Online Discrete Convex Optimization

Jefferson Huang

Assistant Professor
Operations Research Department
Naval Postgraduate School

Joint Work With:

Louis Chen (NPS)

Yasushi Kawase (Univ. Tokyo)

Tasuku Soma (Institute of Statistical Mathematics),

INFORMS Annual Meeting
Phoenix, AZ
16 October, 2023

Online Combinatorial Optimization

Consider a matroid $(N, \mathcal{I})$, where the **ground set** $N = [n] := \{1, \ldots, n\}$ is finite and $\mathcal{I} \subseteq 2^N$ is the associated family of **independent sets**.

Online Setting

At each time $t \in [T] := \{1, \ldots, T\}$:

- Select an independent set $S_t \in \mathcal{I}$ to play.
- The cost incurred by S_t is evaluated according to an *adversarially* chosen function $f_t : \mathcal{I} \to [0, 1]$.

We will focus on the **full-information** setting.

- i.e., f_t's are revealed via a value oracle.

Objective: Approximate Regret Minimization

Do as well, over the time horizon T, as some constant $\alpha \in [0, 1]$ times the best fixed $S \in \mathcal{I}$ in hindsight, i.e. minimize the **α-regret**

$$\text{Regret}_\alpha(T) := \sum_{t=1}^{T} f_t(S_t) - \alpha \min_{S \in \mathcal{F}} \sum_{t=1}^{T} f_t(S)$$

Examples of applications include online versions of:

- Routing on Networks
- Selecting Products to Produce
- Selecting Items to Bid On
- Selecting Locations of Facilities, Sensors, . . .

Online Submodular Optimization

Assumption

The f_t's are **submodular**, i.e.,

$$f_t(S) + f_t(S') \geqslant f_t(S \cup S') + f_t(S \cap S')$$

for all $S, S' \in 2^N$ and $t \in [T]$

- **Special Case:** The f_t's are **modular** (aka linear) if $f_t(S) = \sum_{a \in S} c_t(a)$ for some $c_t : N \to [0, 1]$
- Submodularity models diminishing marginal returns (e.g., profits, coverage).

There are algorithms whose incurred α-**regret grows at most sub-linearly** in T for, e.g.,

- cardinality-constrained modular optimization for $\alpha = 1$ [Audibert Bubeck & Lugosi, 2014]
- combinatorially-constrained monotone submodular minimization [Jegelka & Bilmes, 2010]
- cardinality-constrained monotone submodular maximization for $\alpha = 1 - \frac{1}{e}$ [Streeter & Golovin, 2008], [Harvey Liaw & Soma, 2020]
- matroid-constrained monotone submodular maximization for $\alpha = 1 - \frac{1}{e} - \epsilon$ [Golovin Krause & Streeter 2014], [Harvey Liaw & Soma, 2020]

Motivation: Online Convex Optimization

For optimization problems of the form

$$\min_{x \in \mathbb{R}^n} f(x), \quad f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\},$$

convexity is a useful way to distinguish between "easy" and "hard" problems [Rockafellar, 1970]

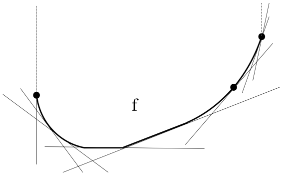

From p. 309 in [Rockafellar & Wets, 2009]

There is a well-developed theory for online minimization when the decisions are vectors in $\mathbb{R}^n$ and the functions f_t are convex [Hazan, 2022].

Examples:

- 1-regret bounds that grow with $\sqrt{T}$
- 1-regret lower bounds that scale with $\sqrt{T}$

Question

Is there a useful notion of convexity in combinatorial optimization?

Discrete Convexity

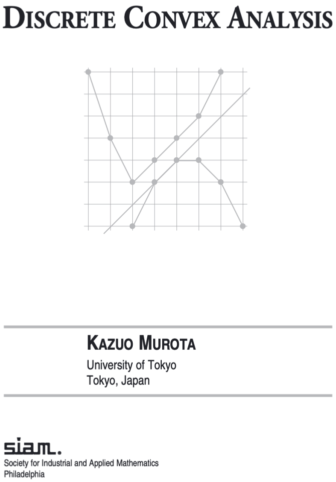

There is more than one way to usefully define convexity for functions on $\mathbb{Z}^n$ [Murota, 2003].

- **Desiderata:** duality between convexity and concavity, separation, . . .
- Not enough to assume that the discrete function (on $\mathbb{Z}^n$) can be extended to a convex function on $\mathbb{R}^n$ [Murota, 2003].

There are two notions of discrete convexity that are useful for modeling certain discrete optimization problems [Murota, 2003]:

- **$L^\natural$-Convexity** ("Lattice Convexity")
- **$M^\natural$-Convexity** ("Matroid Convexity")

Online L$^\natural$-Convex Minimization

Proposition (p. 145 in [Murota, 2003])

A function

$$f : 2^N \to \mathbb{R} \cup \{+\infty\}$$

is **submodular** if and only if its indicator function on $\{0,1\}^n$ is L$^\natural$-convex.

- L$^\natural$-convex minimization is equivalent to **submodular minimization**

Projected Subgradient Descent (PSD)

Given a step size η and an initial $x_1 \in [0,1]^n$, for $t = 1, \ldots, T$ do the following:

1. Round x_t to an indicator vector for a subset S_t.
2. Play S_t.
3. Compute a subgradient p_t of the Lovász extension of f_t at x_t.
4. Take a projected gradient step in the direction of $-p_t$ from x_t to x_{t+1}

Theorem [Jegelka & Bilmes 2010]

If the rounding procedure is α-approximate, then the expected α-regret of PSD is bounded by $\sqrt{nT}$.

Online M$^\natural$-Concave Maximization

Proposition (p. 179 in [Murota, 2003])

If $g : \{0,1\}^n \to \mathbb{R} \cup \{+\infty\}$ is M$^\natural$-concave, then its associated set function

$$f : 2^N \to \mathbb{R} \cup \{+\infty\}$$

is **submodular**.

- M$^\natural$-concave maximization is a special case of **submodular maximization**.

Projected Supergradient Ascent (PSA)

Given a step size η and initial $x_1 \in [0,1]^n$, for $t = 1, \ldots, T$ do the following:

1. For each $i \in N$, draw a threshold τ_i uniformly at random.
2. Play the subset $S_t = \{i : x_t(i) > \tau_i\}$
3. Compute a "supergradient" p_t of f_t at S_t, i.e., a p that minimizes

$$p \cdot x_t - f^\circ(p), \quad p \in \mathbb{Z}^n, \|p\|_\infty \leqslant 2n$$

4. Take a projected gradient step in the direction of p_t from x_t to get x_{t+1}

Theorem [Chen, H. Kawase & Soma]

For $\eta = 1/(2n\sqrt{T})$, the expected $(1 - \frac{1}{e})$-regret of PSA is bounded by $2n^2\sqrt{T}$.

Ongoing Work

Current Work:

- Derive a parameter-free version of the online M$^\natural$-concave maximization algorithm
- Study FTPL-type algorithms [Jegelka & Bilmes, 2010]
- Hardness conjecture for online M$^\natural$-concave maximization.

Future Work: Extension to online optimization of functions on $\mathbb{Z}^n$

- Inventory Control [Chen & Li, 2021]
- Bike Sharing [Freund Henderson & Shmoys 2022], [Shioura 2022]
- Games [Fujishige Goemans Harks Peis & Zenklusen, 2015]