# Online Improper Learning with an Approximation Oracle* 

Elad Hazan<br>Princeton University \& Google AI Princeton ehazan@cs.princeton.edu

Wei Hu<br>Princeton University huwei@cs.princeton.edu

Yuanzhi Li<br>Stanford University<br>yuanzhil@stanford.edu

Zhiyuan Li<br>Princeton University<br>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.

[^0]
### 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 $f_{t}^{\top} \overline{x_{t}}$ 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 $\alpha$-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 $\alpha$-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) | $\boldsymbol{O}(\sqrt{\boldsymbol{T}})$ | $\boldsymbol{O}(\log \boldsymbol{T})$ | PNIP property (Def. 2.4) |
| Alg. 6 (this paper) | $\tilde{\boldsymbol{O}}(\sqrt{\boldsymbol{T}})$ | $\boldsymbol{O}(\log \boldsymbol{T})$ | 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) | $O\left(T^{\frac{2}{3}}\right)$ | $O\left(T^{\frac{4}{3}}\right)$ | general |
| Garber (2017) | $O\left(T^{\frac{2}{3}}\right)$ | $\tilde{O}(T)$ | non-negative |
| Alg. 3 (this paper) | $\boldsymbol{O}\left(T^{\frac{2}{3}}\right)$ | $\tilde{\boldsymbol{O}}\left(T^{\frac{2}{3}}\right)$ | 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 $O\left(T^{2 / 3}\right)$ 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 $\mathcal{S}, \mathcal{S}^{\prime} \subseteq \mathbb{R}^{d}, \beta \in \mathbb{R}, y \in \mathbb{R}^{d}$ and $A \in \mathbb{R}^{d^{\prime} \times d}$, define $\mathcal{S}+\mathcal{S}^{\prime}:=\left\{x+x^{\prime}: x \in \mathcal{S}, x^{\prime} \in \mathcal{S}^{\prime}\right\}$, $\beta \mathcal{S}:=\{\beta x: x \in \mathcal{S}\}, x+\mathcal{S}:=\{x+y: y \in \mathcal{S}\}$, and $A \mathcal{S}:=\{A x: x \in \mathcal{S}\}$. The convex hull of $\mathcal{S} \subseteq \mathbb{R}^{d}$ is denoted by $\operatorname{CH}(\mathcal{S})$. Denote by $\operatorname{Vol}(\mathcal{S})$ the volume (Lebesgue measure) of a set $\mathcal{S} \subseteq \mathbb{R}^{d}$. Denote by $\Delta^{k-1}$ the probability simplex in $\mathbb{R}^{k}$.
A set $\mathcal{C} \subseteq \mathbb{R}^{d}$ is called a cone if for any $\beta \geq 0$ we have $\beta \mathcal{C} \subseteq \mathcal{C}$. For any $\mathcal{S} \subseteq \mathbb{R}^{d}$, define the dual cone of $\mathcal{S}$ as $\mathcal{S}^{\circ}:=\left\{y \in \mathbb{R}^{d}: x^{\top} y \geq 0, \forall x \in \mathcal{S}\right\}$. $\mathcal{S}^{\circ}$ is always a convex cone, even when $\mathcal{S}$ is

```
Algorithm 1 Online Mirror Descent using a Projection-and-Separation Oracle
Input: Learning rate \(\eta>0\), tolerance \(\epsilon>0\), regularizer \(\varphi\), convex cone \(W\), time horizon \(T \in \mathbb{N}_{+}\)
    \(y_{1} \leftarrow \arg \min _{y \in \operatorname{Dom}(\varphi)} \varphi(y)\).
    for \(t=1\) to \(T\) do
        \(\left(x_{t}, V=\left(v_{1}, \ldots, v_{k}\right), p\right) \leftarrow \mathcal{P} \mathcal{A D}\left(y_{t}, \epsilon, W, \varphi\right)\)
        Play \(\tilde{x}_{t}=v_{i}\) with probability \(p_{i}(i \in[k])\), and observe the loss vector \(f_{t}\).
        \(\nabla \varphi\left(y_{t+1}\right) \leftarrow \nabla \varphi\left(x_{t}\right)-\eta f_{t}\)
    end for
```

neither convex nor a cone. For any closed set $\mathcal{S} \subseteq \mathbb{R}^{d}$, define $\Pi_{\mathcal{S}}: \mathbb{R}^{d} \rightarrow \mathcal{S}$ to be the projection onto $\mathcal{S}$, namely $\Pi_{\mathcal{S}}(x):=\arg \min _{x^{\prime} \in \mathcal{S}}\left\|x^{\prime}-x\right\|^{2}$.
Definition 2.1. A strictly convex function $f: \mathcal{A} \rightarrow \mathbb{R}\left(\mathcal{A} \subseteq \mathbb{R}^{d}\right.$ is convex $)$ is Legendre if $\nabla f$ is continuous in $\operatorname{int}(\mathcal{A})$ and for any sequence $x_{1}, x_{2}, \cdots \in \mathcal{A}$ converging to a boundary point of $\mathcal{A}$, $\lim _{n \rightarrow \infty}\left\|\nabla f\left(x_{n}\right)\right\|=\infty$.
Definition 2.2. For a Legendre function $\varphi: \mathcal{A} \rightarrow \mathbb{R}$, the Bregman divergence with respect to $\varphi$ 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 $\mathcal{S} \subseteq \mathbb{R}^{d}, x \in \mathbb{R}^{d}, y \in \mathcal{S}$, and any Legendre function $\varphi: \mathbb{R}^{d} \rightarrow \mathbb{R}$, letting $z=\arg \min _{x^{\prime} \in \mathcal{S}} D_{\varphi}\left(x^{\prime}, x\right)$, we must have $D_{\varphi}(y, x) \geq D_{\varphi}(y, z)+D_{\varphi}(z, x)$.
Definition 2.4 (Pairwise non-negative inner product). For a twice-differentiable Legendre function $\varphi: \mathcal{A} \rightarrow \mathbb{R}$ with domain $\mathcal{A} \subseteq \mathbb{R}^{d}$ and a convex cone $W \subseteq \mathbb{R}^{d}$, we say $(\varphi, W)$ satisfies the pairwise non-negative inner product (PNIP) property, if for all $\bar{w}, w^{\prime} \in W$ and $H \in \mathrm{CH}(\mathcal{H})$, where $\mathcal{H}=\left\{\nabla^{2} \varphi(x): x \in \mathcal{A}\right\}$, it holds that $w^{\top} H^{-1} w^{\prime} \geq 0$.

Examples. $(\varphi, W)$ 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 $L_{d+1}=\left\{(x, z) \in \mathbb{R}^{d} \times \mathbb{R}:\|x\|_{2} \leq z\right\}$.
2. $\varphi(x)=\sum_{i=1}^{d} x_{i}\left(\log x_{i}-1\right)$ (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 $Q=M M^{\top}, M \in \mathbb{R}^{d \times d}$ is an invertible matrix, and $W=\left(M^{\top}\right)^{-1} \mathbb{R}_{+}^{d}$.

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 $\mathcal{S}$ equipped with a membership oracle, there exist polynomial-time algorithms for sampling from any log-concave distribution over $\mathcal{S}$ (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 logconcave 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 $\varphi$, which achieves $O(\sqrt{T})$ regret when the regularizer $\varphi$ 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 $\left\{f_{t}\right\}$ come from a convex cone $W \subseteq \mathbb{R}^{d}$ and $\left\|f_{t}\right\| \leq L(R, L>0)$. Omitted proofs in this section are given in Appendix A.
Theorem 3.1. Suppose $(\varphi, W)$ satisfies the PNIP property (Definition 2.4). Then for any $\epsilon, \eta>0$, Algorithm 1 satisfies the following regret guarantee:

$$
\forall x^{*} \in \mathcal{K}^{*}: \quad \mathbb{E}\left[\sum_{t=1}^{T}\left(f_{t}\left(\tilde{x}_{t}\right)-f_{t}\left(x^{*}\right)\right)\right] \leq \frac{1}{\eta}\left(\varphi\left(x^{*}\right)-\varphi\left(y_{1}\right)+\sum_{t=1}^{T} D_{\varphi}\left(x_{t}, y_{t+1}\right)\right)+\epsilon L T .
$$

In particular, if $\varphi$ is $\mu$-strongly convex and $A \geq \max _{x^{*} \in \mathcal{K}^{*}}\left(\varphi\left(x^{*}\right)-\varphi\left(y_{1}\right)\right)$, setting $\epsilon=\frac{R}{T}$ and $\eta=\frac{1}{L} \sqrt{\frac{2 \mu A}{T}}$, we have

$$
\forall x^{*} \in \mathcal{K}^{*}: \quad \mathbb{E}\left[\sum_{t=1}^{T}\left(f_{t}\left(\tilde{x}_{t}\right)-f_{t}\left(x^{*}\right)\right)\right] \leq L \sqrt{\frac{2 A T}{\mu}}+L R,
$$

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

For the problem of $\alpha$-regret minimization using an $\alpha$-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:

$$
\forall x^{*} \in \mathcal{K}: \quad \mathbb{E}\left[\sum_{t=1}^{T} f_{t}\left(\tilde{x}_{t}\right)-\alpha \sum_{t=1}^{T} f_{t}\left(x^{*}\right)\right]=\mathbb{E}\left[\sum_{t=1}^{T} f_{t}\left(\tilde{x}_{t}\right)-\sum_{t=1}^{T} f_{t}\left(\alpha x^{*}\right)\right] \leq \alpha L R(\sqrt{T}+1) .
$$

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 $\mathcal{K}^{*}, \mathcal{P} \mathcal{A D}(y, \epsilon, W, \varphi)$, is defined as a procedure that given $y \in \mathbb{R}^{d}, \epsilon>0$, a convex cone $W$ and a Legendre function $\varphi$ produces a tuple $\left(y^{\prime}, V, p\right)$, where $y^{\prime} \in \mathbb{R}^{d}, V=\left(v_{1}, \ldots, v_{k}\right) \in \mathbb{R}^{d \times k}$ and $p=\left(p_{1}, \ldots, p_{k}\right)^{\top} \in \Delta^{k-1}$, such that:

1. $y^{\prime}$ is "closer" to $\mathcal{K}^{*}$ than $y$ with respect to the Bregman divergence of $\varphi$ (and hence is an "infeasible projection"): $\forall x^{*} \in \mathcal{K}^{*}, D_{\varphi}\left(x^{*}, y^{\prime}\right) \leq D_{\varphi}\left(x^{*}, y\right)$;
2. $v_{1}, \ldots, v_{k} \in \mathcal{K}$, and $\sum_{i=1}^{k} p_{i} v_{i}$ is a point that "almost dominates" $y^{\prime}$ in all directions in $W$. In other words, there exists $c \in W^{\circ}$ such that $\left\|\sum_{i=1}^{k} p_{i} v_{i}+c-y^{\prime}\right\| \leq \epsilon$.
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 $y^{\prime}$ 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 $y^{\prime}$ 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 $y^{\prime}$ 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 ( $\varphi, W$ ) satisfies PNIP condition (Definition 2.4) and $\varphi$ is $\mu$-strongly convex. Then for any $y \in \mathbb{R}^{d}, \epsilon \in(0, R]$, Algorithm 2 terminates in $\left\lceil 5 d \log \frac{4 R+2 \sqrt{2 \min _{x^{*} \in \mathcal{K}^{*}} D_{\varphi}\left(x^{*}, y\right) / \mu}}{\epsilon}\right\rceil$ iterations, and it correctly implements the projection-and-decomposition oracle $\mathcal{P} \mathcal{A D}(y, \epsilon, W, \varphi)$, i.e., its output $\left(y^{\prime}, V, p\right)$ 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 $O\left(d^{2} \log T\right)$ 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, \(\mathcal{P} \mathcal{A D}(y, \epsilon, W, \varphi)\)
Input: Point \(y \in \mathbb{R}^{d}\), tolerance \(\epsilon>0\), convex cone \(W\), regularizer \(\varphi\),
Output: \(\left(y^{\prime}, V, p\right)\), where \(y^{\prime} \in \mathbb{R}^{d}, V=\left(v_{1}, \ldots, v_{k}\right) \in \mathbb{R}^{d \times k}\) for some \(k\) such that \(v_{i} \in \mathcal{K}\)
    \((\forall i \in[k])\), and \(p=\left(p_{1}, \ldots, p_{k}\right)^{\top} \in \Delta^{k-1}\)
    \(W_{1} \leftarrow W \cap B(0,1), \quad z_{1} \leftarrow y, \quad i \leftarrow 0\)
    while \(i<5 d \log \frac{2\left(R+\left\|z_{i+1}\right\|\right)}{\epsilon}\) do
        \(i \leftarrow i+1, \quad w_{i} \leftarrow \frac{\int_{W_{i}}^{\epsilon} w \mathrm{~d} w}{\operatorname{Vol}\left(W_{i}\right)}, \quad v_{i} \leftarrow \mathcal{O}_{\mathcal{K}, \mathcal{K}^{*}}\left(w_{i}\right)\).
        \(z_{i+1} \leftarrow \underset{z \in \mathbb{R}^{d}, w_{i}^{\top}\left(z-v_{i}\right) \geq 0}{\arg \min } D_{\varphi}\left(z, z_{i}\right), \quad W_{i+1} \leftarrow W_{i} \cap\left\{w \in \mathbb{R}^{d}: w^{\top}\left(v_{i}-z_{i+1}\right) \geq 0\right\}\).
    end while
    \(k \leftarrow i\) and solve \(\min _{p \in \Delta^{k-1}, c \in W^{\circ}}\left\|\sum_{i=1}^{k} p_{i} v_{i}+c-z_{k+1}\right\|\) to get \(p\)
    return \(y^{\prime}=z_{k+1}, V=\left(v_{1}, \ldots, v_{k}\right), p\)
```

We break the proof of Theorem 3.4 into several lemmas.
Lemma 3.5. If $(\varphi, W)$ 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 $z_{i+1}=\underset{z \in \mathbb{R}^{d}: w_{i}^{\top}\left(z-v_{i}\right) \geq 0}{\arg \min } D_{\varphi}\left(z, z_{i}\right)$, by the KKT condition, we have

$$
0=\left.\frac{\partial}{\partial z}\left(D_{\varphi}\left(z, z_{i}\right)-\lambda w_{i}^{\top}\left(z-v_{i}\right)\right)\right|_{z=z_{i+1}}=\nabla \varphi\left(z_{i+1}\right)-\nabla \varphi\left(z_{i}\right)-\lambda w_{i}
$$

for some $\lambda \geq 0$. On the other hand, note that $\nabla \varphi\left(z_{i+1}\right)-\nabla \varphi\left(z_{i}\right)=\int_{0}^{1} \nabla^{2} \varphi\left(\gamma z_{i+1}+(1-\gamma) z_{i}\right)$. $\left(z_{i+1}-z_{i}\right) \mathrm{d} \gamma=H\left(z_{i+1}-z_{i}\right)$, for some $H \in \mathrm{CH}(\mathcal{H})$, where $\mathcal{H}=\left\{\nabla^{2} \varphi(x): x \in \operatorname{Dom}(\varphi)\right\}$. Therefore, for all $w \in W$ we have $w^{\top}\left(z_{i+1}-z_{i}\right)=w^{\top} H^{-1} H\left(z_{i+1}-z_{i}\right)=\lambda w^{\top} H^{-1} w_{i} \geq 0$. 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 $\left\lceil 5 d \log \frac{4 R+2 \sqrt{2 \min _{x^{*} \in \mathcal{K}^{*}} D_{\varphi}\left(x^{*}, y\right) / \mu}}{\epsilon}\right\rceil$ iterations.

Proof. According to the algorithm, for each $i, z_{i+1}$ is the Bregman projection of $z_{i}$ onto a halfspace containing $\mathcal{K}^{*}$, since the oracle $\mathcal{O}_{\mathcal{K}, \mathcal{K}^{*}}$ ensures $w_{i}^{\top} v_{i} \leq w_{i}^{\top} x^{*}$ for all $x^{*} \in \mathcal{K}^{*}$. Then by the generalized Pythagorean theorem (Lemma 2.3) we know $D_{\varphi}\left(x^{*}, z_{i+1}\right) \leq D_{\varphi}\left(x^{*}, z_{i}\right)$ for all $x^{*} \in \mathcal{K}^{*}$ and $i$. Therefore we have $D_{\varphi}\left(x^{*}, z_{i}\right) \leq D_{\varphi}\left(x^{*}, z_{1}\right)=D_{\varphi}\left(x^{*}, y\right)$ for all $x^{*} \in \mathcal{K}^{*}$ and $i$.

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

Lemma 3.7. Under Theorem 3.4's setting, $\forall w \in W,\|w\|=1, \exists i \in[k]$, such that $w^{\top}\left(v_{i}-y^{\prime}\right) \leq \epsilon$.
Proof. We assume for contradiction that there exists a unit vector $h \in W$ such that $\min _{i \in[k]} h^{\top}\left(v_{i}-\right.$ $\left.y^{\prime}\right)>\epsilon$. Note that $\left\|v_{i}-y^{\prime}\right\| \leq\left\|v_{i}\right\|+\left\|y^{\prime}\right\| \leq R+\left\|y^{\prime}\right\|$. Letting $r:=\frac{\epsilon}{2\left(R+\left\|y^{\prime}\right\|\right)}$, we have $\forall w \in \frac{h}{2}+(W \cap B(0, r)): \quad \min _{i \in[k]} w^{\top}\left(v_{i}-y^{\prime}\right)>0$.
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 $w \in W_{1} \backslash W_{k+1}$, there exists $i \in[k]$ such that $w^{\top}\left(v_{i}-z_{i+1}\right) \leq 0$. Notice that from Lemma 3.5 we know $z_{j+1}-z_{j} \in W^{\circ}$ for all $j \in[k]$. Thus for all $w \in W_{1} \backslash W_{k+1}$ there exists $i \in[k]$ such that $w^{\top}\left(v_{i}-y^{\prime}\right)=w^{\top}\left(v_{i}-z_{k+1}\right) \leq w^{T}\left(v_{i}-z_{i+1}\right) \leq 0$. In other words, we have $\forall w \in W_{1} \backslash W_{k+1}: \quad \min _{i \in[k]} w^{\top}\left(v_{i}-y^{\prime}\right) \leq 0$.

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

$$
\begin{aligned}
\operatorname{Vol}\left(W_{1}\right) & =\operatorname{Vol}(W \cap B(0,1))=r^{-d} \operatorname{Vol}(W \cap B(0, r)) \leq r^{-d} \operatorname{Vol}\left(W_{k+1}\right) \\
& \leq r^{-d}(1-1 /(2 e))^{k} \operatorname{Vol}\left(W_{1}\right)<\operatorname{Vol}\left(W_{1}\right)
\end{aligned}
$$

where the last step is due to $k \geq 5 d \log \frac{1}{r}=5 d \log \frac{2\left(R+\left\|y^{\prime}\right\|\right)}{\epsilon}=5 d \log \frac{2\left(R+\left\|z_{k+1}\right\|\right)}{\epsilon}$, 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 $v_{1}, \ldots, v_{k} \in \mathbb{R}^{d}, \epsilon \geq 0$ 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 $p=\left(p_{1}, \ldots, p_{k}\right)^{\top} \in \Delta^{k-1}$ and $c \in W^{\circ}$ such that $\left\|\sum_{i=1}^{k} p_{i} v_{i}+c-x\right\| \leq \epsilon$.
(B) For all $w \in W,\|w\|=1$, there exists $i \in[k]$ such that $w^{\top}\left(v_{i}-x\right) \leq \epsilon$.


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

$$
x \in \mathrm{CH}\left(\left\{v_{i}\right\}_{i=1}^{k}\right) \Longleftrightarrow \text { There is no hyperplane that separates } x \text { and all } v_{i} \text { 's. }
$$

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

$$
x \in \mathrm{CH}\left(\left\{v_{i}\right\}_{i=1}^{k}\right)+W^{\circ} \Longleftrightarrow \text { There is no direction } w \in W \text { such that } w^{\top} x<w^{\top} v_{i} \text { for all } i
$$

Denote $F:=\mathrm{CH}\left(\left\{v_{i}\right\}_{i=1}^{k}\right)+W^{\circ}$. For the " $\Rightarrow$ " side, if $x \in F$, it is clear that for all $w \in W$ we must have $w^{\top} x \geq w^{\top} v_{i}$ 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 $D_{\varphi}\left(x^{*}, z_{i+1}\right) \leq D_{\varphi}\left(x^{*}, z_{i}\right)$ for all $x^{*} \in \mathcal{K}^{*}$ and $i$. This implies $D_{\varphi}\left(x^{*}, y^{\prime}\right)=D_{\varphi}\left(x^{*}, z_{k+1}\right) \leq D_{\varphi}\left(x^{*}, z_{k}\right) \leq \cdots \leq D_{\varphi}\left(x^{*}, z_{1}\right)=D_{\varphi}\left(x^{*}, y\right)$ 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 \(\epsilon>0,\left\{q_{1}, \ldots, q_{d}\right\}\) - a \(\beta\) - \(\mathrm{BS}(\mathcal{K})\) for some \(\beta>0\), exploration
    probability \(\gamma \in(0,1)\), time horizon \(T \in \mathbb{N}_{+}\)
    Instantiate Algorithm 1 with parameters \(\eta, \epsilon, \varphi(x)=\frac{1}{2} x^{\top} Q^{-1} x, W^{\prime}=\left(M^{\top}\right)^{-1} \mathbb{R}_{+}^{d}\), and \(T\)
    for \(t=1\) to \(T\) do
        Receive \(\tilde{x}_{t}\) (the point to play in round \(t\) ) from Algorithm 1
        \(b_{t} \leftarrow \begin{cases}\text { EXPLORE, } & \text { with probability } \gamma \\ \text { EXPLOIT, } & \text { with probability } 1-\gamma\end{cases}\)
        if \(b_{t}=\) EXPLORE then
            Sample \(i_{t} \in[d]\) uniformly at random, and play \(q_{i_{t}}\)
            Receive loss \(l_{t}=q_{i_{t}}^{\top} f_{t}\)
            \(\tilde{f}_{t} \leftarrow \frac{d}{\gamma} l_{t} Q^{-1} q_{i_{t}}\)
        else
            Play \(\tilde{x}_{t}\) and receive loss \(l_{t}=\tilde{x}_{t}^{\top} f_{t}\)
            \(\tilde{f}_{t} \leftarrow 0\)
        end if
        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
        Algorithm 1 can simply play according to the distribution computed in this round without any
        oracle calls.)
    end for
```


## $\alpha$-Regret Minimization in the Bandit Setting

In this section we consider the $\alpha$-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 $\left\{f_{t}\right\}$ come from $\mathbb{R}^{d}$ and $\left\|f_{t}\right\| \leq L$. Similar to (Kakade et al., 2009), we assume we know a $\beta$-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 $\left\{q_{1}, \ldots, q_{d}\right\} \subset \mathbb{R}^{d}$ is a $\beta$-barycentric spanner for a set $\mathcal{K} \subset \mathbb{R}^{d}$, denoted by $\beta-B S(\mathcal{K})$, if $\left\{q_{1}, \ldots, q_{d}\right\} \subseteq \mathcal{K}$ 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 $\left\{q_{1}, \ldots, q_{d}\right\}$ which is a $\beta$ - $\operatorname{BS}(\mathcal{K})$, define $Q:=\sum_{i=1}^{d} q_{i} q_{i}^{\top}$ and $M:=\left(q_{1}, \ldots, q_{d}\right) \in \mathbb{R}^{d \times d}$.
The need for a new regularization. The bandit algorithm of Garber (2017) additionally requires a certain boundedness property of barycentric spanners, namely:

$$
\max _{i \in[d]} q_{i}^{\top} Q^{-2} q_{i} \leq \chi
$$

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 $\eta=\frac{\alpha \beta^{4 / 3}}{L R T^{2 / 3}}, \epsilon=\frac{\alpha R}{T}$ 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

$$
\forall x^{*} \in \mathcal{K}: \quad \mathbb{E}\left[\sum_{t=1}^{T}\left(f_{t}\left(z_{t}\right)-\alpha f_{t}\left(x^{*}\right)\right)\right] \leq \alpha L R\left(3 d(\beta T)^{2 / 3}+1\right)
$$

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 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 $®$ 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.


[^0]:    ${ }^{*}$ The full version of this paper can be found on https://arxiv.org/abs/1804.07837.

