
Improved Regret Bounds for Projection-free Bandit Convex Optimization

Dan Garber¹

Technion - Israel Institute of Technology

Ben Kretzu¹

Technion - Israel Institute of Technology

Abstract

We revisit the challenge of designing online algorithms for the bandit convex optimization problem (BCO) which are also scalable to high dimensional problems. Hence, we consider algorithms that are *projection-free*, i.e., based on the conditional gradient method whose only access to the feasible decision set is through a linear optimization oracle (as opposed to other methods which require potentially much more computationally-expensive subprocedures, such as computing Euclidean projections). We present the first such algorithm that attains $O(T^{3/4})$ expected regret using only $O(T)$ overall calls to the linear optimization oracle, in expectation, where T is the number of prediction rounds. This improves over the $O(T^{4/5})$ expected regret bound recently obtained by Chen et al. (2019), and actually matches the current best regret bound for projection-free online learning in the *full information* setting.

1 INTRODUCTION

In this work we are interested in the design of efficient algorithms for online learning Cesa-Bianchi and Lugosi (2006); Hazan (2016); Bubeck et al. (2012) which lie at the intersection of two families of algorithms, each by its own studied quite extensively in recent years with many new and exciting discoveries. The first, is the family of online learning algorithms for the *bandit convex optimization* problem, and the second

is the family of so-called *projection-free* algorithms, which is a term casually used to refer to algorithms which are based on the conditional gradient method (aka Frank-Wolfe method), a well known first-order method for continuous optimization. These algorithms are called projection-free since, as opposed to popular first-order alternatives such as the projected / proximal / mirror gradient methods, which require in many cases to solve computationally-expensive optimization problems over the feasible domain (i.e., the projection step, which for instance in case of Euclidean projection, amounts to minimizing a quadratic function over the feasible set), the conditional gradient method only requires to minimize a linear function over the feasible set, which is in many cases much more efficient.

The bandit feedback model is well motivated by natural settings in which the online learner, upon making his prediction, only observes the loss associated with his prediction, and cannot infer the loss of different actions. The projection-free model is mostly motivated by large-scale settings which involve high-dimensional decision sets with non-trivial structure, for which computing Euclidean / mirror projections, which are required by standard algorithms (e.g., the celebrated online gradient descent algorithm Zinkevich (2003) and its adaptation to the BCO setting Flaxman et al. (2005)), is computationally impractical (e.g., convex relaxations for sets of low-rank matrices or polytopes with special combinatorial structure, see Jaggi (2013) and Hazan and Kale (2012) for many examples). Thus, the combination of these two basic ingredients, both concerning the possibility of applying online algorithms to large-scale real-world problems, is of interest.

A first attempt to combine these two ingredients was recently made in Chen et al. (2019), who combined the Online Frank-Wolfe method, suggested in Hazan and Kale (2012), with the framework introduced in Flaxman et al. (2005) for reducing BCO to the full-information setting (also known as *online convex optimization* (OCO)), to obtain an algorithm that achieves expected regret of $O(T^{4/5})$ (treating all quantities ex-

¹The authors are arranged according to the alphabet

cept for the number of prediction rounds T as constants), using overall T calls to the linear optimization oracle of the feasible set. Unfortunately, this regret bound is higher than both the expected regret achieved by the original method of Flaxman et al. (2005) (though Flaxman et al. (2005) uses Euclidean projections), which is $O(T^{3/4})$, and the regret obtained by the current state-of-the-art projection-free method (at least for arbitrary convex sets²) for the *full-information* setting Hazan and Kale (2012), which is also $O(T^{3/4})$.

It is thus natural to ask whether there is a price to pay, in terms of the worst-case expected regret bound, for combining these two settings, or alternatively, whether it is possible to obtain the best of both worlds, and get a projection-free algorithm for BCO that matches the state-of-the-art for the full-information setting.

In this work we show that the latter is the case, i.e., we give a projection-free algorithm for BCO which attains $O(T^{3/4})$ expected regret bound, and uses overall only $O(T)$ calls to the linear optimization oracle, in expectation, thus matching the current state-of-the-art for projection-free algorithms even in the full-information setting. See also Table 1.

In terms of techniques, as in Chen et al. (2019), our method is also based on combining the BCO framework of Flaxman et al. (2005) and the online Frank-Wolfe method Hazan and Kale (2012). The main novelty in our algorithm and analysis is based on the simple idea of partitioning the prediction rounds into non-overlapping equally-sized blocks. Surprisingly, by carefully analyzing the variance of the gradient estimator on each block, this simple trick allows us to strike a better and crucial tradeoff between the accuracy to which the subproblems of the Regularized-Follow-the-Leader method (the meta online learning algorithm on which our work, as well as Hazan and Kale (2012); Chen et al. (2019), is based) could be solved (via the conditional gradient method), and the overall regret of the algorithm. This results in meeting the current state-of-the-art bound for projection-free online convex optimization over general sets (even with full information of the loss functions), while maintaining linear (in T) optimization oracle complexity.

1.1 Additional Related Work

As discussed, both the subject of designing projection-free methods for continuous optimization and bandit convex optimization have been studied extensively in

recent years.

Projection-free Methods: the conditional gradient method, which is the basic technique in most so-called projection-free methods, dates back to the classical works of Frank and Wolfe Frank and Wolfe (1956), and Polak Levitin and Polyak (1966). The method has regained interest in recent years, especially in the context of large scale optimization and machine learning, see for instance Jaggi (2013); Jaggi and Sulovský (2010); Lacoste-Julien et al. (2013); Garber and Hazan (2015); Dudík et al. (2012); Harchaoui et al. (2012); Shalev-Shwartz et al. (2011); Laue (2012), just to name a few. There is also a recent effort to prove faster rates for simple variants of the method, usually under the assumption that the objective function is strongly convex (or a slightly weaker assumption) and assuming the feasible set admits certain structure (e.g., polytope, strongly convex set, bounded positive semidefinite cone, etc.), see for instance Garber and Hazan (2016); Lacoste-Julien and Jaggi (2015); Garber and Hazan (2015); Garber (2016); Garber and Meshi (2016); Allen-Zhu et al. (2017). Hazan and Kale (2012) were the first to suggest an algorithm for online convex optimization based on the conditional gradient method. Their method makes a single call to the linear optimization oracle on each round and achieves a regret bound of $O(T^{3/4})$ for convex loss functions with bounded gradients (note that this is worse than the optimal bound of $O(\sqrt{T})$, achievable for instance via the projection-based online gradient descent method Zinkevich (2003)). To date, this regret bound is the state-of-the-art for arbitrary compact and convex feasible sets. Garber and Hazan (2016) presented projection-free algorithms for OCO with optimal dependence on T (i.e., \sqrt{T}), in case the feasible set is a polytope. Very recently, Levy and Krause (2019) suggested a regret-optimal algorithm for OCO in case the feasible set is smooth, however, as opposed to previous works, with an algorithm that is not based on the conditional gradient method.

Bandit Convex Optimization: following the work Flaxman et al. (2005), which presented an algorithm with $O(T^{3/4})$ expected regret bound for convex and Lipschitz loss functions, several other works obtained improved bounds, mostly under an additional smoothness assumption on the losses, see for instance Saha and Tewari (2011); Dekel et al. (2015); Yang and Mohri (2016); Hazan and Levy (2014). In particular, in a recent effort, a series of works accumulated to a regret-optimal algorithm for BCO, achieving $\tilde{O}(\sqrt{T})$ regret Bubeck et al. (2015); Bubeck and Eldan (2016); Bubeck et al. (2017); Hazan and Li (2016), where \tilde{O} suppresses a polylogarithmic term of T . Importantly, all these works which improve upon the $O(T^{3/4})$ bound

²for feasible sets with specific structure such as polytopes or smooth sets there are other algorithms that obtain optimal regret bounds in T Garber and Hazan (2016); Levy and Krause (2019).

Table 1: Comparison of regret bounds and optimization oracle complexity. Only dependence on T is stated.

METHOD	FEEDBACK	PROJECTION-FREE?	ORACLE COMPLEXITY	$\mathbb{E}[\text{regret}]$
Flaxman et al. (2005)	Bandit	x	T projections	$T^{3/4}$
Hazan and Kale (2012)	Full	✓	T linear opt. steps	$T^{3/4}$
Chen et al. (2019)	Bandit	✓	T linear opt. steps	$T^{4/5}$
This work (Thm. 1)	Bandit	✓	T linear opt. steps	$T^{3/4}$

of Flaxman et al. (2005), are based on much more complicated algorithms with running time either exponential in the dimension of the problem and T , or polynomial with a high-degree polynomial, and hence have impractical running times for large-scale problems. On the other-hand, in Flaxman et al. (2005), the only non-trivial operation is that of computing a Euclidean projection, which, as we show in this work, can be roughly speaking, replaced with a linear optimization step.

Finally, a special case of BCO in which all loss functions are linear was also studied extensively due to its special structure, see for instance Auer et al. (2002); Abernethy et al. (2008); Dani et al. (2008); Hazan and Karnin (2016).

We also refer the interested reader to the following excellent introductory books on online learning and online convex optimization Cesa-Bianchi and Lugosi (2006); Bubeck et al. (2012); Hazan (2016).

2 PRELIMINARIES

2.1 Bandit Convex Optimization And Assumptions

We recall that in the bandit convex optimization problem, an online learner is required to iteratively draw actions from a fixed feasible set $\mathcal{K} \subseteq \mathbb{R}^n$.³ After choosing his action $\mathbf{x}_t \in \mathcal{K}$ on round $t \in [T]$ (T is assumed to be known beforehand), he observes his loss given by $f_t(\mathbf{x}_t)$, where $f_t : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex over \mathbb{R}^n and chosen by an adversary. Importantly, besides the value $f_t(\mathbf{x}_t)$, the learner does not gain any additional knowledge of $f_t(\cdot)$. In this work, we assume the adversary is oblivious, i.e., the loss functions f_1, \dots, f_T are chosen beforehand and do not depend on the actions of the learner.

The goal of the learner is to minimize the expected regret which is given by

$$\mathbb{E}[\mathcal{R}_T] := \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_t)] - \min_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T f_t(\mathbf{x}). \quad (1)$$

³For convenience we assume the linear space of interest is \mathbb{R}^n , however naturally, any finite-dimensional Euclidean space will work.

In this work, in addition to assuming the loss functions are convex, we also make some standard assumptions that they are bounded by M in ℓ_∞ norm and have subgradients upper-bounded by G in ℓ_2 norm over the feasible set \mathcal{K} , for some $M, G > 0$. That is, $\forall t \in [T] \forall \mathbf{x} \in \mathcal{K} \forall \mathbf{g} \in \partial f_t(\mathbf{x}): \|f_t\|_\infty = \sup_{\mathbf{x} \in \mathcal{K}} |f_t(\mathbf{x})| \leq M$ and $\|\mathbf{g}\|_2 \leq G$. Also, as in Flaxman et al. (2005) we make the standard assumption that the feasible set \mathcal{K} is full dimensional, contains the origin, and that there exist scalars $r, R > 0$ such that $r\mathcal{B}^n \subseteq \mathcal{K} \subseteq R\mathcal{B}^n$, where \mathcal{B}^n denotes the unit Euclidean ball centered at the origin in \mathbb{R}^n .

2.2 Additional Notation And Definitions

We denote by \mathcal{S}^n the unit sphere in \mathbb{R}^n , and we write $\mathbf{u} \sim \mathcal{S}^n$ and $\mathbf{u} \sim \mathcal{B}^n$ to denote a random vector \mathbf{u} sampled uniformly from \mathcal{S}^n and \mathcal{B}^n , respectively. We denote by $\|\mathbf{x}\|$ the ℓ_2 norm of the vector \mathbf{x} .

Finally, for a compact and convex set $\mathcal{K} \subseteq \mathbb{R}^n$, which satisfies the above assumptions (i.e., $r\mathcal{B}^n \subseteq \mathcal{K} \subseteq R\mathcal{B}^n$), and a scalar $0 < \delta \leq r$, we define the set $\mathcal{K}_\delta := (1 - \delta/r)\mathcal{K} = \{(1 - \delta/r)\mathbf{x} \mid \mathbf{x} \in \mathcal{K}\}$. In particular, it holds that $\mathcal{K}_\delta \subseteq \mathcal{K}$ and for all $\mathbf{x} \in \mathcal{K}_\delta$, $\mathbf{x} + \delta\mathcal{B}^n \subseteq \mathcal{K}$ (see Hazan (2016)).

We now recall some standard definitions from continuous optimization. For all definitions we assume that \mathcal{K} is a convex and compact subset of \mathbb{R}^n .

Definition 1. We say that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is G -Lipschitz over \mathcal{K} if $\forall \mathbf{x}, \mathbf{y} \in \mathcal{K}: |f(\mathbf{x}) - f(\mathbf{y})| \leq G\|\mathbf{x} - \mathbf{y}\|$.

Here we recall, that if f is convex over \mathbb{R}^n with subgradients upper-bounded by G in ℓ_2 -norm over \mathcal{K} , then f is G -Lipschitz over \mathcal{K} .

Definition 2. We say that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is β -smooth over \mathcal{K} if $\forall \mathbf{x}, \mathbf{y} \in \mathcal{K}$:

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{\beta}{2} \|\mathbf{y} - \mathbf{x}\|^2.$$

Definition 3. We say that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is α -strongly convex over \mathcal{K} if $\forall \mathbf{x}, \mathbf{y} \in \mathcal{K}$:

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{\alpha}{2} \|\mathbf{y} - \mathbf{x}\|^2.$$

Let \mathbf{x}^* be the unique minimizer of f , an α -strongly convex function over \mathcal{K} . From the above definition and the first order optimality condition it follows that $\forall \mathbf{x} \in \mathcal{K}$:

$$\frac{\alpha}{2} \|\mathbf{x} - \mathbf{x}^*\|^2 \leq f(\mathbf{x}) - f(\mathbf{x}^*). \quad (2)$$

2.3 Basic Algorithmic Ingredients

In this section we introduce some basic and standard algorithmic tools on which our algorithm is based.

2.3.1 Regularized Follow The Leader

One component of our algorithm is a variant of Regularized-Follow-the-Leader (RFTL), which is a well known algorithm for online convex optimization (in the full information setting) Hazan (2016); Shalev-Shwartz et al. (2012). The prediction in time t is according to the following rule

$$\mathbf{x}_t = \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} \left\{ \sum_{i=1}^{t-1} f_i(\mathbf{x}) + \mathcal{R}(\mathbf{x}) \right\},$$

where $\mathcal{R}(\mathbf{x})$ is a strongly convex function.

Lemma 1 (Lemma 2.3 in Shalev-Shwartz et al. (2012)). For all $t \in [T]$ let $\mathbf{x}_t^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} \left\{ \sum_{i=1}^{t-1} f_i(\mathbf{x}) + \mathcal{R}(\mathbf{x}) \right\}$.

Then, $\forall \mathbf{x} \in \mathcal{K}$ it holds that

$$\begin{aligned} \sum_{t=1}^T (f_t(\mathbf{x}_t^*) - f_t(\mathbf{x})) &\leq \mathcal{R}(\mathbf{x}) - \mathcal{R}(\mathbf{x}_1^*) \\ &+ \sum_{t=1}^T (f_t(\mathbf{x}_t^*) - f_t(\mathbf{x}_{t+1}^*)). \end{aligned}$$

2.3.2 Smoothed Loss Functions

Another standard component of our algorithm is the use of a smoothed version of each loss function. We define the δ -smoothing of a loss function f by $\hat{f}_\delta(\mathbf{x}) = \mathbb{E}_{\mathbf{u} \sim \mathcal{B}^n} [f(\mathbf{x} + \delta \mathbf{u})]$. We now cite several useful lemmas regarding smoothed functions.

Lemma 2 (Lemma 2.1 in Hazan (2016)). Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be convex and G -Lipschitz over a convex and compact set $\mathcal{K} \subset \mathbb{R}^n$. Then \hat{f}_δ is convex and G -Lipschitz over \mathcal{K}_δ , and $\forall \mathbf{x} \in \mathcal{K}_\delta$ it holds that $|\hat{f}_\delta(\mathbf{x}) - f(\mathbf{x})| \leq \delta G$.

Lemma 3 (Lemma 6.5 in Hazan (2016)). $\hat{f}_\delta(\mathbf{x})$ is differentiable and $\nabla \hat{f}_\delta(\mathbf{x}) = \mathbb{E}_{\mathbf{u} \sim \mathcal{S}^n} \left[\frac{n}{\delta} f(\mathbf{x} + \delta \mathbf{u}) \mathbf{u} \right]$.

Lemma 4 (see Bertsekas (1973)). Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be convex and suppose that all subgradients of f are upper-bounded by G in ℓ_2 -norm over a convex and compact set $\mathcal{K} \subset \mathbb{R}^n$. Then, for any $\mathbf{x} \in \mathcal{K}_\delta$ it holds that $\|\nabla \hat{f}_\delta(\mathbf{x})\| \leq G$.

3 ALGORITHM AND ANALYSIS

As in Chen et al. (2019), our algorithm (see Algorithm 1 below) is based on combining the BCO framework of Flaxman et al. (2005) with the Online Frank-Wolfe method of Hazan and Kale (2012). That is, the algorithm applies the Regularized-Follow-the-Leader meta-algorithm with Euclidean regularization

(see Section 2.3.1), and uses the bandit feedback to construct unbiased estimates for the gradients of the smoothed losses, by sampling points in a sphere around the current iterate (see Section 2.3.2). In order to avoid solving the RFTL optimization problem (which with the standard linearization trick of the smoothed losses and using Euclidean regularization, amounts to minimizing a quadratic function over the set \mathcal{K}_δ), we invoke the conditional gradient method (see Algorithm 2), to solve this problem only to sufficient approximation using only linear optimization steps over the feasible domain \mathcal{K}_δ .

Very importantly, different from Chen et al. (2019), we partition the T prediction rounds into non-overlapping blocks of size K (K is a parameter determined in the analysis), where on each block the iterate of the algorithm remains unchanged (though we use fresh samples for exploration on each round within a block). Essentially without losing generality we assume that T/K is an integer. This partition into blocks is important since as we show in the analysis, it allows us to solve the RFTL objective via the conditional gradient method with better accuracy, without incurring any substantial price in the regret or the overall linear oracle complexity.

It is also important to note that our algorithm is structured in a way that on each block m in the run of the algorithm, the point \mathbf{x}_{m-1} used for prediction, only takes into account the loss function revealed up to (and including) block $m-2$ (note \mathbf{x}_{m-1} is an approximate minimizer of $\hat{F}_{m-1}(\mathbf{x})$, which in turn depends only on the estimates $\hat{\mathbf{g}}_1, \dots, \hat{\mathbf{g}}_{m-2}$). Thus, in principle, Algorithm 1 does not have to wait after each block m until the new iterate \mathbf{x}_m is computed for the following block $m+1$ via Algorithm 2. While Algorithm 1 uses \mathbf{x}_{m-1} for prediction on block m , it can run Algorithm 2 in parallel, to simultaneously compute the next iterate \mathbf{x}_m (which is independent of the gradient estimates obtained in block m).

While this self-induced delay in information usage is not important for the theoretical complexity analysis, we believe it is of practical importance, since otherwise without this delay, Algorithm 1 would have to stop after each block and wait for Algorithm 2 to finish its computation, which can be potentially prohibitive in high-frequency prediction settings.

Finally, note that while the conditional gradient method is run over the shrunk set \mathcal{K}_δ , solving the linear optimization problem over \mathcal{K}_δ is identical, up to scaling, to solving it over the original set \mathcal{K} .

In the following, for any iteration (or block) m of the outer-loop in Algorithm 1, we denote by L_m the overall number of iterations performed by the do-while loop of

Algorithm 1: Block Bandit Conditional Gradient Method

Data: horizon T , feasible set \mathcal{K} with parameters r, R , block size K , step size η , smoothing parameter $\delta \in (0, r]$, tolerance parameter ϵ

Result: $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T$

$\mathbf{x}_0 \leftarrow$ arbitrary point in \mathcal{K}_δ , $\mathbf{x}_1 \leftarrow \mathbf{x}_0$

```

for  $m = 1, \dots, \frac{T}{K}$  do
    define  $\hat{F}_m(\mathbf{x}) := \eta \sum_{i=1}^{m-1} \mathbf{x}^\top \hat{\mathbf{g}}_i + \|\mathbf{x} - \mathbf{x}_1\|^2$ 
    if  $m > 1$  then
        run Algorithm 2 with set  $\mathcal{K}_\delta$ , tolerance  $\epsilon$ ,
        initial vector  $\mathbf{x}_{m-1}$ , and function  $\hat{F}_m(\mathbf{x})$ .
        Execute in parallel to the following for loop
        over  $s$ 
    end
    for  $s = 1, \dots, K$  do
         $\mathbf{u}_t \sim S^n$   $\triangleright t = (m-1)K + s$ 
        play  $\mathbf{y}_t \leftarrow \mathbf{x}_{m-1} + \delta \mathbf{u}_t$  and observe  $f_t(\mathbf{y}_t)$ 
         $\mathbf{g}_t \leftarrow \frac{n}{\delta} f_t(\mathbf{y}_t) \mathbf{u}_t$ 
    end
     $\hat{\mathbf{g}}_m \leftarrow \sum_{s=1}^K \mathbf{g}_{(m-1)K+s}$ 
    if  $m > 1$  then
         $\mathbf{x}_m \leftarrow$  output of Algorithm 2
    end
end
    
```

Algorithm 2: Conditional Gradient with Stopping Condition

Data: feasible set \mathcal{K}_δ , error tolerance ϵ , initial vector \mathbf{x}_{in} , objective function $\hat{F}_m(\mathbf{x})$

Result: \mathbf{x}_{out}

$\mathbf{z}_1 \leftarrow \mathbf{x}_{in}$, $\tau \leftarrow 0$

```

do
     $\tau \leftarrow \tau + 1$ 
     $\mathbf{v}_\tau \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}_\delta} \{\nabla \hat{F}_m(\mathbf{z}_\tau)^\top \cdot \mathbf{x}\}$ 
     $\sigma_\tau = \operatorname{argmin}_{\sigma \in [0,1]} \{\hat{F}_m(\mathbf{z}_\tau + \sigma(\mathbf{v}_\tau - \mathbf{z}_\tau))\}$   $\triangleright$  Line-search
     $\mathbf{z}_{\tau+1} = \mathbf{z}_\tau + \sigma_\tau(\mathbf{v}_\tau - \mathbf{z}_\tau)$   $\triangleright \mathbf{z}_{\tau+1} \in \mathcal{K}_\delta$ 
while  $\nabla \hat{F}_m(\mathbf{z}_\tau)^\top (\mathbf{z}_\tau - \mathbf{v}_\tau) > \epsilon$ ;
 $\mathbf{x}_{out} \leftarrow \mathbf{z}_\tau$ 
    
```

Algorithm 2, when invoked on iteration m . In particular, note that $\sum_{i=1}^{T/K} L_i$ is the overall number of calls to the linear optimization oracle of \mathcal{K} throughout the run of Algorithm 1.

Theorem 1 (Main theorem). For all $c > 0$ such that $\frac{cT^{-1/4}}{r} \leq 1$, setting $\eta = \frac{2cR}{nM} T^{-3/4}$, $\delta = cT^{-1/4}$, $\epsilon = 16R^2 T^{-1/2}$, $K = T^{1/2}$ in Algorithm 1, guarantees that the expected regret is upper-bounded by

$$\mathbb{E}[\mathcal{R}_T] \leq \left(3cG + \frac{cRG}{r} + 6GR + 4 \frac{cG^2 R}{nM} + 4 \frac{RnM}{c} \right) T^{3/4},$$

and that the expected overall number of calls to the

linear optimization oracle is upper-bounded by

$$\mathbb{E} \left[\sum_{m=1}^{\frac{T}{K}} L_m \right] \leq \left(\frac{3}{4} + \frac{Gc}{2nM} + \frac{G^2 c^2}{4n^2 M^2} \right) T.$$

In particular, if $\left(\frac{nM}{Gr} \right)^2 \leq T$ then, setting $c = \sqrt{\frac{nMr}{G}}$, we have

$$\mathbb{E}[\mathcal{R}_T] \leq \left(8 \frac{R}{r} \sqrt{nMrG} + 6GR + 4GR \sqrt{\frac{rG}{nM}} \right) T^{3/4},$$

and $\mathbb{E} \left[\sum_{m=1}^{\frac{T}{K}} L_m \right] \leq \left(\frac{3}{4} + \frac{1}{2} \sqrt{\frac{Gr}{nM}} + \frac{1}{4} \sqrt{\frac{Gr}{nM}} \right) T.$

3.1 Analysis

For the purpose of the analysis, we define the auxiliary sequence $\{\mathbf{x}_m^*\}_{m=1}^{T/K}$ as $\mathbf{x}_m^* = \arg \min_{\mathbf{x} \in \mathcal{K}_\delta} \hat{F}_m(\mathbf{x})$, where $\hat{F}_m(\cdot)$ is as defined in Algorithm 1. Note that this sequence corresponds to running the RFTL algorithm in blocks of length K , with respect to the feasible set \mathcal{K}_δ (see Section 2.3.1).

The following lemma, which is crucial to obtain our improved regret bound, shows that the squared norm of the gradient estimator over a block of size K , as a first approximation, grows only linearly with the block size K .

Lemma 5. For any iteration (block) m of the outer-loop in Algorithm 1 it holds that

$$\mathbb{E} [\|\hat{\mathbf{g}}_m\|^2] \leq \mathbb{E} [\|\mathbf{g}_m\|^2] \leq K \left(\frac{nM}{\delta} \right)^2 + K^2 G^2.$$

Proof. Fix some block m . For convenience, we denote $\mathcal{T}_m = \{(m-1)K + 1, \dots, mK\}$ (i.e., the set of all rounds included in block m). It holds that

$$\begin{aligned} \mathbb{E} [\|\hat{\mathbf{g}}_m\|^2] &= \mathbb{E} \left[\left\| \sum_{t \in \mathcal{T}_m} \mathbf{g}_t \right\|^2 \right] \\ &= \mathbb{E} \left[\sum_{t \in \mathcal{T}_m} \|\mathbf{g}_t\|^2 + \sum_{(i,j) \in \mathcal{T}_m^2, i \neq j} \mathbf{g}_i^\top \mathbf{g}_j \right] \\ &= \mathbb{E} \left[\sum_{t \in \mathcal{T}_m} \|\mathbf{g}_t\|^2 \right] + \sum_{(i,j) \in \mathcal{T}_m^2, i \neq j} \mathbb{E} [\mathbf{g}_i^\top \mathbf{g}_j]. \end{aligned}$$

Since, conditioned on the iterate \mathbf{x}_{m-1} , $\forall i \neq j$ $\mathbf{g}_i, \mathbf{g}_j$ are independent random vectors, we have

$$\begin{aligned} \mathbb{E} [\|\hat{\mathbf{g}}_m\|^2] &= \mathbb{E} \left[\sum_{t \in \mathcal{T}_m} \|\mathbf{g}_t\|^2 \right] \\ &\quad + \sum_{(i,j) \in \mathcal{T}_m^2, i \neq j} \mathbb{E} [\mathbb{E}[\mathbf{g}_i^\top | \mathbf{x}_{m-1}] \mathbb{E}[\mathbf{g}_j | \mathbf{x}_{m-1}]]. \end{aligned}$$

Using Lemma 4 we have that for all $t \in \mathcal{T}_m$, $\|\mathbb{E}[\mathbf{g}_t | \mathbf{x}_{m-1}]\| = \|\nabla \hat{f}_{t,\delta}(\mathbf{x}_{m-1})\| \leq G$. Since

$\max_{\mathbf{x} \in \mathcal{K}} \|f_t(\mathbf{x})\| \leq M$, we also have $\|\mathbf{g}_t\| \leq \frac{n}{\delta} \|f_t(\mathbf{y}_t)\| \|\mathbf{u}_t\| \leq \frac{nM}{\delta}$, and thus,

$$\begin{aligned} & \mathbb{E} \left[\sum_{t \in \mathcal{T}_m} \|\mathbf{g}_t\|^2 \right] + \sum_{(i,j) \in \mathcal{T}_m^2, i \neq j} \mathbb{E} \left[\mathbb{E}[\mathbf{g}_i^\top | \mathbf{x}_{m-1}] \mathbb{E}[\mathbf{g}_j | \mathbf{x}_{m-1}] \right] \\ & \leq K \left(\frac{nM}{\delta} \right)^2 + (K^2 - K) G^2 \leq K \left(\frac{nM}{\delta} \right)^2 + K^2 G^2. \end{aligned}$$

Finally, the inequality $\mathbb{E}[\|\hat{\mathbf{g}}_m\|] \leq \mathbb{E}[\|\hat{\mathbf{g}}_m\|^2]$ stated in the lemma follows from using Jensen's inequality. \square

The following lemma combines the RFTL regret bound with the unbiased gradient estimates of the smoothed loss functions, and upper-bounds the expected regret of Algorithm 1.

Lemma 6. Let $\epsilon > 0$. Suppose that throughout the run of Algorithm 1, for all blocks $m = 1, \dots, \frac{T}{K}$ it holds that $\hat{F}_m(\mathbf{x}_m) - \hat{F}_m(\mathbf{x}_m^*) \leq \epsilon$. Then, the expected regret of the algorithm is upper-bounded by

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T] & \leq \left(3\delta G + \delta R G / r + G\sqrt{\epsilon} + \frac{\eta G \sqrt{K} n M}{\delta} \right. \\ & \quad \left. + \eta \left(\frac{nM}{\delta} \right)^2 + 2\eta K G^2 \right) T + \frac{4R^2}{\eta}. \end{aligned}$$

Define $\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T f_t(\mathbf{x})$, $\tilde{\mathbf{x}}^* = (1 - \delta/r)\mathbf{x}^*$, $\forall t$: $m(t) := \lceil \frac{t}{K} \rceil$. Recall that throughout any block m , Algorithm 1 predicts according to \mathbf{x}_{m-1} .

Proof. It holds that

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T] & = \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{y}_t)] - \sum_{t=1}^T f_t(\mathbf{x}^*) \\ & = \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{y}_t)] - \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_{m(t)-1})] + \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_{m(t)-1})] \\ & \quad - \sum_{t=1}^T f_t(\tilde{\mathbf{x}}^*) + \sum_{t=1}^T f_t(\tilde{\mathbf{x}}^*) - \sum_{t=1}^T f_t(\mathbf{x}^*). \end{aligned} \quad (3)$$

f_t is G -Lipschitz, and thus we have that

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{y}_t)] - \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_{m(t)-1})] \\ & = \sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_{m(t)-1} + \delta \mathbf{u}_t) - f_t(\mathbf{x}_{m(t)-1})] \\ & \leq \sum_{t=1}^T \mathbb{E}[G \|\delta \mathbf{u}_t\|] \leq \delta G T. \end{aligned} \quad (4)$$

Also,

$$\begin{aligned} & \sum_{t=1}^T f_t(\tilde{\mathbf{x}}^*) - \sum_{t=1}^T f_t(\mathbf{x}^*) \leq \sum_{t=1}^T G \|\tilde{\mathbf{x}}^* - \mathbf{x}^*\| \\ & = \sum_{t=1}^T G \|(1 - \delta/r)\mathbf{x}^* - \mathbf{x}^*\| \leq \delta R G T / r. \end{aligned} \quad (5)$$

Now, we need to obtain an upper bound on $\sum_{t=1}^T \mathbb{E}[f_t(\mathbf{x}_{m(t)-1})] - \sum_{t=1}^T f_t(\tilde{\mathbf{x}}^*)$. We will first take a few preliminary steps. Define for all $m \in \lceil \frac{T}{K} \rceil$ $\mathcal{F}_m = \{\mathbf{x}_1, \hat{\mathbf{g}}_1, \dots, \mathbf{x}_{m-1}, \hat{\mathbf{g}}_{m-1}\}$ the history of all predictions and gradient estimates. Throughout the sequel we introduce the short notation $\hat{\nabla}_{t,\delta,m(t)-1} = \nabla \hat{f}_{t,\delta}(\mathbf{x}_{m(t)-1})$. Since \mathbf{g}_t is an unbiased estimator of $\nabla \hat{f}_{t,\delta}(\mathbf{x}_{m(t)-1}) = \hat{\nabla}_{t,\delta,m(t)-1}$, then $\mathbb{E}[\mathbf{g}_t | \mathcal{F}_m] = \hat{\nabla}_{t,\delta,m(t)-1}$. Since $\mathbf{x}_m^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}_\delta} \left\{ \hat{F}_m(\mathbf{x}) := \eta \sum_{i=1}^{m-1} \mathbf{x}^\top \hat{\mathbf{g}}_i + \|\mathbf{x} - \mathbf{x}_1\|^2 \right\}$, we have that $\mathbb{E}[\mathbf{x}_m^* | \mathcal{F}_m] = \mathbf{x}_m^*$. From both observations $\forall \mathbf{x} \in \mathcal{K}_\delta$ and $\forall m \in \lceil \frac{T}{K} \rceil$, it holds that

$$\begin{aligned} \mathbb{E}[\hat{\mathbf{g}}_m^\top (\mathbf{x}_m^* - \mathbf{x})] & = \mathbb{E}[\mathbb{E}[\hat{\mathbf{g}}_m | \mathcal{F}_m]^\top (\mathbf{x}_m^* - \mathbf{x})] \\ & = \mathbb{E} \left[\sum_{t=(m-1)K+1}^{mK} \mathbb{E}[\mathbf{g}_t | \mathcal{F}_{m(t)}]^\top (\mathbf{x}_{m(t)}^* - \mathbf{x}) \right] \\ & = \sum_{t=(m-1)K+1}^{mK} \mathbb{E}[\hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)}^* - \mathbf{x})]. \end{aligned}$$

Using Lemma 2 with the regularizer $\mathcal{R}(X) = \frac{\|\mathbf{x} - \mathbf{x}_1\|^2}{\eta}$ and $\{\hat{\mathbf{g}}_m^\top \mathbf{x}\}_{m=1}^{T/K}$ as the (linear) loss functions, we have that $\forall \mathbf{x} \in \mathcal{K}_\delta$,

$$\sum_{m=1}^{\frac{T}{K}} \hat{\mathbf{g}}_m^\top (\mathbf{x}_m^* - \mathbf{x}) \leq \sum_{m=1}^{\frac{T}{K}} \hat{\mathbf{g}}_m^\top (\mathbf{x}_m^* - \mathbf{x}_{m+1}^*) + \frac{1}{\eta} \|\mathbf{x} - \mathbf{x}_1\|^2.$$

Since for all m , $\hat{F}_m(\mathbf{x})$ is 2-strongly convex and $\hat{F}_m(\mathbf{x}_m^*) \leq \hat{F}_m(\mathbf{x}_{m+1}^*)$, using Eq. (2) we have that

$$\begin{aligned} \|\mathbf{x}_m^* - \mathbf{x}_{m+1}^*\|^2 & \leq \hat{F}_{m+1}(\mathbf{x}_m^*) - \hat{F}_{m+1}(\mathbf{x}_{m+1}^*) \\ & = \hat{F}_m(\mathbf{x}_m^*) - \hat{F}_m(\mathbf{x}_{m+1}^*) + \eta \hat{\mathbf{g}}_m^\top (\mathbf{x}_m^* - \mathbf{x}_{m+1}^*) \\ & \leq \eta \|\hat{\mathbf{g}}_m\| \|\mathbf{x}_m^* - \mathbf{x}_{m+1}^*\|. \end{aligned} \quad (6)$$

From the above inequality we obtain $\|\mathbf{x}_m^* - \mathbf{x}_{m+1}^*\| \leq \eta \|\hat{\mathbf{g}}_m\|$. From these three observations and the Cauchy-Schwarz inequality, we have

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}[\hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)}^* - \tilde{\mathbf{x}}^*)] = \mathbb{E} \left[\sum_{m=1}^{\frac{T}{K}} \hat{\mathbf{g}}_m^\top (\mathbf{x}_m^* - \tilde{\mathbf{x}}^*) \right] \\ & \leq \mathbb{E} \left[\sum_{m=1}^{\frac{T}{K}} \hat{\mathbf{g}}_m^\top (\mathbf{x}_m^* - \mathbf{x}_{m+1}^*) \right] + \frac{1}{\eta} \|\tilde{\mathbf{x}}^* - \mathbf{x}_1\|^2 \\ & \leq \eta \sum_{m=1}^{\frac{T}{K}} \mathbb{E}[\|\hat{\mathbf{g}}_m\|^2] + \frac{4R^2}{\eta} \\ & \stackrel{(a)}{\leq} \eta T \left(\frac{nM}{\delta} \right)^2 + \eta K T G^2 + \frac{4R^2}{\eta}. \end{aligned} \quad (7)$$

Inequality (a) is due to Lemma 5. Using Lemma 4 we have that for all $t \in [T]$, $\|\nabla \hat{f}_{t,\delta}(\mathbf{x}_t)\| \leq G$.

Also, using Eq. (2) w.r.t. \hat{F}_{m-1} and our assumption, $\hat{F}_{m-1}(\mathbf{x}_{m-1}) - \hat{F}_{m-1}(\mathbf{x}_{m-1}^*) \leq \epsilon$, we have that

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{E} \left[\hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)-1} - \mathbf{x}_{m(t)-1}^*) \right] \\
 & \leq G \sum_{t=1}^T \mathbb{E} \left[\|\mathbf{x}_{m(t)-1} - \mathbf{x}_{m(t)-1}^*\| \right] \\
 & \leq GK \sum_{m=1}^{\frac{T}{K}} \mathbb{E} \left[\sqrt{\hat{F}_{m-1}(\mathbf{x}_{m-1}) - \hat{F}_{m-1}(\mathbf{x}_{m-1}^*)} \right] \\
 & \leq GT\sqrt{\epsilon}. \tag{8}
 \end{aligned}$$

Using Eq. (6), we have

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{E} \left[\hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)-1}^* - \mathbf{x}_{m(t)}^*) \right] \\
 & \leq G \sum_{t=1}^T \mathbb{E} \left[\|\mathbf{x}_{m(t)-1}^* - \mathbf{x}_{m(t)}^*\| \right] \\
 & \leq GK\eta \sum_{m=1}^{\frac{T}{K}} \mathbb{E} [\|\hat{\mathbf{g}}_{m-1}\|] \\
 & \stackrel{(b)}{\leq} GT\eta \left(\frac{\sqrt{K}nM}{\delta} + KG \right). \tag{9}
 \end{aligned}$$

Inequality (b) is due to Lemma 5 and the fact that for all $a, b \in \mathbb{R}^+$ it holds that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$. One last step before we will achieve the target bound, we require an upper bound on the regret w.r.t. the δ -smoothed losses. Combining the results of Eq. (7), (8) and (9), using the convexity of $\hat{f}_{t,\delta}$, $\hat{f}_{t,\delta}(\mathbf{x}) - \hat{f}_{t,\delta}(\mathbf{y}) \leq \nabla \hat{f}_{t,\delta}^\top(\mathbf{x})(\mathbf{x} - \mathbf{y})$, we obtain

$$\begin{aligned}
 & \sum_{t=1}^T \left(\mathbb{E} \left[\hat{f}_{t,\delta}(\mathbf{x}_{m(t)-1}) \right] - \hat{f}_{t,\delta}(\tilde{\mathbf{x}}^*) \right) \\
 & \leq \sum_{t=1}^T \mathbb{E} \left[\hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)-1} - \tilde{\mathbf{x}}^*) \right] \\
 & = \sum_{t=1}^T \mathbb{E} \left[\hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)-1} - \mathbf{x}_{m(t)-1}^*) \right] \\
 & \quad + \hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)-1}^* - \mathbf{x}_{m(t)}^*) \\
 & \quad + \hat{\nabla}_{t,\delta,m(t)-1}^\top (\mathbf{x}_{m(t)}^* - \tilde{\mathbf{x}}^*) \\
 & \leq \left(G\sqrt{\epsilon} + \frac{\eta nM}{\delta} \left(G\sqrt{K} + \frac{nM}{\delta} \right) + 2\eta KG^2 \right) T \\
 & \quad + \frac{4R^2}{\eta}. \tag{10}
 \end{aligned}$$

Using Lemma 2 and the above equation, we have

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{E} [f_t(\mathbf{x}_{m(t)-1})] - \sum_{t=1}^T f_t(\tilde{\mathbf{x}}^*) = \\
 & \sum_{t=1}^T \mathbb{E} \left[f_t(\mathbf{x}_{m(t)-1}) - \hat{f}_{t,\delta}(\mathbf{x}_{m(t)-1}) \right] + \left(\hat{f}_{t,\delta}(\tilde{\mathbf{x}}^*) - f_t(\tilde{\mathbf{x}}^*) \right) \\
 & \quad + \sum_{t=1}^T \left(\mathbb{E} \left[\hat{f}_{t,\delta}(\mathbf{x}_{m(t)-1}) \right] - \hat{f}_{t,\delta}(\tilde{\mathbf{x}}^*) \right) \leq \frac{4R^2}{\eta} \\
 & \quad + \left(2\delta G + G\sqrt{\epsilon} + \frac{\eta nMG\sqrt{K}}{\delta} + \frac{\eta^2 M^2}{\delta^2} + 2\eta KG^2 \right) T. \tag{11}
 \end{aligned}$$

Combining the last equation with Eq. (4), (5) and Eq. (3), we obtain the required bound. \square

The following lemma is used to upper-bound the number of iterations required by the conditional gradient method, Algorithm 2, to terminate on each invocation.

Lemma 7. Let $\epsilon > 0$. Given a function $F(\mathbf{x})$, 2-smooth and 2-strongly convex, and $\mathbf{x}_1 \in \mathcal{K}_\delta$ such that $h_1 := F(\mathbf{x}_1) - F(\mathbf{x}^*)$, where $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}_\delta} F(\mathbf{x})$, Algorithm 2 produces a point $\mathbf{x}_{L+1} \in \mathcal{K}_\delta$ such that $F(\mathbf{x}_{L+1}) - F(\mathbf{x}^*) \leq \epsilon$ after at most $L = \max \left\{ \frac{16R^2}{\epsilon^2} (h_1 - \epsilon), \frac{2}{\epsilon} (h_1 - \epsilon) \right\}$ iterations.

Proof. For any iteration τ of Algorithm 2, define $h_\tau := F(\mathbf{x}_\tau) - F(\mathbf{x}^*)$ and denote $\nabla_\tau := \nabla F(\mathbf{x}_\tau)$. From the choice of \mathbf{v}_τ and the convexity of $F(\cdot)$, it follows that

$$\begin{aligned}
 \nabla_\tau^\top (\mathbf{x}_\tau - \mathbf{v}_\tau) & \geq \nabla_\tau^\top (\mathbf{x}_\tau - \mathbf{x}^*) \\
 & \geq F(\mathbf{x}_\tau) - F(\mathbf{x}^*) = h_\tau \tag{12}
 \end{aligned}$$

Now, we establish the convergence rate of Algorithm 2. It holds that

$$\begin{aligned}
 h_{\tau+1} & = F(\mathbf{x}_{\tau+1}) - F(\mathbf{x}^*) \\
 & = F(\mathbf{x}_\tau + \sigma_\tau(\mathbf{v}_\tau - \mathbf{x}_\tau)) - F(\mathbf{x}^*).
 \end{aligned}$$

For our analysis we define the step-size $\hat{\sigma}_\tau = \min \left\{ \frac{\nabla_\tau^\top (\mathbf{x}_\tau - \mathbf{v}_\tau)}{8R^2}, 1 \right\}$. Since σ_τ is chosen via line-search, we have that

$$\begin{aligned}
 h_{\tau+1} & = F(\mathbf{x}_\tau + \sigma_\tau(\mathbf{v}_\tau - \mathbf{x}_\tau)) - F(\mathbf{x}^*) \\
 & \leq F(\mathbf{x}_\tau + \hat{\sigma}_\tau(\mathbf{v}_\tau - \mathbf{x}_\tau)) - F(\mathbf{x}^*).
 \end{aligned}$$

Since $F(\mathbf{x})$ is 2-smooth it holds that

$$\begin{aligned}
 F(\mathbf{x}_\tau + \hat{\sigma}_\tau(\mathbf{v}_\tau - \mathbf{x}_\tau)) & \leq F(\mathbf{x}_\tau) + \hat{\sigma}_\tau \nabla_\tau^\top (\mathbf{v}_\tau - \mathbf{x}_\tau) \\
 & \quad + \hat{\sigma}_\tau^2 \|\mathbf{v}_\tau - \mathbf{x}_\tau\|^2,
 \end{aligned}$$

and we obtain

$$h_{\tau+1} \leq h_\tau + \hat{\sigma}_\tau^2 (2R)^2 - \hat{\sigma}_\tau \nabla_\tau^\top (\mathbf{x}_\tau - \mathbf{v}_\tau).$$

We now consider several cases.

Case 1: If $\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau) \leq \epsilon$ for some $\tau < L$, the algorithm will stop after less than L iterations. Moreover, from Eq. (12) we have $h_\tau \leq \epsilon$.

Case 2: Else, $\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau) \geq \epsilon$ for all $\tau < L$. We have two cases:

Case 2.1: If $\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau) \geq 8R^2$ then $\hat{\sigma}_\tau = 1$ and we have

$$\begin{aligned} h_{\tau+1} &\leq h_\tau + \hat{\sigma}_\tau^2(2R)^2 - \hat{\sigma}_\tau \nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau) \\ &\leq h_\tau - \frac{\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau)}{2}. \end{aligned}$$

Case 2.2: Else, $\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau) \leq 8R^2$, and then $\hat{\sigma}_\tau = \frac{\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau)}{8R^2}$, and we have

$$\begin{aligned} h_{\tau+1} &\leq h_\tau + \hat{\sigma}_\tau^2(2R)^2 - \hat{\sigma}_\tau \nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau) \\ &\leq h_\tau - \left(\frac{\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau)}{4R} \right)^2. \end{aligned}$$

From both cases, we have

$$\begin{aligned} h_{\tau+1} &\leq h_\tau - \min \left\{ \left(\frac{\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau)}{4R} \right)^2, \frac{\nabla_\tau^\top(\mathbf{x}_\tau - \mathbf{v}_\tau)}{2} \right\} \\ &\leq h_1 - \tau \min_{i=1, \dots, \tau} \left\{ \left(\frac{\nabla_i^\top(\mathbf{x}_i - \mathbf{v}_i)}{4R} \right)^2, \frac{\nabla_i^\top(\mathbf{x}_i - \mathbf{v}_i)}{2} \right\} \\ &\leq h_1 - \tau \min \left\{ \left(\frac{\epsilon}{4R} \right)^2, \frac{\epsilon}{2} \right\}. \end{aligned} \quad (13)$$

Thus, for all cases, after a maximum of L iterations, when

$$L = \max \left\{ \frac{16R^2}{\epsilon^2}(h_1 - \epsilon), \frac{2}{\epsilon}(h_1 - \epsilon) \right\},$$

we obtain $h_{L+1} \leq \epsilon$. \square

We can now finally prove our main theorem, Theorem 1.

Proof of Theorem 1. We first upper bound the expected overall number of calls to the linear optimization oracle throughout the run of the algorithm, and then we upper-bound the expected regret.

Let $\mathbf{z}_{m,\tau}$ be the iterate of Algorithm 2 after completing $\tau - 1$ iterations of the do-while loop, when invoked on iteration (block) m of Algorithm 1. Also, for all m, τ , define $h_{m,\tau} := \hat{F}_m(\mathbf{z}_{m,\tau}) - \hat{F}_m(\mathbf{x}_m^*)$. Recall that for any iteration m of Algorithm 1, we have $\mathbf{z}_{m,1} = \mathbf{x}_{m-1}$.

Using the triangle inequality and the fact $\hat{F}_m(\mathbf{x}_{m+1}^*) \geq \hat{F}_m(\mathbf{x}_m^*)$, we have

$$\begin{aligned} \mathbb{E}[h_{m+1,1}] &= \mathbb{E}[\hat{F}_{m+1}(\mathbf{z}_{m+1,1}) - \hat{F}_{m+1}(\mathbf{x}_{m+1}^*)] \\ &\leq \mathbb{E}[\hat{F}_m(\mathbf{x}_m) - \hat{F}_m(\mathbf{x}_m^*) + \eta \|\hat{\mathbf{g}}_m\| \|\mathbf{x}_m - \mathbf{x}_{m+1}^*\|]. \end{aligned}$$

Since $h_{m,L_m} = \hat{F}_m(\mathbf{x}_m) - \hat{F}_m(\mathbf{x}_m^*) \leq \epsilon$, using the triangle inequality, we have

$$\begin{aligned} \mathbb{E}[h_{m+1,1}] &\leq \epsilon + \eta \mathbb{E}[\|\hat{\mathbf{g}}_m\| \|\mathbf{x}_m - \mathbf{x}_m^* + \mathbf{x}_m^* - \mathbf{x}_{m+1}^*\|] \\ &\leq \epsilon + \eta \mathbb{E}[\|\hat{\mathbf{g}}_m\| (\|\mathbf{x}_m - \mathbf{x}_m^*\| + \|\mathbf{x}_m^* - \mathbf{x}_{m+1}^*\|)]. \end{aligned}$$

Since $\hat{F}_m(\mathbf{x})$ is 2-strongly convex and $h_{m,L_m} \leq \epsilon$, using Eq. (2), we have that $\|\mathbf{x}_m - \mathbf{x}_m^*\| \leq \sqrt{\epsilon}$. Also, from Eq. (6), we have $\|\mathbf{x}_m^* - \mathbf{x}_{m+1}^*\| \leq \eta \|\hat{\mathbf{g}}_m\|$. Thus, we have

$$\mathbb{E}[h_{m+1,1}] \leq \epsilon + \eta\sqrt{\epsilon} \mathbb{E}[\|\hat{\mathbf{g}}_m\|] + \eta^2 \mathbb{E}[\|\hat{\mathbf{g}}_m\|^2].$$

Using Lemma 5 and the fact that for all $a, b \in \mathbb{R}^+$ it holds that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, we have

$$\begin{aligned} \mathbb{E}[h_{m+1,1}] &\leq \epsilon + \eta\sqrt{\epsilon} \left(\sqrt{K} \left(\frac{nM}{\delta} \right) + KG \right) \\ &\quad + \eta^2 \left(K \left(\frac{nM}{\delta} \right)^2 + K^2 G^2 \right). \end{aligned} \quad (14)$$

Using Lemma 7 with $h_1 = h_{m,1} \leq \epsilon + \eta\sqrt{\epsilon} \left(\sqrt{K} \left(\frac{nM}{\delta} \right) + KG \right) + \eta^2 \left(K \left(\frac{nM}{\delta} \right)^2 + K^2 G^2 \right)$ for $m = 1, \dots, \frac{T}{K}$, we have that on each iteration (block) m , the number of calls to the linear optimization oracle is $L_m \leq \max \left\{ \frac{16R^2}{\epsilon^2}(h_{m,1} - \epsilon), \frac{2}{\epsilon}(h_{m,1} - \epsilon) \right\}$. Plugging in ϵ , we have $L_m \leq \frac{16R^2}{\epsilon^2}(h_{m,1} - \epsilon)$. Following Eq. (14) we have

$$\begin{aligned} \mathbb{E}[L_m] &\leq \frac{16R^2}{\epsilon^2}(\mathbb{E}[h_{m,1}] - \epsilon) \\ &\leq \frac{16R^2}{\epsilon^2} \eta\sqrt{\epsilon} \left(\sqrt{K} \left(\frac{nM}{\delta} \right) + KG \right) \\ &\quad + \frac{16R^2}{\epsilon^2} \eta^2 \left(K \left(\frac{nM}{\delta} \right)^2 + K^2 G^2 \right) \\ &\stackrel{(c)}{=} \left(\frac{3}{4} + \frac{Gc}{2nM} + \frac{G^2 c^2}{4n^2 M^2} \right) T^{\frac{1}{2}}. \end{aligned}$$

Equality (c) is due to plugging-in $\eta, \delta, \epsilon, K$. Thus, overall on all blocks, we obtain

$$\mathbb{E} \left[\sum_{m=1}^{\frac{T}{K}} L_m \right] \leq \left(\frac{3}{4} + \frac{Gc}{2nM} + \frac{G^2 c^2}{4n^2 M^2} \right) T.$$

We now turn to upper-bound the expected regret of the algorithm. Using Lemma 6 we have that

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T] &\leq \left(3\delta G + \delta R G / r + G\sqrt{\epsilon} + \frac{\eta G \sqrt{K} n M}{\delta} \right. \\ &\quad \left. + \eta \left(\frac{nM}{\delta} \right)^2 + 2\eta K G^2 \right) T + \frac{4R^2}{\eta} \\ &\stackrel{(d)}{=} \left(3cG + \frac{cRG}{r} + 6GR + 4 \frac{cG^2 R}{nM} + 4 \frac{RnM}{c} \right) T^{\frac{3}{4}}. \end{aligned}$$

Equality (d) is due to plugging-in $\eta, \delta, \epsilon, K$. \square

4 Acknowledgments

This research was supported by the ISRAEL SCIENCE FOUNDATION (grant No. 1108/18).

References

- Abernethy, J., Hazan, E. E., and Rakhlin, A. (2008). Competing in the dark: An efficient algorithm for bandit linear optimization. In *21st Annual Conference on Learning Theory, COLT 2008*, pages 263–273.
- Allen-Zhu, Z., Hazan, E., Hu, W., and Li, Y. (2017). Linear convergence of a frank-wolfe type algorithm over trace-norm balls. In *Advances in Neural Information Processing Systems*, pages 6191–6200.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77.
- Bertsekas, D. P. (1973). Stochastic optimization problems with nondifferentiable cost functionals. *Journal of Optimization Theory and Applications*, 12(2):218–231.
- Bubeck, S., Cesa-Bianchi, N., et al. (2012). Regret analysis of stochastic and nonstochastic multiarmed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122.
- Bubeck, S., Dekel, O., Koren, T., and Peres, Y. (2015). Bandit convex optimization: \sqrt{T} regret in one dimension. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 266–278.
- Bubeck, S. and Eldan, R. (2016). Multi-scale exploration of convex functions and bandit convex optimization. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, pages 583–589.
- Bubeck, S., Lee, Y. T., and Eldan, R. (2017). Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 72–85.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- Chen, L., Zhang, M., and Karbasi, A. (2019). Projection-free bandit convex optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, pages 2047–2056.
- Dani, V., Kakade, S. M., and Hayes, T. P. (2008). The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, pages 345–352.
- Dekel, O., Eldan, R., and Koren, T. (2015). Bandit smooth convex optimization: Improving the bias-variance tradeoff. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2926–2934.
- Dudík, M., Harchaoui, Z., and Mallick, J. (2012). Lifted coordinate descent for learning with trace-norm regularization. *Journal of Machine Learning Research - Proceedings Track*, 22:327–336.
- Flaxman, A. D., Kalai, A. T., Kalai, A. T., and McMahan, H. B. (2005). Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 385–394. Society for Industrial and Applied Mathematics.
- Frank, M. and Wolfe, P. (1956). An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3:149–154.
- Garber, D. (2016). Faster projection-free convex optimization over the spectrahedron. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 874–882.
- Garber, D. and Hazan, E. (2015). Faster rates for the frank-wolfe method over strongly-convex sets. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 541–549.
- Garber, D. and Hazan, E. (2016). A linearly convergent variant of the conditional gradient algorithm under strong convexity, with applications to online and stochastic optimization. *SIAM Journal on Optimization*, 26(3):1493–1528.
- Garber, D. and Meshi, O. (2016). Linear-memory and decomposition-invariant linearly convergent conditional gradient algorithm for structured polytopes. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1001–1009.
- Harchaoui, Z., Douze, M., Paulin, M., Dudík, M., and Mallick, J. (2012). Large-scale image classification with trace-norm regularization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*.
- Hazan, E. (2016). Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325.
- Hazan, E. and Karnin, Z. (2016). Volumetric spanners: an efficient exploration basis for learning. *The*

- Journal of Machine Learning Research*, 17(1):4062–4095.
- Hazan, E. and Levy, K. (2014). Bandit convex optimization: Towards tight bounds. In *Advances in Neural Information Processing Systems*, pages 784–792.
- Hazan, E. and Li, Y. (2016). An optimal algorithm for bandit convex optimization. *arXiv preprint arXiv:1603.04350*.
- Hazan, E. E. and Kale, S. (2012). Projection-free online learning. In *29th International Conference on Machine Learning, ICML 2012*, pages 521–528.
- Jaggi, M. (2013). Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML (1)*, pages 427–435.
- Jaggi, M. and Sulovský, M. (2010). A simple algorithm for nuclear norm regularized problems. In *Proceedings of the 27th International Conference on Machine Learning, ICML*.
- Lacoste-Julien, S. and Jaggi, M. (2015). On the global linear convergence of Frank-Wolfe optimization variants. In *Advances in Neural Information Processing Systems*, pages 496–504.
- Lacoste-Julien, S., Jaggi, M., Schmidt, M. W., and Pletscher, P. (2013). Block-coordinate frank-wolfe optimization for structural svms. In *Proceedings of the 30th International Conference on Machine Learning, ICML*.
- Laue, S. (2012). A hybrid algorithm for convex semidefinite optimization. In *Proceedings of the 29th International Conference on Machine Learning, ICML*.
- Levitin, E. S. and Polyak, B. T. (1966). Constrained minimization methods. *USSR Computational mathematics and mathematical physics*, 6:1–50.
- Levy, K. and Krause, A. (2019). Projection free online learning over smooth sets. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1458–1466.
- Saha, A. and Tewari, A. (2011). Improved regret guarantees for online smooth convex optimization with bandit feedback. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 636–642.
- Shalev-Shwartz, S. et al. (2012). Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194.
- Shalev-Shwartz, S., Gonen, A., and Shamir, O. (2011). Large-scale convex minimization with a low-rank constraint. In *Proceedings of the 28th International Conference on Machine Learning, ICML*.
- Yang, S. and Mohri, M. (2016). Optimistic bandit convex optimization. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 2297–2305.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 928–936.