No-regret algorithms for online k-submodular maximization

Tasuku Soma
The University of Tokyo
tasuku_soma@mist.i.u-tokyo.ac.jp

Abstract

We present a polynomial time algorithm for online maximization of k-submodular maximization. For online (nonmonotone) k-submodular maximization, our algorithm achieves a tight approximate factor in the approximate regret. For online monotone k-submodular maximization, our approximate-regret matches to the best-known approximation ratio, which is tight asymptotically as k tends to infinity. Our approach is based on the Blackwell approachability theorem and online linear optimization, and provides simpler and clearner analysis.

1 Introduction

Submodular functions have a wide veriety of applications in combinatorial optimization, economics, communication, and machine learning (Fujishige, 2005; Krause and Golovin, 2014). A set function $f : 2^V \to \mathbb{R}$ on a ground set V is called a submodular function if it satisfies $f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y)$ for all $X, Y \subseteq V$. Equivalently, f is submodular if it satisfies the *diminishing return property*: $f(X \cup \{j\}) - f(X) \geq f(Y \cup \{j\}) - f(Y)$ for all $X \subseteq Y$ and $j \in V \setminus Y$. In the last two decades, *submodular maximization* has been studied extensively in theoretical computer science (Calinescu et al., 2011; Buchbinder et al., 2015), machine learning (Krause and Golovin, 2014), and viral marketing (Kempe et al., 2003). Although submodular maximization is NP-hard in general, constant-factor approximation algorithms have been devised for various constraints (Calinescu et al., 2011; Buchbinder et al., 2015).

Recently, the paradigm of *"optimization as a process"* has been proposed in the context of *online learning* (Hazan, 2016; Cesa-Bianchi and Lugosi, 2006). The goal of online learning is making a better decision in the face of uncertainty. Formally, let us consider the following repeated two-player game between a player and an adversary. At each tth round ($t \in [T] := \{1, \dots, T\}$), the player must select an action $x_t \in K$ (possibly in a randomized manner). After the choice of x_t, the adversary reveals a reward function $f_t : K \to [0, 1]$ in the round, and the player gains $f_t(x_t)$. The performance metric of the player's algorithm is the *regret*:

$$\mathrm{regret}(f_1, \dots, f_T) = \max_{x \in K} \sum_{t \in [T]} f_t(x) - \sum_{t \in [T]} f_t(x_t).$$

That is, the regret is the difference between the player's total gain and the gain of the best fixed action in hindsight. A player's algorithm is said to be *no regret* if the expectation of the regret is sublinear: $\mathbf{E}[\mathrm{regret}(f_1, \dots, f_T)] = o(T)$, where the expectation is taken under the randomness in the player.

Online submodular maximization is an online learning problem in which the action set is a set family $\mathcal{C} \subseteq 2^V$ and the reward functions f_t are submodular functions on V. Since submodular maximization is NP-hard even in the offline setting, it is reasonable to relax the definition of the regret to the *α-regret*:

$$\mathrm{regret}_\alpha(f_1, \dots, f_T) = \alpha \max_{X \in \mathcal{C}} \sum_{t \in [T]} f_t(X) - \sum_{t \in [T]} f_t(X_t),$$

where $\alpha > 0$ is a constant. Intuitively, α corresponds to the offline approximation ratio. A player's algorithm is said to be *no α-regret* if $\mathbf{E}[\mathrm{regret}_\alpha(f_1, \dots, f_T)] = o(T)$. Streeter and Golovin (2009) presented the first no $(1 - 1/e)$-regret algorithm for online monotone submodular maximization under a cardinality constraint ($\mathcal{C}$ is the set of subsets satisfying the cardinality constraint and f_t are monotone submodular functions). Golovin et al. (2014) extended this algorithm to a matroid constraint, generalizing a well-known *continuous greedy algorithm* (Calinescu et al., 2011). Recently, Roughgarden and Wang

Proceedings of the 22[nd] International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89.

for $t = 1, \ldots, T$ **do**
- A player (randomly) plays $\mathbf{x}_t \in (k+1)^V$.

- An adversary reveals a k-submodular function $f_t : (k+1)^V \to [0,1]$ to the player as a value oracle.
- The player gains reward $f_t(\mathbf{x}_t)$.

end for

Figure 1: The online k-submodular maximization protocol.

(2018) proposed a no 1/2-regret algorithm for (unconstrained) online nonmonotone submodular maximization. Their algorithm is based on the *double greedy algorithm* (Buchbinder et al., 2015); at its core, they designed an online learning algorithm with two actions with a stronger regret guarantee.

1.1 Our contribution

This paper examines online maximization of *k-submodular* functions. k-submodular functions are generalizations of submodularity and bisubmodularity, introduced by Huber and Kolmogolov (Huber and Kolmogorov, 2012). Formally, k-submodular functions are defined on $(k+1)^V = \{0, 1, \ldots, k\}^V$. A function $f : (k+1)^V \to \mathbb{R}$ is k-submodular if for any $\mathbf{x}, \mathbf{y} \in (k+1)^V$, $f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} \sqcup \mathbf{y}) + f(\mathbf{x} \sqcap \mathbf{y})$, where $\sqcup$ and $\sqcap$ are generalized "union" and "intersection" in $(k+1)^V$, respectively (see Section 2 for the formal definition). Indeed, if $k = 1, 2$, k-submodularity is equivalent to submodularity and bisubmodularity, respectively. The concepts of bisubmodularity and k-submodularity have numerous applications in valued CSP, delta matroids, generalized influence maximization, and image segmentation (Huber and Kolmogorov, 2012; Fujishige, 2005; Fujishige and Iwata, 2005; Ohsaka and Yoshida, 2015; Hirai and Oki, 2017).

For offline k-submodular maximization, Iwata et al. (2016) gave a 1/2-approximation algorithm. The approximation ratio is tight even for $k = 1$, i.e., submodular maximization (Feige et al., 2011). They also devised a $\frac{k}{2k-1}$-approximation algorithm for *monotone* k-submodular maximzation and the approximation ratio is asymptotically tight.

The main results of this paper are as follows:

- For online k-submodular maximization, we devise a polynomial-time algorithm whose expected 1/2-regret is bounded by $O(nk\sqrt{T})$, where $n = |V|$. This result generalizes the previous algorithm of Roughgarden and Wang (2018) for online submodular maximization.

- For online monotone k-submodular maximization, we present a polynomial-time algorithm whose expected $\frac{k}{2k-1}$-regret is $O(nk\sqrt{T})$.

See Table 1 for comparison of our results and previous results.

Technical Ideas To extend the algorithm of Iwata et al. (2016) to the online setting, we must consider an auxiliary online learning problem, which we call a *k-submodular selection game*. We show that it is sufficient to design an online algorithm for k-submodular selection games with a stronger regret guarantee, which is not obtained by using a standard online learning algorithm such as multiplicative weight update (Arora et al., 2012). To this end, we exploit *Blackwell's approachability theorem*[1] (Blackwell, 1956) and *online linear optimization (OLO)*. The Blackwell approachability theorem is a powerful generalization of von Neumann's minimax theorem for finite two-player games. In the online learning literature, the Blackwell approachability theory has been exploited to demonstrate the existence of no-regret algorithms for various problems, such as online learning with the internal and generalized regret, and well-calibrated forecasters (see Cesa-Bianchi and Lugosi (2006) and references therein). We exploit the Blackwell approachability theorem to design an algorithm with the desired stronger regret guarantee. To obtain a concrete regret bound, we use a beautiful duality result between approachability and OLO (Abernethy et al., 2011). More precisely, we use their framework to obtain an online algorithm for k-submodular selection games by converting an OLO algorithm.

To demonstrate the flexibility of our approach based on Blackwell's theorem, we show that the algorithm for the nonmonotone case can be easily modified for the monotone case with a stronger approximation ratio $\frac{k}{2k-1}$. Furthermore, our algorithm and analysis work even for an *adaptive adversary*. An *oblivious adversary* fixes f_t ($t \in [T]$) before the first round, whereas an adaptive adversary can select f_t after seeing $\mathbf{x}_t$. Since our approach is conceptually simpler than previous work (Roughgarden and Wang, 2018), it almost immediately extends to an adaptive adversary.

[1]The possibility of using of Blackwell's approachability theorem was mentioned in Roughgarden and Wang (2018) without detail in a footnote. They designed an alternative algorithm for a similar problem without using Blackwell's theorem.

Table 1: Summary of previous results and our results.

	offline	online
$k = 1$	$\frac{1}{2}$-approx Buchbinder et al. (2015)	$\frac{1}{2}$-regret $O(n\sqrt{T})$ Roughgarden and Wang (2018)
$k = 2$	$\frac{1}{2}$-approx Ward and Živný (2016)	**$\frac{1}{2}$-regret $O(n\sqrt{T})$ [this work]**
general k	$\frac{1}{2}$-approx Iwata et al. (2016)	**$\frac{1}{2}$-regret $O(nk\sqrt{T})$ [this work]**
general k (monotone)	$\frac{k}{2k-1}$-approx Iwata et al. (2016)	**$\frac{k}{2k-1}$-regret $O(nk\sqrt{T})$ [this work]**

1.2 Related work

An important special case of k-submodular functions is the *bisubmodular* function. Singh et al. (2012) studied maximizing a bisubmodular function[2]. General k-submodular maximization was first studied by Ward and Živný (2016). They devised a $1/(1 + \sqrt{k/2})$-approximation algorithm for k-submodular maximization. Iwata et al. (2016) presented a randomized algorithm with an improved and tight approximation factor of 1/2 for k-submodular maximization. A derandomized version of their algorithm (for the monotone case) was developed by Oshima (2017). Ohsaka and Yoshida (2015) studied monotone k-submodular maximization under a cardinality constraint. Later, Sakaue (2017) generalized it to a matroid constraint. A similar but different setting, namely *streaming* submodular optimization, also has been actively studied, e.g., see Feldman et al. (2018); Mirzasoleiman et al. (2018). Recenetly, there has been a series of work considering *continous* submodular maximization (Bian et al., 2017a,b; Chen et al., 2018a,b; Niazadeh et al., 2018). To the best of our knowledge, a similar continous analogue of k-submodular functions is unknown.

Online learning of discrete structure is called *online structured learning*. Efficient online algorithms were developed for various discrete structures, such as shortest paths and matroid basis (Takimoto and Warmuth, 2003; Suehiro et al., 2012). Most of these studies focused on optimizing *linear* reward/loss functions (under a constraint), whereas our paper studies *nonlinear* functions (without a constraint).

1.3 Organization

The reminder of this paper is organized as follows. Section 2 introduces k-submodularity, Blackwell's approachability theorem, and OLO. Section 3 describes our algorithm for online k-submodular maximization along with k-submodular selection games. Section 4 presents our algorithm for online monotone k-submodular maximization. Section 5 provides several examples of online k-submodular maximization in machine learning. We conclude the paper in Section 6.

[2]Note that they used different terminology, *directed bisubmodular* functions, to describe such functions.

2 Preliminaries

2.1 Notation

For a positive integer n, we denote the set $\{1, \ldots, n\}$ by $[n]$. The probability simplex in $\mathbb{R}^k$ is denoted by Δ_k. The sets of nonnegative and nonpositive reals are denoted by $\mathbb{R}_+$ and $\mathbb{R}_-$, respectively. The Euclidian norm is denoted by $\|\cdot\|$. The jth standard unit vector is denoted by $\mathbf{e}_j$. The distance between a point $\mathbf{x}$ and a set S is defined as $\mathrm{dist}(\mathbf{x}, S) := \inf_{\mathbf{y}\in S}\|\mathbf{x} - \mathbf{y}\|$. The orthogonal projection of a point $\mathbf{x}$ onto a set S is denoted by $\mathrm{proj}_S(\mathbf{x})$.

2.2 k-submodular functions

Let k be a positive integer. Throughout the paper, let $V = [n]$ be a ground set. Define $(k + 1)^V = \{0, 1, \ldots, k\}^V$. For $\mathbf{x} \in (k+1)^V$, we denote $\mathrm{supp}(\mathbf{x}) = \{j \in V : x(j) \neq 0\}$. For a function $f : (k + 1)^V \to \mathbb{R}$, $\mathbf{x} \in (k + 1)^V$, and $j \notin \mathrm{supp}(\mathbf{x})$, we define

$$\Delta_{j,i} f(\mathbf{x}) := f(\mathbf{x} + i\mathbf{e}_j) - f(\mathbf{x}),$$

where $\mathbf{x} + i\mathbf{e}_j$ is a vector obtained by setting the jth entry of $\mathbf{x}$ to i. Since $x(j) = 0$, this is the standard addition in $\mathbb{R}^V$. Let us define binary operators $\sqcup$ and $\sqcap$ on $\{0, 1, \ldots, k\}$ as

$$i \sqcup i' = \begin{cases} \max\{i, i'\} & \text{if either } i = 0,\ i' = 0 \text{ or } i = i' \\ 0 & \text{otherwise} \end{cases}$$

$$i \sqcap i' = \begin{cases} \min\{i, i'\} & \text{if either } i = 0,\ i' = 0 \text{ or } i = i' \\ 0 & \text{otherwise} \end{cases}$$

We extend these binary operations to $(k+1)^V$ so that the operations are applied entry-wise: for $\mathbf{x}, \mathbf{y} \in (k+1)^V$, define $\mathbf{x} \sqcup \mathbf{y}, \mathbf{x} \sqcap \mathbf{y} \in (k+1)^V$ as

$$(\mathbf{x} \sqcup \mathbf{y})(j) = x(j) \sqcup y(j) \quad (j \in V)$$
$$(\mathbf{x} \sqcap \mathbf{y})(j) = x(j) \sqcap y(j) \quad (j \in V).$$

A function $f : (k+1)^V \to \mathbb{R}$ is *k-submodular* if

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} \sqcup \mathbf{y}) + f(\mathbf{x} \sqcap \mathbf{y}) \tag{1}$$

for arbitrary $\mathbf{x}, \mathbf{y} \in (k+1)^V$.

Define a partial order on $(k+1)^V$ by $\mathbf{x} \leq \mathbf{y}$ if $\mathbf{x} \sqcap \mathbf{y} = \mathbf{x}$. We say that $f : (k+1)^V \to \mathbb{R}$ is *monotone* if $f(\mathbf{x}) \leq f(\mathbf{y})$ for arbitrary $\mathbf{x} \leq \mathbf{y}$.

Lemma 2.1 (Ward and Živný (2016))**.** *The k-submodularity is equivalent to the following two conditions:*

Pairwise monotonicity $\Delta_{j,i} f(\mathbf{x}) + \Delta_{j,i'} f(\mathbf{x}) \geq 0$ *for $i \neq i'$, $\mathbf{x} \in (k+1)^V$, and $j \notin \mathrm{supp}(\mathbf{x})$.*

Orthant submodularity $\Delta_{j,i} f(\mathbf{x}) \geq \Delta_{j,i} f(\mathbf{y})$ *for i, $\mathbf{x} \leq \mathbf{y}$, and $j \notin \mathrm{supp}(\mathbf{y})$.*

A vector $\mathbf{x} \in (k+1)^V$ can be regarded as a k-subpartition of V. That is, $(k+1)^V$ can be regarded as the set of $(X_1, \ldots, X_k)$ ($X_i \subseteq V$, $X_i \cap X_{i'} = \emptyset$ if $i \neq i'$). The correspondence is given by $x(j) = i$ if and only if $j \in X_i$ (we conventionally regard that $x(j) = 0$ if and only if j is in none of X_i). For $k = 1$, k-submodularity (1) is equivalent to submodularity, $f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y)$ for $X, Y \in 2^V$. For $k = 2$, it is equivalent to bisubmodularity (Fujishige, 2005),

$$\begin{aligned} &f(X_1, X_2) + f(Y_1, Y_2) \\ &\geq f((X_1 \cup Y_1) \setminus (X_1 \cap Y_1), (X_2 \cup Y_2) \setminus (X_2 \cap Y_2)) \\ &\quad + f(X_1 \cap Y_1, X_2 \cap Y_2), \end{aligned}$$

for $(X_1, X_2), (Y_1, Y_2) \in 3^V$. Ward and Živný (2016) showed that a submodular function $g : 2^V \to \mathbb{R}_+$ can be embedded into a bisubmodular function $f : 3^V \to \mathbb{R}_+$ as

$$f(S, T) = g(S) + g(V \setminus T) - g(T) \tag{2}$$

preserving the approximation ratio. That is, an α-approximate maximizer of f corresponds to an α-approximate maximizer of g, for arbitrary $\alpha > 0$. This embedding demonstrates that our algorithm for online k-submodular maximization corresponds to the algorithm of Roughgarden and Wang (2018) for online submodular maximization.

A useful fact of k-submodular maximization is that there always exists a maximizer corresponding to a partition of V.

Lemma 2.2 (Ward and Živný (2016))**.** *Let $k \geq 2$. For any k-submodular function f, there exists $\mathbf{o} \in \mathrm{argmax}_{\mathbf{x} \in (k+1)^V} f(\mathbf{x})$ such that $\mathrm{supp}(\mathbf{o}) = V$.*

2.3 Blackwell's approachability theorem

The celebrated Blackwell approachability theorem (Blackwell, 1956) is a powerful generalization of the von Neumann minimax theorem for two-player zero-sum games. Our presentation mostly follows Abernethy et al. (2011). Let $X \subseteq \mathbb{R}^m$ and $Y \subseteq \mathbb{R}^n$ be convex sets. Let $\boldsymbol{\ell} : X \times Y \to \mathbb{R}^k$ be a biaffine function, i.e, $\boldsymbol{\ell}(\cdot, \mathbf{y})$ is affine for any $\mathbf{y} \in Y$ and vice versa. Let $S \subseteq \mathbb{R}^k$ be a closed convex set. We call a tuple $(X, Y, \boldsymbol{\ell}, S)$ a *Blackwell instance*. We say that:

- S is *satisfiable* if $\exists \mathbf{x} \in X \forall \mathbf{y} \in Y : \boldsymbol{\ell}(\mathbf{x}, \mathbf{y}) \in S$.
- S is *response-satisfiable* if $\forall \mathbf{y} \in Y \exists \mathbf{x} \in X : \boldsymbol{\ell}(\mathbf{x}, \mathbf{y}) \in S$.
- S is *halfspace-satisfiable* if an arbitrary halfspace H containing S is satisfiable.
- S is *approachable* if there exists an algorithm $\mathcal{A}$ which outputs an element in X, such that for any sequence $(\mathbf{y}_t)_{t \in [T]} \subseteq Y$, $\mathrm{dist}\left(\frac{1}{T} \sum_{t \in [T]} \ell(\mathbf{x}_t, \mathbf{y}_t), S\right) \to 0$ as $T \to \infty$, where $\mathbf{x}_t := \mathcal{A}(\mathbf{y}_1, \ldots, \mathbf{y}_{t-1})$ ($t \in [T]$).

Theorem 2.3 (The Blackwell approachability theorem (Blackwell, 1956))**.** *For a Blackwell instance $(X, Y, \boldsymbol{\ell}, S)$, the following conditions are equivalent:*

1. *S is approachable.*
2. *S is halfspace-satisfiable.*
3. *S is response-satisfiable.*

A *halfspace oracle* $\mathcal{O}$ is an oracle that takes a halfspace H with $S \subseteq H$ as input and returns $\mathcal{O}(H) = \mathbf{x}_H \in X$. A halfspace oracle is said to be *valid* if $\boldsymbol{\ell}(\mathbf{x}_H, \mathbf{y}) \in H$ for any $\mathbf{y} \in Y$. Note that the existence of a valid halfspace oracle is equivalent to the halfspace-satisfiability of S. Even if a valid halfspace oracle exists, its efficient computation depends on the geometry of the feasible regions X and Y. If X and Y are polytopes, then a halfspace oracle can be constructed by linear programming (LP) as follows.

Let $H := \{\mathbf{z} : \boldsymbol{\theta}^\top \mathbf{z} \geq \beta\}$ be a halfspace. Since $\boldsymbol{\ell}$ is biaffine, $\boldsymbol{\theta}^\top \boldsymbol{\ell}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top P \mathbf{y} + \mathbf{b}^\top \mathbf{y} + c$ for some matrix P, a vector $\mathbf{b}$, and a constant c. For computing a valid halfspace oracle, we can assume that $c = 0$ without loss of generality. Then, finding $\mathbf{x}_H$ is a responce of a valid halfspace oracle if

$$\min_{\mathbf{y} \in Y} (\mathbf{x}_H^\top P \mathbf{y} + \mathbf{b}^\top \mathbf{y}) \geq \beta.$$

Algorithm 1 Online Gradient Descent for OLO (Zinkevich, 2003)

Input: a compact convex set $K \subseteq \mathbb{R}^k$ and learning rate $\eta > 0$.

Let $\mathbf{x}_0 \in K$ be an arbitrary point.
for $t \in [T]$ **do**
    Play $\mathbf{x}_t$ and observe $\mathbf{f}_t$.
    Let $\mathbf{y}_{t+1} = \mathbf{x}_t - \eta \mathbf{f}_t$ and $\mathbf{x}_{t+1} = \mathrm{proj}_K(\mathbf{y}_{t+1})$.
end for

Let $Y = \{\mathbf{y} : A\mathbf{y} \geq \mathbf{c}\}$. By the LP duality, the inner minimization $\min_{\mathbf{y}\in Y}(P^\top \mathbf{x}+\mathbf{b})^\top \mathbf{y}$ is equivalent to the following dual:

$$\max \mathbf{c}^\top \mathbf{q} \quad \text{s.t.} \quad A^\top \mathbf{q} = P^\top \mathbf{x} + \mathbf{b},\, \mathbf{q} \geq \mathbf{0}.$$

Since X is also a polytope, after adding a constraint $\mathbf{x} \in X$, we still have an LP, and therefore we can find $\mathbf{x}_H$ by solving single maximization LP.

2.3.1 Online linear optimization and approachability

The beauty of Blackwell's approachability theory is that it provides an algorithm for finding an approaching sequence, given a valid halfspace oracle. Abernethy et al. (2011) connected the approachability and OLO. In OLO, we are given a fixed compact convex set $K \subseteq \mathbb{R}^k$. In each tth round of OLO, a player selects $\mathbf{x}_t \in K$. Then an adversary reveals a vector $\mathbf{f}_t$. The goal of the player is to minimize the regret:

$$\mathrm{regret}(\mathbf{f}_1, \dots, \mathbf{f}_T) = \sum_{t\in[T]} \mathbf{f}_t^\top \mathbf{x}_t - \min_{\mathbf{x}\in K} \sum_{t\in[T]} \mathbf{f}_t^\top \mathbf{x} \qquad (3)$$

They devised an elegant algorithm for approachability, given a valid halfspace oracle $\mathcal{O}$ and an algorithm $\mathcal{A}$ for OLO, under the assumption that S is a cone.

Theorem 2.4 (Abernethy et al. (2011)). *Given a valid halfspace oracle $\mathcal{O}$, a value oracle of $\boldsymbol{\ell}$, a cone S, and an OLO algorithm $\mathcal{A}$ on the polar cone S°, there exists an algorithm $\mathcal{B}$ that given a sequence $(\mathbf{y}_t)_{t\in[T]}$, computes a sequence $(\mathbf{x}_t)_{t\in[T]}$ satisfying*

$$\mathrm{dist}\left(\frac{1}{T}\sum_{t\in[T]} \boldsymbol{\ell}(\mathbf{x}_t, \mathbf{y}_t), S\right) \leq \frac{1}{T}\,\mathrm{regret}_{\mathcal{A}}(\mathbf{f}_1, \dots, \mathbf{f}_T),$$

where $\mathbf{x}_t = \mathcal{B}(\mathbf{y}_1, \dots, \mathbf{y}_{t-1})$ and $\mathbf{f}_t = -\boldsymbol{\ell}(\mathbf{x}_t, \mathbf{y}_t)$ ($t \in [T]$).

We use *online gradient descent* (Zinkevich, 2003) as a standard OLO algorithm. See Algorihm 1 for the detail.

Theorem 2.5 (Zinkevich (2003)). *Online gradient descent with learning rate $\eta > 0$ satisfies*

$$\mathrm{regret}(\mathbf{f}_1, \dots, \mathbf{f}_T) \leq \frac{1}{\eta} D^2 + \eta \sum_{t\in[T]} \|\mathbf{f}_t\|^2,$$

where D is the diameter of K.

3 No 1/2-regret algorithm for k-submodular maximization

In this section, we present our algorithm for online k-submodular maximization. In the following, we assume that $k \geq 2$ and f_t is nonnegative and bounded, i.e., $f_t : (k+1)^V \to [0,1]$ for $t \in [T]$.

3.1 k-submodular selection game

Let us consider the following online learning problem, which we call a *k-submodular selection game*. In the tth round of the game, a player predicts a probability vector $\mathbf{p}_t \in \Delta_k$. An adversary's play is $\mathbf{y}_t = (\mathbf{a}_t, \mathbf{b}_t) \in Y$, where Y is the set of $(\mathbf{a}, \mathbf{b}) \in [-1,1]^k \times [-1,1]^k$ such that

$$\begin{aligned} a(i) + a(i') &\geq 0 && (i \neq i') \\ b(i) + b(i') &\geq 0 && (i \neq i') \\ b(i) &\geq a(i) && (i \in [k]). \end{aligned}$$

The feedback to the player is only $\mathbf{b}_t$. For a fixed $\mathbf{b}$, we denote $Y(\mathbf{b}) = \{\mathbf{a} \in [-1,1]^k : (\mathbf{a}, \mathbf{b}) \in Y\}$.

Definition 3.1. Let $\alpha > 0$. An online algorithm $\mathcal{A}$ is an α-selection algorithm for a k-submodular selection game with rate $g(k,T)$ if it satisfies

$$\begin{aligned} &\max_{i^*\in[k]} \sum_{t\in[T]} a_t(i^*) - \sum_{t\in[T]} \sum_{i\in[k]} (\alpha \cdot b_t(i) + a_t(i)) p_t(i) \\ &\quad \leq g(k,T), \end{aligned}$$

where $g(k,T)$ is sublinear in T.

Our main result is as follows.

Theorem 3.2. *There exists a 1-selection algorithm for a k-submodular selection game with rate $g(k,T) = O(k\sqrt{T})$.*

To prove this theorem, we appeal to the Blackwell approachability theorem. First, we define a biaffine vector reward function $\boldsymbol{\ell}$: For $\mathbf{p} \in \Delta_k$ and $\mathbf{y} = (\mathbf{a}, \mathbf{b}) \in Y$, let

$$\ell(\mathbf{p}, \mathbf{y})(i) = a(i) - \sum_{i'\in[k]} (b(i') + a(i')) p(i').$$

Then, $S = \mathbb{R}^k_-$ is approachable in a Blackwell instance $(\Delta_k, Y, \boldsymbol{\ell}, S)$ if and only if a 1-selection algorithm exists for a k-submodular selection game. We now show

Algorithm 2 A 1-selection algorithm for a k-submodular selection game

Input: An OLO algorithm $\mathcal{A}$ with feasible region $K := \{\boldsymbol{\theta} \in \mathbb{R}^k_+ : \|\boldsymbol{\theta}\| \leq 1\}$.
1: Set up $\mathcal{A}$.
2: **for** $t \in [T]$ **do**
3: $\boldsymbol{\theta}_t \leftarrow \mathcal{A}(\mathbf{f}_1, \ldots, \mathbf{f}_{t-1})$, where $\mathbf{f}_s := -\hat{\boldsymbol{\ell}}_s$ ($s \in [t-1]$).
4: Solve LP

$$\mathbf{p}_t \in \operatorname*{argmin}_{\mathbf{p} \in \Delta_k} \max_{\mathbf{y} \in Y} \boldsymbol{\theta}^\top \boldsymbol{\ell}(\mathbf{p}, \mathbf{y}) \tag{4}$$

to obtain $\mathbf{p}_t$.
5: Play $\mathbf{p}_t$ and observe $\mathbf{b}_t$.
6: For $i \in [k]$, let $\hat{\boldsymbol{\ell}}_t$ be a vector such that $\hat{\ell}_t(i) := \max_{\mathbf{a}_t \in Y(\mathbf{b}_t)} \ell(\mathbf{p}_t, (\mathbf{a}_t, \mathbf{b}_t))(i)$.
7: **end for**

that S is approachable. By the Blackwell approachability theorem, it suffices to show that S is response-satisfiable. Indeed, this fact is already observed in Iwata et al. (2016).

Lemma 3.3 (Iwata et al. (2016, Theorem 2.1))**.** *For a fixed adversary's play* $(\mathbf{a}, \mathbf{b})$*, there exists* $\mathbf{p} \in \Delta_k$ *that only depends on* $\mathbf{b}$ *and satisfies*

$$\max_{i^* \in [k]} a(i^*) - \sum_{i \in [k]} (b(i) + a(i)) p(i) \leq 0.$$

Therefore, the Blackwell approachability theorem implies the existence of a no-regret algorithm for a k-submodular selection game. In particular, exploiting the result of Abernethy et al. (2011), we obtain Algorithm 2 for a k-submodular selection game.

Lemma 3.4. *Algorithm 2 satisfies*

$$\max_{i^* \in [k]} \sum_{t \in [T]} a_t(i^*) - \sum_{t \in [T]} \sum_{i \in [k]} (b_t(i) + a_t(i)) p_t(i) \leq \mathrm{regret}_{\mathcal{A}}(\mathbf{f}_1, \ldots, \mathbf{f}_T),$$

for any $(\mathbf{a}_t, \mathbf{b}_t) \in Y$ $(t \in [T])$*, where* $\mathrm{regret}_{\mathcal{A}}(\mathbf{f}_1, \ldots, \mathbf{f}_T) = \sum_{t \in [T]} \mathbf{f}_t^\top \boldsymbol{\theta}_t - \min_{\boldsymbol{\theta} \in K} \sum_{t \in [T]} \mathbf{f}_t^\top \boldsymbol{\theta}$ *is the regret of the OLO algorithm* $\mathcal{A}$*.*

Proof. The proof mostly follows from Abernethy et al. (2011), but we provide the full proof for the sake of completeness. Since S is halfspace-satisfiable, LP (4) has a solution. Indeed, solving LP (4) simply computes an output of a valid halfspace oracle for a halfspace $H_t = \{\mathbf{x} \in \mathbb{R}^k : \boldsymbol{\theta}_t^\top \mathbf{x} \leq 0\}$. Let us fix arbitrary $\mathbf{y}_t = (\mathbf{a}_t, \mathbf{b}_t) \in Y$ $(t \in [T])$. Then,

$$\begin{aligned}
&\mathrm{dist}\left(\frac{1}{T} \sum_{t \in [T]} \boldsymbol{\ell}(\mathbf{p}_t, \mathbf{y}_t), S\right) \\
&= \max\left\{\max_{\boldsymbol{\theta} \in K} \frac{1}{T} \sum_{t \in [T]} \boldsymbol{\ell}(\mathbf{p}_t, \mathbf{y}_t)^\top \boldsymbol{\theta}, 0\right\} \\
&\leq \max_{\boldsymbol{\theta} \in K} \left[\frac{1}{T} \sum_{t \in [T]} \hat{\boldsymbol{\ell}}_t^\top \boldsymbol{\theta}\right] \\
&= \max_{\boldsymbol{\theta} \in K} \left[-\frac{1}{T} \sum_{t \in [T]} \mathbf{f}_t^\top \boldsymbol{\theta}\right] \\
&\leq \frac{1}{T} \max_{\boldsymbol{\theta} \in K} \left[\sum_{t \in [T]} \mathbf{f}_t^\top \boldsymbol{\theta}_t - \sum_{t \in [T]} \mathbf{f}_t^\top \boldsymbol{\theta}\right] \\
&= \frac{\mathrm{regret}_{\mathcal{A}}(\mathbf{f}_1, \ldots, \mathbf{f}_T)}{T},
\end{aligned}$$

where the second inequality follows since $\mathbf{f}_t^\top \boldsymbol{\theta}_t = -\boldsymbol{\theta}_t^\top \hat{\boldsymbol{\ell}}_t \geq 0$ by the valid halfspace oracle property. Now the claim of the lemma is immediate from the following:

$$\begin{aligned}
&\frac{1}{T}\left[\max_{i^* \in [k]} \sum_{t \in [T]} a_t(i^*) - \sum_{t \in [T]} \sum_{i \in [k]} (b_t(i) + a_t(i)) p_t(i)\right] \\
&\leq \mathrm{dist}\left(\frac{1}{T} \sum_{t \in [T]} \boldsymbol{\ell}(\mathbf{p}_t, \mathbf{y}_t), S\right)
\end{aligned}$$

□

Proof of Theorem 3.2. We can use online gradient descent as an internal OLO algorithm $\mathcal{A}$, which satisfies

$$\begin{aligned}
\mathrm{regret}_{\mathcal{A}}(\mathbf{f}_1, \ldots, \mathbf{f}_T) &\leq \frac{1}{\eta} D^2 + \eta \sum_{t \in [T]} \|\mathbf{f}_t\|^2 \\
&\leq \frac{1}{\eta} O(k) + \eta O(kT)
\end{aligned}$$

where we used that $D = O(\sqrt{k})$ is the diameter of Δ_k and $\|\mathbf{f}_t\|^2 = O(k)$ for $t \in [T]$ in the second inequality. Setting $\eta = O(1/\sqrt{T})$, we obtain the regret bound $O(k\sqrt{T})$. Combined with Lemma 3.4, we see that Algorithm 2 is a 1-selection algorithm with rate $O(k\sqrt{T})$. □

Remark 3.5. Since Algorithm 2 is deterministic if we use online gradient descent as an internal OLO algorithm, the guarantee in Theorem 3.2 holds even for an adaptive adversary.

3.2 Main algorithm

Now we present our main algorithm for online k-submodular maximization.

Theorem 3.6. *Given* α*-selection algorithms* $\mathcal{A}_j$ $(j \in [n])$ *for* k*-submodular selection games with rate*

Algorithm 3 No $1/(\alpha+1)$-regret algorithm for k-submodular maximziation

Input: α-selection algorithms $\mathcal{A}_j$ for a k-submodular selection game ($j \in [n]$).

Set up $\mathcal{A}_j$ ($j \in [n]$).
for $t = 1, \ldots, T$ **do**
  Set $\mathbf{x}_t^{(0)} := \mathbf{0}$.
  for $j \in [n]$ **do**
    Receive $\mathbf{p}_t^{(j)} \in \Delta_k$ from $\mathcal{A}_j$.
    Sample $i \in [k]$ from the probability distribution $\mathbf{p}_t^{(j)}$, and set $\mathbf{x}_t^{(j)} := \mathbf{x}_t^{(j-1)} + i\mathbf{e}_j$.
  end for
  Play $\mathbf{x}_t = \mathbf{x}_t^{(n)}$ and receive f_t.
  for $j \in [n]$ **do**
    Feedback $b_t^{(j)}(i) := \Delta_{j,i} f_t(\mathbf{x}^{(j-1)})$ ($i \in [k]$) to $\mathcal{A}_j$.
  end for
end for

$g(k,T)$, Algorithm 3 achieves

$$\mathbf{E}\left[\frac{1}{\alpha+1}\max_{\mathbf{o}\in(k+1)^V}\sum_{t\in[T]} f_t(\mathbf{o}) - \sum_{t\in[T]} f_t(\mathbf{x}_t)\right] \le \frac{ng(k,T)}{\alpha+1}, \tag{5}$$

where the expectation is taken under the randomness in Algorithm 3.

Proof. Let $\mathbf{o} \in (k+1)^V$ be an optimal solution such that $\mathrm{supp}(\mathbf{o}) = [n]$ (such an optimal solution exists by Lemma 2.2). For each $t \in [T]$ and $j = 0, 1, \ldots, n$, let $\mathbf{o}_t^{(j)} := (\mathbf{o} \sqcup \mathbf{x}_t^{(j)}) \sqcup \mathbf{x}_t^{(j)}$. Note that $\mathbf{o}_t^{(0)} = \mathbf{o}$ and $\mathbf{o}_t^{(n)} = \mathbf{x}_t^{(n)}$. Let $\mathbf{s}_t^{(j-1)}$ be a vector obtained by setting the jth element of $\mathbf{o}_t^{(j-1)}$ to 0 for $j \in [n]$. Define $a_t^{(j)}(i) := \Delta_{j,i} f_t(\mathbf{s}_t^{(j-1)})$ and $b_t^{(j)}(i) := \Delta_{j,i} f_t(\mathbf{x}_t^{(j-1)})$. By orthant submodularity and pairwise monotonicity, we have

$$\begin{aligned} a_t^{(j)}(i) + a_t^{(j)}(i') &\ge 0 && (i \ne i') \\ b_t^{(j)}(i) + b_t^{(j)}(i') &\ge 0 && (i \ne i') \\ b_t^{(j)}(i) &\ge a_t^{(j)}(i) && (i \in [k]). \end{aligned}$$

Therefore, $\mathbf{b}_t^{(j)}$ is valid feedback to $\mathcal{A}_j$ ($j \in [n]$). Let us fix $j \in [n]$ and let $i^* := o(j)$. Note that $i^* \in [k]$, since $\mathrm{supp}(\mathbf{o}) = [n]$. Since $\mathcal{A}_j$ is an α-selection algorithm, we have

$$\begin{aligned} &\sum_{t\in[T]}\sum_{i\in[k]} (a_t^{(j)}(i^*) - a_t^{(j)}(i)) p_t^{(j)}(i) \\ &\le \alpha \sum_{t\in[T]}\sum_{i\in[k]} b_t^{(j)}(i) p_t^{(j)}(i) + g(k,T), \end{aligned} \tag{6}$$

conditioned on $\mathbf{x}_t$ and f_t ($t \in [T]$). We note that (6) is valid for an adaptive adversary, since Algorithm 2 is deterministic (if we use OGD). Taking the expectation of both sides over these fixed random variables, we obtain

$$\begin{aligned} &\mathbf{E}\left[\sum_{t\in[T]} (f_t(\mathbf{o}_t^{(j-1)}) - f_t(\mathbf{o}_t^{(j)}))\right] \\ &\le \alpha \mathbf{E}\left[\sum_{t\in[T]} (f_t(\mathbf{x}_t^{(j)}) - f_t(\mathbf{x}_t^{(j-1)}))\right] + g(k,T). \end{aligned}$$

Summing these inequalities for $j \in [n]$, we arrive at

$$\begin{aligned} &\mathbf{E}\left[\sum_{t\in[T]} (f_t(\mathbf{o}) - f_t(\mathbf{x}_t))\right] \\ &\le \alpha \mathbf{E}\left[\sum_{t\in[T]} (f_t(\mathbf{x}_t) - f_t(\mathbf{0}))\right] + ng(k,T) \\ &\le \alpha \mathbf{E}\left[\sum_{t\in[T]} f_t(\mathbf{x}_t)\right] + ng(k,T), \\ &\qquad\qquad (\text{since } f_t(\mathbf{0}) \ge 0 \ (t \in [T])) \end{aligned}$$

which proves the theorem. ∎

Combining this theorem with Lemma 3.4, we obtain the main result.

Corollary 3.7. *There exists a polynomial-time algorithm for online k-submodular maximization whose $1/2$-regret is bounded by $O(kn\sqrt{T})$.*

4 Online monotone k-submodular maximization

To demonstrate the flexibility of our method with the Blackwell approachability theory, we present a no $\frac{k}{2k-1}$-regret algorithm for online monotone k-submodular maximization. To this end, we define a modified version of a k-submodular selection game, which we call a *monotone k-submodular selection game.* The only difference in the monotone case is that the set of the adversary's play is further restricted to $Y_+ := Y \cap (\mathbb{R}_+^k \times \mathbb{R}_+^k)$, which means that $\mathbf{y}_t \ge \mathbf{0}$.

Lemma 4.1. *There exists a $(1 - 1/k)$-selection algorithm for a monotone k-submodular selection game with rate $g(k,T) = O(k\sqrt{T})$.*

Proof. Again, we use the Blackwell approachability theorem. We define a slightly modified vector reward function $\boldsymbol{\ell}'$ as follows:

$$\ell'(\mathbf{p}, \mathbf{y})(i) = a(i) - \sum_{i'\in[k]} (\alpha \cdot b(i') + a(i')) p(i'),$$

where $\alpha = 1 - 1/k$. It suffices to show that $S = \mathbb{R}^k_-$ is response-satisfiable for a Blackwell instance $(X, Y_+, \boldsymbol{\ell}', S)$. In Iwata et al. (2016, Theorem 2.2), it is shown that for fixed $\mathbf{y} = (\mathbf{a}, \mathbf{b}) \in Y_+$, there exists $\mathbf{p} \in \Delta_k$ such that $\boldsymbol{\ell}'(\mathbf{p}, \mathbf{y}) \leq \mathbf{0}$. Therefore, by the Blackwell approchability theorem, there exists an online algorithm for producing an approaching sequence. Indeed, such an algorithm can be constructed by a slight modification of Algorithm 2: instead of $\boldsymbol{\ell}$ and Y, we use $\boldsymbol{\ell}'$ and Y_+, respectively. It is easy to see that the modified algorithm produces a sequence $\mathbf{p}_t$ ($t \in [T]$) with the same guarantee as in Lemma 3.4:

$$\max_{i^* \in [k]} \sum_{t \in [T]} a_t(i^*) - \sum_{t \in [T]} \sum_{i \in [k]} (\alpha \cdot b_t(i) + a_t(i)) p_t(i) \leq \mathrm{regret}_{\mathcal{A}}(\mathbf{f}_1, \ldots, \mathbf{f}_T),$$

for any $(\mathbf{a}_t, \mathbf{b}_t) \in Y$ ($t \in [T]$), where $\mathcal{A}$ is an internal OLO algorithm. Again, using online gradient descent as $\mathcal{A}$, we obtain the same bound as before, which completes the proof. □

Combining this result with Theorem 3.6, we obtain the following.

Theorem 4.2. *There exists a polynomial-time algorithm for online monotone k-submodular maximization whose $\frac{k}{2k-1}$-regret is bounded by $O(kn\sqrt{T})$.*

Proof. We use the same notation as in the proof of Theorem 3.6. Since f_t is monotone ($t \in [T]$), we have $\mathbf{a}_t^{(j)}, \mathbf{b}_t^{(j)} \geq \mathbf{0}$ ($t \in [T]$, $j \in [n]$). Therefore, $\mathbf{b}_t^{(j)}$ is valid feedback to an algorithm for a monotone k-submodular selection game. Since $\alpha = 1 - 1/k$, we have the same bound for the $\frac{k}{2k-1}$-regret. □

5 Applications

In this section, we briefly describe applications of online k-submodular maximization.

5.1 Coupled feature selection

This application is inspired by Singh et al. (2012). Suppose that we are to predict k variables $Z_1, \ldots, Z_k$ (say, weather of k different areas) using n features $Y_1, \ldots, Y_n$. We have a communication constraints such that each feature can be used for prediction of only one of Z_i ($i \in [k]$). Then, the problem is to find a subpartition $(X_1, \ldots, X_k)$ of features into k groups that maximizes the mutual information:

$$\begin{aligned} f(X_1, \ldots, X_k) &= I(X_1, \ldots, X_k; Z) \\ &= H(X_1, \ldots, X_k) - H(X_1, \ldots, X_k \mid Z), \end{aligned}$$

where $H(\cdot)$ is the entropy function. We make the following assumption: the features $Y_1, \ldots, Y_n$ are paiwise independent given Z. Then we obtain $f(X_1, \ldots, X_k) = H(X_1, \ldots, X_k) - \sum_{i \in [k]} \sum_{j \in X_i} H(Y_j | Z_i)$. One can show that this function is indeed k-submodular (see Appendix A).

Online setting naturally corresponds the following scenario: we repeatedly predict weather over T days and we want to minimize the approximate regret.

5.2 Sensor placement with k different sensors

This applications is inspired by Singh et al. (2012); Ohsaka and Yoshida (2015); Iwata et al. (2016). Suppose that we have k types of sensors each of which corresponds to different physical observation. We are to place sensors to n spots but we can only place one type of sensors in each spot. Let us assume that a global information gain function g exists and g is symmetric (i.e., $g(X) = g(V \setminus X)$ for all X) and submodular (see Golovin et al. (2014) for examples of such functions). Our objective is to maximize the total information gain of k different observations: $f(X_1, \ldots, X_k) = \sum_{i \in [k]} g(X_i)$. One can check that f is a k-submodular function (see Appendix A). We note that this example does not reduce to monotone submodular maximization under partition matroid constraint (Calinescu et al., 2011), since f is not necessarily monotone. Again, online scenario naturally corresponds to repeated observations over T periods.

6 Conclusion

In this paper, we developed polynomial time algorithms for online k-submodular maximization and online monotone k-submodular maximization with the sublinear approximate regret.

For open problems, we may consider constrained online k-submodular maximization such as a size constraint (Ohsaka and Yoshida, 2015) and a matroid constraint (Sakaue, 2017) . We hope that our framework with the Blackwell approachability theorem is useful for these settings. Another open problem is improving the dependence on n and k in the approximate regret bound. Our algorithms provide the approximate regret bound $O(nk\sqrt{T})$, but we do not know any lower bound on n and k.

Acknowledgement

The author thanks Takanori Maehara, Shinsaku Sakaue, Yuichi Yoshida, and Kaito Fujii for valuable discussions. The author also thanks Tim Roughgarden and Joshua R. Wang for sharing a draft of their paper. This work was supported by ACT-I, JST, and CREST, JST.

References

Abernethy, J., Bartlett, P. L., and Hazan, E. (2011). Blackwell approachability and no-regret learning are equivalent. In *Proceedings of the 24th Annual Conference on Learning Theory (COLT)*, volume 19, pages 27–46.

Arora, S., Hazan, E., and Kale, S. (2012). The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8:121–164.

Bian, A., Levy, K., Krause, A., and Buhmann, J. M. (2017a). Continuous DR-submodular maximization: Structure and algorithms. In *Advances in Neural Information Processing Systems (NIPS)*, pages 486–496.

Bian, A. A., Mirzasoleiman, B., Buhmann, J. M., and Krause, A. (2017b). Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 111–120.

Blackwell, D. (1956). An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8.

Buchbinder, N., Feldman, M., Seffi, J., and Schwartz, R. (2015). A tight linear time (1/2)-approximation for unconstrained submodular maximization. *SIAM Journal on Computing*, 44(5):1384–1402.

Calinescu, G., Chekuri, C., Pál, M., and Vondrák, J. (2011). Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games.* Cambridge university press.

Chen, L., Harshaw, C., Hassani, H., and Karbasi, A. (2018a). Projection-free online optimization with stochastic gradient: From convexity to submodularity. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 814–823.

Chen, L., Hassani, H., and Karbasi, A. (2018b). Online continuous submodular maximization. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 84, pages 1896–1905.

Feige, U., Mirrokni, V., and Vondrák, J. (2011). Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153.

Feldman, M., Karbasi, A., and Kazemi, E. (2018). Do less, get more: Streaming submodular maximization with subsampling. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 730–740.

Fujishige, S. (2005). *Submodular Functions and Optimization.* Elsevier, 2nd edition.

Fujishige, S. and Iwata, S. (2005). Bisubmodular function minimization. *SIAM Journal on Discrete Mathematics*, 19(4):1065–1073.

Golovin, D., Krause, A., and Streeter, M. (2014). Online submodular maximization under a matroid constraint with application to learning assignments. *arxiv*.

Hazan, E. (2016). *Introduction to Online Convex Optimization*, volume 2. Foundations and Trends in Optimization.

Hirai, H. and Oki, T. (2017). A compact representation for minimizers of k-submodular functions. *Journal of Combinatorial Optimization.*

Huber, A. and Kolmogorov, V. (2012). Towards minimizing k-submodular functions. In *Proceedings of the International Symposium on Combinatorial Optimization*, pages 451–462.

Iwata, S., Tanigawa, S., and Yoshida, Y. (2016). Improved approximation algorithms for k-submodular function maximization. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 404–413.

Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 137–146.

Krause, A. and Golovin, D. (2014). Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems*, pages 71–104. Cambridge University Press.

Mirzasoleiman, B., Jegelka, S., and Krause, A. (2018). Streaming non-monotone submodular maximization: Personalized video summarization on the fly. In *The proceedings of the 32nd AAAI Conference on Artificial Intelligence.*

Niazadeh, R., Roughgarden, T., and Wang, J. (2018). Optimal algorithms for continuous non-monotone submodular and DR-submodular maximization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9617–9627.

Ohsaka, N. and Yoshida, Y. (2015). Monotone k-submodular function maximization with size constraints. In *Advances in Neural Information Processing Systems (NIPS)*, pages 694–702.

Oshima, H. (2017). Derandomization for k-submodular maximization. In *Proceedings of the International Workshop on Combinatorial Algorithms (IWOCA)*, pages 88–99.

Roughgarden, T. and Wang, J. R. (2018). An optimal algorithm for online unconstrained submodular maximization. In *Proceedings of the 31st Annual Conference on Learning Theory (COLT)*, pages 1307–1325.

Sakaue, S. (2017). On maximizing a monotone k-submodular function subject to a matroid constraint. *Discrete Optimization*, 23:105–113.

Singh, A., Guillory, A., and Bilmes, J. (2012). On bisubmodular maximization. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 22, pages 1055–1063.

Streeter, M. and Golovin, D. (2009). An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1577–1584.

Suehiro, D., Hatano, K., Kijima, S., Takimoto, E., and Nagano, K. (2012). Online prediction under submodular constraints. In *Proceedings of the International Conference on Algorithmic Learning Theory (ICML)*, pages 260–274.

Takimoto, E. and Warmuth, M. K. (2003). Path kernels and multiplicative updates. *Journal of Machine Learning Research*, 4:773–818.

Ward, J. and Živný, S. (2016). Maximizing k-submodular functions and beyond. *ACM Transactions of Algorithms*, 12(4):47:1–47:26.

Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936.