Threshold Policies with Tight Guarantees for Online Selection with Convex Costs

xiaoqi.tan@ualberta.ca

Joint work with Siyuan Yu, Raouf Boutaba, and Alberto Leon-Garcia









WINE 2023 - The 19th Conference on Web and Internet Economics, Shanghai, China

Xiaoqi Tan



UNIVERSITY OF WATERLOO





- Selling k items for a sequence of buyers
- Each buyer $t = 1, 2, \dots, T$ makes an offer v_t
- Accept it and sell 1 item, or wait for next one
- Once sold, cannot be reclaimed

$$\alpha = \max_{\sigma} \frac{\mathsf{Opt}(\sigma)}{\mathbb{E}[\mathsf{Alg}(\sigma)]}$$



e.g., k = 2, T = 14

Online Selection Problems

- A set of T items (one at a time)
- Select a subset $S \subset T$ of items (with possible constraints)
- Maximize objective v(S)

Optimal Search and One-Way Trading Online Algorithms

R. El-Yaniv,¹ A. Fiat,² R. M. Karp,³ and G. Turpin⁴

Optimal Algorithms for k-Search with **Application in Option Pricing**

Julian Lorenz¹, Konstantinos Panagiotou², and Angelika Steger

Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland {jlorenz,panagick,steger}@inf.ethz.ch

Algorithmica 2001

WINE 2008

Budget Constrained Bidding in Keyword Auctions and Online Knapsack Problems

Deeparnab Chakrabarty College of Computing Georgia Tech Atlanta, GA 30332 deepc@cc.gatech.edu

Yunhong Zhou HP Labs 1501 Page Mill Rd Palo Alto CA 95304 vunhong.zhou@hp.com

Rajan Lukose HP Labs 1501 Page Mill Rd Palo Alto CA 95304 rajan.lukose@hp.com

Online Selection Problems against Constrained Adversary

Zhihao Jiang¹ Pinyan Lu² Zhihao Gavin Tang² Yuhao Zhang³

ICML 2021

Pareto-Optimal Learning-Augmented Algorithms for **Online Conversion Problems**

Bo Sun ECE, HKUST bsunaa@connect.ust.hk

Russell Lee CICS, UMass Amherst rclee@cs.umass.edu

Mohammad Hajiesmaili CICS, UMass Amherst hajiesmaili@cs.umass.edu

Adam Wierman CMS, Caltech adamw@caltech.edu Danny H.K. Tsang ECE, HKUST eetsang@ust.hk

NeurIPS 2021

ESA 2007



OSCC: Online Selection w/ Convex Costs

- A set of T items (one at a time)
- Select a subset $S \subset T$ of items (with possible constraints)
- Maximize objective v(S)

 $f(y) = \begin{cases} \text{convex \& monotone} & \text{if } y \in [0, 15], \\ +\infty & \text{otherwise.} \end{cases}$





So...how long have you been circling for a spot?

3.0000





- A set of T items (one at a time)
- Select a subset $S \subset T$ of items (with possible constraints)
- Maximize objective v(S)

convex & monotone production costs

At most one bundle is selected

 $\sum_{b} x_t^b \le 1, \forall t$ $x_t^b = \{0,1\}, \forall t, b$

 $\sum_{t} \sum_{h} v_t^b x_t^b$

 $\max_{x_t \in \mathbb{X}_t}$



OSCC: Basic Setting







OSCC: Basic Setting





$Opt(\sigma)$ $\max_{\sigma \in \Omega_{\mathcal{S}}} \mathbb{E}[Alg(\sigma)]$ $\alpha = \max$

$\mathcal{S} = (k, f, v_{\min}, v_{\max})$

Time

 $\sum_{t=1}^{T} v_t x_t - f\left(\sum_{t=1}^{T} x_t\right)$ maximize $\{x_t\}_{\forall t}$ T $\sum x_t \le k,$ subject to t=1 $x_t = \{0, 1\}, \forall t$.







 $\rho = v_{\rm max}/v_{\rm min}$

 $\lim \operatorname{Gap}(k) = 0$







Theorem 1: Optimal Threshold (Deterministic)

Theorem 1 (Optimal Threshold). Given a setup $S = \{f, p_{\min}, p_{\max}, k\}$, TOS_{λ^*} achieves the optimal competitive ratio of all deterministic online algorithms, denoted by $CR_f^*(\rho, k)$, if and only if $\boldsymbol{\lambda}^* = \left\{\lambda_0^*, \lambda_1^*, \cdots, \lambda_{\tau}^*, \cdots, \lambda_{\overline{k}}^*\right\}$ is an admission threshold such that

- The lower and upper limits: $\lambda_0^* = \lambda_1^* = \cdots = \lambda_{\tau}^* = p_{\min}$ and $\lambda_{\bar{k}}^* = p_{\max}$.
- The turning point τ of the admission threshold is given by

$$\tau = g^{inv} \left(\frac{f^*(p_{\min})}{\mathsf{CR}_f^*(\rho, k)} \right) - 1.$$

• The optimal competitive ratio $CR_f^*(\rho, k)$ and $\{\lambda_{\tau+1}^*, \lambda_{\tau+2}^*, \cdots, \lambda_{\bar{k}-1}^*, \lambda_{\bar{k}}^*\}$ satisfy:

(SoSE):
$$\operatorname{CR}_{f}^{*}(\rho,k) = \frac{f^{*}(\lambda_{\tau+1}^{*})}{g(\tau+1)} = \frac{f^{*}(\lambda_{\tau+2}^{*}) - f^{*}(\lambda_{\tau+1}^{*})}{\lambda_{\tau+1}^{*} - c_{\tau+2}} = \dots = \frac{f^{*}(\lambda_{\bar{k}}^{*}) - f^{*}(\lambda_{\bar{k}-1}^{*})}{\lambda_{\bar{k}-1}^{*} - c_{\bar{k}}}$$

- **Existence** of many competitive TPs

- Uniqueness of optimal (deterministic) TP
- Optimality among all deterministic algs









$$f^{*}(p) \triangleq \max_{i \in \{0,1,\cdots,k\}} pi - j$$
- Initial flat phase
$$v_{\min}(\tau+1) - f(\tau+1) \ge \frac{1}{\alpha} f^{*}(v_{\max})$$
- Strictly increasing phase
$$\Delta_{\text{primal}} \ge \frac{1}{\alpha} \cdot \Delta_{\text{dual}}$$







Theorem 2 (Lower Bound). For any given setup $S = \{f, p_{\min}, p_{\max}, k\}$, no online algorithms (possibly randomized) is $(CR_f^{lb}(\rho, k) - \epsilon)$ -competitive for any $\epsilon > 0$, where $CR_f^{lb}(\rho, k)$ is given by

$$\mathsf{CR}^{\mathsf{lb}}_f(
ho,k) = rac{p_{\min} \underline{k} - f(\underline{k})}{p_{\min} \gamma^{(1)} - f(\gamma^{(1)})},$$

where $\gamma^{(1)}$ is a real value within $(0, \underline{k}]$. Specifically, let us define the right-hand-side of Eq. (19) as $F(\gamma^{(1)})$ to indicate that it is an explicit function of $\gamma^{(1)} \in (0, \underline{k}]$. Together with $\{\gamma^{(\ell)}\}_{\ell=\{2, \cdots, \overline{k}-\underline{k}+2\}}$, they form a unique set of increasing positive real numbers (i.e., $0 < \gamma^{(1)} < \gamma^{(2)} < \cdots < \gamma^{(\overline{k}-\underline{k}+1)} < \overline{(1-\alpha)}$. $\gamma^{(\bar{k}-\bar{k}+2)}=\bar{k}$) that satisfy

$$\frac{q^{(\ell+1)}(\underline{k}+\ell-1)}{\exp\left(\frac{F(\gamma^{(1)})}{\underline{k}+\ell-1}\gamma^{(\ell+1)}\right)} - \frac{q^{(\ell)}(\underline{k}+\ell-1)}{\exp\left(\frac{F(\gamma^{(1)})}{\underline{k}+\ell-1}\gamma^{(\ell)}\right)} = \int_{\gamma^{(\ell)}}^{\gamma^{(\ell+1)}} \frac{F(\gamma^{(1)})f'(y)}{\exp\left(\frac{F(\gamma^{(1)})}{\underline{k}+\ell-1}y\right)} dy, \quad \forall \ell = [\bar{k}-\underline{k}+1],$$

where $q^{(\ell)} = c_{k+\ell-1}$ for $\ell = \{2, 3, \dots, \bar{k} - \underline{k} + 1\}$, $q^{(1)} = p_{\min}$, and $q^{(\bar{k} - \underline{k} + 2)} = p_{\max}$.

Piece-wise continuous threshold w/ fractional end points

$$0 < \gamma^{(1)} < \gamma^{(2)} < \dots < \gamma^{(\bar{k}-\underline{k}+1)} < \gamma^{(\bar{k}-\underline{k}+2)} = q^{(1)} < q^{(2)} < \dots < q^{(\bar{k}-\underline{k}+1)} < q^{(\bar{k}-\underline{k}+2)} = q^{(1)} < q^{(2)} < \dots < q^{(\bar{k}-\underline{k}+1)} < q^{(\bar{k}-\underline{k}+2)} = q^{(\bar{k}-\underline{k}+1)} < q^{(\bar{k}-\underline{k}+2)} = q^{(\bar{k}$$











Theorem 3 (Asymptotic Lower Bound). For any given setup $S = \{f, p_{\min}, p_{\max}, k\}$, TOS_{λ^*} is asymptotically optimal among all online algorithms (including those with randomization) when \bar{k} is sufficiently large, namely,

$$\lim_{\bar{k} \to \infty} \mathsf{CR}_f^*(\rho, k) = \lim_{\bar{k} \to \infty} \mathsf{CR}_f^{\mathsf{lb}}(\rho, k) = \underline{\mathsf{CR}}_f(\rho), \tag{2}$$

where $\underline{CR}_f(\rho)$ is the asymptotic lower bound that depends on f and ρ only.

A two-segment threshold is enough if k is sufficiently large!!







- $1 + \ln \rho$ **Strongly convex**

Higher value fluctuations \Longrightarrow worse guarantees (larger α^*)

Faster cost growth \Longrightarrow better guarantees (smaller α^*)





Generalization and Variations

Action

Convex
Box
Ramping
Inventory
Deadlines
Binary

Value

 $\max_{x_t \in \mathbb{X}_t} \sum_{t \neq t} v_t$

Linear Concave Strongly concave Separable Combinatorial

Cost

Linear Quadratic Polynomial Convex Strongly convex Separable Non-convex

$s = (s_0, s_1, s_2, \dots, s_t, \dots, s_T)$ $s_t \sim s_{t-1}, x_t$

Dynamics

Linear / Nonlinear Stationary / Non-stationary Too complex





S+

*S*₂

Proceedings of Machine Learning Research vol 75:1-21, 2018

Smoothed Online Convex Optimization in High Dimensions via Online Balanced Descent

Niangjun Chen ^a* Gautam Goel ^b* Adam Wierman ^b ^a Institute of High Performance Computing

CHENNJ@IHPC.A-STAR.EDU.SG GGOEL@CALTECH.EDU ADAMW@CALTECH.EDU ^b California Institute of Technology

COLT 2018

Chasing Convex Bodies Optimally

Mark Sellke Stanford University

Abstract

In the chasing convex bodies problem, an online player receives a request sequence of N convex sets K_1, \ldots, K_N contained in a normed space X of dimension d. The player starts at $x_0 = 0 \in X$, and at time n observes the set K_n and then moves to a new point $x_n \in K_n$, paying a cost $||x_n - x_{n-1}||$. The player aims to ensure the total cost exceeds the minimum possible total cost by at most a bounded factor α_d independent of N, despite x_n being chosen without knowledge of the future sets K_{n+1}, \ldots, K_N . The best possible α_d is called the competitive ratio. Finiteness of the competitive ratio for convex body chasing was proved for d = 2 in [FL93] and conjectured for all d. [BLLS19] recently resolved this conjecture, proving an exponential $2^{O(d)}$ upper bound on the competitive ratio.

We give an improved algorithm achieving competitive ratio d in any normed space, which is *exactly* tight for ℓ^{∞} . In Euclidean space, our algorithm also achieves competitive ratio $O(\sqrt{d \log N})$, nearly matching a \sqrt{d} lower bound when N is subexponential in d. Our approach extends that of [BKL⁺20] for *nested* convex bodies, which is based on the classical Steiner point of a convex body. We define the *functional* Steiner point of a convex function and apply it to the associated work function.

SODA 2020

Chasing Convex Bodies with Linear Competitive Ratio

C.J. Argue¹, Anupam Gupta¹, Guru Guruganesh², and Ziye Tang¹

We study the problem of chasing convex bodies online: given a sequence of convex bodies $K_t \subseteq \mathbb{R}^d$ the algorithm must respond with points $x_t \in K_t$ in an online fashion (i.e., x_t is chosen before K_{t+1} is revealed). The objective is to minimize the sum of distances between successive points in this sequence. Bubeck et al. (STOC 2019) gave a $2^{O(d)}$ -competitive algorithm for this problem. We give an algorithm that is $O(\min(d, \sqrt{d \log T}))$ -competitive for any sequence of length T.



Regularized Online Allocation Problems: Fairness and Beyond

Santiago Balseiro^{*12} Haihao Lu^{*32} Vahab Mirrokni²

ICML 2021

Many more

Optimal Regularized Online Allocation by Adaptive Re-Solving

Wanteng Ma¹, Ying Cao², Danny H.K. Tsang², Dong Xia^{1*}

¹Department of Mathematics, HKUST ²Department of Electronic and Computer Engineering, HKUST

(July 18, 2023)

 $\bullet \quad \bullet \quad \bullet$

¹Carnegie Mellon University ²Google Research

January 8, 2020

Abstract

SODA 2020



17



Concluding Remarks

Online selection with convex costs

- Simple algorithms; Strong guarantees
- Variations; Open questions



- A set of T items (one at a time)
- Select a subset $S \subset T$ of items (with possible constraints)
- Maximize objective v(S) f(S)







References

- of Operations Research, 34(2):270–286, May 2009.
- * Y. Azar et al. Online algorithms for covering and packing problems with convex objectives. FOCS 2016.
- R. El-Yaniv, A. Fiat, R. M. Karp, and G. Turpin. Optimal search and one-way trading online algorithms. Algorithmica, 30(1):101–139, 2001.
- in option pricing. **ESA 2007**.
- adversary, ICML 2021.
- dependent competitive ratios. **Operations Research**, 68(6):1787–1803, 2020.
- allocation: A unified approach. SIGMETRICS 2020.
- and online knapsack problems. WINE 2008.

A Stripping and Anthony Kim. Welfare maximization with production costs: A primal dual approach. SODA 2015. * Niv Buchbinder and Joseph (Seffi) Naor. Online Primal-Dual Algorithms for Covering and Packing. Mathematics

* Julian Lorenz, Konstantinos Panagiotou, and Angelika Steger. Optimal algorithms for k-search with application

Zhihao Jiang, Pinyan Lu, Zhihao Gavin Tang, and Yuhao Zhang. Online selection problems against constrained
 A

* Will Ma and David Simchi-Levi. Algorithms for online matching, assortment, and pricing with tight weight-

* Xiaoqi Tan, Bo Sun, Alberto Leon-Garcia, Yuan Wu, and D.H.K Tsang. Mechanism design for online resource

Yunhong Zhou, Deeparnab Chakrabarty, and Rajan Lukose. Budget constrained bidding in keyword auctions







