, we do not require the cost functions to be Lipschitz continuous while the condition stated in (7) serves as its replacement. In online optimization, the cost functions are revealed to the decision maker sequentially over time, and we have no control over them. In contrast, we have control over the choice of the regularization function, which can benefit from a careful design to allow properties such as Lipschitz continuity. Therefore, the condition in (7) is indeed a milder requirement compared with the assumption of Lipschitz continuity of the cost functions.

A. Relatively Smooth Cost Functions

In this section, we consider relatively smooth cost functions that may not be Lipschitz continuous or uniformly smooth. The notion of relative smoothness is proposed in [19], which measures the smoothness relative to a user-specific function. Thus, it does not require the specification of any norm. We provide below the definitions of uniform smoothness and relative smoothness.

Definition 1: A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is uniformly smooth with modulus β , i.e., β -smooth, with respect to some norm $\|\cdot\|$, if there exists a positive constant β such that

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\beta}{2} \|y - x\|^2, \quad \forall x, y \in \mathcal{X}.$$

An equivalent definition is that $\nabla f(\cdot)$ is Lipschitz continuous, i.e., $\|\nabla f(x) - \nabla f(y)\| \leq \beta \|x - y\|$, $\forall x, y \in \mathcal{X}$.

Definition 2: A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is β -smooth relative to $r(\cdot)$, if there exists a positive constant β such that

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \beta D_r(y, x), \quad \forall x, y \in \mathcal{X}.$$

We note that the relative smoothness in Definition 2 generalizes the uniform smoothness in Definition 1, which is commonly assumed in the literature of online convex optimization [8], [9], [18], [25]. In Definition 2, the Bregman divergence naturally serves as a distance measure. It replaces the norm squared in Definition 1. Therefore, the smoothness of $f(\cdot)$ does not depend on any norm and is instead measured with respect to the function $r(\cdot)$. In particular, by setting $r(x) = \frac{1}{2} \|x\|^2$, relative smoothness specializes to the uniform smoothness. Furthermore, from *Proposition 1.1* in [19], an equivalent form of Definition 2 is

$$\nabla^2 f(x) \preceq \beta \nabla^2 r(x), \quad \forall x \in \mathcal{X}, \quad (8)$$

which establishes a simple condition on the Hessian matrices of the two functions.

We make the following assumption for the analysis in this subsection.

Assumption 2. The cost functions $f_t(\cdot)$ are convex and β -smooth relative to the regularization function $r(\cdot)$.

As an example, the D-optimal design problem [19] satisfies the condition stated in Assumption 2, but it is neither Lipschitz continuous nor uniformly smooth. There are many other important relatively smooth functions that arise in various application domains, such as the Poisson inverse problem [20] and minimum-volume covering ellipsoid [19], [22]. A systematic way of choosing a proper function $r(\cdot)$ is presented

in [19] for any objective function whose norm of subgradients are bounded by a polynomial in either $\sum_{i=1}^d \frac{1}{x_i}$ or $\|x\|_2$. It is a useful construction to reveal the relative smoothness of a wide range of cost functions.

We are now ready to upper bound the dynamic regret of online mirror descent under relative smoothness.

Theorem 1: Under Assumptions 1 and 2, the dynamic regret of online mirror descent with fixed step size $\eta = 1/\beta$ satisfies

$$\sum_{t=1}^T \left(f_t(x_t) - f_t(u_t) \right) \leq \beta R + f_1(x_1) - f_{T+1}(x_{T+1}) + \gamma \beta C_T + V_T,$$

for any feasible sequence $\{u_1, u_2, \dots, u_T\}$, where $R = \max_{x, y \in \mathcal{X}} D_r(x, y)$, γ is the Lipschitz constant associated with the Bregman divergence, and C_T and V_T denote the path length and functional variation, as defined in (3) and (5), respectively.

The proof of Theorem 1 is given in App. IV-B.

Remark 1. The dynamic regret of standard online mirror descent under uniform smoothness has not been studied in prior works. However, the works [10], [11] have shown that for convex and Lipschitz continuous cost functions, the dynamic regret of online mirror descent is bounded by $O(\sqrt{T}(1+C_T))$. Theorem 1 shows that even when the Lipschitz continuity requirement is replaced by relative smoothness, the dynamic regret can still be upper bounded. Furthermore, the new $O(1+C_T+V_T)$ bound removes the dependency on \sqrt{T} and relates only to the regularity measures C_T and V_T .

B. Relatively Smooth and Strongly Convex Cost Functions

In this part, we consider cost functions that are relatively smooth and strongly convex with respect to the same function $r(\cdot)$. We provide the formal definitions of uniform strong convexity and relative strong convexity below:

Definition 3: A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is uniformly strongly convex with modulus λ , i.e., λ -strongly convex, with respect to some norm $\|\cdot\|$, if there exists a positive constant λ such that

$$f(x) + \langle \nabla f(x), y - x \rangle + \frac{\lambda}{2} \|y - x\|^2 \leq f(y), \quad \forall x, y \in \mathcal{X}.$$

Definition 4: A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is λ -strongly convex relative to $r(\cdot)$, if there exists a positive constant λ such that

$$f(x) + \langle \nabla f(x), y - x \rangle + \lambda D_r(y, x) \leq f(y), \quad \forall x, y \in \mathcal{X}.$$

We make the following assumption in this subsection.

Assumption 3. The cost functions $f_t(\cdot)$ are β -smooth and λ -strongly convex, both relative to the regularization function $r(\cdot)$. Furthermore, $r(\cdot)$ is 1-strongly convex with respect to some norm $\|\cdot\|$.

We note that the conditions stated in Assumption 3 imply $\lambda \leq \beta$. Furthermore, since $r(\cdot)$ is strongly convex with respect to a norm $\|\cdot\|$, Assumption 3 implies the strong convexity of $f_t(\cdot)$ with respect to the same norm. We also note that the strong convexity of the regularization function is a standard

assumption, commonly used in the analysis of online mirror descent [4], [16], [26].

Theorem 2: Under Assumptions 1 and 3, the dynamic regret of online mirror descent with fixed step size $\eta = 1/\beta$ satisfies

$$\sum_{t=1}^T \left(f_t(x_t) - f_t(u_t) \right) \leq (\beta - \lambda)D_r(u_1, x_0) + (\beta - \lambda)\gamma C_T + 2MG_T,$$

for any feasible sequence $\{u_1, u_2, \dots, u_T\}$, where M represents the diameter of the feasible set, i.e., $M = \max_{x \in \mathcal{X}} \|x\|$, and C_T and G_T denote the path length and gradient variation, as defined in (3) and (4), respectively.

The proof of Theorem 2 is given in App. IV-C.

Theorem 2 states that the dynamic regret of online mirror descent is upper bounded by $O(1 + C_T + G_T)$. Together with Theorem 1 this immediately leads to the following result.

Corollary 3: Under the conditions stated in Theorem 2, the dynamic regret of online mirror descent has an upper bound of $O(1 + C_T + \min(V_T, G_T))$.

Remark 2. The regularity measures V_T and G_T represent different aspects of an online learning problem, i.e., variation in the functions and gradients. Each of these quantities can be small in an environment that does not change too fast. The resultant bound of $O(1 + C_T + \min(V_T, G_T))$ combines the advantage of these two regularity measures.

Remark 3. When the comparator sequence is fixed over time, i.e., $u_1 = \dots = u_T = \operatorname{argmin}_{x \in \mathcal{X}} \sum_{t=1}^T f_t(x)$, Theorems 1 and 2 also bound the static regret. In this case, the term involving the path length C_T disappears, and we obtain static regret bounds of $O(1 + V_T)$ and $O(1 + \min(V_T, G_T))$.

III. NUMERICAL EXPERIMENTS

In this section, we present numerical examples to demonstrate the performance of online mirror descent on cost functions that arise in practice, which are relatively smooth with respect to a carefully chosen regularization function.

In the first experiment, we proceed with an application of our algorithmic results to a broad class of D-optimal design problems. The cost functions of interest are $f_t(x) = -\ln \det(H_t \mathbf{D}(x) H_t^T)$, where $\mathbf{D}(x) = \operatorname{diag}(x)$, and $H_t \in \mathbb{R}^{m \times d}$. We note that $f_t(x)$ is neither Lipschitz nor uniformly smooth, but with respect to the Burg regularization function $r(x) = -\sum_{i=1}^d \ln(x_i)$ [20], it is relatively smooth. In our experiment, we consider the simplex feasible set, and set $m = 5$ and $d = 10$. In every round t , the matrix H_t is selected from a set of randomly generated matrices with independent entries distributed uniformly in $[0, 1]$. Fig. 1 shows the accumulated cost of online mirror descent with the Burg regularization function $r(x)$, as well as two of the most commonly used alternatives, namely the ℓ_1 -squared and KL regularization functions. The figure highlights that setting $r(x)$ so that the cost functions enjoy relative smoothness, results in substantially lower accumulated cost compared with the other common choices of regularization function.

Next, we study the application of our results to the Poisson linear inverse problem, which arises in the domain of image science. The cost functions are $f_t(x) =$




Fig. 1. Accumulated cost of D-optimal design problem for online mirror descent with the Burg, ℓ_1 -squared, and KL regularization functions.

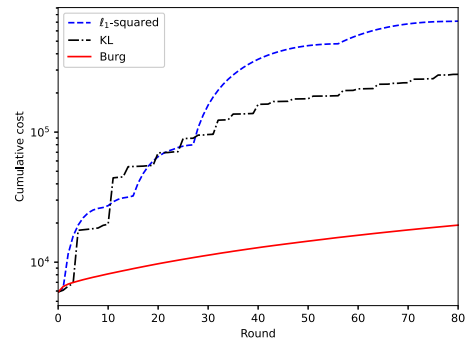


Fig. 2. Accumulated cost of Poisson linear inverse problem for online mirror descent with the Burg, ℓ_1 -squared, and KL regularization functions.

$\sum_{i=1}^m \left(b_{t,i} \log\left(\frac{b_{t,i}}{(A_t x)_i}\right) + (A_t x)_i - b_{t,i} \right)$, where $A_t \in \mathbb{R}_+^{m \times d}$ models some experimental protocol, and $b_t \in \mathbb{R}_+^m$ is the vector of measurements at round t . The goal is to reconstruct the signal $x \in \mathbb{R}_+^d$ from the measurements b_t such that $A_t x \simeq b_t$. We randomly generate A_t and b_t from uniform distribution in $[0, 1]$ and set $m = 1500$ and $d = 10$. Since the gradient of $f_t(x)$ is in the order $O(1/x)$, its norm cannot be bounded by a constant, so $f_t(x)$ is neither Lipschitz nor uniformly smooth. However, it can be shown that it is relatively smooth with respect to the Burg regularization function. Fig. 2 shows the accumulated cost versus the number of rounds for online mirror descent with the Burg, ℓ_1 -squared, and KL regularization functions. We again observe that Burg regularization performs significantly better than the two commonly used alternatives.

IV. CONCLUSION

In this letter, we provide a new analysis on the dynamic regret of online convex optimization with mirror descent under relative smoothness. The cost functions do not need to be Lipschitz continuous or uniformly smooth. When the cost functions are relatively smooth, we show that the dynamic regret is bounded by $O(1 + C_T + V_T)$, which depends only on the path length C_T and functional variation V_T . In addition, when the cost functions are also relatively strongly convex, we show that the dynamic regret bound can be tightened

to $O(1 + C_T + \min(V_T, G_T))$. A main observation in these results is that the Hessians of the cost functions can be upper bounded by the Hessian of a carefully designed regularization function. Our numerical experiments show significant gain in the performance of online mirror descent after choosing an appropriate regularization function with respect to which the cost functions are relatively smooth.

APPENDIX

A. Helper Lemmas

Lemma 4: For any $x, y, z \in \mathcal{X}$,

$$\langle \nabla r(z) - \nabla r(y), x - y \rangle = D_r(x, y) - D_r(x, z) + D_r(y, z).$$

The proof of Lemma 4 is given in earlier studies [26], [27].

Lemma 5: Let the function $f(\cdot)$ be smooth relative to the function $r(\cdot)$. Then, the following inequality holds for any $x, y, z \in \mathcal{X}$:

$$f(x) - f(y) \leq \langle \nabla f(z), x - y \rangle + \beta D_r(x, z).$$

The proof of Lemma 5 is given in [22].

Lemma 6: Let \mathcal{X} be a convex set and $r(\cdot)$ be a 1-strongly convex function on \mathcal{X} with respect to some norm. Then, any update of the form

$$x^* = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle a, x \rangle + D_r(x, c) \right\},$$

satisfies the following inequality:

$$\langle x^* - u, a \rangle \leq D_r(u, c) - D_r(u, x^*) - D_r(x^*, c), \forall u \in \mathcal{X}.$$

The proof of Lemma 6 is given in [24].

B. Proof of Theorem 1

By the optimality condition of the update (6), for any $u \in \mathcal{X}$ we have

$$\begin{aligned} 0 &\leq \langle \eta \nabla f_t(x_t) + \nabla r(x_{t+1}) - \nabla r(x_t), u - x_{t+1} \rangle \\ &= D_r(u, x_t) - D_r(u, x_{t+1}) - D_r(x_{t+1}, x_t) \\ &\quad + \eta \langle \nabla f_t(x_t), u - x_{t+1} \rangle, \end{aligned} \quad (9)$$

where the equality follows from Lemma 4. On the other hand, by applying Lemma 5 we obtain

$$\begin{aligned} \eta \left(f_t(x_{t+1}) - f_t(u) \right) &\leq \\ &\eta \left(\langle \nabla f_t(x_t), x_{t+1} - u \rangle + \beta D_r(x_{t+1}, x_t) \right). \end{aligned} \quad (10)$$

By combining (9) and (10), we have

$$\begin{aligned} \eta \left(f_t(x_{t+1}) - f_t(u) \right) &\leq \\ D_r(u, x_t) - D_r(u, x_{t+1}) - (1 - \beta\eta) D_r(x_{t+1}, x_t). \end{aligned} \quad (11)$$

We then set $\eta = 1/\beta$, so that the last term in the above inequality disappears. We add $f_t(x_t)$ to both sides of (11) and set $u = u_t$ to obtain

$$\begin{aligned} f_t(x_t) - f_t(u_t) &\leq \frac{1}{\eta} \left(D_r(u_t, x_t) - D_r(u_t, x_{t+1}) \right) \\ &\quad + f_t(x_t) - f_t(x_{t+1}). \end{aligned} \quad (12)$$

Then, we sum (12) over time to obtain

$$\begin{aligned} \sum_{t=1}^T \left(f_t(x_t) - f_t(u_t) \right) &\leq \sum_{t=1}^T \left(\frac{D_r(u_t, x_t)}{\eta} - \frac{D_r(u_t, x_{t+1})}{\eta} \right) \\ &\quad + \sum_{t=1}^T \left(f_t(x_t) - f_t(x_{t+1}) \right). \end{aligned} \quad (13)$$

We now separately bound the terms on the right-hand side of (13). In order to bound the first term of the above inequality, we add and subtract several terms as follows:

$$\begin{aligned} \sum_{t=1}^T \left(\frac{D_r(u_t, x_t)}{\eta} - \frac{D_r(u_t, x_{t+1})}{\eta} \right) &= \\ &\sum_{t=1}^T \left(\frac{D_r(u_t, x_t)}{\eta} - \frac{D_r(u_{t+1}, x_{t+1})}{\eta} \right) \\ &\quad + \sum_{t=1}^T \left(\frac{D_r(u_{t+1}, x_{t+1})}{\eta} - \frac{D_r(u_t, x_{t+1})}{\eta} \right) \\ &\leq \frac{D_r(u_1, x_1)}{\eta} + \sum_{t=1}^T \frac{\gamma \|u_{t+1} - u_t\|}{\eta}, \end{aligned} \quad (14)$$

where the last line follows from the fact that $D_r(x, y)$ is non-negative when $r(x)$ is convex, and the Lipschitz condition stated in (7).

We now proceed to bound the other term on the right-hand side of (13). We add and subtract $f_{t+1}(x_{t+1})$ to obtain

$$\begin{aligned} \sum_{t=1}^T \left(f_t(x_t) - f_t(x_{t+1}) \right) &= \\ \sum_{t=1}^T \left(f_t(x_t) - f_{t+1}(x_{t+1}) \right) &+ \sum_{t=1}^T \left(f_{t+1}(x_{t+1}) - f_t(x_{t+1}) \right) \\ &\leq f_1(x_1) - f_{T+1}(x_{T+1}) + \sum_{t=1}^T \sup_{x \in \mathcal{X}} |f_{t+1}(x) - f_t(x)|. \end{aligned} \quad (15)$$

Substituting (14) and (15) into (13) completes the proof. \square

C. Proof of Theorem 2

The smoothness of $f_t(\cdot)$ relative to $r(\cdot)$ implies

$$\begin{aligned} f_t(x_t) &\leq f_t(x_{t-1}) + \langle \nabla f_t(x_{t-1}), x_t - x_{t-1} \rangle + \beta D_r(x_t, x_{t-1}) \\ &\leq f_t(x_{t-1}) + \langle \nabla f_{t-1}(x_{t-1}), x_t - x_{t-1} \rangle \\ &\quad + \langle \nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1}), x_t - x_{t-1} \rangle \\ &\quad + \beta D_r(x_t, x_{t-1}). \end{aligned} \quad (16)$$

To bound the second term on the right hand-side of (16), we have

$$\begin{aligned} \eta \langle \nabla f_{t-1}(x_{t-1}), x_t - x_{t-1} \rangle &= \\ \eta \langle \nabla f_{t-1}(x_{t-1}), u - x_{t-1} \rangle &+ \eta \langle \nabla f_{t-1}(x_{t-1}), x_t - u \rangle \leq \\ \eta \langle \nabla f_{t-1}(x_{t-1}), u - x_{t-1} \rangle &+ D_r(u, x_{t-1}) - D_r(u, x_t) \\ &\quad - D_r(x_t, x_{t-1}), \forall u \in \mathcal{X}. \end{aligned} \quad (17)$$

where the last line follows from Lemma 6. By combining (16) and (17), and setting $u = u_t$ we obtain

$$\begin{aligned} f_t(x_t) &\leq f_t(x_{t-1}) + \langle \nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1}), x_t - x_{t-1} \rangle \\ &\quad + \langle \nabla f_{t-1}(x_{t-1}), u_t - x_{t-1} \rangle + \frac{D_r(u_t, x_{t-1})}{\eta} \\ &\quad - \frac{D_r(u_t, x_t)}{\eta} + \left(\beta - \frac{1}{\eta} \right) D_r(x_t, x_{t-1}). \end{aligned} \quad (18)$$

Furthermore, the strong convexity of $f_t(\cdot)$ relative to $r(\cdot)$ implies

$$f_t(x_{t-1}) + \langle \nabla f_t(x_{t-1}), u_t - x_{t-1} \rangle + \lambda D_r(u_t, x_{t-1}) \leq f_t(u_t). \quad (19)$$

By adding the last inequality to (18), and setting $\eta = \frac{1}{\beta}$ we have

$$\begin{aligned} f_t(x_t) &\leq f_t(u_t) + \langle \nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1}), x_t - x_{t-1} \rangle \\ &\quad + (\beta - \lambda) D_r(u_t, x_{t-1}) - \beta D_r(u_t, x_t) \\ &\leq f_t(u_t) + \|\nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1})\|_* \|x_t - x_{t-1}\| \\ &\quad + (\beta - \lambda) D_r(u_t, x_{t-1}) - (\beta - \lambda) D_r(u_t, x_t), \end{aligned} \quad (20)$$

where the second inequality follows from the fact that $\lambda D_r(u_t, x_t) \geq 0$. We now re-arrange the terms on (20) and sum over time to obtain

$$\begin{aligned} &\sum_{t=1}^T (f_t(x_t) - f_t(u_t)) \\ &\leq (\beta - \lambda) \sum_{t=1}^T (D_r(u_t, x_{t-1}) - D_r(u_t, x_t)) \\ &\quad + \sum_{t=1}^T \|x_t - x_{t-1}\| \|\nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1})\|_* \\ &\leq (\beta - \lambda) D_r(u_1, x_0) \\ &\quad + (\beta - \lambda) \sum_{t=1}^T (D_r(u_{t+1}, x_t) - D_r(u_t, x_t)) \\ &\quad + \sum_{t=1}^T \|x_t - x_{t-1}\| \|\nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1})\|_* \\ &\leq (\beta - \lambda) D_r(u_1, x_0) + (\beta - \lambda) \sum_{t=1}^T \gamma \|u_{t+1} - u_t\| \\ &\quad + 2M \sum_{t=1}^T \sup_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|_*, \end{aligned} \quad (21)$$

where the last inequality follows from the Bregman Lipschitz condition in (7). \square

REFERENCES

- [1] X. Cao, J. Zhang, and H. V. Poor, "Constrained online convex optimization with feedback delays," *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5049–5064, 2021.
- [2] S. M. Fosson, "Centralized and distributed online learning for sparse time-varying optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, pp. 2542–2557, 2021.
- [3] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the International Conference on Machine Learning*, pp. 928–936, 2003.
- [4] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.
- [5] A. Lesage-Landry, J. A. Taylor, and D. S. Callaway, "Online convex optimization with binary constraints," *IEEE Transactions on Automatic Control*, vol. 66, no. 12, pp. 6164–6170, 2021.
- [6] 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 International Conference on Machine Learning*, pp. 449–457, 2016.
- [7] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, "Online optimization in dynamic environments: Improved regret rates for strongly convex problems," in *Proceedings of the IEEE Conference on Decision and Control*, pp. 7195–7201, 2016.
- [8] L. Zhang, T. Yang, J. Yi, R. Jin, and Z.-H. Zhou, "Improved dynamic regret for non-degenerate functions," in *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, 2017.
- [9] T.-J. Chang and S. Shahrampour, "On online optimization: Dynamic regret analysis of strongly convex and smooth problems," in *Proceedings of the International Conference on Artificial Intelligence*, vol. 35, 2021.
- [10] E. Hall and R. Willett, "Dynamical models and tracking regret in online convex programming," in *Proceedings of the International Conference on Machine Learning*, pp. 579–587, 2013.
- [11] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 647–662, 2015.
- [12] Y. Zhao, S. Qiu, K. Li, L. Luo, J. Yin, and J. Liu, "Proximal online gradient is optimum for dynamic regret: A general lower bound," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–10, 2021.
- [13] X. Yi, X. Li, T. Yang, L. Xie, T. Chai, and K. Johansson, "Regret and cumulative constraint violation analysis for online convex optimization with long term constraints," in *Proceedings of the International Conference on Machine Learning*, pp. 11998–12008, 2021.
- [14] N. Eshraghi and B. Liang, "Dynamic regret bounds without Lipschitz continuity: online convex optimization with multiple mirror descent steps," in *Proceedings of the American Control Conference*, 2022.
- [15] O. Besbes, Y. Gur, and A. Zeevi, "Non-stationary stochastic optimization," *Operations Research*, vol. 63, no. 5, pp. 1227–1244, 2015.
- [16] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online optimization: Competing with dynamic comparators," in *Proc. Int. Conference on Artificial Intelligence and Statistics*, pp. 398–406, 2015.
- [17] S. Rakhlin and K. Sridharan, "Optimization, learning, and games with predictable sequences," in *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, pp. 3066–3074, 2013.
- [18] C.-K. Chiang, T. Yang, C.-J. Lee, M. Mahdavi, C.-J. Lu, R. Jin, and S. Zhu, "Online optimization with gradual variations," in *Proceedings of the Annual Conference on Learning Theory*, pp. 1–20, 2012.
- [19] H. Lu, R. M. Freund, and Y. Nesterov, "Relatively smooth convex optimization by first-order methods, and applications," *SIAM Journal on Optimization*, vol. 28, no. 1, pp. 333–354, 2018.
- [20] H. H. Bauschke, J. Bolte, and M. Teboulle, "A descent lemma beyond Lipschitz gradient continuity: first-order methods revisited and applications," *Mathematics of Operations Research*, vol. 42, no. 2, pp. 330–348, 2017.
- [21] H. Lu, "Relative continuity for non-Lipschitz nonsmooth convex optimization using stochastic or deterministic mirror descent," *INFORMS Journal on Optimization*, vol. 1, no. 4, pp. 288–303, 2019.
- [22] H. H. Bauschke, J. Bolte, J. Chen, M. Teboulle, and X. Wang, "On linear convergence of non-Euclidean gradient methods without strong convexity and Lipschitz gradient continuity," *Journal of Optimization Theory and Applications*, vol. 182, no. 3, pp. 1068–1087, 2019.
- [23] Y. Zhou, V. S. Portella, M. Schmidt, and N. J. Harvey, "Regret bounds without Lipschitz continuity: online learning with relative-Lipschitz losses," in *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, pp. 15823–15833, 2020.
- [24] S. Shahrampour and A. Jadbabaie, "Distributed online optimization in dynamic environments using mirror descent," *IEEE Transactions on Automatic Control*, vol. 63, no. 3, pp. 714–725, 2018.
- [25] P. Zhao, Y.-J. Zhang, L. Zhang, and Z.-H. Zhou, "Dynamic regret of convex and smooth functions," in *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, pp. 12510–12520, 2020.
- [26] J. C. Duchi, S. Shalev-Shwartz, Y. Singer, and A. Tewari, "Composite objective mirror descent," in *Proceedings of the Conference on Learning Theory*, pp. 14–26, 2010.
- [27] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003.