Bandit Convex Optimization: $\sqrt{T}$ Regret in One Dimension

Sébastien Bubeck **Ofer Dekel**

Tomer Koren **Yuval Peres**

Technion Israel Institute of Technology

Microsoft Research

Bandit Convex Optimization

- **Adversary** picks convex $f_1, f_2, \ldots, f_T : K \to [0,1]$ over convex $K \subseteq \mathbf{R}^n$
- On rounds $t=1,\ldots,T$, **player**:
  1. chooses $x_t \in K$ at random
  2. incurs loss $f_t(x_t)$
  3. observes $f_t(x_t)$ and nothing else

- Player's goal: minimize regret

$$R_T = \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t(x)\right]$$

Bandit Convex Optimization

- **Adversary** picks convex $f_1, f_2, \ldots, f_T : K \to [0,1]$ over convex $K \subseteq \mathbf{R}^n$
- On rounds $t=1,\ldots,T$, **player**:
  1. chooses $x_t \in K$ at random
  2. incurs loss $f_t(x_t)$
  3. observes $f_t(x_t)$ and nothing else

- Player's goal: minimize regret

$$R_T = \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t(x)\right]$$

Thm: for n=1, there exists algorithm for which $R_T = \tilde{O}(\sqrt{T})$

Note: no additional assumptions (not even Lipschitz)

Previous results

- $T^{5/6}$ for general convex [Flaxman, Kalai, McMahan '05]

- $T^{3/4}$ for Lipschitz [Flaxman, Kalai, McMahan '05]
- $T^{1/2}$ for linear + Lipschitz [Dani et al '08, Abernethy et al '08]
- $T^{2/3}$ for strongly convex + Lipschitz [Agarwal et al '10]
- $T^{2/3}$ for smooth [Saha & Tewari' 11]
- $T^{1/2}$ for stochastic i.i.d. + Lipschitz [Agarwal et al '11]
- $T^{1/2}$ for smooth + strongly convex [Hazan & Levy '14]

☛ **Our result: first tight bound** without further assumptions

Main ideas

1. Non-constructive analysis: don’t give an algorithm; instead, use **minimax duality** and attack the dual

Main ideas

1. Non-constructive analysis: don't give an algorithm; instead, use **minimax duality** and attack the dual
2. Dual is **Bayesian BCO:** $f_1,\ldots,f_T$ are RVs drawn from adversarial, yet **known joint** (non-i.i.d.) distribution

Main ideas

1. Non-constructive analysis: don't give an algorithm; instead, use **minimax duality** and attack the dual

2. Dual is **Bayesian BCO:** $f_1,\ldots,f_T$ are RVs drawn from adversarial, yet **known joint** (non-i.i.d.) distribution

3. Solve Bayesian problem with a variant of **Thompson Sampling**

Main ideas

1. Non-constructive analysis: don't give an algorithm; instead, use **minimax duality** and attack the dual

2. Dual is **Bayesian BCO:** $f_1, \ldots, f_T$ are RVs drawn from adversarial, yet **known joint** (non-i.i.d.) distribution

3. Solve Bayesian problem with a variant of **Thompson Sampling**

4. Bayesian regret analysis via information-theoretic arguments [**Russo & van Roy '14**]

Main ideas

1. Non-constructive analysis: don't give an algorithm; instead, use **minimax duality** and attack the dual
2. Dual is **Bayesian BCO:** $f_1,\ldots,f_T$ are RVs drawn from adversarial, yet **known joint** (non-i.i.d.) distribution
3. Solve Bayesian problem with a variant of **Thompson Sampling**
4. Bayesian regret analysis via information-theoretic arguments [Russo & van Roy '14]
5. Novel use of convexity: new **"local-to-global" property** of convex functions