Online Improper Learning with an Approximation Oracle*

Elad Hazan Princeton University & Google AI Princeton ehazan@cs.princeton.edu

Yuanzhi Li Stanford University yuanzhil@stanford.edu Wei Hu Princeton University huwei@cs.princeton.edu

Zhiyuan Li Princeton University zhiyuanli@cs.princeton.edu

Abstract

We study the following question: given an efficient approximation algorithm for an optimization problem, can we learn efficiently in the same setting? We give a formal affirmative answer to this question in the form of a reduction from online learning to offline approximate optimization using an efficient algorithm that guarantees near optimal regret. The algorithm is efficient in terms of the number of oracle calls to a given approximation oracle – it makes only logarithmically many such calls per iteration. This resolves an open question by Kalai and Vempala, and by Garber. Furthermore, our result applies to the more general improper learning problems.

1 Introduction

A fundamental question in learning theory is whether one can *efficiently* learn a given problem using an optimization oracle. Namely, does efficient offline optimization for a certain problem imply efficient learning algorithm for the same setting?

For online learning in games, it was shown by Kalai and Vempala (2005) that an optimization oracle giving the best decision in hindsight is sufficient for attaining optimal regret. However, in many non-convex settings, such an optimization oracle is either unavailable or NP-hard to compute. In the face of NP-hardness, algorithm designers resort to approximation algorithms that are guaranteed to return a solution within a certain multiplicative factor of the optimum. We give numerous examples in Section 1.2.

Kakade et al. (2009) considered the question of whether such an approximation algorithm is sufficient to obtain vanishing regret compared with an approximation to the best solution in hindsight. They gave an algorithm for this offline-to-online conversion. However, their reduction is inefficient in the number of per-iteration queries to the approximation oracle, which grows linearly with time. Ideally, an efficient reduction should call the oracle only a constant number of times per iteration and guarantee optimal regret at the same time, and this was considered an open question in the literature.

Various authors have improved upon this original online-to-offline reduction under certain cases, as we survey below. Recently, Garber (2017) made significant progress by giving a more efficient reduction, which improves the number of oracle calls in both full information and bandit settings. He explicitly asked whether a near-optimal reduction with only logarithmically many calls per iteration exists.

^{*}The full version of this paper can be found on https://arxiv.org/abs/1804.07837.

1.1 Problem Setting and Our Results

In this paper we resolve this question on the positive side and in a more general setting, which we formally define now.

Formal description of problem setting. We consider the standard setting of online linear optimization which is known to generalize statistical learning (Hazan, 2016; Shalev-Shwartz, 2012). In a repeated game, in round t a player chooses a point x_t from a decision set $\mathcal{K} \subseteq \mathbb{R}^d$ while an adversary chooses a loss vector $f_t \in \mathbb{R}^d$, which determines the loss of the player $f_t^{\top} x_t$ in this round. The loss vector f_t is revealed to the player *after* her choice x_t is made. We sometimes treat f_t as a function on \mathbb{R}^d , i.e., $f_t(x) := f_t^{\top} x$.

Since we consider computationally intractable problems like maximum cut or minimum-rank matrix completion, we assume that the player has access to an offline optimization oracle. This oracle may return a point which does not belong to the target set \mathcal{K}^* , but rather to a different set \mathcal{K} . For example, in matrix completion the oracle may return a low-trace-norm matrix rather than a low-rank matrix.

This notion is formally captured by an optimization oracle $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$. Given an input $v \in \mathbb{R}^d$, this oracle outputs a point $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}(v) \in \mathcal{K}$ which dominates all points in \mathcal{K}^* in the direction v, that is

$$v^{\top}\mathcal{O}_{\mathcal{K},\mathcal{K}^{*}}(v) \leq \min_{x^{*}\in\mathcal{K}^{*}} v^{\top}x^{*}$$

The goal for the player is to minimize her *regret*, which is the difference between her cumulative loss and that of the best single decision (in \mathcal{K}^*) in hindsight:

$$\operatorname{Reg}_{\mathcal{K},\mathcal{K}^*}(T) := \sum_{t=1}^T f_t^\top x_t - \min_{x^* \in \mathcal{K}^*} \sum_{t=1}^T f_t^\top x^*.$$

We remark that the above problem setting is similar in spirit to the notion of *improper learning*, where one is allowed to output a hypothesis not from the target set. Therefore, we view the problem setting described above as an *online* version of improper learning.

In the special case of $\mathcal{K}^* = \alpha \mathcal{K}$ ($\alpha > 1$), $\mathcal{O}_{\mathcal{K},\alpha\mathcal{K}}$ becomes an α -approximation oracle on \mathcal{K} , and the setting and the notation of regret are the same with those studied in (Kakade et al., 2009; Garber, 2017), i.e., $\operatorname{Reg}_{\mathcal{K},\alpha\mathcal{K}}(T) := \sum_{t=1}^{T} f_t^{\top} x_t - \alpha \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t^{\top} x$. This is called the α -regret.

Our results. In this setting, we give two different algorithms, one based on the online mirror descent (OMD) method and another based on the continuous multiplicative weight update (CMWU) algorithm. Both of them give nearly optimal regret as well as oracle efficiency, while applying to general loss vectors. Our results are summarized in Table 1 below. We present these two algorithms and their guarantees in Sections 3 and Appendix B.

Algorithm	Regret over T rounds	Oracle calls per round	Loss vectors
Kakade et al. (2009)	$O(\sqrt{T})$	O(T)	general
Garber (2017)	$O(\sqrt{T})$	$\tilde{O}(\sqrt{T})$	non-negative
Alg. 1 (this paper)	$O(\sqrt{T})$	$O(\log T)$	PNIP property (Def. 2.4)
Alg. 6 (this paper)	$ ilde{O}(\sqrt{T})$	$O(\log T)$	general

Table 1: Summary of results in the full information setting. The \tilde{O} notation hides constant and logarithmic factors.

Algorithm	Regret over T rounds	Oracle calls in T rounds	Loss vectors
Kakade et al. (2009)	$O(T^{\frac{2}{3}})$	$O(T^{\frac{4}{3}})$	general
Garber (2017)	$O(T^{\frac{2}{3}})$	$ ilde{O}(T)$	non-negative
Alg. 3 (this paper)	$O(T^{rac{2}{3}})$	$ ilde{O}(T^{rac{2}{3}})$	non-negative

Table 2: Summary of results in the bandit setting.

In addition to these two algorithms, we give an improved result in the bandit setting. In this more difficult setting, the player cannot observe f_t , but rather only the loss she has suffered, namely the

scalar $f_t^{\top} x_t$. We show how to extend our mirror descent-based algorithm to the bandit setting and obtain the same $O(T^{2/3})$ regret as in (Kakade et al., 2009; Garber, 2017), but with a significantly lower computational cost. See Table 2 for a comparison. We present our bandit result in Section 4.

1.2 Applications

The setting of online learning with approximation algorithms has been well studied since (Kalai and Vempala, 2005) with numerous applications.

For example, in the online max-cut problem, a learner iteratively predicts a cut over a set of vertices V, and afterwards the adjacency information for two vertices is revealed. The loss is zero or one, depending on whether the learner correctly predicted the connectivity of the two vertices. The offline version of this problem is NP-hard, but admits SDP-based approximation algorithms such as the famous 0.878-approximation by Goemans and Williamson (1995). Our results imply an online algorithm that can predict as accurate as the best 0.878-approximation to the maximum cut in hindsight, and calls the SDP relaxation only logarithmically many times per iteration.

Numerous other examples exist for combinatorial graph optimization problems such as the traveling salesman problem, sparsest graph cut, etc. Other applications include prominent machine learning problems whose offline optimization problem is NP-hard, for example, matrix completion and recommendation systems. The reader is referred to (Kakade et al., 2009; Garber, 2017) for more detailed exposition of applications.

1.3 Related Work

The reduction from online learning to offline approximation algorithms was already considered by Kalai and Vempala (2005). Their scheme, based on the follow-the-perturbed-leader (FTPL) algorithm, requires very strong approximation guarantee from the approximation oracle, namely, a fully polynomial time approximation scheme (FPTAS), and requires an approximation that improves with time. Balcan and Blum (2006) used the same approach in the context of mechanism design.

Kalai and Vempala (2005) also proposed a specialized reduction that works under certain conditions on the approximation oracle, satisfied by some known algorithms for problems such as MAX-CUT. Fujita et al. (2013) further gave more general reductions that apply to problems whose approximation algorithms are based on convex relaxations of mathematical programs. Their scheme is also based on the FTPL method.

Recent advancements on black-box online-to-offline reductions were made in (Kakade et al., 2009; Dudík et al., 2016; Garber, 2017). Hazan and Koren (2016) showed that efficient reductions are in general impossible, unless special structure is present. In the settings we consider this special structure is a linear cost function over the space.

Our algorithms fall into one of two templates. The first is the online mirror descent method, which is an adaptive version of the follow-the-regularized-leader (FTRL) algorithm. The second is the continuous multiplicative weight update method, which dates back to Cover's portfolio selection method (Cover, 1991) and Vovk's aggregating algorithm (Vovk, 1990). The reader is referred to the books (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2012; Hazan, 2016) for details and background on these prediction frameworks. We also make use of polynomial-time algorithms for sampling from log-concave distributions (Lovász and Vempala, 2007).

2 Preliminaries

For $x \in \mathbb{R}^d$ and r > 0, denote by B(x, r) the Euclidean ball in \mathbb{R}^d of radius r centered at x. For $S, S' \subseteq \mathbb{R}^d$, $\beta \in \mathbb{R}$, $y \in \mathbb{R}^d$ and $A \in \mathbb{R}^{d' \times d}$, define $S + S' := \{x + x' : x \in S, x' \in S'\}$, $\beta S := \{\beta x : x \in S\}$, $x + S := \{x + y : y \in S\}$, and $AS := \{Ax : x \in S\}$. The convex hull of $S \subseteq \mathbb{R}^d$ is denoted by CH(S). Denote by Vol(S) the volume (Lebesgue measure) of a set $S \subseteq \mathbb{R}^d$. Denote by Δ^{k-1} the probability simplex in \mathbb{R}^k .

A set $C \subseteq \mathbb{R}^d$ is called a *cone* if for any $\beta \ge 0$ we have $\beta C \subseteq C$. For any $S \subseteq \mathbb{R}^d$, define the *dual cone* of S as $S^\circ := \{y \in \mathbb{R}^d : x^\top y \ge 0, \forall x \in S\}$. S° is always a convex cone, even when S is

Algorithm 1 Online Mirror Descent using a Projection-and-Separation Oracle

Input: Learning rate $\eta > 0$, tolerance $\epsilon > 0$, regularizer φ , convex cone W, time horizon $T \in \mathbb{N}_+$ 1: $y_1 \leftarrow \arg\min_{y \in \text{Dom}(\varphi)} \varphi(y)$. 2: **for** t = 1 **to** T **do** 3: $(x_t, V = (v_1, \dots, v_k), p) \leftarrow \mathcal{PAD}(y_t, \epsilon, W, \varphi)$ 4: Play $\tilde{x}_t = v_i$ with probability p_i $(i \in [k])$, and observe the loss vector f_t . 5: $\nabla \varphi(y_{t+1}) \leftarrow \nabla \varphi(x_t) - \eta f_t$

6: end for

neither convex nor a cone. For any closed set $S \subseteq \mathbb{R}^d$, define $\Pi_S : \mathbb{R}^d \to S$ to be the *projection* onto S, namely $\Pi_S(x) := \arg \min_{x' \in S} \|x' - x\|^2$.

Definition 2.1. A strictly convex function $f : \mathcal{A} \to \mathbb{R}$ ($\mathcal{A} \subseteq \mathbb{R}^d$ is convex) is Legendre if ∇f is continuous in $int(\mathcal{A})$ and for any sequence $x_1, x_2, \dots \in \mathcal{A}$ converging to a boundary point of \mathcal{A} , $\lim_{n\to\infty} \|\nabla f(x_n)\| = \infty$.

Definition 2.2. For a Legendre function $\varphi : \mathcal{A} \to \mathbb{R}$, the Bregman divergence with respect to φ is defined as $D_{\varphi}(x, y) := \varphi(x) - \varphi(y) - \nabla \varphi(y)^{\top}(x - y)$ ($\forall x, y \in \mathcal{A}$).

Lemma 2.3 (Generalized Pythagorean theorem, see e.g. Lemma 11.3 in (Cesa-Bianchi and Lugosi, 2006)). For any closed convex set $S \subseteq \mathbb{R}^d$, $x \in \mathbb{R}^d$, $y \in S$, and any Legendre function $\varphi : \mathbb{R}^d \to \mathbb{R}$, letting $z = \arg \min_{x' \in S} D_{\varphi}(x', x)$, we must have $D_{\varphi}(y, x) \ge D_{\varphi}(y, z) + D_{\varphi}(z, x)$.

Definition 2.4 (Pairwise non-negative inner product). For a twice-differentiable Legendre function $\varphi : \mathcal{A} \to \mathbb{R}$ with domain $\mathcal{A} \subseteq \mathbb{R}^d$ and a convex cone $W \subseteq \mathbb{R}^d$, we say (φ, W) satisfies the pairwise non-negative inner product (PNIP) property, if for all $w, w' \in W$ and $H \in CH(\mathcal{H})$, where $\mathcal{H} = \{\nabla^2 \varphi(x) : x \in \mathcal{A}\}$, it holds that $w^\top H^{-1}w' \geq 0$.

Examples. (φ, W) satisfies the PNIP property if:

- 1. $\varphi(x) = \frac{1}{2} ||x||^2$, $x \in \mathbb{R}^d$ and $W \subseteq W^\circ$, such as the non-negative orthant \mathbb{R}^d_+ , the positive semidefinite matrix cone, and the Lorentz cone $L_{d+1} = \{(x, z) \in \mathbb{R}^d \times \mathbb{R} : ||x||_2 \le z\}$.
- 2. $\varphi(x) = \sum_{i=1}^{d} x_i (\log x_i 1)$ (with domain \mathbb{R}^d_+) and $W = \mathbb{R}^d_+$;
- 3. $\varphi(x) = \frac{1}{2}x^{\top}Q^{-1}x$ (with domain \mathbb{R}^d), where $Q = MM^{\top}$, $M \in \mathbb{R}^{d \times d}$ is an invertible matrix, and $W = (M^{\top})^{-1}\mathbb{R}^d_+$.

This is useful in our bandit algorithm in Section 4.

Log-concave distributions. A distribution over \mathbb{R}^d with a density function f is *log-concave* if $\log(f)$ is a concave function. For a convex set S equipped with a membership oracle, there exist polynomial-time algorithms for sampling from any log-concave distribution over S (Lovász and Vempala, 2007). This can be used to approximately compute the mean of any log-concave distribution. For ease of presentation, we will assume that we can compute the mean of bounded-supported log-concave distributions exactly. Detailed explanation is provided in Appendix D.

3 Mirror Descent with an Approximation Oracle

In this section, we give an efficient online improper linear optimization algorithm (Algorithm 1) in the full information setting based on online mirror descent (OMD) equipped with a strongly convex regularizer φ , which achieves $O(\sqrt{T})$ regret when the regularizer φ and the domain of linear loss functions W satisfy the *pairwise non-negative inner product (PNIP)* property (Definition 2.4).

We suppose $\mathcal{K}, \mathcal{K}^* \subseteq B(0, R)$, and the loss vectors $\{f_t\}$ come from a convex cone $W \subseteq \mathbb{R}^d$ and $\|f_t\| \leq L(R, L > 0)$. Omitted proofs in this section are given in Appendix A.

Theorem 3.1. Suppose (φ, W) satisfies the PNIP property (Definition 2.4). Then for any $\epsilon, \eta > 0$, Algorithm 1 satisfies the following regret guarantee:

$$\forall x^* \in \mathcal{K}^* : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le \frac{1}{\eta} \left(\varphi(x^*) - \varphi(y_1) + \sum_{t=1}^T D_\varphi(x_t, y_{t+1})\right) + \epsilon LT.$$

In particular, if φ is μ -strongly convex and $A \ge \max_{x^* \in \mathcal{K}^*} (\varphi(x^*) - \varphi(y_1))$, setting $\epsilon = \frac{R}{T}$ and $\eta = \frac{1}{L} \sqrt{\frac{2\mu A}{T}}$, we have

$$\forall x^* \in \mathcal{K}^* : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le L\sqrt{\frac{2AT}{\mu}} + LR,$$

and in this case, Algorithm 1 makes at most $\left\lceil 5d \log \left(\left(6\sqrt{T} + \frac{4}{R}\sqrt{\frac{A}{\mu}} + 4 \right)T \right) \right\rceil$ calls of $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ per round.

For the problem of α -regret minimization using an α -approximation oracle, we have the following regret guarantee, which is an immediate corollary of Theorem 3.1.

Corollary 3.2. If $W \subseteq \mathbb{R}^d_+$, $\mathcal{K} \subseteq B(0, R)$, $\mathcal{K}^* = \alpha \mathcal{K}$, $\varphi(x) = \frac{1}{2} \|x\|^2$, setting $\epsilon = \frac{\alpha R}{T}$, $\eta = \frac{\alpha R}{L\sqrt{T}}$, Algorithm 1 has the following regret guarantee:

$$\forall x^* \in \mathcal{K} : \quad \mathbb{E}\left[\sum_{t=1}^T f_t(\tilde{x}_t) - \alpha \sum_{t=1}^T f_t(x^*)\right] = \mathbb{E}\left[\sum_{t=1}^T f_t(\tilde{x}_t) - \sum_{t=1}^T f_t(\alpha x^*)\right] \le \alpha LR(\sqrt{T} + 1).$$

Algorithm 1 is a variant of the OMD algorithm that makes use of a *projection-and-decomposition* (*PAD*) oracle, defined as follows:

Definition 3.3 (Projection-and-decomposition oracle). A projection-and-decomposition (PAD) oracle onto \mathcal{K}^* , $\mathcal{PAD}(y, \epsilon, W, \varphi)$, is defined as a procedure that given $y \in \mathbb{R}^d$, $\epsilon > 0$, a convex cone Wand a Legendre function φ produces a tuple (y', V, p), where $y' \in \mathbb{R}^d$, $V = (v_1, \ldots, v_k) \in \mathbb{R}^{d \times k}$ and $p = (p_1, \ldots, p_k)^\top \in \Delta^{k-1}$, such that:

- 1. y' is "closer" to \mathcal{K}^* than y with respect to the Bregman divergence of φ (and hence is an "infeasible projection"): $\forall x^* \in \mathcal{K}^*, \ D_{\varphi}(x^*, y') \leq D_{\varphi}(x^*, y);$
- 2. $v_1, \ldots, v_k \in \mathcal{K}$, and $\sum_{i=1}^k p_i v_i$ is a point that "almost dominates" y' in all directions in W. In other words, there exists $c \in W^\circ$ such that $\|\sum_{i=1}^k p_i v_i + c y'\| \le \epsilon$.

The purpose of the PAD oracle is the following. Suppose the OMD algorithm tells us to play a point y. Since y might not be in the feasible set \mathcal{K} , we can call the PAD oracle to find another point y' as well as a distribution p over points $v_1, \ldots, v_k \in \mathcal{K}$. The first property in Definition 3.3 is sufficient to ensure that playing y' also gives low regret, and the second property further ensures that we have a distribution of points in \mathcal{K} that suffers less loss than y' for every possible loss function so we can play according to that distribution.

Assuming the availability of a PAD oracle, one can use a standard analysis of OMD to prove a regret bound for Algorithm 1 as in Theorem 3.1. The proof is given in Appendix A.

Next we show how to construct a PAD oracle using the optimization oracle $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$. Our construction is given in Algorithm 2. Theorem 3.4 gives its guarantee.

Theorem 3.4. Suppose (φ, W) satisfies PNIP condition (Definition 2.4) and φ is μ -strongly convex. Then for any $y \in \mathbb{R}^d$, $\epsilon \in (0, R]$, Algorithm 2 terminates in $\left[5d \log \frac{4R + 2\sqrt{2 \min_{x^* \in \mathcal{K}^*} D_{\varphi}(x^*, y)/\mu}}{\epsilon} \right]$

iterations, and it correctly implements the projection-and-decomposition oracle $\mathcal{PAD}(y, \epsilon, W, \varphi)$, *i.e., its output* (y', V, p) satisfies the two properties in Definition 3.3.

Remark. We can use random walk methods to compute an $\frac{1}{T}$ -approximation of the gravity center (line 3 in Algorithm 2) in poly(T) time, which is enough for the purpose of bounding regret. We can also replace the center of gravity method with the ellipsoid method, or any other optimization method with a similar "optimization interface" (i.e., any method that is based on separation queries and guarantees similar bounds on the number of iterations required to find a feasible point), as pointed out by Garber (2017). Specifically, using the ellipsoid method, we can significantly reduce the computational complexity to depend only polynomially in log T (rather than T), at the cost of a slightly higher oracle complexity, namely $O(d^2 \log T)$ calls to the oracle per round. We choose center of gravity over other optimization methods only because it has the best oracle complexity, which is the main focus of this paper.

Algorithm 2 Projection-and-Decomposition Oracle, $\mathcal{PAD}(y, \epsilon, W, \varphi)$

Input: Point $y \in \mathbb{R}^d$, tolerance $\epsilon > 0$, convex cone W, regularizer φ , Output: (y', V, p), where $y' \in \mathbb{R}^d$, $V = (v_1, \dots, v_k) \in \mathbb{R}^{d \times k}$ for some k such that $v_i \in \mathcal{K}$ $(\forall i \in [k])$, and $p = (p_1, \dots, p_k)^\top \in \Delta^{k-1}$ 1: $W_1 \leftarrow W \cap B(0, 1), \quad z_1 \leftarrow y, \quad i \leftarrow 0$ 2: while $i < 5d \log \frac{2(R+||z_{i+1}||)}{\epsilon}$ do 3: $i \leftarrow i+1, \quad w_i \leftarrow \frac{\int_{W_i} w dw}{\operatorname{Vol}(W_i)}, \quad v_i \leftarrow \mathcal{O}_{\mathcal{K},\mathcal{K}^*}(w_i).$ 4: $z_{i+1} \leftarrow \underset{z \in \mathbb{R}^d, w_i^\top(z-v_i) \geq 0}{\operatorname{min}} D_{\varphi}(z, z_i), \quad W_{i+1} \leftarrow W_i \cap \{w \in \mathbb{R}^d : w^\top(v_i - z_{i+1}) \geq 0\}.$ 5: end while 6: $k \leftarrow i$ and solve $\underset{p \in \Delta^{k-1}, c \in W^\circ}{\min} \|\sum_{i=1}^k p_i v_i + c - z_{k+1}\|$ to get p7: return $y' = z_{k+1}, V = (v_1, \dots, v_k), p$

We break the proof of Theorem 3.4 into several lemmas.

Lemma 3.5. If (φ, W) satisfies the PNIP condition (Definition 2.4), then z_1, \ldots, z_{k+1} computed in Algorithm 2 satisfy $z_{i+1} - z_i \in W^\circ$ for all $i \in [k]$.

Proof. Since we have $z_{i+1} = \underset{z \in \mathbb{R}^d: w_i^{\top}(z-v_i) \ge 0}{\operatorname{arg\,min}} D_{\varphi}(z, z_i)$, by the KKT condition, we have

$$0 = \frac{\partial}{\partial z} \left(D_{\varphi}(z, z_i) - \lambda w_i^{\top}(z - v_i) \right) \Big|_{z = z_{i+1}} = \nabla \varphi(z_{i+1}) - \nabla \varphi(z_i) - \lambda w_i^{\top}(z - v_i) \right)$$

for some $\lambda \geq 0$. On the other hand, note that $\nabla \varphi(z_{i+1}) - \nabla \varphi(z_i) = \int_0^1 \nabla^2 \varphi(\gamma z_{i+1} + (1-\gamma)z_i) \cdot (z_{i+1} - z_i) d\gamma = H(z_{i+1} - z_i)$, for some $H \in CH(\mathcal{H})$, where $\mathcal{H} = \{\nabla^2 \varphi(x) : x \in Dom(\varphi)\}$. Therefore, for all $w \in W$ we have $w^\top(z_{i+1} - z_i) = w^\top H^{-1}H(z_{i+1} - z_i) = \lambda w^\top H^{-1}w_i \geq 0$. This means $z_{i+1} - z_i \in W^\circ$.

Lemma 3.6. Under the setting of Theorem 3.4, Algorithm 2 terminates in at most $\left[5d \log \frac{4R+2\sqrt{2\min_{x^* \in \mathcal{K}^*} D_{\varphi}(x^*,y)/\mu}}{\epsilon} \right]$ iterations.

Proof. According to the algorithm, for each i, z_{i+1} is the Bregman projection of z_i onto a half-space containing \mathcal{K}^* , since the oracle $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ ensues $w_i^\top v_i \leq w_i^\top x^*$ for all $x^* \in \mathcal{K}^*$. Then by the generalized Pythagorean theorem (Lemma 2.3) we know $D_{\varphi}(x^*, z_{i+1}) \leq D_{\varphi}(x^*, z_i)$ for all $x^* \in \mathcal{K}^*$ and i. Therefore we have $D_{\varphi}(x^*, z_i) \leq D_{\varphi}(x^*, z_1) = D_{\varphi}(x^*, y)$ for all $x^* \in \mathcal{K}^*$ and i.

Let $P := \min_{x^* \in \mathcal{K}^*} D_{\varphi}(x^*, y)$. Then there exists $x^* \in \mathcal{K}^*$ such that $P = D_{\varphi}(x^*, y) \geq D_{\varphi}(x^*, z_i) \geq \frac{\mu}{2} \|x^* - z_i\|^2$ for all i, where the last inequality is due to the μ -strong convexity of φ . This implies $\|z_i\| \leq \|x^*\| + \sqrt{\frac{2P}{\mu}} \leq R + \sqrt{\frac{2P}{\mu}}$ for all i. Therefore, when $i \geq 5d \log \frac{4R + 2\sqrt{2P/\mu}}{\epsilon}$, we must have $i \geq 5d \log \frac{2(R+\|z_{i+1}\|)}{\epsilon}$, which means the loop must have terminated at this time. \Box

Lemma 3.7. Under Theorem 3.4's setting, $\forall w \in W, ||w|| = 1, \exists i \in [k]$, such that $w^{\top}(v_i - y') \leq \epsilon$.

Proof. We assume for contradiction that there exists a unit vector $h \in W$ such that $\min_{i \in [k]} h^{\top}(v_i - y') > \epsilon$. Note that $||v_i - y'|| \le ||v_i|| + ||y'|| \le R + ||y'||$. Letting $r := \frac{\epsilon}{2(R+||y'||)}$, we have $\forall w \in \frac{h}{2} + (W \cap B(0, r)) : \min_{i \in [k]} w^{\top}(v_i - y') > 0$.

Since $r \leq \frac{1}{2}$ for $\epsilon \leq R$, we have $\frac{h}{2} + (W \cap B(0, r)) \subseteq \frac{h}{2} + (W \cap B(0, 1/2)) \subseteq W \cap B(0, 1) = W_1$. By the algorithm, we know that for all $w \in W_1 \setminus W_{k+1}$, there exists $i \in [k]$ such that $w^{\top}(v_i - z_{i+1}) \leq 0$. Notice that from Lemma 3.5 we know $z_{j+1} - z_j \in W^\circ$ for all $j \in [k]$. Thus for all $w \in W_1 \setminus W_{k+1}$ there exists $i \in [k]$ such that $w^{\top}(v_i - y') = w^{\top}(v_i - z_{k+1}) \leq w^T(v_i - z_{i+1}) \leq 0$. In other words, we have $\forall w \in W_1 \setminus W_{k+1}$: $\min_{i \in [k]} w^{\top}(v_i - y') \leq 0$.

Therefore, we must have $\frac{h}{2} + (W \cap B(0,r)) \subseteq W_{k+1}$. We also have $Vol(W_{i+1}) \leq (1 - 1/(2e))Vol(W_i)$ for each $i \in [k]$ from Lemma D.2, since W_{i+1} is the intersection of W_i with a half-space that does not contain W_i 's centroid w_i in the interior. Then we have

$$Vol(W_1) = Vol(W \cap B(0, 1)) = r^{-d}Vol(W \cap B(0, r)) \le r^{-d}Vol(W_{k+1}) \le r^{-d}(1 - 1/(2e))^k Vol(W_1) < Vol(W_1),$$

where the last step is due to $k \ge 5d \log \frac{1}{r} = 5d \log \frac{2(R+||y'||)}{\epsilon} = 5d \log \frac{2(R+||z_{k+1}||)}{\epsilon}$, which is true according to the termination condition of the loop. Therefore we have a contradiction.

The following lemma is a more general version of Lemma 6 in (Garber, 2017).

Lemma 3.8. Given $v_1, \ldots, v_k \in \mathbb{R}^d$, $\epsilon \ge 0$ and a convex cone $W \in \mathbb{R}^d$, for any $x \in \mathbb{R}^d$, the following two statements are equivalent:

- (A) There exists $p = (p_1, \ldots, p_k)^\top \in \Delta^{k-1}$ and $c \in W^\circ$ such that $\|\sum_{i=1}^k p_i v_i + c x\| \le \epsilon$.
- (B) For all $w \in W$, ||w|| = 1, there exists $i \in [k]$ such that $w^{\top}(v_i x) \leq \epsilon$.



(a) Convex cone W and its dual cone W°

(b) An example for $CH(V) + W^{\circ}$, where $V = \{v_i\}_{i=1}^4$.

Figure 1: Geometric Interpretation of Lemma 3.8.

Geometric interpretation of Lemma 3.8. We defer the proof of Lemma 3.8 to Appendix A, and discuss its geometric intuition here. For simplicity of illustration, we only consider $\epsilon = 0$ here (Figure 1). First we look at the case where $W = \mathbb{R}^d$, $W^\circ = \{0\}$. In this case the lemma simply degenerated to the fact

 $x \in CH(\{v_i\}_{i=1}^k) \iff$ There is no hyperplane that separates x and all v_i 's.

In the general case where $W \subseteq \mathbb{R}^d$ is an arbitrary convex cone, lemma 3.8 becomes

$$x \in CH(\{v_i\}_{i=1}^k) + W^{\circ} \iff$$
 There is no direction $w \in W$ such that $w^{\top} x < w^{\top} v_i$ for all i.

Denote $F := CH(\{v_i\}_{i=1}^k) + W^\circ$. For the " \Rightarrow " side, if $x \in F$, it is clear that for all $w \in W$ we must have $w^\top x \ge w^\top v_i$ for some *i*. For the " \Leftarrow " side, if $x \notin F$, then $w = \Pi_F(x) - x$ satisfies $w^\top x < w^\top v_i$ for all *i*. Moreover it is easy to see $\Pi_F(x) - x \in W$, which completes the proof.

Theorem 3.4 can be proved now using the above lemmas.

Proof of Theorem 3.4. The upper bound on the number of iterations is proved in Lemma 3.6. In the proof of Lemma 3.6, we have shown $D_{\varphi}(x^*, z_{i+1}) \leq D_{\varphi}(x^*, z_i)$ for all $x^* \in \mathcal{K}^*$ and *i*. This implies $D_{\varphi}(x^*, y') = D_{\varphi}(x^*, z_{k+1}) \leq D_{\varphi}(x^*, z_k) \leq \cdots \leq D_{\varphi}(x^*, z_1) = D_{\varphi}(x^*, y)$ for all $x^* \in \mathcal{K}^*$, which verifies the first property in Definition 3.3. The second property is a direct consequence of combining Lemmas 3.7 and 3.8.

Algorithm 3 Online Stochastic Mirror Descent with Barycentric Regularization

Input: Learning rate $\eta > 0$, tolerance $\epsilon > 0$, $\{q_1, \ldots, q_d\}$ - a β -BS(\mathcal{K}) for some $\beta > 0$, exploration probability $\gamma \in (0, 1)$, time horizon $T \in \mathbb{N}_+$

- 1: Instantiate Algorithm 1 with parameters η , ϵ , $\varphi(x) = \frac{1}{2}x^{\top}Q^{-1}x$, $W' = (M^{\top})^{-1}\mathbb{R}^d_+$, and T
- 2: for t = 1 to T do
- Receive \tilde{x}_t (the point to play in round t) from Algorithm 1 3:
- $b_t \leftarrow \begin{cases} \text{EXPLORE,} & \text{with probability } \gamma \\ \text{EXPLOIT,} & \text{with probability } 1 \gamma \end{cases}$ 4:
- if $b_t = EXPLORE$ then 5:
- 6: Sample $i_t \in [d]$ uniformly at random, and play q_{i_t}
- Receive loss $l_t = q_{i_t}^\top f_t$ $\tilde{f}_t \leftarrow \frac{d}{\gamma} l_t Q^{-1} q_{i_t}$ 7:
- 8:
- 9:
- Play \tilde{x}_t and receive loss $l_t = \tilde{x}_t^\top f_t$ 10:
- $\tilde{f}_t \leftarrow 0$ 11:
- end if 12:
- Feed \tilde{f}_t to Algorithm 1 as the loss vector for round t (Note that when $\tilde{f}_t = 0$, in the next round 13: Algorithm 1 can simply play according to the distribution computed in this round without any oracle calls.)
- 14: end for

α -Regret Minimization in the Bandit Setting 4

In this section we consider the α -regret minimization problem in the bandit setting, where $W = \mathbb{R}^d_+$, $\mathcal{K} \subseteq \mathbb{R}^d_+ \cap B(0, R)$ and $\mathcal{K}^* = \alpha \mathcal{K}$. Suppose the loss vectors $\{f_t\}$ come from \mathbb{R}^d and $\|f_t\| \leq L$. Similar to (Kakade et al., 2009), we assume we know a β -barycentric spanner for \mathcal{K} . This concept was first introduced by Awerbuch and Kleinberg (2004).

Definition 4.1 (Barycentric spanner). A set of d linearly independent vectors $\{q_1, \ldots, q_d\} \subset \mathbb{R}^d$ is a β -barycentric spanner for a set $\mathcal{K} \subset \mathbb{R}^d$, denoted by β -BS(\mathcal{K}), if $\{q_1, \ldots, q_d\} \subseteq \mathcal{K}$ and for all $x \in \mathcal{K}$, there exist $\beta_1, \ldots, \beta_d \in [-\beta, \beta]$ such that $x = \sum_{i=1}^d \beta_i q_i$.

Given $\{q_1, \ldots, q_d\}$ which is a β -BS(\mathcal{K}), define $Q := \sum_{i=1}^d q_i q_i^{\top}$ and $M := (q_1, \ldots, q_d) \in \mathbb{R}^{d \times d}$.

The need for a new regularization. The bandit algorithm of Garber (2017) additionally requires a certain boundedness property of barycentric spanners, namely:

$$\max_{i \in [d]} q_i^\top Q^{-2} q_i \le \chi.$$

However, for certain bounded sets this quantity may be unbounded, such as the two-dimensional axisaligned rectangle with one axis being of size unity, and the other arbitrarily small. This unboundedness creates problems with the unbiased estimator of loss vector, whose variance can be as large as certain geometric properties of the decision set. To circumvent this issue, we design a new regularizer called barycentric regularizer, which gives rise to an unbiased estimator coupled with an online mirror descent variant that automatically ensures constant variance.

Similar to (Kakade et al., 2009; Garber, 2017), our bandit algorithm also simulates the full information algorithm with estimated loss vectors. Namely, our algorithm implements Algorithm 1 with a specific *barycentric regularizer* $\varphi(x) = \frac{1}{2}x^{\top}Q^{-1}x$. The algorithm is detailed in Algorithm 3, and its regret guarantee is given in Theorem 4.2. We prove Theorem 4.2 in Appendix C.

Theorem 4.2. Denote by z_t the point played by Algorithm 3 in round t.

Suppose we set $\eta = \frac{\alpha \beta^{4/3}}{LRT^{2/3}}$, $\epsilon = \frac{\alpha R}{T}$ and $\gamma = \frac{\beta^{2/3} d}{T^{1/3}}$ in Algorithm 3 (assuming $T > \beta^2 d^3$ so $\gamma < 1$). Then we have

$$\forall x^* \in \mathcal{K} : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(z_t) - \alpha f_t(x^*))\right] \le \alpha LR\left(3d(\beta T)^{2/3} + 1\right)$$

and the expected total number of oracle calls to $\mathcal{O}_{\mathcal{K},\alpha\mathcal{K}}$ in T rounds is at most $O\left(d^2(\beta T)^{2/3}\log T\right)$.

5 Conclusion and Open Problems

We have described two different algorithmic approaches to reducing regret minimization to offline approximation algorithms and maintaining optimal regret and poly-logarithmic oracle complexity per iteration, resolving previously stated open questions.

An intriguing open problem remaining is to find an efficient algorithm in the bandit setting that guarantees both $\tilde{O}(\sqrt{T})$ regret and $\operatorname{poly}(\log T)$ oracle complexity per iteration (at least on average).

References

- Awerbuch, B. and Kleinberg, R. D. (2004). Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 45–53. ACM.
- Balcan, M.-F. and Blum, A. (2006). Approximation algorithms and online mechanisms for item pricing. In Proceedings of the 7th ACM Conference on Electronic Commerce, pages 29–35. ACM.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- Cover, T. M. (1991). Universal portfolios. *Mathematical Finance*, 1(1):1–29.
- Dudík, M., Haghtalab, N., Luo, H., Schapire, R. E., Syrgkanis, V., and Vaughan, J. W. (2016). Oracle-efficient online learning and auction design. *arXiv preprint arXiv:1611.01688*.
- Fujita, T., Hatano, K., and Takimoto, E. (2013). Combinatorial online prediction via metarounding. In *International Conference on Algorithmic Learning Theory*, pages 68–82. Springer.
- Garber, D. (2017). Efficient online linear optimization with approximation algorithms. In Advances in Neural Information Processing Systems, pages 627–635.
- Goemans, M. X. and Williamson, D. P. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. J. ACM, 42(6):1115–1145.
- Hazan, E. (2016). Introduction to online convex optimization. Foundations and Trends® in Optimization, 2(3-4):157–325.
- Hazan, E. and Koren, T. (2016). The computational power of optimization in online learning. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 128–141. ACM.
- Kakade, S. M., Kalai, A. T., and Ligett, K. (2009). Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106.
- Kalai, A. and Vempala, S. (2005). Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307.
- Lovász, L. and Vempala, S. (2007). The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358.
- Prékopa, A. (1973). On logarithmic concave measures and functions. *Acta Scientiarum Mathematicarum*, 34:335–343.
- Shalev-Shwartz, S. (2012). Online learning and online convex optimization. *Foundations and Trends*(R) *in Machine Learning*, 4(2):107–194.
- Vovk, V. G. (1990). Aggregating strategies. In Proceedings of the Third Annual Workshop on Computational Learning Theory, COLT '90, pages 371–386.

Appendix

A Omitted Proofs in Section 3

Theorem A.1 (Restatement of Theorem 3.1). Suppose (φ, W) satisfies the PNIP property (Definition 2.4). Then for any $\epsilon, \eta > 0$, Algorithm 1 satisfies the following regret guarantee:

$$\forall x^* \in \mathcal{K}^* : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le \frac{1}{\eta} \left(\varphi(x^*) - \varphi(y_1) + \sum_{t=1}^T D_\varphi(x_t, y_{t+1})\right) + \epsilon LT.$$

In particular, if φ is μ -strongly convex and $A \ge \max_{x^* \in \mathcal{K}^*} (\varphi(x^*) - \varphi(y_1))$, setting $\epsilon = \frac{R}{T}$ and $\eta = \frac{1}{L} \sqrt{\frac{2\mu A}{T}}$, we have

$$\forall x^* \in \mathcal{K}^* : \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le L\sqrt{\frac{2AT}{\mu}} + LR.$$

Moreover, Algorithm 1 makes at most $\left\lceil 5d \log \left(\left(6\sqrt{T} + \frac{4}{R}\sqrt{\frac{A}{\mu}} + 4 \right) T \right) \right\rceil$ calls of $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ per round.

Proof. First, for any fixed round $t \in [T]$, let (x_t, V, p) be the output of $\mathcal{PAD}(y_t, \epsilon, W, \varphi)$ in this round. We know by the second property of the PAD oracle that there exists $c \in W^\circ$ such that $\|\sum_i p_i v_i + c - x_t\| \le \epsilon$. Since \tilde{x}_t is equal to v_i with probability p_i , letting $\overline{x}_t := \mathbb{E}[\tilde{x}_t] = \sum_i p_i v_i$, we have

$$f_t(\overline{x}_t) - f_t(x_t) = \mathbb{E}\left[f_t(\widetilde{x}_t) - f_t(x_t)\right] = f_t\left(\sum_i p_i v_i - x_t\right) \le f_t\left(\sum_i p_i v_i - x_t + c\right) \le \epsilon L.$$
(1)

We make use of the following properties of Bregman divergence, which can be verified easily (see e.g. Section 11.2 in (Cesa-Bianchi and Lugosi, 2006)):

$$\forall x, y, z: \quad (x-y)^{\top} (\nabla \varphi(z) - \nabla \varphi(y)) = D_{\varphi}(x, y) - D_{\varphi}(x, z) + D_{\varphi}(y, z).$$
(2)

Consider any $x^* \in \mathcal{K}^*$. We have

$$\sum_{t=1}^{T} (f_t(x_t) - f_t(x^*))$$

$$= \sum_{t=1}^{T} \frac{1}{\eta} (\nabla \varphi(x_t) - \nabla \varphi(y_{t+1}))^{\top} (x_t - x^*) \qquad \text{(by algorithm definition)}$$

$$= \frac{1}{\eta} \sum_{t=1}^{T} (D_{\varphi}(x^*, x_t) - D_{\varphi}(x^*, y_{t+1}) + D_{\varphi}(x_t, y_{t+1})) \qquad \text{(by (2))}$$

$$\leq \frac{1}{\eta} \sum_{t=1}^{T} (D_{\varphi}(x^*, y_t) - D_{\varphi}(x^*, y_{t+1}) + D_{\varphi}(x_t, y_{t+1})) \qquad \text{(by property of the PAD oracle)}$$

$$= \frac{1}{\eta} \left(D_{\varphi}(x^*, y_t) - D_{\varphi}(x^*, y_{t+1}) + D_{\varphi}(x_t, y_{t+1}) \right) \qquad \text{(by talesconing)}$$

$$= \frac{1}{\eta} \left(D_{\varphi}(x^*, y_1) - D_{\varphi}(x^*, y_{T+1}) + \sum_{t=1}^T D_{\varphi}(x_t, y_{t+1})) \right).$$
 (by telescoping) (3)

Combining (1) and (3), we can bound the expected improper regret of Algorithm 1 as

$$\forall x^* \in \mathcal{K}^* : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] = \sum_{t=1}^T (f_t(\overline{x}_t) - f_t(x^*)) \\ \leq \frac{1}{\eta} \left(D_{\varphi}(x^*, y_1) - D_{\varphi}(x^*, y_{T+1}) + \sum_{t=1}^T D_{\varphi}(x_t, y_{t+1})) \right) + \epsilon LT.$$
(4)

By the optimality condition $\nabla \varphi(y_1)^{\top}(x^* - y_1) \ge 0$, we have

$$D_{\varphi}(x^*, y_1) \le \varphi(x^*) - \varphi(y_1).$$
(5)

Plugging (5) into (4) and noting $D_{\varphi}(x^*, y_{T+1}) \ge 0$, we finish the proof of the first regret bound.

When φ is μ -strongly convex, for any fixed y, $f(x) = \varphi(x) - \nabla \varphi(y)^{\top} x$ is also strongly convex, achieving its unique minimum at y. From strong convexity we have

$$f(z) \ge f(x) + \nabla f(x)^{\top} (z - x) + \frac{\mu}{2} ||x - z||^2, \ \forall x, z.$$

Minimizing both sides of the above inequality over z, we get $f(y) \ge f(x) - \frac{1}{2\mu} \|\nabla f(x)\|^2$. Thus,

$$D_{\varphi}(x,y) = f(x) - f(y) \le \frac{1}{2\mu} \|\nabla f(x)\|^2 = \frac{1}{2\mu} \|\nabla \varphi(x) - \nabla \varphi(y)\|^2.$$

Then by the definition in Algorithm 1 we have

$$\forall t \in [T]: \quad D_{\varphi}(x_t, y_{t+1}) \le \frac{1}{2\mu} \|\nabla\varphi(x_t) - \nabla\varphi(y_{t+1})\|^2 = \frac{1}{2\mu} \|\eta f_t\|^2 \le \frac{\eta^2 L^2}{2\mu}. \tag{6}$$

From the above inequality and the choices of parameters $\epsilon = \frac{R}{T}$ and $\eta = \frac{1}{L}\sqrt{\frac{2\mu A}{T}}$, we have

$$\mathbb{E}\left[\sum_{t=1}^{T} (f_t(\tilde{x}_t) - f_t(x^*))\right] \le \frac{A}{\eta} + \frac{\eta L^2 T}{2\mu} + LR \le L\sqrt{\frac{2AT}{\mu}} + LR.$$

Next we bound the number of oracle calls made by Algorithm 1.

According to Theorem 3.4, round t of Algorithm 1 calls $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ for at most

$$\left[5d \log \frac{4R + 2\sqrt{\frac{2}{\mu}} \min_{x^* \in \mathcal{K}^*} D_{\varphi}(x^*, y_t)}{\epsilon} \right]$$

times. Hence it suffices to obtain an upper bound on $\min_{x^* \in \mathcal{K}^*} D_{\varphi}(x^*, y_t)$. According to (4) (substituting T with t), we have:

$$\forall t \in [T], \ \forall x^* \in \mathcal{K}^*: \quad D_{\varphi}(x^*, y_{t+1}) \le D_{\varphi}(x^*, y_1) + \sum_{j=1}^t D_{\varphi}(x_j, y_{j+1}) - \eta \sum_{j=1}^t (f_j(\overline{x}_j) - f_j(x^*)) + \epsilon \eta L t$$

Plug (5) and (6) into the above inequality, we have

$$\begin{aligned} \forall t \in [T], \ \forall x^* \in \mathcal{K}^*: \quad D_{\varphi}(x^*, y_{t+1}) \leq D_{\varphi}(x^*, y_1) + \sum_{j=1}^t D_{\varphi}(x_j, y_{j+1}) - \eta \sum_{j=1}^t (f_j(\overline{x}_j) - f_j(x^*)) + \epsilon \eta L t \\ \leq \varphi(x^*) - \varphi(y_1) + \frac{\eta^2 L^2}{2\mu} t - \eta \sum_{j=1}^t (f_j(\overline{x}_j) - f_j(x^*)) + \epsilon \eta L t \\ \leq A + \frac{\eta^2 L^2}{2\mu} t + \eta \sum_{j=1}^t \|f_j\| \cdot \|\overline{x}_j - x^*\| + \epsilon \eta L t \\ \leq A + \frac{\eta^2 L^2}{2\mu} T + 2\eta L R T + \epsilon \eta L T. \end{aligned}$$

For t = 1 we also have $D_{\varphi}(x^*, y_1) \leq A$. Therefore Algorithm 1 calls $\mathcal{O}_{\mathcal{K}, \mathcal{K}^*}$ for at most

$$\left[5d \log \frac{4R + 2\sqrt{\frac{2}{\mu} \left(A + \left(\frac{\eta^2 L^2}{2\mu} + 2\eta LR + \epsilon \eta L\right)T\right)}}{\epsilon} \right]$$

times per round.

When $\epsilon = \frac{R}{T}$ and $\eta = \frac{1}{L}\sqrt{\frac{2\mu A}{T}}$, we have

$$\frac{2}{\mu} \left(A + \left(\frac{\eta^2 L^2}{2\mu} + 2\eta L R + \epsilon \eta L \right) T \right) \le \frac{4A}{\mu} + 6R \sqrt{\frac{2AT}{\mu}} \le \left(2\sqrt{\frac{A}{\mu}} + 3R\sqrt{T} \right)^2,$$

so the number of oracle calls per iteration is at most $\left[5d \log\left(\left(6\sqrt{T} + \frac{4}{R}\sqrt{\frac{A}{\mu}} + 4\right)T\right)\right]$.

Next we prove Lemma 3.8. First we need the following basic properties of projections. **Lemma A.2** (Pythagorean theorem). For any closed convex set $S \subseteq \mathbb{R}^d$, $x \in \mathbb{R}^d$ and $y \in S$, we have $(\Pi_S(x) - x)^\top (\Pi_S(x) - y) \leq 0$, or equivalently, $||x - y||^2 \geq ||\Pi_S(x) - x||^2 + ||\Pi_S(x) - y||^2$.

Lemma A.3. For any closed convex cone $W \subseteq \mathbb{R}^d$ and any $x \in \mathbb{R}^d$, we have $\Pi_W(x) - x \in W^\circ$.

Proof. Since W is a convex cone and $\Pi_W(x) \in W$, we have $w + \Pi_W(x) \in W$ ($\forall w \in W$). By Pythagorean theorem, we have $(\Pi_W(x) - x)^\top (y - \Pi_W(x)) \ge 0$ ($\forall y \in W$). Letting $y = w + \Pi_W(x) \in W$ ($\forall w \in W$), we get $(\Pi_W(x) - x)^\top w \ge 0$, which means $\Pi_W(x) - x \in W^\circ$. \Box

Lemma A.4 (Restatement of Lemma 3.8). *Given* $v_1, \ldots, v_k \in \mathbb{R}^d$, $\epsilon \ge 0$ and a convex cone $W \in \mathbb{R}^d$, for any $x \in \mathbb{R}^d$, the following two statements are equivalent:

(A) There exists $p = (p_1, \dots, p_k)^\top \in \Delta^{k-1}$ and $c \in W^\circ$ such that $\|\sum_{i=1}^k p_i v_i + c - x\| \le \epsilon$. (B) For all $w \in W$, $\|w\| = 1$, there exists $i \in [k]$ such that $w^\top (v_i - x) \le \epsilon$.

Proof. Suppose (A) holds. Then for any $w \in W$, ||w|| = 1, we have

$$\min_{i \in [k]} w^{\top}(v_i - x) \le w^{\top} \left(\sum_{i=1}^k p_i v_i - x \right) \le w^{\top} \left(\sum_{i=1}^k p_i v_i + c - x \right) \le \|w\| \cdot \left\| \sum_{i=1}^k p_i v_i + c - x \right\| \le \epsilon.$$

So we have $(A) \Longrightarrow (B)$.

Now suppose (B) holds. Let $(p^*, c^*) = \underset{p \in \Delta^{k-1}, c \in W^{\circ}}{\operatorname{arg\,min}} \|\sum_{i=1}^k p_i v_i + c - x\|$. Since $0 \in W$, by the Pythagorean theorem (Lemma A.2), we have

$$\left\|\sum_{i=1}^{k} p_{i}^{*} v_{i} - x + c^{*}\right\| \geq \left\|\Pi_{W}\left(\sum_{i=1}^{k} p_{i}^{*} v_{i} - x + c^{*}\right)\right\|,$$

where the equality holds only when $\sum_{i=1}^{k} p_i^* v_i - x + c^* \in W$. Now we claim $\sum_{i=1}^{k} p_i^* v_i - x + c^* \in W$. Otherwise, letting $c' = c^* + \prod_W \left(\sum_{i=1}^{k} p_i^* v_i - x + c^* \right) - \left(\sum_{i=1}^{k} p_i^* v_i - x + c^* \right)$, by Lemma A.3 we have $c' \in W^\circ$, and furthermore

$$\left|\sum_{i=1}^{k} p_{i}^{*} v_{i} - x + c'\right| = \left\|\Pi_{W}\left(\sum_{i=1}^{k} p_{i}^{*} v_{i} - x + c^{*}\right)\right\| < \left\|\sum_{i=1}^{k} p_{i}^{*} v_{i} - x + c^{*}\right\|,$$

which contradicts the optimality of (p^*, c^*) .

Thus we have $\sum_{i=1}^{k} p_i^* v_i + c^* - x \in W$. Let $w = \sum_{i=1}^{k} p_i^* v_i + c^* - x$ and $G = \operatorname{CH}(\{v_1, \ldots, v_k\}) + W^\circ + \{-x\}$. Then we have $w = \Pi_G(0)$ by the definition of (p^*, c^*) . Since G is convex and $v_i - x \in G$ for all $i \in [k]$, by the Pythagorean theorem (Lemma A.2) we have $w^\top (v_i - x - w) \ge 0$ for all $i \in [k]$, which implies $||w||^2 \le \min_{i \in [k]} w^\top (v_i - x) \le \epsilon ||w||$, i.e., $||w|| \le \epsilon$. Hence we have (B) \Longrightarrow (A).

B Efficient Online Improper Linear Optimization via Continuous Multiplicative Weight Update (CMWU)

In this section, we design our second online improper linear optimization algorithm (in the full information setting) based on the continuous multiplicative weight update (CMWU) method. Same as in Section 3, we suppose $\mathcal{K}, \mathcal{K}^* \subseteq B(0, R)$, and the loss vectors $\{f_t\}$ come from a convex cone $W \subseteq \mathbb{R}^d$ and $||f_t|| \leq L$ (R, L > 0). Compared with Algorithm 1, the CMWU-based algorithm allows loss vectors to come from a general convex cone W and does not require the PNIP condition (Definition 2.4).

B.1 Separation-or-Decomposition Oracle

We first construct a *separation-or-decomposition (SOD)* oracle (Algorithm 4) using $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$, which we will use to design the online improper linear optimization algorithm later in this section. Given a point $x \in B(0, R)$, the SOD oracle either outputs a *separating hyperplane* between x and \mathcal{K}^* , or outputs a *distribution of points in* \mathcal{K} which approximately dominates x in every direction in W. The guarantee of the SOD oracle is summarized in Theorem B.1.

Algorithm 4 Separation-or-Decomposition Oracle, $SOD(x, \epsilon, W)$

Input: Point $x \in B(0, R)$, tolerance $\epsilon > 0$, convex cone $W \subseteq \mathbb{R}^d$ **Output:** Decomposition $V = (v_1, \ldots, v_k) \in \mathbb{R}^{d \times k}$, $p = (p_1, \ldots, p_k)^\top \in \Delta^{k-1}$, such that $v_i \in \mathcal{K}$ $(\forall i \in [k])$ and $\|\sum_{i=1}^k p_i v_i - x + c\| \leq 3\epsilon$ for some $c \in W^\circ$. **Or:** Separating hyperplane $(w, b) \in \mathbb{R}^d \times \mathbb{R}$, such that $\|w\| = 1$ and $w^\top x \leq b - \epsilon \leq C$ $\min_{x^* \in K^*} w^\top x^* - \epsilon.$ 1: $k \leftarrow \left\lceil 5d \log \frac{4R}{\epsilon} \right\rceil$ 2: $W_1 \leftarrow W \cap B(0,1)$ 3: for i = 1 to k do 4: $w_i \leftarrow \frac{\int_{W_i} w dw}{\operatorname{Vol}(W_i)}$ $v_i \leftarrow \mathcal{O}_{\mathcal{K},\mathcal{K}^*}(w_i)$ if $w_i^\top x \leq w_i^\top v_i - \epsilon$ then 5: 6: **return** Separating hyperplane $\left(\frac{w_i}{\|w_i\|}, \frac{w_i^{\top} v_i}{\|w_i\|}\right)$ 7: 8: else $W_{i+1} \leftarrow W_i \cap \{ w \in \mathbb{R}^d : w^\top (v_i - x) \ge \epsilon \}$ 9: 10: end if 11: end for 12: Solve $\min_{p \in \Delta^{k-1}, c \in W^{\circ}} \|\sum_{i=1}^{k} p_i v_i + c - x\|$ to get p13: **return** $V = (v_1, \dots, v_k), p$

Theorem B.1. For any $x \in B(0, R)$ and $\epsilon \in (0, 2R]$, the separation-or-decomposition oracle in Algorithm 4, $SOD(x, \epsilon, W)$, returns one of the two followings, using at most $k = \lfloor 5d \log \frac{4R}{\epsilon} \rfloor$ calls of $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$:

- 1. a decomposition $V = (v_1, \ldots, v_k) \in \mathbb{R}^{d \times k}$, $p = (p_1, \ldots, p_k)^\top \in \Delta^{k-1}$, such that $v_i \in \mathcal{K}$ $(\forall i \in [k])$ and $\|\sum_{i=1}^k p_i v_i - x + c\| \leq 3\epsilon$ for some $c \in W^\circ$.
- 2. a separating hyperplane $(w,b) \in \mathbb{R}^d \times \mathbb{R}$, where ||w|| = 1 and $w^{\top}x \leq b \epsilon \leq \min_{x^* \in \mathcal{K}^*} w^{\top}x^* \epsilon$.

Before proving Theorem B.1, we first prove the following lemma using a similar argument in Lemma 3.7.

Lemma B.2. For $x \in B(0, R)$ and $\epsilon \in (0, 2R]$, if $SOD(x, \epsilon, W)$ returns a decomposition (V, p), then for all unit vector $w \in W$, we have $\min_{i \in [k]} w^{\top}(v_i - x) \leq 3\epsilon$.

Proof. Suppose that there exists a unit vector $h \in W$ such that $\min_{i \in [k]} h^{\top}(v_i - x) > 3\epsilon$. Note that $||v_i - x|| \le ||v_i|| + ||x|| \le 2R$. Denoting $r = \frac{\epsilon}{4R}$, we have

$$\forall h' \in \frac{h}{2} + (W \cap B(0, r)): \quad \min_{i \in [k]} h'^{\top}(v_i - x) > \epsilon.$$

Since $r \leq \frac{1}{2}$ for $\epsilon \leq 2R$, we have $\frac{h}{2} + (W \cap B(0, r)) \subseteq \frac{h}{2} + (W \cap B(0, 1/2)) \subseteq W \cap B(0, 1) = W_1$. Because the algorithm returns a decomposition, we have that after the last iteration,

 $\forall w \in W_1 \setminus W_{k+1}: \quad \exists i \in [k], \text{ s.t. } w^\top (v_i - x) \le \epsilon.$

Therefore, we must have $\frac{h}{2} + (W \cap B(0, r)) \subseteq W_{k+1}$. We also have $\operatorname{Vol}(W_{i+1}) \leq (1 - 1/2e) \operatorname{Vol}(W_i)$ from Lemma D.2 since W_{i+1} is the intersection of W_i with a half-space that does not contain W_i 's centroid. Then we have

$$Vol(W_1) = Vol(W \cap B(0, 1)) = r^{-d} Vol(W \cap B(0, r)) \le r^{-d} Vol(W_{k+1}) \le r^{-d} (1 - 1/2e)^k Vol(W_1) < Vol(W_1),$$

where the last step is due to $k \ge 5d \log \frac{1}{r} = 5d \log \frac{4R}{\epsilon}$. Therefore we have a contradiction.

Proof of Theorem B.1. If $SOD(x, \epsilon, W)$ returns a decomposition (V, p), by Lemmas B.2 and 3.8, we know that there exists $c \in W^{\circ}$ such that $\|\sum_{i=1}^{k} p_i v_i + c - x\| \leq 3\epsilon$.

If $\mathcal{SOD}(x, \epsilon, W)$ returns a separating hyperplane (w, b) at iteration $i \in [k]$, we know $w_i^{\top} x \leq w_i^{\top} v_i - \epsilon$. Since $w = \frac{w_i}{\|w_i\|}$, $b = w^{\top} v_i$ and $\|w_i\| \leq 1$, we have $w^{\top} x \leq w^{\top} v_i - \frac{\epsilon}{\|w_i\|} \leq b - \epsilon$. By the guarantee of $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$, we have $b - \epsilon = w^{\top} v_i - \epsilon \leq \min_{x^* \in \mathcal{K}^*} w^{\top} x^* - \epsilon$.

The number of calls of $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ is clearly upper bounded by $k = \lfloor 5d \log \frac{4R}{\epsilon} \rfloor$ since there are at most k iterations and each iteration only calls $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ once.

B.2 CMWU with Refining Domains

Now we look at a general online learning setting where the feasible domain is shrinking over time while being a superset of the target domain. Namely, suppose \mathcal{K}^* is the target domain and \mathcal{K}_t is the feasible domain in the *t*-th round. We assume $B(0, R) \supseteq \mathcal{K}_0 \supseteq \mathcal{K}_1 \supseteq \mathcal{K}_2 \cdots \supseteq \mathcal{K}_T \supseteq (1-\gamma)\mathcal{K}^* + \gamma\mathcal{K}_0$ for some $\gamma \in (0, 1]$. In round *t*, the player only knows $\mathcal{K}_1, \ldots, \mathcal{K}_t$ and does not know \mathcal{K}_j for all j > t. We can still run CMWU in this setting, using the knowledge of \mathcal{K}_t at iteration *t* - the algorithm is given in Algorithm 5. Theorem B.3 bounds the regret of Algorithm 5 in this setting.

Algorithm 5 Continuous Multiplicative Weight Update (CMWU) with Refining Domains

Input: Learning rate $\eta > 0$, time horizon $T \in \mathbb{N}_+$ 1: for t = 1 to T do

2: Receive current domain \mathcal{K}_t

3: Play
$$x_t = \frac{\int_{\mathcal{K}_t} e^{-\eta \sum_{i=1}^{t} \int_{i}^{t} (x) x d}}{1 + 1 + 1 + 1 + 1 + 1 + 1}$$

$$\int_{\mathcal{K}_{\star}} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} \mathrm{d}x$$

4: Receive loss vector
$$f_t$$

5: end for

Theorem B.3. Suppose $B(0, R) \supseteq \mathcal{K}_0 \supseteq \mathcal{K}_1 \supseteq \mathcal{K}_2 \supseteq \cdots \supseteq \mathcal{K}_T \supseteq (1 - \gamma)\mathcal{K}^* + \gamma \mathcal{K}_0$ for $\gamma \in (0, 1]$. Then for any $0 < \eta \leq \frac{1}{LR}$, Algorithm 5 has the following regret guarantee:

$$\forall x^* \in \mathcal{K}^* : \quad \sum_{t=1}^{T} (f_t(x_t) - f_t(x^*)) \le \frac{d \log \frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + \gamma L R T - \frac{\sum_{t=1}^{T} \delta_t}{\eta},$$

where

$$\delta_t := \log \frac{\int_{\mathcal{K}_{t-1}} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} \mathrm{d}x}{\int_{\mathcal{K}_t} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} \mathrm{d}x}.$$

In particular, setting $\gamma = 1/T$ and $\eta = \frac{1}{LR} \min\left\{1, \sqrt{\frac{d \log T}{T}}\right\}$, we have

$$\forall x^* \in \mathcal{K}^*: \quad \sum_{t=1}^T (f_t(x_t) - f_t(x^*)) \le LR\left(1 + 2\max\left\{\sqrt{dT\log T}, d\log T\right\}\right).$$

Proof. We fix any $x^* \in \mathcal{K}^*$ and denote $\overline{\mathcal{K}} := (1 - \gamma)x^* + \gamma \mathcal{K}_0$. Since $(1 - \gamma)\mathcal{K}^* + \gamma \mathcal{K}_0 \subseteq \mathcal{K}_T$, we have $\overline{\mathcal{K}} \subseteq \mathcal{K}_T$. We define

$$z_t(x) := e^{-\eta \sum_{i=1}^{t-1} f_i(x)}, \quad Z_t := \int_{\mathcal{K}_t} z_t(x) \mathrm{d}x, \quad Z'_t := \int_{\mathcal{K}_{t-1}} z_t(x) \mathrm{d}x.$$

A straightforward calculation gives us:

$$\log \frac{Z'_{T+1}}{Z'_1} = \log \left(\frac{\int_{\mathcal{K}_T} e^{-\eta \sum_{t=1}^T f_t(x)} dx}{\int_{\mathcal{K}_0} 1 dx} \right)$$

$$\geq \log \left(\frac{\int_{\bar{\mathcal{K}}} e^{-\eta \sum_{t=1}^T f_t(x)} dx}{\int_{\mathcal{K}_0} 1 dx} \right)$$

$$= \log \left(\frac{\int_{\mathcal{K}_0} e^{-\eta \sum_{t=1}^T f_t((1-\gamma)x^* + \gamma x)\gamma^d dx}}{\int_{\mathcal{K}_0} 1 dx} \right)$$

$$= \log \left(\frac{\int_{\mathcal{K}_0} e^{-\eta \sum_{t=1}^T ((1-\gamma)f_t(x^*) + \gamma f_t(x))\gamma^d dx}}{\int_{\mathcal{K}_0} 1 dx} \right)$$

$$\geq \log \left(\frac{\int_{\mathcal{K}_0} e^{-\eta \sum_{t=1}^T (f_t(x^*) + \gamma LR)\gamma^d dx}}{\int_{\mathcal{K}_0} 1 dx} \right)$$

$$= d \log \gamma - \eta \sum_{t=1}^T f_t(x^*) - \eta \gamma LRT.$$

On the other hand, we have

$$\begin{split} \log \frac{Z'_{t+1}}{Z_t} &= \log \left(\int_{\mathcal{K}_t} \frac{z_t(x)}{Z_t} e^{-\eta f_t(x)} \mathrm{d}x \right) \\ &\leq \log \left(\int_{\mathcal{K}_t} \frac{z_t(x)}{Z_t} \left(1 - \eta f_t(x) + (\eta f_t(x))^2 \right) \mathrm{d}x \right) \\ &\leq \left(\int_{\mathcal{K}_t} \frac{z_t(x)}{Z_t} \left(1 - \eta f_t(x) + (\eta f_t(x))^2 \right) \mathrm{d}x \right) - 1 \\ &\leq \int_{\mathcal{K}_t} \frac{z_t(x)}{Z_t} \left(-\eta f_t(x) + \eta^2 L^2 R^2 \right) \mathrm{d}x \\ &= -\eta f_t \left(\int_{\mathcal{K}_t} \frac{z_t(x)}{Z_t} \mathrm{xd}x \right) + \eta^2 L^2 R^2 \\ &= -\eta f_t(x_t) + \eta^2 L^2 R^2, \end{split}$$

where the first inequality is due to $e^a \leq 1 + a + a^2$ ($\forall a \leq 1$) and $|\eta f_t(x)| \leq \eta LR \leq 1$, the second inequality is due to $\log a \leq a - 1$ ($\forall a > 0$), and the third inequality is due to $|\eta f_t(x)| \leq \eta LR$.

Note that $\delta_t = \log \frac{Z'_t}{Z_t}$. Combining the two bounds above, we get:

$$\sum_{t=1}^{T} \left(-\eta f_t(x_t) + \eta^2 L^2 R^2 \right) \ge \sum_{t=1}^{T} \log \frac{Z'_{t+1}}{Z_t} = \log \frac{Z'_{T+1}}{Z'_1} + \sum_{t=1}^{T} \log \frac{Z'_t}{Z_t}$$
$$\ge d \log \gamma - \eta \sum_{t=1}^{T} f_t(x^*) - \eta \gamma L R T + \sum_{t=1}^{T} \delta_t.$$

In other words,

$$\sum_{t=1}^{T} (f_t(x_t) - f_t(x^*)) \le \frac{d \log \frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + \gamma L R T - \frac{\sum_{t=1}^{T} \delta_t}{\eta}$$

The second regret bound in the theorem follows directly by plugging in the values of γ and η stated in the theorem. Note that $\delta_t \ge 0$ for all $t \in [T]$ since $\mathcal{K}_t \subseteq \mathcal{K}_{t-1}$.

B.3 Online Improper Linear Optimization via CMWU

Now we are ready to present our CMWU-based online improper learning algorithm. At a high level, this algorithm is a specialized implementation of Algorithm 5 for the online improper linear optimization problem. The algorithm starts with an initial convex domain \mathcal{K}_0 which is a superset of \mathcal{K}^* , and maintains a convex domain \mathcal{K}_t at iteration t. In iteration t, the algorithm first computes the mean x_t of a log-linear distribution over \mathcal{K}_t using random walk, as in Line 3 of Algorithm 5. Then the algorithm calls the SOD oracle on x_t . If the SOD oracle returns a distribution of points in \mathcal{K} , then we can play according to that distribution, since the SOD oracle ensures that the expected loss of this distribution is not much larger than that of x_t . If the SOD oracle returns a separating hyperplane between x_t and \mathcal{K}^* , then the algorithm replaces \mathcal{K}_t with the intersection of the original \mathcal{K}_t and the half-space given by this hyperplane that contains \mathcal{K}^* , and repeats the same process for the new \mathcal{K}_t until a decomposition is not in the new \mathcal{K}_t , which according to Lemma D.2 implies that a constant probability mass is removed. This allows us to bound the total number of oracle calls. We detail our algorithm in Algorithm 6 and its regret bound in Theorem B.4.

Algorithm 6 CMWU for Online Improper Linear Optimization

Input: Learning rate $\eta > 0$, tolerance $\gamma > 0$, initial convex domain \mathcal{K}_0 , convex cone W, time horizon $T \in \mathbb{N}_+$

1: $\mathcal{K}_1 \leftarrow \mathcal{K}_0$ 2: **for** t = 1 **to** T **do** 3: $x_t \leftarrow \frac{\int_{\mathcal{K}_t} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} x dx}{\int_{\mathcal{K}_t} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} dx}$ while $\mathcal{SOD}(x_t, 2\gamma R, W)$ returns a separating hyperplane $(w, b) \in \mathbb{R}^d \times \mathbb{R}$ do 4: $\mathcal{K}_t \leftarrow \mathcal{K}_t \cap \{x \in \mathbb{R}^d : w^\top x \ge b - 2\gamma R\}$ $x_t \leftarrow \frac{\int_{\mathcal{K}_t} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} x dx}{\int_{\mathcal{K}_t} e^{-\eta \sum_{i=1}^{t-1} f_i(x)} dx}$ 5: 6: end while 7: Let $(V, p) \in \mathbb{R}^{d \times k} \times \Delta^{k-1}$ be the output of $\mathcal{SOD}(x_t, 2\gamma R, W)$ Play $\tilde{x}_t = v_i$ with probability p_i (i = 1, ..., k), where $V = (v_1, ..., v_k)$ 8: 9: 10: $\mathcal{K}_{t+1} \leftarrow \mathcal{K}_t$ Receive loss vector f_t 11: 12: end for

Theorem B.4. Suppose that the initial convex domain \mathcal{K}_0 satisfies $\mathcal{K}^* \subseteq \mathcal{K}_0 \subseteq B(0, R)$. Then for any $\gamma \in (0, 1]$ and $\eta \in (0, \frac{1}{LR}]$, Algorithm 6 satisfies the following regret guarantee:

$$\forall x^* \in \mathcal{K}^* : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le \frac{d\log\frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + 7\gamma LRT - \frac{s}{5\eta},$$

where $s = \sum_{t=1}^{T} s_t$, and s_t is the number of separating hyperplanes returned by the SOD oracle during round t.

In particular, if we set $\gamma = \frac{1}{T}, \eta = \frac{1}{LR} \min\left\{1, \sqrt{\frac{d \log T}{T}}\right\}$, then we have $\forall x^* \in \mathcal{K}^*: \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le LR\left(7 + 2\max\left\{\sqrt{dT \log T}, d \log T\right\}\right),$

and in this case Algorithm 6 calls $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ for $O(dT \log T)$ times in T rounds.

Proof. In the proof, we use $\overline{\mathcal{K}}_t$ and \overline{x}_t to denote the values of \mathcal{K}_t and x_t at the end of iteration t $(\overline{\mathcal{K}}_0 = \mathcal{K}_0)$. We define

$$z_t(x) := e^{-\eta \sum_{i=1}^{t-1} f_i(x)}, \quad Z_t := \int_{\bar{\mathcal{K}}_t} z_t(x) \mathrm{d}x, \quad Z'_t := \int_{\bar{\mathcal{K}}_{t-1}} z_t(x) \mathrm{d}x, \quad \delta_t := \log \frac{Z'_t}{Z_t}.$$

We first prove the following two claims:

- (i) For all $t \in \{0, 1, \dots, T\}$, we have $(1 \gamma)\mathcal{K}^* + \gamma \overline{\mathcal{K}}_0 \subseteq \overline{\mathcal{K}}_t$.
- (ii) For all $t \in [T]$, we have $\delta_t \geq \frac{s_t}{5}$.

We use induction to prove (i). It holds for t = 0 since $\mathcal{K}^* \subseteq \mathcal{K}_0 = \overline{\mathcal{K}}_0$ and \mathcal{K}_0 is convex. Suppose it holds for t - 1. If $\overline{\mathcal{K}}_t = \overline{\mathcal{K}}_{t-1}$, then it already holds for t. Otherwise, consider any separating hyperplane $(w, b) \in \mathbb{R}^d \times \mathbb{R}$ obtained in round t, which is the output of $SOD(x', 2\gamma R, W)$ for some x'. By the guarantee of the SOD oracle, we have

$$w^{\top}x' \le b - 2\gamma R \le \min_{x^* \in \mathcal{K}^*} w^{\top}x^* - 2\gamma R.$$

This implies

$$(1-\gamma)\mathcal{K}^* + \gamma \bar{\mathcal{K}}_0 \subseteq (1-\gamma)\mathcal{K}^* + \gamma B(0,R) \subseteq \mathcal{K}^* + B(0,2\gamma R) \subseteq \{x \in \mathbb{R}^d : w^\top x \ge b - 2\gamma R\}.$$

Note that $\{x \in \mathbb{R}^d : w^\top x \ge b - 2\gamma R\}$ is exactly the half-space to intersect with when updating \mathcal{K}_t . Hence we know that during the execution of the algorithm, \mathcal{K}_t is always a superset of $(1-\gamma)\mathcal{K}^* + \gamma \overline{\mathcal{K}}_0$. This proves (i).

For (ii), note that each time \mathcal{K}_t is updated, the mean of the distribution over \mathcal{K}_t with density proportional to $z_t(x)$ is not included in the interior of the new \mathcal{K}_t . By Lemma D.2, this implies $\int_{\text{new }\mathcal{K}_t} z_t(x) dx \leq (1 - \frac{1}{2e}) \int_{\text{old }\mathcal{K}_t} z_t(x) dx$. Hence we have $-\delta_t = \log \frac{Z_t}{Z'_t} \leq \log (1 - \frac{1}{2e})^{s_t}$, which gives $\delta_t \geq \frac{s_t}{5}$.

Now we show the regret bound. From (i) we know

$$(1-\gamma)\mathcal{K}^* + \gamma \mathcal{\bar{K}}_0 \subseteq \mathcal{\bar{K}}_T \subseteq \mathcal{\bar{K}}_{T-1} \subseteq \cdots \subseteq \mathcal{\bar{K}}_1 \subseteq \mathcal{\bar{K}}_0 \subseteq B(0,R).$$

Therefore, we can apply Theorem B.3 to get

$$\begin{aligned} \forall x^* \in \mathcal{K}^* : \quad \sum_{t=1}^T (f_t(\bar{x}_t) - f_t(x^*)) &\leq \frac{d \log \frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + \gamma L R T - \frac{\sum_{t=1}^T \delta_t}{\eta} \\ &\leq \frac{d \log \frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + \gamma L R T - \frac{s}{5\eta}, \end{aligned}$$

where the second inequality is due to (ii).

The actual algorithm does not play \bar{x}_t , but a random \tilde{x}_t . Namely, letting $(V, p) \in \mathbb{R}^{d \times k} \times \Delta^{k-1}$ be the output of $SOD(\bar{x}_t, 2\gamma R, W)$, we have that \tilde{x}_t is equal to v_i with probability p_i (i = 1, ..., k), where $V = (v_1, ..., v_k)$. By Theorem B.1 we know that there exists $c \in W^\circ$ such that $\|\sum_{i=1}^k p_i v_i + c - \bar{x}_t\| \leq 6\gamma R$, which implies (note $f_t \in W \cap B(0, L)$)

$$\mathbb{E}[f_t(\tilde{x}_t)] = f_t\left(\sum_{i=1}^k p_i v_i\right) \le f_t\left(\sum_{i=1}^k p_i v_i + c\right) \le f_t(\bar{x}_t) + 6\gamma LR.$$

Therefore we have

$$\begin{aligned} \forall x^* \in \mathcal{K}^* : \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] &= \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(\bar{x}_t))\right] + \sum_{t=1}^T (f_t(\bar{x}_t) - f_t(x^*)) \\ &\leq 6\gamma LRT + \frac{d\log\frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + \gamma LRT - \frac{s}{5\eta} \end{aligned}$$

$$= \frac{d\log\frac{1}{\gamma}}{\eta} + \eta L^2 R^2 T + 7\gamma LRT - \frac{s}{5\eta}$$

Setting $\gamma = \frac{1}{T}$ and $\eta = \frac{1}{LR} \min\left\{1, \sqrt{\frac{d \log T}{T}}\right\}$, the above bound becomes

$$\forall x^* \in \mathcal{K}^* : \mathbb{E}\left[\sum_{t=1}^T (f_t(\tilde{x}_t) - f_t(x^*))\right] \le LR\left(7 + 2\max\left\{\sqrt{dT\log T}, d\log T\right\}\left(1 - \frac{s}{10d\log T}\right)\right)$$

We can use the above regret bound to bound the number of oracle calls in Algorithm 6. Since the regret is always lower bounded by -2LRT, the above regret upper bound implies s = O(T). Therefore Algorithm 6 calls the SOD oracle for s + T = O(T) times. Note that each implementation of SOD needs to call $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ for $O\left(d\log\frac{4R}{2\gamma R}\right) = O(d\log T)$ times (Theorem B.1). We conclude that the total number of calls to $\mathcal{O}_{\mathcal{K},\mathcal{K}^*}$ in Algorithm 6 is at most $O(dT\log T)$.

Remark. Intuitively, when the SOD oracle is called in Algorithm 6, between the two outcomes (separation and decomposition) we should prefer decomposition, since this means we can make the play and move on to the next iteration. However, Theorem B.4 shows an interesting trade-off between oracle complexity and regret: the more oracle calls, the less the regret. This means obtaining separating hyperplanes helps the regret. Interestingly, we obtain our upper bound on the oracle calls by observing that regret can never be lower by -2LRT.

C **Proof for the Bandit Setting (Theorem 4.2)**

We prove the following more general theorem than Theorem 4.2.

Theorem C.1. Denote by z_t the point played by Algorithm 3 in round t. Then for any $\gamma \in (0, 1)$, $\epsilon \in (0, \alpha R]$ and $\eta > 0$, Algorithm 3 satisfies the following regret guarantee:

$$\forall x^* \in \mathcal{K} : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(z_t) - \alpha f_t(x^*))\right] \le \frac{\alpha^2 \beta^2 d}{2\eta} + \frac{\eta L^2 R^2 d^2}{2\gamma} T + 2\gamma \alpha L R T + \epsilon L T,$$

and the expected total number of calls to the oracle $\mathcal{O}_{\mathcal{K},\alpha\mathcal{K}}$ in T rounds is at most

$$(1+\gamma T)\left(1+5d\log\frac{4\alpha R+2R\sqrt{\alpha^2\beta^2 d^2+\frac{\eta^2 L^2 R^2 d^3}{\gamma}T+6\eta\alpha LRdT}}{\epsilon}\right)$$

In particular, setting $\eta = \frac{\alpha \beta^{4/3}}{LRT^{2/3}}$, $\epsilon = \frac{\alpha R}{T}$ and $\gamma = \frac{\beta^{2/3}d}{T^{1/3}}$ (assuming $T > \beta^2 d^3$ so $\gamma < 1$), we have

$$\forall x^* \in \mathcal{K}: \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(z_t) - \alpha f_t(x^*))\right] \le \alpha LR\left(3d(\beta T)^{2/3} + 1\right),$$

and the expected total number of oracle calls in T rounds is at most $O\left(d^2(\beta T)^{2/3}\log T\right)$.

Proof. Let x_t, y_t and \tilde{x}_t be the same x_t, y_t and \tilde{x}_t appearing in Algorithm 1 during our implementation. We define $\overline{x}_t := \mathbb{E}[\tilde{x}_t|y_t]$ similarly to the proof of Theorem 3.1. It is easy to see that (φ, W') satisfies the PNIP property (Definition 2.4) and $\tilde{f}_t \in W'$ for all $t \in [T]$, where $W' = (M^{\top})^{-1} \mathbb{R}^d_+$.

Note that $\nabla\!\varphi(x)=Q^{-1}x,$ which implies

$$D_{\varphi}(x_{t}, y_{t+1}) = \frac{1}{2} (x_{t} - y_{t+1})^{\top} Q^{-1} (x_{t} - y_{t+1}) = \frac{1}{2} (\nabla \varphi(x_{t}) - \nabla \varphi(y_{t+1}))^{\top} Q (\nabla \varphi(x_{t}) - \nabla \varphi(y_{t+1}))$$
$$= \frac{\eta^{2}}{2} \tilde{f}_{t}^{\top} Q \tilde{f}_{t}.$$

Using the regret bound (3) in the proof of Theorem 3.1, for any $x^* \in \mathcal{K}$ we have

$$\sum_{t=1}^{T} \left(\tilde{f}_t(x_t) - \alpha \tilde{f}_t(x^*) \right) \leq \frac{1}{\eta} \left(D_{\varphi}(\alpha x^*, y_1) - D_{\varphi}(\alpha x^*, y_{T+1}) + \sum_{t=1}^{T} D_{\varphi}(x_t, y_{t+1}) \right)$$

$$\leq \frac{1}{\eta} \left(\varphi(\alpha x^*) - \min_{y \in \mathbb{R}^d} \varphi(y) + \sum_{t=1}^{T} D_{\varphi}(x_t, y_{t+1}) \right)$$

$$= \frac{1}{\eta} \varphi(\alpha x^*) + \frac{\eta}{2} \sum_{t=1}^{T} \tilde{f}_t^{\mathsf{T}} Q \tilde{f}_t.$$
(7)

Since $\{q_i\}_{i=1}^d$ is a β -BS(\mathcal{K}), there exist $\beta_1, \ldots, \beta_d \in [-\beta, \beta]$ such that $x^* = \sum_{i=1}^d \beta_i q_i$. Then we have

$$\varphi(\alpha x^*) = \frac{1}{2} (\alpha x^*)^\top Q^{-1}(\alpha x^*) = \frac{\alpha^2}{2} \|M^{-1} x^*\|^2 = \frac{\alpha^2}{2} \left\| \sum_{i=1}^d \beta_i M^{-1} q_i \right\|^2 = \frac{\alpha^2}{2} \left\| \sum_{i=1}^d \beta_i e_i \right\|^2 \le \frac{\alpha^2 \beta^2 d}{2}$$

We also have

$$\mathbb{E}\left[\tilde{f}_t^\top Q \tilde{f}_t\right] = \gamma \sum_{i=1}^d \frac{1}{d} \left(\frac{d}{\gamma} q_i^\top f_t\right)^2 q_i^\top Q^{-1} Q Q^{-1} q_i = \frac{d}{\gamma} \sum_{i=1}^d \left(q_i^\top f_t\right)^2 q_i^\top Q^{-1} q_i$$
$$\leq \frac{L^2 R^2 d}{\gamma} \sum_{i=1}^d q_i^\top (M M^\top)^{-1} q_i = \frac{L^2 R^2 d}{\gamma} \sum_{i=1}^d e_i^\top e_i = \frac{L^2 R^2 d^2}{\gamma}.$$

Hence by taking expectation on (7) we get

$$\mathbb{E}\left[\sum_{t=1}^{T} \left(\tilde{f}_t(x_t) - \alpha \tilde{f}_t(x^*)\right)\right] \le \frac{\alpha^2 \beta^2 d}{2\eta} + \frac{\eta L^2 R^2 d^2}{2\gamma} T.$$
(8)

Note that $\mathbb{E}[\tilde{f}_t|x_t] = \gamma \sum_{i=1}^d \frac{1}{d} \frac{d}{\gamma} Q^{-1} q_i q_i^\top f_t = Q^{-1} \left(\sum_{i=1}^d q_i q_i^\top \right) f_t = f_t$. Therefore (8) becomes

$$\mathbb{E}\left[\sum_{t=1}^{T} \left(f_t(x_t) - \alpha f_t(x^*)\right)\right] \le \frac{\alpha^2 \beta^2 d}{2\eta} + \frac{\eta L^2 R^2 d^2}{2\gamma} T.$$
(9)

Next, by the guarantee of the PAD oracle, for any $t \in [T]$ we know that there exists $c_t \in (W')^\circ$ such that $\|\overline{x}_t + c_t - x_t\| \le \epsilon$. It is easy to see that $\mathbb{R}^d_+ \subseteq W'$, which implies $(W')^\circ \subseteq \mathbb{R}^d_+$, so we know $c_t \in \mathbb{R}^d_+$. Then we have

$$\forall t \in [T]: \quad \mathbb{E}\left[f_t(\overline{x}_t) - f_t(x_t)\right] = \mathbb{E}\left[f_t^\top(\overline{x}_t - x_t)\right] \le \mathbb{E}\left[f_t^\top(\overline{x}_t + c_t - x_t)\right] \le \epsilon L.$$

Thus (9) implies

$$\mathbb{E}\left[\sum_{t=1}^{T} (f_t(\overline{x}_t) - \alpha f_t(x^*))\right] \le \frac{\alpha^2 \beta^2 d}{2\eta} + \frac{\eta L^2 R^2 d^2}{2\gamma} T + \epsilon LT.$$
(10)

Finally, since the point played in round t, z_t , is equal to \tilde{x}_t (whose expectation is \bar{x}_t) with probability $1 - \gamma$, we have

$$\mathbb{E}\left[\sum_{t=1}^{T} (f_t(z_t) - \alpha f_t(x^*))\right] \le (1 - \gamma) \mathbb{E}\left[\sum_{t=1}^{T} (f_t(\widetilde{x}_t) - \alpha f_t(x^*))\right] + \gamma \cdot 2\alpha LRT$$
$$= (1 - \gamma) \mathbb{E}\left[\sum_{t=1}^{T} (f_t(\overline{x}_t) - \alpha f_t(x^*))\right] + 2\gamma \alpha LRT$$
$$\le \frac{\alpha^2 \beta^2 d}{2\eta} + \frac{\eta L^2 R^2 d^2}{2\gamma} T + \epsilon LT + 2\gamma \alpha LRT.$$

Oracle complexity. Using Theorem 3.4, we know that when $b_t = \text{EXPLORE}$, the number of calls to the oracle $\mathcal{O}_{\mathcal{K},\alpha\mathcal{K}}$ in round t is at most $\left[5d \log \frac{4\alpha R + 2\sqrt{\frac{2}{\mu}\min_{x^*\in\mathcal{K}} D_{\varphi}(\alpha x^*, y_t)}}{\epsilon} \right]$, where μ is the strong convexity parameter of φ .

In the above proof of the regret bound, we have ignored the term $D_{\varphi}(\alpha x^*, y_{T+1})$ in (7). If we instead keep this term, the regret bound (10) will become

$$\forall x^* \in \mathcal{K} : \quad \mathbb{E}\left[\sum_{t=1}^T (f_t(\overline{x}_t) - \alpha f_t(x^*))\right] \le \frac{\alpha^2 \beta^2 d}{2\eta} + \frac{\eta L^2 R^2 d^2}{2\gamma} T + \epsilon LT - \frac{1}{\eta} \mathbb{E}\left[D_{\varphi}(\alpha x^*, y_{T+1})\right].$$

In the above inequality, substituting T with t, we have

$$\begin{aligned} \forall t \in [T], \forall x^* \in \mathcal{K}: \quad \mathbb{E}\left[D_{\varphi}(\alpha x^*, y_{t+1})\right] &\leq \frac{\alpha^2 \beta^2 d}{2} + \frac{\eta^2 L^2 R^2 d^2}{2\gamma} t + \epsilon \eta L t - \eta \mathbb{E}\left[\sum_{j=1}^t (f_j(\overline{x}_j) - \alpha f_j(x^*))\right] \\ &\leq \frac{\alpha^2 \beta^2 d}{2} + \frac{\eta^2 L^2 R^2 d^2}{2\gamma} T + \epsilon \eta L T + \eta \cdot 2\alpha L R T \\ &\leq \frac{\alpha^2 \beta^2 d}{2} + \frac{\eta^2 L^2 R^2 d^2}{2\gamma} T + 3\eta \alpha L R T. \end{aligned}$$

The above upper bound is also clearly valid for $D_{\varphi}(\alpha x^*, y_1)$.

Since $\varphi(x) = \frac{1}{2}x^{\top}Q^{-1}x$ is quadratic, we know that $\mu = \lambda_{\min}(Q^{-1}) = \frac{1}{\lambda_{\max}(Q)} = \frac{1}{\max_{u \in \mathbb{R}^d, \|u\|=1} \|Qu\|} = \frac{1}{\max_{u \in \mathbb{R}^d, \|u\|=1} \|\sum_{i=1}^d q_i q_i^{\top} u\|} \geq \frac{1}{\sum_{i=1}^d \|q_i\|^2} \geq \frac{1}{R^2d}$, where $\lambda_{\min}(P)$ and $\lambda_{\max}(P)$ are respectively the smallest and the largest eigenvalues of a symmetric matrix P.

Note that $\log(a + \sqrt{x})$ is a concave function in x for a > 0. By Jensen's inequality, the expected number of calls to the oracle $\mathcal{O}_{\mathcal{K}}^{\alpha}$ in round t when $b_t = \text{EXPLORE}$ is upper bounded by:

$$\begin{split} & \mathbb{E}\left[1+5d\log\frac{4\alpha R+2\sqrt{\frac{2}{\mu}}\min_{x^*\in\mathcal{K}}D_{\varphi}(\alpha x^*,y_t)}{\epsilon}\right] \\ & \leq 1+\min_{x^*\in\mathcal{K}}\mathbb{E}\left[5d\log\frac{4\alpha R+2\sqrt{\frac{2}{\mu}}D_{\varphi}(\alpha x^*,y_t)}{\epsilon}\right] \\ & \leq 1+\min_{x^*\in\mathcal{K}}5d\log\frac{4\alpha R+2\sqrt{\frac{2}{\mu}\mathbb{E}\left[D_{\varphi}(\alpha x^*,y_t)\right]}}{\epsilon} \\ & \leq 1+5d\log\frac{4\alpha R+2\sqrt{2R^2d\left(\frac{\alpha^2\beta^2d}{2}+\frac{\eta^2L^2R^2d^2}{2\gamma}T+3\eta\alpha LRT\right)}}{\epsilon} \\ & = 1+5d\log\frac{4\alpha R+2R\sqrt{\alpha^2\beta^2d^2+\frac{\eta^2L^2R^2d^3}{\gamma}T+6\eta\alpha LRdT}}{\epsilon}. \end{split}$$

Therefore the expected total number of calls to the oracle $\mathcal{O}_{\mathcal{K}}^{\alpha}$ in T rounds is at most

$$(1+\gamma T)\left(1+5d\log\frac{4\alpha R+2R\sqrt{\alpha^2\beta^2 d^2+\frac{\eta^2 L^2 R^2 d^3}{\gamma}T+6\eta\alpha LRdT}}{\epsilon}\right)$$

The second part of the theorem can be directly verified using the specific choices of η , ϵ and γ and noting $\log(\text{poly}(\beta d)) = O(\log T)$ since $T > \beta^2 d^3$.

D Log-Concave Distributions

A distribution over \mathbb{R}^d with a density function f is *log-concave* if $\log(f)$ is a concave function. For a convex set S equipped with a membership oracle, there exist polynomial-time algorithms for sampling from any log-concave distribution over S (Lovász and Vempala, 2007). This can be used to approximately compute the mean of any log-concave distribution.

We have the following classical result which says that every half-space close enough to the mean of a log-concave distribution must contain at least constant probability mass. For simplicity, we only state and prove the result for isotropic (i.e., identity covariance) log-concave distributions, but the result can be easily generalized to allow arbitrary covariance.

Lemma D.1. Consider any isotropic (identity covariance) log-concave distribution p over \mathbb{R}^d with mean x^* . Then for any half-space H such that $||x^* - \Pi_H(x^*)|| \leq \frac{1}{2e}$, we have $\int_H p(x) dx \geq \frac{1}{2e}$.

Proof. Let $H = \{x \in \mathbb{R}^d : w^\top x \ge b\}$ for a unit vector $w \in \mathbb{R}^d$ and $b \in \mathbb{R}$, and assume without loss of generality that $x^* = 0$. Consider the one-dimensional random variable $Y := w^\top X - b$, where $X \sim p$. Denote by $q : \mathbb{R} \to \mathbb{R}$ the density function of y. Then we have

$$\int_{H} p(x) \mathrm{d}x = \int_{0}^{\infty} q(y) \mathrm{d}y.$$

Let $y^* := \mathbb{E}[Y] = w^\top x^* - b = -b$. By our assumption, we know $|y^*| \leq \frac{1}{2e}$. Moreover, since log-concavity is preserved under linear transformations (Prékopa, 1973), we know that y also follows a log-concave distribution, and it is easy to see that it is also isotropic. Using Lemma 5.4 in (Lovász and Vempala, 2007), we know $\int_{y^*}^{\infty} q(y) dy \geq \frac{1}{e}$. In addition, from Lemma 5.5 in (Lovász and Vempala, 2007) we know $q(y) \leq 1$ ($\forall y \in \mathbb{R}$). Therefore, we have

$$\frac{1}{e} - \int_0^\infty q(y) \mathrm{d}y \le \int_{y^*}^\infty q(y) \mathrm{d}y - \int_0^\infty q(y) \mathrm{d}y \le |y^*| \sup_{y \in \mathbb{R}} q(y) \le \frac{1}{2e},$$

which implies $\int_0^\infty q(y) dy \ge \frac{1}{2e}$, completing the proof.

As an implication, we have the following lemma regarding mean computation of a log-concave distribution, which is useful in this paper.

Lemma D.2. For any log-concave distribution p in \mathbb{R}^d with mean x^* , whose support supp(p) is in B(0, R) (R > 0), and any $\epsilon > 0$ and $0 < \delta < 1$, it is possible to compute a point \tilde{x}^* in poly $\left(d, \frac{1}{\epsilon}, \log \frac{1}{\delta}\right)$ time such that with probability at least $1 - \delta$ we have: (1) $\|\tilde{x}^* - x^*\| \leq R\epsilon$; (2) for any half space H containing \tilde{x}^* , $\int_H p(x) dx \geq \frac{1}{2\epsilon}$.

For our purpose in this paper, it always suffices to choose $\epsilon = \frac{1}{T}$ and $\delta = \frac{1}{\text{poly}(T)}$ (*T* being the total number of rounds) without hurting our regret bounds.