No-regret algorithms for online k-submodular maximization

相馬 輔

東京大学大学院 情報理工学系研究科

This talk examines online maximization of *k-submodular* functions. k-submodular functions are generalizations of submodularity and bisubmodularity, introduced by Huber and Kolmogolov (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. If $k = 1, 2$, k-submodularity is equivalent to submodularity and bisubmodularity, respectively. Iwata, Tanigawa, and Yoshida (2016) devised a $1/2$-approximation algorithm for maximizing k-submodular functions.

Online k-submodular maximization is a two-player game between a player and an adversary (see Figure 1). The performance measure of the player is the *α-regret*:

$$\mathrm{regret}_\alpha(f_1, \ldots, f_T) = \alpha \max_{\mathbf{x} \in \mathcal{C}} \sum_{t \in [T]} f_t(\mathbf{x}) - \sum_{t \in [T]} f_t(\mathbf{x}_t).$$

For $t = 1, \ldots, T$

- 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 $f_t(\mathbf{x}_t)$.

Figure 1 The online k-submodular maximization protocol

We show that:

- 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 $O(nk\sqrt{T})$.