Online Submodular Maximization via Online Convex Optimization

Tareq Si Salem* Gözde Özcan* Iasonas Nikolaou† Evimaria Terzi† Stratis Ioannidis*

January 9, 2024

Abstract

We study monotone submodular maximization under general matroid constraints in the online setting. We prove that online optimization of a large class of submodular functions, namely, weighted threshold potential functions, reduces to online convex optimization (OCO). This is precisely because functions in this class admit a concave relaxation; as a result, OCO policies, coupled with an appropriate rounding scheme, can be used to achieve sublinear regret in the combinatorial setting. We show that our reduction extends to many different versions of the online learning problem, including the dynamic regret, bandit, and optimistic-learning settings.

---
*Northeastern University, Boston, MA, USA `{sisalem, gozcan, ioannidis}@ece.neu.edu`
†Boston University, Boston, MA, USA, `{nikolaou, evimaria}@bu.edu`

1 Introduction

In online submodular optimization (OSM) [1], submodular reward functions chosen by an adversary are revealed over several rounds. In each round, a decision maker first commits to a set satisfying, e.g., matroid constraints. Subsequently, the reward function is revealed and evaluated over this set. The objective is to minimize α-regret, i.e., the difference of the cumulative reward accrued from the one attained by a static set, selected by an α-approximation algorithm operating in hindsight. OSM has received considerable interest recently, both in the full information [2, 3, 1, 4] and bandit setting [4, 3, 5], where only the reward values (rather than the entire functions) are revealed.

Online convex optimization (OCO) studies a similar online setting in which reward functions are concave, and decisions are selected from a compact convex set. First proposed by Zinkevich [6], who showed that projected gradient ascent attains sublinear regret, OCO generalizes previous online problems like prediction with expert advice [7], and has become widely influential in the learning community [8, 9, 10]. Its success is evident from the multitude of OCO variants in literature: in *dynamic regret* OCO [6, 11, 12, 13], the regret is evaluated w.r.t. an optimal comparator sequence instead of an optimal static decision in hindsight. *Optimistic* OCO [14, 15, 16] takes advantage of benign sequences, in which reward functions are predictable: the decision maker attains tighter regret bounds when predictions are correct, falling back to the existing OCO regret guarantees when predictions are unavailable or inaccurate. *Bandit* OCO algorithms [17, 18, 19] study the aforementioned bandit setting, where again only rewards (i.e., function evaluations) are observed.

We make the following contributions:

- We provide a methodology for reducing OSM to OCO, when submodular functions selected by the adversary are bounded from above and below by concave relaxations and coupled with an opportune rounding. We prove that algorithms and regret guarantees in the OCO setting transfer to α-regret guarantees in the OSM setting, via a transformation that we introduce. Ratio α is determined by how well concave relaxations approximate the original submodular functions.
- We show that the above condition is satisfied by a wide class of submodular functions, namely, *weighted threshold potential* (WTP) functions. This class strictly generalizes weighted coverage functions [20, 21], and includes many important applications, including influence maximization [22], facility location [23, 24], cache networks [25, 26], similarity caching [27], demand forecasting [28], and team formation [29], to name a few.
- We show our reduction also extends to the dynamic regret and optimistic settings, reducing such full-information OSM settings to the respective OCO variants. To the best of our knowledge, our resulting algorithms are the first to come with guarantees for the dynamic regret and optimistic settings in the context of OSM with general matroid constraints. Finally, we also provide a different reduction for the bandit setting, again for all three (static, dynamic regret, and optimistic) variants; this reduction applies to general submodular functions, but is restricted to partition matroids.

The remainder of the paper is organized as follows. We present related work and a technical preliminary in Sec. 2 and Sec. 3, respectively, and our setting in Sec. 4. We state our main results in Sec. 5, and discuss extensions in Sec. 6. Our experimental results are in Sec. 7; we conclude in Sec. 8.

2 Related Work

Offline Submodular Maximization and Relaxations of Submodular Functions. Continuous relaxations of submodular functions play a prominent role in submodular maximization. The so-called continuous greedy algorithm [34] maximizes the multilinear relaxation of a submodular objective over the convex hull of a matroid, using a variant of the Frank-Wolfe algorithm [35]. The fractional solution

Paper	Prob. Class	(1 − 1/e)-regret (Full Information)									Time
		Static			Dynamic			Optimistic			
		Uni.	Part.	Gen.	Uni.	Part.	Gen.	Uni.	Part.	Gen.	
[4]	GS	$r\sqrt{\log(n)T}$	×	×	×	×	×	×	×	×	$T^4 O_b$
[2]	GS	$\sqrt{r\log\left(\frac{n}{r}\right)T}$			×	×	×	×	×	×	$nr^2 + O_m \cdot n^4/\epsilon^3$ $\log(n^3 T/\epsilon)$
[3]	GS	$r\sqrt{r\log(nT)T}$	×	×	$\sqrt{r(r\log(nT)+P_T)T}$	×	×	×	×	×	nr
[1]	GS	$r^{\frac{3}{2}}\sqrt{\log(n)T}$		×	×	×	×	×	×	×	$n^2 c_p$
[30]	DR-S	$\sqrt{rnT}$			×	×	×	×	×	×	$\sqrt{T}O_{oco}\cdot O_m + nr^2$
[31]	DR-S	$T^{\frac{4}{5}}$			×	×	×	×	×	×	$T^{\frac{3}{5}}O_{oco}\cdot O_m + nr^2$
[32]	DR-S	$\sqrt{rnT}$			×	×	×	×	×	×	$O_{oco}\cdot O_m + nr^2$
[33]	LWD	$n(\alpha+2)\sqrt{T}$			×	×	×	×	×	×	TO_α
This work	WTP	$r\sqrt{\log(\frac{n}{r})T}$			$\sqrt{r(r\log(\frac{n}{r})+\log(n)P_T)T}$			$\sqrt{r(r\log(\frac{n}{r})+\log(n)P_T}$ $\cdot\sqrt{\sum_{t=1}^{T}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\infty^2}$			nr^2

Table 1: Order $(\mathcal{O}(\cdot))$ comparison of regrets and time complexities attained by different online submodular optimization algorithms for general submodular (GS) and others for continuous DR-submodular (DR-S) functions. We also include Kakade et al. [33], who operate on Linearly Weighted Decomposable (LWD) functions, and our work on Weighted Threshold Potential (WTP) functions. We also indicate whether algorithms operate over uniform, partition, or general matroid constraints, in the static, bandit, and optimistic settings. Regret and complexity are characterized in terms of the time horizon T, the ground set size n, and the matroid rank r. Additional, algorithm-specific parameters are specified in Appendix C.VI, and the derivation of the regret constants in this work is provided in Appendix C. The SoTA general submodular+general matroid algorithm [2] has a tighter regret than us by a factor $\sqrt{r}$, but has a much higher computational complexity. We attain the same or better regret than the DR-S [30, 31, 32] and remaining algorithms, that also either operate on restricted constraint sets [4, 3, 1] or on the much more restrictive LWD class [33]. Most importantly, our work readily generalizes to the dynamic and optimistic settings. With the sole exception of Matsuoka et al. [3], who study dynamic regret restricted to uniform matroids, our work is the first to provide guarantees for online submodular optimization in the dynamic and optimistic settings under general matroid constraints. Leveraging the concave relaxation and eschewing computing the multilinear relaxation also yields significant computational complexity dividends.

is then rounded via pipage [36] or swap rounding [37], which we also use. The multilinear relaxation is not convex but is continuous DR-submodular [38, 39], and continuous greedy comes with a $1 - \frac{1}{e}$ approximation guarantee. However, the multilinear relaxation is generally not tractable and is usually estimated via sampling. Ever since the seminal paper by Ageev and Sviridenko [36], several works have exploited the existence of concave relaxations of weighted coverage functions (e.g., [20, 25]), a strict subset of the threshold potential functions we consider here. For coverage functions, a version of our "sandwich" property (Asm. 2) follows directly by the Goemans & Williamson inequality [40], which we also use. Both in the standard [36] and stochastic offline [20, 25] submodular maximization setting, in which the objective is randomized, exploiting concave relaxations of coverage functions yields significant computational dividends, as it eschews any sampling required for estimating the multilinear relaxation. We depart from both by considering a much broader class than coverage functions and studying the online/no-regret setting.

OSM via Regret Minimization. Several online algorithms have been proposed for maximizing general submodular functions [4, 2, 3, 1] under different matroid constraints. There has also been recent work [30, 31, 32] on the online maximization of continuous DR-submodular functions [38]. Proposed algorithms are applicable to our setting, because the multi-linear relaxation is DR-submodular, and guarantees can be extended to matroid constraint sets again through rounding [37], akin to the approach we follow here. Also pertinent is the work by Kakade et al. [33]: their proposed online algorithm operates over reward functions that can be decomposed as the weighted sum of finitely many (non-parametric)

reference functions—termed Linearly Weighted Decomposable (LWD); the adversary selects only the weights. Applied to OSM, this is a more restrictive function class than the ones we study.

We compare these algorithms to our work in Table 1. In the full information setting, the OSM algorithm by Harvey et al. [2] has a slightly tighter α-regret than us, but also a much higher computational complexity. We attain the same or better regret than DR-S [30, 31, 32] and remaining algorithms that either operate on restricted constraint sets [4, 3, 1] or on the much more restrictive LWD class [33]. Most importantly, our work generalizes to the dynamic and optimistic settings. With the exception of Matsuoka et al. [3], who study dynamic regret restricted to uniform matroids, our work provides the first guarantees for OSM in the dynamic and optimistic settings under general matroid constraints.

OSM in the Bandit Setting. Our reduction to OCO in the bandit setting extends the analysis by Wan et al. [5], who provide a reduction to just FTRL in the static setting, under general submodular functions and partition matroid constraints. We generalize this to any OCO algorithm and to the dynamic and optimistic settings. Interestingly, Wan et al. [5] conjecture that no sublinear regret algorithm exists for general submodular functions under general matroid constraints in the bandit setting. We compare to bounds attained by Wan et al. [5] and other bandit algorithms for OSM [4, 1, 31, 33] in Table 4 in Appendix G.III. Our main contribution is again the extension to the dynamic and optimistic settings.

3 Technical Preliminary

Submodularity and Matroids. Given a ground set $V = [n] \triangleq \{1, 2, \ldots, n\}$, a set function $f : 2^V \to \mathbb{R}$ is *submodular* if $f(S \cup \{i, j\}) - f(S \cup \{j\}) \leq f(S \cup \{i\}) - f(S)$ for all $S \subseteq V$ and $i, j \in V \setminus S$ and *monotone* if $f(A) \leq f(B)$ for all $A \subseteq B \subseteq 2^V$. A *matroid* is a pair $\mathcal{M} = (V, \mathcal{I})$, where $\mathcal{I} \subseteq 2^V$, for which the following holds: (1) if $B \in \mathcal{I}$ and $A \subseteq B$, then $A \in \mathcal{I}$, (2) if $A, B \in \mathcal{I}$ and $|A| < |B|$, then there exists an $x \in B \setminus A$ s.t. $A \cup \{x\} \in \mathcal{I}$. The *rank* $r \in \mathbb{N}$ of $\mathcal{M}$ is the cardinality of the largest set in $\mathcal{I}$. With slight abuse of notation, we represent set functions $f : 2^V \to \mathbb{R}$ as functions over $\{0, 1\}^n$: given a set function f and an $\mathbf{x} \in \{0, 1\}^n$, we denote by $f(\mathbf{x})$ as the value $f(\text{supp}(\mathbf{x}))$, where $\text{supp}(\mathbf{x}) \triangleq \{i \in V : x_i \neq 0\} \subseteq V$ is the support of $\mathbf{x}$. Similarly, we treat matroids as subsets of $\{0, 1\}^n$.

Online Learning. In the general protocol of *online learning* [41], a decision-maker makes sequential decisions and incurs rewards as follows: at timeslot $t \in [T]$, where $T \in \mathbb{N}$ is the time horizon, the decision-maker first commits to a decision $\mathbf{x}_t \in \mathcal{X}$ from some set $\mathcal{X}$. Then, a reward function $f_t : \mathcal{X} \to \mathbb{R}_{\geq 0}$ is selected by an adversary from a set of functions $\mathcal{F}$ and revealed to the decision-maker, who accrues a reward $f_t(\mathbf{x}_t)$. The decision $\mathbf{x}_t$ is determined according to a (potentially random) mapping $\mathcal{P}_{\mathcal{X}, t} : \mathcal{X}^{t-1} \times \mathcal{F}^{t-1} \to \mathcal{Y}$, i.e.,

$$\mathbf{x}_t = \mathcal{P}_{\mathcal{X}, t}\left((\mathbf{x}_s)_{s=1}^{t-1}, (f_s)_{s=1}^{t-1}\right). \tag{1}$$

Let $\boldsymbol{\mathcal{P}}_{\mathcal{X}} = (\mathcal{P}_{\mathcal{X}, t})_{t \in [T]}$ be the *online policy* of the decision-maker. We define the *regret* of $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ at horizon T as:

$$\text{regret}_T(\boldsymbol{\mathcal{P}}_{\mathcal{X}}) \triangleq \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \max_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) - \mathbb{E}_{\boldsymbol{\mathcal{P}}_{\mathcal{X}}}\left[\sum_{t=1}^T f_t(\mathbf{x}_t)\right] \right\}. \tag{2}$$

We seek policies that attain a sublinear (i.e., $o(T)$) regret; intuitively, such policies perform on average as well as the static optimum in hindsight. Note that the regret is defined w.r.t. the optimal *fixed* decision $\mathbf{x}$, i.e., the time-invariant decision $\mathbf{x} \in \mathcal{X}$ that would be optimal in hindsight, after the sequence $(f_t)_{t=1}^T$ is revealed. When selecting $\mathbf{x}_t$, the decision-maker has no information about the upcoming reward f_t. Finally, this is the *full-information* setting: at each timeslot, the decision maker observes the entire reward function $f_t(\cdot)$, rather than just the reward $f_t(\mathbf{x}_t) \in \mathbb{R}_{\geq 0}$.

Deviating from these assumptions is of both practical and theoretical interest. In the *dynamic regret* setting [6], the regret is measured w.r.t. a time-variant optimum, appropriately constrained so that changes from one timeslot to the next do not vary significantly. In *learning with optimism* [14], additional information is assumed to be available w.r.t. f_t, in the form of so-called predictions. In the *bandit* setting [42], the online policy $\mathcal{P}_\mathcal{X}$ of the decision maker only has access to rewards $f_t(\mathbf{x}_t) \in \mathbb{R}_{\geq 0}$, as opposed to the entire reward function.

Online Convex Optimization. The *online convex optimization* (OCO) framework [6, 8] follows the above online learning protocol (1), where (a) the decision space $\mathcal{X}$ is a convex set in $\mathbb{R}^n$, and (b) the set of reward functions $\mathcal{F}$ is a subset of concave functions over $\mathcal{X}$. Formally, OCO operates under the following assumption:

Assumption 1. *Set $\mathcal{X} \subset \mathbb{R}^n$ is convex and compact. The reward functions in $\mathcal{F}$ are all L-Lipschitz concave functions w.r.t. a norm $\|\cdot\|$ over $\mathcal{X}$, for some common $L \in \mathbb{R}_{>0}$.*

There is a rich literature on OCO policies [8]; examples include Online Gradient Ascent (OGA) [8], Online Mirror Ascent (OMA) [43], and Follow-The-Regularized-Leader (FTRL) [10]. All three enjoy sublinear regret:

Theorem 1. *Under Asm. 1 OGA, OMA, and FTRL attain $O(\sqrt{T})$ regret.*

Details on all three algorithms and the regret they attain are in Appendix A. Most importantly, the OCO framework generalizes to the dynamic, learning-with-optimism, as well as bandit settings (see Sec. 6).

Weighted Threshold Potentials. A *threshold potential* [21] $\Psi_{b,\mathbf{w},S} : \{0,1\}^n \to \mathbb{R}_{\geq 0}$, also known as a *budget-additive function* [44, 45, 46], is defined as:

$$\Psi_{b,\mathbf{w},S}(\mathbf{x}) \triangleq \min\left\{b, \sum_{j \in S} x_j w_j\right\}, \text{ for } \mathbf{x} \in \{0,1\}^n, \tag{3}$$

where $b \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ is a threshold, $S \subseteq V$ is a subset of $V = [n]$, and $\mathbf{w} = (w_j)_{j \in S} \in [0, b]^{|S|}$ is a weight vector bounded by b.[1] The linear combination of threshold potential functions defines the rich class of *weighted threshold potentials* (WTP) [21], defined as:

$$f(\mathbf{x}) \triangleq \sum_{\ell \in C} c_\ell \Psi_{b_\ell, \mathbf{w}_\ell, S_\ell}(\mathbf{x}), \quad \text{for } \mathbf{x} \in \{0,1\}^n, \tag{4}$$

where C is an arbitrary index set and $c_\ell \in \mathbb{R}_{\geq 0}$, for $\ell \in C$. WTP functions are submodular and monotone (see Appendix B). We define the *degree* of a WTP function $\Delta_f = \max_{\ell \in C} |S_\ell|$ as the maximum number of variables that a threshold potential Ψ in f depends on.

We give several examples of WTP functions in Appendix B. In short, classic problems such as influence maximization [22] and facility location [23, 24], resource allocation problems like cache networks [25, 26] and similarity caching [27], as well as demand forecasting [28] and team formation [29] can all be expressed using WTP functions. Overall, the WTP class is very broad: there exists a hierarchy among submodular functions, including weighted coverage functions [20], weighted cardinality truncations [47], and sums of concave functions composed with non-negative modular functions [21]; all of them are strictly dominated by the WTP class (see Stobbe and Krause [21], as well as Appendix B).

4 Problem Formulation

We consider a combinatorial version of the online learning protocol defined in Eq. (1). In particular, we focus on the case where (a) the decision set is $\mathcal{X} \subseteq \{0,1\}^n$, i.e., the vectorized representation of subsets

---

[1]Assumption $w_j \leq b$, $j \in S$, is w.l.o.g., as replacing w_j with $\min\{w_j, b\}$ preserves values of f over $\{0,1\}^n$.

Algorithm 1 Rounding-Augmented OCO (RAOCO) policy

Require: OCO policy $\boldsymbol{\mathcal{P}_\mathcal{Y}}$, randomized rounding $\Xi : \mathcal{Y} \to \mathcal{X}$
1: **for** $t = 1, 2, \ldots, T$ **do**
2:      $\mathbf{y}_t \leftarrow \mathcal{P}_{\mathcal{Y},t} \left((\mathbf{y}_s)_{s=1}^{t-1}, (\tilde{f}_s)_{s=1}^{t-1} \right)$
3:      $\mathbf{x}_t \leftarrow \Xi(\mathbf{y}_t)$
4:      Receive reward $f_t(\mathbf{x}_t)$
5:      Reward function f_t is revealed
6:      Construct $\tilde{f}_t$ from f_t satisfying Asm. 2
7: **end for**

---

of $V = [n]$, and (b) the set $\mathcal{F}$ of reward functions comprises set functions over $\mathcal{X}$. Though some of our results (e.g., Thm. 2) pertain to this general combinatorial setting, we are particularly interested in the case where (a) $\mathcal{X}$ is a matroid, and (b) $\mathcal{F}$ is the WTP class, i.e., the set of functions whose form is given by Eq. (4).

Both in the general combinatorial setting and for WTP functions, evaluating the best fixed decision may be computationally intractable even in hindsight, i.e., when all reward functions were revealed. As is customary (see, e.g., [23]), instead of the regret in Eq. (2), we consider the so-called α-regret:

$$\alpha\text{-regret}_T(\boldsymbol{\mathcal{P}_\mathcal{X}}) \triangleq \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \max_{\mathbf{x} \in \mathcal{X}} \alpha \sum_{t=1}^T f_t(\mathbf{x}) - \mathbb{E}_{\boldsymbol{\mathcal{P}_\mathcal{X}}} \left[\sum_{t=1}^T f_t(\mathbf{x}_t) \right] \right\}. \tag{5}$$

Intuitively, we compare the performance of the policy $\boldsymbol{\mathcal{P}_\mathcal{X}}$ w.r.t. the best polytime(n) α-approximation of the static offline optimum in hindsight. For example, the approximation ratio would be $\alpha = 1 - 1/e$ in the case of submodular set functions maximized over matroids.

5 Online Submodular Optimization via Online Convex Optimization

5.1 The Case of General Set Functions

First, we show how OCO can be leveraged to tackle online learning in the general combinatorial setting, i.e., when $\mathcal{X} \subseteq \{0,1\}^n$ and $\mathcal{F}$ comprises general functions defined over $\mathcal{X}$.

Rounding Augmented OCO Policy. We begin by stating a "sandwich" property that functions in $\mathcal{F}$ should satisfy, so that the reduction to OCO holds. To do so, we first need to introduce the notion of randomized rounding. Let $\mathcal{Y} \triangleq \text{conv}(\mathcal{X})$ be the convex hull of $\mathcal{X}$. A *randomized rounding* is a random map $\Xi : \mathcal{Y} \to \mathcal{X}$, i.e., a map from a fractional $\mathbf{y} \in \mathcal{Y}$ and, possibly, a source of randomness to an integral variable $\mathbf{x} \in \mathcal{X}$. We assume that the set $\mathcal{F}$ satisfies the following:

Assumption 2. *(Sandwich Property) There exists an $\alpha \in (0,1]$, an $L \in \mathbb{R}_{>0}$, and a randomized rounding $\Xi : \mathcal{Y} \to \mathcal{X}$ such that, for every $f : \mathcal{X} \to \mathbb{R}_{\geq 0} \in \mathcal{F}$ there exists a L-Lipschitz concave function $\tilde{f} : \mathcal{Y} \to \mathbb{R}$ s.t.*

$$\tilde{f}(\mathbf{x}) \geq f(\mathbf{x}), \quad \text{for all } \mathbf{x} \in \mathcal{X}, \text{ and} \tag{6}$$

$$\mathbb{E}_\Xi [f(\Xi(\mathbf{y}))] \geq \alpha \cdot \tilde{f}(\mathbf{y}), \quad \text{for all } \mathbf{y} \in \mathcal{Y}. \tag{7}$$

We refer to $\tilde{f}$ as the *concave relaxation* of f. Intuitively, Asm. 2 postulates the existence of such a concave relaxation $\tilde{f}$ that is not "far" from f: Eqs. (6) and (7) imply that $\tilde{f}$ bounds f both from above and below, up to the approximation factor α. Moreover, the upper bound (Eq. (6)) needs to only hold for integral values, while the lower bound (Eq. (7)) needs to only hold in expectation, under an appropriately-defined randomized rounding Ξ.

Armed with this assumption, we can convert any OCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}}$ operating over $\mathcal{Y} = \text{conv}\,(\mathcal{X})$ to a *Randomized-rounding Augmented OCO* (RAOCO) policy, denoted by $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$, operating over $\mathcal{X}$. This transformation (see Alg. 1) uses both the randomized-rounding Ξ, as well as the concave relaxations $(\tilde{f}_s)_{s=1}^{t-1}$ of the functions $(f_s)_{s=1}^{t-1}$ observed so far. At $t \in [T]$, the RAOCO policy amounts to:

$$\mathbf{y}_t = \mathcal{P}_{\mathcal{Y},t}\left((\mathbf{y}_s)_{s=1}^{t-1}, (\tilde{f}_s)_{s=1}^{t-1}\right), \tag{8a}$$

$$\mathbf{x}_t = \Xi(\mathbf{y}_t) \in \mathcal{X}. \tag{8b}$$

In short, the OCO policy $\mathcal{P}_{\mathcal{Y}}$ is first used to generate a new fractional state $\mathbf{y}_t \in \mathcal{Y}$ by applying $\mathcal{P}_{\mathcal{Y},t}$ to the history of concave relaxations. Then, this fractional decision $\mathbf{y}_t$ is randomly mapped to an integral decision $\mathbf{x}_t \in \mathcal{X}$ according to the rounding scheme Ξ. Then, the reward $f(\mathbf{x}_t)$ is received and f_t is revealed, at which point a concave function $\tilde{f}_t$ is constructed from f_t and added to the history. Our first main result is the following:

Theorem 2. *Under Asm. 2, given an OCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}}$, the RAOCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ described by Alg. 1 satisfies α-$\text{regret}_T (\boldsymbol{\mathcal{P}}_{\mathcal{X}}) \le \alpha \cdot \text{regret}_T (\boldsymbol{\mathcal{P}}_{\mathcal{Y}})$.*

The proof is in Appendix D. As a result, *any regret guarantee obtained by an OCO algorithm over $\mathcal{Y}$, immediately transfers to an α-regret for RAOCO, where α is determined by Asm. 2*. In particular, the decision set $\mathcal{Y}$ is closed, bounded, and convex by construction. Combined with the fact that concave relaxations $\tilde{f}$ are L-Lipschitz (by Asm. 2), Thms. 1 and 2 yield the following corollary:

Corollary 1. *Under Asm. 2, RAOCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ in Alg. 1 equipped with OGA, OMA, or FTRL as OCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}}$ has sublinear α-regret. That is, α-$\text{regret}_T (\boldsymbol{\mathcal{P}}_{\mathcal{X}}) = \mathcal{O}\left(\sqrt{T}\right)$.*

To use this result, Asm. 2 should hold, and both the randomized rounding and the concave relaxations used in RAOCO should be poly-time: all are true for WTP functions optimized over matroid constraints, which we turn to next.

5.2   The Case of Weighted Threshold Potentials

We now consider the case where the decision set is a matroid, and reward functions belong to the class of WTP functions, defined by Eq. (4). We will show that, under appropriate definitions of a randomized rounding and concave relaxations, the class $\mathcal{F}$ satisfies Asm. 2 and, thus, online learning via RAOCO comes with the regret guarantees of Corollary 1. For $\mathcal{Y} = \text{conv}\,(\mathcal{X})$, consider the map $f \mapsto \tilde{f}$ of WTP functions $f : \mathcal{X} \to \mathbb{R}$ to concave relaxations $\tilde{f} : \mathcal{Y} \to \mathbb{R}$ of the form:

$$\tilde{f}(\mathbf{y}) \triangleq f(\mathbf{y}) = \sum_{\ell \in C} c_\ell \min\left\{b_\ell, \sum_{j \in S_\ell} y_j w_{\ell,j}\right\}, \tag{9}$$

for $\mathbf{y} \in \mathcal{Y}$. In other words, *the relaxation of f is itself*: it has the same functional form, allowing integral variables to become fractional.[2] This is clearly concave, as the minimum of affine functions is concave, and the positively weighted sum of concave functions is concave. Finally, all such functions are Lipschitz, with a parameter that depends on c_ℓ, b_ℓ, $\mathbf{w}_\ell$, $\ell \in C$. Let $\tilde{\mathcal{F}}$ be the image of $\mathcal{F}$ under the map (9). We make the following assumption, which is readily satisfied if, e.g., all constituent parameters (c_ℓ, b_ℓ, $\mathbf{w}_\ell$, $\ell \in C$) are uniformly bounded, or the set $\mathcal{F}$ is finite, etc.:

Assumption 3. *There exists an $L > 0$ such that all functions in $\tilde{\mathcal{F}}$ are L-Lipschitz.*

Next, we turn our attention to the randomized rounding Ξ. We can in fact characterize the property that Ξ must satisfy for Asm. 2 to hold for relaxations given by Eq. (9):

---

[2]In Appendix G.IV, we provide an example where the functional form of concave relaxations differs.

Definition 1. *A randomized rounding* $\Xi : \mathcal{Y} \to \mathcal{X}$ *is* negatively correlated *if, for* $\mathbf{x} = \Xi(\mathbf{y}) \in \mathcal{X}$ *(a) the coordinates of* $\mathbf{x}$ *are negatively correlated [3] and (b)* $\mathbb{E}_\Xi [\mathbf{x}] = \mathbf{y}$.

Our next result immediately implies that *any negatively correlated rounding can be used in RAOCO*:

Lemma 1. *Let* $\Xi : \mathcal{Y} \to \mathcal{X}$ *be a negatively correlated randomized rounding, and consider the concave relaxations* $\tilde{f}$ *constructed from* $f \in \mathcal{F}$ *via Eq. (9). Then, if Asm. 3 holds for some* $L > 0$*, the set* $\mathcal{F}$ *satisfies Asm. 2 with* $\alpha = \left(1 - \frac{1}{\Delta}\right)^\Delta$ *where* $\Delta = \sup_{f \in \mathcal{F}} \Delta_f$.

The proof is in Appendix E. As $\Delta \to \infty$, the approximation ratio α approaches $1 - 1/e$ from above, recovering the usual approximation guarantee. However, for finite Δ, we in fact obtain an improved approximation ratio; an example (quadratic submodular functions) is described in Appendix B. Finally, and most importantly, *a negatively-correlated randomized rounding can always be constructed if* $\mathcal{X}$ *is a matroid*. Chekuri et al. [37] provide two polynomial-time randomized rounding algorithms that satisfy this property:

Lemma 2. *(Chekuri et al. [37, Theorem 1.1.]) Given a matroid* $\mathcal{X} \subset \{0,1\}^n$*, let* $\mathbf{y} \in \text{conv}(\mathcal{X})$ *and* Ξ *be either* swap rounding *or* randomized pipage rounding*. Then,* Ξ *is negatively correlated.*

We review swap rounding in Appendix F. Interestingly, the existence of a negatively correlated rounding is inherently linked to matroids: a negatively-correlated rounding exists *if and only if* $\mathcal{X}$ *is a matroid* (see Thm. I.1. in Chekuri et al. [37]). Lemma 1 thus implies that the reduction of RAOCO to OCO policies is also linked to matroids. Putting everything together, Lemmas 1–2 and Corollary 1 yield the following:

Theorem 3. *Let* $\mathcal{X} \subseteq \{0,1\}^n$ *be a matroid, and* $\mathcal{F}$ *be a subset of the WTP class for which Asm. 3 holds. Then, the RAOCO policy* $\mathcal{P}_\mathcal{X}$ *defined by Alg. 1 using swap rounding or randomized pipage rounding as* Ξ *and OGA, OMA, or FTRL as OCO policy* $\mathcal{P}_\mathcal{Y}$*, and concave relaxations in Eq. (9) has sublinear* α*-regret. In particular,* $\alpha\text{-regret}(\mathcal{P}_\mathcal{X}) = \mathcal{O}\left(\sqrt{T}\right)$.

Note that, though all algorithms yield $O(\sqrt{T})$ regret, the dependence of constants on problem parameters (such as n and the matroid rank r), as reported in Table 1, is optimized under OMA (see Appendix C).

5.3   Computational Complexity

OCO Policy. OCO policies are polytime (see, e.g., [8]). Taking gradient-based OCO policies (e.g., OMA in Alg. 2 in the Appendix), their computational complexity is dominated by a projection operation to the convex set $\mathcal{Y} = \text{conv}(\mathcal{X})$. The exact time complexity of this projection depends on $\mathcal{Y}$, however, given a membership oracle that decides $\mathbf{y} \in \text{conv}(\mathcal{X})$, the projection can be computed efficiently in polynomial time [8]. Moreover, the projection problem is a convex problem that can be computed efficiently (e.g., iteratively to an arbitrary precision) and can also be computed in strongly polynomial time [48, Theorem 3].

Concave Relaxations and Randomized Rounding. Concave relaxations are linear in the representation of the function f, but in practice come "for free", once parameters in Eq. (4) are provided. Swap rounding over a general matroid is $\mathcal{O}\left(nr^2\right)$, where r is the rank of the matroid [37]. This dominates the remaining operations (including OCO), and thus determines the overall complexity of our algorithm. This $\mathcal{O}\left(nr^2\right)$ term assumes access to the decomposition of a fractional point $\mathbf{y} \in \mathcal{Y}$ to bases of the matroid. Carathéodory's theorem implies the existence of such decomposition of at most $n+1$ points

---

[3]A set of random variables $x_i \in \{0,1\}, i \in [n]$, are negatively correlated if $\mathbb{E}\left[\prod_{i \in S} x_i\right] \leq \prod_{i \in S} \mathbb{E}[x_i]$ for all $S \subseteq [n]$.

Datasets & Constr.		$F^\star$	t	RAOCO-OGA			RAOCO-OMA				FSF*				TabularGreedy				Random	
				$\bar{F}_\mathcal{X}/F^\star$	std. dev.	η	$\bar{F}_\mathcal{X}/F^\star$	std. dev.	η	γ	$\bar{F}_\mathcal{X}/F^\star$	std. dev.	η	γ	$\bar{F}_\mathcal{X}/F^\star$	std. dev.	η	c_p	$\bar{F}_\mathcal{X}/F^\star$	std. dev.
ZKC (uniform)		0.234	33	0.902	1.85×10^{-2}	2.5	**0.965**	6.02×10^{-3}	10	0.05	0.839	5.02×10^{-3}	75	0.0	0.833	8.12×10^{-3}	160	1	0.642	3.03×10^{-2}
			66	0.924	1.60×10^{-2}		**0.967**	5.65×10^{-3}			0.896	3.83×10^{-3}			0.894	2.55×10^{-3}			0.624	2.80×10^{-2}
			99	0.945	7.70×10^{-3}		**0.982**	5.28×10^{-3}			0.933	4.17×10^{-3}			0.931	1.47×10^{-3}			0.622	1.91×10^{-2}
ZKC (partition)		0.83	33	0.994	2.52×10^{-3}	8	**0.997**	6.95×10^{-4}	10	0.1	✗				0.985	5.65×10^{-3}	10	1	0.953	4.88×10^{-3}
			66	0.99	8.75×10^{-4}		**0.994**	3.48×10^{-4}							0.987	2.93×10^{-3}			0.950	2.87×10^{-3}
			99	0.993	9.66×10^{-4}		**0.997**	3.37×10^{-4}							0.995	2.70×10^{-3}			0.953	1.77×10^{-3}
Epinions (uniform)		0.171	50	0.845	1.61×10^{-2}	4	**0.853**	2.46×10^{-2}	10	0.01	0.703	6.03×10^{-2}	75	0.0	0.694	3.44×10^{-2}	160	1	0.632	2.99×10^{-2}
			100	0.865	1.13×10^{-2}		**0.906**	1.25×10^{-2}			0.776	1.28×10^{-2}			0.768	2.87×10^{-2}			0.615	2.32×10^{-2}
			149	0.88	7.97×10^{-3}		**0.925**	9.09×10^{-3}			0.807	1.52×10^{-2}			0.805	2.48×10^{-2}			0.629	2.36×10^{-2}
Epinions (partition)		0.171	50	0.826	1.35×10^{-2}	3	**0.861**	1.82×10^{-2}	10	0.001	✗				0.720	2.28×10^{-2}	160	1	0.620	6.08×10^{-3}
			100	0.854	6.73×10^{-3}		**0.908**	9.42×10^{-3}							0.786	1.27×10^{-2}			0.619	4.70×10^{-3}
			149	0.88	2.40×10^{-3}		**0.927**	6.34×10^{-3}							0.818	1.3×10^{-2}			0.625	1.10×10^{-2}
MovieLens (uniform)		0.407	98	0.749	3.21×10^{-2}	0.5	**0.792**	2.60×10^{-2}	1.0	0.05	0.681	8.41×10^{-2}	1.0	0.001	0.69	1.11×10^{-2}	160	1	0.748	5.26×10^{-2}
			196	**0.786**	3.82×10^{-2}		0.781	1.36×10^{-2}			0.713	7.44×10^{-2}			0.676	1.03×10^{-2}			0.7	4.76×10^{-2}
			293	0.846	3.98×10^{-2}		**0.866**	1.02×10^{-2}			0.756	6.42×10^{-2}			0.769	1.65×10^{-2}			0.711	3.11×10^{-2}
MovieLens (partition)		0.419	98	0.889	5.20×10^{-2}	0.5	**0.948**	9.70×10^{-3}	10	0.001	✗				0.908	2.39×10^{-2}	160	8	0.829	5.54×10^{-2}
			196	0.872	2.12×10^{-2}		**0.908**	5.14×10^{-3}							0.902	1.98×10^{-2}			0.814	1.51×10^{-2}
			293	0.926	2.34×10^{-2}		0.948	3.13×10^{-3}							**0.964**	2.01×10^{-2}			0.874	2.84×10^{-3}
SynthTF (uniform)		200	33	0.984	7.05×10^{-3}	4	**0.987**	4.03×10^{-3}	0.05	0.1	0.845	4.73×10^{-2}	1	0	0.844	2.21×10^{-2}	1	2	0.605	3.13×10^{-2}
			66	0.994	3.87×10^{-3}		**0.994**	2.83×10^{-3}			0.868	2.14×10^{-2}			0.886	1.88×10^{-2}			0.603	2.01×10^{-2}
			99	0.995	2.01×10^{-3}		**0.998**	2.55×10^{-3}			0.869	2.95×10^{-2}			0.902	2.28×10^{-2}			0.612	1.18×10^{-2}
SynthTF (partition)		400	33	0.98	3.48×10^{-3}	3	**0.983**	2.33×10^{-3}	0.1	0.001	✗				0.834	4.01×10^{-3}	1	1	0.611	9.15×10^{-3}
			66	0.990	1.75×10^{-3}		**0.991**	1.18×10^{-3}							0.852	2.04×10^{-3}			0.607	5.44×10^{-3}
			99	0.993	1.11×10^{-3}		**0.994**	7.75×10^{-4}							0.857	1.36×10^{-3}			0.601	7.71×10^{-3}

Table 2: Average cumulative reward $\bar{F}_\mathcal{X}$ ($t = T/3, 2T/3, T$), normalized by fractional optimal $F^\star$, of integral policies across different datasets and constraints. Optimal hyperparameters $(\eta, \gamma, c_\mathrm{p})$ are reported with each algorithm (see Appendix H for value ranges explored). RAOCO combined with OGA or OMA outperforms competitors almost reaching one, with the exception of MovieLens, where TabularGreedy does better. As Random also performs well on MovieLens, this indicates that the (static) offline optimal is quite poor for this reward sequence. By Property 2, fractional solutions strictly dominate the integral optimal, which implies that in all cases RAOCO outperformed the $1 - 1/e$ approximation, attaining an almost optimal value.

in $\mathcal{X}$; moreover, there exists a decomposition algorithm [49] for general matroids with running time $\mathcal{O}\left(n^6\right)$. However, in all practical cases we consider (including uniform and partition matroids) the complexity is significantly lower. More specifically, for partition matroids, swap rounding reduces to an algorithm with linear time complexity, namely, $\mathcal{O}\left(mn\right)$ for m partitions [50].

6 Extensions

Dynamic Setting. In the dynamic setting, the decision maker compares its performance to the best sequence of decisions $(\mathbf{y}_t^\star)_{t \in [T]}$ with a path length regularity condition [6, 51]. I.e., let $\Lambda_\mathcal{X}(T, P_T) \triangleq \left\{ (\mathbf{x}_t)_{t=1}^T \in \mathcal{X}^T : \sum_{t=1}^T \|\mathbf{x}_{t+1} - \mathbf{x}_t\| \leq P_T \right\} \subset \mathcal{X}^T$ be the set of sequences of decision in a set $\mathcal{X}$ with path length less than $P_T \in \mathbb{R}_{\geq 0}$ over a time horizon T. We extend the definition of the regret (Eq. (5)) to that of *dynamic α-regret*:

$$\alpha\text{-regret}_{T, P_T}(\boldsymbol{\mathcal{P}}_\mathcal{X}) \triangleq \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \max_{(\mathbf{x}_t^\star)_{t=1}^T \in \Lambda_\mathcal{X}(T, P_T)} \alpha \sum_{t=1}^T f_t(\mathbf{x}_t^\star) - \sum_{t=1}^T f_t(\mathbf{x}_t) \right\}.$$

When $\alpha\text{-regret}_{T, P_T}(\boldsymbol{\mathcal{P}}_\mathcal{X})$ is sublinear in T, the policy attains average rewards that asymptotically compete with the optimal decisions *of bounded path length*, in hindsight.

Through our reduction to OCO, we can leverage OGA [6] or meta-learning algorithms over OGA [52] to obtain dynamic regret guarantees in OSM. As an additional technical contribution, we provide the
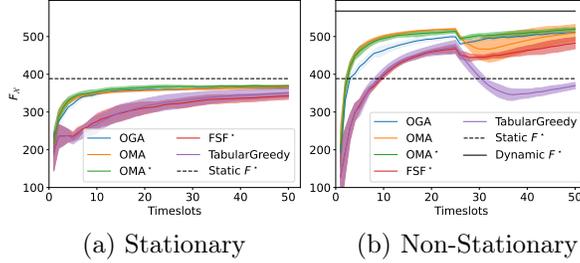
(a) Stationary (b) Non-Stationary

Figure 1: Average cumulative reward $\bar{F}_{\mathcal{X}}$ of the different policies under `SynthWC` dataset under different setups: stationary in (a) and non-stationary in (b). Non-stationarity in (b) is applied by changing the objective at $t = 25$ (see Appendix H). The area depicts the standard deviation over 5 runs.

first sufficient and necessary conditions for OMA to admit a dynamic regret guarantee (see Appendix A). This allows to extend a specific instance of OMA operating on the simplex, the so-called fixed-share algorithm [51, 53], to matroid polytopes. This yields a tighter regret guarantee than OGA [6, 52] (see Theorem 7 in Appendix A). Putting everything together, we get:

Theorem 4. *Under Asm. 2, RAOCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ described by Alg. 1 equipped with an OMA policy in Appendix A as OCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}}$ has sublinear dynamic α-regret, i.e., $\alpha\text{-regret}_{T,P_T}(\boldsymbol{\mathcal{P}}_{\mathcal{X}}) = \mathcal{O}\left(\sqrt{P_T T}\right).$*

This follows from Thm. 2 and the dynamic regret guarantee for OMA in Appendix A.

Optimistic Setting. In the optimistic setting, the decision maker has access to additional information available in the form of predictions: a function $\pi_{t+1} : [0,1]^n \to \mathbb{R}_{\geq 0}$, serving as a prediction of the reward function $f_{t+1}(\mathbf{x})$ is made available before committing to a decision $\mathbf{x}_{t+1}$ at timeslot $t \in [T]$. The prediction π_t encodes prior information available to the decision maker at timeslot t.[4] Let $\mathbf{g}_t$ and $\mathbf{g}_t^\pi$ be supergradients of $\tilde{f}_t$ and π_t at point $\mathbf{y}_t$, respectively. We can define an optimistic OMA policy (see Alg. 3 in Appendix A) that leverages both the prediction and the observed rewards. Applying again our reduction of OSM to this setting we get:

Theorem 5. *Under Asm. 2, RAOCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ in Alg. 1 equipped with OOMA in Appendix A as policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}}$ yields $\alpha\text{-regret}_{T,P_T}(\boldsymbol{\mathcal{P}}_{\mathcal{X}}) = \mathcal{O}\left(\sqrt{P_T \sum_{t=1}^T \|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\infty^2}\right).$*

This theorem follows from Theorem 2 and the optimistic regret guarantee established for OMA policies in Theorem 7 in Appendix A. The optimistic regret guarantee in Theorem 5 shows that the regret of a policy can be reduced to 0 when the predictions are perfect, while providing $\mathcal{O}\left(\sqrt{T}\right)$ guarantee in Thm. 3 when the predictions are arbitrarily bad (with bounded gradients). To the best of our knowledge, ours is the first work to provide guarantees for optimistic OSM.

Bandit Setting. Recall that in the bandit setting the decision maker only has access to the reward $f_t(\mathbf{x}_t)$ after committing to a decision $\mathbf{x}_t \in \mathcal{X}$ at $t \in [T]$; i.e., the reward function is *not* revealed. Our reduction to OCO does not readily apply to the bandit setting; however, we show that the bandit algorithm by Wan et al. [5] can be used to construct such a reduction. The main challenge is to estimate gradients of inputs in $\mathcal{Y}$ only from bandit feedback; this can be done via a perturbation method (see also Hazan and Levy [17]). This approach, described in Appendix G, yields the following theorem:

Theorem 6. *Under bounded submodular monotone rewards and partition matroid constraint sets, LI-RAOCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ in Alg. 5 in Appendix G equipped with an OCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta}$ yields $\alpha\text{-regret}_{T,P_T,W}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq$*

---

[4]Function π_t can extend over fractional values, e.g., be the multi-linear relaxation of a set function.

$W \cdot \mathrm{regret}_{T/W, P_T}(\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta}) + \frac{T}{W} + 2\alpha\delta r^2 nT$, where δ, W, are tuneable parameters of the algorithm and $\mathrm{regret}_{T/W, P_T}(\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta})$ is the regret of an OCO policy executed for T/W timeslots.

Thm. 6 applies to general submodular functions, but is restricted to partition matroids. The theorem also extends to dynamic and optimistic settings: we provide the full description in Appendix G. Our analysis generalizes Wan et al. [5] in that (a) we show that a reduction can be performed to any OCO algorithm, rather than just FTRL, as well as (b) in extending it to the dynamic and optimistic settings (see also Table 4 in the Appendix).

7   Experiments

Datasets and Problem Instances. We consider five different OSM problems: two influence maximization problems, over the `ZKC` [54] and `Epinions` [55] graphs, respectively, a facility location problem over the `MovieLens` dataset [56], a team formation, and a weighted coverage problem over synthetic datasets. Reward functions are generated over a finite horizon and optimized online over both uniform and partition matroid constraints. Details are provided in Appendix H.

Algorithms. We implement the policy in Alg. 1 (`RAOCO`), coupled with `OGA` (`RAOCO-OGA`) and `OMA` (`RAOCO-OMA`) as OCO policies. As competitors, we also implement the *fixed share forecaster* (`FSF`⋆) policy by Matsuoka et al. [3], which only applies to uniform matroids, and the `TabularGreedy` policy by Streeter et al. [1], as well as a `Random` algorithm, which selects a decision u.a.r. from the bases of $\mathcal{X}$. Details and hyperparameters explored are reported in Appendix H. Our code is publicly available.[5]

Metrics. For each setting, we compute $F^\star = \max_{\mathbf{y} \in \mathcal{Y}} \frac{1}{T} \sum_{t=1}^T \tilde{f}(\mathbf{y})$, the value of the optimal fractional solution in hindsight, assuming rewards are replaced by their concave relaxations. Note that this involves solving a convex optimization problem, and overestimates the (intractable) offline optimal, i.e., $F^\star \geq \max_{\mathbf{x} \in \mathcal{X}} \frac{1}{T} \sum_{t=1}^T f(\mathbf{x})$. For each online policy, we compute both the integral and fractional average cumulative reward at different timeslots t, given by $\bar{F}_{\mathcal{X}} = \frac{1}{t} \sum_{s=1}^t f(\mathbf{x}_s)$, $\bar{F}_{\mathcal{Y}} = \frac{1}{t} \sum_{s=1}^t \tilde{f}(\mathbf{y}_s)$, respectively. We repeat each experiment for 5 different random seeds, and use this to report standard deviations.

OSM Policy Comparison. A comparison between the two versions of `RAOCO` (`OGA` and `OMA`) with the three competitors is shown in Table 2. We observe that both `OGA` and `OMA` significantly outperform competitors, with the exception of `MovieLens`, where `TabularGreedy` does better. `OMA` is almost always better than `OGA`. Most importantly, we significantly outperform both `TabularGreedy` and `FSF`⋆ w.r.t. running time, being $2.15 - 250$ times faster (see Appendix H).

Dynamic Regret and Optimistic Learning. Fig. 1 shows the performance of the different policies in a dynamic scenario. All the policies have similar performance in the stationary setting, however in the non-stationary setting only `OGA`, `OMA`, and `FSF`⋆ show robustness. We further experiment with an optimistic setting in Fig. 4 under a non-stationary setting where we provide additional information about future rewards to the optimistic policies, which can leverage this information and yield better results (see Appendix H).

8   Conclusion

We provide a reduction of online maximization of a large class of submodular functions to online convex optimization and show that our reduction extends to many different versions of the online learning problem. There are many possible extensions. As our framework does not directly apply to general

---

[5] https://github.com/neu-spiral/OSMviaOCO

submodular functions, it would be interesting to derive a general method to construct the concave relaxations which are the building block of our reduction in Sec. 5.1. It is also meaningful to investigate the applicability of our reduction to monotone submodular functions with bounded curvature [57], and non-monotone submodular functions [23].

9   Acknowledgements

References

[1] Matthew Streeter, Daniel Golovin, and Andreas Krause. Online Learning of Assignments. In *NeurIPS*, 2009.

[2] Nicholas Harvey, Christopher Liaw, and Tasuku Soma. Improved algorithms for online submodular maximization via first-order regret bounds. *NeurIPS*, 2020.

[3] Tatsuya Matsuoka, Shinji Ito, and Naoto Ohsaka. Tracking regret bounds for online submodular optimization. In *AISTATS*, 2021.

[4] Rad Niazadeh, Negin Golrezaei, Joshua R Wang, Fransisca Susan, and Ashwinkumar Badanidiyuru. Online learning via offline greedy algorithms: Applications in market design and optimization. In *EC*, 2021.

[5] Zongqi Wan, Jialin Zhang, Wei Chen, Xiaoming Sun, and Zhijie Zhang. Bandit Multi-linear DR-Submodular Maximization and Its Applications on Adversarial Submodular Bandits. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, July 2023.

[6] Martin Zinkevich. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. In *ICML*, 2003.

[7] N. Littlestone and M. K. Warmuth. The Weighted Majority Algorithm. *Information and computation*, 1994.

[8] Elad Hazan. Introduction to Online Convex Optimization. *Foundations and Trends® in Optimization*, 2016.

[9] Shai Shalev-Shwartz. Online Learning and Online Convex Optimization. *Foundations and Trends® in Machine Learning*, 2012.

[10] H Brendan McMahan. A survey of algorithms and analysis for adaptive online learning. *JMLR*, 2017.

[11] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.

[12] Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *AISTATS*, 2015.

[13] Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *CDC*, 2016.

[14] Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *COLT*, 2013.

[15] Ofer Dekel, Arthur Flajolet, Nika Haghtalab, and Patrick Jaillet. Online learning with a hint. *NeurIPS*, 2017.

[16] Mehryar Mohri and Scott Yang. Accelerating online convex optimization via adaptive prediction. In *AISTATS*, 2016.

[17] Elad Hazan and Kfir Levy. Bandit convex optimization: Towards tight bounds. *NeurIPS*, 2014.

[18] Robert Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. *NeurIPS*, 2004.

[19] Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online Convex Optimization in the Bandit Setting: Gradient Descent without a Gradient. In *SODA*, 2005.

[20] Mohammad Karimi, Mario Lucic, Hamed Hassani, and Andreas Krause. Stochastic Submodular Maximization: The Case of Coverage Functions. *NeurIPS*, 2017.

[21] Peter Stobbe and Andreas Krause. Efficient Minimization of Decomposable Submodular Functions. *NeurIPS*, 2010.

[22] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD*, 2003.

[23] Andreas Krause and Daniel Golovin. Submodular Function Maximization. In *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press, 2014.

[24] Alan M Frieze. A cost function property for plant location problems. *Mathematical Programming*, 1974.

[25] Stratis Ioannidis and Edmund Yeh. Adaptive caching networks with optimality guarantees. *PER*, 2016.

[26] Yuanyuan Li, Tareq Si Salem, Giovanni Neglia, and Stratis Ioannidis. Online caching networks with adversarial guarantees. *POMACS*, 2021.

[27] Tareq Si Salem, Giovanni Neglia, and Damiano Carra. Ascent Similarity Caching With Approximate Indexes. *ToN*, 2022.

[28] Shinji Ito and Ryohei Fujimaki. Large-scale price optimization via network flow. *NeurIPS*, 2016.

[29] Yao Li, Minhao Cheng, Kevin Fujii, Fushing Hsieh, and Cho-Jui Hsieh. Learning from group comparisons: exploiting higher order interactions. *NeurIPS*, 2018.

[30] Lin Chen, Hamed Hassani, and Amin Karbasi. Online Continuous Submodular Maximization. In *AISTATS*, 2018.

[31] Mingrui Zhang, Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization: From full-information to bandit feedback. *NeurIPS*, 2019.

[32] Qixin Zhang, Zengde Deng, Zaiyi Chen, Haoyuan Hu, and Yu Yang. Stochastic continuous submodular maximization: Boosting via non-oblivious function. In *ICML*, 2022.

[33] Sham M Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. In *STOC*, 2007.

[34] Gruia Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SICOMP*, 2011.

[35] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 1956.

[36] Alexander A Ageev and Maxim I Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 2004.

[37] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, 2010.

[38] Francis Bach. Submodular functions: from discrete to continuous domains. *Mathematical Programming*, 2019.

[39] Andrew An Bian, Baharan Mirzasoleiman, Joachim Buhmann, and Andreas Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *AISTATS*, 2017.

[40] Michel X Goemans and David P Williamson. New $\frac{3}{4}$-approximation algorithms for the maximum satisfiability problem. *SIDMA*, 1994.

[41] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[42] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *FOCS*, 1995.

[43] Sébastien Bubeck. Introduction to online optimization. *Lecture notes*, 2011.

[44] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In *SWAT*, 2004.

[45] Shahar Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *STOC*, 2016.

[46] Dave Buchfuhrer, Shaddin Dughmi, Hu Fu, Robert Kleinberg, Elchanan Mossel, Christos Papadimitriou, Michael Schapira, Yaron Singer, and Chris Umans. Inapproximability for vcg-based combinatorial auctions. In *SODA*, 2010.

[47] Brian W Dolhansky and Jeff A Bilmes. Deep submodular functions: Definitions and learning. *NeurIPS*, 2016.

[48] Swati Gupta, Michel Goemans, and Patrick Jaillet. Solving combinatorial games using products, projections and lexicographically optimal bases. *arXiv preprint arXiv:1603.00522*, 2016.

[49] William H Cunningham. Testing membership in matroid polyhedra. *Journal of Combinatorial Theory, Series B*, 1984.

[50] A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *FOCS*, 2001.

[51] Nicolo Cesa-Bianchi, Pierre Gaillard, Gábor Lugosi, and Gilles Stoltz. Mirror descent meets fixed share (and feels no regret). *NeurIPS*, 2012.

[52] Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. *NeurIPS*, 2020.

[53] Mark Herbster and Manfred K Warmuth. Tracking the best expert. *Machine learning*, 1998.

[54] Wayne W Zachary. An information flow model for conflict and fission in small groups. *Journal of anthropological research*, 1977.

[55] Matthew Richardson, Rakesh Agrawal, and Pedro Domingos. Trust Management for the Semantic Web. In *ISWC*, 2003.

[56] F. Maxwell Harper and Joseph A. Konstan. The MovieLens Datasets: History and Context. *TiiS*, 2015.

[57] Moran Feldman. Guess free maximization of submodular and linear sums. *Algorithmica*, 2021.

[58] Yawei Zhao, Qian Zhao, Xingxing Zhang, En Zhu, Xinwang Liu, and Jianping Yin. Understand Dynamic Regret with Switching Cost for Online Decision Making. *ACM Trans. Intell. Syst. Technol.*, 11(3), apr 2020. ISSN 2157-6904.

[59] Amir Beck and Marc Teboulle. Mirror Descent and Nonlinear Projected Subgradient Methods for Convex Optimization. *Operations Research Letters*, 2003.

[60] Krzysztof C. Kiwiel. Proximal Minimization Methods with Generalized Bregman Functions. *SIAM Journal on Control and Optimization*, 1997.

[61] Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. *NeurIPS*, 2013.

[62] Jacob Goldenberg, Barak Libai, and Eitan Muller. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters*, 2001.

[63] Mark Granovetter. Threshold models of collective behavior. *American journal of sociology*, 1978.

[64] Matthew Streeter and Daniel Golovin. An online algorithm for maximizing submodular functions. *NeurIPS*, 2008.

[65] W. F. King. Analysis of Paging Algorithms. In *Proceedings of the IFIP congress on Information Processing*, 1972.

[66] Philippe Flajolet, Danièle Gardy, and Loÿs Thimonier. Birthday Paradox, Coupon Collectors, Caching Algorithms and Self-Organizing Search. *Discrete Applied Mathematics*, 1992.

[67] Martin Skutella. Convex Quadratic and Semidefinite Programming Relaxations in Scheduling. *J. ACM*, page 206–242, mar 2001.

[68] Zhi-quan Luo, Wing-kin Ma, Anthony Man-cho So, Yinyu Ye, and Shuzhong Zhang. Semidefinite Relaxation of Quadratic Optimization Problems. *IEEE Signal Processing Magazine*, 2010.

[69] Theodoros Lappas, Kun Liu, and Evimaria Terzi. Finding a team of experts in social networks. In *KDD*, 2009.

[70] Bart H Boon and Gerard Sierksma. Team formation: Matching quality supply and quality demand. *EJOR*, 2003.

[71] Francis Bach. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends® in machine learning*, 6(2-3):145–373, 2013.

[72] Kris J Ferreira, Sunanda Parthasarathy, and Shreyas Sekar. Learning to rank an assortment of products. *Management Science*, 68(3):1828–1848, 2022.

[73] Ryan A. Rossi and Nesreen K. Ahmed. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*, 2015.

[74] Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. *Advances in neural information processing systems*, 31, 2018.

[75] Peter Auer, Nicolo Cesa-Bianchi, and Claudio Gentile. Adaptive and self-confident on-line learning algorithms. *JCSS*, 2002.

Appendices

Algorithm 2 Online Mirror Ascent: $\text{OMA}_{\eta,\Phi,\mathcal{Y}}(\mathbf{y}_t, \tilde{f}_t)$

Require: Learning rate $\eta \in \mathbb{R}_{\geq 0}$, mirror map $\Phi : \mathcal{D} \to \mathbb{R}$, decision set $\mathcal{Y}$, decision $\mathbf{y}_t$, reward function
$\tilde{f}_t : \mathcal{Y} \to \mathbb{R}$,

1: $\hat{\mathbf{y}}_t \leftarrow \nabla\Phi(\mathbf{y}_t)$ $\triangleright$ Map the primal point to a dual point
2: $\hat{\mathbf{z}}_{t+1} \leftarrow \hat{\mathbf{y}}_t + \eta\mathbf{g}_t$ $\triangleright$ Take supergradient step in the dual space $(\mathbf{g}_t \in \partial\tilde{f}_t(\mathbf{y}_t))$
3: $\mathbf{z}_{t+1} \leftarrow (\nabla\Phi)^{-1}(\hat{\mathbf{z}}_{t+1})$ $\triangleright$ Map the dual point to a primal point
4: $\mathbf{y}_{t+1} \leftarrow \Pi^{\Phi}_{\mathcal{Y} \cap \mathcal{D}}(\mathbf{z}_{t+1})$ $\triangleright$ Project new point onto feasible region $\mathcal{X}$
5: **return** $\mathbf{y}_{t+1}$
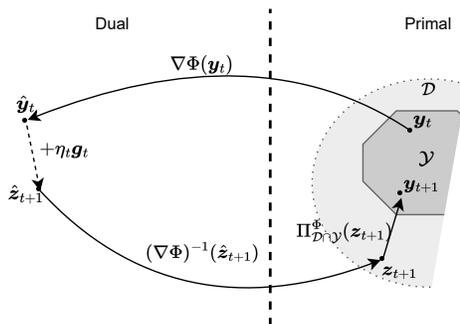
Figure 2: Online Mirror Ascent in Alg. 2.

A   Online Convex Optimization

In this section, we describe two powerful families of OCO policies: online mirror ascent (OMA, in Appendix A.I) and follow the regularized leader (FTRL, in Appendix A.II). Online gradient ascent (OGA) can be described as a special case of OMA, under a Euclidean mirror map. Our exposition below also simultaneously covers the guarantees that come with OMA in the dynamic and optimistic settings; both follow from Theorem 7 (see Appendix G for the derivation). Results for the static setting (e.g., Theorem 1) also follow from Theorem 7 by appropriately restricting the corresponding parameters (see Appendix A.III).

The setting, algorithms, and analysis we present here are a combination of the ones provided by Rakhlin and Sridharan [14], Zinkevich [6], and Zhao et al. [58]. In particular, Rakhlin and Sridharan [14] present OMA in the optimistic setting, while Zinkevich [6] and Zhao et al. [58] studied OGA and OMA, respectively, in the dynamic setting. In the "master theorem" (Theorem 7) that we present below, we combine the proof techniques by all three, to give an algorithm that comes with simultaneous guarantees in all settings (static or dynamic, optimistic or non-optimistic, and combinations thereof).

A.I   Online Mirror Ascent

Online Mirror Ascent [43, Sec. 5.3] is the online version of the mirror ascent (MA) algorithm [59] for convex optimization of a fixed, known function; OMA is a class of policies that generalize OGA. The main premise behind mirror ascent is that variables and supergradients live in two distinct spaces: the primal space, for variables, and the dual space, for supergradients. The two are linked via a function

Algorithm 3 Optimistic Online Mirror Ascent: $\text{OOMA}_{\eta,\Phi,\mathcal{Y}}(\mathbf{y}_t, \tilde{f}_t, \pi_{t+1})$

Require: Learning rate $\eta \in \mathbb{R}_{\geq 0}$, mirror map $\Phi : \mathcal{D} \to \mathbb{R}$, decision set $\mathcal{Y}$, decision $\mathbf{y}_t$, reward function $\tilde{f}_t : \mathcal{Y} \to \mathbb{R}$, prediction function π_{t+1}
  1: $\mathbf{y}_{t+1}^{\pi} \leftarrow \text{OMA}_{\eta,\Phi,\mathcal{Y}}(\mathbf{y}_t, \tilde{f}_t)$ ▷ Adapt secondary decision according to the reward f_t and previous primary decision $\mathbf{y}_t$
  2: $\mathbf{y}_{t+1} \leftarrow \text{OMA}_{\eta,\Phi,\mathcal{Y}}(\mathbf{y}_t^{\pi}, \pi_{t+1})$ ▷ Adapt primary decision according to the precision π_{t+1} of the future reward $\tilde{f}_{t+1}$ and previous secondary decision $\mathbf{y}_t^{\pi}$
  3: **return** $\mathbf{y}_{t+1}$

---

known as a mirror map. Contrary to standard supergradient ascent, updates using the gradient occur on the dual space; the mirror map is used to invert this update to a change on the primal variables. For several constrained optimization problems of interest, mirror ascent leads to faster convergence compared to gradient ascent [43, Sec. 4.3].

The primal and dual spaces are $\mathbb{R}^n$. To disambiguate between the two, we denote primal points by $\mathbf{y}, \mathbf{z} \in \mathbb{R}^n$ and dual points by $\hat{\mathbf{y}}, \hat{\mathbf{z}} \in \mathbb{R}^n$, respectively. Formally, OMD is parameterized by (1) a learning rate $\eta \in \mathbb{R}_+$, and (2) a differentiable map $\Phi : \mathcal{D} \to \mathbb{R}$ where $\mathcal{D}$ is its domain.

Assumption 4. *The map $\Phi : \mathcal{D} \to \mathbb{R}$ satisfies the following properties:*

  1. *The domain $\mathcal{D}$ of Φ is a convex and open set such that the decision set $\mathcal{Y}$ is included in its closure, i.e., $\mathcal{Y} \subseteq \text{closure}(\mathcal{D})$, and their intersection is nonempty $\mathcal{Y} \cap \mathcal{D} \neq \emptyset$.*

  2. *The map Φ is ρ strongly-convex over $\mathcal{D}$ w.r.t. a norm $\|\cdot\|$ and differentiable over $\mathcal{D}$.*

  3. *The map $\nabla\Phi(\mathbf{x}) : \mathcal{D} \to \mathbb{R}^n$ is surjective.*

  4. *The gradient of Φ diverges on the boundary of $\mathcal{D}$, i.e., $\lim_{\mathbf{x} \to \partial\mathcal{D}} \|\nabla\Phi(\mathbf{x})\| = +\infty$, where $\partial\mathcal{D} = \text{closure}(\mathcal{D}) \setminus \mathcal{D}$.*

When the map Φ satisfies the above properties is said to be a *mirror map.*

OMA takes the form described in Alg. 2. Given η and Φ, an OMA iteration proceeds as follows: after receiving the reward $\tilde{f}_t(\mathbf{y}_t)$ and the reward function $\tilde{f}_t(\cdot)$ is revealed, the current state $\mathbf{y}_t$ is first mapped from the primal to the dual space via:

$$\hat{\mathbf{y}}_t = \nabla\Phi(\mathbf{y}_t). \tag{10}$$

Then, a regular supergradient ascent step is performed in the *dual space* to obtain an updated dual point:

$$\hat{\mathbf{z}}_{t+1} = \hat{\mathbf{y}}_t + \eta\mathbf{g}_t, \tag{11}$$

where $\mathbf{g}_t \in \partial\tilde{f}_t(\mathbf{y}_t)$ is the supergradient of $\tilde{f}_t$ at point $\mathbf{y}_t$. This updated dual point is then mapped back to the primal space using the inverse of mapping $\nabla\Phi$, i.e.:

$$\mathbf{z}_{t+1} = (\nabla\Phi)^{-1}(\hat{\mathbf{z}}_{t+1}). \tag{12}$$

The resulting primal point $\mathbf{z}_{t+1}$ may lie outside the decision set $\mathcal{Y}$. To obtain the final feasible point $\mathbf{y}_{t+1} \in \mathcal{Y}$, a projection is made using the Bregman divergence associated with the mirror map Φ. The final result becomes:

$$\mathbf{y}_{t+1} = \Pi_{\mathcal{Y}\cap\mathcal{D}}^{\Phi}(\mathbf{z}_{t+1}), \tag{13}$$

where $\Pi^{\Phi}_{\mathcal{Y}\cap\mathcal{D}}(\cdot)$ is the Bregman projection (15) associated to Φ onto the set $\mathcal{Y}\cap\mathcal{D}$. The steps described Eqs. (10)–(13) define OMA. The illustration in Fig. 2 summarizes these steps. Note that $\mathbf{y}_{t+1}$ is a function of $\{(\mathbf{y}_t, \tilde{f}_t)\} \subseteq \{(\mathbf{y}_s, \tilde{f}_s)\}_{s=1}^{t}$, hence OMA is indeed an OCO policy.

We conclude our description of OMA by providing a formal definition of supergradients and the Bregman projection.

Definition 2. *The superdifferential of a concave function $\tilde{f} : \mathbb{R}^n \to \mathbb{R}$ at point $\mathbf{y} \in \mathbb{R}^n$ is defined as the set $\partial\tilde{f}(\mathbf{y}) \subset \mathbb{R}^n$ of supergradients s.t. every $\mathbf{g}(\mathbf{y}) \in \partial\tilde{f}(\mathbf{y})$ satisfies the following inequality*

$$\tilde{f}(\mathbf{y}) - \tilde{f}(\mathbf{y}') \geq \mathbf{g}(\mathbf{y}) \cdot (\mathbf{y} - \mathbf{y}'), \quad \text{for every } \mathbf{y}' \in \mathbb{R}^n. \tag{14}$$

When the function $\tilde{f}$ is differentiable, then by definition $\partial\tilde{f}(\mathbf{y}) = \{\nabla\tilde{f}(\mathbf{y})\}$, where $\nabla\tilde{f}(\mathbf{y})$ is the gradient of $\tilde{f}$ at point $\mathbf{y} \in \mathbb{R}^n$.

Definition 3. *The Bregman projection [60] associated to a map Φ onto a convex set $\mathcal{S}$ is denoted by $\Pi^{\Phi}_{\mathcal{S}} : \mathbb{R}^n \to \mathcal{S}$, is defined as*

$$\Pi^{\Phi}_{\mathcal{S}}(\mathbf{y}') = \underset{\mathbf{y}\in\mathcal{S}}{\arg\min}\, D_{\Phi}(\mathbf{y}, \mathbf{y}'), \quad \text{where} \quad D_{\Phi}(\mathbf{y}, \mathbf{y}') = \Phi(\mathbf{y}) - \Phi(\mathbf{y}') - \nabla\Phi(\mathbf{y}') \cdot (\mathbf{y} - \mathbf{y}'). \tag{15}$$

OMA considers that the decision set is bounded w.r.t. the Bregman divergence associated to Φ. Formally, we assume the following.

Assumption 5. *Consider a decision set $\mathcal{Y}$ and a mirror map Φ and $\mathbf{y}_1 \in \mathcal{Y}$. The Bregman divergence D_{Φ} associated to Φ is bounded over $\mathcal{Y}$, i.e., there exits $D < \infty$ s.t.*

$$D_{\Phi}(\mathbf{y}, \mathbf{y}_1) \leq D^2, \quad\quad\quad\quad \text{for any } \mathbf{y} \in \mathcal{Y}. \tag{16}$$

Assumption 6. *Consider a decision set $\mathcal{Y}$ and a mirror map Φ. The dual norm $\|\cdot\|_{\star}$ of the gradient of the mirror map Φ is bounded by $L_{\Phi} \in \mathbb{R}_{\geq 0}$, i.e.,*

$$\|\nabla\Phi(\mathbf{y})\|_{\star} \leq L_{\Phi}, \quad\quad\quad\quad \text{for every } \mathbf{y} \in \mathcal{Y}. \tag{17}$$

Note that Asm. 6 is typically not required to establish static regret guarantees (2) (see, e.g., Hazan [8]). However, as we remark in Appendix C, policies with sublinear static regret do not necessarily perform well in dynamic settings, e.g., this the case for OMA configured with the standard negative entropy mirror map [43]. In Theorem 7, we show that Asm. 6 is sufficient to establish sublinear dynamic regret against a comparator sequence (see Sec. 6) with path length $P_T > 0$.

Regret Guarantee. Here, we provide a general regret bound for *Optimistic* OMA in Alg. 3 configured with a general mirror map Φ; note that the same algorithm applies to both the optimistic and dynamic settings, and so does the formal guarantee that we provide below. The non-optimistic regret bounds follow from the optimistic regret bounds by setting the predictions to 0 ($\pi_t = 0$), and Alg. 3 reduces to Alg. 2. Static (i.e., non-dynamic) bounds follow by setting $P_T = 0$.

Theorem 7. *Consider an OCO setting with a convex decision space $\mathcal{Y}$, concave rewards $\tilde{f} \in \hat{\mathcal{F}}$ satisfying Asm. 1 from a set $\hat{\mathcal{F}}$ selected by an adversary, and predictions $\pi_t : \mathcal{Y} \to \mathbb{R}$ for $t \in [T]$. The regret of OOMA in Alg. 3 configured with Φ satisfying Asms. 4–6 against a sequence of decision $(\mathbf{y}^{\star}_t)_{t\in[T]} \in \Lambda_{\mathcal{Y}}(T, P_T)$ is upper bounded as follows*

$$\text{regret}_{T,P_T}(\mathcal{P}_{\mathcal{Y}}) \leq \eta\sum_{t=1}^{T}\|\mathbf{g}_t - \mathbf{g}^{\pi}_t\|^2_{\star}/2\rho + \left(D^2 + 2L_{\Phi}P_T\right)/\eta. \tag{18}$$

18

When the learning rate $\eta_\star = \sqrt{2\rho\left(D^2 + 2L_\Phi P_T\right)/\sum_{t\in[T]}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2}$ is selected, the following holds

$$\text{regret}_{T,P_T}(\mathcal{P}_\mathcal{Y}) \leq \sqrt{2/\rho\left(D^2 + 2L_\Phi P_T\right)\sum_{t=1}^{T}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2} \tag{19}$$

Proof. We adapt the proof of Rakhlin and Sridharan [61, Lemma 1] to support dynamic regret.

$$\sum_{t\in[T]}(\mathbf{g}_t)\cdot(\mathbf{y}_t^\star - \mathbf{y}_t) \tag{20}$$

$$\leq \sum_{t\in[T]}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star\|\mathbf{y}_t - \mathbf{y}_t^\pi\|_\star + \frac{1}{\eta}\left(D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_{t-1}^\pi) - D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_t^\pi) - \frac{\rho}{2}\|\mathbf{y}_t^\pi - \mathbf{y}_t\|^2\right) \tag{21}$$

In order to upper bound the r.h.s. term in the above inequality. We first bound the term $\sum_{t\in[T]}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star\|\mathbf{y}_t - \mathbf{y}_t^\pi\|$, and then, secondly, we bound the term $\sum_{t\in[T]}\left(D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_{t-1}^\pi) - D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_t^\pi)\right)$.

First, note that it holds that $ab = \inf_{\eta'}\frac{\eta'}{2}a^2 + \frac{1}{2\eta'}b^2$, so we have

$$\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star\|\mathbf{y}_t - \mathbf{y}_t^\pi\| = \inf_{\eta'}\left\{\frac{\eta'}{2}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2 + \frac{1}{2\eta'}\|\mathbf{y}_t - \mathbf{y}_t^\pi\|^2\right\} \tag{22}$$

$$\leq \frac{\eta}{2\rho}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2 + \frac{\rho}{2\eta}\|\mathbf{y}_t^\pi - \mathbf{y}_t\|^2 \qquad \text{(for some } \eta/\rho\text{).} \tag{23}$$

Second, we employ the Cauchy-Schwarz inequality and the bounds on D_Φ and $\nabla\Phi$ from the assumptions, so we have the following:

$$\sum_{t\in[T]}\left(D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_{t-1}^\pi) - D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_t^\pi)\right) \tag{24}$$

$$\leq \underbrace{D_\Phi(\mathbf{y}_1^\star, \mathbf{y}_0^\pi)}_{\leq D^2} + \sum_{t\in[T]}\left(D_\Phi(\mathbf{y}_{t+1}^\star, \mathbf{y}_t^\pi) - D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_t^\pi)\right) \tag{25}$$

$$\leq D^2 + \sum_{t\in[T]}\underbrace{\left(\nabla\Phi(\mathbf{y}_{t+1}^\star) - \nabla\Phi(\mathbf{y}_t^\pi)\right)\cdot\left(\mathbf{y}_{t+1}^\star - \mathbf{y}_t^\star\right)}_{\text{Cauchy-Schwarz's Ineq.}} - \underbrace{D_\Phi(\mathbf{y}_t^\star, \mathbf{y}_{t+1}^\star)}_{\geq 0} \tag{26}$$

$$\leq D^2 + \sum_{t\in[T]}\underbrace{\|\nabla\Phi(\mathbf{y}_{t+1}^\star) - \nabla\Phi(\mathbf{y}_t^\pi)\|_\star}_{\leq L_\Phi \ (17)}\|\mathbf{y}_{t+1}^\star - \mathbf{y}_t^\star\| \tag{27}$$

$$\leq D^2 + 2L_\Phi\sum_{t\in[T]}\|\mathbf{y}_{t+1}^\star - \mathbf{y}_t^\star\|. \tag{28}$$

Combine Eqs. (21)–(28) to obtain the final upper bound.

$$\sum_{t\in[T]}\tilde{f}_t(\mathbf{y}_t^\star) - \tilde{f}_t(\mathbf{y}_t) \leq \sum_{t\in[T]}\mathbf{g}_t\cdot\mathbf{y}_t^\star - \mathbf{y}_t \tag{29}$$

$$\leq \frac{1}{\eta}\left(D^2 + 2L_\Phi\underbrace{\sum_{t\in[T]}\|\mathbf{y}_{t+1}^\star - \mathbf{y}_t^\star\|}_{P_T}\right) + \frac{\eta}{2\rho}\sum_{t\in[T]}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2 \tag{30}$$

$$= \frac{1}{\eta}\left(D^2 + 2L_\Phi P_T\right) + \frac{\eta}{2\rho}\sum_{t\in[T]}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2. \tag{31}$$

The learning rate $\eta_\star = \sqrt{2\rho\left(D^2 + 2L_\Phi P_T\right)/\sum_{t\in[T]}\|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2}$ yields the tightest upper bound given by the theorem. $\qquad\square$

Algorithm 4 Follow the Regularized Leader: $\mathrm{FTRL}_{\eta,R,\mathcal{Y}}((\mathbf{y}_s)_{s=1}^t, (\tilde{f}_s)_{s=1}^t)$

Require: Learning rate $\eta \in \mathbb{R}_{\geq 0}$, regularization function $R : \mathcal{X} \to \mathbb{R}$, decision set $\mathcal{Y}$, decisions $\mathbf{y}_s$,
reward functions $\tilde{f}_s : \mathcal{Y} \to \mathbb{R}$ for $s \in [t]$
1: $\mathbf{g}_s \leftarrow \partial_\mathbf{y} \tilde{f}_s(\mathbf{y}_s)$
2: $\mathbf{y}_{t+1} \leftarrow \mathrm{argmin}_{\mathbf{y} \in \mathcal{Y}} \, \eta \sum_{s=1}^t \mathbf{g}_s \cdot \mathbf{y} + R(\mathbf{y})$ ▷ Project new point onto feasible region $\mathcal{X}$
3: **return** $\mathbf{y}_{t+1}$

Note that, to obtain the tightest upper bound on the regret, we assume that when setting the learning rate we have access to quantities ρ, D, L_Φ, P_T, as well as quantity $l_T \triangleq \sqrt{\sum_{t \in [T]} \|\mathbf{g}_t - \mathbf{g}_t^\pi\|_\star^2}$. The first three are problem-specific bounds and are usually easy to determine given the problem instance (see also Appendix C). Parameter P_T is a usual problem input, as it bounds the power of the (dynamic) adversary; assuming that it is known is standard (see, e.g., Zinkevich [6], Besbes et al. [11]). Guarantees however can be still provided even if it is not a priori known using the same meta-learning approach as the one discussed next (see, e.g., Zhao et al. [52]).

Quantity l_T can readily be bounded by $L\sqrt{T}$ in the non-optimistic setting, by Assumption 1. However, to obtain a tighter bound, but also when operating in the optimistic setting, this parameter can be learned through a meta-algorithm (see also Hazan [8], Shalev-Shwartz [9], McMahan [10]). In short, one can execute in parallel multiple instances of OMA algorithms configured for different l_T with an appropriately chosen range and resolution, and frame an expert problem to learn the best expert (policy). The meta problem is a standard prediction with expert advice problem, and can be tackled with well understood parameter-free and computationally efficient learning algorithms [8, 9, 10]. Alternatively, one could consider adaptive online algorithms [9] that adapt to the observed gradients with a slight degradation in performance.

Note that in our implementation and experiments, we perform a grid search for η values (see Sec. 7). We also implemented the meta-learning algorithm for the optimistic setting (see Sec. H), also with a grid search for η.

A.II   Follow the Regularized Leader

It is well known that *follow the leader* (FTL) policy, also known as *fictitious play* in economics, fails to provide sublinear regret guarantee [8]. The FTL policy selects naïvely the decision that would minimize the past costs, i.e., $\mathbf{x}_{t+1} \in \mathrm{argmin}_{\mathbf{y} \in \mathcal{Y}} \sum_{s=1}^t f_s(\mathbf{y})$. This policy fails against an adversary; an adversary can drive the policy to change severely its decision from one timeslot to another. The policy can be modified to exhibit more *stability* [10] through regularization, by the introduction of regularization function $R : \mathcal{X} \to \mathbb{R}$. The update rule is then modified to the following:

$$\mathbf{x}_{t+1} = \mathrm{argmin}_{\mathbf{y} \in \mathcal{Y}} \, \eta \sum_{s=1}^t \mathbf{g}_s \cdot \mathbf{y} + R(\mathbf{y}). \tag{32}$$

We assume that the regularization function satisfies the following properties:

Assumption 7. *The map $R : \mathcal{X} \to \mathbb{R}$ is 1-strongly convex w.r.t. a norm $\|\cdot\|$, smooth, and twice differentiable.*

The pseudocode of the algorithm is provided in Alg. 4.

Regret Guarantee. The regret of the FTRL policy is given by:

Theorem 8. *([8, Theorem 5.2.]) Consider an OCO setting with a convex decision space $\mathcal{Y}$, concave rewards $\tilde{f} \in \hat{\mathcal{F}}$ satisfying Asm. 1 from a set $\hat{\mathcal{F}}$ selected by an adversary. The regret of FTRL in Alg. 4*

against a fixed decision $\mathbf{y}_\star \in \mathcal{Y}$ is upper bounded as follows

$$\sum_{t=1}^{T} \tilde{f}_t(\mathbf{y}^\star) - \sum_{t=1}^{T} \tilde{f}_t(\mathbf{y}_t) \leq 2\eta \sum_{t \in [T]} \|\mathbf{g}_t\|_\star^2 + \frac{R(\mathbf{y}^\star) - R(\mathbf{y}_1)}{\eta}. \tag{33}$$

When the learning rate rate $\eta_\star = \sqrt{\frac{2(R(\mathbf{y}^\star) - R(\mathbf{y}_1))}{\sum_{t \in [T]} \|\mathbf{g}_t\|_\star^2}}$ is selected, the following holds:

$$\sum_{t=1}^{T} \tilde{f}_t(\mathbf{y}^\star) - \sum_{t=1}^{T} \tilde{f}_t(\mathbf{y}_t) \leq 2\sqrt{2\left(R(\mathbf{y}^\star) - R(\mathbf{y}_1)\right) \sum_{t \in [T]} \|\mathbf{g}_t\|_\star^2}. \tag{34}$$

A.III   Derivation of Theorem 1

Assumption 1 implies that the supergradients of the reward functions are bounded. Note that when the reward functions are Lipschitz continuous their supergradients are bounded at any point [9]. Moreover, the diameter of the constraint set is bounded by a constant. Thus, Theorem 7 implies that in the case of OMA in Alg. 2 configured with an appropriately selected mirror map satisfying Asms. 4–6 (e.g. OGA by selecting $\Phi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$) the regret bound is given by

$$\text{regret}_T(\mathcal{P}_\mathcal{Y}) \leq D\sqrt{2/\rho \sum_{t=1}^{T} \|\mathbf{g}_t\|_\star^2} \leq DL\frac{2T}{\rho} = \mathcal{O}\left(\sqrt{T}\right). \tag{35}$$

Furthermore, the regret of FTRL for an appropriately selected regularizer satisfying Asm 7 (e.g., Euclidean regularizer $R(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$) in Alg. 4 is given by

$$\text{regret}_T(\mathcal{P}_\mathcal{Y}) \leq 2\sqrt{2\left(R(\mathbf{y}^\star) - R(\mathbf{y}_1)\right) \sum_{t \in [T]} \|\mathbf{g}_t\|_\star^2} \leq 2L\sqrt{2\left(R(\mathbf{y}^\star) - R(\mathbf{y}_1)\right)T} = \mathcal{O}\left(\sqrt{T}\right). \tag{36}$$

B   Weighted Threshold Potential Functions

B.I   Submodularity and Monotonicity

Consider a threshold potential $\Psi_{b,\mathbf{w},S}(A)$ defined in Eq. (3) for $A \subseteq V = [n]$. We omit the parameters $b, \mathbf{w}, S$ when they are clear from the context. We split our proof into two parts, by first proving submodularity and then monotonicity.

Submodularity. Given a finite set V, the function is said to be submodular if and only if for every $A \subseteq V$ and $i, j \in V \setminus A$ it holds

$$f(A \cup \{i, j\}) - f(A \cup \{j\}) \leq f(A \cup \{i\}) - f(A). \tag{37}$$

We have the following

$$\Psi(A \cup \{i\}) - \Psi(A) = \begin{cases} 0 & \text{for } \sum_{k \in S \cap A} w_k \geq b, \\ \min\left\{b - \sum_{k \in S \cap A} w_k, w_i\right\} & \text{otherwise.} \end{cases} \tag{38}$$

It also holds

$$\Psi(A \cup \{i, j\}) - \Psi(A \cup \{j\}) = \begin{cases} 0 & \text{for } \sum_{k \in S \cap A} w_k \geq b - w_j, \\ \min\left\{b - \sum_{k \in S \cap A} w_k - w_j, w_i\right\} & \text{otherwise.} \end{cases} \tag{39}$$

21

Combine Eqs. (38) and (39) to obtain

$$\Psi(A \cup \{i, j\}) - \Psi(A \cup \{j\}) \le \Psi(A \cup \{i\}) - \Psi(A). \tag{40}$$

Thus, we proved that the set function g is submodular. The function of interest in Eq. (4) can be expressed for every $S \subseteq V$ as

$$f(A) = \sum_{\ell \in C} c_\ell \Psi_{b_\ell, \mathbf{w}_\ell, S_\ell}(A). \tag{41}$$

The function f is a weighted sum with positive weights of submodular function; therefore, from the definition f is also submodular.

Monotonicity. Consider the sets $A \subseteq B \subseteq V$. We have the following

$$\Psi(B) = \min\left\{b, \sum_{k \in S \cap B} w_k\right\} = \min\left\{b, \sum_{k \in S \cap A} w_k + \sum_{k \in S \cap (B \setminus A)} w_k\right\}. \tag{42}$$

Note that $\sum_{k \in S \cap (B \setminus A)} w_k \ge 0$. Thus, it holds

$$\Psi(B) \ge \min\left\{b, \sum_{k \in S \cap A} w_k\right\} = \Psi(A). \tag{43}$$

We conclude the proof by noting that f is a weighted sum with positive weights of monotone functions (43) as expressed in Eq. (41).

B.II   Applications

B.II.1   Weighted Coverage Functions

For $V = [n]$, let $\{S_\ell\}_{\ell \in C}$ be a collection of subsets of V, and $c_\ell \in \mathbb{R}_{\ge 0}$ be a weight associated to each subset S_ℓ. A weighted set coverage function $f : \{0, 1\}^n \to \mathbb{R}_{\ge 0}$ receives as input an $\mathbf{x} \in \{0, 1\}^n$ representing a subset of V and accrues reward c_ℓ if S_ℓ is "covered" by an element in $\mathbf{x}$: formally,

$$f(\mathbf{x}) = \sum_{\ell \in C} c_\ell \left(1 - \prod_{j \in S_\ell}(1 - x_j)\right). \tag{44}$$

Observation 1. *The weighted coverage function f belongs to the class of WTP functions described by Eq. (4); in particular:*

$$f(\mathbf{x}) = \sum_{\ell \in C} c_\ell \Psi_{1, \mathbf{1}, S_\ell}(\mathbf{x}) \qquad \qquad \text{for every } \mathbf{x} \in \{0, 1\}^n. \tag{45}$$

This follows from the simple fact that

$$1 - \prod_{i \in S}(1 - x_i) = \min\left\{1, \sum_{i \in S} x_i\right\}, \qquad \text{for all } \mathbf{x} \in \{0, 1\}^n, S \subseteq V. \tag{46}$$

Classic problems such as influence maximization [22] and facility location [23, 24] can all be expressed via the maximization of weighted coverage functions subject to matroid constraints (see also [20]). Another such example is cache networks [25], which have also been studied in the online setting [26].

Our online setting naturally captures different epidemics over the same graph in influence maximization and dynamic tasks in facility location [20], as well as time-variant content item requests in cache networks [26]. We describe each of these settings in more detail below.

Influence Maximization [22, 20]. In the classic paper by Kempe et al. [22], a network is represented as a graph $G(V, E)$ and a probabilistic model like Independent Cascades [62] or Linear Threshold [63] is used to emulate the propagation of an epidemic over G, starting from a seed set $S \subseteq V$. The objective function is then the total number of nodes in V reached by the epidemic. This is hard to compute so, in practice, the objective is generated by sampling different *reachability graphs* from the underlying probabilistic epidemic model, and estimating the objective in expectation. This estimation is indeed a WC function: in particular, $C = V$, and there is one set S_i for every node i, comprising the nodes that would infect i if they are placed in the seed set, and $c_i = 1$.

The influence maximization setting maps very well to the online submodular maximization problem we consider, particularly in the bandit setting. Different reachability graphs/coverage functions correspond to different instances the epidemic propagation model at each timeslot or, more broadly, different epidemics, potentially with time-varying statistics. The decision maker gets to pick the seeds subject to, e.g., a cardinality/uniform matroid constraint, and observes either the entire reachability graph (in the full information setting), or just the final number of nodes infected (in the bandit setting).

Facility Location [23, 24, 20]. In the classic facility location problem, we are given a complete weighted bipartite graph $G(V \cup V', E)$, where $V = [n]$, $E = V \times V'$, with non-negative weights $w_{v,v'} \geq 0$, $v \in V, v \in V'$. The decision maker selects a subset $S \subset V$ and each $v' \in V$ responds by selecting the $v \in S$ with the highest weight $w_{v,v'}$. The goal is to maximize the average weight of these selected edges, i.e. to maximize

$$f(S) = \frac{1}{|V'|} \sum_{v' \in V} \max_{v \in S} w_{v,v'} \tag{47}$$

given some matroid constraints on S. Set V can be considered a set of facilities and V' a set of customers or tasks to be served, while $w_{v,v'}$ is the utility accrued if facility v serves tasks v'. Karimi et al. [20] note that this is also an instance of the *Exemplar-based Clustering* problem, in which $V = V'$ is a set of objects and $w_{v,v}$ is the similarity (or inverted distance) between objects v and v, and one tries to find a "summary" subset S of exemplars/centroids highly similar to all elements in V. Karimi et al. also show that this objective is a weighted coverage function, of the form (44), because, for $\mathbf{x} \in \{0, 1\}^n$ the characteristic vector of S, we have

$$\max_{v \in S} w_{v,v'} = \sum_{i=1}^{n-1} \left(w_{\pi_j, v'} - w_{\pi_{i+1}, v'} \right) \left(1 - \prod_{j=1}^{i} (1 - x_{\pi_j}) \right) + w_{\pi_n, v'} \left(1 - \prod_{j=1}^{n} (1 - x_{\pi_n}) \right)$$

where $\boldsymbol{\pi} : V \to V$ is a permutation such that: $w_{\pi_1, v'} \geq w_{\pi_2, v'} \geq \ldots \geq w_{\pi_n, v'}$.

This also maps well to our online setting. In the context of facility location, we wish to choose facilities per timeslot in the presence of time-varying, online arrivals of customers/tasks, and in the context of exemplar clustering we wish to choose summaries/exemplars that are good across a potentially time-varying sequence of sets of points. Both fit very well in the full-information setting, as all distances/similarities are immediately observable/computable whenever the tasks/points are revealed.

Cache Networks [25, 26]. Ioannidis and Yeh [25] consider a network of caches, each capable of storing at most a constant number of content items of equal size. Item requests are routed over paths in the graph, and are satisfied upon hitting the first cache that contains the requested item. The objective is to determine a mapping of items to caches that minimizes the aggregate retrieval cost.

Formally, the network is represented as a directed graph $G(V, E)$, with $V = [n]$, Let $\mathcal{C}$ be the set of content items available. Each node $v \in V$ has capacity $c_v \in \mathbb{N}$: exactly c_v content items in $\mathcal{C}$ can be stored in v. Let $x_{vi} \in \{0, 1\}$ the variable indicating whether $v \in V$ stores item $i \in \mathcal{C}$, and let $\mathbf{x} = [x_{vi}]_{v \in V, i \in \mathcal{C}} \in \{0, 1\}^{|V| \times |\mathcal{C}|}$. Note that the capacity constraints imply that $\sum_{i \in \mathcal{C}} x_{vi} \leq c_v$, for all $v \in V$. This set of constraints defines a partition matroid [25]. Each item i in the catalog $\mathcal{C}$ is associated with a fixed set of *designated servers* $\mathcal{S}_i \subseteq V$, that always store i; w.l.o.g., it is assumed that memory used is outside the cache capacity (as it is fixed and outside the optimization problem). A request *request* r is a pair (i, p) where $i \in \mathcal{C}$ is the item requested, and p is the path traversed to serve this request. The terminal node in the path is always a designated source node for i, i.e., if $|p| = K$, $p_K \in S_i$. An incoming request (i, p) is routed over the network G following path p, until it reaches a cache that stores i. The objective is to find an allocation $\mathbf{x}$ that maximizes the *caching gain*, i.e., the reduction in routing costs compared to serving the request from the final designated server. Given request $r = (i, p)$, this is captured by the objective:

$$f(\mathbf{x}) = \sum_{k=1}^{|p|-1} w_{p_{k+1}p_k} \left(1 - \prod_{k'=1}^{k} (1 - x_{p_{k'}i}) \right), \tag{48}$$

where $w_{u,v} \geq 0$ is the cost of transfering an item over edge (u, v). This objective is clearly of the form in Eq. (44).

This problem also maps very well to the online submodular maximization setting: different functions, varying through time, naturally correspond to different requests $r = (i, p)$ for different contents, traversing different paths across the network. Again, the full information (rather than bandit) setting is natural here, as the entire function is revealed once the item and the path are revealed; both are needed at execution time, to identify what is requested and how the request should be routed. The online submodular maximization setting was recently studied by Li et al. [26], who used a distributed variant of the algorithm by Streeter et al. [64] to obtain α-regret guarantees (see Table 1 for a comparison of [64] with RAOCO).

B.II.2 Similarity Caching

Similarity caching [27] is an extension of the so-called paging problem [65, 66]. The objective is to efficiently use local storage memory and provide approximate (similar) answers to a query in order to reduce retrieval costs from some distant server. Formally, consider a catalog of files represented by the set $V = [n]$, and an augmented catalog represented by the set $\hat{V} = [2n]$. Item $i \in V$ is the local copy of a remote file $i + n \in \hat{V}$. Let $c(q, i) \in \mathbb{R}$ be the cost of responding to query $q \in V$ with a similar file $i \in \hat{V}$: this can be determined, e.g., by the similarity or distance between the files, in some appropriately defined feature space. Si Salem et al. [27] show that the reward received under this setting is given by

$$f(\mathbf{x}) = \sum_{l=1}^{C_q} (c(q, \pi_{q,l+1}) - c(q, \pi_{q,l})) \min \left\{ k - |S_l|, \sum_{j \in [l] \setminus S_l} x_{\pi_{q,j}} \right\}, \tag{49}$$

for some $C_q \in \mathbb{N}$, subset $S_l \subset V$ where $|S_l| < k$, and permutation $\boldsymbol{\pi}_q : \hat{V} \to \hat{V}$ satisfying $c(q, \pi_{q,l+1}) - c(q, \pi_{q,l}) \geq 0$ for $l \in [C_q]$. This is indeed a WTP function:

Observation 2. *The similarity caching objective belongs to the class of WTP functions defined by Eq. (4). In particular,* $f(\mathbf{x}) = \sum_{l=1}^{C_q} (c(q, \pi_{q,l+1}) - c(q, \pi_{q,l})) \Psi_{k-|S_l|, \mathbf{1}, S_l}(\mathbf{x})$.

Note that the above functions *cannot* be described as weighted coverage functions (Eq. (44)), and serve as an example of how WTP generalizes the WCF class (see also the discussion in Section B.III below). Moreover, this problem also fits the full information OSM scenario very well: queries q arrive online and, once revealed, the full reward function is determined as well.

B.II.3 Quadratic Set Functions

Motivated by applications such as price optimization on the basis of demand forecasting [28], scheduling [67], detection in multiple-input multiple-output channels in communication systems [68] and team formation [69, 70], we consider quadratic monotone submodular functions $f : \{0,1\}^n \to \mathbb{R}_{\geq 0}$ of the form

$$f(\mathbf{x}) = \mathbf{h}^\mathsf{T}\mathbf{x} + \frac{1}{2}\mathbf{x}^\mathsf{T}\mathbf{H}\mathbf{x}, \tag{50}$$

where, w.l.o.g., $\mathbf{H}$ is symmetric and has zeros in the diagonal. Note that f is monotone and submodular if and only if : (a) $\mathbf{H} \leq \mathbf{0}_{n \times n}$, and (b) $\mathbf{h} + \mathbf{H}\mathbf{x} \geq \mathbf{0}_{n \times 1}$. Thinking of this as the performance of a team, the linear/modular part of the function captures the strength of each team member, and the quadratic part captures the complementarity/overlaps between their skills.

Observation 3. *The quadratic function f belongs to the class of WTP functions* (4), *i.e.,* $f(\mathbf{x}) = \Psi_{\infty, \mathbf{h}+\mathbf{H}\mathbf{1}, [n]}(x) + \sum_{i=1}^{n-1}\sum_{j=i+1}^{n}(-H_{i,j})\Psi_{1,\mathbf{1},\{i,j\}}(\mathbf{x})$, *for every* $\mathbf{x} \in \{0,1\}^n$.

To see this, observe that the following equality holds for all $x, y \in \{0,1\}$

$$1 - (1-x)(1-y) = x + y - xy = \min\{1, x+y\} \tag{51}$$

We have that

$$\mathbf{x}^\mathsf{T}\mathbf{H}\mathbf{x} = \sum_{i=1}^{n}\sum_{j=1}^{n}x_i x_j H_{i,j} \overset{(51)}{=} \sum_{i=1}^{n}\sum_{j=1}^{n}(x_i + x_j - \min\{1, x_i + x_j\})H_{i,j}$$

$$= 2\mathbf{x}^\mathsf{T}\mathbf{H}\mathbf{1} + 2\sum_{i=1}^{n-1}\sum_{j=i+1}^{n}(-H_{i,j})\min\{1, x_i + x_j\}$$

The last equality holds because $\mathbf{H}$ is symmetric and hollow. Substituting the above expression for $\mathbf{x}^\mathsf{T}\mathbf{H}\mathbf{x}$ into $f(\mathbf{x}) = \mathbf{h}^\mathsf{T}\mathbf{x} + \frac{1}{2}\mathbf{x}^\mathsf{T}\mathbf{H}\mathbf{x}$ yields the observation.

This also matches the OSM setting well. For example, different team performance functions can correspond to the score/quality attained by the team at a sequence of progressively revealed tasks. In the bandit setting, only the final performance score is revealed, while in the full information setting, this is decomposed into the individual contributions and their pairwise interactions.

Finally, note that, by definition, the degree Δ_f of quadratic set functions is at most 2. As a result, the α-regret we attain in this setting has $\alpha = \frac{3}{4}$, by Lemma 1.

B.III Generality of the WTP Class

SCMM $\subset$ WTP. Stobbe and Krause [21] defined the class of sums of concave functions composed of non-negative modular functions plus an arbitrary modular function (SCMM) taking the form

$$f(\mathbf{x}) = \mathbf{c} \cdot \mathbf{x} + \sum_{j \in J}\phi_j(\mathbf{w}_j \cdot \mathbf{x}) \tag{52}$$

where $\mathbf{c}, \mathbf{w}_j \in \mathbb{R}^n$ and $\mathbf{0} \leq \mathbf{w}_j \leq \mathbf{1}$ (element-wise) and $\phi_j : [0, \mathbf{w}_j \cdot \mathbf{1}] \to \mathbb{R}$ are arbitrary concave functions. They then show that the SCMM class can be expressed in terms of threshold potentials.

Proposition 1. *(Stobbe and Krause [21]) For $\phi \in C^2([0,M])$,*

$$\phi(v) = \phi(v) + \phi'(M)v - \int_0^M \min\{v, y\}\,\phi''(y)dy \quad \text{for } v \in [0, M]. \tag{53}$$

Discretized versions of Eq. (53) can be efficiently computed when v belongs to a finite set, e.g.,

$$\phi\left(\sum_{i\in S} x_i\right) = \phi(0) + (\phi(|S|) - \phi(|S|-1))\left(\sum_{i\in S} x_i\right)$$
$$+ \sum_{k=1}^{|S|-1}(2\phi(k) - \phi(k-1) - \phi(k+1))\min\left\{k, \sum_{i\in S} x_i\right\}. \tag{54}$$

A Hierarchy of Submodular Functions. Consider now a finite set $V = [n]$, $c_i \in \mathbb{R}_{\geq 0}$, $b_i \in \mathbb{R}_{\geq 0}\cup\{\infty\}$, $w_{i,j} \in \mathbb{R}_{\geq 0}$, and $S_i \subseteq V$ for $i \in C$, $j \in V$. Consider the following classes of submodular functions:

1. Weighted cardinality truncations (WCT) functions which take the form $f(\mathbf{x}) = \sum_{i\in C} c_i \Psi_{b_i,\mathbf{1},S_i}(\mathbf{x})$.

2. Weighted coverage (WC) functions which take the form $f(\mathbf{x}) = \sum_{i\in C} c_i \Psi_{\mathbf{1},\mathbf{1},S_i}(\mathbf{x})$.

3. Facility location (FL) functions. $f(\mathbf{x}) = \sum_{i\in C} \max\{x_j w_{i,j} : j \in [n]\}$. It is a subclass of weighted coverage functions [20].

The following chain relationship between these classes of functions holds:

$$\text{FL} \subset \text{WC} \subset \text{WCT} \subset \text{SCMM} \subset \text{WTP (3)} \subset \text{Submodular-Functions.} \tag{55}$$

The strict inclusions $\text{WC} \subset \text{WCT} \subset \text{SCMM}$ was shown by Dolhansky and Bilmes [47], and the strict inclusion $\text{SCMM} \subset \text{WTP}$ was shown by Stobbe and Krause [21].

C OMA under WTP Functions and Matroid Polytopes and Derivation of Bounds in Table 1

For completeness, in this section, we consider set functions in the class of WTP functions (4), and discuss implementation details (i.e., how OMA is instantiated) over the concave relaxations of functions in this class. We note that, to obtain the most favorable dependence of regret constants on problem parameters, we need to use an OMA policy coupled with a negative entropy mirror map. However, we prove a negative result: OMA with negative entropy has linear dynamic regret when applied "out-of-the-box" (see Proposition 2 in Section C.II). Nevertheless, we are able to overcome this issue by extending the so-called fixed share update rule, originally proposed for optimization over the simplex [51, 53], to matroid polytopes (see Proposition 3).

We provide a summary of notation specific to this section in Sec. C.I, the negative result in Section C.II, the two setups of OMA in Alg. 2: the shifted negative entropy in Sec. C.III and the Euclidean in Sec. C.IV (OGA) setup, the characterization of the supergradient of the concave relaxation of WTP functions in Sec. C.VI, and the derivation of the regret bounds and a comparison with related work in Sec. C.V.

C.I Notation Summary

We denote by $\mathbf{1}_S = (\mathbb{1}(i \in S))_{i\in V} \in \{0,1\}^n$ the characteristic vector of a subset $S \subseteq V$ for $\mathbb{1}(\chi) \in \{0,1\}$ takes the value 1 when χ is true, and 0 otherwise. We consider the convex set $\mathcal{Y}$ as the matroid basis polytope of some matroid $\mathcal{M}(V,\mathcal{I})$ with rank $r \in \mathbb{N}$, i.e., the set $\mathcal{Y}$ is given by

$$\mathcal{Y} = \text{conv}\left(\{\mathbf{1}_S : S \in \mathcal{I}\}\right) \subset \{\mathbf{y} \in [0,1]^n : \|\mathbf{y}\|_1 = r\}. \tag{56}$$

We denote by d_∞ the upper bound on discrete derivatives of f, which is given by

$$d_\infty \triangleq \max_{i\in[n]}\{f(\{i\}) - f(\emptyset)\} = \max_{i\in[n]}\{f(\{i\})\}. \tag{57}$$

26

C.II   A Negative Result

OMA configured with the negative entropy mirror map $\Phi(\mathbf{y}) = \sum_{y \in [n]} y_i \log(y_i)$ which enjoys sublinear static regret over the simplex $\mathcal{Y} = \Delta_n \subset [0,1]^n$ does not maintain such a guarantee for dynamic regret. We prove this negative result in Proposition 2 through a simple construction; since $L_\Phi = \infty$ in this setting, then Asm. 6 is also necessary for OMA to yield a sublinear dynamic regret policy. Formally:

Proposition 2. *OMA in Alg. 2 configured with negative entropy mirror map* $\Phi(\mathbf{y}) = \sum_{y \in [n]} y_i \log(y_i)$ *has linear dynamic regret* ($P_T > 0$) *under simplex constraint set* $\mathcal{Y} = \Delta_n \subset [0,1]^n$ *for any* $\eta \in \mathbb{R}_{>0}$.

Proof. Consider $n = 2$, and the following rewards $f_t(\mathbf{x}) = x_1$ for $t \leq T/2$, and $f_t(\mathbf{x}) = x_2$ for $T/2 \leq t \leq T$. The dynamic optimum selects $\mathbf{x}_t = (1, 0)$ for $t \leq T/2$, and $\mathbf{x}_t = (0, 1)$ for $T/2 \leq t \leq T$ incurring a total reward of T. It is easy to verify that the policy's state is given by $\mathbf{x}_t = \left(e^{(t-1)\eta}/(e^{(t-1)\eta} + 1), 1/(e^{(t-1)\eta} + 1) \right)$ for $t \leq T/2$ and $\mathbf{x}_t = \left(e^{(T/2-1)\eta}/(e^{(t-1-T/2)\eta} + e^{(T/2-1)\eta}), e^{(t-1-T/2)\eta}/(e^{(t-1-T/2)\eta} + e^{(T/2-1)\eta}) \right)$ for $T/2 + 1 \leq t \leq T$. Thus, the regret is given by

$$\text{regret}_{T,2}(\boldsymbol{\mathcal{P}}_{\Delta_2}) = T - \sum_{t=1}^{\frac{T}{2}} \underbrace{\frac{e^{(t-1)\eta}}{e^{(t-1)\eta} + 1}}_{\leq 1} + \underbrace{\frac{e^{(t-1)\eta}}{e^{(\frac{T}{2}-1)\eta} + e^{(t-1)\eta}}}_{\leq 1/2} \geq T/4 = \Omega(T). \tag{58}$$

This concludes the proof. $\qquad\square$

We overcome this negative result next, by extending the so-called fixed share update rule, originally proposed for optimization over the simplex [51, 53], to matroid polytopes (see Proposition 3).

C.III   Shifted Negative Entropy Setup

Let $\gamma \in [0, \sqrt{e^{-2} + 1/4} - 1/2]$ be a shifting parameter where $\sqrt{e^{-2} + 1/4} - 1/2 \approx 0.12$. The shifted negative entropy setup of OMA, defined as

$$\Phi(\mathbf{y}) \triangleq \sum_{i \in [n]} (y_i + \gamma) \log(y_i + \gamma) \qquad\qquad \text{for } \mathbf{y} \in \mathcal{D} \triangleq (-\gamma, \infty)^n. \tag{59}$$

Note that it holds $\nabla\Phi(\mathbf{y}) = (1 + \log(y_i + \gamma))_{i \in [n]}$ and $(\nabla\Phi)^{-1}(\hat{\mathbf{y}}) = (\exp(\hat{y}_i - 1) - \gamma)$ for $\mathbf{y} \in \mathcal{D}$ and $\hat{\mathbf{y}} \in \mathbb{R}^n$. Thus, at timeslot t, the steps in lines 2–4 in Alg. 2 correspond to the following update rule

$$z_{t+1,i} = y_{t,i} e^{\eta_t g_{t,i}} + \gamma (e^{\eta_t g_{t,i}} - 1) \qquad\qquad \text{for } i \in [n]. \tag{60}$$

Note that this update rule has the same form as the fixed share update rule [51, 53]. However, this policy extends beyond the simplex.

It is easy to see why the *pure* multiplicative update rule obtained by OMA in Alg. 2 configured with the negative-entropy mirror map fails to provide sublinear dynamic regret guarantees through the constructed scenario in the proof of Proposition 2. The adversary can drive the state of the algorithm to arbitrary close to $(1, 0)$, but when the best decision changes to $(0, 1)$ the algorithm fails to adapt quickly purely from a multiplicative update. This is why the additive term controlled by γ is needed, which also allows us to control $L_\Phi \leq \log(1/\gamma)$.

We summarize the properties satisfied by this setup in the following proposition.

Proposition 3. *Consider a matroid basis polytope* $\mathcal{Y}$ *with rank* r *defined in Sec. C and* $\mathbf{y}_1 = \Pi^\Phi_{\mathcal{Y} \cap \mathcal{D}}(r/n\mathbf{1})$. *Then, Asms. 4–5 are satisfied under the mirror map defined in Eq. (59)* ($\gamma \in [0, \sqrt{e^{-2} + 1/4} - 1/2]$) *under the bounds in Table 3.*

Proof. Strong Convexity. Since Φ is twice differentiable, it is sufficient [9] to show that $\mathbf{y}^\top \nabla^2 \Phi(\mathbf{w}) \mathbf{y} \geq \frac{1}{r+\gamma n} \|\mathbf{y}\|_1^2$ for every $\mathbf{w} \in \mathcal{Y}$ and $\mathbf{y} \in \mathbb{R}^n$, so that Φ is $\frac{1}{r+\gamma n}$ strongly convex w.r.t. the norm $\|\cdot\|_1$ over the set $\mathcal{Y}$.

$$\mathbf{y}^\top \nabla^2 \Phi(\mathbf{w}) \mathbf{y} = \sum_{i \in [n]} \frac{y_i^2}{w_i + \gamma} = \frac{1}{\|\mathbf{w}+\gamma\mathbf{1}\|_1} \left(\sum_{i\in[n]} w_i + \gamma \right) \left(\sum_{i\in[n]} \frac{y_i^2}{w_i+\gamma} \right) \tag{61}$$

$$\geq \frac{1}{\|\mathbf{w}+\gamma\mathbf{1}\|_1} \left(\sum_{i\in[n]} \sqrt{w_i+\gamma} \frac{y_i}{\sqrt{w_i+\gamma}} \right)^2 \geq \frac{1}{r+\gamma n} \|\mathbf{y}\|_1^2. \tag{62}$$

The inequality follows from Cauchy–Schwarz inequality.

Bregman Divergence Bound. Consider a point $\mathbf{y}_1 = \Pi_{\mathcal{Y}\cap\mathcal{D}}^\Phi(r/n\mathbf{1}) \in \mathcal{Y}$, thus for $\mathbf{y} \in \mathcal{Y}$ the Bregman divergence $D_\Phi(\mathbf{y}, \mathbf{y}_1)$ is given by

$$D_\Phi(\mathbf{y}, \mathbf{y}_1) \leq D_\Phi(\mathbf{y}, r/n\mathbf{1}) = \sum_{i\in[n]} (y_i + \gamma) \log \left(\frac{y_i+\gamma}{r/n+\gamma} \right) + r - y_i \leq \sum_{i\in[n]} (y_i+\gamma) \log \left(\frac{y_i+\gamma}{r/n+\gamma} \right)$$

$$\leq r(1+\gamma) \log \left(\frac{(1+\gamma)n}{r+\gamma n} \right). \tag{63}$$

The upper bound is obtained by setting $\mathbf{y} \in \mathcal{Y} \cap \{0,1\}^n$.

Mirror Map's Lipschitzness. The gradient of the mirror map Φ is upper bounded by $L_\Phi = \log(1/\gamma) - 1$ under the $\|\cdot\|_\infty$ norm.

It is straightforward to check that all the properties in Asm 4 are satisfied. $\qquad\square$

Corollary 2. *Under WTP reward functions in Eq. (4) and concave functions in Eq. (9), RAOCO policy in Alg. 1 $\boldsymbol{\mathcal{P}_X}$ equipped with an OOMA policy in Alg. 3 configured with shifted negative entropy mirror map (59) for $\gamma \in [0, \sqrt{e^{-2}+1/4}-1/2]$ and a fixed learning rate $\eta = \sqrt{\frac{2\left(r(1+\gamma)\log\left(\frac{(1+\gamma)n}{r+\gamma n}\right)+2(\log(1/\gamma)-1)P_T\right)}{r\sum_{t\in[T]}\|\mathbf{g}_t-\mathbf{g}_t^\pi\|_\star^2}}$ has the following α-regret*

$$\alpha-\text{regret}(\boldsymbol{\mathcal{P}_X}) \leq \sqrt{2(r+\gamma n)\left(r(1+\gamma)\log\left(\frac{(1+\gamma)n}{r+\gamma n}\right)+2(\log(1/\gamma)-1)P_T\right)\sum_{t\in[T]}\|\mathbf{g}_t-\mathbf{g}_t^\pi\|_\infty^2}. \tag{64}$$

The corollary follows directly from Theorem 7 and Proposition 3. Note that when $P_T = 0$ and $\gamma = 0$, the regret is $\alpha-\text{regret}(\boldsymbol{\mathcal{P}_X}) = \mathcal{O}\left(r\sqrt{\log(n/r)T}\right)$. When $P_T > 0$ and $\gamma = 1/n$, the regret is $\alpha-\text{regret}(\boldsymbol{\mathcal{P}_X}) = \mathcal{O}\left(\sqrt{r(r\log(n/r)+\log(n)P_T)T}\right)$.

C.IV   Euclidean Setup

The Euclidean setup is defined for the squared Euclidean norm mirror map, i.e.,

$$\Phi(\mathbf{y}) \triangleq \frac{1}{2}\|\mathbf{y}\|_2^2 \qquad\qquad\qquad \text{for } \mathbf{y} \in \mathcal{D} \triangleq \mathbb{R}^n. \tag{65}$$

Note that it holds $\nabla\Phi(\mathbf{y}) = \mathbf{y}$ and $(\nabla\Phi)^{-1}(\hat{\mathbf{y}}) = \hat{\mathbf{y}}$ for $\mathbf{y} \in \mathcal{D}$ and $\hat{\mathbf{y}} \in \mathbb{R}^n$. Thus, at timeslot t, the steps in lines 2–4 in Alg. 2 correspond to the Online Gradient Ascent update rule

$$\mathbf{z}_{t+1} = \mathbf{y}_t + \eta\mathbf{g}_t \tag{66}$$

The Bregman projection step in line 5 Alg. 2 corresponds to the Euclidean projection.

$$\mathbf{y}_{t+1} = \Pi_{\mathcal{Y}\cap\mathcal{D}}^{\Phi}(\mathbf{z}_{t+1}) = \Pi_{\mathcal{Y}}^{\Phi}(\mathbf{z}_{t+1}) = \mathrm{argmin}_{\mathbf{y}\in\mathcal{Y}}\|\mathbf{z}_{t+1} - \mathbf{y}\|_2^2. \tag{67}$$

We summarize the properties satisfied by this setup in the following proposition.

Proposition 4. *Consider a matroid basis polytope $\mathcal{Y}$ with rank r defined in Sec. C and $\mathbf{y}_1 = \Pi_{\mathcal{Y}}^{\Phi}(\mathbf{0})$. Then, Asms. 4–6 are satisfied for the mirror map defined in Eq. (65) under the bounds in Table 3.*

Proof. Under this setup the Bregman divergence is given by $D_{\Phi}(\mathbf{y},\mathbf{y}) = \frac{1}{2}\|\mathbf{y}'-\mathbf{y}\|_2^2$ and thus Φ is 1-strongly convex w.r.t. $\|\cdot\|_2$ (from the definition of strong convexity). The upper bound on $D_{\Phi}(\mathbf{y},\mathbf{y}_1)$ for $\mathbf{y}_1 = \Pi_{\mathcal{Y}}^{\Phi}(\mathbf{0})$ is given by $\frac{1}{2}\|\mathbf{y}\|^2 \leq \frac{r}{2}$. The gradient of the mirror map Φ is upper bounded by $L_{\Phi} = \max_{\mathbf{y}\in\mathcal{Y}}\|\mathbf{y}\|_{\infty} = 1$. It is straightforward to check that all the properties in Asm 4 are satisfied. $\square$

Corollary 3. *Under WTP reward functions in Eq. (4) and concave functions in Eq. (9), RAOCO policy in Alg. 1 $\mathcal{P}_{\mathcal{X}}$ equipped with an OMA policy in Alg. 2 configured with Euclidean mirrr map (65) and a fixed learning rate $\eta = \sqrt{\frac{r+2P_T}{\sum_{t\in[T]}\|\mathbf{g}_t-\mathbf{g}_t^{\pi}\|_2^2}}$ has the following α-regret*

$$\alpha-\mathrm{regret}(\mathcal{P}_{\mathcal{X}}) \leq \sqrt{(r+4P_T)\sum_{t=1}^{T}\|\mathbf{g}_t - \mathbf{g}_t^{\pi}\|_2^2}. \tag{68}$$

The corollary follows directly from Theorem 7 and Proposition 3.

Note that all the OCO regret results in this paper follow from Corollaries 2 and 3. In particular, the static non-optimistic regret guarantees are obtained for $P_T = 0$ and $\mathbf{g}_t^{\pi} = \mathbf{0}$. The dynamic non-optimistic regret guarantees are obtained for $P_T > 0$ and $\mathbf{g}_t^{\pi} = \mathbf{0}$. Finally, the non-dynamic optimistic regret guarantees are obtained for $P_T = 0$ and $\mathbf{g}_t^{\pi} \neq 0$.

C.V Supergradient Computation

We conclude our description of the different setups of OMA by providing an explicit expression of the supergradient of the concave relaxation of WTP functions.

Proposition 5. *A supergradient of $\tilde{f} : [0,1]^n \to \mathbb{R}_{\geq 0}$ in Eq. (9) at point $\mathbf{y} \in [0,1]^n$ is given as*

$$\mathbf{g}(\mathbf{y}) \triangleq \left(\sum_{\ell\in C} c_l w_{\ell,j}\mathbb{1}\left(j\in S_{\ell} \wedge \sum_{j'\in S_{\ell}} y_{j'}w_{j'} \leq b_{\ell}\right)\right)_{j\in V}, \tag{69}$$

where $\wedge$ is a logical and *operator.*

Proof. Consider the concave relaxation $\tilde{\Psi}$ of a threshold potential Ψ in Eq. (3) given by

$$\tilde{\Psi}_{b,\mathbf{w},S}(\mathbf{y}) \triangleq \min\left\{b, \sum_{j\in S} w_j y_j\right\}, \qquad \text{for } \mathbf{y} \in [0,1]^n. \tag{70}$$

We omit the parameters $b, \mathbf{w}, S$ when they are clear from the context. We verify that the vector $\mathbf{g}'(\mathbf{y}) = \left(w_j\mathbb{1}\left(j\in S \wedge \sum_{j'\in S} y_{j'}w_{j'} \leq b\right)\right)_{j\in V}$ is a supergradient of $\tilde{\Psi}(\cdot)$ at point $\mathbf{y}$, i.e., it holds (14)

$$\tilde{\Psi}(\mathbf{y}) - \tilde{\Psi}(\mathbf{y}') \geq \mathbf{g}'(\mathbf{y})\cdot(\mathbf{y}-\mathbf{y}'), \quad \text{for every } \mathbf{y}' \in \mathbb{R}^n. \tag{71}$$

Mirror Map Φ	Domain $\mathcal{D}$	Dual-primal update step	Str. convexity ρ	Bregman div. bound D^2	Lip. const. L_Φ
$\sum_{i=1}^n (y_i+\gamma)\log(y_i+\gamma)$	$(-\gamma,\infty)^n$	$z_{t+1,i} = (y_{t,i}+\gamma)e^{\eta_t g_{t,i}} - \gamma, i\in[n]$	$1/(r+\gamma n)$	$r(1+\gamma)\log\left(\frac{(1+\gamma)n}{r+\gamma n}\right)$	$\log(1/\gamma)-1$
$\sum_{i=1}^n y_i\log(y_i)$	$\mathbb{R}_{>0}^n$	$z_{t+1,i} = y_{t,i}e^{\eta_t g_{t,i}}, i\in[n]$	$1/r$	$r\log\left(\frac{n}{r}\right)$	∞
$\frac{1}{2}\|\mathbf{y}\|_2^2$	$\mathbb{R}^n$	$\mathbf{z}_{t+1} = \mathbf{y}_t + \eta_t \mathbf{g}_t$	1	$\min\{r, n-r\}/2$	$1/2$

Table 3: Properties of the different mirror map choices of OMA in Alg. 2.

Consider $\sum_{j\in S} w_j y_j > b$, then it holds

$$\tilde{\Psi}(\mathbf{y}) - \tilde{\Psi}(\mathbf{y}') = b - \min\left\{b, \sum_{j\in S} w_k j_j'\right\} \geq 0. \tag{72}$$

$$\mathbf{g}'(\mathbf{y})\cdot(\mathbf{y}-\mathbf{y}') = \sum_{j\in V} g_j(\mathbf{y})(y_j - y_j') = \sum_{j\in S} w_j(y_j - y_j')\mathbb{1}\left(\sum_{j'\in S} w_{j'} y_{j'} \leq b\right) = 0 \tag{73}$$

$$\leq \tilde{\Psi}(\mathbf{y}) - \tilde{\Psi}(\mathbf{y}'). \tag{74}$$

Thus, when $\sum_{j\in S} w_j y_j > b$ the inequality in Eq. (71) holds. It remains to check if it holds when $\sum_{j\in S} w_j y_j \leq b$. We have

$$\mathbf{g}'(\mathbf{y})\cdot(\mathbf{y}-\mathbf{y}') = \sum_{j\in V} g_j(\mathbf{y})(y_j - y_j') = \sum_{j\in S} w_j(y_j - y_j') = \sum_{j\in S} w_j y_j - \sum_{j\in S} w_j y_j' \tag{75}$$

$$\leq \sum_{j\in S} w_j y_j - \min\left\{b, \sum_{j\in S} w_j y_j'\right\} \leq \min\left\{b, \sum_{j\in S} w_j y_j\right\} - \min\left\{b, \sum_{j\in S} w_j y_j'\right\} \tag{76}$$

$$= \tilde{\Psi}(\mathbf{y}) - \tilde{\Psi}(\mathbf{y}'). \tag{77}$$

The vector $\mathbf{g}'(\mathbf{y})$ is a supergradient of $\tilde{\Psi}(\cdot)$ at point $\mathbf{y}$. We conclude the proof that by noting that $\tilde{f}(\mathbf{y}) = \sum_{\ell\in C} c_\ell \tilde{\Psi}_{b_\ell, \mathbf{w}_\ell, S_\ell}(\mathbf{y})$, and it is straightforward to check that $\partial\left(\sum_{\ell\in C} c_\ell \tilde{\Psi}_{b_\ell, \mathbf{w}_\ell, S_\ell}(\mathbf{y})\right) = \sum_{\ell\in C} c_\ell \partial\left(\tilde{\Psi}_{b_\ell, \mathbf{w}_\ell, S_\ell}(\mathbf{y})\right)$ from the definition of superdifferentials. $\qquad\square$

C.VI Regret Bounds Derivation and Comparison with Related Work

In this section, we first discuss the derivation of the regret bounds attained by this work in Table 1. Then, we discuss algorithm-specific parameters of related work.

Regret Bounds. We discuss the regret bound appearing in Table 1. We obtain the tightest regret bounds by employing the negative entropy setup of OMA in Sec. C, allowing us to optimize the regret constants. In particular, the regret of OMA when configured with the negative-entropy mirror map grows logarithmically w.r.t. the problem dimension (i.e., $\mathcal{O}\left(r\sqrt{\log(n/r)T}\right)$ in Corollary 2) instead of sublinearly as in the case of OGA (e.g., $\mathcal{O}\left(\sqrt{rnT}\right)$ in Corollary 3). We note that the related work assumes that the reward functions are bounded by 1. We make the same assumption to better compare our bounds in Table 3. However, note that our bounds are further scaled by $d_\infty = \max_{i\in[n]}\{f(\{i\}) - f(\emptyset)\} = \max_{i\in[n]}\{f(\{i\})\} \in [1/n, 1]$ which can further improve our bounds under the static and dynamic full-information settings.

Algorithm-specific Parameters in Related Work. Niazadeh et al. [4] require a Blackwell approachability policy. We denote by O_a its time complexity. Streeter et al. [1] require a color palette size

parameter $c_{\mathrm{p}} \in \mathbb{N}$. This parameter creates c_{p} expert algorithms (OCO algorithms operating over the simplex) for a single slot, e.g., over partition matroids the algorithm instantiates rc_{p} experts. Harvey et al. [2] require a discretization parameter $\epsilon > 0$ to convert their continuous-time algorithm to a discrete algorithm, an oracle that gives the multi-linear extension with time-complexity O_{m}, and swap rounding F. Works that operate on continuous DR-submodular objectives [30, 31, 32] require OCO policies with a time complexity O_{oco}; moreover, these algorithms operate over the multilinear extension with time complexity O_{m}. The decisions can be rounded through a randomized-rounding algorithm to operate over matroids (e.g., swap rounding F). We assume that these algorithms have access to a decomposition oracle; then the rounding step has an additional time complexity of $\mathcal{O}\left(r^2 n\right)$. Kakade et al. [33] operate on Linearly Weighted Decomposable (LWD) functions, i.e., reward functions of the form $f_t(\mathbf{x}) = \mathbf{\Phi}(\mathbf{x}) \cdot \mathbf{w}_t$, where $\mathbf{\Phi}(\mathbf{x}) \in \mathbb{R}^s$ is a vector-valued set function known to the decision maker, and $\mathbf{w}_t \in \mathbb{R}^s$ are time-varying vector-valued weights selected by an adversary for some $s \in \mathbb{N}$. This algorithm requires access to an α-approximation oracle with time-complexity O_α.

D   Proof of Theorem 2

Proof. Consider the sequence of reward functions $\{f_1, f_2, \ldots, f_T\} \in \mathcal{F}^T$ and the associated sequence of concave relaxations $\left\{\tilde{f}_1, \tilde{f}_2, \ldots, \tilde{f}_T\right\} \in \hat{\mathcal{F}}^T$. Then,

$$\max_{\mathbf{x} \in \mathcal{X}} \sum_{t \in [T]} f_t(\mathbf{x}) \leq \max_{\mathbf{x} \in \mathcal{X}} \sum_{t \in [T]} \tilde{f}_t(\mathbf{x}) \leq \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t \in [T]} \tilde{f}_t(\mathbf{y}). \tag{78}$$

The first inequality is implied by Eq. (6) in Assumption 2, and the second inequality holds because maximizing over a superset $\mathcal{Y} = \mathrm{conv}\left(\mathcal{X}\right) \supseteq \mathcal{X}$ can only increase the objective value attained. The total expected reward obtained by the RAOCO policy is given by

$$\mathbb{E}_{\Xi}\left[\sum_{t \in [T]} f_t(\mathbf{x}_t)\right] = \sum_{t \in [T]} \mathbb{E}_{\Xi}\left[f_t(\mathbf{x}_t)\right] = \sum_{t \in [T]} \mathbb{E}_{\Xi}\left[f_t(\Xi(\mathbf{y}_t))\right] \overset{(7)}{\geq} \alpha \sum_{t \in [T]} \tilde{f}_t(\mathbf{y}_t). \tag{79}$$

From Eqs. (78) and (79), it follows that:

$$\alpha \max_{\mathbf{x} \in \mathcal{X}} \sum_{t \in [T]} f_t(\mathbf{x}) - \mathbb{E}_{\Xi}\left[\sum_{t \in [T]} f_t(\mathbf{x}_t)\right] \leq \alpha \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t \in [T]} \tilde{f}_t(\mathbf{y}) - \alpha \sum_{t \in [T]} \tilde{f}_t(\mathbf{y}_t). \tag{80}$$

We conclude the proof by noting that the above inequality holds for arbitrary sequences of reward functions $\{f_1, f_2, \ldots, f_T\} \in \mathcal{F}^T$. $\square$

E   Proof of Lemma 1

We begin by proving a series of auxiliary lemmas.

Lemma 3. *Consider $n \in \mathbb{N}$, $\mathbf{y} \in [0,1]^n$, $b \in \mathbb{R}_{>0}$, and $\mathbf{w} \in [0,b]^n$. The following holds*

$$\min\left\{b, \sum_{i \in V} y_i w_i\right\} \geq b - b \prod_{i \in V}(1 - y_i w_i/b). \tag{81}$$

Proof. We define $A_n \triangleq b - b \prod_{i \in V} (1 - y_i w_i / b)$ and $B_n \triangleq \min \left\{ b, \sum_{i \in V} y_i w_i \right\}$.

We first show by induction that, if $A_n \leq B_n$, then this inequality holds also for $n + 1$.

Base Case (n = 1).

$$A_1 = b - b + y_1 w_1 = y_1 w_1 = \min\{b, w_1 y_1\} = B_1. \tag{82}$$

Induction Step.

$$A_{n+1} = b - b \prod_{i \in [n+1]} (1 - y_i w_i / b) \tag{83}$$

$$= b - b \prod_{i \in [n]} (1 - y_i w_i / b)(1 - y_{n+1} w_{n+1} / b) \tag{84}$$

$$= b - b \prod_{i \in [n]} (1 - y_i w_i / b) + (b y_{n+1} w_{n+1} / b) \prod_{i \in [n]} (1 - y_i w_i / b) \tag{85}$$

$$= A_n + y_{n+1} w_{n+1} \prod_{i \in [n]} (1 - y_i w_i / b) \tag{86}$$

$$\leq A_n + y_{n+1} w_{n+1}. \tag{87}$$

The last inequality holds since by construction $w_i \leq b$ and thus $0 \leq y_i w_i / b \leq 1$, and $0 \leq \prod_{i \in V} (1 - y_i w_i / b) \leq 1$. For the same reason, $0 \leq \prod_{i \in [n+1]} (1 - y_i w_i / b) \leq 1$ and thus, by (83), we have $A_{n+1} \leq b$. Moreover, note that if $B_n = b$ then $B_{n+1} = b$. Therefore:

$$A_{n+1} \leq \min \{b, A_n + y_{n+1} w_{n+1}\} \tag{88}$$

$$\leq \min \{b, B_n + y_{n+1} w_{n+1}\} \tag{89}$$

$$= \begin{cases} \min \left\{ b, \sum_{i=1}^{n+1} y_i w_i \right\} = B_{n+1}, & \text{if} \quad B_n \leq b, \\ \min \{b, b + y_{n+1} w_{n+1}\} = b = B_{n+1}, & \text{if} \quad B_n = b, \end{cases} \tag{90}$$

and the proof by induction is completed. $\qquad\square$

Lemma 4. *Consider $n \in \mathbb{N}$, $\mathbf{y} \in [0,1]^n$, $b \in \mathbb{R}_{>0}$, and $\mathbf{w} \in [0,b]^n$. The following holds*

$$b - b \prod_{i \in [n]} (1 - y_i w_i / b) \geq (1 - 1/n)^n \min \left\{ b, \sum_{i \in [n]} y_i w_i \right\}, \tag{91}$$

Proof. Consider $\mathbf{y}' = (y_i w_i / b)_{i \in [n]} \in [0,1]^n$. Then, it follows from [40, Lemma 3.1].

$$1 - 1 \prod_{i \in [n]} (1 - y_i') \geq (1 - 1/n)^n \min \left\{ 1, \sum_{i \in [n]} y_i' \right\}. \tag{92}$$

Replacing $\mathbf{y}'$ by its value and multiplying both sides by b yields Eq. (91). $\qquad\square$

Lemma 5. *Consider $n \in \mathbb{N}$, $\mathbf{y} \in [0,1]^n$, and $\mathbf{w} \in [0,1]^n$. The following holds*

$$\prod_{i \in [n]} (1 - w_i y_i) = \sum_{S \in 2^{[n]}} \prod_{i \in S} (1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n] \setminus S} (1 - w_i). \tag{93}$$

Proof. We show by induction that, if Eq. (93) holds for some $n \in \mathbb{N}$, then this equality holds also for $n + 1$. Define $A_n(\mathbf{y}) \triangleq \prod_{i \in [n]}(1 - w_i y_i)$ and $B_n(\mathbf{y}) \triangleq \sum_{S \in 2^{[n]}} \prod_{i \in S}(1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n] \setminus S}(1 - w_i)$ for $n \in \mathbb{N}$ and $y \in [0, 1]^n$.

Base case ($n = 1$). In this setting the powerset of $\{1\}$ is $2^{\{1\}} = \{\emptyset, \{1\}\}$.

$$1 - w_1 y_1 = \sum_{S \in \{\emptyset, \{1\}\}} \prod_{i \in S}(1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n] \setminus S}(1 - w_i) \tag{94}$$

$$= 1 - w_1 + (1 - y_1) w_1 = 1 - w_1 y_1. \tag{95}$$

Induction Step.

$$B_{n+1}(\mathbf{y}) = \sum_{\substack{S \in 2^{[n+1]} \\ n+1 \notin S}} \prod_{i \in S}(1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n+1] \setminus S}(1 - w_i)$$

$$+ \sum_{\substack{S \in 2^{[n+1]} \\ n+1 \in S}} \prod_{i \in S}(1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n] \setminus S}(1 - w_i) \tag{96}$$

$$= \sum_{S \in 2^{[n]}} \prod_{i \in S}(1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n] \setminus S}(1 - w_i)(1 - w_{n+1})$$

$$+ \sum_{S \in 2^{[n]}} \prod_{i \in S}(1 - y_i) \prod_{i \in S} w_i \prod_{i \in [n] \setminus S}(1 - w_i) w_{n+1}(1 - y_{n+1}) \tag{97}$$

$$= A_n(\mathbf{y})(1 - w_{n+1})(1 - w_{n+1} + w_{n+1}(1 - y_{n+1})) = A_{n+1}(\mathbf{y}). \tag{98}$$

The first equality Eq. (96) is obtained from the definition of $B_n(\mathbf{y})$ and splitting the summation over the set $2^{[n+1]}$ into a summation over two subsets $\{S \subseteq 2^{[n+1]} : n + 1 \notin S\}$ and $\{S \subseteq 2^{[n+1]} : n + 1 \in S\}$. Equation (97) is obtained by factoring the terms that depend on $n + 1$. Equation (97) follows from the definition of $A_n(\mathbf{y})$. This concludes the proof.

$\square$

Lemma 6. *Consider random variables $x_i \in \{0, 1\}, i \in V$ with expected values $\mathbb{E}[x_i] = y_i, i \in V$ that are negatively correlated. The random variables $w_i x_i \in \{0, w_i\}$ are also negatively correlated for non-negative weights $w_i \in [0, 1], i \in V$. The following holds for every $S \subseteq V$*

$$\mathbb{E}\left[\prod_{i \in S}(1 - w_i x_i)\right] \leq \prod_{i \in S}(1 - w_i y_i), \qquad \mathbb{E}\left[\prod_{i \in S}(1 + w_i - x_i)\right] \leq \prod_{i \in S}(1 + w_i - y_i). \tag{99}$$

Proof. We exploit the alternative expression of $\prod_{i \in B}(1 - w_i x_i)$ developed in Lemma 5 and the negative correlation property to prove this proposition. Consider l.h.s. term in Eq. (99)

$$\mathbb{E}\left[\prod_{i \in S}(1 - w_i x_i)\right] \stackrel{(93)}{=} \mathbb{E}\left[\sum_{S' \in 2^S} \prod_{i \in S'}(1 - x_i) \prod_{i \in S'} w_i \prod_{i \in S \setminus S'}(1 - w_i)\right] \tag{100}$$

$$= \sum_{S' \in 2^S} \mathbb{E}\left[\prod_{i \in S'}(1 - x_i)\right] \prod_{i \in S'} w_i \prod_{i \in S \setminus S'}(1 - w_i) \tag{101}$$

$$\leq \sum_{S' \in 2^S} \prod_{i \in S'}(1 - y_i) \prod_{i \in S'} w_i \prod_{i \in S \setminus S'}(1 - w_i) \tag{102}$$

$$= \prod_{i \in S}(1 - w_i y_i). \tag{103}$$

$\square$

Putting everything together, consider $\mathbf{y} \in [0,1]^n$ and $\mathbf{x} = \Xi(\mathbf{y})$ stated in the Proposition. From Lemmas 3–6 we have:

$$\mathbb{E}_\Xi\left[f(\mathbf{x})\right] \overset{(3)}{=} \mathbb{E}_\Xi\left[\sum_{i \in C} c_i \min\left\{b_i, \sum_{j \in S_i} w_{i,j} x_j\right\}\right] \tag{104}$$

$$\overset{(81)}{\geq} \mathbb{E}_\Xi\left[\sum_{i \in C} c_i \left(b_i - b_i \prod_{j \in S_i}\left(1 - \frac{x_i w_{i,j}}{b_i}\right)\right)\right] \tag{105}$$

$$\overset{(103)}{\geq} \sum_{i \in C} c_i \left(b_i - b_i \prod_{j \in S_i}\left(1 - \frac{y_i w_{i,j}}{b_i}\right)\right) \tag{106}$$

$$\overset{(91)}{\geq} \left(1 - \frac{1}{\Delta}\right)^\Delta \sum_{i \in C} c_i \min\left\{b_i, \sum_{j \in S_i} w_{i,j} y_j\right\}. \tag{107}$$

Note that when $b_i = \infty$, then $\mathbb{E}_\Xi\left[\min\left\{b_i, \sum_{j \in S_i} w_{i,j} x_j\right\}\right] = \min\left\{b_i, \sum_{j \in S_i} w_{i,j} y_j\right\}$ due to linearity. This concludes the proof of Lemma 1. $\qquad\square$

F   Swap Rounding

In swap rounding [37], we are given a fractional solution $\mathbf{y} \in \mathcal{Y} = \mathrm{conv}(\mathcal{X})$. We assume that $\mathbf{y}$ is decomposed as a convex combination of bases (i.e., maximal elements) of $\mathcal{X}$, i.e.,

$$\mathbf{y} = \sum_{k=1}^K \gamma_k \mathbf{z}_k, \qquad \text{where} \qquad \gamma_k \in [0,1], \sum_{k=1}^K \gamma_k = 1, \text{ and } \mathbf{z}_k \in \mathcal{X}, k \in [K], \text{ are bases of } \mathcal{X}. \tag{108}$$

We note again that Carathéodory's theorem implies the existence of such decomposition of at most $n+1$ points in $\mathcal{X}$; moreover, there exists a decomposition algorithm [49] for general matroids with a running time $\mathcal{O}\left(n^6\right)$. However, in all practical cases we consider (including uniform and partition matroids) the complexity is significantly lower. More specifically, on partition matroids, swap rounding reduces to an algorithm with linear time complexity, namely, $\mathcal{O}\left(mn\right)$ for m partitions [50].

In short, swap rounding randomly picks between pairs of basis vectors over several iterations, producing a random final integral solution. Let $\beta_k \triangleq \sum_{k'=1}^k \gamma_k$, $k \in [K]$. Then:

- Set $\mathbf{x}_1 = \mathbf{z}_1$.

- Then, at iteration, a random selection is made between vector $\mathbf{x}_k \in \mathcal{X}$ and basis $\mathbf{z}_{k+1} \in \mathcal{X}$ to produce a new integral $\mathbf{x}_{k+1} \in \mathcal{X}$ as follows:

$$\mathbf{x}_{k+1} = \begin{cases} \mathbf{x}_k, & \text{with probability } \frac{\beta_k}{\beta_{k+1}}, \\ \mathbf{z}_{k+1}, & \text{with probability } \frac{\gamma_{k+1}}{\beta_{k+1}}. \end{cases} \tag{109}$$

  The authors refer to this randomized selection as a "merge" between $\mathbf{x}_k$ and $\mathbf{z}_{k+1}$.

The final vector $\mathbf{x}_K$ is the output of the algorithm and satisfies, by construction:

$$\mathbb{E}[\mathbf{x}_K] = \sum_{k=1}^K \gamma_k \mathbf{z}_k = \mathbf{y}. \tag{110}$$

G Extensions

G.I Dynamic Setting

The derivation of Theorem 4 follows the same steps as the proof of Theorem 2. Recall that $\Lambda_{\mathcal{X}}(T, P_T) \subset \mathcal{X}^T$ is the set of sequences of decision in a set $\mathcal{X}$ with path length less than $P_T \in \mathbb{R}_{\geq 0}$ over a time horizon T defined in Sec.6. Consider the sequence of reward functions $\{f_1, f_2, \ldots, f_T\} \in \mathcal{F}^T$ and the associated sequence of concave relaxations $\left\{ \tilde{f}_1, \tilde{f}_2, \ldots, \tilde{f}_T \right\} \in \hat{\mathcal{F}}^T$. Then, Asm. 2 implies

$$\max_{(\mathbf{x}_t)_{t=1}^T \in \Lambda_{\mathcal{X}}(T,P_T)} \sum_{t \in [T]} f_t(\mathbf{x}_t) \leq \max_{(\mathbf{x}_t)_{t=1}^T \in \Lambda_{\mathcal{X}}(T,P_T)} \sum_{t \in [T]} \tilde{f}_t(\mathbf{x}_t) \leq \max_{(\mathbf{y}_t)_{t=1}^T \in \Lambda_{\mathcal{Y}}(T,P_T)} \sum_{t \in [T]} \tilde{f}_t(\mathbf{y}_t). \tag{111}$$

The first inequality is implied by Eq. (6) in Asm. 2, and the second inequality holds because maximizing over a superset $\Lambda_{\mathcal{Y}}(T, P_T) \supseteq \Lambda_{\mathcal{X}}(T, P_T)$ can only increase the attained objective value. The total expected reward obtained by the RAOCO policy is given by

$$\mathbb{E}_{\Xi}\left[\sum_{t \in [T]} f_t(\mathbf{x}_t) \right] \overset{(79)}{\geq} \alpha \sum_{t \in [T]} \tilde{f}_t(\mathbf{y}_t). \tag{112}$$

Eqs. (111) and (112) imply that

$$\alpha\text{-regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq \alpha \cdot \text{regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{Y}}\right). \tag{113}$$

The OCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}}$ is OMA in Alg. 2 configured with an appropriately selected mirror map satisfying Asms. 4–6 (e.g. OGA by selecting $\Phi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$). Then under Asm. 1, we conclude from Theorem 7 the following

$$\alpha\text{-regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq \alpha \cdot \text{regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{Y}}\right) \leq \sqrt{2/\rho \left(D^2 + 2L_{\Phi}P_T\right) L^2 T} = \mathcal{O}\left(\sqrt{P_T T}\right), \tag{114}$$

where L, ρ^{-1}, D are all bounded from above by a finite quantity.

G.II Optimistic Setting

Under Asm. 2, the following inequality holds from Eq. (113) in the dynamic setting.

$$\alpha\text{-regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq \alpha \cdot \text{regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{Y}}\right). \tag{115}$$

Combining the above inequality and Theorem 7 yields the following upper bound:

$$\alpha\text{-regret}_{T,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq \alpha \cdot \sqrt{2/\rho \left(D^2 + 2L_{\Phi}P_T\right) \sum_{t=1}^T \|\mathbf{g}_t - \mathbf{g}_t^{\pi}\|_{\star}^2} = \mathcal{O}\left(\sqrt{P_T \sum_{t=1}^T \|\mathbf{g}_t - \mathbf{g}_t^{\pi}\|_{\star}^2}\right) \tag{116}$$

where L, ρ^{-1}, D, P_T are all bounded from above by a finite quantity.

G.III Bandit setting

In this section, we provide a full description of the Limited-Information Rounding-augmented OCO (LI-RAOCO) policy in Alg. 5. The policy decomposes the timeslots $[T]$ to T/W (assuming w.l.o.g. $T/W \in \mathbb{N}$) sequences of equally-sized contiguous timeslots denoted by $\mathcal{W}_l \triangleq \{(l-1)W+1, (l-1)W+2, \ldots, lW\}$ for $l \in [T/W]$. The algorithm takes as input exploration matrices $\mathbf{H}_l$ and shrinkage parameter $\delta > 0$.

Paper	Prob. Class	(1 − 1/e)-regret (Bandit) Static — Uni.	Part.	Gen.	Dynamic Part.	Optimistic Part.	Time
[33]	LWD	$n^{\frac{5}{3}}T^{\frac{2}{3}}$			×	×	TO_α
[4]	GS	$rn^{\frac{2}{3}}\log^{\frac{1}{3}}(n)T^{\frac{2}{3}}$	×	×	×	×	T^4O_b
[1]	GS	$\log^{\frac{1}{3}}(n)(nc_{\mathrm p}r)^{\frac{2}{3}}T^{\frac{2}{3}}$	×	×	×	×	$n^2c_{\mathrm p}$
[5]	GS	$r^{\frac{4}{3}}n^{\frac{4}{3}}T^{\frac{2}{3}}$	×	×	×	×	nr
This work	GS	$\alpha r^{4/3}n^{4/3}T^{2/3}$		×	$\sqrt{P_T}T^{4/5}$	$W\sqrt{P_T\sum_{l=1}^{T/W}\|\tilde{\mathbf g}_l - \mathbf g_l^\pi\|_2^2}+T(\delta+1/W)$	nr

Table 4: Order $(\mathcal{O}(\,\cdot\,))$ comparison of regrets and time complexities attained by different online submodular optimization algorithms for general submodular (GS) under the bandit setting. We also include Kakade et al. [33], who operate on Linearly Weighted Decomposable (LWD) functions. Competing algorithms that outperform us either only operate on limited constraint sets [4, 1] or the LWD class [33]; with the exception of the work of [5] which attains a tighter regret in the static setting over GS class, however when we employ a specialized static OCO algorithm we attain the same bound (see the discussion in Sec. G). Our work also generalizes to the dynamic and optimistic settings in the bandit setting.

In what follows we consider $\mathbf{H}_l = \delta\mathbf{I}$. Moreover, the algorithm utilizes an OCO algorithm $\mathcal{P}_{\mathcal{Y}_\delta}$ operating over a properly designed convex subset $\mathcal{Y}_\delta$ of a set $\mathcal{Y}$. The decisions are sampled through a rounding scheme $\hat{\tilde{\Xi}} : \mathcal{Y} \to \mathcal{X}$ at a given timeslot. The policy computes gradient estimates to explore the best actions by an OCO algorithm. The regret bound of the OCO policy is exploited by freezing the decision obtained by the algorithm for a time window of size $W \in \mathbb{N}$. Overall, the algorithm attains the bound specified in Theorem 6. In what follows, we describe the decision set $\mathcal{X}$, the construction of the subset $\mathcal{Y}_\delta \subseteq \mathcal{Y}$, and the rounding scheme $\tilde{\Xi}$. Moreover, we introduce additional definitions that include the continuous extension of the reward function induced by the randomized rounding scheme, and its auxiliary function. We finally show that the gradient estimates in Alg. 5 (line 21) are unbiased estimates of the gradients of the auxiliary function; this reduction relies on the reduction of Wan et al. [5]. We restate Theorem 6, and discuss how the regret bounds under dynamic and optimistic settings are obtained when OOMA in Alg. 3 is selected as the OCO policy.

Under this setting, we restrict the decision set $\mathcal{X}$ to partition matroids. Given a ground set $V = [n]$, partitions $V_i \subseteq V$, cardinalities $r_i \in \mathbb{N}$ for $i \in [m]$, the set $\mathcal{X}$ is given by

$$\mathcal{X} \triangleq \left\{\mathbf{x} \in \{0,1\}^n : \sum_{j\in V_i} x_j \leq r_i, i \in [m]\right\}. \tag{117}$$

Let $n' \triangleq \sum_{i\in[m]} |V_i| \times r_j$, we define the alternative representation of the partition matroid $\mathcal{X}$ given as

$$\mathcal{X}' \triangleq \left\{\mathbf{x} \in \{0,1\}^{n'} : \sum_{k\in V_i} x_{i,j,k} \leq 1, j \in [r_i], i \in [m]\right\}, \tag{118}$$

where $x_{i,j,k}$ indicates whether item $k \in V$ is selected from partition V_i at slot j. Consider the mapping $T : \mathcal{X}' \to \mathcal{X}$ defined by

$$T(\mathbf{x}') \triangleq \left(\min\left\{1, \sum_{i\in[m]:k\in V_i} \sum_{j\in[r_i]} \mathbf{x}'_{i,j,k}\right\}\right)_{k\in V}, \qquad \text{for } \mathbf{x} \in \mathcal{X}'. \tag{119}$$

It is easy to see that maximizing a function $f : \mathcal{X} \to \mathbb{R}_{\geq 0}$ over the sets $\mathcal{X}$ or $\{T(\mathbf{x}') : \mathbf{x}' \in \mathcal{X}'\}$ attains the same value, i.e.,

$$\max_{\mathbf{x}'\in\mathcal{X}'} f(T(\mathbf{x}')) = \max_{\mathbf{x}\in\mathcal{X}} f(\mathbf{x}). \tag{120}$$

Thus, we could simply optimize over the set $\mathcal{X}'$ and map to points in $\mathcal{X}$ through the mapping T.

Algorithm 5 Limited-Information Rounding-augmented OCO (LIRAOCO) policy
---
Require: Exploration parameter $\delta \in (0, 1/2n]$, OCO policy $\mathcal{P}_{\mathcal{Y}_\delta}$, $\mathbf{H}_l, l \in [T/W]$
  1: **for** $\mathcal{W}_l = \mathcal{W}_1, \mathcal{W}_2, \ldots, \mathcal{W}_{T/W}$ **do** $\triangleright$ $\mathcal{W}_l$ is a sequence of $|\mathcal{W}_l| = W$ contiguous timeslots s.t.
       $[T] = \bigcup_{l \in [T/W]} \mathcal{W}_l$
  2:      Draw t_l u.a.r. from $\mathcal{W}_l$
  3:      **for** $t \in \mathcal{W}_l$ **do**
  4:          **if** $t = t_l$ **then**
  5:              $\mathbf{y}_l^\delta \leftarrow \mathcal{P}_{\mathcal{Y}_\delta, l}\left((\mathbf{y}_s^\delta)_{s=1}^{l-1}, (r_s)_{s=1}^{l-1} \right)$
  6:              Sample z_l from Z where $\mathbb{P}(Z < z) = \int_0^z \frac{\exp(s-1)}{1-\exp(-1)} \mathbb{1}\,(s \in [0,1])\,ds$
  7:              Sample $\mathbf{v}_l$ u.a.r. from $\mathbb{S}_{n-1}$
  8:              **if** $z_l \geq \frac{1}{2}$ **then**
  9:                  Sample $\mathbf{u}_l$ from $\{\mathbf{0}, \mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n\}$ w.p. $\mathbb{P}(\mathbf{u}_l = \mathbf{0}) = 1/2$ and $\mathbb{P}(\mathbf{u}_l = \mathbf{e}_i) = 1/(2n)$
 10:                  $\mathbf{y}_l \leftarrow z_l \cdot (\mathbf{y}_l^\delta + ((\mathbf{H}_l\mathbf{v}_l)^\mathsf{T}\mathbf{u}_l)\,\mathbf{u}_l)$
 11:                  Play $\mathbf{x}_t \leftarrow \hat{\Xi}(\mathbf{y}_l)$
 12:                  Receive reward $f_t(\mathbf{x}_t)$
 13:                  $\tilde{l}_l \leftarrow 2\alpha \frac{n}{z_l} f_t(\mathbf{x}_t)(1 - 2\mathbb{1}\,(\mathbf{u}_l = \mathbf{0}))$
 14:              **else**
 15:                  Sample $\mathbf{u}_l$ u.a.r. from $\{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n\}$
 16:                  Assign u.a.r. $\mathbf{y}_l \leftarrow z_l\mathbf{y}_l^\delta + \frac{1}{2}\mathbf{u}_l$ or $\mathbf{y}_l \leftarrow z_l\mathbf{y}_l^\delta$
 17:                  Play $\mathbf{x}_t \leftarrow \hat{\Xi}(\mathbf{y}_l)$
 18:                  Receive reward $f_t(\mathbf{x}_t)$
 19:                  $\tilde{l}_l \leftarrow 4\alpha n\,((\mathbf{H}_l\mathbf{v}_l)^\mathsf{T}\mathbf{u}_l)\,f_t(\mathbf{x}_t)(1 - 2\mathbb{1}\,(\mathbf{y}_l = z_l\mathbf{y}_l^\delta))$
 20:              **end if**
 21:              Construct an estimate of supergradient $\tilde{\mathbf{g}}_l \leftarrow n\tilde{l}_l\mathbf{H}_l^{-1}\mathbf{v}_l$
 22:              Construct a linear function $r_l(\mathbf{y}) = \mathbf{y} \cdot \tilde{\mathbf{g}}_l$
 23:          **else**
 24:              Play $\mathbf{x}_t \leftarrow \hat{\Xi}(\mathbf{y}_l^\delta)$
 25:              Receive reward $f_t(\mathbf{x}_t)$
 26:          **end if**
 27:      **end for**
 28: **end for**

    We define the fractional decision set for the OCO policy $\mathcal{P}_{\mathcal{Y}_\delta}$ employed by LIRAOCO policy in Alg. 5 to be the convex set $\mathcal{Y} \triangleq \mathrm{conv}\,(\mathcal{X}')$. We define an appropriate subset of $\mathcal{Y}$ denoted by $\mathcal{Y}_\delta$ s.t. for every point $\mathbf{y} \in \mathcal{Y}_\delta$ we can form a ball of radius $\delta \in (0,1)$ contained in $\mathcal{Y}$. This set permits the decision maker to explore (perturb) without exiting the original decision set $\mathcal{Y}$. We provide an illustration in Fig. 3. Formally,

Proposition 6. *Given the convex set $\mathcal{Y} = \mathrm{conv}\,(\mathcal{X}')$ in Eq. (118) and $\delta \in [0, 1/\,(2\max\{|V_i| : i \in [m]\})]$, every ball formed around points in*

$$\mathcal{Y}_\delta \triangleq \left\{ \left((1 - 2\delta|V_i|)\mathbf{y}_{i,j} + \mathbf{1}_{|V_i|}\delta \right)_{i\in[m], j\in[r_i]} : \mathbf{y} \in \mathcal{Y} \right\} \tag{121}$$

are contained in $\mathcal{Y}$, i.e.,

$$\mathbf{y} + \delta\mathbf{v} \in \mathcal{Y}, \qquad \text{for } \mathbf{y} \in \mathcal{Y}_\delta,\ \mathbf{v} \in \mathbb{B}_n. \tag{122}$$

Proof. Note that the set $\mathcal{Y}$ is a product of simplexes, i.e., $\mathcal{Y} = \times_{i\in[m]} (\Delta_{V_i})^{r_i}$, where Δ_{V_i} is a simplex with support V_i. We consider the subset $\Delta_{\delta, V_i} = \left\{ \delta\mathbf{1}_{|V_i|} + (1 - 2|V_i|\delta)\mathbf{y} : \mathbf{y} \in \Delta_{V_i} \right\}$ for $\delta \in (0, 1/(2|V_i|))$.
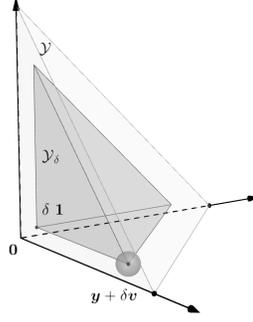
Figure 3: Illustration of Proposition 6: under the shrunken set $\mathcal{Y}_\delta$, any ball formed around points $\mathbf{y} \in \mathcal{Y}_\delta$ of radius δ is contained in $\mathcal{Y}$.

Let $\mathbf{y}_{i,j} \in \Delta_{\delta,|V_i|}$ for $j \in [r_i]$ and $i \in [m]$. It holds that $\mathbf{y}_{i,j} + \delta\mathbf{v}_{i,j} = \delta\left(\mathbf{1}_{|V_i|} + \mathbf{v}_{i,j}\right) + (1 - 2|V_i|\delta)\mathbf{y}'_{i,j}$ for some point $\mathbf{y}'_{i,j}$ in Δ_{V_i}. The point $\frac{1}{2|V_i|}(\mathbf{1}_{|V_i|} + \mathbf{v}_{i,j})$ satisfies $\mathbf{0} \leq \frac{1}{2|V_i|}(\mathbf{v}_{i,j} + \mathbf{1}_{|V_i|}) \leq \frac{1}{|V_i|}\mathbf{1}_{|V_i|}$ (element-wise). Thus, it holds $\frac{1}{2|V_i|}(\mathbf{v}_{i,j} + \mathbf{1}_{|V_i|}) \in \text{conv}\left(\left\{\mathbf{0}, \mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_{|V_i|}\right\}\right) \subseteq \Delta_{V_i}$. Thus, we have

$$\mathbf{y}_{i,j} + \delta\mathbf{v}_{i,j} = (1 - 2|V_i|\delta)\underbrace{\mathbf{y}'_{i,j}}_{\in \Delta_{V_i}} + 2|V_i|\delta\underbrace{(\mathbf{1}_{|V_i|} + \mathbf{v}_{i,j})/2|V_i|}_{\in \Delta_{V_i}}. \tag{123}$$

It follows that $\mathbf{y}_{i,j} + \delta\mathbf{v}_{i,j} \in \Delta_{V_i}$ because Δ_{V_i} is a convex set, and $\mathbf{y}_{i,j}$ is a convex combinations of two points in Δ_{V_i}. Since $\mathcal{Y}$ is a product of simplexes, it holds that

$$\mathcal{Y}_\delta = \left\{\left((1 - 2\delta|V_i|)\mathbf{y}_{k,l} + \mathbf{1}_{|V_i|}\delta\right)_{i\in[m],j\in[r_i]} : \mathbf{y} \in \mathcal{Y}\right\}.$$

$\square$

We provide a definition of the rounding scheme $\hat{\Xi} : \mathcal{Y} \to \mathcal{X}$ employed by LIRAOCO in Alg. 5.

Definition 4. *[5, Extention Mapping] The randomized rounding scheme $\hat{\Xi} : \mathcal{Y} \to \mathcal{X}$ defined for sets $\mathcal{Y} = \text{conv}(\mathcal{X}')$, and $\mathcal{Y}$ in Eqs. (118) and (117). Given a point $\mathbf{y} \in \mathcal{Y}$, an intermediate randomized rounding scheme $\xi : \mathcal{Y} \to \mathcal{X}'$ outputs a point $\boldsymbol{\omega} \in \mathcal{X}'$, by sampling a single item $k \in V_i$ from the distribution $\mathbf{y}_{i,j}/\|\mathbf{y}_{i,j}\|_1$ with support V_i w.p. $\|\mathbf{y}_{i,j}\|_1$ and assigns $\mathbf{e}_k$ to $\boldsymbol{\omega}_{i,j}$, otherwise it assigns $\mathbf{0}$ w.p. $1 - \|\mathbf{y}_{i,j}\|_1$ to $\boldsymbol{\omega}_{i,j}$, for every $i \in [m]$ and $j \in [r_i]$. The rounding scheme $\hat{\Xi}$ is then defined as the composition mapping $T \circ \xi : \mathcal{Y} \to \mathcal{X}$ where T is provided in Eq. (119).*

Definition 5. *A function $\hat{f} : \mathcal{Y} \to \mathbb{R}_{\geq 0}$ is a DR-submodular function if for any $\mathbf{y}, \mathbf{y}' \in \mathcal{Y}$*

$$\hat{f}(\mathbf{y} + t\mathbf{e}_i) - \hat{f}(\mathbf{y}) \geq \hat{f}(\mathbf{y}' + t\mathbf{e}_i) - \hat{f}(\mathbf{y}'), \tag{124}$$

where $\mathbf{y} \leq \mathbf{y}'$, $\mathbf{y} + t\mathbf{e}_i, \mathbf{y}' + t\mathbf{e}_i \in \mathcal{Y}$. When f is differentiable we have $\nabla\hat{f}(\mathbf{y}) \geq \nabla\hat{f}(\mathbf{y}')$ for $\mathbf{y} \leq \mathbf{y}'$. When f is twice differentiable, the DR-submodularity is equivalent to

$$\frac{\partial^2 \hat{f}(\mathbf{y})}{\partial y_i \partial y_j} \leq 0, \qquad\qquad \text{for all } i, j \in [n] \text{ and } \mathbf{y} \in \mathcal{Y}. \tag{125}$$

Moreover, f is monotone if $\hat{f}(\mathbf{y}) \leq \hat{f}(\mathbf{y}')$ when $\mathbf{y} \leq \mathbf{y}'$.

Definition 6. *[32, Auxillary Function] For a monotone DR-submodular function $\hat{f} : \mathcal{Y} \to \mathbb{R}_{\geq 0}$ satisfying $\hat{f}(\mathbf{0}) = 0$, its auxiliary function $\hat{F} : \mathcal{Y} \to \mathbb{R}_{\geq 0}$ is defined as follows*

$$\hat{F}(\mathbf{y}) \triangleq \int_0^1 \frac{\exp(z - 1)}{z}\hat{f}(z \cdot \mathbf{y})dz., \qquad \text{for } \mathbf{y} \in \mathcal{Y}. \tag{126}$$

38

Definition 7. *For continuous function $l : \mathbb{R}^n \to \mathbb{R}$ and invertible matrix $\mathbf{H} \in \mathbb{R}^{d \times d}$, an $\mathbf{H}$-smoothed version of l denoted by $l^{\mathbf{H}}$ is defined as*

$$l^{\mathbf{H}}(\mathbf{x}) \triangleq \mathbb{E}_{\mathbf{v} \sim \mathbb{B}_d} \left[l(\mathbf{x} + \mathbf{H}\mathbf{v}) \right], \tag{127}$$

where $\mathbf{v} \sim \mathbb{B}_n$ indicates that $\mathbf{v}$ is sampled from the n-dimensional unit ball $\mathbb{B}_d$ uniformly at random.

Remark that when $l : \mathbb{R}^n \to \mathbb{R}$ is a linear function, then $l = l^{\mathbf{H}}$.

Lemma 7. *[5, Lemma 5.1] Given a submodular monotone function $f : \mathcal{X} \to \mathbb{R}_{\geq 0}$ s.t. $f(\mathbf{x}) < L$ for $\mathbf{x} \in \mathcal{X}$ and $f(\mathbf{0}) = 0$. The randomized rounding scheme $\hat{\Xi} : \mathcal{Y} \to \mathcal{X}$ in Def. 4 defines a multi-linear, monotone-increasing, DR-submodular, and L-lipschitz (w.r.t. $\| \cdot \|_\infty$) function $\hat{f} : \mathcal{Y} \to \mathbb{R}_{\geq 0}$ given by*

$$\hat{f}(\mathbf{y}) \triangleq \mathbb{E}_{\mathbf{x} \sim \hat{\Xi}(\mathbf{y})} \left[f(\mathbf{x}) \right], \qquad for \, \mathbf{y} \in \mathcal{Y}, \tag{128}$$

Proof. The proof follows the same lines as the proof of [5, Lemma 5.3].

Multilinearity. Given a point $\mathbf{y} \in \mathcal{Y}$, we have

$$\hat{f}(\mathbf{y}) = \mathbb{E}_{\mathbf{x} \sim \hat{\Xi}(\mathbf{y})} \left[f(\mathbf{x}) \right] = \mathbb{E}_{\boldsymbol{\omega} \sim \xi(\mathbf{y})} \left[f(T(\omega)) \right] = \sum_{\mathbf{x}' \in \mathcal{X}'} \mathbb{P}(\boldsymbol{\omega} = \mathbf{x}') f(T(\mathbf{x}')). \tag{129}$$

Note that $\mathbb{P}(\boldsymbol{\omega} = \mathbf{x}') = \prod_{i \in [m]} \prod_{j \in [r_i]} \mathbb{P}(\boldsymbol{\omega}_{i,j} = \mathbf{x}'_{i,j})$; moreover, $\mathbb{P}(\boldsymbol{\omega}_{i,j} = \mathbf{x}'_{i,j}) = y_{i,j,k}$ when $\mathbf{x}'_{i,j} = \mathbf{e}_k$, otherwise $\mathbb{P}(\boldsymbol{\omega}_{i,j} = \mathbf{x}'_{i,j}) = 1 - \sum_{k \in V_i} y_{i,j,k}$ when $\mathbf{x}'_{i,j} = \mathbf{0}$, for $i \in [m]$ and $j \in [r_i]$. The factor $f(T(\mathbf{x}'))$ is a constant independent from $\mathbf{y}$, so we conclude $\hat{f}$ is a multilinear function.

Monotonicity and lipschitzness. Let $\mathbf{e}_{i,j,k} = (\mathbb{1}((i', j', k') = (i, j, k)))_{i' \in [m], j' \in [r_i], k' \in V_i}$ and $\bar{\mathbf{e}}_{i,j,k} = \mathbf{1} - \mathbf{e}_{i,j,k}$. Since the function $\hat{f}$ is multilinear, it holds that

$$\frac{\partial \hat{f}(\mathbf{y})}{\partial y_{i,j,k}} = \frac{\hat{f}(\mathbf{y} \oplus \lambda \mathbf{e}_{i,j,k}) - \hat{f}(\mathbf{y} \ominus \bar{\mathbf{e}}_{i,j,k})}{\lambda}. \tag{130}$$

where $\lambda = 1 - \sum_{k' \in V_i : k' \neq k} y_{i,j,k'} > 0$, and the binary operators $\oplus$ and $\ominus$ are component-wise maximum and minimum operations, respectively, i.e., $\mathbf{y} \oplus \mathbf{y}' = \left(\max \left\{ y_{i,j,k}, y'_{i,j,k} \right\} \right)_{i \in [m], j \in [r_i], k \in V_i}$ and $\mathbf{y} \ominus \mathbf{y}' = \left(\min \left\{ y_{i,j,k}, y'_{i,j,k} \right\} \right)_{i \in [m], j \in [r_i], k \in V_i}$. Consider the two mapping $L_{i,j,k}^{\oplus} : \mathcal{X}' \to \mathcal{X}'$ and $L_{i,j,k}^{\ominus} : \mathcal{X}' \to \mathcal{X}'$ defined as follows:

$$\left(L_{i,j,k}^{\oplus}(\mathbf{x}) \right)_{i',j'} = \begin{cases} \mathbf{x}_{i',j'} + \mathbf{e}_k \mathbb{1} \left(\mathbf{x}_{i',j'} = \mathbf{0} \right) & \text{if } (i', j') = (i, j), \\ \mathbf{x}_{i',j'} & \text{otherwise,} \end{cases} \tag{131}$$

$$\left(L_{i,j,k}^{\ominus}(\mathbf{x}) \right)_{i',j'} = \begin{cases} \mathbf{x}_{i',j'} \mathbb{1} \left(\mathbf{x}_{i',j'} \neq \mathbf{e}_k \right) & \text{if } (i', j') = (i, j), \\ \mathbf{x}_{i',j'} & \text{otherwise,} \end{cases} \tag{132}$$

for $i' \in [m], j' \in [r_i]$ and $\mathbf{x} \in \mathcal{X}'$. One could check that

$$\hat{f}(\mathbf{y} \oplus \lambda \mathbf{e}_{i,j,k}) = \mathbb{E}_{\boldsymbol{\omega} \sim \xi(\mathbf{y})} \left[f(T(L_{i,j,k}^{\oplus}(\boldsymbol{\omega}))) \right], \quad \hat{f}(\mathbf{y} \ominus \lambda \mathbf{e}_{i,j,k}) = \mathbb{E}_{\boldsymbol{\omega} \sim \xi(\mathbf{y})} \left[f(T(L_{i,j,k}^{\ominus}(\boldsymbol{\omega}))) \right]. \tag{133}$$

Thus, it holds

$$\frac{\partial \hat{f}(\mathbf{y})}{\partial y_{i,j,k}} = \frac{\hat{f}(\mathbf{y} \oplus \lambda \mathbf{e}_{i,j,k}) - \hat{f}(\mathbf{y} \ominus \bar{\mathbf{e}}_{i,j,k})}{\lambda} = \frac{\mathbb{E}_{\boldsymbol{\omega} \sim \xi(\mathbf{y})} \left[f(T(L_{i,j,k}^{\oplus}(\boldsymbol{\omega}))) - f(T(L_{i,j,k}^{\ominus}(\boldsymbol{\omega}))) \right]}{\lambda} \geq 0. \tag{134}$$

Because $\lambda \geq 0$, $L_{i,j,k}^{\ominus}(\boldsymbol{\omega}) \leq L_{i,j,k}^{\oplus}(\boldsymbol{\omega})$ (element-wise) by construction, and f is monotone increasing. Note that $L_{i,j,k}^{\oplus}(\boldsymbol{\omega}) \neq L_{i,j,k}^{\ominus}(\boldsymbol{\omega})$ only when $\boldsymbol{\omega}_{i,j} = \mathbf{e}_k$ or $\boldsymbol{\omega}_{i,j} = \mathbf{0}$. So, we have

$$\frac{\partial \hat{f}(\mathbf{y})}{\partial y_{i,j,k}} \leq L \frac{\mathbb{P}(\boldsymbol{\omega}_{i,j} \in \{\mathbf{e}_k, \mathbf{0}\})}{1 - \sum_{k' \in V_i : k' \neq k} y_{i,j,k'}} = L \frac{1 - \sum_{k' \in V_i : k' \neq k} y_{i,j,k'}}{1 - \sum_{k' \in V_i : k' \neq k} y_{i,j,k'}} = L. \tag{135}$$

The function $\hat{f}$ is M-lipschitz w.r.t. $\|\cdot\|_{\infty}$.

DR-submodularity. The partial derivative of a multilinear function is also multilinear. So, we have the following

$$\frac{\partial^2 \hat{f}(\mathbf{y})}{\partial y_{i,j,k} \partial y_{i',j',k'}} = \frac{\partial}{\partial y_{i',j',k'}} \left(\frac{\hat{f}(\mathbf{y} \oplus \lambda \mathbf{e}_{i,j,k}) - \hat{f}(\mathbf{y} \ominus \bar{\mathbf{e}}_{i,j,k})}{\lambda} \right) \tag{136}$$

$$= \frac{\hat{f}(\mathbf{y} \oplus \lambda \mathbf{e}_{i,j,k} \oplus \lambda' \mathbf{e}_{i',j',k'}) - \hat{f}(\mathbf{y} \ominus \bar{\mathbf{e}}_{i',j',k'} \oplus \lambda \mathbf{e}_{i,j,k}) - f(\mathbf{y} \oplus \lambda' \mathbf{e}_{i',j',k'} \ominus \mathbf{e}_{i,j,k}) + f(\mathbf{y} \ominus \mathbf{e}_{i,j,k} \ominus \mathbf{e}_{i',j',k'})}{\lambda \lambda'}. \tag{137}$$

As established in Eq. (133), it holds

$$\hat{f}(\mathbf{y} \oplus \lambda \mathbf{e}_{i,j,k} \oplus \lambda' \mathbf{e}_{i',j',k'}) - \hat{f}(\mathbf{y} \ominus \bar{\mathbf{e}}_{i',j',k'} \oplus \lambda \mathbf{e}_{i,j,k}) - f(\mathbf{y} \oplus \lambda' \mathbf{e}_{i',j',k'} \ominus \mathbf{e}_{i,j,k}) + f(\mathbf{y} \ominus \mathbf{e}_{i,j,k} \ominus \mathbf{e}_{i',j',k'}) =$$

$$\tag{138}$$

$$\mathbb{E}_{\boldsymbol{\omega} \sim \xi(\mathbf{y})} \left[f(T(L_{i,j,k}^{\oplus}(L_{i',j',k'}^{\oplus}(\boldsymbol{\omega})))) - f(T(L_{i,j,k}^{\oplus}(L_{i',j',k'}^{\ominus}(\boldsymbol{\omega})))) - f(T(L_{i,j,k}^{\ominus}(L_{i',j',k'}^{\oplus}(\boldsymbol{\omega})))) \right. \tag{139}$$

$$\left. + f(T(L_{i,j,k}^{\ominus}(L_{i',j',k'}^{\ominus}(\boldsymbol{\omega})))) \right]. \tag{140}$$

Considering that f is submodular and the second-order differences definition of submodularity [71], we obtain $f(T(L_{i,j,k}^{\oplus}(L_{i',j',k'}^{\oplus}(\boldsymbol{\omega})))) - f(T(L_{i,j,k}^{\oplus}(L_{i',j',k'}^{\ominus}(\boldsymbol{\omega})))) - f(T(L_{i,j,k}^{\ominus}(L_{i',j',k'}^{\oplus}(\boldsymbol{\omega})))) \geq 0$ for any $\boldsymbol{\omega} \in \mathcal{X}'$. It follows that the second-partial derivative is non-positive

$$\frac{\partial^2 \hat{f}(\mathbf{y})}{\partial y_{i,j,k} \partial y_{i',j',k'}} \leq 0. \tag{141}$$

Thus, $\hat{f}$ is DR-submodular. This concludes the proof. $\qquad\square$

Lemma 8. *[8, Corollary 6.8] Let $\mathbf{H} \in \mathbb{R}^{n \times n}$ be an invertible matrix, $l : \mathbb{R}^n \to \mathbb{R}$ be a continuous function. Then*

$$\nabla l^{\mathbf{H}}(\mathbf{x}) = n \mathbb{E}_{\mathbf{v} \sim \mathbb{S}_{n-1}} \left[l(\mathbf{x} + \mathbf{H}\mathbf{v}) \mathbf{H}^{-1} \mathbf{v} \right], \tag{142}$$

where $\mathbf{v} \sim \mathbb{S}_{n-1}$ indicates that $\mathbf{v}$ is sampled from the $(n-1)$-dimensional unit sphere $\mathbb{S}_{n-1}$ u.a.r.

Proof. This proof follows the same lines as the proof of [Lemma 6.7][8].

Part I. Consider the case when $\mathbf{H} = \delta \mathbf{I}$. Using Calculus Stoke's theorem, we have

$$\nabla_{\mathbf{x}} \left(\int_{\mathbb{B}_{\delta}} l(\mathbf{x} + \mathbf{v}) d\mathbf{v} \right) = \int_{\mathbb{S}_{\delta}} l(\mathbf{x} + \mathbf{u}) \frac{\mathbf{u}}{\|\mathbf{u}\|} d\mathbf{u}. \tag{143}$$

We have the following

$$l^{\delta \mathbf{I}}(\mathbf{x}) = \mathbb{E}_{\mathbf{v} \sim \mathbb{B}} \left[l(\mathbf{x} + \delta \mathbf{v}) \right] = \frac{1}{\text{vol}(\mathbb{B}_{\delta})} \int_{\mathbb{B}_{\delta}} l(\mathbf{x} + \mathbf{v}) d\mathbf{v}, \tag{144}$$

where $\mathrm{vol}(\mathbb{B}_\delta)$ is the volume of the n-dimensional ball of radius δ. Similarly, we also have

$$\mathbb{E}_{\mathbf{u}\sim\mathbb{S}_{n-1}}\left[l(\mathbf{x}+\delta\mathbf{u})\mathbf{u}\right] = \frac{1}{\mathrm{vol}(\mathbb{S}_{n-1})}\int_{\mathbb{S}_\delta} l(\mathbf{x}+\mathbf{u})\frac{\mathbf{u}}{\|\mathbf{u}\|}d\mathbf{u}. \tag{145}$$

The ratio of the volume of a ball in n dimensions and the sphere of dimension $n-1$ of radii δ is $\mathrm{vol}(\mathbb{B}_n)/\mathrm{vol}(\mathbb{S}_{n-1}) = \delta/n$. Combining these facts with Eq. (143), we showed that:

$$\nabla l^{\delta\mathbf{I}}(\mathbf{x}) = \frac{n}{\delta}\mathbb{E}_{\mathbf{v}\sim\mathbb{S}_{n-1}}\left[l(x+\delta\mathbf{v})\mathbf{v}\right]. \tag{146}$$

Part II. Let $g(\mathbf{x}) = l(\mathbf{H}\mathbf{x})$, and $g^{\mathbf{I}}(\mathbf{x}) = \mathbb{E}_{\mathbf{v}\in\mathbb{B}_n}\left[g(\mathbf{x}+\mathbf{v})\right]$. We have the following

$$n\mathbb{E}_{\mathbf{v}\sim\mathbb{S}_{n-1}}\left[l(\mathbf{x}+\mathbf{H}\mathbf{v})\mathbf{H}^{-1}\mathbf{v}\right] = n\mathbf{H}^{-1}\left[l(\mathbf{x}+\mathbf{H}\mathbf{v})\mathbf{v}\right] \tag{147}$$

$$= n\mathbf{H}^{-1}\left[g(\mathbf{H}^{-1}\mathbf{x}+\mathbf{v})\mathbf{v}\right] \tag{148}$$

$$= \mathbf{H}^{-1}\nabla g^{\mathbf{I}}(\mathbf{H}^{-1}\mathbf{x}) \tag{149}$$

$$= \mathbf{H}^{-1}\mathbf{H}\nabla l^{\mathbf{H}}(\mathbf{x}) = \nabla l^{\mathbf{H}}(\mathbf{x}). \tag{150}$$

This concludes the proof. □

Lemma 9. *[32, Lemma 2] Let $\hat{f} : \mathbb{R}^n \to \mathbb{R}$ be a monotone DR-submodular function and $\hat{f}(\mathbf{0}) = 0$, $\mathbf{x}$, $\mathbf{y} \in \mathbb{R}^n$. Let $\hat{F}$ be the auxiliary function defined in (126). Then*

$$\nabla\hat{F}(\mathbf{x}) = \int_0^1 \exp(z-1)\nabla f(z\cdot\mathbf{x})dz, \qquad (\mathbf{y}-\mathbf{x})\cdot\nabla\hat{F}(\mathbf{x}) \geq (1-1/e)\hat{f}(\mathbf{y})-\hat{f}(\mathbf{x}), \tag{151}$$

Proof. This proof follows the same lines as the proof of [32, Lemma 2]. It is simplified for the case when $\theta(w) = 1$ and $(w(z) = \exp(z-1))$. It follows from the DR-submodularity of $\hat{f}$:

$$\hat{f}(\mathbf{y})-\hat{f}(\mathbf{x}) = \int_0^1 (\mathbf{y}-\mathbf{x})\cdot\nabla\hat{f}(\mathbf{x}+z(\mathbf{y}-\mathbf{x}))dz \in [(\mathbf{y}-\mathbf{x})\cdot\nabla\hat{f}(\mathbf{y}), (\mathbf{y}-\mathbf{x})\cdot\nabla\hat{f}(\mathbf{x})], \tag{152}$$

for $\mathbf{x} \leq \mathbf{y}$. The two inequalities follow from $\mathbf{y} \geq \mathbf{x}+z(\mathbf{y}-\mathbf{x}) \geq \mathbf{x}$ s.t. $\nabla\hat{f}(\mathbf{x}) \geq \nabla\hat{f}(\mathbf{x}+z(\mathbf{y}-\mathbf{x})) \geq \nabla\hat{f}(\mathbf{y})$ for any $z \in [0,1]$; a direct result from the DR-submodularity of $\hat{f}$, see Eq. (124).

The gradient of $\hat{F}$ is given by $\nabla\hat{F}(\mathbf{x}) = \nabla_{\mathbf{x}}\left(\int_0^1 \frac{\exp(z-1)}{z}\hat{f}(z\cdot\mathbf{x})dz\right) = \int_0^1 \frac{z\exp(z-1)}{z}\nabla\hat{f}(z\cdot\mathbf{x})dz = \int_0^1 \exp(z-1)\nabla\hat{f}(z\cdot\mathbf{x})dz$. Let $w(z) = \exp(z-1)$. Consider the following:

$$\mathbf{x}\cdot\nabla\hat{F}(\mathbf{x}) = \int_0^1 w(z)\mathbf{x}\cdot\nabla\hat{f}(z\mathbf{x})dz = \int_0^1 w(z)\frac{\partial\hat{f}(z\mathbf{x})}{\partial z}dz = \int_0^1 w(z)d\hat{f}(z\mathbf{x}) \tag{153}$$

$$= \left[w(z)\hat{f}(z\mathbf{x})\right]_0^1 - \int_0^1 w'(z)\hat{f}(z\mathbf{x})dz = w(1)\hat{f}(\mathbf{x}) - \int_0^1 w'(z)\hat{f}(z\mathbf{x})dz \tag{154}$$

$$= \hat{f}(\mathbf{x}) - \int_0^1 w(z)\hat{f}(z\mathbf{x})dz. \tag{155}$$

We use integration by part and the fact that $w'(z) = w(z)$.

$$\mathbf{y}\cdot\nabla\hat{F}(\mathbf{y}) = \int_0^1 w(z)\mathbf{y}\cdot\nabla\hat{f}(z\mathbf{x})dz \geq \int_0^1 w(z)(\mathbf{y}\oplus z\mathbf{x} - z\mathbf{x})\cdot\nabla\hat{f}(z\mathbf{x})dz \tag{156}$$

$$\geq \int_0^1 w(z)(\hat{f}(\mathbf{y}\oplus z\mathbf{x}) - \hat{f}(z\mathbf{x}))\cdot\nabla\hat{f}(z\mathbf{x})dz \tag{157}$$

$$\geq \underbrace{\left(\int_0^1 w(z)dz\right)}_{1-1/e}\hat{f}(\mathbf{y}) - \int_0^1 w(z)\hat{f}(z\mathbf{x})dz. \tag{158}$$

41

The first inequality is obtained considering $\mathbf{y} \geq \mathbf{y} \oplus z\mathbf{x} - z\mathbf{x} \geq \mathbf{0}$ and $\nabla \hat{f}(z\mathbf{x}) \geq 0$ Combine Eqs. (155) and (158) to obtain

$$(\mathbf{y} - \mathbf{x}) \cdot \nabla \hat{F}(\mathbf{x}) \geq (1 - 1/e)\hat{f}(\mathbf{y}) - \hat{f}(\mathbf{x}). \tag{159}$$

We conclude the proof. $\qquad\qquad \square$

We define the following quantities which will be used in what follows.

$$\mathcal{H}_l \triangleq \{(\mathbf{v}_s, \mathbf{u}_s, z_s, t_s) : s \leq l\}, \qquad \hat{f}_l^{\mathrm{avg}}(\mathbf{y}) \triangleq \frac{1}{W} \sum_{t \in \mathcal{W}_l} \hat{f}_t(\mathbf{y}), \qquad \hat{F}_l^{\mathrm{avg}}(\mathbf{y}) \triangleq \frac{1}{W} \sum_{t \in \mathcal{W}_l} \hat{F}_t(\mathbf{y}), \qquad \text{for } \mathbf{y} \in \mathcal{Y}. \tag{160}$$

Lemma 10. *[5, Lemma E.1] Consider $\mathcal{H}_l$ and $\hat{F}_t^{\mathrm{avg}}$ defined in Eq. (160), the estimator $\tilde{l}_l$ in Alg. 5 is an unbiased estimator for $\delta \mathbf{v}_l \cdot \nabla \hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta)$, i.e.,*

$$\mathbb{E}\left[\tilde{l}_l \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l \right] = \delta \mathbf{v}_l \cdot \nabla \hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta). \tag{161}$$

Proof. This proof follows the same lines as the proof of [5, Lemma E.1].
Part I ($z_t \geq 1/2$). Condition on $\mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l$. If $z_l \geq 1/2$ and $\mathbf{u}_l = \mathbf{0}$, we have

$$\mathbb{E}\left[\tilde{l}_l \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] = \mathbb{E}\left[-2\alpha \frac{n}{z_l} f_{t_l}(\mathbf{x}_{t_l}) \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] = -\frac{2\alpha n}{z_l} \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta). \tag{162}$$

If $z_l \geq 1/2$ and $\mathbf{u}_l \neq \mathbf{0}$, we have

$$\mathbb{E}\left[\tilde{l}_l \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] = \mathbb{E}\left[-2\alpha \frac{n}{z_l} f_{t_l}(\mathbf{x}_{t_l}) \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] = \frac{2\alpha n}{z_l} \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta + z_l \delta (\mathbf{v}_l \odot \mathbf{u}_l)). \tag{163}$$

Condition on $\mathcal{H}_{l-1}, \mathbf{v}_l, z_l, t_l$, to obtain

$$\mathbb{E}\left[\tilde{l}_l \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, z_l, t_l \right] = \frac{1}{2}\left(-\frac{2\alpha n}{z_l} \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta) \right) + \frac{1}{2n}\left(\sum_{i=1}^n \frac{2\alpha n}{z_l} \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta + z_l \delta (\mathbf{v}_l \odot \mathbf{e}_i)) \right) \tag{164}$$

$$= \alpha \sum_{i=1}^n \frac{1}{z_l}\left(\hat{f}_{t_l}(z_l \mathbf{y}_l^\delta + z_l \delta (\mathbf{v}_l \odot \mathbf{e}_i)) - \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta) \right) \tag{165}$$

$$= \alpha \sum_{i=1}^n \frac{z_l}{z_l} \delta \mathbf{v}_l \cdot \mathbf{e}_i \frac{\partial \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta)}{\partial x_i} = \alpha \delta \mathbf{v}_l \cdot \nabla \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta). \tag{166}$$

Part II ($z_t < 1/2$). Condition on $\mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l$. If $z_l < 1/2$ and $\mathbf{y}_l = z_l \mathbf{y}_l^\delta$, we have

$$\mathbb{E}\left[\tilde{l}_l \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] = \mathbb{E}\left[-4\alpha n\delta (\mathbf{v}_l \cdot \mathbf{u}_l) f_{t_l}(\mathbf{x}_{t_l}) \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] \tag{167}$$

$$= -4\alpha n\delta (\mathbf{v}_l \cdot \mathbf{u}_l) \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta). \tag{168}$$

If $z_l < 1/2$ and $\mathbf{y}_l = z_l \mathbf{y}_l^\delta + 1/2\mathbf{u}_l$, we have

$$\mathbb{E}\left[\tilde{l}_l \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] = \mathbb{E}\left[4\alpha n\delta (\mathbf{v}_l \cdot \mathbf{u}_l) f_{t_l}(\mathbf{x}_{t_l}) \,\middle|\, \mathcal{H}_{l-1}, \mathbf{v}_l, \mathbf{u}_l, z_l, t_l \right] \tag{169}$$

$$= 4\alpha n\delta (\mathbf{v}_l \cdot \mathbf{u}_l) \hat{f}_{t_l}(z_l \mathbf{y}_l^\delta + 1/2\mathbf{u}_l). \tag{170}$$

Condition on $\mathcal{H}_{l-1}, \mathbf{v}_l, z_l, t_l$, to obtain

$$\mathbb{E}\left[\tilde{l}_l \mid \mathcal{H}_{l-1}, \mathbf{v}_l, z_l, t_l\right] = \sum_{i=1}^{n} \frac{1}{n}\left(1/2(4\alpha n\delta\left(\mathbf{v}_l \cdot \mathbf{e}_i\right)\hat{f}_{t_l}(z_l\mathbf{y}_l^\delta + 1/2\mathbf{e}_i) - 4\alpha n\delta\left(\mathbf{v}_l \cdot \mathbf{e}_i\right)\hat{f}_{t_l}(z_l\mathbf{y}_l^\delta))\right) \quad (171)$$

$$= \sum_{i=1}^{n} 2\alpha\delta\left(\mathbf{v}_l \cdot \mathbf{e}_i\right)\left(\hat{f}_{t_l}(z_l\mathbf{y}_l^\delta + 1/2\mathbf{e}_i) - \hat{f}_{t_l}(z_l\mathbf{y}_l^\delta)\right) \quad (172)$$

$$= \sum_{i=1}^{n} 2/2\alpha\delta\left(\mathbf{v}_l \cdot \mathbf{e}_i\right)\frac{\partial\hat{f}_{t_l}(z_l\mathbf{y}_l^\delta)}{\partial x_i} \quad (173)$$

$$= \alpha\delta\mathbf{v}_l \cdot \nabla\hat{f}_{t_l}(z_l\mathbf{y}_l^\delta). \quad (174)$$

Part III. $(z_l \in [0,1])$ Condition on $\mathcal{H}_{l-1}, \mathbf{v}_l$, to obtain

$$\mathbb{E}\left[\tilde{l}_l \mid \mathcal{H}_{l-1}, \mathbf{v}_l\right] = \int_0^1 \frac{\exp(z-1)}{\alpha}\frac{1}{W}\sum_{t\in\mathcal{W}_l}\left(\alpha\delta\mathbf{v}_l \cdot \nabla\hat{f}_t(z_l\mathbf{y}_l^\delta)\right)dz \quad (175)$$

$$= \int_0^1 \frac{\exp(z-1)}{\alpha}\left(\alpha\delta\mathbf{v}_l \cdot \nabla\hat{f}_l^{\mathrm{avg}}(z_l\mathbf{y}_l^\delta)\right)dz \quad (176)$$

$$= \delta\mathbf{v}_l \cdot \nabla\hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta). \quad (177)$$

$\square$

Lemma 11. *[5, Lemma 3.2] Consider $\mathcal{H}_l$ and $\hat{F}_t^{\mathrm{avg}}$ defined in Eq. (160), the following holds for the supergradient estimator $\tilde{\mathbf{g}}_l$ in Alg. (5):*

$$\mathbb{E}\left[\tilde{\mathbf{g}}_l \mid \mathcal{H}_{l-1}\right] = \nabla\hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta), \qquad \mathbb{E}\left[\|\tilde{\mathbf{g}}_l\|_2^2 \mid \mathcal{H}_{l-1}\right] \leq 16\frac{\alpha^2 n^4 L^2}{\delta^2}. \quad (178)$$

Part I. We have the following

$$\mathbb{E}\left[\tilde{\mathbf{g}}_l \mid \mathcal{H}_{l-1}\right] = \mathbb{E}_{\mathbf{v}\sim\mathbb{S}_{n'-1}}\left[\mathbb{E}\left[\tilde{\mathbf{g}}_l \mid \mathcal{H}_{l-1}, \mathbf{v}\right]\right] = \mathbb{E}_{\mathbf{v}\sim\mathbb{S}_{n'-1}}\left[\frac{n'}{\delta}l_l\mathbf{v}\right] = \mathbb{E}_{\mathbf{v}\sim\mathbb{S}_{n'-1}}\left[n'\mathbf{v}\cdot\nabla\hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta)\mathbf{v}\right] \quad (179)$$

$$= \nabla l_l^{\mathbf{I}}(\mathbf{0}) = \nabla l_l(\mathbf{0}) = \nabla\hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta). \qquad (l_l(\mathbf{x}) \triangleq \mathbf{x}\cdot\nabla\hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta)). \quad (180)$$

Recall that $l^{\mathbf{I}}$ is the **I**-smoothed version of l in Def. 7, which are identical since l is linear.

Proof. This proof follows the same lines as the proof of [5, Lemma 3.2].

Part II. We have the following $|\tilde{l}_l| \leq 4\alpha L$. It follows that $\|\tilde{\mathbf{g}}_l\|_2 \leq 4\frac{\alpha n'^2 L}{\delta}$. Thus, we have

$$\mathbb{E}\left[\|\tilde{\mathbf{g}}_t\|_2^2 | \mathcal{H}_{l-1}\right] \leq 16\frac{\alpha^2 n'^4 L^2}{\delta^2}. \quad (181)$$

We conclude the proof. $\square$

Recall that in the dynamic setting, the decision maker compares its performance to the best sequence of decisions $R(\mathbf{x}_t^\star)_{t\in[T]}$ with a path length P_T from the set

$$\Lambda_{\mathcal{X}}(T, P_T) = \left\{(\mathbf{x}_t)_{t=1}^T \in \mathcal{X}^T : \sum_{t=1}^T \|\mathbf{x}_{t+1} - \mathbf{x}_t\| \leq P_T\right\} \subset \mathcal{X}^T.$$

In the bandit setting, we define the following regularity condition that further restricts the movements of the comparator sequence:

$$\Lambda_{\mathcal{X}}(T, P_T, W) \triangleq \left\{ (\mathbf{x}_t)_{t=1}^T \in \mathcal{X}^T : \sum_{t=1}^T \|\mathbf{x}_{t+1} - \mathbf{x}_t\| \le P_T, \mathbf{x}_t = \mathbf{x}_{t'}, t, t' \in \mathcal{W}_l, l \in [T/W] \right\} \subset \mathcal{X}^T. \quad (182)$$

We extend the definition of the dynamic α-regret in Sec. 6:

$$\alpha - \mathrm{regret}_{T,P_T,W}(\boldsymbol{\mathcal{P}_X}) \triangleq \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \max_{(\mathbf{x}_t^\star)_{t=1}^T \in \Lambda_{\mathcal{X}}(T,P_T,W)} \alpha \sum_{t=1}^T f_t(\mathbf{x}_t^\star) - \sum_{t=1}^T f_t(\mathbf{x}_t) \right\}$$

Note that this dynamic regret definition reduces to the static regret for $P_T = 0$, and the definition in the full-information setting in Sec. 6 for $W = 1$ (i.e., $\alpha - \mathrm{regret}_T(\boldsymbol{\mathcal{P}_X}) = \alpha - \mathrm{regret}_{T,0,W}(\boldsymbol{\mathcal{P}_X})$ and $\alpha - \mathrm{regret}_{T,P_T,1}(\boldsymbol{\mathcal{P}_X}) = \alpha - \mathrm{regret}_{T,P_T}(\boldsymbol{\mathcal{P}_X})$).

Theorem 6 (Restatement). *Under uniformly L-bounded submodular monotone rewards and partition matroid constraint sets, LIRAOCO policy $\boldsymbol{\mathcal{P}_X}$ in Alg. 5 equipped with an OCO policy $\boldsymbol{\mathcal{P}_{\mathcal{Y}_\delta}}$ yields $\alpha - \mathrm{regret}_{T,P_T,W}(\boldsymbol{\mathcal{P}_X}) \le W \cdot \mathrm{regret}_{T/W,P_T}(\boldsymbol{\mathcal{P}_{\mathcal{Y}_\delta}}) + \frac{LT}{W} + 2\alpha\delta r^2 nLT$, where $\mathrm{regret}_{T/W,P_T}(\boldsymbol{\mathcal{P}_{\mathcal{Y}_\delta}})$ is the regret of an OCO policy executed for T/W timeslots.*

Proof. The performance of $\boldsymbol{\mathcal{P}_X}$ is measured against the comparator sequence

$$(\mathbf{x}_{\star,t})_{t=1}^T \in \mathrm{argmax}_{(\mathbf{x}_t)_{t=1}^T \in \Lambda_{\mathcal{X}}(T,P_T,W)} \sum_{t=1}^T f_t(\mathbf{x}_t).$$

From the definition of the extension $\hat{f}$ in Lemma 7, we can always find a point $\mathbf{y}_{\star,t} \in \mathcal{Y}$ such that $f_t(\mathbf{x}_{\star,t}) = \hat{f}_t(\mathbf{y}_{\star,t})$ for every $t \in [T]$, $\sum_{t=1}^T \|\mathbf{y}_{\star,t} - \mathbf{y}_{\star,t}\| = \sum_{t=1}^T \|\mathbf{x}_{\star,t+1} - \mathbf{x}_{\star,t}\| \le P_T$, and $\mathbf{y}_{\star,t} = \mathbf{y}_{\star,t'}$ for $t, t' \in \mathcal{W}_l$ for $l \in [T/W]$ since $\mathbf{x}_{\star,t}$ is an integral point in $\{0,1\}^n$. We define $\mathbf{y}_l^\star = \mathbf{y}_t^\star$ for some $t \in \mathcal{W}_l$, $l \in [T/W]$.

$$\mathbb{E}\left[\sum_{t=1}^T (1 - 1/e) f_t(\mathbf{x}_{\star,t}) - f_t(\mathbf{x}_t) \right] \quad (183)$$

$$= \mathbb{E}\left[\sum_{t=1}^T (1 - 1/e) \hat{f}_t(\mathbf{y}_{\star,t}) - \mathbb{E}\left[f_t(\mathbf{x}_t) \,\middle|\, \mathcal{H}_{t-1} \right] \right] \quad (184)$$

$$= \mathbb{E}\left[\sum_{t=1}^T (1 - 1/e) \hat{f}_t(\mathbf{y}_{\star,t}) - \hat{f}_t(\mathbf{y}_t) \right] \qquad (\, \hat{f}_t(\mathbf{y}) = \mathbb{E}_{\mathbf{x} \sim \hat{\Xi}(\mathbf{y})}[f_t(\mathbf{x})] \text{ in Lemma 7}) \quad (185)$$

$$= \mathbb{E}\left[\sum_{t=1}^T (1 - 1/e) \hat{f}_t(\mathbf{y}_{\star,t}) - \hat{f}_t(\mathbf{y}_{\lceil t/W \rceil}) \right] + \mathbb{E}\left[\sum_{t=1}^T \hat{f}_t(\mathbf{y}_{\lceil t/W \rceil}) - \hat{f}_t(\mathbf{y}_t) \right] \quad (186)$$

$$= \mathbb{E}\left[W \sum_{l=1}^{T/W} (1 - 1/e) \hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{\star,l}) - \hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{\lceil t/W \rceil}^\delta) \right] + \mathbb{E}\left[\sum_{l=1}^{T/W} \hat{f}_t(\mathbf{y}_l^\delta) - \hat{f}_t(\mathbf{y}_{t_l}) \middle| \mathcal{H}_{l-1} \right] \quad (187)$$

$$\le \mathbb{E}\left[W \sum_{l=1}^{T/W} (1 - 1/e) \hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{l,\star}^\delta) - \hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{\lceil t/W \rceil}^\delta) \right] + \frac{LT}{W} + 2\alpha\delta r^2 nLT. \quad (188)$$

The last inequality is obtained using $\hat{f}_t(\mathbf{y}_l^\delta) - \hat{f}_t(\mathbf{y}_{t_l}) \le L$ to yield $\frac{LT}{W}$ term; moreover, consider that $\|\mathbf{y}_\star - \mathbf{y}_{t,\star}^\delta\|_2 \le \delta D$ $(D \le 2r)$ and $\hat{f}$ is $n'L$-Lipschitz (w.r.t. $\|\cdot\|_2$) (see Lemma 7), this yields an

44

additional $2\alpha\delta nr^2 LT$ term ($n' \leq rn$). It remains to bound $\mathbb{E}\left[\sum_{l=1}^{T/W}(1-1/e)\hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{l,\star}^\delta) - \hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{\lceil t/W \rceil}^\delta)\right]$ to complete the proof:

$$\mathbb{E}\left[\sum_{l=1}^{T/W}(1-1/e)\hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{\star,l}^\delta) - \hat{f}_l^{\mathrm{avg}}(\mathbf{y}_{\lceil t/W \rceil}^\delta)\right] \leq \sum_{l=1}^{T/W}\mathbb{E}\left[\nabla\hat{F}_l^{\mathrm{avg}}(\mathbf{y}_l^\delta)\cdot(\mathbf{y}_{l,\star}^\delta - \mathbf{y}_l^\delta)\,\big|\,\mathcal{H}_{l-1}\right] \quad (189)$$

$$\leq \mathrm{regret}_{T/W,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta}\right). \quad (190)$$

The first inequality follows from Lemma 9. The second inequality holds because the gradient estimates $\tilde{\mathbf{g}}_l$ are unbiased estimates of $\nabla\hat{F}_l^{\mathrm{avg}}$ and the reduction of the regret bound of the first-order policy $\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta}$ [8, Lemma 6.5]. We finally obtain the following:

$$\mathbb{E}\left[\sum_{t=1}^{T}(1-1/e)f_t(\mathbf{x}_{\star,t}) - f_t(\mathbf{x}_t)\right] \leq W\cdot\mathrm{regret}_{T/W,P_T}\left(\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta}\right) + \frac{LT}{W} + 2\alpha\delta r^2 nLT. \quad (191)$$

$\square$

LIRAOCO policy $\boldsymbol{\mathcal{P}}_{\mathcal{X}}$ in Alg. 5 equipped with OOMA in Alg. 3 configured with Φ satisfying Asms. 4–6 as $\boldsymbol{\mathcal{P}}_{\mathcal{Y}_\delta}$ has the following regret:[6]

$$\alpha\text{-regret}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq \frac{W}{\eta}\left(4r^2 + 2L_\Phi P_T\right) + \frac{W\eta}{2\rho}\left(\sum_{l\in T/W}\mathbb{E}\left[\|\mathbf{g}_l - \mathbf{g}_l^\pi\|_\star^2\right]\right) + \frac{LT}{W} + 2\alpha\delta r^2 nLT. \quad (192)$$

This bound is obtained considering the regret bound in Theorem 5 and OOMA policy regret bound in Theorem 7. Consider an Euclidean mirror map $\Phi(\mathbf{y}) = \frac{1}{2}\|\mathbf{y}\|_2^2$ with the corresponding bounds in Table 3.

Static non-optimistic setting. In the static non-optimistic setting ($P_T = 0, \mathbf{g}_l^\pi = \mathbf{0}$), we have:

$$\alpha\text{-regret}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) \leq \frac{4r^2 W}{\eta} + \eta\left(8\frac{\alpha^2 r^4 n^4 L^2}{\delta^2}T\right) + \frac{LT}{W} + 2\alpha\delta r^2 nLT \quad (193)$$

$$\leq (12\alpha r^{5/2}n^{3/2} + 1)LT^{4/5} \quad (194)$$

For $\eta = \frac{32^{1/4}}{2^{1/2}\alpha r^{1/2}n^{3/2}L}T^{-3/5}$, where $\delta = (32)^{1/4}(r^{1/2}n^{1/2})T^{-1/5}$, $W = T^{1/5}$.

Dynamic non-optimistic setting. In the static non-optimistic setting ($P_T > 0$ and $\mathbf{g}_l^\pi = \mathbf{0}$). We have the following:

$$\alpha\text{-regret}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) = \mathcal{O}\left(\sqrt{P_T}T^{4/5}\right). \quad (195)$$

For $\eta = \Theta\left(\sqrt{P_T}T^{-3/5}\right)$, $\delta = \Theta\left(T^{-1/5}\right)$, $W = \Theta\left(T^{1/5}\right)$.

Dynamic optimistic setting. In the dynamic optimistic setting ($P_T > 0$ and $\mathbf{g}_l^\pi \neq \mathbf{0}$). We have the following:

$$\alpha\text{-regret}\left(\boldsymbol{\mathcal{P}}_{\mathcal{X}}\right) = \mathcal{O}\left(T^\alpha\sqrt{P_T\sum_{l=1}^{T/W}\|\tilde{\mathbf{g}}_l - \mathbf{g}_l^\pi\|_2^2} + T^{1-\beta} + T^{1-\alpha}\right). \quad (196)$$

---

[6]Note that the result in Theorem 6 readily extends to dynamic regret as in the proof of Theorem 4.

For $\eta = \Theta\left(\sqrt{P_T / \sum_{l=1}^{T/W} \|\tilde{\mathbf{g}}_l - \mathbf{g}_l^\pi\|_2^2}\right)$, $\delta = \Theta\left(T^{-\beta}\right)$, $W = \Theta\left(T^\alpha\right)$. To interpret the above bound, consider for simplicity that $P_T = \Theta\left(1\right)$ and that the optimistic gradients $\mathbf{g}_t^\pi$ are inaccurate (but bounded by some constant). Then, when we take $\eta = \Theta\left(T^{-3/5}\right)$, $\delta = \Theta\left(T^{-1/5}\right)$, $W = \Theta\left(T^{1/5}\right)$ and recover the worst-case regret α-regret$\left(\boldsymbol{\mathcal{P}}_\mathcal{X}\right) = \mathcal{O}\left(T^{4/5}\right)$. However, when the optimistic gradients are accurate the above bound can be much tighter. For example, in the extreme scenario when $\sqrt{\sum_{l=1}^{T/W} \|\tilde{\mathbf{g}}_l - \mathbf{g}_l^\pi\|_2^2} = \epsilon$ for $W = \Theta\left(T\right)$, selecting $\delta = \Theta\left(T^{-1}\right)$, and $\eta = \Theta\left(\frac{1}{\epsilon}\right)$, yields α-regret $= \mathcal{O}\left(\epsilon\right)$.

Specialized Algorithms for Static Regret. The regret attained through the reduction to an arbitrary OCO algorithm (α-regret$\left(\boldsymbol{\mathcal{P}}_\mathcal{X}\right) = \mathcal{O}\left(T^{4/5}\right)$) is higher than with the specific instance of FTRL used by Wan et al. [5] ($\mathcal{O}\left(T^{2/3}\right)$), which however does not extend to the dynamic and optimistic settings.

In particular, in the static regret setting, we can recover the bound by Wan et al. [5] by instantiating the generic reduction in the manner they do, by employing a specialized OCO algorithm with appropriate exploration matrices $\mathbf{H}_l$. Wan et al. [5] showed that when the OCO policy $\mathcal{P}_\mathcal{Y}$ is selected to be FTRL in Alg. 4 with a *self-concordant regularizer* [8] $R(\mathbf{y}) : \mathbb{R}^n \to \mathbb{R}$, the exploration matrix can be configured as $\mathbf{H}_l = \left(\nabla^2 R(\mathbf{y}_l)\right)^{1/2}$. Then, the sampled action $\mathbf{y}_l + \mathbf{H}_l \mathbf{v}_l$ is located on the surface of the *Dikin ellipsoid* centered at $\mathbf{y}_l$ defined as

$$\mathcal{E}(\mathbf{y}_l) \triangleq \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{y} - \mathbf{y}_l\|_l \leq 1\}, \tag{197}$$

where $\|\mathbf{y}\|_l$ is a local norm given by $\|\mathbf{y}\|_l = \sqrt{\mathbf{y}^\intercal \nabla^2 R(\mathbf{y}_l)\mathbf{y}}$. The Dikin ellipsoid has the nice property that it is fully contained in $\mathcal{Y}$. This permits selecting $\delta = 0$, and this reduces the variance of the gradient estimates due to $1/\delta$ scaling. Wan et al. [5] demonstrate that this configuration attains a regret α-regret$\left(\boldsymbol{\mathcal{P}}_\mathcal{X}\right) = \mathcal{O}\left(T^{2/3}\right)$.

G.IV   Alternate Concave Relaxation to Reduce Computational Complexity

In our setup so far, the concave relaxation $\tilde{f}$ was the WTP function "itself". As an additional extension, we present a setting in which a concave relaxation is different from f, while still satisfying Asm. 2. This example demonstrates that our framework and, in particular, Asm. 2 provide additional flexibility to design algorithms with low time complexity.

Consider the following set function

$$f(\mathbf{x}) = \sum_{\ell \in C} c_\ell \left(b_\ell - b_\ell \prod_{j \in S_\ell} (1 - (w_{\ell,j}/b_\ell)x_j)\right), \qquad \mathbf{x} \in \{0,1\}^n, \tag{198}$$

where $b_\ell \in \mathbb{R}_{\geq 0}$, $S_\ell \subseteq V$ is a subset of $V = [n]$, $\mathbf{w}_\ell = (w_{\ell,j})_{j \in S} \in [0,b]^{|S|}$ bounded by b_ℓ, and $c_\ell \in \mathbb{R}_{\geq 0}$ for $\ell \in C$ for some index set C.

This objective appears in many applications such as product ranking [72, 4] and influence maximization [22, 20]. Note for $\mathbf{w}_\ell = \mathbf{1}, \ell \in C$, this set function reduces to the weighted coverage (44).

Lemma 5 implies that the function f belongs to the class of WTP functions; in particular, it can be expressed as follows:

$$f(\mathbf{x}) = \sum_{\ell \in C} c_\ell \left(b_\ell - b_\ell \sum_{S \subseteq S_\ell} \prod_{i \in S} (1 - x_i) \prod_{i \in S} w_{\ell,i}/b_\ell \prod_{j \in S_\ell \setminus S} (1 - w_{\ell,j}/b_\ell)\right). \tag{199}$$

We obtain this equality by considering that $1 - \prod_{i \in S}(1 - x_i) = \Psi_{1,\mathbf{1},S}(\mathbf{x})$ for all $\mathbf{x} \in \{0,1\}^{|S|}$ and, hence:

$$f(\mathbf{x}) = \sum_{\ell \in C} c_\ell \left(b_\ell - b_\ell \sum_{S \subset S_\ell} \prod_{i \in S} w_{\ell,i}/b_\ell \prod_{j \in S_\ell \setminus S} (1 - w_{\ell,j}/b_\ell) \left(\Psi_{1,\mathbf{1},S}(\mathbf{x}) - 1 \right) \right) \tag{200}$$

$$= \sum_{\ell \in C} c_\ell b_\ell \sum_{S \subset S_\ell} \left(\prod_{i \in S} w_{\ell,i}/b_\ell \prod_{j \in S_\ell \setminus S} (1 - w_{\ell,j}/b_\ell) \right) \Psi_{1,\mathbf{1},S}(\mathbf{x}) \triangleq h(\mathbf{x}) \qquad \text{for } \mathbf{x} \in \{0,1\}^n. \tag{201}$$

The second inequality follows from the fact that $1 = \sum_{S \subset S_\ell} \prod_{i \in S} w_{\ell,i}/b_\ell \prod_{j \in S_\ell \setminus S}(1 - w_{\ell,j}/b_\ell)$ for $\ell \in C$ (see Lemma 5).

A straightforward approach to construct a concave relaxation of f is to extend $h(\mathbf{x})$ to fractional values $\mathbf{y} \in [0,1]^n$, i.e., $\tilde{f}(\mathbf{y}) = h(\mathbf{y}), \forall \mathbf{y} \in [0,1]^n$ as in Sec. 5.2. However, *this extension is intractable, since it involves a summation of potentially 2^n terms.* Alternatively, we can consider the concave relaxation given by

$$\tilde{f}(\mathbf{y}) = \sum_{\ell \in C} c_\ell \Psi_{b_\ell, \mathbf{w}_\ell, S_\ell}(\mathbf{y}) \qquad \text{for } \mathbf{y} \in [0,1]^n. \tag{202}$$

It is easy to check that Lemmas 3–4 emply that this is a valid concave function satisfying Asm. 2. This relaxation is of interest for two reasons: (1) it provides an example of the "sandwich" property under which the two functions do not necessarily exactly match over integral values, and (2) it illustrates that carefully designing the concave relaxation can greatly change the computational complexity of the proposed method in Sec. 5.1.

H   Additional Experimental Details & Results

H.I   Datasets & Experiment Setup

We provide additional details about the datasets we use and our experiment setup below. For all experiments reported in Tables 2 and 6, we repeat experiments 5 times per policy; we ensure that the adversary presents the same five sequences across all competitor policies.

Influence Maximization. In the case of influence maximization (see Appendix B.II.1), we use the Zachary Karate Club (ZKC) and the Epinions (Epinions) datasets [73]. We denote by m the number of partitions of the partition matroid. We also sort the number of nodes of the original graph by their degrees and we divide the nodes into $m = 2$ partitions where sorted nodes with even indices are assigned to one partition and the odd indices are assigned to the other. At every timeslot t, the online policy selects $\frac{r_{\text{part}}}{2} = 2$ seeds from each partition. Then, the adversary generates a cascade reachability graph following the independent cascade (IC) model, [22] by independently sampling the edges of the original graph with probability $p = 0.1$. We repeat the above for $T = 100$ timeslots. We follow a similar procedure with the Epinions dataset. This time, we generate a subgraph of the original dataset by sorting the nodes by their outdegrees and keeping the top 200 nodes. Then, we sample the edges of this subgraph with probability $p = 0.1$ and generate $T = 150$ instances. We divide the dataset into $m = 2$ equal-sized partitions u.a.r and select at most $\frac{r_{\text{part}}}{2} = 5$ from each partition.

Facility Location. For this application, we experiment on a subset of the MovieLens 10M (MovieLens) dataset [56]. The original dataset has more than $10M$ ratings in total from 71567 users for 10681 movies. We sort the users based on the number of movies they have rated and keep the top $T = 294$ of these users. Then, we take the user with the most number of ratings and only keep the movies that have been rated by this user in our subset. This leaves us with $|C| = 21$ movies. In the facility location

context, we treat movies as facilities, users as customers, and normalized ratings $w_{i,j}$ as the utilities. For the partition matroid constraint, we divide the movies into $m = 6$ partitions based on the first genre listed for each movie. At each timeslot, the adversary selects a user in the order they show up in the original dataset. The online algorithm maximizes the total utility by selecting $\frac{r_{\text{part}}}{6} = 1$ movie from each genre. For the uniform matroid constraint, we set $r_{\text{uni}} = 6$.

Team Formation. Our experiments for team-formation focus on objective functions that are monotone, submodular quadratic functions (see Appendix B.II.3). We generate 5 monotone and submodular quadratic functions as follows: we generate the h vector by sampling n numbers from a normal distribution (with $\mu = 60$ and $\sigma = 20$) and enforcing every coordinate of $\mathbf{h}$ to be between 0 and 100. We generate the $\mathbf{H}$ (symmetric) matrix by sampling each entry $H_{i,j} = H_{j,i}$ from a normal distribution (with $\mu = -20$ and $\sigma = 10$). We enforce every entry to be less than or equal to 0 and the main diagonal of H to be equal to 0. We shrink the rows (and columns) of $\mathbf{H}$ by multiplying them with a constant factor until the function satisfies the monotonicity property. We consider a pool of $n = 100$ individuals. At timeslot t, the adversary selects one of the five functions uniformly at random and the online algorithm forms a team maximizing the function. We consider a total of $T = 100$ timeslots. For the uniform matroid constraint, the online algorithm can select at most $r_{\text{uni}} = 2$ individuals. For the partition matroid constraint, we divide the individuals into $m = 2$ partitions and the online algorithm selects at most $\frac{r_{\text{part}}}{2} = 2$ individuals from each partition.

Synthetic Weighted Coverage. We generate two variants of the weighted-coverage problem with a ground set size $n = 20$, designed to study reward function distribution shifts in the dynamic regret setting (Appendix H.III), and an additional variant designed to study optimistic learning (Appendix H.III). In all cases, we consider a uniform matroid constraint with rank $r_{\text{uni}} = 5$. The objectives are designed such that the maximizing decisions are disjoint under the two objectives. Each objective is the sum of $|C| = 63$ threshold potentials. Our first (stationary setting) problem considers a fixed function equal to the average of the two objectives. In the second (time-varying setting) problem, we set the rewards to be equal to the first objective for the first $\frac{T}{2} = 25$ timeslots, and then abruptly set the rewards to be the second objective for the remaining $\frac{T}{2} = 25$ timeslots. Finally, in the third (also time-varying) setting, we alternate between the two objectives at each timeslot.

H.II   Policy Implementation

All timing experiments were run on a machine with a Broadwell CPU and 128GB RAM. All policies, including ones by other researchers (`FSF★` and `TabularGreedy`), were implemented by us in Python. For solving convex optimization problems, we used the CVXPY package, which is an open-source python library.

Even though hyperparameters differ per algorithm, for a fair comparison, we explore the same number of configurations across all our algorithms and baselines. For the `RAOCO-OGA` algorithm we try η values from the set $\{0.001, 0.01, 0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 6, 8, 10\}$. For the `RAOCO-OMA`, `FSF★`, and `TabularGreedy` algorithms; we perform grid search where $(\eta, \gamma) \in \{0.05, 0.1, 6.5, 10\} \times \{0.001, 0.01, 0.05, 0.1\}$, $(\eta, \gamma) \in \{1, 10, 75, 100\} \times \{0, 0.001, 0.01, 0.1\}$, and $(\eta, c_p) \in \{0.1, 1, 10, 160\} \times \{1, 2, 4, 8\}$ respectively. In order to choose these values, we first perform an exponential search and then conduct a linear search on the interval with maximum reward.

H.III   Additional Results

Fractional policies. We report the fractional rewards $\bar{F}_\mathcal{Y}$ in Table 6. We observe that both `RAOCO-OGA` and `RAOCO-OMA` significantly outperform `Random`. Note that in the `MovieLens` dataset all algorithms, including `Random`, attain almost optimal rewards. In comparison with the integral rewards (Table 2),

Dataset	Problem	T	n	$|C|$	$\min(|S_\ell|)$	$\max(S_\ell)$	$\mathrm{avg}(|S_\ell|)$	m	r_{part}	r_{uni}
ZKC	Inf. Max.	100	34	34	1	8	1.27	2	4	4
Epinions	Inf. Max.	150	200	200	1	47	2.80	2	10	10
MovieLens	Fac. Loc.	294	21	21	1	21	1.80	6	6	6
SynthTF	Team Form.	100	100	4951	2	100	2.02	2	4	2
SynthWC	Weight. Cov.	50	20	24	1	4	2.04	4	8	5

Table 5: Dataset properties and experiment parameters. The columns T, n, $|C|$ specify the number of timeslots, the number of elements of the matroid, and the number of threshold potentials in each objective function, respectively. The operators avg, min, and max are applied on possible values of the input parameters S_ℓ of the WTP class in Eq. (4). The columns m, r_{part}, and r_{uni} specify the number of partitions in the partition matroid, the rank of the partition matroid, and the rank of the uniform matroid, respectively, used across the different problem instances.

datasets & constr.		$F^\star$	t	RAOCO-OGA $\bar{F}_\mathcal{Y}/F^\star$	RAOCO-OGA avg. time per timeslot (s)	RAOCO-OMA $\bar{F}_\mathcal{Y}/F^\star$	RAOCO-OMA avg. time per timeslot (s)	FSF* avg. time per timeslot (s)	TabularGreedy avg. time per timeslot (s)	Random $\bar{F}_\mathcal{Y}/F^\star$	Random std. dev.	Random avg. time per timeslot (s)
ZKC	Uni.	0.234	33	0.912		**0.957**				0.646	0	
			66	0.929	7.62×10^{-3}	**0.965**	7.71×10^{-3}	0.158	7.63×10^{-2}	0.643	0	3.32×10^{-3}
			99	0.948		**0.981**				0.640	0	
	Part.	0.83	33	0.995		**0.996**				0.957	0	
			66	0.99	6.85×10^{-3}	**0.993**	7.09×10^{-3}	$\times$	1.97×10^{-2}	0.953	1.34×10^{-16}	2.84×10^{-3}
			99	0.993		**0.996**				0.955	0	
Epinions	Uni.	0.171	50	**0.919**		0.905				0.724	0	
			100	**0.937**	7.27×10^{-2}	0.932	7.4×10^{-2}	14.2	6.69	0.711	0	1.93×10^{-2}
			149	**0.946**		0.943				0.720	0	
	Part.	0.171	50	**0.918**		0.902				0.724	0	
			100	**0.938**	7.14×10^{-2}	0.929	7.41×10^{-2}	$\times$	3.3	0.711	8.12×10^{-17}	1.86×10^{-2}
			149	**0.949**		0.941				0.720	0	
MovieLens	Uni.	0.407	98	0.838		**0.875**				0.833	0	
			196	0.817	1.26×10^{-2}	**0.820**	1.25×10^{-2}	1.29×10^{-2}	7.83×10^{-3}	0.774	0	1.18×10^{-3}
			293	0.868		**0.893**				0.799	0	
	Part.	0.419	98	0.896		**0.96**				0.882	0	
			196	0.874	1.2×10^{-2}	**0.911**	4.69×10^{-3}	$\times$	1.71×10^{-2}	0.852	0	1.04×10^{-3}
			293	0.940		**0.949**				0.903	0	
SynthTF	Uni.	200	33	0.983		**0.983**				0.611	0	
			66	0.992	0.515	**0.995**	0.517	45.5	44.4	0.609	0	0.224
			99	0.995		**0.999**				0.611	0	
	Part.	400	33	0.982		**0.985**				0.611	0	
			66	0.99	0.526	**0.992**	0.518	$\times$	22.2	0.609	0	0.228
			99	0.993		**0.995**				0.611	0	

Table 6: Average cumulative reward $\bar{F}_\mathcal{Y}$ ($t = T/3, 2T/3, T$), normalized by fractional optimal $F^\star$, of fractional policies (RAOCO and Random) across different datasets and constraints, along with average execution time per timeslot (in seconds). Highest $\bar{F}_\mathcal{Y}/F^\star$ values are indicated in bold. The policies FSF* and TabularGreedy by construction only produce integral solutions (whose performance is reported in Table 2). The policy FSF* only operates on uniform matroids. The standard deviation of the rewards of RAOCO policies is 0, since the fractional rewards are deterministic, and thus omitted from the table. RAOCO combined with OGA or OMA outperforms Random, while almost reaching the optimal value 1. As Random also performs well on MovieLens, this indicates that the (static) offline optimal is quite poor for this reward sequence. With respect to the execution time, RAOCO policies are consistently faster by orders of magnitude than competing policies FSF* and TabularGreedy, and pretty close to (trivial to execute) policy Random; an exception is TabularGreedy on MovieLens under a partition matroid constraint, where its execution time is slightly better, but still comparable to RAOCO policies.

RAOCO-OGA slightly outperforms RAOCO-OMA in the Epinions dataset. Overall, in all other datasets, fractional results are on par with the integral results reported in Table 2.

Execution times. We report the average execution time per timeslot (avg. time per timeslot) in Table 6. We observe that RAOCO-OGA and RAOCO-OMA, significantly outperform all other policies, except Random. Particularly, in the ZKC dataset RAOCO algorithms are faster than the competitors by one order of magnitude, while in the Epinions dataset RAOCO algorithms are faster than the competitors by two orders of magnitude.

Dynamic Regret. Recall that we consider two settings for the `SynthWC` dataset, which we designed to study dynamic regret: one with static rewards, and one with a distribution shift. Recall that optimal decision under the two objectives are disjoint; therefore, the different policies need to relearn for the different objective. We consider the following configurations for the different policies: `OGA` ($\eta = 0.001$), `OMA` ($\eta = 0.05, \gamma = 0$) configured with negative-entropy, `OMA`* ($\eta = 0.05, \gamma = 0.02$) configured with negative-entropy, `FSF`* ($\eta = 0.05, \gamma = 0.02$), `TabularGreedy` ($\eta = 0.05, c_p = 1$).

The results are provided in Fig. 1. We observe that all the policies have similar performance in the stationary setting and are able to match the performance of the static optimum. In the non-stationary setting, we observe that that `OGA`, `OMA`, and `FSF`* are robust to to non-stationarity and are able to match the performance of a dynamic optimum. `TabularGreedy` is less robust to non-stationarity and is only able to match the performance of a static optimum. We observe that the shifted version of `OMA` increases its robustness and provides a slightly better performance compared to `OGA`.

Optimistic Learning. We now consider the third variant of `SynthWC` dataset, in which we alternate between the two objectives at each timeslot. Under this setting, it is difficult to track the dynamic optimum since $P_T = \Omega(T)$ (see our theoretical result in Sec. 6; moreover, the regret $\mathcal{O}\left(\sqrt{P_T T}\right)$ is order-wise optimal [74]). However, if we incorporate predictions, the large P_T can be dampened when the predictions are accurate. We provide as predictions the future supergradient $\mathbf{g}_t$ with added Gaussian noise with mean $\mathbf{0}$ and standard deviation $n_\sigma \in \{10, 200\}$.

The results are provided in Fig. 4. We observe that when the predictions are accurate `Optimistic OGA` is able to exploit these predictions to match the performance of the dynamic optimum. We note that for larger learning rate, the performance improves as the policy can follow aggressively the predictions instead of the past rewards, whereas classical `OGA` diverges for this learning rate configuration. The main difficulty with optimistic leaning is that the learning rate selection depends on the quality of the received predictions. This motivates the next experiment that considers a meta-learning setup that learns the appropriate learning rates in a parameter-free fashion.

Meta-Policies. We consider the `SynthWC` dataset and a fixed function equal to the average of the two objectives. We construct a parameter-free `OGA` that does not require a learning rate. The learning rate is adapted through the received gradients; such a learning rate schedule is known as "self-confident" [75]. This policy acts as a meta-policy that learns over different policies configured with different learning rates $\eta \in \left\{5 \times 10^{-4}, 1 \times 10^{-3}, 2 \times 10^{-3}, 4 \times 10^{-3}\right\}$. The results are provided in Fig. 5. We observe that the meta-policy quickly learns the best learning rate.
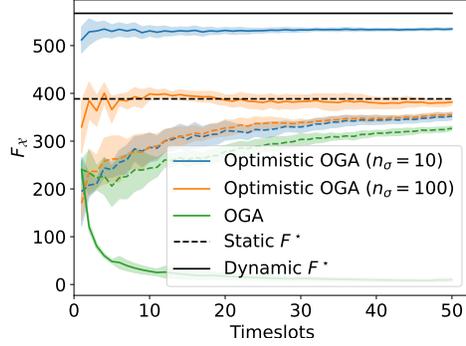
Figure 4: Average cumulative reward $\bar{F}_{\mathcal{X}}$ of the different policies under `SynthWC` dataset under a non-stationary setup: the objective is changed at every timeslot. The algorithms `Optimistic OGA` and `OGA` are executed with different learning rates under different prediction accuracy (noise with std. dev. $n_\sigma \in \{10, 100\}$). The larger learning rate is depicted by a solid line. The area depicts the standard deviation over 5 runs.
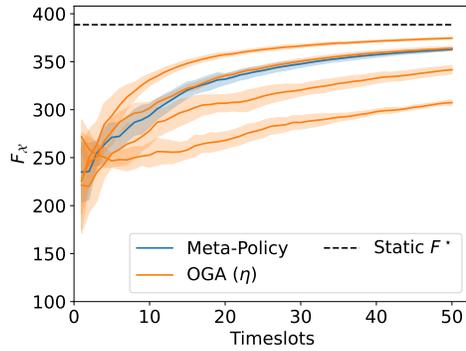

Figure 5: Average cumulative reward $\bar{F}_{\mathcal{X}}$ of `OGA` under `SynthWC` dataset for different learning rates. The meta-policy can learn the best configuration of OGA without tuning the learning rate. The area depicts the standard deviation over 5 runs.