



Universiteit
Leiden
The Netherlands

The many faces of online learning

Hoeven, D. van der

Citation

Hoeven, D. van der. (2021, March 4). *The many faces of online learning*. Retrieved from <https://hdl.handle.net/1887/3147345>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3147345>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <https://hdl.handle.net/1887/3147345> holds various files of this Leiden University dissertation.

Author: Hoeven, D. van der

Title: The many faces of online learning

Issue Date: 2021-03-4

Comparator-Adaptive Convex Bandits

This chapter is based on: Van der Hoeven, D., Cutkosky, A., and Luo, H. (2020). Comparator-adaptive convex bandits. *To Appear in Advances in Neural Information Processing Systems* 33.¹

Abstract

We study bandit convex optimization methods that adapt to the norm of the comparator, a topic that has only been studied before for its full-information counterpart. Specifically, we develop convex bandit algorithms with regret bounds that are small whenever the norm of the comparator is small. We first use techniques from the full-information setting to develop comparator-adaptive algorithms for linear bandits. Then, we extend the ideas to convex bandits with Lipschitz or smooth loss functions, using a new single-point gradient estimator and carefully designed surrogate losses.

¹The author of this dissertation performed the following tasks: co-deriving the theoretical results and co-writing the paper. Part of the work on this Chapter was done while the author of this dissertation was visiting Haipeng Luo at the University of Southern California.

4.1 Introduction

In many situations, information is readily available. For example, if a gambler were to bet on the outcome of a football game, he can observe the outcome of the game regardless of what bet he made. In other situations, information is scarce. For example, the gambler could be deciding what to eat for dinner: should I eat a salad, a pizza, a sandwich, or not at all? These actions will result in different and unknown outcomes, but the gambler will only see the outcome of the action he actually takes, with one notable exception: not eating results in a predetermined outcome of being very hungry.

These two situations are instantiations of two different settings in online convex optimization: the full information setting and the bandit setting. More formally, both settings are sequential decision making problems where in each round $t = 1, \dots, T$, a learner has to make a prediction $\mathbf{w}_t \in \mathcal{W} \subseteq \mathbb{R}^d$ and an adversary provides a convex loss function $\ell_t : \mathcal{W} \rightarrow \mathbb{R}$. Afterwards, in the full information setting (Zinkevich, 2003) the learner has access to the loss function ℓ_t , while in the bandit setting (Kleinberg, 2005; Flaxman et al., 2005) the learner only receives the loss evaluated at the prediction, that is, $\ell_t(\mathbf{w}_t)$. In both settings the goal is to minimize the regret with respect to some benchmark point \mathbf{u} in hindsight, referred to as the *comparator*. More specifically, the regret against \mathbf{u} is the difference between the total loss incurred by the predictions of the learner and that of the comparator:

$$\mathcal{R}_T(\mathbf{u}) = \sum_{t=1}^T (\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})).$$

When the learner's strategy is randomized, we measure the performance by the expected regret $\mathbb{E}[\mathcal{R}_T(\mathbf{u})]$.

Standard algorithms in both the full information setting and the bandit setting assume that the learner's decision space \mathcal{W} is a convex *compact* set and achieve sublinear regret against the optimal comparator in this set: $\mathbf{u} = \arg \min_{\mathbf{u}^* \in \mathcal{W}} \sum_{t=1}^T \ell_t(\mathbf{u}^*)$. To tune these standard algorithms optimally, however, one requires knowledge of the norm of the comparator $\|\mathbf{u}\|$, which is unknown. A common work-around is to simply tune the algorithms in terms of the worst-case norm: $\max_{\mathbf{u} \in \mathcal{W}} \|\mathbf{u}\|$, assumed to be 1 without loss of generality. This results in worst-case bounds that do not take advantage of the case when $\|\mathbf{u}\|$ is small. For example, when the loss functions are L -Lipschitz, classic Online Gradient Descent (Zinkevich, 2003) guarantees $\mathcal{R}_T(\mathbf{u}) = O(L\sqrt{T})$ in the full information setting, while the algorithm of (Flaxman et al., 2005) guarantees

$\mathbb{E}[\mathcal{R}_T(\mathbf{u})] = O(d\sqrt{LT}^{3/4})$ in the bandit setting, both of which are independent of $\|\mathbf{u}\|$.

Recently, there has been a series of works in the full information setting that addresses this problem by developing *comparator-adaptive* algorithms, whose regret against \mathbf{u} depends on $\|\mathbf{u}\|$ for *all* $\mathbf{u} \in \mathcal{W}$ *simultaneously* (see for example McMahan and Orabona (2014); Orabona and Pál (2016); Foster et al. (2017); Cutkosky and Boahen (2017); Kotłowski (2017); Cutkosky and Orabona (2018); Foster et al. (2018b); Jun and Orabona (2019); Van der Hoeven (2019)). These bounds are often not worse than the standard worst-case bounds, but could be much smaller in the case when there exists a comparator with small norm and reasonably small total loss. Moreover, most of these results also hold for the so-called *unconstrained* setting where $\mathcal{W} = \mathbb{R}^d$, that is, both the learner's predictions and the comparator can be any point in \mathbb{R}^d . For example, Cutkosky and Orabona (2018) achieve $\mathcal{R}_T(\mathbf{u}) = \tilde{O}(\|\mathbf{u}\|L\sqrt{T})$ for all \mathbf{u} , in both the constrained and unconstrained settings, under full information feedback.²

While developing comparator-adaptive algorithms is relatively well-understood at this point in the full information setting, to the best of our knowledge, this has not been studied at all for the more challenging bandit setting. In this work, we take the first attempt in this direction and develop comparator-adaptive algorithms for several situations, including learning with linear losses, general convex losses, and convex and smooth losses, for both the constrained and unconstrained settings. Our results are summarized in Table 4.1. Ignoring other parameters for simplicity, for the linear case, we achieve $\tilde{O}(\|\mathbf{u}\|\sqrt{T})$ regret (Section 4.3.2); for the general convex case, we achieve $\tilde{O}(\|\mathbf{u}\|T^{\frac{3}{4}})$ regret in both the constrained and unconstrained setting (Sections 4.4.1 and 4.4.2); and for the convex and smooth case, we achieve $\tilde{O}\left(\max\{\|\mathbf{u}\|^2, \|\mathbf{u}\|\}\beta(dLT)^{\frac{2}{3}}\right)$ regret in the unconstrained setting (Section 4.4.1).

In order to achieve our results for the convex case, we require an assumption on the loss, namely that the value of $\ell_t(\mathbf{0})$ is known for all t .³ While restrictive at first sight, we believe that there are abundant applications where this assumption holds. As one instance, in control or reinforcement learning problems, $\mathbf{0}$ may represent some nominal action which has a known outcome: not eating results in hunger, or buying zero inventory will result in zero revenue. Another application is a classification problem where the features are not revealed to the learner. For example, end-users of a prediction service may not feel comfortable revealing their information to the

²Throughout the chapter, the notation \tilde{O} hides logarithmic dependence on parameters T, L , and $\|\mathbf{u}\|$.

³For the linear case, this clearly holds since $\ell_t(\mathbf{0}) = 0$.

4. Comparator-Adaptive Convex Bandits

Table 4.1: Summary of main results. Regret is measured with respect to the total loss of an arbitrary point $\mathbf{u} \in \mathbb{R}^d$ in the unconstrained setting, or an arbitrary point $\mathbf{u} \in \mathcal{W}$ in the constrained setting with a decision space \mathcal{W} contained in the unit ball. T is the total number of rounds and $1/c$ is radius of the largest ball contained by \mathcal{W} . c is bounded by $O(d)$.

| Loss functions (L -Lipschitz) | Regret for unconstrained settings | Regret for constrained settings |
|--|---|--|
| Linear (Section 4.3.2) | $\tilde{O}(\ \mathbf{u}\ dL\sqrt{T})$ | $\tilde{O}(\ \mathbf{u}\ cdL\sqrt{T})$ |
| Convex (Section 4.4.1 and 4.4.2) | $\tilde{O}(\ \mathbf{u}\ L\sqrt{dT}^{\frac{3}{4}})$ | $\tilde{O}(\ \mathbf{u}\ cL\sqrt{dT}^{\frac{3}{4}})$ |
| Convex and β -smooth (Section 4.4.2) | $\tilde{O}(\max\{\ \mathbf{u}\ ^2, \ \mathbf{u}\ \}\beta(dLT)^{\frac{2}{3}})$ | - |

service. Instead, they may be willing to do some local computation and report the loss of the service’s model. Most classification models (e.g. logistic regression) have the property that the loss of the $\mathbf{0}$ parameter is a known constant regardless of the data, and so this situation would also fit into our framework. Common loss functions that satisfy this assumption are linear loss, logistic loss, and hinge loss.

Techniques Our algorithms are based on sophisticated extensions of the black-box reduction introduced by Cutkosky and Orabona (2018), which separately learns the magnitude and the direction of the prediction. To make the reduction work in the bandit setting, however, new ideas are required, including designing an appropriate surrogate loss function and a new one-point gradient estimator with time-varying parameters. Note that (Cutkosky and Orabona, 2018) also proposes a method to convert any unconstrained algorithm to a constrained one in the full information setting, but this does not work in the bandit setting for technical reasons. Instead, we take a different approach by constraining the magnitude of the prediction directly.

Related work As mentioned, there has been a line of recent works on comparator-adaptive algorithms for the full information setting. Most of them do not transfer to the bandit setting, except for the approach of Cutkosky and Orabona (2018) from which we draw heavy inspiration. To the best of our knowledge, comparator-adaptive bandit algorithms have not been studied before. Achieving “adaptivity” in a broader sense is generally hard for problems with bandit feedback; see negative results such as (Daniely et al., 2015; Lattimore, 2015) as well as recent progress such as (Chen et al., 2019; Foster et al., 2019).

In terms of worst-case (non-adaptive) regret, the seminal work of (Abernethy et al., 2008) is the first to achieve $O(\sqrt{T})$ regret for bandit with linear losses, and (Kleinberg, 2005; Flaxman et al., 2005) are the first to achieve sublinear regret for general

convex case. Over the past decade, the latter result has been improved in many different ways (Agarwal et al., 2010; Saha and Tewari, 2011; Agarwal et al., 2011; Hazan and Levy, 2014), and regret of order $O(\sqrt{T})$ under no extra assumptions was recently achieved (Bubeck et al., 2015; Bubeck and Eldan, 2016; Bubeck et al., 2017). However, these $O(\sqrt{T})$ bounds are achieved by very complicated algorithms that incur a huge dependence on the dimension d . Our algorithms are more aligned with the simpler ones with milder dimension-dependence (Abernethy et al., 2008; Flaxman et al., 2005; Saha and Tewari, 2011) and achieve the same dependence on T in different cases. How to achieve comparator-adaptive regret of order $O(\sqrt{T})$ for the general convex case is an important future direction.

4.2 Preliminaries

In this section, we describe our notation, state the definitions we use, and introduce the bandit convex optimization setting formally. We also describe the black-box reduction of Cutkosky and Orabona (2018) we will use throughout the chapter.

Notation and definitions The inner product between vectors $\mathbf{g} \in \mathbb{R}^d$ and $\mathbf{w} \in \mathbb{R}^d$ is denoted by $\langle \mathbf{w}, \mathbf{g} \rangle$. \mathbb{R}_+ denotes the set of positive numbers. The Fenchel conjugate F^* of a convex function F is defined as $F^*(\mathbf{w}) = \sup_{\mathbf{g}} \langle \mathbf{w}, \mathbf{g} \rangle - F(\mathbf{g})$. $\|\cdot\|$ denotes a norm and $\|\mathbf{g}\|_* = \sup_{\mathbf{w}: \|\mathbf{w}\| \leq 1} \langle \mathbf{w}, \mathbf{g} \rangle$ denotes the dual norm of \mathbf{g} . The Bregman divergence associated with convex function F between points \mathbf{x} and \mathbf{y} is denoted by $B_F(\mathbf{x} \|\mathbf{y}) = F(\mathbf{x}) - F(\mathbf{y}) - \langle \nabla F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$, where $\nabla F(\mathbf{x})$ denotes the gradient of F evaluated at \mathbf{x} . The unit ball equipped with norm $\|\cdot\|$ is denoted by $\mathcal{B} = \{\mathbf{w} : \|\mathbf{w}\| \leq 1\}$. The unit sphere with norm $\|\cdot\|$ is denoted by $\mathcal{S} = \{\mathbf{w} : \|\mathbf{w}\| = 1\}$. The unit ball and sphere with norm $\|\cdot\|_2$ are denoted by \mathbb{B} and \mathbb{S} respectively. $\mathbf{x} \sim U(\mathcal{Z})$ denotes that \mathbf{x} follows the uniform distribution over \mathcal{Z} . We say a function f is β -smooth over the set \mathcal{W} if the following holds:

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|_2^2, \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{W}.$$

We say a function f is L -Lipschitz over the set \mathcal{W} if the following holds:

$$|f(\mathbf{y}) - f(\mathbf{x})| \leq L \|\mathbf{y} - \mathbf{x}\|_2, \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{W}.$$

Throughout the chapter we will assume that $\beta, L \geq 1$. Also, by mild abuse of notation, we use $\partial f(\mathbf{x})$ to indicate an arbitrary subgradient of a convex function f at \mathbf{x} .

All of our algorithms are reductions that use prior algorithms in disparate ways to obtain our new results. In order for these reductions to work, we need some

Algorithm 4 Black-Box Reduction with Full Information

- 1: **Input:** “Direction” algorithm $\mathcal{A}_{\mathcal{Z}}$ and “scaling” algorithm $\mathcal{A}_{\mathcal{V}}$
 - 2: **for** $t = 1 \dots T$ **do**
 - 3: Get $z_t \in \mathcal{Z}$ from $\mathcal{A}_{\mathcal{Z}}$
 - 4: Get $v_t \in \mathbb{R}$ from algorithm $\mathcal{A}_{\mathcal{V}}$
 - 5: Play $w_t = v_t z_t$, receive g_t
 - 6: Send g_t to algorithm $\mathcal{A}_{\mathcal{Z}}$ as the t -th loss vector
 - 7: Send $\langle z_t, g_t \rangle$ to algorithm $\mathcal{A}_{\mathcal{V}}$ as the t -th loss value
 - 8: **end for**
-

assumptions on the base algorithms. We will encapsulate these assumptions in *interfaces* that describe inputs, outputs, and guarantees described by an algorithm rather than its actual operation (see Interfaces 6 and 7 for examples). We can use specific algorithms from the literature to implement these interfaces, but our results depend only on the properties described in the interfaces.

4.2.1 Bandit Convex Optimization

The bandit convex optimization protocol proceeds in rounds $t = 1, \dots, T$. In each round t the learner plays $w_t \in \mathcal{W} \subseteq \mathbb{R}^d$. Simultaneously, the environment picks an L -Lipschitz convex loss function $\ell_t : \mathcal{W} \rightarrow \mathbb{R}$, after which the learner observes $\ell_t(w_t)$. Importantly, the learner only observes the loss function evaluated at w_t , not the function itself. This forces the learner to play random points and estimate the feedback he wants to use to update w_t . Therefore, in the bandit feedback setting, the goal is to bound the *expected* regret $\mathbb{E}[\mathcal{R}_T(u)]$, where the expectation is with respect to randomisation of the learner.

We make a distinction between linear bandits, where $\ell_t(w) = \langle w, g_t \rangle$, and convex bandits, where ℓ_t can be any L -Lipschitz convex function. Throughout the chapter, if $\mathcal{W} \neq \mathbb{R}^d$ we assume that \mathcal{W} is compact, has a non-empty interior, and contains $\mathbf{0}$. Without loss of generality we assume that $\frac{1}{c}\mathbb{B} \subseteq \mathcal{W} \subseteq \mathbb{B}$ for some $c \geq 1$. Some of our bounds depend on c , which, without loss of generality, is always bounded by d , due to a reshaping trick discussed in (Flaxman et al., 2005).

4.2.2 Black-Box Reductions with Full Information

Our algorithms are based on a black-box reduction from (Cutkosky and Orabona, 2018) for the full information setting (see Algorithm 4). The reduction works as follows. In each round t the algorithms plays $w_t = v_t z_t$, where $z_t \in \mathcal{Z}$ for some domain \mathcal{Z} , is the prediction of a constrained algorithm $\mathcal{A}_{\mathcal{Z}}$, and v_t is the prediction

of a one-dimensional algorithm $\mathcal{A}_{\mathcal{V}}$. The goal of $\mathcal{A}_{\mathcal{Z}}$ is to learn the directions of the comparator while the goal of $\mathcal{A}_{\mathcal{V}}$ is to learn the norm of the comparator. Let \mathbf{g}_t be the gradient of ℓ_t at \mathbf{w}_t , which is known to the algorithm in the full information setting. We feed \mathbf{g}_t as feedback to $\mathcal{A}_{\mathcal{Z}}$ and $\langle \mathbf{z}_t, \mathbf{g}_t \rangle$ as feedback to $\mathcal{A}_{\mathcal{V}}$. Although the original presentation considers only $\mathcal{Z} = \mathcal{B}$, we will need to extend the analysis to more general domains.

As outlined by Cutkosky and Orabona (2018), the regret of Algorithm 4 decomposes into two parts. The first part of the regret is for learning the norm of \mathbf{u} , and is controlled by Algorithm $\mathcal{A}_{\mathcal{V}}$. The second part of the regret is for learning the direction of \mathbf{u} and is controlled by $\mathcal{A}_{\mathcal{Z}}$. The proof is provided in Section 4.6 for completeness.

Lemma 12. *Let $\mathcal{R}_T^{\mathcal{V}}(\|\mathbf{u}\|) = \sum_{t=1}^T (v_t - \|\mathbf{u}\|) \langle \mathbf{z}_t, \mathbf{g}_t \rangle$ be the regret for learning $\|\mathbf{u}\|$ by Algorithm $\mathcal{A}_{\mathcal{V}}$ and let $\mathcal{R}_T^{\mathcal{Z}}\left(\frac{\mathbf{u}}{\|\mathbf{u}\|}\right) = \sum_{t=1}^T \langle \mathbf{z}_t - \frac{\mathbf{u}}{\|\mathbf{u}\|}, \mathbf{g}_t \rangle$ be the regret for learning $\frac{\mathbf{u}}{\|\mathbf{u}\|}$ by $\mathcal{A}_{\mathcal{Z}}$. Then Algorithm 4 satisfies*

$$\mathcal{R}_T(\mathbf{u}) = \mathcal{R}_T^{\mathcal{V}}(\|\mathbf{u}\|) + \|\mathbf{u}\| \mathcal{R}_T^{\mathcal{Z}}\left(\frac{\mathbf{u}}{\|\mathbf{u}\|}\right). \quad (4.2.1)$$

Cutkosky and Orabona (2018) provide an algorithm to ensure $\mathcal{R}_T^{\mathcal{V}}(\|\mathbf{u}\|) \leq 1 + \|\mathbf{u}\| 8L \sqrt{T \log(\|\mathbf{u}\|T + 1)}$, given that $\|\mathbf{g}_t\|_* \leq L$. This algorithm satisfies the requirements described later in Interface 6, and will be used throughout this chapter.

4.3 Comparator-Adaptive Linear Bandits

Now, we apply the reduction of section 4.2.2 to develop comparator-adaptive algorithms for linear bandits. We will see that in the unconstrained case, the reduction works almost without modification. In the constrained case we will need to be more careful to enforce the constraints.

4.3.1 Unconstrained Linear Bandits

We begin by discussing the unconstrained linear bandit setting, which turns out to be the easiest setting we consider. Following Algorithm 4, we will still play $\mathbf{w}_t = v_t \mathbf{z}_t$. However, instead of taking a fixed \mathbf{z}_t from a full-information algorithm, we take a random \mathbf{z}_t from a *bandit* algorithm. Importantly, we can recover $\langle \mathbf{z}_t, \mathbf{g}_t \rangle$ exactly since $\langle \mathbf{w}_t, \mathbf{g}_t \rangle \frac{1}{v_t} = \langle \mathbf{z}_t, \mathbf{g}_t \rangle$. This means that we have enough information to send appropriate feedback to both $\mathcal{A}_{\mathcal{V}}$ and $\mathcal{A}_{\mathcal{Z}}$ and apply the argument of Lemma 12. Interestingly, we use a full-information one-dimensional algorithm for $\mathcal{A}_{\mathcal{V}}$, and

Algorithm 5 Black-Box Reduction for Linear Bandits

- 1: **Input:** Constrained Linear Bandit Algorithm $\mathcal{A}_{\mathcal{Z}}$ and unconstrained 1-d Algorithm $\mathcal{A}_{\mathcal{V}}$
 - 2: **for** $t = 1 \dots T$ **do**
 - 3: Get $\mathbf{z}_t \in \mathcal{Z}$ from $\mathcal{A}_{\mathcal{Z}}$
 - 4: Get $v_t \in \mathbb{R}$ from $\mathcal{A}_{\mathcal{V}}$
 - 5: Play $\mathbf{w}_t = v_t \mathbf{z}_t$
 - 6: Receive loss $\langle \mathbf{w}_t, \mathbf{g}_t \rangle$
 - 7: Compute $\mathcal{L}_t = \frac{1}{v_t} \langle \mathbf{w}_t, \mathbf{g}_t \rangle = \langle \mathbf{z}_t, \mathbf{g}_t \rangle$.
 - 8: Send \mathcal{L}_t to Algorithm $\mathcal{A}_{\mathcal{Z}}$ as t -th loss value.
 - 9: Send \mathcal{L}_t to Algorithm $\mathcal{A}_{\mathcal{V}}$ as t -th loss value.
 - 10: **end for**
-

only need $\mathcal{A}_{\mathcal{Z}}$ to take bandit input. This is because $\mathcal{A}_{\mathcal{V}}$ gets full information in the form of $\langle \mathbf{z}_t, \mathbf{g}_t \rangle$.

The algorithm $\mathcal{A}_{\mathcal{Z}}$ for learning the direction, on the other hand, now must be a bandit algorithm because intuitively we do not immediately get the full direction information \mathbf{g}_t from the value of the loss alone. We will need this algorithm to fulfill the requirements described by Interface 7. An important requirement is that the expected regret of the direction learning algorithm is bounded by $dL\tau\sqrt{T\log(T+1)}$, where τ is a constant. One such algorithm is given by continuous Exponential Weights on a constrained set (see Van der Hoeven et al. (2018, section 6) for details).

Our unconstrained linear bandit algorithm then is constructed from Algorithm 5 by choosing an algorithm that implements Interface 7 as $\mathcal{A}_{\mathcal{Z}}$ and Interface 6 with $l = \mathbb{R}$ as $\mathcal{A}_{\mathcal{V}}$. Plugging in the guarantees of the individual algorithms and taking the expectation of (4.2.1), the total expected regret is $\tilde{O}(1 + \|\mathbf{u}\|dL\sqrt{T})$. Compared to the full information setting we have lost a factor d in the regret bound, which is unavoidable given the bandit feedback (Dani et al., 2008). The formal result is below.

Theorem 13. *Suppose $\mathcal{A}_{\mathcal{Z}}$ implements Interface 7 with domain $\mathcal{Z} = \mathcal{B}$ and $\mathcal{A}_{\mathcal{V}}$ implements Interface 6 with $l = \mathbb{R}_+$. Then Algorithm 5 satisfies for all $\mathbf{u} \in \mathbb{R}^d$:*

$$\mathbb{E}[\mathcal{R}(\mathbf{u})] = 1 + \|\mathbf{u}\|(dL\tau\sqrt{T\log(T+1)} + L\sqrt{TB(\hat{v})}).$$

4.3.2 Constrained Linear Bandits

The algorithm in the previous section only works for $\mathcal{W} = \mathbb{R}^d$. In this section, we consider a compact set $\mathcal{W} \subset \mathbb{R}^d$.

Interface 6 Scale Learning Interface (see example implementation in Cutkosky and Orabona (2018))

- 1: **Input:** A line segment $l \subseteq \mathbb{R}$
 - 2: **for** $t = 1 \dots T$ **do**
 - 3: Play $v_t \in l$
 - 4: Receive loss value g_t such that $|g_t| \leq L_V$
 - 5: **end for**
 - 6: **Ensure:** for all $\hat{v} \in l$, $\sum_{t=1}^T (v_t - \hat{v})g_t \leq 1 + |\hat{v}|L_V\sqrt{TB(\hat{v})}$
-

Interface 7 Direction Learning Interface for Linear Bandits (see example implementation in Van der Hoeven et al. (2018))

- 1: **Input:** Domain \mathcal{Z}
 - 2: **for** $t = 1 \dots T$ **do**
 - 3: Play $z_t \in \mathcal{Z}$
 - 4: Receive loss value $\langle z_t, g_t \rangle$ such that $|\langle z_t, g_t \rangle| \leq L$
 - 5: **end for**
 - 6: **Ensure:** for all $u \in \mathcal{Z}$, $\mathbb{E} \left[\sum_{t=1}^T \langle z_t - u, g_t \rangle \right] \leq dL\tau\sqrt{T\log(T+1)}$
-

In the full-information setting, Cutkosky and Orabona (2018) provide a projection technique for producing constrained algorithms from unconstrained ones. Unfortunately, this technique does not translate directly to the bandit setting, and we must be more careful in designing our constrained linear bandit algorithm. The key idea is to constrain the internal scaling algorithm \mathcal{A}_V , rather than attempting to constrain the final predictions w_t . Enforcing constraints on the scaling algorithm's outputs v_t will naturally translate into a constraint on the final predictions w_t .

To produce a constrained linear bandit algorithm, we again use Algorithm 5, but now we instantiate \mathcal{A}_V implementing Interface 6 with $l = [0, 1]$ rather than $l = \mathbb{R}_+$, and instantiate \mathcal{A}_Z implementing Interface 7 with $\mathcal{Z} = \mathcal{W}$ rather than $\mathcal{Z} = \mathcal{B}$. As in the unconstrained setting, this allows us to feed full information feedback to \mathcal{A}_V . At the same time, restricting Interface 6 to $l = [0, 1]$ also guarantees that $w_t \in \mathcal{W}$. The regret bound of this algorithm is given in Theorem 14. The proof follows from combining Lemma 12 with the guarantees of Interfaces 6 and 7 and can be found in Section 4.7.

Theorem 14. *Suppose \mathcal{A}_Z implements 7 with domain $\mathcal{Z} = \mathcal{W}$ and \mathcal{A}_V implements Interface 6 with $l = [0, 1]$. Then Algorithm 5 satisfies for all $u \in \mathcal{W}$,*

$$\mathbb{E}[\mathcal{R}_T(u)] \leq 1 + \|u\|cL \left(d\tau\sqrt{T\log(T+1)} + \sqrt{TB(c\|u\|)} \right).$$

If \mathcal{W} is a unit ball, then $c = 1$. For other shapes of \mathcal{W} , recall that c is at most d , which leads to a expected regret bound of $\tilde{O}\left(1 + \|\mathbf{u}\|d^2L\sqrt{T}\right)$.

4.4 Comparator-Adaptive Convex Bandits

In the general convex bandit problem, it is not clear how to use the single evaluation point feedback $\ell_t(\mathbf{w}_t)$ to derive any useful information about ℓ_t . Fortunately, Flaxman et al. (2005) solved this problem by using randomness to extract the gradients of a smoothed version of ℓ_t . To adapt to the norm of the comparator, we employ the following tweaked version of smoothing used by Flaxman et al. (2005):

$$\ell_t^v(\mathbf{w}) = \mathbb{E}_{\mathbf{b} \sim U(\mathbb{B})}[\ell_t(\mathbf{w} + v\delta\mathbf{b})], \quad (4.4.1)$$

where $v, \delta > 0$. In contrast to prior work using this framework, our smoothing now depends on the scaling parameter v . Lemma 13 gives the gradient of $\ell_t^v(\mathbf{w})$ and is a straightforward adaptation of Lemma 2.1 by Flaxman et al. (2005).

Lemma 13. For $\delta \in (0, 1]$, $v > 0$:

$$\nabla \ell_t^v(\mathbf{w}) = \frac{d}{v\delta} \mathbb{E}_{\mathbf{s} \sim U(\mathbb{S})}[\ell_t(\mathbf{w} + v\delta\mathbf{s})\mathbf{s}]. \quad (4.4.2)$$

With this lemma, we can estimate the gradient of the smoothed version of ℓ_t by evaluating ℓ_t at a random point, essentially converting the convex problem to a linear problem, except that one also needs to control the bias introduced by smoothing. Note that this estimate scales with $\frac{1}{v}$, which can be problematic if v is small. To deal with this issue, we require one extra assumption: the value of $\ell_t(\mathbf{0})$ is known to the learner. As discussed in section 4.1, this assumption holds for several applications, including some control or reinforcement learning problems, where $\mathbf{0}$ represents a nominal action with a known outcome. Furthermore, certain loss functions satisfy the second assumption by default, such as linear loss, logistic loss, and hinge loss. Without loss of generality we assume that $\ell_t(\mathbf{0}) = 0$, as we can always shift ℓ_t without changing the regret.

Our general algorithm template is provided in Algorithm 8. It incorporates the ideas of Algorithm 5, but adds new smoothing and regularization elements in order to deal with the present more general situation. More specifically, it again makes use of subroutine \mathcal{A}_V , which learns the scaling. The direction is learned by Online Gradient Descent (Zinkevich, 2003), as was also done by Flaxman et al. (2005). Given \mathbf{z}_t and v_t , our algorithm plays the point $\mathbf{w}_t = v_t(\mathbf{z}_t + \delta\mathbf{s}_t)$ for some

parameter δ and \mathbf{s}_t drawn uniformly at random from \mathbb{S} . By equation (4.4.2), we have

$$\mathbb{E} \left[\frac{d}{v_t \delta} \ell_t(\mathbf{w}_t) \mathbf{s}_t \right] = \nabla \ell_t^{v_t}(\mathbf{v}_t \mathbf{z}_t). \quad (4.4.3)$$

This means that we can use $\hat{\mathbf{g}}_t = \frac{d}{v_t \delta} \ell_t(\mathbf{w}_t) \mathbf{s}_t$ as an approximate gradient estimate, and we send this $\hat{\mathbf{g}}_t$ to Online Gradient Descent as the feedback. In other words, Online Gradient Descent itself is essentially dealing with a full-information problem with gradient feedback and is required to ensure a regret bound $\mathbb{E}[\sum_{t=1}^T \langle \mathbf{z}_t - \mathbf{u}, \hat{\mathbf{g}}_t \rangle] = \tilde{O}(\frac{dL}{\delta} \sqrt{T})$ for all \mathbf{u} in some domain \mathcal{Z} . For technical reasons, we will also need to enforce $\mathbf{z}_t \in (1 - \alpha)\mathcal{Z}$ for some $\alpha \in [0, 1]$. This restriction will be necessary in the constrained setting to ensure $v_t(\mathbf{z}_t + \delta \mathbf{s}_t) \in \mathcal{W}$.

Next, to specify the feedback to the scaling learning black-box $\mathcal{A}_{\mathcal{V}}$, we define a surrogate loss function $\bar{\ell}_t(v)$ which contains a linear term $v \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle$ and also a regularization term (see Algorithm 8 for the exact definition). The feedback to $\mathcal{A}_{\mathcal{V}}$ is then $\partial \bar{\ell}_t(v_t)$. Therefore, $\mathcal{A}_{\mathcal{V}}$ is essentially learning these surrogate losses, also with full gradient information. The regularization term is added to deal with the bias introduced by smoothing. This term does not appear in prior work on convex bandits, and it is one of the key components needed to ensure that the final regret is in terms of the unknown $\|\mathbf{u}\|$.

Algorithm 8 should be seen as the analogue of the black-box reduction of Algorithm 4, but for bandit feedback instead of full information. The expected regret guarantee of Algorithm 8 is shown below, and the proof can be found in Section 4.8.

Lemma 14. *Suppose $\mathcal{A}_{\mathcal{V}}$ implements Interface 6 with $l \subseteq \mathbb{R}_+$. Suppose $\mathbf{w}_t \in \mathcal{W}$ for all t , $\ell_t(\mathbf{0}) = 0$, and let $L_{\mathcal{V}} = \max_t \partial \bar{\ell}_t(v_t)$. Then Algorithm 8 with $\delta, \alpha \in (0, 1]$ and $\eta = \sqrt{\frac{\delta^2}{4(dL)^2 T}}$ satisfies for all $\|\mathbf{u}\| \in l$ and $r > 0$ with $\frac{ur}{\|\mathbf{u}\|} \in \mathcal{Z}$,*

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T(\mathbf{u})] \leq & 1 + 2T\delta L \frac{\|\mathbf{u}\|}{r} + \frac{\|\mathbf{u}\|}{r} L_{\mathcal{V}} \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} \\ & + \frac{2\|\mathbf{u}\|dL}{r\delta} \sqrt{T} + \alpha \|\mathbf{u}\|_2 TL. \end{aligned}$$

In addition, if ℓ_t is also β -smooth for all t , then we have

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T(\mathbf{u})] \leq & 1 + T\beta\delta^2 \left(\frac{\|\mathbf{u}\|}{r} \right)^2 + \frac{\|\mathbf{u}\|}{r} L_{\mathcal{V}} \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} \\ & + \frac{2\|\mathbf{u}\|dL}{r\delta} \sqrt{T} + \alpha \|\mathbf{u}\|_2 TL. \end{aligned}$$

Algorithm 8 Black-Box Comparator-Adaptive Convex Bandit Algorithm

```

1: Input: Scaling algorithm  $\mathcal{A}_\mathcal{V}$ ,  $\delta \in (0, 1]$ ,  $\alpha \in [0, 1]$ , domain  $\mathcal{Z} \subseteq \mathbb{B}$ , and
   learning rate  $\eta$ 
2: Set  $\mathbf{z}_1 = \mathbf{0}$ 
3: for  $t = 1 \dots T$  do
4:   Get  $v_t$  from  $\mathcal{A}_\mathcal{V}$ 
5:   Sample  $\mathbf{s}_t \sim U(\mathbb{S})$ 
6:   Set  $\mathbf{w}_t = v_t(\mathbf{z}_t + \delta \mathbf{s}_t)$ 
7:   Play  $\mathbf{w}_t$ 
8:   Receive  $\ell_t(\mathbf{w}_t)$ 
9:   Set  $\hat{\mathbf{g}}_t = \frac{d}{v_t \delta} \ell_t(\mathbf{w}_t) \mathbf{s}_t$ 
10:  if  $\ell_t$  is  $\beta$ -smooth then
11:    Set  $\bar{\ell}_t(v) = v \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle + \beta \delta^2 v^2$ 
12:  else
13:    Set  $\bar{\ell}_t(v) = v \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle + 2\delta L|v|$ 
14:  end if
15:  Send  $\partial \bar{\ell}_t(v_t)$  to algorithm  $\mathcal{A}_\mathcal{V}$  as the  $t$ -th loss value
16:  Update  $\mathbf{z}_{t+1} = \arg \min_{\mathbf{z} \in (1-\alpha)\mathcal{Z}} \eta \langle \mathbf{z}, \hat{\mathbf{g}}_t \rangle + \|\mathbf{z}_t - \mathbf{z}\|_2^2$ 
17: end for
    
```

This bound has two main points not obviously under our direct control: the assumption that the \mathbf{w}_t lie in \mathcal{W} , and the value of $L_\mathcal{V}$, which is a bound on $|\partial \bar{\ell}_t(v_t)|$. In the remainder of this section we will specify the various settings of Algorithm 8 that guarantee that $\mathbf{w}_t \in \mathcal{W}$ and that $L_\mathcal{V}$ is suitably bounded: two settings for the unconstrained setting and one for the constrained setting. The $\alpha \|\mathbf{u}\|TL$ term due to $\mathbf{z}_t \in (1 - \alpha)\mathcal{Z}$ rather than $\mathbf{z}_t \in \mathcal{Z}$, which induces a small amount of bias. The r in Lemma 14 is to ensure that we satisfy the requirements for Online Gradient Descent to have a suitable regret bound. For unconstrained convex bandits $r = 1$. For constrained convex bandits we will find that $\frac{1}{r} = c$ (recall that we assume that $\frac{1}{c}\mathbb{B} \subseteq \mathcal{W} \subseteq \mathbb{B}$).

4.4.1 Unconstrained Convex Bandits

In this section we instantiate Algorithm 8 and derive regret bounds for either general convex losses or convex and smooth losses. We start with general convex losses. Since $\mathcal{W} = \mathbb{R}^d$, we do not need to ensure that $\mathbf{z}_t + \delta \mathbf{s}_t \in \mathcal{W}$ and we can safely set $\alpha = 0$. This choice guarantees that $\mathbf{z}_t + \delta \mathbf{s}_t \in 2\mathbb{B}$ and that $|\partial \bar{\ell}_t(v_t)| \leq \frac{2dL}{\delta} + 2\delta L$. Then, Lemma 14 directly leads to Theorem 15 (the proof is deferred to Section 4.8.1).

Theorem 15. Suppose \mathcal{A}_V implements Interface 6 with $l = \mathbb{R}_+$ and that $\ell_t(\mathbf{0}) = 0$. Then Algorithm 8 with $\delta = \min\{1, \sqrt{dT}^{-\frac{1}{4}}\}$, $\mathcal{Z} = \mathbb{B}$, $\alpha = 0$, and $\eta = \sqrt{\frac{\delta^2}{4(dL)^2T}}$ satisfies for all $\mathbf{u} \in \mathbb{R}^d$,

$$\mathbb{E}[\mathcal{R}_T(\mathbf{u})] = 2 + \left(4\|\mathbf{u}\|Ld\sqrt{T} + 5\|\mathbf{u}\|L\sqrt{dT}^{\frac{3}{4}}\right)(1 + \sqrt{B(\|\mathbf{u}\|)}).$$

For unconstrained smooth bandits, we face an extra challenge. To bound the regret of Algorithm 8, $|\partial\bar{\ell}_t(v_t)| = |\langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle + \beta 2\delta^2 v_t|$ must be bounded. Now in contrast to the linear or Lipschitz cases, in the smooth case $\bar{\ell}_t(v_t)$ is not Lipschitz over \mathbb{R}_+ . We will address this by artificially constraining v_t . Specifically, we ensure that $v_t \leq \frac{1}{\delta^3}$, which implies $|\delta^2 v_t| = O(\frac{1}{\delta})$. This makes the Lipschitz constant of $\bar{\ell}_t$ to be dominated by the gradient estimate $\hat{\mathbf{g}}_t$ rather than the regularization. To see how this affects the regret bound, consider two cases, $\|\mathbf{u}\|_2 \leq \frac{1}{\delta^3}$ and $\|\mathbf{u}\|_2 > \frac{1}{\delta^3}$. If $\|\mathbf{u}\|_2 \leq \frac{1}{\delta^3}$ then we have not hurt anything by constraining v_t since $\|\mathbf{u}\|_2$ satisfies the same constraint. If instead $\|\mathbf{u}\|_2 > \frac{1}{\delta^3}$ then the consequences for the regret bound are not immediately clear. However, following a similar technique in Cutkosky (2019), we use the fact that the regret against $\mathbf{0}$ is $O(1)$ and the Lipschitz assumption to show that we have added a penalty of only $O(\|\mathbf{u}\|_2 LT)$:

$$\mathbb{E}[\mathcal{R}_T(\mathbf{u})] = \mathbb{E}[\mathcal{R}_T(\mathbf{0})] + \sum_{t=1}^T \mathbb{E}[\ell_t(\mathbf{0}) - \ell_t(\mathbf{u})] = O(1 + \|\mathbf{u}\|_2 LT).$$

Since $\|\mathbf{u}\|_2 > \frac{1}{\delta^3}$ the penalty for constraining v_t is $O(\|\mathbf{u}\|_2 LT) = O(\|\mathbf{u}\|_2^2 L \delta^3 T)$, which is $O(\|\mathbf{u}\|_2^2 L \sqrt{T})$ if we set $\delta = O(T^{-1/6})$. The formal result can be found below and its proof can be found in Section 4.8.1.

Theorem 16. Suppose \mathcal{A}_V implements Interface 6 with $l = (0, \frac{1}{\delta^3}]$, that $\ell_t(\mathbf{0}) = 0$, and that ℓ_t is β -smooth for all t . Then Algorithm 8 with $\delta = \min\{1, (dL)^{1/3}T^{-1/6}\}$, $\mathcal{Z} = \mathbb{B}$, $\alpha = 0$, and $\eta = \sqrt{\frac{\delta^2}{4(dL)^2T}}$ satisfies for all $\mathbf{u} \in \mathbb{R}^d$,

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u}) \right] \\ & \leq 2 + 2 \left(\|\mathbf{u}\|_2^2 dL \sqrt{T} + \|\mathbf{u}\|^2 (dLT)^{2/3} \beta \right) \\ & \quad + 6\|\mathbf{u}\| \left(\beta \sqrt{B(\|\mathbf{u}\|)} + 1 \right) \left((dL + 1)\sqrt{T} + ((dL)^{2/3} + (dL)^{-1/3})T^{2/3} \right). \end{aligned}$$

4.4.2 Constrained convex bandits

For the constrained setting we will set $\mathcal{Z} = \mathcal{W}$ and $\alpha = \delta$. This ensures that $v_t(\mathbf{z}_t + \delta \mathbf{s}_t) \in \mathcal{W}$ and we can apply Lemma 14 to find the regret bound in Theorem

17 below. Compared to the unconstrained setting, the regret bound now scales with c , which is due to the reshaping trick discussed in Flaxman et al. (2005).

Theorem 17. *Suppose \mathcal{A}_γ implements Interface 6 with $l = (0, 1]$ and that $\ell_t(\mathbf{0}) = 0$. Then Algorithm 8 with $\mathcal{Z} = \mathcal{W}$, $\alpha = \delta = \min\{1, \sqrt{dT}^{-1/4}\}$, and $\eta = \sqrt{\frac{\delta^2}{4(dL)^2T}}$ satisfies for all $\mathbf{u} \in \mathcal{W}$,*

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T(\mathbf{u})] = & 2 + \left(3\|\mathbf{u}\|Ld\sqrt{T} + 4\|\mathbf{u}\|L\sqrt{dT}^{\frac{3}{4}}\right) (1 + \sqrt{B(\|\mathbf{u}\|)}) \\ & + \|\mathbf{u}\|_2 dLT^{3/4}. \end{aligned}$$

4.5 Conclusion

In this chapter, we develop the first algorithms that have comparator-adaptive regret bounds for various bandit convex optimization problems. The regret bounds of our algorithms scale with $\|\mathbf{u}\|$, which may yield smaller regret in favourable settings.

For future research, there are a number of interesting open questions. First, our current results do not encompass improved rates for smooth losses on constrained domains. At first blush, one might feel this is relatively straightforward via methods based on self-concordance (Saha and Tewari, 2011), but it turns out that while such techniques provide good direction-learning algorithms, they may cause the gradients provided to the *scaling* algorithm to blow-up. Secondly, there is an important class of loss functions for which we did not obtain norm adaptive regret bounds: smooth and strongly convex losses. It is known that in this case an expected regret bound of $O(d\sqrt{T})$ can be efficiently achieved (Hazan and Levy, 2014). However, to achieve this regret bound the algorithm of Hazan and Levy (2014) uses a clever exploration scheme, which unfortunately leads to sub-optimal regret bounds for our algorithms.

4.6 Details from section 4.2

Proof of Lemma 12. By definition we have

$$\begin{aligned} \mathcal{R}_T(\mathbf{u}) &= \sum_{t=1}^T \langle \mathbf{w}_t - \mathbf{u}, \mathbf{g}_t \rangle = \sum_{t=1}^T \langle \mathbf{z}_t, \mathbf{g}_t \rangle (v_t - \|\mathbf{u}\|) + \|\mathbf{u}\| \sum_{t=1}^T \langle \mathbf{z}_t - \frac{\mathbf{u}}{\|\mathbf{u}\|}, \mathbf{g}_t \rangle \\ &= \mathcal{R}_T^{\mathcal{V}}(\|\mathbf{u}\|) + \|\mathbf{u}\| \mathcal{R}_T^{\mathcal{Z}}\left(\frac{\mathbf{u}}{\|\mathbf{u}\|}\right). \end{aligned}$$

□

4.7 Details from section 4.3

Proof of Theorem 14. For any fixed $\mathbf{u} \in \mathcal{W}$, let $r = \max_{\frac{r'\mathbf{u}}{\|\mathbf{u}\|} \in \mathcal{W}} r'$. Note that by definition we have $\frac{\|\mathbf{u}\|}{r} \in [0, 1]$ and $\frac{r\mathbf{u}}{\|\mathbf{u}\|} \in \mathcal{W}$. Therefore, similar to the proof of Lemma 12, we decompose the regret against \mathbf{u} as:

$$\begin{aligned} \mathcal{R}_T(\mathbf{u}) &= \sum_{t=1}^T \langle \mathbf{w}_t - \mathbf{u}, \mathbf{g}_t \rangle \\ &= \sum_{t=1}^T \langle \mathbf{z}_t, \mathbf{g}_t \rangle \left(v_t - \frac{\|\mathbf{u}\|}{r} \right) + \frac{\|\mathbf{u}\|}{r} \sum_{t=1}^T \langle \mathbf{z}_t - \frac{r\mathbf{u}}{\|\mathbf{u}\|}, \mathbf{g}_t \rangle, \end{aligned}$$

which, by the guarantees of $\mathcal{A}_{\mathcal{V}}$ and $\mathcal{A}_{\mathcal{Z}}$,⁴ is bounded in expectation by

$$\frac{\|\mathbf{u}\|}{r} L \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} + \frac{\|\mathbf{u}\|}{r} dL \sqrt{T \log(T+1)}.$$

Finally noticing $\frac{1}{c} \leq r$ by the definition of c finishes the proof. \square

4.8 Details from section 4.4

Proof of Lemma 14. Denote by $\tilde{\mathbf{w}}_t = v_t \mathbf{z}_t$. By Jensen's inequality we have

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] &= \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t(\mathbf{u})] + \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{w}_t)] \\ &\leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t(\mathbf{u})]. \end{aligned} \tag{4.8.1}$$

We now continue under the assumption that ℓ_t is L -Lipschitz. After completing the proof of the first equation of Lemma 14 we use the β -smoothness assumption to prove the second equation of Lemma 14.

⁴Note that the condition $|\langle \mathbf{z}_t, \mathbf{g}_t \rangle| \leq 1$ in Algorithm 7 indeed holds in this case since $\mathcal{Z} = \mathcal{W} \subseteq \mathbb{B}$ and $\|\mathbf{g}_t\|_2 \leq L$ by the Lipschitzness condition.

Using the L -Lipschitz assumption we proceed:

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{u})] + \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{u}) - \ell_t(\mathbf{u})] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E}[L|v_t| \|\delta \mathbf{s}_t\|_2] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E}[\delta L |v_t|] \\
 & = \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\tilde{\mathbf{w}}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E}[\delta L |v_t|] \\
 & \quad + \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\tilde{\mathbf{w}}_t)] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\tilde{\mathbf{w}}_t) - \ell_t^{v_t}(\mathbf{u})] + 2 \mathbb{E}[\delta L |v_t|].
 \end{aligned}$$

Now, by using the L -Lipschitz assumption once more we find that

$$\sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}((1-\alpha)\mathbf{u}) - \ell_t^{v_t}(\mathbf{u})] \leq \alpha \|\mathbf{u}\|_2 T L \quad (4.8.2)$$

By using equation (4.8.2), the convexity of $\ell_t^{v_t}$, and Lemma 13 we continue with:

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
& \leq \sum_{t=1}^T \mathbb{E} [\langle \tilde{\mathbf{w}}_t - (1 - \alpha)\mathbf{u}, \hat{\mathbf{g}}_t \rangle] + 2 \mathbb{E} [\delta L |v_t|] + \alpha \|\mathbf{u}\|_2 T L \\
& = \sum_{t=1}^T \mathbb{E} \left[\left(v_t - \frac{\|\mathbf{u}\|}{r} \right) \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle \right] + \mathbb{E} \left[\frac{\|\mathbf{u}\|}{r} \langle \mathbf{z}_t - \tilde{\mathbf{u}}, \hat{\mathbf{g}}_t \rangle \right] \\
& \quad + \sum_{t=1}^T 2 \mathbb{E} [\delta L |v_t|] + \alpha \|\mathbf{u}\|_2 T L \\
& = \sum_{t=1}^T \mathbb{E} \left[\bar{\ell}_t(v_t) - \bar{\ell}_t \left(\frac{\|\mathbf{u}\|}{r} \right) \right] + \sum_{t=1}^T \frac{\|\mathbf{u}\|}{r} \mathbb{E} [\langle \mathbf{z}_t - \tilde{\mathbf{u}}, \hat{\mathbf{g}}_t \rangle] \\
& \quad + 2T\delta L \frac{\|\mathbf{u}\|}{r} + \alpha \|\mathbf{u}\|_2 T L
\end{aligned}$$

where $\bar{\ell}_t(v) = v \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle + 2\delta L |v|$ as defined in Algorithm 8, $\tilde{\mathbf{u}} = \frac{r}{\|\mathbf{u}\|} (1 - \alpha)\mathbf{u}$, and $r > 0$ is such that $\frac{ur}{\|\mathbf{u}\|} \in \mathcal{Z}$.

Finally, by using the convexity of $\bar{\ell}_t$, plugging in the guarantee of \mathcal{A}_V , and using Theorem 18 we conclude the proof of the first equation of Lemma 14:

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
& \leq 2T\delta L \frac{\|\mathbf{u}\|}{r} + \mathbb{E} \left[\sum_{t=1}^T \left(v_t - \frac{\|\mathbf{u}\|}{r} \right) \partial \bar{\ell}_t(v_t) \right] \\
& \quad + \frac{\|\mathbf{u}\|}{r} \mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{z}_t - \tilde{\mathbf{u}}, \hat{\mathbf{g}}_t \rangle \right] + \alpha \|\mathbf{u}\|_2 T L \\
& \leq 1 + 2T\delta L \frac{\|\mathbf{u}\|}{r} + \frac{\|\mathbf{u}\|}{r} L_V \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} + \frac{2\|\mathbf{u}\|dL}{r\delta} \sqrt{T} + \alpha \|\mathbf{u}\|_2 T L.
\end{aligned}$$

Next, we continue from equation (4.8.1) under the smoothness condition. Using

the definition of smoothness we find

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{u})] + \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{u}) - \ell_t(\mathbf{u})] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E} \left[\frac{1}{2} \beta |v_t|^2 \|\delta \mathbf{s}_t\|_2^2 \right] \\
 & = \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E} \left[\frac{1}{2} \delta^2 |v_t|^2 \beta \right] \\
 & = \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\tilde{\mathbf{w}}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E} \left[\frac{1}{2} \delta^2 |v_t|^2 \beta \right] \\
 & \quad + \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\mathbf{w}_t) - \ell_t^{v_t}(\tilde{\mathbf{w}}_t)] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\ell_t^{v_t}(\tilde{\mathbf{w}}_t) - \ell_t^{v_t}(\mathbf{u})] + \mathbb{E} [\beta \delta^2 |v_t|^2].
 \end{aligned}$$

Using equation (4.8.2), the convexity of $\ell_t^{v_t}$, and Lemma 13 we continue with:

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
 & \leq \sum_{t=1}^T \mathbb{E} [\langle \tilde{\mathbf{w}}_t - (1 - \alpha)\mathbf{u}, \hat{\mathbf{g}}_t \rangle] + \mathbb{E} [\beta \delta^2 |v_t|^2] + \alpha \|\mathbf{u}\|_2 T L \\
 & = \sum_{t=1}^T \mathbb{E} \left[\left(v_t - \frac{\|\mathbf{u}\|}{r} \right) \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle \right] + \mathbb{E} [\beta \delta^2 |v_t|^2] \\
 & \quad + \sum_{t=1}^T \frac{\|\mathbf{u}\|}{r} \mathbb{E} [\langle \mathbf{z}_t - \tilde{\mathbf{u}}, \hat{\mathbf{g}}_t \rangle] + \alpha \|\mathbf{u}\|_2 T L \\
 & = T \beta \delta^2 \left(\frac{\|\mathbf{u}\|}{r} \right)^2 + \sum_{t=1}^T \mathbb{E} \left[\bar{\ell}_t(v_t) - \bar{\ell}_t \left(\frac{\|\mathbf{u}\|}{r} \right) \right] \\
 & \quad + \sum_{t=1}^T \frac{\|\mathbf{u}\|}{r} \mathbb{E} [\langle \mathbf{z}_t - \tilde{\mathbf{u}}, \hat{\mathbf{g}}_t \rangle] + \alpha \|\mathbf{u}\|_2 T L,
 \end{aligned}$$

where $\bar{\ell}_t(v) = v \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle + \beta \delta^2 v^2$ as defined in Algorithm 8. Finally, by using

the convexity of $\bar{\ell}_t$, plugging in the guarantee of \mathcal{A}_V , and using Theorem 18 we conclude the proof:

$$\begin{aligned}
& \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
& \leq T\beta\delta^2 \left(\frac{\|\mathbf{u}\|}{r} \right)^2 + \mathbb{E} \left[\sum_{t=1}^T \left(v_t - \frac{\|\mathbf{u}\|}{r} \right) \partial \bar{\ell}_t(v_t) \right] \\
& \quad + \frac{\|\mathbf{u}\|}{r} \mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{z}_t - \tilde{\mathbf{u}}, \hat{\mathbf{g}}_t \rangle \right] + \alpha \|\mathbf{u}\|_2 TL \\
& \leq 1 + T\beta\delta^2 \left(\frac{\|\mathbf{u}\|}{r} \right)^2 + \frac{\|\mathbf{u}\|}{r} L_V \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} + \frac{2\|\mathbf{u}\| dL}{r\delta} \sqrt{T} + \alpha \|\mathbf{u}\|_2 TL.
\end{aligned}$$

□

Theorem 18. Suppose that $\ell_t(\mathbf{0}) = 0$, that ℓ_t is L -Lipschitz for all t , and that $\mathcal{Z} \subseteq \mathbb{B}$. For $\mathbf{u} \in (1 - \alpha)\mathcal{Z}$, Online Gradient Descent on $(1 - \alpha)\mathcal{Z}$ with learning rate $\eta = \sqrt{\frac{\delta^2}{(dL)^2 4T}}$ satisfies

$$\mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{z}_t - \mathbf{u}, \hat{\mathbf{g}}_t \rangle \right] \leq 2 \frac{dL}{\delta} \sqrt{T}.$$

Proof. The proof essentially follows from the work of Zinkevich (2003); Flaxman et al. (2005) and using the assumptions that $\ell_t(\mathbf{0}) = 0$ and that ℓ_t is L -Lipschitz. We start by bounding the norm of the gradient estimate:

$$\begin{aligned}
\|\hat{\mathbf{g}}_t\|_2 &= \frac{d}{v_t \delta} |\ell_t(\mathbf{w}_t)| \|\mathbf{s}_t\|_2 \\
&= \frac{d}{v_t \delta} |\ell_t(v_t(\mathbf{z}_t + \delta \mathbf{s}_t)) - \ell_t(\mathbf{0})| \\
&\leq \frac{dL \|\mathbf{z}_t + \delta \mathbf{s}_t\|_2}{\delta} \leq \frac{dL(1 - \alpha + \delta)}{\delta}
\end{aligned} \tag{4.8.3}$$

By using equation (4.8.3) and the regret bound of Online Gradient Descent

(Zinkevich, 2003) we find that

$$\begin{aligned}
 \sum_{t=1}^T \langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle - \min_{\mathbf{z} \in (1-\alpha)\mathcal{Z}} \sum_{t=1}^T \langle \mathbf{z}, \hat{\mathbf{g}}_t \rangle &\leq \frac{(1-\alpha)}{2\eta} + \frac{\eta}{2} \sum_{t=1}^T \|\hat{\mathbf{g}}_t\|_2^2 \\
 &\leq \frac{(1-\alpha)}{2\eta} + \frac{\eta}{2} \left(\frac{dL(1-\alpha+\delta)}{\delta} \right)^2 T \\
 &\leq \frac{1}{2\eta} + 2\eta \left(\frac{dL}{\delta} \right)^2 T
 \end{aligned}$$

Plugging in $\eta = \sqrt{\frac{\delta^2}{(dL)^2 4T}}$ completes the proof. \square

4.8.1 Details of section 4.4.1

Proof of Theorem 15. First, since $\ell_t(\mathbf{0}) = 0$, ℓ_t is L -Lipschitz, and $\mathbf{z}_t \in (1-\alpha)\mathcal{Z} = (1-\alpha)\mathbb{B}$ we have that

$$\langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle \leq \|\mathbf{z}_t\|_2 \|\hat{\mathbf{g}}_t\|_2 \leq (1-\alpha) \frac{dL(1-\alpha+\delta)}{\delta} \leq \frac{2dL}{\delta}, \quad (4.8.4)$$

where the first inequality is the Cauchy-Schwarz inequality and the second is due to equation (4.8.3). Since $|\partial \bar{\ell}_t(v_t)| \leq |\langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle| + 2\delta L \leq \frac{4dL}{\delta} = L_{\mathcal{V}}$ we can use Lemma 14 to find

$$\begin{aligned}
 \mathbb{E}[\mathcal{R}_T(\mathbf{u})] &\leq 1 + 2T\delta L \frac{\|\mathbf{u}\|}{r} + \frac{\|\mathbf{u}\|}{r} \frac{4dL}{\delta} \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} \\
 &\quad + \frac{2\|\mathbf{u}\|dL}{r\delta} \sqrt{T} + \alpha \|\mathbf{u}\|_2 TL.
 \end{aligned}$$

Plugging in $\alpha = 0$ and $\delta = \min\{1, \sqrt{dT}^{-\frac{1}{4}}\}$ completes the proof. \square

Proof of Theorem 16. By equation (4.8.4) $|\langle \mathbf{z}_t, \hat{\mathbf{g}}_t \rangle| \leq \frac{2dL}{\delta}$. Since $v_t \leq \frac{1}{\delta^3}$ we have that

$$|\partial \bar{\ell}_t(v_t)| \leq \frac{2dL}{\delta} + 2|v_t|\beta\delta^2 \leq 2\frac{dL+\beta}{\delta} \leq \frac{2\beta(dL+1)}{\delta}$$

If $\|\mathbf{u}\|_2 \leq \frac{1}{\delta^3}$ applying Lemma 14 with $\alpha = 0$ gives us

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{E}[\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\
 \leq 1 + T\beta\delta^2 \|\mathbf{u}\|^2 + 2\|\mathbf{u}\| \frac{(dL+1)\beta}{\delta} \sqrt{TB(\|\mathbf{u}\|)} + \frac{2\|\mathbf{u}\|dL}{\delta} \sqrt{T}.
 \end{aligned} \quad (4.8.5)$$

If $\|\mathbf{u}\|_2 > \frac{1}{\delta^3}$ then using the Lipschitz assumption on ℓ_t and equation (4.8.5) with $\mathbf{u} = \mathbf{0}$ gives us

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] &= \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{0}) + \ell_t(\mathbf{0}) - \ell_t(\mathbf{u})] \\ &\leq 1 + \|\mathbf{u}\|_2 LT \\ &\leq 1 + \|\mathbf{u}\|_2^2 \delta^3 LT, \end{aligned} \tag{4.8.6}$$

where we used that $\|\mathbf{u}\|_2 \geq \frac{1}{\delta^3}$. Adding equations (4.8.5) and (4.8.6) gives

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\ \leq 2 + \|\mathbf{u}\|_2^2 \delta^3 LT + T\beta\delta^2 \|\mathbf{u}\|^2 + 2\|\mathbf{u}\| \frac{(dL+1)\beta}{\delta} \sqrt{TB(\|\mathbf{u}\|)} + \frac{2\|\mathbf{u}\|dL}{\delta} \sqrt{T} \end{aligned}$$

Setting $\delta = \min\{1, (dL)^{1/3} T^{-1/6}\}$ gives us

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\ \leq 2 + 2 \left(\|\mathbf{u}\|_2^2 dL \sqrt{T} + \|\mathbf{u}\|^2 (dLT)^{2/3} \beta \right) \\ + 6\|\mathbf{u}\| \left(\beta \sqrt{B(\|\mathbf{u}\|)} + 1 \right) \left((dL+1) \sqrt{T} + ((dL)^{2/3} + (dL)^{-1/3}) T^{2/3} \right) \end{aligned}$$

□

4.8.2 Details of section 4.4.2

Proof of Theorem 17. First, to see that $\mathbf{z}_t + \delta \mathbf{s}_t \in \mathcal{W}$ recall that by assumption $\mathcal{W} \subseteq \mathbb{B}$. Since $\alpha = \delta$ we have that $\mathbf{z}_t + \delta \mathbf{s}_t \in (1-\alpha)\mathcal{W} + \delta \mathbb{S} \subseteq (1-\delta)\mathcal{W} + \delta \mathcal{W} = \mathcal{W}$. For any fixed $\mathbf{u} \in \mathcal{W}$, let $r = \max_{\frac{r'\mathbf{u}}{\|\mathbf{u}\|} \in \mathcal{W}} r'$. Note that by definition we have $\frac{\|\mathbf{u}\|}{r} \in [0, 1]$ and $\frac{r\mathbf{u}}{\|\mathbf{u}\|} \in \mathcal{W}$. By using equation (4.8.4) we can see that $|\partial \bar{\ell}_t(v_t)| \leq \frac{dL}{\delta} + 2\delta L$. By definition, $\frac{1}{r} \leq c$. This implies that the regret of \mathcal{A}_V is $\tilde{O} \left(1 + \frac{\|\mathbf{u}\|}{r} \frac{dL}{\delta} \sqrt{T} \right)$. Applying Lemma 14 with the parameters above we find

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \\ = 1 + 2T\delta L \frac{\|\mathbf{u}\|}{r} + \frac{3dL\|\mathbf{u}\|}{r} \sqrt{TB \left(\frac{\|\mathbf{u}\|}{r} \right)} + \frac{2\|\mathbf{u}\|dL}{r\delta} \sqrt{T} + \alpha \|\mathbf{u}\|_2 TL. \end{aligned}$$

Finally, setting $\delta = \min\{1, \sqrt{d}T^{-1/4}\}$ completes the proof:

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} [\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{u})] \leq & 2 + \left(3\|\mathbf{u}\|Ld\sqrt{T} + 4\|\mathbf{u}\|L\sqrt{dT}^{\frac{3}{4}} \right) (1 + \sqrt{B(\|\mathbf{u}\|)}) \\ & + \|\mathbf{u}\|_2 dLT^{3/4}. \end{aligned}$$

□