



Online Non-Monotone DR-Submodular Maximization

Kim Thang Nguyen, Abhinav Srivastav

► To cite this version:

Kim Thang Nguyen, Abhinav Srivastav. Online Non-Monotone DR-Submodular Maximization. 35th AAAI Conference on Artificial Intelligence (AAAI-2021), Feb 2021, Vancouver (virtual), Canada. hal-03170034

HAL Id: hal-03170034

<https://hal.science/hal-03170034>

Submitted on 15 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Online Non-Monotone DR-Submodular Maximization*

Nguyễn Kim Thắng

Abhinav Srivastav

IBISC, Univ. Evry, University Paris-Saclay, France

Abstract

In this paper, we study fundamental problems of maximizing DR-submodular continuous functions that have real-world applications in the domain of machine learning, economics, operations research and communication systems. It captures a subclass of non-convex optimization that provides both theoretical and practical guarantees. Here, we focus on minimizing regret for online arriving non-monotone DR-submodular functions over down-closed and general convex sets.

First, we present an online algorithm that achieves a $1/e$ -approximation ratio with the regret of $O(T^{3/4})$ for maximizing DR-submodular functions over any down-closed convex set. Note that, the approximation ratio of $1/e$ matches the best-known offline guarantee. Next, we give an online algorithm that achieves an approximation guarantee (depending on the search space) for the problem of maximizing non-monotone continuous DR-submodular functions over a *general* convex set (not necessarily down-closed). Finally we run experiments to verify the performance of our algorithms on problems arising in machine learning domain with the real-world datasets.

1 Introduction

Continuous DR-submodular optimization is a subclass of non-convex optimization that is an upcoming frontier in machine learning. Roughly speaking, a differentiable non-negative bounded function $F : [0, 1]^n \rightarrow [0, 1]$ is *DR-submodular* if $\nabla F(\mathbf{x}) \geq \nabla F(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in [0, 1]^n$ where $x_i \leq y_i$ for every $1 \leq i \leq n$. (Note that, w.l.o.g. after normalization, assume that F has values in $[0, 1]$.) Intuitively, continuous DR-submodular functions represent the diminishing returns property or the economy of scale in continuous domains. DR-submodularity has been of great interest [3, 1, 10, 22, 31, 35]. Many problems arising in machine learning and statistics, such as Non-definite Quadratic Programming [24], Determinantal Point Processes [26], log-submodular models [12], to name a few, have been modelled using the notion of continuous DR-submodular functions.

In the past decade, online computational framework has been quite successful for tackling a wide variety of challenging problems and capturing many real-world problems with uncertainty. In this computational framework, we focus on the model of online learning. In online learning, at any time step, given a history of actions and a set of associated reward functions, online algorithm first chooses an action from a set of feasible actions; then, an adversary subsequently selects a reward function. The objective is to perform as good as the best fixed action in hindsight. This setting have been extensively explored in the literature, especially in context of convex functions [23].

In fact, several algorithms with theoretical approximation guarantees are known for maximizing (offline and online) DR-submodular functions. However, these guarantees hold under the assumptions that are

*This work is supported by the ANR project OATA no ANR-15-CE40-0015-01

based on the monotonicity of functions coupled with the structure of convex sets such as unconstrained hypercube and down-closed. Though, a majority of real-world problems can be formulated as *non-monotone* DR-submodular functions over convex sets that might not be necessarily down-closed. For example, the problems of Determinantal Point Processes, log-submodular models, etc can be viewed as non-monotone DR-submodular maximization problems. Although several algorithms for non-monotone (set) submodular maximization are known and they are built on maximizing the corresponding multilinear extensions (a particular class of DR-submodular functions), they do not extend to general DR-submodular functions since the optimal solutions of the former are integer points while the ones of the latter are not. Besides, general convex sets include conic convex sets, up-closed convex sets, mixture of covering and packing linear constraints, etc. which appear in many applications. Among others, conic convex sets play a crucial role in convex optimization. Conic programming [4] — an important subfield of convex optimization — consists of optimizing an objective function over a conic convex set. Conic programming reduces to linear programming and semi-definite programming when the objective function is linear and the convex cones are the positive orthant \mathbb{R}_+^n and positive semidefinite matrices \mathbb{S}_+^n , respectively. Optimizing non-monotone DR-submodular functions over a (bounded) conic convex set (not necessarily downward-closed) is an interesting and important problem both in theory and in practice. To best of our knowledge, no prior work has been done on maximizing DR-submodular functions over conic sets in online setting. The limit of current theory [3, 1, 10, 22, 31, 38] motivates us to develop online algorithms for non-monotone functions.

In this work, we explore the *online* problem of maximizing *non-monotone* DR-submodular functions over a hypercube and down-closed¹ and over a *general* convex sets. Formally, we consider the following setting for DR-submodular maximization: given a convex domain $\mathcal{K} \subseteq [0, 1]^n$ in advance, at each time step $t = 1, 2, \dots, T$, the online algorithm first selects a vector $\mathbf{x}^t \in \mathcal{K}$. Subsequently, the adversary reveals a non-monotone continuous DR-submodular function F^t and the algorithm receives a reward of $F^t(\mathbf{x}^t)$. The objective is also to maximize the total reward. We say that an algorithm achieves a $(r, R(T))$ -*regret* if

$$\sum_{t=1}^T F^t(\mathbf{x}^t) \geq r \cdot \max_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}) - R(T)$$

In other words, r is the approximation ratio that measures the quality of the online algorithm compared to the best fixed solution in hindsight and $R(T)$ represents the regret in the classic terms. Equivalently, we say that the algorithm has r -*regret* at most $R(T)$. Our goal is to design online algorithms with $(r, R(T))$ -regret where $0 < r \leq 1$ is as large as possible, and $R(T)$ is sub-linear in T , i.e., $R(T) = o(T)$.

1.1 Our contributions and techniques

In this paper, we provide algorithms with performance guarantees for the problem over down-closed and general convex sets. Our contributions and techniques are summarized as follows (also see Table 1, the entries in the red correspond to our contribution).

1.1.1 Maximizing online non-monotone DR-submodular functions over down-closed convex sets

The natures of online DR-submodular maximization for monotone and non-monotone functions are quite different. In the setting of monotone functions, the best approximation guarantee is $(1 - 1/e)$ [10] but one can prove that the natural gradient descent algorithm attains 1/2-approximation [10]. However, for online non-monotone DR-submodular maximization, no constant guarantee is known and it is likely that

¹ \mathcal{K} is *down-closed* if for every $\mathbf{z} \in \mathcal{K}$ and $\mathbf{y} \leq \mathbf{z}$ then $\mathbf{y} \in \mathcal{K}$

		Monotone	Non-monotone
Offline	hypercube		$\frac{1}{2}$ -approx [3] [31]
	down-closed	$(1 - \frac{1}{e})$ -approx [2]	$(\frac{1}{e})$ -approx [1]
	general	$(1 - \frac{1}{e})$ -approx [29]	$(\frac{1 - \min_{\mathbf{x} \in \mathcal{K}} \ \mathbf{x}\ _\infty}{3\sqrt{3}})$ -approx [14]
Online	hypercube		$(1/2, O(\sqrt{T}))$
	down-closed		$(\frac{1}{e}, O(T^{3/4}))$
	general	$(1 - \frac{1}{e}, O(\sqrt{T}))$ [10]	$(\frac{1 - \min_{\mathbf{x} \in \mathcal{K}} \ \mathbf{x}\ _\infty}{3\sqrt{3}}, O(\frac{T}{\ln T}))$

Table 1: Summary of results on DR-submodular maximization. The entries represent the best-known $(r, R(T))$ -regret; our results are shown in red. The entry in blue is not explicitly stated in literature but can be deduced from [17] or from our technique. The guarantees in blank cells can be deduced from a more general one.

natural algorithms such as gradient descent, mirror descent, etc do not guarantee to achieve any constant approximation, even small one. An illustrative example of poor local optima (for coordinate ascent algorithm) can be found in [3, Appendix A].

The main result of our paper is the *first* online algorithm that achieves a constant approximation ratio with the regret of $O(T^{3/4})$ for maximizing non-monotone DR-submodular functions over any down-closed convex set where T is number of time steps. Moreover, the approximation is $1/e$ which matches to the best guarantee in the offline setting.

Our algorithm is built on the Meta-Frank-Wolfe algorithm introduced by Chen et al. [10] for *monotone* DR-submodular maximization. Their meta Frank-Wolfe algorithm combines the framework of meta-actions proposed in [36] with a variant of the Frank-Wolfe proposed in [2] for maximizing monotone DR-submodular functions. Informally, at every time t , our algorithm starts from the origin $\mathbf{0}$ ($\mathbf{0} \in \mathcal{K}$ since \mathcal{K} is down-closed) and executes L steps of the Frank-Wolfe algorithm where every update vector at iteration $1 \leq \ell \leq L$ is constructed by combining the output of an optimization oracle \mathcal{E}_ℓ and the current vector so that we can exploit the *concavity property in positive directions* of DR-submodular functions. The solution \mathbf{x}^t is produced at the end of the L -th step. Subsequently, after observing the function F^t , the algorithm subtly defines a vector \mathbf{d}_ℓ^t and feedbacks $\langle \cdot, \mathbf{d}_\ell^t \rangle$ as the reward function at time t to the oracle \mathcal{E}_ℓ for $1 \leq \ell \leq L$. The most important distinguishing point of our algorithm compared to that in [10] relies on the oracles. In their algorithm, they used linear oracles which are appropriate for maximizing monotone DR-submodular functions. However, it is not clear whether *linear or even convex oracles* are strong enough to achieve a constant approximation for non-monotone functions.

While aiming for a $1/e$ -approximation — the best known approximation in the offline setting — we consider the following online non-convex problem, referred throughout this paper as *online vee learning* problem. At each time t , the online algorithm knows $\mathbf{c}^t \in \mathcal{K}$ in advance and selects a vector $\mathbf{x}^t \in \mathcal{K}$. Subsequently, the adversary reveals a vector $\mathbf{a}^t \in \mathbb{R}^n$ and the algorithm receives a reward $\langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle$.

Given two vectors \mathbf{x} and \mathbf{y} , the vector $\mathbf{x} \vee \mathbf{y}$ has the i^{th} coordinate $(\mathbf{x} \vee \mathbf{y})_i = \max\{x(i), y(i)\}$. The goal is to maximize the total reward over a time horizon T .

In order to bypass the non-convexity of the online vee learning problem, we propose the following methodology. Unlike dimension reduction techniques that aim at reducing the dimension of the search space, we *lift* the search space and the functions to a higher dimension so as to benefit from their properties in this new space. Concretely, at a high level, given a convex set \mathcal{K} , we define a “sufficiently dense” lattice \mathcal{L} in $[0, 1]^n$ such that every point in \mathcal{K} can be approximated by a point in the lattice. The term “sufficiently dense” corresponds to the fact that the values of any Lipschitz function can be approximated by the function values at lattice points. As the next step, we lift all lattice points in \mathcal{K} to a high dimension space so that they form a subset of vertices (corners points) $\tilde{\mathcal{K}}$ of the unit hypercube in the newly defined space. Interestingly, the reward function $\langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle$ can be transformed to a linear function in the new space. Hence, our algorithm for the online vee problem consists of updating at every time step a solution in the high-dimension space using a linear oracle and projecting the solution back to the original space.

Once the solution is projected back to original space, we construct appropriate update vectors for the original DR-submodular maximization and feedback rewards for the online vee oracle. Exploiting the underlying properties of DR-submodularity, we show that our algorithm achieves the regret bound of $(1/e, O(T^{3/4}))$. A useful feature of our algorithm is that it is projection-free if using appropriate projection-free oracles.

Over the hypercube. Restricting on the DR-submodular maximization problem over the hypercube $[0, 1]^n$, a particular down-closed convex set that has been widely studied, our approach leads to an $1/2$ -approximation algorithm with regret bound of $O(\sqrt{T} \log T)$. Note that this result can be deduced from [17] but as it is not explicitly stated in literature, for completeness, we present the proof in the appendix.

1.1.2 Maximizing online non-monotone DR-submodular functions over general convex sets

We go beyond the down-closed structure by considering general convex sets. Building upon our salient ideas from previous algorithm, we prove that for general convex sets, the Meta-Frank-Wolfe algorithm with adapted step-sizes guarantees $\left(\frac{1 - \min_{\mathbf{x} \in \mathcal{K}} \|\mathbf{x}\|_\infty}{3\sqrt{3}}, O\left(\frac{T}{\ln T}\right)\right)$ -regret in expectation. Notably, if the set \mathcal{K} contains $\mathbf{0}$ (for example, \mathcal{K} is the intersection of the semi-definite positive matrix cone and the hypercube $[0, 1]^n$) then the algorithm guarantees in expectation a $\left(\frac{1}{3\sqrt{3}}, O\left(\frac{T}{\log T}\right)\right)$ -regret. To the best of our knowledge, this is the *first* constant-approximation algorithm for the problem of online non-monotone DR-submodular maximization over a non-trivial convex set that goes beyond down-closed sets. Note that, any algorithm for the problem over a non-down-closed convex set (in particular arbitrary convex set containing the origin) that guarantees a constant approximation must require in general exponentially many value queries to the function [37]. Remark that the quality of the solution, specifically the approximation ratio, depends on the initial solution \mathbf{x}_0 . This confirms an observation in various contexts that initialization plays an important role in non-convex optimization.

1.2 Related Work

In this section, we give a summary on best-known results on DR-submodular maximization. The domain has been investigated more extensively in recent years due to its numerous applications in the field of statistics and machine learning, for example active learning [19], viral marketing [25], network monitoring [20], document

summarization [27], crowd teaching [33], feature selection [16], deep neural networks [15], diversity models [13] and recommender systems [21].

Offline setting. Bian et al. [2] considered the problem of maximizing monotone DR-functions subject to a down-closed convex set and showed that the greedy method proposed by [7], a variant of well-known Frank-Wolfe algorithm in convex optimization, guarantees a $(1 - 1/e)$ -approximation. However, it has been observed by [22] that the greedy method is not robust in stochastic settings (where only unbiased estimates of gradients are available). Subsequently, Hassani et al. [29] proposed a $(1 - 1/e)$ -approximation algorithm for maximizing monotone DR-submodular functions over general convex sets in stochastic settings by a new variance reduction technique.

The problem of maximizing non-monotone DR-submodular functions is much harder. Bian et al. [3] and Niazadeh et al. [31] have independently presented algorithms with the same approximation guarantee of $(1/2)$ for the problem of non-monotone DR-submodular maximization over the hypercube ($\mathcal{K} = [0, 1]^n$). These algorithms are inspired by the bi-greedy algorithms in [6, 5]. Bian et al. [1] made a further step by providing a $(1/e)$ -approximation algorithm where the convex sets are down-closed. Mokhtari et al. [29] also presented an algorithm that achieve $(1/e)$ for non-monotone continuous DR-submodular function over a down-closed convex domain that uses only stochastic gradient estimates. Remark that when aiming for approximation algorithms (in polynomial time), the restriction to down-closed polytopes is unavoidable. Specifically, Vondrak [37] proved that any algorithm for the problem over a non-down-closed set that guarantees a constant approximation must require in general exponentially many value queries to the function. Recently, Durr et al. [14] presented an algorithm that achieves $((1 - \min_{\mathbf{x} \in \mathcal{K}} \|\mathbf{x}\|_\infty)/3\sqrt{3})$ -approximation for non-monotone DR-submodular function over general convex sets that are not necessarily down-closed.

Online setting. The results for online DR-submodular maximization are known *only* for monotone functions. Chen et al. [10] considered the online problem of maximizing monotone DR-submodular functions over general convex sets and provided an algorithm that achieves a $(1 - 1/e, O(\sqrt{T}))$ -regret. Subsequently, Zhang et al. [38] presented an algorithm that reduces the number of per-function gradient evaluations from $T^{3/2}$ in [10] to 1 and achieves the same approximation ratio of $(1 - 1/e)$. Leveraging the idea of one gradient per iteration, they presented a bandit algorithm for maximizing monotone DR-submodular function over a general convex set that achieves in expectation $(1 - 1/e)$ approximation ratio with regret $O(T^{8/9})$. Note that in the discrete setting, Roughgarden and Wang [32] studied the non-monotone (discrete) submodular maximization over the hypercube and gave an algorithm which guarantees the tight $(1/2)$ -approximation and $O(\sqrt{T})$ regret. Recently and independently, Niazadeh et al. [30] have provided an $(1/2, O(\sqrt{T}))$ -regret algorithm for DR-submodular maximization over the hypercube. Their algorithm is built on an interesting framework that transform offline greedy algorithms to online ones using Blackwell approachability. Our approach, based on the lifting procedure, is completely different to theirs.

2 Preliminaries and Notations

Throughout the paper, we use bold face letters, e.g., \mathbf{x}, \mathbf{y} to represent vectors. For every $S \subseteq \{1, 2, \dots, n\}$, vector $\mathbf{1}_S$ is the n -dim vector whose i^{th} coordinate equals to 1 if $i \in S$ and 0 otherwise. Given two n -dimensional vectors \mathbf{x}, \mathbf{y} , we say that $\mathbf{x} \leq \mathbf{y}$ iff $x_i \leq y_i$ for all $1 \leq i \leq n$. Additionally, we denote by $\mathbf{x} \vee \mathbf{y}$, their coordinate-wise maximum vector such that $(\mathbf{x} \vee \mathbf{y})_i = \max\{x_i, y_i\}$. Moreover, the symbol \circ represents the element-wise multiplication that is, given two vectors \mathbf{x}, \mathbf{y} , vector $\mathbf{x} \circ \mathbf{y}$ is the such that the i -th coordinate $(\mathbf{x} \circ \mathbf{y})_i = x_i y_i$. The scalar product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_i x_i y_i$ and the norm $\|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{x} \rangle^{1/2}$. In the

paper, we assume that $\mathcal{K} \subseteq [0, 1]^n$. We say that \mathcal{K} is the *hypercube* if $\mathcal{K} = [0, 1]^n$; \mathcal{K} is *down-closed* if for every $\mathbf{z} \in \mathcal{K}$ and $\mathbf{y} \leq \mathbf{z}$ then $\mathbf{y} \in \mathcal{K}$; and \mathcal{K} is *general* if \mathcal{K} is a convex subset of $[0, 1]^n$ without any special property.

Definition 1. A function $F : [0, 1]^n \rightarrow \mathbb{R}^+ \cup \{0\}$ is *diminishing-return (DR) submodular* if for all vector $\mathbf{x} \geq \mathbf{y} \in [0, 1]^n$, any basis vector $\mathbf{e}_i = (0, \dots, 0, 1, 0, \dots, 0)$ and any constant $\alpha > 0$ such that $\mathbf{x} + \alpha \mathbf{e}_i \in [0, 1]^n$, $\mathbf{y} + \alpha \mathbf{e}_i \in [0, 1]^n$, it holds that

$$F(\mathbf{x} + \alpha \mathbf{e}_i) - F(\mathbf{x}) \leq F(\mathbf{y} + \alpha \mathbf{e}_i) - F(\mathbf{y}). \quad (1)$$

Note that if function F is differentiable then the diminishing-return (DR) property (1) is equivalent to

$$\nabla F(\mathbf{x}) \leq \nabla F(\mathbf{y}) \quad \forall \mathbf{x} \geq \mathbf{y} \in [0, 1]^n. \quad (2)$$

Moreover, if F is twice-differentiable then the DR property is equivalent to all of the entries of its Hessian being non-positive, i.e., $\frac{\partial^2 F}{\partial x_i \partial x_j} \leq 0$ for all $1 \leq i, j \leq n$.

Besides, a differentiable function $F : \mathcal{K} \subseteq [0, 1]^n \rightarrow \mathbb{R}^+ \cup \{0\}$ is said to be β -smooth if for any $\mathbf{x}, \mathbf{y} \in \mathcal{K}$, the following holds,

$$F(\mathbf{y}) \leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\beta}{2} \|\mathbf{y} - \mathbf{x}\|^2 \quad (3)$$

or equivalently, the gradient is β -Lipschitz, i.e.,

$$\|\nabla F(\mathbf{x}) - \nabla F(\mathbf{y})\| \leq \beta \|\mathbf{x} - \mathbf{y}\|. \quad (4)$$

Properties of DR-submodularity In the following, we present a property of DR-submodular functions that are crucial in our analyses. The most important property is the concavity in positive directions, i.e., for every $\mathbf{x}, \mathbf{y} \in \mathcal{K}$ and $\mathbf{x} \geq \mathbf{y}$, it holds that

$$F(\mathbf{x}) \leq F(\mathbf{y}) + \langle \nabla F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle.$$

Additionally, the following lemmas are useful in our analysis.

Lemma 1 ([22]). For every $\mathbf{x}, \mathbf{y} \in \mathcal{K}$ and any DR-submodular function F , it holds that

$$\langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \geq F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y}) - 2F(\mathbf{x})$$

Lemma 2 ([18, 8, 1]). For every $\mathbf{x}, \mathbf{y} \in \mathcal{K}$ and for any DR-submodular function $F : [0, 1]^n \rightarrow \mathbb{R}^+ \cup \{0\}$, it holds that $F(\mathbf{x} \vee \mathbf{y}) \geq (1 - \|\mathbf{x}\|_\infty)F(\mathbf{y})$.

Variance Reduction Technique. Our algorithm and its analysis relies on a recent variance reduction technique proposed by [29]. This theorem has also been used in the context of online monotone submodular maximization [9, 38].

Lemma 3 ([29], Lemma 2). Let $\{\mathbf{a}^t\}_{t=0}^T$ be a sequence of points in \mathbb{R}^n . such that $\|\mathbf{a}^t - \mathbf{a}^{t-1}\| \leq C/(t+s)$ for all $1 \leq t \leq T$ with fixed constants $C \geq 0$ and $s \geq 3$. Let $\{\tilde{\mathbf{a}}\}_{t=1}^T$ be a sequence of random variables such that $\mathbb{E}[\tilde{\mathbf{a}} | \mathcal{H}^{t-1}] = \mathbf{a}^t$ and $\mathbb{E}[\|\tilde{\mathbf{a}} - \mathbf{a}^t\|^2 | \mathcal{H}^{t-1}] \leq \sigma^2$ for every $t \geq 0$, where \mathcal{H}^{t-1} is the history up to

$t - 1$. Let $\{\mathbf{d}^t\}_{t=0}^T$ be a sequence of random variables where \mathbf{d}^0 is fixed and subsequent \mathbf{d}^t are obtained by the recurrence

$$\mathbf{d}^t = (1 - \rho^t)\mathbf{d}^{t-1} + \rho^t \tilde{\mathbf{a}}^t$$

with $\rho^t = \frac{2}{(t+s)^{2/3}}$. Then we have

$$\mathbb{E} \left[\left\| \mathbf{a}^t - \mathbf{d}^t \right\|^2 \right] \leq \frac{Q}{(t+s+1)^{2/3}},$$

where $Q = \max\{\|\mathbf{a}^0 - \mathbf{d}^0\|^2 (s+1)^{2/3}, 4\sigma^2 + 3C^2/2\}$.

3 Online DR-Submodular Maximization over Down-Closed Convex Sets

3.1 Online Vee Learning

In this section, we give an algorithm for an online problem that will be the main building block in the design of algorithms for online DR-submodular maximization over down-closed convex sets. In the online vee learning problem, given a down-closed convex set \mathcal{K} , at every time t , the online algorithm receives $\mathbf{c}^t \in \mathcal{K}$ at the beginning of the step and needs to choose a vector $\mathbf{x}^t \in \mathcal{K}$. Subsequently, the adversary reveals a vector $\mathbf{a}^t \in \mathbb{R}^n$ and the algorithm receives a reward $\langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle$. The goal is to maximize the total reward $\sum_{t=1}^T \langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle$ over a time horizon T .

The main issue in the online vee problem is that the reward functions are non-concave. In order to overcome this obstacle, we consider a novel approach that consists of discretizing and lifting the corresponding functions to a higher dimensional space.

Discretization and lifting. Let \mathcal{L} be a lattice such that $\mathcal{L} = \{0, \frac{1}{M}, \frac{2}{M}, \dots, \frac{\ell}{M}, \dots, 1\}^n$ where $0 \leq \ell \leq M$ for some parameter M (to be defined later). Note that each $x_i = \frac{\ell}{M}$ for $0 \leq \ell \leq M$ can be uniquely represented by a vector $\mathbf{y}_i \in \{0, 1\}^{M+1}$ such that $y_{i0} = \dots = y_{i\ell} = 1$ and $y_{i,\ell+1} = \dots = y_{i,M} = 0$. Based on this observation, we lift the discrete set $\mathcal{K} \cap \mathcal{L}$ to the $(n \times (M+1))$ -dim space. Specifically, define a *lifting* map $m : \mathcal{K} \cap \mathcal{L} \rightarrow \{0, 1\}^{n \times (M+1)}$ such that each point $(x_1, \dots, x_n) \in \mathcal{K} \cap \mathcal{L}$ is mapped to a unique point $(y_{10}, \dots, y_{1M}, \dots, y_{n0}, \dots, y_{nM}) \in \{0, 1\}^{n \times (M+1)}$ where $y_{i0} = \dots = y_{i\ell} = 1$ and $y_{i,\ell+1} = \dots = y_{i,M} = 0$ iff $x_i = \frac{\ell}{M}$ for $0 \leq \ell \leq M$. Define $\tilde{\mathcal{K}}$ be the set $\{\mathbf{1}_X \in \{0, 1\}^{n \times (M+1)} : \mathbf{1}_X = m(\mathbf{x}) \text{ for some } \mathbf{x} \in \mathcal{K} \cap \mathcal{L}\}$. Observe that $\tilde{\mathcal{K}}$ is a subset of discrete points in $\{0, 1\}^{n \times (M+1)}$. Let $\mathcal{C} := \text{conv}(\tilde{\mathcal{K}})$ be the convex hull of $\tilde{\mathcal{K}}$.

Algorithm description. In our algorithm, at every time t , we will output $\mathbf{x}^t \in \mathcal{L} \cap \mathcal{K}$. In fact, we will reduce the online vee problem to a online linear optimization problem in the $(n \times (M+1))$ -dim space. Given $\mathbf{c}^t \in \mathcal{K}$, we round every coordinate c_i^t for $1 \leq i \leq n$ to the largest multiple of $\frac{1}{M}$ which is smaller than c_i^t . In other words, the rounded vector $\bar{\mathbf{c}}^t$ has the i^{th} -coordinate $\bar{c}_i^t = \frac{\ell}{M}$ where $\frac{\ell}{M} \leq c_i^t < \frac{\ell+1}{M}$. Vector $\bar{\mathbf{c}}^t \in \mathcal{L}$ and also $\bar{\mathbf{c}}^t \in \mathcal{K}$ (since $\bar{\mathbf{c}}^t \leq \mathbf{c}^t$ and \mathcal{K} is down-closed). Denote $\mathbf{1}_{C^t} = m(\bar{\mathbf{c}}^t)$. Besides, for each vector $\mathbf{a}^t \in \mathbb{R}^n$, define its correspondence $\tilde{\mathbf{a}}^t \in \mathbb{R}^{n \times (M+1)}$ such that $\tilde{a}_{i,j}^t = \frac{1}{M} a_i^t$ for all $1 \leq i \leq n$ and $0 \leq j \leq M$. Observe that $\langle \mathbf{a}^t, \bar{\mathbf{c}}^t \rangle = \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle$ (where the second scalar product is taken in the space of dimension $n \times (M+1)$).

The formal description is given in Algorithm 1. In the algorithm, we use a procedure `update`. This procedure takes arguments as the polytope \mathcal{C} , the current vector \mathbf{y}^t , the gradients of previous time steps $\tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t})$ and outputs the next vector \mathbf{y}^{t+1} . One can use different update strategies, in particular the gradient descent or the follow-the-perturbed-leader algorithm (if aiming for a projection-free algorithm).

$$\mathbf{y}^{t+1} \leftarrow \text{Proj}_{\mathcal{C}} \left(\mathbf{y}^t - \eta \cdot \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}) \right) \quad (\text{Gradient Descent})$$

Algorithm 1 Online algorithm for vee learning

Input: A convex set \mathcal{K} , a time horizon T .

- 1: Let \mathcal{L} be a lattice in $[0, 1]^n$ and denote the polytope $\mathcal{C} := \text{conv}(\tilde{\mathcal{K}})$.
 - 2: Initialize arbitrarily $\mathbf{x}^1 \in \mathcal{K} \cap \mathcal{L}$ and $\mathbf{y}^1 = \mathbf{1}_{X^1} = m(\mathbf{x}^1) \in \tilde{\mathcal{K}}$.
 - 3: **for** $t = 1$ to T **do**
 - 4: Observe \mathbf{c}^t , compute $\mathbf{1}_{C^t} = m(\bar{\mathbf{c}}^t)$.
 - 5: Play \mathbf{x}^t .
 - 6: Observe \mathbf{a}^t (so compute $\tilde{\mathbf{a}}^t$) and receive the reward $\langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle$.
 - 7: Set $\mathbf{y}^{t+1} \leftarrow \text{update}_{\mathcal{C}}(\mathbf{y}^t; \tilde{\mathbf{a}}^{t'} \circ (\mathbf{1} - \mathbf{1}_{C^{t'}}) : t' \leq t)$ where \circ denotes the element-wise multiplication.
 - 8: Round: $\mathbf{1}_{X^{t+1}} \leftarrow \text{round}(\mathbf{y}^{t+1})$.
 - 9: Compute $\mathbf{x}^{t+1} = m^{-1}(\mathbf{1}_{X^{t+1}})$.
 - 10: **end for**
-

$$\mathbf{y}^{t+1} \leftarrow \arg \min_{\mathcal{C}} \left\{ \eta \sum_{t'=1}^t \langle \tilde{\mathbf{a}}^{t'} \circ (\mathbf{1} - \mathbf{1}_{C^{t'}}), \mathbf{y} \rangle + \mathbf{n}^T \cdot \mathbf{y} \right\} \quad (\text{Follow-the-perturbed-leader})$$

Besides, in the algorithm, we use additionally a procedure `round` in order to transform a solution \mathbf{y}^t in the polytope \mathcal{C} of dimension $n \times (M + 1)$ to an integer point in $\tilde{\mathcal{K}}$. Specifically, given $\mathbf{y}^t \in \text{conv}(\tilde{\mathcal{K}})$, we round \mathbf{y}^t to $\mathbf{1}_{X^t} \in \tilde{\mathcal{K}}$ using an efficient polynomial-time algorithm given in [28] for the approximate Carathéodory's theorem w.r.t the ℓ_2 -norm. Specifically, given $\mathbf{y}^t \in \mathcal{C}$ and an arbitrary $\epsilon > 0$, the algorithm in [28] returns a set of $k = O(1/\epsilon^2)$ integer points $\mathbf{1}_{X_1^t}, \dots, \mathbf{1}_{X_k^t} \in \tilde{\mathcal{K}}$ in $O(1/\epsilon^2)$ -time such that $\|\mathbf{y}^t - \frac{1}{k} \sum_{j=1}^k \mathbf{1}_{X_j^t}\| \leq \epsilon$. Hence, given $\mathbf{1}_{X_1^t}, \dots, \mathbf{1}_{X_k^t}$ which have been computed, the procedure `round` simply consists of rounding \mathbf{y}^t to $\mathbf{1}_{X_j^t}$ with probability $1/k$. Finally, the solution \mathbf{x}^t in the original space is computed as $m^{-1}(\mathbf{1}_{X^{t+1}})$.

Lemma 4. Assume that $\|\mathbf{a}^t\| \leq G$ for all $1 \leq t \leq T$ and the update scheme is chosen either as the gradient ascent update or other procedure with regret of $O(\sqrt{T})$. Then, using the lattice \mathcal{L} with parameter $M = (T/n)^{1/4}$ and $\epsilon = 1/\sqrt{T}$ (in the round procedure), Algorithm 1 achieves a regret bound of $O(G(nT)^{3/4})$. The total running time of the algorithm is $O(T^2)$.

Proof. Observe that \mathbf{y}^t be the solution of the online linear optimization problem in dimension $n \times (M + 1)$ at time t where the reward function at time t is $\langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \cdot \rangle$. Moreover, recall that $\mathbf{x}^t \in \mathcal{L} \cap \mathcal{K}$ for every time t . First, by the construction, we have:

$$\begin{aligned} m(\bar{\mathbf{c}}^t \vee \mathbf{x}^t) &= m(\bar{\mathbf{c}}^t) \vee m(\mathbf{x}^t) = \mathbf{1}_{C^t} \vee \mathbf{1}_{X^t}, \\ \langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \mathbf{x}^t \rangle &= \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \vee \mathbf{1}_{X^t} \rangle. \end{aligned}$$

Therefore,

$$\sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \mathbf{x}^t \rangle] = \sum_{t=1}^T \mathbb{E}[\langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \vee \mathbf{1}_{X^t} \rangle]$$

Besides, for every vector $\mathbf{1}_C$ and vector $\mathbf{z} \in [0, 1]^{n \times M}$, the following identity always holds: $\mathbf{1}_C \vee \mathbf{z} = \mathbf{z} + \mathbf{1}_C \circ (\mathbf{1} - \mathbf{z})$. Therefore, for every vector $\mathbf{z} \in [0, 1]^{n \times M}$,

$$\langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \vee \mathbf{y} \rangle = \langle \tilde{\mathbf{a}}^t, \mathbf{z} + \mathbf{1}_{C^t} \circ (\mathbf{1} - \mathbf{z}) \rangle = \langle \tilde{\mathbf{a}}^t, \mathbf{z} \circ (\mathbf{1} - \mathbf{1}_{C^t}) \rangle + \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle$$

$$= \langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbf{z} \rangle + \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle \quad (5)$$

where note that the last term $\langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle$ is independent of the decision at time t .

By linearity of expectation, we have

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \mathbf{x}^t \rangle] &= \sum_{t=1}^T \mathbb{E}[\langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbf{1}_{X^t} \rangle + \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle] \\ &= \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbf{y}^t \rangle + \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbb{E}[\mathbf{1}_{X^t}] - \mathbf{y}^t \rangle + \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle \\ &\geq \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbf{y}^t \rangle + \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle - \sum_{t=1}^T \|\tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t})\| \cdot \|\mathbb{E}[\mathbf{1}_{X^t}] - \mathbf{y}^t\| \\ &\geq \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbf{y}^t \rangle + \sum_{t=1}^T \langle \tilde{\mathbf{a}}^t, \mathbf{1}_{C^t} \rangle - TG\epsilon. \end{aligned} \quad (6)$$

The first inequality follows Cauchy-Schwarz inequality. The last inequality is due to the property of `round` and $\|\tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t})\| \leq G$ since $\|\mathbf{a}^t\| \leq G$ and the construction of $\tilde{\mathbf{a}}^t$.

Let \mathbf{x}^* be the optimal solution in hindsight for the online vee learning problem. Let $\bar{\mathbf{x}}^*$ be the rounded solution of \mathbf{x}^* onto the lattice \mathcal{L} . Denote $\mathbf{1}_{X^*} = m(\bar{\mathbf{x}}^*)$. By (5) and (6) and the choice of $\epsilon = O(1/\sqrt{T})$, we have

$$\sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \mathbf{x}^t \rangle] - \sum_{t=1}^T \langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \bar{\mathbf{x}}^* \rangle \geq \sum_{t=1}^T \mathbb{E}[\langle \tilde{\mathbf{a}}^t \circ (\mathbf{1} - \mathbf{1}_{C^t}), \mathbf{y}^t - \mathbf{1}_{X^*} \rangle] \geq -O(nMG\sqrt{T})$$

where the last inequality is due to the regret bound of the gradient descent or the follow-the-perturbed-leader algorithms (in dimension nM). As $\|\mathbf{a}^t\| \leq G$, function $\langle \mathbf{a}^t, \cdot \rangle$ is G -Lipschitz. Hence,

$$\begin{aligned} &\sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle] - \sum_{t=1}^T \langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^* \rangle \\ &= \sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^t \rangle] - \sum_{t=1}^T \langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \mathbf{x}^t \rangle + \sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \mathbf{x}^t \rangle] - \sum_{t=1}^T \langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \bar{\mathbf{x}}^* \rangle \\ &\quad + \sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \bar{\mathbf{c}}^t \vee \bar{\mathbf{x}}^* \rangle] - \sum_{t=1}^T \langle \mathbf{a}^t, \mathbf{c}^t \vee \bar{\mathbf{x}}^* \rangle + \sum_{t=1}^T \mathbb{E}[\langle \mathbf{a}^t, \mathbf{c}^t \vee \bar{\mathbf{x}}^* \rangle] - \sum_{t=1}^T \langle \mathbf{a}^t, \mathbf{c}^t \vee \mathbf{x}^* \rangle \\ &\geq -O(nMG\sqrt{T}) - O(TG\frac{\sqrt{n}}{M}). \end{aligned}$$

where $\|\mathbf{c}^t \vee \mathbf{x}^t - \bar{\mathbf{c}}^t \vee \mathbf{x}^t\| \leq \|\mathbf{c}^t - \bar{\mathbf{c}}^t\| \leq \sqrt{n}/M$ and similarly $\|\bar{\mathbf{c}}^t \vee \bar{\mathbf{x}}^* - \mathbf{c}^t \vee \bar{\mathbf{x}}^*\|$ and $\|\mathbf{c}^t \vee \bar{\mathbf{x}}^* - \bar{\mathbf{c}}^t \vee \mathbf{x}^*\|$ are bounded by \sqrt{n}/M . Choosing $M = (T/n)^{1/4}$, one gets the regret bound of $O(G(nT)^{3/4})$.

Finally, the procedure `round` has average running time $O(1/\epsilon^2) = O(T)$ per time step; that leads to the complexity of $O(T^2)$. \square

3.2 Algorithm for online non-monotone DR-submodular maximization

In this section, we will provide an algorithm for the problem of online non-monotone DR-submodular maximization. The algorithm maintains L online vee learning oracles $\mathcal{E}_1, \dots, \mathcal{E}_L$. At each time step t , the

algorithm starts from $\mathbf{0}$ (origin) and executes L steps of the Frank-Wolfe algorithm where the update vector \mathbf{v}_ℓ^t is constructed by combining the output \mathbf{u}_ℓ^t of the online vee optimization oracle \mathcal{E}_ℓ with the current vector \mathbf{x}_ℓ^t . In particular, \mathbf{v}_ℓ^t is set to be $\mathbf{x}_\ell^t \vee \mathbf{u}_\ell^t - \mathbf{x}_\ell^t$. Then, the solution \mathbf{x}^t is produced at the end of the L -th step. Subsequently, after observing the DR-submodular function F^t , the algorithm defines a vector \mathbf{d}_ℓ^t and feedbacks $\langle \mathbf{d}_\ell^t, \mathbf{x}_\ell^t \vee \mathbf{u}_\ell^t \rangle$ as the reward function at time t to the oracle \mathcal{E}_ℓ for $1 \leq \ell \leq L$. The pseudocode is presented in Algorithm 2.

Algorithm 2 Online algorithm for down-closed convex sets

Input: A convex set \mathcal{K} , a time horizon T , online vee optimization oracles $\mathcal{E}_1, \dots, \mathcal{E}_L$, step sizes $\rho_\ell \in (0, 1)$, and $\eta_\ell = 1/L$ for all $\ell \in [L]$

- 1: Initialize vee optimizing oracle \mathcal{E}_ℓ for all $\ell \in \{1, \dots, L\}$.
 - 2: Initialize $\mathbf{x}_1^t \leftarrow \mathbf{0}$ and $\mathbf{d}_0^t \leftarrow \mathbf{0}$ for every $1 \leq t \leq T$.
 - 3: **for** $t = 1$ to T **do**
 - 4: Compute $\mathbf{u}_\ell^t \leftarrow$ output of oracle \mathcal{E}_ℓ in round t for all $\ell \in \{1, \dots, L\}$.
 - 5: **for** $1 \leq \ell \leq L$ **do**
 - 6: Set $\mathbf{v}_\ell^t \leftarrow \mathbf{x}_\ell^t \vee \mathbf{u}_\ell^t - \mathbf{x}_\ell^t$
 - 7: Set $\mathbf{x}_{\ell+1}^t \leftarrow \mathbf{x}_\ell^t + \eta_\ell \mathbf{v}_\ell^t$.
 - 8: **end for**
 - 9: Set $\mathbf{x}^t \leftarrow \mathbf{x}_{L+1}^t$.
 - 10: Play \mathbf{x}^t , observe the function F^t and get the reward of $F^t(\mathbf{x}^t)$.
 - 11: Compute a gradient estimate \mathbf{g}_ℓ^t such that $\mathbb{E}[\mathbf{g}_\ell^t | \mathbf{x}_\ell^t] = \nabla F^t(\mathbf{x}_\ell^t)$ for all $\ell \in \{1, 2, \dots, L\}$.
 - 12: Compute $\mathbf{d}_\ell^t = (1 - \rho_\ell) \mathbf{d}_{\ell-1}^t + \rho_\ell \mathbf{g}_\ell^t$, for every $\ell \in \{1, \dots, L\}$
 - 13: Feedback $\langle \mathbf{d}_\ell^t, \mathbf{x}_\ell^t \vee \mathbf{u}_\ell^t \rangle$ to oracle \mathcal{E}_ℓ for all $\ell \in \{1, 2, \dots, L\}$,
 - 14: **end for**
-

We first prove a technical lemma.

Lemma 5. *Let \mathbf{x}_ℓ^t be the vectors constructed in Algorithm 2. Then, it holds that $1 - \|\mathbf{x}_{\ell+1}^t\|_\infty \geq \prod_{\ell'=1}^\ell (1 - \eta_{\ell'})$ for every $t \in \{1, 2, \dots, T\}$ and for every $\ell \in \{1, 2, \dots, L\}$.*

Proof. We show that inequality holds for every t . Fix an arbitrary $t \in \{1, 2, \dots, T\}$. For the ease of exposition, we drop the index t in the following proof. Let $(x_\ell)_i$, $(u_\ell)_i$ and $(v_\ell)_i$ be the i^{th} coordinates of vectors \mathbf{x}_ℓ , \mathbf{u}_ℓ and \mathbf{v}_ℓ for $1 \leq i \leq n$ and for $1 \leq \ell \leq L$, respectively. We first obtain the following recursion on fixed $(x_\ell)_i$.

$$1 - (x_{\ell+1})_i = 1 - ((x_\ell)_i + \eta_\ell (v_\ell)_i) \geq 1 - ((x_\ell)_i + \eta_\ell (1 - (x_\ell)_i)) = (1 - \eta_\ell)(1 - (x_\ell)_i)$$

where the inequality holds since $0 \leq (x_\ell \vee u_\ell)_i \leq 1$. Since $(x_1)_i = 0$, we have $1 - (x_{\ell+1})_i \geq \prod_{\ell'=1}^\ell (1 - \eta_{\ell'})$. The lemma holds since the inequality holds for every $1 \leq i \leq n$. \square

We are now ready to prove the main theorem of the section.

Theorem 1. *Let $\mathcal{K} \subseteq [0, 1]^n$ be a down-closed convex set with diameter D . Assume that functions F^t 's are G -Lipschitz, β -smooth and $\mathbb{E}[\|\tilde{\mathbf{g}}_\ell^t - \nabla F^t(\mathbf{x}_\ell^t)\|^2] \leq \sigma^2$ for every t and ℓ . Then, the step-sizes $\eta_\ell = 1/L$ and $\rho_\ell = 2/(\ell + 3)^{2/3}$ for all $1 \leq \ell \leq L$, Algorithm 2 achieves the following guarantee:*

$$\sum_{t=1}^T \mathbb{E}[F^t(\mathbf{x}^t)] \geq \left(\frac{1}{e} - O\left(\frac{1}{L}\right) \right) \sum_{t=1}^T F^t(\mathbf{x}^*) - O(n^{3/4} G T^{3/4}) - O\left(\frac{(\beta D + G + \sigma) D T}{L^{1/3}} \right).$$

Choosing $L = T^{3/4}$ and note that $D \leq \sqrt{n}$, we get

$$\sum_{t=1}^T \mathbb{E}[F^t(\mathbf{x}^t)] \geq \left(\frac{1}{e} - O\left(\frac{1}{T^{3/4}}\right) \right) \max_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}) - O((G + \sigma)n^{5/4}T^{3/4}). \quad (7)$$

Proof. Let \mathbf{x}^* be an optimal solution in hindsight, i.e., $\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x})$. Fix a time step $t \in \{1, 2, \dots, T\}$. For every $2 \leq \ell \leq L + 1$, we have following:

$$F^t(\mathbf{x}_\ell^t) - F^t(\mathbf{x}_{\ell-1}^t) \geq \langle \nabla F^t(\mathbf{x}_\ell^t), (\mathbf{x}_\ell^t - \mathbf{x}_{\ell-1}^t) \rangle - (\beta/2) \|\mathbf{x}_\ell^t - \mathbf{x}_{\ell-1}^t\|^2 \quad (\text{using } \beta\text{-smoothness})$$

$$= \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_\ell^t), \mathbf{v}_{\ell-1}^t \rangle - \frac{\beta \eta_{\ell-1}^2}{2} \|\mathbf{v}_{\ell-1}^t\|^2 \quad (\text{using the update step in the algorithm})$$

$$= \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t \rangle + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_\ell^t) - \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t \rangle - \frac{\beta \eta_{\ell-1}^2}{2} \|\mathbf{v}_{\ell-1}^t\|^2$$

$$\geq \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t \rangle - \eta_{\ell-1} \|\nabla F^t(\mathbf{x}_\ell^t) - \nabla F^t(\mathbf{x}_{\ell-1}^t)\| \cdot \|\mathbf{v}_{\ell-1}^t\| - \frac{\beta \eta_{\ell-1}^2}{2} \|\mathbf{v}_{\ell-1}^t\|^2 \quad (\text{Cauchy-Schwarz})$$

$$\geq \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t \rangle - \eta_{\ell-1} \beta \|\mathbf{x}_\ell^t - \mathbf{x}_{\ell-1}^t\| \cdot \|\mathbf{v}_{\ell-1}^t\| - \frac{\beta \eta_{\ell-1}^2}{2} \|\mathbf{v}_{\ell-1}^t\|^2 \quad (\text{using } \beta\text{-smoothness})$$

$$\geq \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t \rangle - 2\beta \eta_{\ell-1}^2 \|\mathbf{v}_{\ell-1}^t\|^2$$

$$\geq \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t \rangle - 2\beta \eta_{\ell-1}^2 D^2 \quad (\text{diameter of } \mathcal{K} \text{ is bounded})$$

$$= \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t \rangle + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{v}_{\ell-1}^t - (\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t) \rangle - 2\beta \eta_{\ell-1}^2 D^2$$

$$= \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t), \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t \rangle + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{v}_{\ell-1}^t - (\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t) \rangle + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t, \mathbf{v}_{\ell-1}^t - (\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t) \rangle - 2\beta \eta_{\ell-1}^2 D^2$$

$$\geq \eta_{\ell-1} (F^t(\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^*) - F^t(\mathbf{x}_{\ell-1}^t)) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{v}_{\ell-1}^t - (\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t) \rangle + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t, \mathbf{v}_{\ell-1}^t - (\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* - \mathbf{x}_{\ell-1}^t) \rangle - 2\beta \eta_{\ell-1}^2 D^2 \quad (\text{using concavity in positive direction})$$

$$\geq \eta_{\ell-1} (F^t(\mathbf{x}_{\ell-1}^t \vee \mathbf{x}^*) - F^t(\mathbf{x}_{\ell-1}^t)) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle - 2\beta \eta_{\ell-1}^2 D^2 \quad (\text{using the definition of } \mathbf{v}_{\ell-1}^t)$$

$$\begin{aligned}
&\geq \eta_{\ell-1} \left((1 - \|\mathbf{x}_{\ell-1}^t\|_\infty) F^t(\mathbf{x}^*) - F^t(\mathbf{x}_{\ell-1}^t) \right) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle \\
&\quad + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle - 2\beta\eta_{\ell-1}^2 D^2 \quad (\text{using Lemma 2}) \\
&\geq \eta_{\ell-1} \prod_{\ell'=1}^{\ell-2} (1 - \eta_{\ell'}) \cdot F^t(\mathbf{x}^*) - \eta_{\ell-1} F^t(\mathbf{x}_{\ell-1}^t) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle \\
&\quad + \eta_{\ell-1} \langle \nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle - 2\beta\eta_{\ell-1}^2 D^2 \quad (\text{using Lemma 5}) \\
&\geq \eta_{\ell-1} \prod_{\ell'=1}^{\ell-2} (1 - \eta_{\ell'}) \cdot F^t(\mathbf{x}^*) - \eta_{\ell-1} F^t(\mathbf{x}_{\ell-1}^t) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle \\
&\quad - \frac{\eta_{\ell-1}}{2} \left(\frac{1}{\gamma_{\ell-1}} \|\nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t\|^2 + \gamma_{\ell-1} \|\mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^*\|^2 \right) \\
&\quad - 2\beta\eta_{\ell-1}^2 D^2 \quad (\text{Young's Inequality, parameter } \gamma_{\ell-1} \text{ will be defined later}) \\
&\geq \eta_{\ell-1} \prod_{\ell'=1}^{\ell-2} (1 - \eta_{\ell'}) \cdot F^t(\mathbf{x}^*) - \eta_{\ell-1} F^t(\mathbf{x}_{\ell-1}^t) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle \\
&\quad - \frac{\eta_{\ell-1}}{2} \left(\frac{1}{\gamma_{\ell-1}} \|\nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t\|^2 + \gamma_{\ell-1} D^2 \right) - 2\beta\eta_{\ell-1}^2 D^2. \\
&\quad (\text{since } \|\mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^*\| \leq \|\mathbf{u}_{\ell-1}^t - \mathbf{x}^*\| \leq D.)
\end{aligned}$$

From the above we have

$$\begin{aligned}
F^t(\mathbf{x}_\ell^t) &\geq (1 - \eta_{\ell-1}) F^t(\mathbf{x}_{\ell-1}^t) + \eta_{\ell-1} \prod_{\ell'=1}^{\ell-2} (1 - \eta_{\ell'}) \cdot F^t(\mathbf{x}^*) + \eta_{\ell-1} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle \\
&\quad - \frac{\eta_{\ell-1}}{2} \left(\frac{1}{\gamma_{\ell-1}} \|\nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t\|^2 + \gamma_{\ell-1} D^2 \right) - 2\beta\eta_{\ell-1}^2 D^2 \\
&= \left(1 - \frac{1}{L} \right) F^t(\mathbf{x}_{\ell-1}^t) + \frac{1}{L} \left(1 - \frac{1}{L} \right)^{\ell-2} F^t(\mathbf{x}^*) + \frac{1}{L} \langle \mathbf{d}_{\ell-1}^t, \mathbf{x}_{\ell-1}^t \vee \mathbf{u}_{\ell-1}^t - \mathbf{x}_{\ell-1}^t \vee \mathbf{x}^* \rangle \\
&\quad - \frac{1}{2L} \left(\frac{1}{\gamma_{\ell-1}} \|\nabla F^t(\mathbf{x}_{\ell-1}^t) - \mathbf{d}_{\ell-1}^t\|^2 + \gamma_{\ell-1} D^2 \right) - \frac{2\beta D^2}{L^2} \quad (\text{replacing } \eta_\ell = 1/L)
\end{aligned}$$

Applying recursively on ℓ , we get:

$$\begin{aligned}
F^t(\mathbf{x}_\ell^t) &\geq \frac{\ell-1}{L} \left(1 - \frac{1}{L} \right)^{\ell-2} F^t(\mathbf{x}^*) + \frac{1}{L} \sum_{\ell'=1}^{\ell-1} \left(\left(1 - \frac{1}{L} \right)^{\ell-1-\ell'} \langle \mathbf{d}_{\ell'}^t, \mathbf{x}_{\ell'}^t \vee \mathbf{u}_{\ell'}^t - \mathbf{x}_{\ell'}^t \vee \mathbf{x}^* \rangle \right) \\
&\quad - \frac{1}{2L} \sum_{\ell'=1}^{\ell-1} \left(\left(1 - \frac{1}{L} \right)^{\ell-1-\ell'} \left(\frac{1}{\gamma_{\ell'}} \|\nabla F^t(\mathbf{x}_{\ell'}^t) - \mathbf{d}_{\ell'}^t\|^2 + \gamma_{\ell'} D^2 \right) \right) - \frac{2\beta D^2}{L} \quad (8)
\end{aligned}$$

Summing above inequality (with $\ell = L + 1$) over all $t \in [T]$ and using online vee maximization oracle with regret \mathcal{R}^T , we obtain

$$\begin{aligned}
& \sum_{t=1}^T F^t(\mathbf{x}_{L+1}^t) \\
& \geq \left(1 - \frac{1}{L}\right)^{L-1} \sum_{t=1}^T F^t(\mathbf{x}^*) + \frac{1}{L} \sum_{t=1}^T \sum_{\ell'=1}^L \left(\left(1 - \frac{1}{L}\right)^{L-\ell'} \langle \mathbf{d}_{\ell'}^t, \mathbf{x}_{\ell'}^t \vee \mathbf{u}_{\ell'}^t - \mathbf{x}_{\ell'}^t \vee \mathbf{x}^* \rangle \right) \\
& \quad - \frac{1}{2L} \sum_{t=1}^T \sum_{\ell'=1}^L \left(\left(1 - \frac{1}{L}\right)^{L-\ell'} \left(\frac{1}{\gamma_{\ell'}} \|\nabla F^t(\mathbf{x}_{\ell'}^t) - \mathbf{d}_{\ell'}^t\|^2 + \gamma_{\ell'} D^2 \right) \right) - \frac{2\beta T D^2}{L} \\
& \geq \left(1 - \frac{1}{L}\right)^{L-1} \sum_{t=1}^T F^t(\mathbf{x}^*) + \frac{1}{L} \sum_{\ell'=1}^L \left(1 - \frac{1}{L}\right)^{L-\ell'} \cdot \left(\sum_{t=1}^T \langle \mathbf{d}_{\ell'}^t, \mathbf{x}_{\ell'}^t \vee \mathbf{u}_{\ell'}^t - \mathbf{x}_{\ell'}^t \vee \mathbf{x}^* \rangle \right) \\
& \quad - \frac{1}{2L} \sum_{t=1}^T \sum_{\ell'=1}^L \left(\frac{1}{\gamma_{\ell'}} \|\nabla F^t(\mathbf{x}_{\ell'}^t) - \mathbf{d}_{\ell'}^t\|^2 + \gamma_{\ell'} D^2 \right) - \frac{2\beta T D^2}{L} \quad (\text{since } 1 - 1/L < 1) \\
& \geq \left(1 - \frac{1}{L}\right)^{L-1} \sum_{t=1}^T F^t(\mathbf{x}^*) - \frac{1}{L} \sum_{\ell'=1}^L \left(1 - \frac{1}{L}\right)^{L-\ell'} \mathcal{R}^T - \frac{2\beta T D^2}{L} \\
& \quad - \frac{1}{2L} \sum_{t=1}^T \sum_{\ell'=1}^L \left(\frac{1}{\gamma_{\ell'}} \|\nabla F^t(\mathbf{x}_{\ell'}^t) - \mathbf{d}_{\ell'}^t\|^2 + \gamma_{\ell'} D^2 \right) \quad (\mathcal{R}^T \text{ is the regret of a single oracle}) \\
& \geq \left(\frac{1}{e} - O\left(\frac{1}{L}\right) \right) \sum_{t=1}^T F^t(\mathbf{x}^*) - \mathcal{R}^T - \frac{2\beta T D^2}{L} - \frac{1}{2L} \sum_{t=1}^T \sum_{\ell'=1}^L \left(\frac{1}{\gamma_{\ell'}} \|\nabla F^t(\mathbf{x}_{\ell'}^t) - \mathbf{d}_{\ell'}^t\|^2 + \gamma_{\ell'} D^2 \right) \quad (9)
\end{aligned}$$

Claim 1. Choosing $\gamma_\ell = \frac{Q^{1/2}}{D(\ell+3)^{1/3}}$ for $1 \leq \ell \leq L$ where $Q = \{\max_{1 \leq t \leq T} \|\nabla F^t(\mathbf{0})\|^2 4^{2/3}, 4\sigma^2 + 3(\beta D)^2/2\}$, it holds that

$$\sum_{t=1}^T \sum_{\ell=1}^L \left(\frac{1}{\gamma_\ell} \mathbb{E}[\|\nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t\|^2] + \gamma_\ell D^2 \right) = O\left((\beta D + G + \sigma) D T L^{2/3}\right)$$

Proof of claim. Fix a time step t . We apply Lemma 3 to the left hand side of the claim inequality where, for $1 \leq \ell \leq L$, \mathbf{a}_ℓ denotes the gradients $\nabla F^t(\mathbf{x}_\ell^t)$, $\tilde{\mathbf{a}}_\ell$ denote the stochastic gradient estimate \mathbf{g}_ℓ^t . Additionally, $\rho_\ell = \frac{2}{(\ell+3)^{2/3}}$ for $1 \leq \ell \leq L$. Hence we get:

$$\mathbb{E}[\|\nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t\|^2] \leq \frac{Q}{(\ell+4)^{2/3}}. \quad (10)$$

With the choice $\gamma_\ell = \frac{Q^{1/2}}{D(\ell+4)^{1/3}}$, we obtain:

$$\begin{aligned} \sum_{\ell=1}^L \left(\frac{1}{\gamma_\ell} \mathbb{E}[\|\nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t\|^2] + \gamma_\ell D^2 \right) &\leq 2DQ^{1/2} \sum_{\ell=1}^L \frac{1}{(\ell+4)^{1/3}} \\ &\leq 2DQ^{1/2} \int_{\ell=1}^L \frac{1}{\ell^{1/3}} d\ell \leq 3DQ^{1/2} L^{2/3} = O\left((\beta D + G + \sigma)DL^{2/3}\right) \end{aligned}$$

The claim follows by summing the inequality above over all $1 \leq t \leq T$. \square

Combining Claim 1 and Inequality (9) we get:

$$\sum_{t=1}^T \mathbb{E}[F^t(\mathbf{x}_{L+1}^t)] \geq \left(\frac{1}{e} - O\left(\frac{1}{L}\right) \right) \sum_{t=1}^T F^t(\mathbf{x}^*) - \mathcal{R}^T - \frac{2\beta TD^2}{L} - O\left(\frac{(\beta D + G + \sigma)DT}{2L^{1/3}}\right).$$

By Lemma 4, the regret of the online vee oracle $\mathcal{R}^T = O(G(nT)^{3/4})$. Thus, the theorem follows. \square

Remark 1. • If a function F^t is not smooth then one can consider \hat{F}_δ^t defined as $\hat{F}_\delta^t(\mathbf{x}) := \mathbb{E}_{\mathbf{r} \sim \mathbb{B}}[F^t(\mathbf{x} + \delta \mathbf{r})]$ where \mathbb{B} is the n -dim unit ball and $\mathbf{r} \sim \mathbb{B}$ denotes an uniformly random variable taken over \mathbb{B} . It is known that \hat{F}_δ^t is nG/δ -smooth and $|\hat{F}_\delta^t(\mathbf{x}) - F_\delta^t(\mathbf{x})| \leq \delta G$ (for example, see [23, Lemma 2.6]). By considering \hat{F}_δ^t (instead of F^t) with $\delta = 1/\sqrt{T}$ and choosing $L = T^{9/4}$, one can achieve the same guarantee (7) stated in Theorem 1.

- In Algorithm 2, we assume the knowledge of the time horizon T . That can be avoided by the standard doubling trick (shown in the appendix) where Algorithm 2 is invoked repeatedly with a doubling time horizon.

4 Maximizing Non-Monotone DR-Submodular Functions over General Convex Sets

In this section, we consider the problem of maximizing non-monotone continuous DR-submodular functions over general convex sets. We show that beyond down-closed convex sets, the Meta-Frank-Wolfe algorithm, studied by [10] for monotone DR-submodular functions, with appropriately chosen step sizes provides indeed meaningful guarantees for the problem of online non-monotone DR-submodular maximization over general convex sets. Note that no algorithm with performance guarantee was known in the online setting.

Online linear optimization oracles. In our algorithms, we use multiple online linear optimization oracles to estimate the gradient of online arriving functions. This idea was originally developed for maximizing monotone DR-submodular functions [9]. Before presenting algorithms, we recall the online linear optimization problems and corresponding oracles. In the online linear optimization problem, at every time $1 \leq t \leq T$, the oracle selects $\mathbf{u}^t \in \mathcal{K}$. Subsequently, the adversary reveals a vector \mathbf{d}^t and feedbacks the function $\langle \cdot, \mathbf{d}^t \rangle$ (and a reward $\langle \mathbf{u}^t, \mathbf{d}^t \rangle$) to the oracle. The objective is to minimize the regret. There exists several oracles that guarantee sublinear regret, for example the gradient descent algorithm has the regret of $O(n\sqrt{T})$ (see for example [23]).

Algorithm description. At a high level, at every time t , our algorithm produces a solution \mathbf{x}^t by running L steps of the Frank-Wolfe algorithm that uses the outputs of L linear optimization oracles as update vectors. After the algorithm plays \mathbf{x}^t , it observes stochastic gradient estimates at L points in the convex domain. Subsequently, these estimates are averaged with the estimates from the previous round and are fed to the reward functions of L online linear oracles. The pseudocode is presented in the Algorithm 3.

Algorithm 3 Meta-Frank-Wolfe for general convex domains

Input: A convex set \mathcal{K} , a time horizon T , online linear optimization oracles $\mathcal{E}_1, \dots, \mathcal{E}_L$, step sizes $\rho_\ell \in (0, 1)$, and $\eta_\ell \in (0, 1)$

- 1: Initialize online linear optimization oracle \mathcal{E}_ℓ for all $\ell \in \{1, \dots, L\}$.
 - 2: Initialize $\mathbf{x}_1^t \leftarrow \mathbf{x}_0$ for some $\mathbf{x}_0 \in \mathcal{K}$ and $\mathbf{d}_0^t \leftarrow 0$ for every $1 \leq t \leq T$. Let $\mathbf{x}_0 \leftarrow \arg \min_{\mathbf{x} \in \mathcal{K}} \|\mathbf{x}\|_\infty$.
 - 3: **for** $t = 1$ to T **do**
 - 4: Set $\mathbf{v}_\ell^t \leftarrow$ output of oracle \mathcal{E}_ℓ in round $t - 1$ for all $\ell \in \{1, \dots, L\}$.
 - 5: Set $\mathbf{x}_{\ell+1}^t \leftarrow (1 - \eta_\ell)\mathbf{x}_\ell^t + \eta_\ell\mathbf{v}_\ell^t$.
 - 6: Play $\mathbf{x}^t = \mathbf{x}_{L+1}^t$.
 - 7: Observe \mathbf{g}_ℓ^t such that $\mathbb{E}[\mathbf{g}_\ell^t | \mathbf{x}_\ell^t] = \nabla F^t(\mathbf{x}_\ell^t)$.
 - 8: Set $\mathbf{d}_\ell^t \leftarrow (1 - \rho_\ell) \cdot \mathbf{d}_{\ell-1}^t + \rho_\ell \cdot \mathbf{g}_\ell^t$ for $\ell = \{1, \dots, L\}$.
 - 9: Feedback the reward $\langle \mathbf{v}_\ell^t, \mathbf{d}_\ell^t \rangle$ to \mathcal{E}_ℓ for $\ell = \{1, \dots, L\}$.
 - 10: **end for**
-

We prove first a technical lemma which is in the same line of Lemma 5.

Lemma 6. Let $(x_\ell^t)_i$ be the i^{th} coordinate of vector \mathbf{x}_ℓ^t in Algorithm 3. Setting $\kappa = \left(\frac{\ln 3}{2}\right)$ and $\eta_\ell = \left(\frac{\kappa}{\ell H_L}\right)$, $\forall \ell \in \{1, 2, \dots, L\}$ where H_L is the L^{th} Harmonic number, the following invariant holds true for every $t \in \{1, \dots, T\}$ and for every coordinate $i \in \{1, \dots, n\}$

$$1 - (x_{L+1}^t)_i \geq e^{-\kappa(1+O(1/\ln^2 L))} \cdot (1 - (x_1^t)_i)$$

Proof. Fix an arbitrary time step $t \in \{1, \dots, T\}$. For the ease of exposition, we drop the index t in the following proof. Using the update step 5 from Algorithm 3, for every $1 \leq \ell \leq L$ and for every coordinate $1 \leq i \leq n$, we have

$$\begin{aligned} 1 - (x_{\ell+1})_i &= 1 - (1 - \eta_\ell)(x_\ell)_i - \eta_\ell(v_\ell)_i \geq 1 - (1 - \eta_\ell)(x_\ell)_i - \eta_\ell = (1 - \eta_\ell)(1 - (x_\ell)_i) \\ &\geq e^{-\eta_\ell - \eta_\ell^2} \cdot (1 - (x_\ell)_i) \end{aligned}$$

where we use inequalities $(v_\ell)_i \leq 1$ and $1 - u \geq e^{-u-u^2}$ for $0 \leq u < 1/2$; and $\eta_\ell = \frac{\kappa}{\ell H_L}$. Therefore, we get

$$1 - (x_{L+1})_i \geq e^{-\sum_{\ell'=1}^L \eta_{\ell'} - \sum_{\ell'=1}^L \eta_{\ell'}^2} \cdot (1 - (x_1)_i) \geq e^{-\kappa(1+O(1/\ln^2 L))} \cdot (1 - (x_1)_i)$$

where the last inequality is due to the fact that $\sum_{\ell'=1}^L \eta_{\ell'} = \kappa$ and $\sum_{\ell'=1}^L \eta_{\ell'}^2 = \sum_{\ell'=1}^L \frac{\kappa^2}{\ell'^2 H_L^2} = O(1/\ln^2 L)$. \square

Theorem 2. Let $\mathcal{K} \subseteq [0, 1]^n$ be a convex set with the diameter D . Assume that for every $1 \leq t \leq T$,

1. F^t is β -smooth DR-submodular function and the norm of the gradient ∇F^t is bounded by G , i.e., $\|\nabla F(\mathbf{x})\| \leq G, \forall \mathbf{x} \in \mathcal{K}$,
2. the variance of the unbiased stochastic gradients is bounded by σ^2 , i.e., $\mathbb{E} \left[\|\mathbf{g}_\ell^t - \nabla F^t(\mathbf{x}_\ell^t)\|^2 \right] \leq \sigma^2$ for every $1 \leq \ell \leq L$; and
3. the online linear optimization oracles used in Algorithm 3 have regret at most $\mathcal{R}_T^\mathcal{E}$.

Then for every $1 \leq \ell \leq L$ setting $\kappa = \frac{\ln 3}{2}$, $\rho_\ell = \frac{2}{(\ell+3)^{2/3}}$ and $\eta_\ell = \frac{\kappa}{\ell H_L}$ where H_L is the L^{th} harmonic number, the following inequality holds true for Algorithm 3:

$$\sum_{t=1}^T \mathbb{E} \left[F^t(\mathbf{x}^t) \right] \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)DT}{\ln L} \right) - O(\mathcal{R}_T^\mathcal{E}).$$

In particular, if $\mathbf{0} \in \mathcal{K}$ (for e.g. conic convex sets) then choosing $\mathbf{x}_0 = \mathbf{0}$, $L = O(T)$ and online linear algorithm with guarantee regret $\mathcal{R}_T^\mathcal{E} = O(\sqrt{T})$ (for eg., Online Gradient Descent), Algorithm 3 achieves

$$\sum_{t=1}^T \mathbb{E} \left[F^t(\mathbf{x}^t) \right] \geq \frac{1}{3\sqrt{3}} \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)DT}{\ln T} \right) - O(\sqrt{T}).$$

Proof. Let $\mathbf{x}^* \in \mathcal{K}$ be a solution that maximizes $\sum_{t=1}^T F^t(\mathbf{x})$. Let

$$r = e^{-\kappa(1+O(1/\ln^2 L))} \cdot (1 - \max_i (x_1^t)_i) = e^{-\kappa(1+O(1/\ln^2 L))} \cdot (1 - \max_i (x_0)_i).$$

Note that from Lemma 6, we have that $(1 - \|\mathbf{x}_\ell^t\|_\infty) \geq r$ for every $1 \leq t \leq T$ and every $1 \leq \ell \leq L+1$.

We prove the following claim which is crucial in our analysis.

Claim 2. For every $1 \leq t \leq T$, it holds that

$$\begin{aligned} & 2F^t(\mathbf{x}_{L+1}^t) - rF^t(\mathbf{x}^*) \\ & \geq \left(\prod_{\ell=1}^L (1 - 2\eta_\ell) \right) \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right) - \sum_{\ell=1}^L \eta_\ell \left(\frac{1}{\gamma_\ell} \left\| \nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t \right\|^2 + \gamma_\ell D^2 \right) \\ & + \sum_{\ell=1}^L 2\eta_\ell \left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \langle \mathbf{d}_\ell^t, (\mathbf{x}^* - \mathbf{v}_\ell^t) \rangle - 3\beta D^2 \sum_{\ell=1}^L \eta_\ell^2 \end{aligned}$$

where D is the diameter of \mathcal{K} and γ_ℓ is any constant greater than 0.

Proof of claim. Fix a time step $t \in \{1, \dots, T\}$. For the ease of exposition, we drop the time index t in equations. For every $1 \leq \ell \leq L$, we have

$$\begin{aligned} & 2F(\mathbf{x}_{\ell+1}) - rF(\mathbf{x}^*) \\ & = 2F((1 - \eta_\ell)\mathbf{x}_\ell + \eta_\ell \mathbf{v}_\ell) - rF(\mathbf{x}^*) \quad (\text{the step 5 of Algorithm 3}) \\ & \geq 2 \left(F(\mathbf{x}_\ell) - \eta_\ell \langle \nabla F((1 - \eta_\ell)\mathbf{x}_\ell + \eta_\ell \mathbf{v}_\ell), \mathbf{x}_\ell - \mathbf{v}_\ell \rangle - \frac{\beta}{2} (\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 \right) - rF(\mathbf{x}^*) \\ & \quad (\text{using } \beta\text{-smoothness}) \end{aligned}$$

$$\begin{aligned}
&= 2F(\mathbf{x}_\ell) - 2\eta_\ell \langle \nabla F((1 - \eta_\ell)\mathbf{x}_\ell + \eta_\ell \mathbf{v}_\ell), \mathbf{x}_\ell - \mathbf{v}_\ell \rangle - \beta(\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 - rF(\mathbf{x}^*) \\
&= (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) - 2\eta_\ell \langle \nabla F((1 - \eta_\ell)\mathbf{x}_\ell + \eta_\ell \mathbf{v}_\ell) - \nabla F(\mathbf{x}_\ell), \mathbf{x}_\ell - \mathbf{v}_\ell \rangle \\
&\quad - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), \mathbf{x}_\ell - \mathbf{v}_\ell \rangle - \beta(\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 \\
&\geq (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) - 2\eta_\ell \|\nabla F((1 - \eta_\ell)\mathbf{x}_\ell + \eta_\ell \mathbf{v}_\ell) - \nabla F(\mathbf{x}_\ell)\| \|\mathbf{x}_\ell - \mathbf{v}_\ell\| \\
&\quad - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), \mathbf{x}_\ell - \mathbf{v}_\ell \rangle - \beta(\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 \quad (\text{Applying Cauchy-Schwarz}) \\
&\geq (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) - 2\eta_\ell \beta \|(1 - \eta_\ell)\mathbf{x}_\ell + \eta_\ell \mathbf{v}_\ell - \mathbf{x}_\ell\| \|\mathbf{x}_\ell - \mathbf{v}_\ell\| \\
&\quad - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), (\mathbf{x}_\ell - \mathbf{v}_\ell) \rangle - \beta(\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 \quad (\beta\text{-smoothness as gradient Lipschitz}) \\
&\geq (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), (\mathbf{x}_\ell - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 \\
&= (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) + 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), (\mathbf{v}_\ell - \mathbf{x}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{x}_\ell - \mathbf{v}_\ell\|^2 \\
&= (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) + 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), (\mathbf{x}^* - \mathbf{x}_\ell) \rangle + 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), (\mathbf{v}_\ell - \mathbf{x}^*) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \\
&= (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) + 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell), (\mathbf{x}^* - \mathbf{x}_\ell) \rangle - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle \\
&\quad - 2\eta_\ell \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \\
&\geq (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) + 2\eta_\ell (F(\mathbf{x}^* \vee \mathbf{x}_\ell) - 2F(\mathbf{x}_\ell)) - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle \\
&\quad - 2\eta_\ell \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \quad (\text{Lemma 1}) \\
&\geq (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) + 2\eta_\ell ((1 - \|\mathbf{x}_\ell\|_\infty)F(\mathbf{x}^*) - 2F(\mathbf{x}_\ell)) - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle \\
&\quad - 2\eta_\ell \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \quad (\text{Lemma 2}) \\
&\geq (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) + 2\eta_\ell (rF(\mathbf{x}^*) - 2F(\mathbf{x}_\ell)) - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle \\
&\quad - 2\eta_\ell \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \quad (\text{Lemma 6}) \\
&= (1 - 2\eta_\ell) (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) - 2\eta_\ell \langle \nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle \\
&\quad - 2\eta_\ell \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \\
&\geq (1 - 2\eta_\ell) (2F(\mathbf{x}_\ell) - rF(\mathbf{x}^*)) - \eta_\ell \left(\frac{1}{\gamma_\ell} \|\nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell\|^2 + \gamma_\ell \|\mathbf{v}_\ell - \mathbf{x}^*\|^2 \right) \\
&\quad - 2\eta_\ell \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta(\eta_\ell)^2 \|\mathbf{v}_\ell - \mathbf{x}_\ell\|^2 \quad (\text{using Young's Inequality})
\end{aligned}$$

where in the last inequality, η_ℓ 's are parameters to be defined later (specifically in Claim 4). Applying the

above inequality for $1 \leq \ell \leq L$ and note that the diameter of \mathcal{K} is bounded by D , we get

$$\begin{aligned}
2F(\mathbf{x}_{L+1}) - rF(\mathbf{x}^*) &\geq \prod_{\ell=1}^L (1 - 2\eta_\ell) (2F(\mathbf{x}_1) - rF(\mathbf{x}^*)) \\
&\quad - \sum_{\ell=1}^L \left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \eta_\ell \left(\frac{1}{\gamma_\ell} \|\nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell\|^2 + \gamma_\ell D^2 \right) \\
&\quad - \sum_{\ell=1}^L 2\eta_\ell \left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta D^2 \sum_{\ell=1}^L \eta_\ell^2 \\
&\geq \prod_{\ell=1}^L (1 - 2\eta_\ell) (2F(\mathbf{x}_1) - rF(\mathbf{x}^*)) \\
&\quad - \sum_{\ell=1}^L \eta_\ell \left(\frac{1}{\gamma_\ell} \|\nabla F(\mathbf{x}_\ell) - \mathbf{d}_\ell\|^2 + \gamma_\ell D^2 \right) \\
&\quad - \sum_{\ell=1}^L 2\eta_\ell \left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \langle \mathbf{d}_\ell, (\mathbf{x}^* - \mathbf{v}_\ell) \rangle - 3\beta D^2 \sum_{\ell=1}^L \eta_\ell^2
\end{aligned}$$

where the second inequality holds since $\left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \leq 1$. Hence, the claim follows \square

Summing the inequality in Claim 2 over all $t = \{1, 2, \dots, T\}$, we get

$$\begin{aligned}
&\sum_{t=1}^T \left(2F^t(\mathbf{x}_{L+1}^t) - rF^t(\mathbf{x}^*) \right) \\
&\geq \sum_{t=1}^T \prod_{\ell=1}^L (1 - 2\eta_\ell) \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right) - \sum_{t=1}^T \sum_{\ell=1}^L \eta_\ell \left(\frac{1}{\gamma_\ell} \|\nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t\|^2 + \gamma_\ell D^2 \right) \\
&\quad - \sum_{\ell=1}^L 2\eta_\ell \left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \sum_{t=1}^T \langle \mathbf{d}_\ell^t, (\mathbf{x}^* - \mathbf{v}_\ell^t) \rangle - 3\beta D^2 T \sum_{\ell=1}^L \eta_\ell^2 \tag{11}
\end{aligned}$$

Next, we bound the terms on the right hand side of Inequation (11) separately by the following claims.

Claim 3. *It holds that*

$$\sum_{t=1}^T \prod_{\ell=1}^L (1 - 2\eta_\ell) \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right) \geq \sum_{t=1}^T e^{-2\kappa(1+O(1/\ln^2 L))} \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right).$$

Proof of claim. Using the inequality $1 - u \geq e^{-(u+u^2)}$, we have that

$$\begin{aligned}
\sum_{t=1}^T \prod_{\ell=1}^L (1 - 2\eta_\ell) \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right) &\geq \sum_{t=1}^T e^{-2 \sum_{\ell=1}^L \eta_\ell - 4 \sum_{\ell=1}^L \eta_\ell^2} \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right) \\
&\geq \sum_{t=1}^T e^{-2\kappa(1+O(1/\ln^2 L))} \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right)
\end{aligned}$$

where the last inequality is due to the facts that $\sum_{\ell=1}^L \eta_\ell = \sum_{\ell=1}^L \frac{\kappa}{\ell H_L} = \kappa$ and $\sum_{\ell=1}^L \eta_\ell^2 = \sum_{\ell=1}^L \frac{\kappa^2}{\ell^2 H_L^2} = O(1/\ln^2 L)$ for a constant κ . \square

Claim 4. Choose $\gamma_\ell = \frac{Q^{1/2}}{D(\ell+3)^{1/3}}$ where $Q = \max\{\max_{1 \leq t \leq T} \|\nabla F^t(\mathbf{x}_0)\|^2 4^{2/3}, 4\sigma^2 + (3/2)\beta^2 D^2\} = O(G^2 + \sigma^2 + \beta^2 D^2)$ for $1 \leq \ell \leq L-1$, it holds that

$$\sum_{t=1}^T \sum_{\ell=1}^L \eta_\ell \left(\frac{1}{\gamma_\ell} \mathbb{E} \left[\left\| \nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t \right\|^2 \right] + \gamma_\ell D^2 \right) \leq \frac{3TDQ^{1/2}}{H_L}.$$

Proof of claim. Fix a time step t . We apply Lemma 3 to left hand side of the inequality where, for $1 \leq \ell \leq L$, \mathbf{a}_ℓ denotes the gradients $\nabla F(\mathbf{x}_\ell^t)$, $\tilde{\mathbf{a}}_\ell$ denote the stochastic gradient estimate \mathbf{g}_ℓ^t , and \mathbf{d}_ℓ denotes the vector \mathbf{d}_ℓ^t in the algorithm. Additionally, $\rho_\ell = \left(\frac{2}{(\ell+3)^{2/3}} \right)$ for $1 \leq \ell \leq L$.

We verify the conditions of Lemma 3. By the algorithm,

$$\begin{aligned} \left\| \nabla F(\mathbf{x}_\ell^t) - \nabla F(\mathbf{x}_{\ell-1}^t) \right\| &= \left\| \nabla F((1 - \eta_\ell)\mathbf{x}_\ell^t + \eta_\ell \mathbf{v}_\ell^t) - \nabla F(\mathbf{x}_{\ell-1}^t) \right\| \\ &\leq \beta \eta_\ell \|\mathbf{v}_\ell^t - \mathbf{x}_\ell^t\| \leq \beta D \left(\frac{\kappa}{\ell H_L} \right) \leq \beta D \left(\frac{\kappa}{\ell + 3} \right). \end{aligned}$$

Moreover, by the theorem assumption, $\mathbb{E} \left[\left\| \mathbf{g}_\ell^t - \nabla F(\mathbf{x}_\ell^t) \right\|^2 \right] \leq \sigma^2$. Hence, applying Lemma 3, we get

$$\mathbb{E} \left[\left\| \nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t \right\|^2 \right] \leq \frac{Q}{(\ell + 4)^{2/3}}. \quad (12)$$

Summing Equation (12) for $1 \leq t \leq T$ and $1 \leq \ell \leq L$ and for $\gamma_\ell = \frac{Q^{1/2}}{D(\ell+3)^{1/3}}$ and $\eta_\ell = \frac{\kappa}{\ell H_L}$, we obtain

$$\begin{aligned} \sum_{t=1}^T \sum_{\ell=1}^L \eta_\ell \left(\frac{1}{\gamma_\ell} \mathbb{E} \left[\left\| \nabla F^t(\mathbf{x}_\ell^t) - \mathbf{d}_\ell^t \right\|^2 \right] + \gamma_\ell D^2 \right) &\leq \frac{TDQ^{1/2}}{H_L} \cdot \sum_{\ell=1}^L \frac{1}{\ell(\ell + 4)^{1/3}} \\ &\leq \frac{TDQ^{1/2}}{H_L} \cdot \int_1^{L+1} \frac{1}{\ell^{4/3}} d\ell \leq \frac{TDQ^{1/2}}{H_L} \left(3 - \frac{3}{(L+1)^{1/3}} \right) = O\left(\frac{(\beta D + G + \sigma)DT}{\log L} \right) \end{aligned}$$

\square

Claim 5. It holds that

$$\sum_{\ell=1}^L 2\eta_\ell \left(\prod_{\ell'=\ell+1}^L (1 - 2\eta_{\ell'}) \right) \sum_{t=1}^T \langle \mathbf{d}_\ell^t, (\mathbf{x}^* - \mathbf{v}_\ell^t) \rangle \leq 2\kappa \mathcal{R}_T^\mathcal{E}.$$

Proof of claim. The claim follows from the definition of the regret for the online linear optimization oracles that is, $\sum_{t=1}^T \langle \mathbf{d}_\ell^t, (\mathbf{x}^* - \mathbf{v}_\ell^t) \rangle \leq \mathcal{R}_T^\mathcal{E}$, and $\sum_{\ell=1}^L \eta_\ell \leq \kappa$ and $\prod_{\ell'=\ell}^L (1 - 2\eta_{\ell'}) \leq 1$. \square

Taking the expectation on Inequality (11) and using Claims 3, 4, 5, we have

$$\begin{aligned} \sum_{t=1}^T \left(2\mathbb{E} \left[F^t(\mathbf{x}_{L+1}^t) \right] - rF^t(\mathbf{x}^*) \right) &\geq \sum_{t=1}^T e^{-2\kappa(1+O(1/\ln^2 L))} \left(2F^t(\mathbf{x}_1^t) - rF^t(\mathbf{x}^*) \right) \\ &\quad - O\left(\frac{(\beta D + G + \sigma)DT}{\ln L} \right) - 2\kappa\mathcal{R}_T^\mathcal{E} - 3\beta D^2 T \sum_{\ell=1}^{L-1} \eta_\ell^2 \end{aligned}$$

Rearranging terms and note that $\sum_{\ell=1}^{L-1} \eta_\ell^2 = O(1/\ln^2 L)$ and $F^t(\mathbf{x}_1^t) \geq 0$, we get

$$\begin{aligned} \sum_{t=1}^T \mathbb{E} \left[F^t(\mathbf{x}_{L+1}^t) \right] &\geq \frac{r(1 - e^{-2\kappa(1+O(1/\ln^2 L))})}{2} \sum_{t=1}^T F^t(\mathbf{x}^*) \\ &\quad - O\left(\frac{(\beta D + G + \sigma)DT}{\ln L} \right) - 2\kappa\mathcal{R}_T^\mathcal{E} - O(\beta D^2 T / \ln^2 L) \end{aligned}$$

For L sufficiently large (so $O(1/\ln^2 L)$ becomes negligible compared to a constant), the factor

$$r(1 - e^{-2\kappa(1+O(1/\ln^2 L))}) = e^{-\kappa(1+O(1/\ln^2 L))}(1 - e^{-2\kappa(1+O(1/\ln^2 L))})(1 - \|\mathbf{x}_0\|_\infty)$$

attains the maximum value of $\frac{1}{3\sqrt{3}}(1 - \|\mathbf{x}_0\|_\infty)$ at $\kappa = \frac{\ln 3}{2}$. Hence, we deduce that

$$\sum_{t=1}^T \mathbb{E} \left[F^t(\mathbf{x}_{L+1}^t) \right] \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)DT}{\ln L} \right) - O(\mathcal{R}_T^\mathcal{E}).$$

The theorem follows. \square

Remark 2. • The regret guarantee, in particular the approximation ratio, depends on the initial solution \mathbf{x}_0 . This confirms the observation that initialization plays an important role in non-convex optimization. For particular case where \mathcal{K} is the intersection of a cone and the hypercube $[0, 1]^n$ (so $\mathbf{0} \in \mathcal{K}$), Algorithm 3 provides vanishing $\left(\frac{1}{3\sqrt{3}} \right)$ -regret. Note that this is the first constant approximation for non-monotone DR-submodular maximization over a non-trivial convex domain beyond the class of down-closed convex domains.

- Assume that $\mathbf{0} \in \mathcal{K}$ and F^t 's are identical, i.e., $F^t = F \forall t$, the algorithm guarantees a convergence rate of $O(1/\log T)$. It means that to be ϵ -close to a solution which is $\frac{1}{3\sqrt{3}}$ -approximation to the optimal solution of F , the algorithm requires $T = O(2^{1/\epsilon})$ iterations. Note that the exponential complexity is unavoidable as any constant approximation algorithm for the multilinear extension of a submodular function (so DR-submodular) over a general convex set requires necessarily an exponential number of value queries [37].
- The assumptions on the smoothness of functions F^t 's and the knowledge of T in Algorithm 3 can be removed by the argument as Remark 1.

5 Experiments

In this section, we validate our online algorithms for non-monotone DR submodular optimization on a set of real-world dataset. We show the performance of our algorithms for maximizing non-monotone DR-submodular function over a down-closed polytope and over a general convex polytope. All experiments are performed in MATLAB using CPLEX optimization tool on MAC OS version 10.15.

In the revenue maximization problem, the goal of a company is to offer for free or advertise a product to users so that the revenue increases through their “word-of-mouth” effect on others. Here, we are given an undirected social network graph $G = (V, W)$, where $w_{ij} \in W$ represents the weight of the edge between vertex i and vertex j . If the company invests x_i unit of cost on an user $i \in V$ then user i becomes an advocate of the product with probability $1 - (1 - p)^{x_i}$ where $p \in (0, 1)$ is a parameter. Intuitively, this signifies that for investing a unit cost to i , we have an extra chance that the user i becomes an advocate with probability p [34].

Let $S \subset V$ be a random set of users who advocate for the product. Then the revenue with respect to S is defined as $\sum_{i \in S} \sum_{j \in V \setminus S} w_{ij}$. Let $F : [0, 1]^{|E|} \rightarrow \mathbb{R}$ be the expected revenue obtained in this model, that is

$$F(\mathbf{x}) = \mathbb{E}_S \left[\sum_{i \in S} \sum_{j \in V \setminus S} w_{ij} \right] = \sum_i \sum_{j: i \neq j} w_{ij} (1 - (1 - p)^{x_i}) (1 - p)^{x_j}$$

It has been shown that F is a non-monotone DR-submodular function [34].




Figure 1: Performance ratio for online revenue maximization over (a) a down-closed polytope; and (b) a general polytope.

In our setting, we consider the online variant of the revenue maximization on a social network where at time t the weight of an edge is given $w_{ij}^t \in \{0, 1\}$. The experiments are performed on the Facebook dataset that contains 20K users (vertices) and 1M relationships (edges). We choose the number of time steps to be $T = 1000$. At each time $t \in 1, \dots, T$, we randomly uniformly select 2000 vertices $V^t \subset V$, independently of V^1, \dots, V^{t-1} , and construct a batch B_t with edge-weights $w_{ij}^t = 1$ if and only if $i, j \in V^t$ and edge (i, j) exists in the Facebook dataset. In case if i or j do not belong to V^t , $w_{ij}^t = 0$. We set $p = 0.0001$. In the first experiment, we impose a maximum investment constraint on the problem such that $\sum_{i \in V} x_i \leq 1$. This, in

addition to $x_i \geq 0, \forall i \in V$ constitutes a *down-closed* feasible convex set. For the general convex set, we impose an additional minimum investment constraint on the problem such that $\sum_{i \in V} x_i \geq 0.1$.

For comparison purposes, we consider (offline) Frank-Wolfe variants as the benchmark. Variants of Frank-Wolfe algorithm are shown to be competitive for maximizing (offline) non-monotone DR-submodular functions over down-closed convex sets [1] and over general convex sets [14]. Moreover, as observed in [1, 14], variants of Frank-Wolfe algorithm behave well in practice with approximation ratio better than 0.9 in several experiment settings. For these reasons, we adopt Frank-Wolfe variants as our benchmark.

In Figure 1 we show the ratio of between the objective value achieved by the online algorithm and that of the benchmark over the down-closed convex set (Figure 1(a)) and the general convex set (Figure 1(b)). These experiments conform with the theoretical guarantees. Specifically, the offline Frank-Wolfe variant performs almost as well as the optimum solution (approximation ratio larger than 0.92) and our online algorithms attain a ratio larger than 0.44 compared with the offline benchmark. Hence, the approximation ratios of the online algorithms is larger than $0.44 * 0.92 > 0.4 > 1/e$.)

6 Conclusion

In this paper, we considered the regret minimization problems while maximization online arriving non-monotone DR-submodular functions over convex sets. We presented the first online algorithm that achieve a constant $1/e$ -approximation ratio with the regret of $O(T^{3/4})$ over any down-closed convex. Moreover, we presented an online algorithm that achieves an approximation guarantee (depending on the search space) for the problem of maximizing non-monotone continuous DR-submodular functions over a *general* convex set (not necessarily down-closed). Finally, we run experiments to verify the performance of our algorithms on a social revenue maximization problem on a Facebook user-relationship dataset. Some interesting directions are worth to be considered. The first direction is to improve the regret bound of $(1/e, O(T^{3/4}))$ for down-closed convex sets. The design of an $(1/e, O(\sqrt{T}))$ -regret algorithm would be challenging as it has to figure out the right update vector (as provided by vee oracles) at every time step in order to obtain the approximation ratio of $1/e$ while maintaining the regret of $O(\sqrt{T})$ as guaranteed by natural algorithms (such as gradient descent, etc) without being trapped into poor local minima. The second direction is to study the trade-off between the approximation ratio and the regret bound for DR-submodular maximization over general convex sets, for example the existence of an algorithm with a reasonable non-constant approximation ratio and regret of $O(\sqrt{T})$. Another direction for future work is to consider the non-monotone DR-submodular maximization in the bandit setting by combining our approach with algorithms for monotone DR-submodular functions in [38].

References

- [1] A. A. Bian, K. Levy, A. Krause, and J. M. Buhmann. Non-monotone continuous DR-submodular maximization: Structure and algorithms. In *Neural Information Processing Systems (NIPS)*, 2017.
- [2] A. A. Bian, B. Mirzasoleiman, J. Buhmann, and A. Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Artificial Intelligence and Statistics*, pages 111–120, 2017.
- [3] Y. Bian, J. Buhmann, and A. Krause. Optimal continuous DR-submodular maximization and applications to provable mean field inference. In *Proceedings of the 36th International Conference on Machine Learning*, pages 644–653, 2019.

- [4] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [5] N. Buchbinder and M. Feldman. Deterministic algorithms for submodular maximization problems. *ACM Transactions on Algorithms (TALG)*, 14(3):32, 2018.
- [6] N. Buchbinder, M. Feldman, J. Seffi, and R. Schwartz. A tight linear time $(1/2)$ -approximation for unconstrained submodular maximization. *SIAM Journal on Computing*, 44(5):1384–1402, 2015.
- [7] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
- [8] C. Chekuri, T. Jayram, and J. Vondrák. On multiplicative weight updates for concave and submodular function maximization. In *Proc. Conference on Innovations in Theoretical Computer Science*, pages 201–210, 2015.
- [9] L. Chen, C. Harshaw, H. Hassani, and A. Karbasi. Projection-free online optimization with stochastic gradient: From convexity to submodularity. In *Proceedings of the 35th International Conference on Machine Learning*, pages 814–823, 2018.
- [10] L. Chen, H. Hassani, and A. Karbasi. Online continuous submodular maximization. In *Proc. 21st International Conference on Artificial Intelligence and Statistics (AISTAT)*, 2018.
- [11] L. Chen, M. Zhang, and A. Karbasi. Projection-free bandit convex optimization. In *Proc. 22nd Conference on Artificial Intelligence and Statistics*, pages 2047–2056, 2019.
- [12] J. Djolonga and A. Krause. From map to marginals: Variational inference in bayesian submodular models. In *Advances in Neural Information Processing Systems*, pages 244–252, 2014.
- [13] J. Djolonga, S. Tschitschek, and A. Krause. Variational inference in mixed probabilistic submodular models. In *Advances in Neural Information Processing Systems*, pages 1759–1767, 2016.
- [14] C. Dürr, N. K. Thang, A. Srivastav, and L. Tible. Non-monotone dr-submodular maximization: Approximation and regret guarantees. In *Proc. 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [15] E. Elenberg, A. G. Dimakis, M. Feldman, and A. Karbasi. Streaming weak submodularity: Interpreting neural networks on the fly. In *Advances in Neural Information Processing Systems*, pages 4044–4054, 2017.
- [16] E. R. Elenberg, R. Khanna, A. G. Dimakis, S. Negahban, et al. Restricted strong convexity implies weak submodularity. *The Annals of Statistics*, 46(6B):3539–3568, 2018.
- [17] A. Ene and H. L. Nguyen. A reduction for optimizing lattice submodular functions with diminishing returns. *arXiv:1606.08362*, 2016.
- [18] M. Feldman, J. Naor, and R. Schwartz. A unified continuous greedy algorithm for submodular maximization. In *Proc. 52nd Symposium on Foundations of Computer Science (FOCS)*, pages 570–579, 2011.
- [19] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.

- [20] M. Gomez Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *Proc. 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1019–1028, 2010.
- [21] A. Guillory and J. A. Bilmes. Simultaneous learning and covering with adversarial noise. In *ICML*, volume 11, pages 369–376, 2011.
- [22] H. Hassani, M. Soltanolkotabi, and A. Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, pages 5841–5851, 2017.
- [23] E. Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [24] S. Ito and R. Fujimaki. Large-scale price optimization via network flow. In *Advances in Neural Information Processing Systems*, pages 3855–3863, 2016.
- [25] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proc. 9th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.
- [26] A. Kulesza, B. Taskar, et al. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2–3):123–286, 2012.
- [27] H. Lin and J. Bilmes. A class of submodular functions for document summarization. In *Proc. 49th Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 510–520, 2011.
- [28] V. Mirrokni, R. P. Leme, A. Vladu, and S. C.-w. Wong. Tight bounds for approximate carathéodory and beyond. In *Proc. 34th International Conference on Machine Learning*, pages 2440–2448, 2017.
- [29] A. Mokhtari, H. Hassani, and A. Karbasi. Conditional gradient method for stochastic submodular maximization: Closing the gap. In *Procs of the International Conference on Artificial Intelligence and Statistics*, volume 84, pages 1886–1895, 2018.
- [30] R. Niazadeh, N. Golrezaei, J. Wang, F. Susan, and A. Badanidiyuru. Online learning via offline greedy: Applications in market design and optimization. *Available at SSRN 3613756*, 2020.
- [31] R. Niazadeh, T. Roughgarden, and J. R. Wang. Optimal algorithms for continuous non-monotone submodular and dr-submodular maximization. In *Neural Information Processing Systems (NIPS)*, 2018.
- [32] T. Roughgarden and J. R. Wang. An optimal learning algorithm for online unconstrained submodular maximization. In *Conference On Learning Theory*, pages 1307–1325, 2018.
- [33] A. Singla, I. Bogunovic, G. Bartók, A. Karbasi, and A. Krause. Near-optimally teaching the crowd to classify. In *ICML*, pages 154–162, 2014.
- [34] T. Soma and Y. Yoshida. Non-monotone dr-submodular function maximization. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [35] M. Staib and S. Jegelka. Robust budget allocation via continuous submodular functions. In *International Conference on Machine Learning*, pages 3230–3240, 2017.

- [36] M. Streeter and D. Golovin. An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 1577–1584, 2009.
- [37] J. Vondrák. Symmetry and approximability of submodular maximization problems. *SIAM Journal on Computing*, 42(1):265–304, 2013.
- [38] M. Zhang, L. Chen, H. Hassani, and A. Karbasi. Online continuous submodular maximization: From full-information to bandit feedback. In *Advances in Neural Information Processing Systems*, pages 9206–9217, 2019.

Appendix

A Removing Knowledge of T

In this section, we show how to remove the assumption on the knowledge of T for Algorithm 3. The procedure for Algorithms 2 is similar. In particular, we use the standard doubling trick (for example [11]) where Algorithm 3 is invoked repeatedly with a doubling time horizon.

Algorithm 4 Meta-Frank-Wolfe with Doubling Trick

Input: A convex set \mathcal{K} and Algorithm 3

- 1: **for** $m = 0, 1, 2, \dots$ **do**
 - 2: $L := 2^{m+1}$
 - 3: Run Algorithm 3 with horizon 2^m , from the $(2^m + 1)$ -th iteration to the 2^{m+1} -th iteration.
 - 4: Let \mathbf{x}^t for $2^m + 1 \leq t \leq 2^{m+1}$ be the solution fo Algorithm 3.
 - 5: **end for**
-

Theorem 3. *Given Theorem 2, the following inequality holds true for Algorithm 4:*

$$\sum_{t=1}^T \mathbb{E} [F^t(\mathbf{x}^t)] \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)DT}{\ln T} \right)$$

Proof. From Theorem 2, for each m , it follows that

$$\sum_{t=2^m+1}^{2^{m+1}} \mathbb{E} [F^t(\mathbf{x}^t)] \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=2^m+1}^{2^{m+1}} F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)D2^{m+1}}{(m+1) \ln 2} \right).$$

Summing this quantity of $m = 0, 1, 2, \dots, \lceil \log_2(T+1) \rceil - 1$, we have that

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E} [F^t(\mathbf{x}^t)] \\ & \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - \sum_{m=0}^{\lceil \log_2(T+1) \rceil - 1} O\left(\frac{(\beta D + G + \sigma)D2^{m+1}}{(m+1) \ln 2} \right) \\ & \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - \sum_{m=0}^{\lceil \log_2(T+1) \rceil - 1} O\left(\frac{(\beta D + G + \sigma)D2^{m+1}}{\ln 2} \right) \\ & \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)D}{\ln 2} \right) \sum_{m=1}^{\lceil \log_2(T+1) \rceil} \left(\frac{2^m}{m} \right) \\ & \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)D}{\ln 2} \right) \frac{2.2^{\lceil \log_2(T+1) \rceil}}{\lceil \log_2(T+1) \rceil} \\ & \geq \left(\frac{1}{3\sqrt{3}} \right) (1 - \|\mathbf{x}_0\|_\infty) \max_{\mathbf{x}^* \in \mathcal{K}} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{(\beta D + G + \sigma)DT}{\ln T} \right). \end{aligned}$$

□

B Online DR-Submodular Maximization over Hypercube $[0, 1]^n$

B.1 Reduction to maximizing online discrete submodular functions

Building upon the salient ideas of the discretization and lifting the problem into a higher dimension space in Section 3.1, we will transform DR-submodular functions to discrete submodular functions in higher space and use an online algorithm in [32] to produce solutions in higher space. Subsequently, we project the latter to the original space and prove that, by this unusual continuous-to-discrete approach (in contrast to usual discrete-to-continuous relaxations), one can achieve a similar regret guarantee as the one proved in [32] for online discrete submodular maximization. Note that, the following discretization is different to that for online vee learning.

Discretization and Lifting. Let $F : [0, 1]^n \rightarrow [0, 1]$ be a non-negative DR-submodular function. Let \mathcal{L} be a lattice $\mathcal{L} = \{0, 2^{-M}, 2 \cdot 2^{-M}, \dots, \ell \cdot 2^{-M}, \dots, 1\}^n$ where $0 \leq \ell \leq 2^M$ for some parameter M . Note that each $x_i \in \{0, 2^{-M}, 2 \cdot 2^{-M}, \dots, \ell \cdot 2^{-M}, \dots, 1\}$ can be uniquely decomposed as $x_i = \sum_{j=0}^M 2^{-j} y_{ij}$ where $y_{ij} \in \{0, 1\}$. We lift the set $[0, 1]^n \cap \mathcal{L}$ to the $(n \times (M+1))$ -dim space. Specifically, define a *lifting* map $m' : [0, 1]^n \cap \mathcal{L} \rightarrow \{0, 1\}^{n \times (M+1)}$ such that each point $(x_1, \dots, x_n) \in \mathcal{L} \cap [0, 1]^n$ is mapped to the unique point $(y_{10}, \dots, y_{1M}, \dots, y_{n0}, \dots, y_{nM}) \in \{0, 1\}^{n \times (M+1)}$ where $x_i = \sum_{j=0}^M 2^{-j} y_{ij}$. Define function $\tilde{f} : \{0, 1\}^{n \times (M+1)} \rightarrow [0, 1]$ such that $\tilde{f}(\mathbf{1}_S) := F(m'^{-1}(\mathbf{1}_S))$; in other words, $\tilde{f}(\mathbf{1}_S) = F(\mathbf{x})$ where $\mathbf{x} \in [0, 1]^n \cap \mathcal{L}$ and $\mathbf{1}_S = m'(\mathbf{x})$.

Lemma 7. *If F is DR-submodular then \tilde{f} is a submodular function.*

Proof. We will prove that for any vectors $\mathbf{1}_S$ and $\mathbf{1}_T$ such that $S \subset T$ and every element k ,

$$\begin{aligned} \tilde{f}(\mathbf{1}_{T \cup k}) - \tilde{f}(\mathbf{1}_T) &\leq \tilde{f}(\mathbf{1}_{S \cup k}) - \tilde{f}(\mathbf{1}_S) \\ \Leftrightarrow F(m'^{-1}(\mathbf{1}_{T \cup k})) - F(m'^{-1}(\mathbf{1}_T)) &\leq F(m'^{-1}(\mathbf{1}_{S \cup k})) - F(m'^{-1}(\mathbf{1}_S)). \end{aligned} \quad (13)$$

The inequality is trivial if $k \in S \subset T$. Assume that $k \notin S \subset T$. By the definition of the map m' , as $k \notin S \subset T$, $m'^{-1}(\mathbf{1}_{S \cup k}) = m'^{-1}(\mathbf{1}_S) + m'^{-1}(\mathbf{1}_k)$, $m'^{-1}(\mathbf{1}_{T \cup k}) = m'^{-1}(\mathbf{1}_T) + m'^{-1}(\mathbf{1}_k)$, and $m'^{-1}(\mathbf{1}_{S \cup k}) \leq m'^{-1}(\mathbf{1}_{T \cup k})$. Hence, the inequality (13) holds by the DR-submodularity of F . Therefore, \tilde{f} is a submodular function. \square

B.2 Algorithm

Fix a lattice $\mathcal{L} = \{0, 2^{-M}, 2 \cdot 2^{-M}, \dots, \ell \cdot 2^{-M}, \dots, 1\}^n$ where $0 \leq \ell \leq 2^M$ for some parameter M (to be defined later). Let RW be the online $(1/2, O(\sqrt{T}))$ -regret randomized algorithm [32] for online discrete submodular functions on $\{0, 1\}^{n \times (M+1)}$. Initially, set $\mathbf{x}^1 \in [0, 1]^n$ arbitrarily. At every time $t \geq 1$,

1. Play \mathbf{x}^t .
2. Observe function $F^t : [0, 1]^n \rightarrow [0, 1]$. Let $\tilde{f}^t : \{0, 1\}^{n \times (M+1)} \rightarrow [0, 1]$ be the corresponding submodular function by the construction above. Let $\mathbf{1}_{S^{t+1}} \in \{0, 1\}^{n \times (M+1)}$ be the random solution returned by the algorithm $\text{RW}(\tilde{f}^1, \dots, \tilde{f}^t)$. Set $\mathbf{x}^{t+1} = m'^{-1}(\mathbf{1}_{S^{t+1}})$.

Theorem 4. *Assume that functions F^t 's are G -Lipschitz. Then, by choosing $M = \log T$, the above algorithm achieves*

$$\sum_{t=1}^T F^t(\mathbf{x}^t) \geq \frac{1}{2} \max_{\mathbf{x} \in [0, 1]^n} \sum_{t=1}^T F^t(\mathbf{x}) - O(nG\sqrt{T} \log T).$$

Proof. Let \mathbf{x}^* be the optimal solution in hindsight, i.e., $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in [0,1]^n} \sum_{t=1}^T F^t(\mathbf{x})$. Let $\bar{\mathbf{x}}^*$ be rounded solution of \mathbf{x}^* onto the lattice \mathcal{L} , i.e., \bar{x}_i^* is the largest multiple of $1/2^M$ which is smaller than x_i^* for every $1 \leq i \leq n$. Note that $\|\mathbf{x}^* - \bar{\mathbf{x}}^*\| \leq \frac{\sqrt{n}}{2^M}$. Denote $\mathbf{1}_{S^*} = m(\bar{\mathbf{x}}^*)$. Algorithm RW [32, Corollary 3.2] guarantees that:

$$\sum_{t=1}^T \tilde{f}^t(\mathbf{1}_{S^t}) \geq \frac{1}{2} \sum_{t=1}^T \tilde{f}^t(\mathbf{1}_{S^*}) - O(nM\sqrt{T}).$$

By the construction of functions \tilde{f}^t 's and the Lipschitz property of F^t 's, we get

$$\begin{aligned} \sum_{t=1}^T F^t(\mathbf{x}^t) &= \sum_{t=1}^T \tilde{f}^t(\mathbf{1}_{S^t}) \geq \frac{1}{2} \sum_{t=1}^T \tilde{f}^t(\mathbf{1}_{S^*}) - O(nM\sqrt{T}) \\ &= \frac{1}{2} \sum_{t=1}^T F^t(\bar{\mathbf{x}}^*) - O(nM\sqrt{T}) \geq \frac{1}{2} \sum_{t=1}^T F^t(\mathbf{x}^*) - O\left(\frac{\sqrt{n}G}{2^M}T\right) - O(nM\sqrt{T}). \end{aligned}$$

Choose $M = \log T$, the theorem follows. □