Contextual Search in the Presence of Irrational Agents

Akshay Krishnamurthy (MSR NYC), Thodoris Lykouris (MSR NYC → MIT),
Chara Podimata (Harvard), and Robert Schapire (MSR NYC)

STOC 2021

Full version: SCAN ME

Main Question

Contextual search ≐ generalization of single dimensional binary search where a context x_t is given per round to guide the search

Can model many real-life settings, e.g., AirBnB's smart pricing algorithm

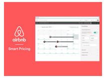

How do we construct contextual search algorithms that are robust to the presence of some **model misspecifications**?

TCS: multidimensional binary search when nature is allowed to lie a fixed number of times
ML: learning with 1-sided feedback in the presence of corruptions
Econ: feature-based pricing against adversarially irrational agents (i.e., deviating from prescribed behavioral model)

Model [Cohen, Lobel, Paes Leme, EC16]

Total corruption (unknown to learner): $C = \sum_t c_t$.
Nature selects hidden parameter $\theta^\star \in \mathbb{R}^d$.
For round $t \in [T]$:

1. Nature chooses observable context $x_t \in \mathbb{R}^d$ with $||x_t||_2 \leq 1$.
   and corruption $c_t \in \{0, 1\}$
2. Learner selects *a price* $p_t \geq 0$.
3. Agent makes purchase if $\langle x_t, \theta^\star \rangle \geq p_t$ **if $c_t = 0$**
   or arbitrarily if $c_t = 1$
4. Learner observes if purchase occurred *(one-sided feedback).*
5. Learner collects revenue $p_t \cdot 1\{\langle x_t, \theta^\star \rangle \geq p_t\}$.

$$Regret = \sum_t \langle x_t, \theta^\star \rangle - \sum_t p_t \cdot 1\{\langle x_t, \theta^\star \rangle \geq p_t\}$$

[Lobel, Paes Leme, Vladu, EC17]: $\boldsymbol{Regret = O(d \log T)}$

Main Result

Robust contextual search algorithm **agnostic to C** achieving $\boldsymbol{Regret = O(Cd^3 \log^3 T)}$ (w. high prob.)

[Ulam 76], [Rivest, Meyer, Kleitman, Winklmann, Spencer, JCSS80], [Karp and Kleinberg, SODA07]
For Ulam's game ($d = 1$) tight bound of $\log T + C \log C + C \log\log T$ [RMKWS, JCSS80]

Extends classical **Ulam's game** (aka twenty questions with a liar) in 2 ways:

1. Multidimensional
2. Agnostic to C

Complications of Adversarial Irrationality

Binary search in 1d: after *at most* $\log\left(\frac{1}{\epsilon}\right)$ *queries* identify the hidden parameter $\theta^\star$ within an ϵ interval.

What if one of the adversary's responses is corrupted?

Completely wrong interval returned after $\log\left(\frac{1}{\epsilon}\right)$

A way out: Repeat queries!

Non-Robust Contextual Search Algorithm

[Lobel, Paes Leme, Vladu]: $Regret = O(d \log T)$

- Maintain active knowledge set with feasible values for $\theta^\star$.
- Set price $\langle \kappa_1, x_1 \rangle$ (κ_1: centroid).
- Eliminate inconsistent side.

Important properties of cut
1. Never eliminate $\theta^\star$.
2. Volumetric progress

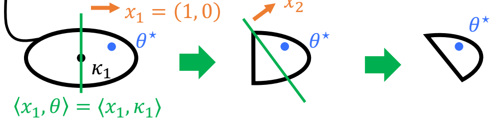

Challenges and Solutions

Known C

1. Cannot repeat same query (adversarially chosen x_t's)

→ Keep "penalty" for each θ. Make cut once region with penalty $\leq C$ (*protected region,* $\mathcal{P}(C)$) is fully on one side

2. May never have a context cut with $\mathcal{P}(C)$ fully on one side (even for $C = 1$ and infinitely many contexts).

→ Run algorithm in epochs.
→ Combine context cuts to a valid cut (w. properties in green box).
→ Epoch size to guarantee *existence* of valid cut: $2d\,(d + 1)C + 1$.

[Caratheodory's theorem]: every point $p \in \mathrm{ConvHull}(\mathcal{P}(C))$ can be written as the convex combination of $\leq d + 1$ points in $\mathcal{P}(C)$.

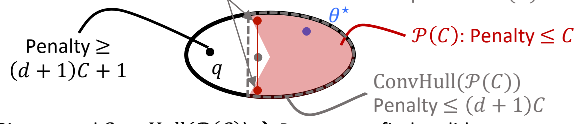

Given q and $\mathrm{ConvHull}(\mathcal{P}(C))$ → Perceptron finds valid cut.

3. Perceptron needs a margin γ to converge after $O\left(\frac{1}{\gamma^2}\right)$ mistakes.

- Sample random point q from ball (enough penalty for its center).
- W. const. prob., q has enough margin.

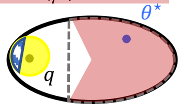

→ Use Perceptron *as a witness*:

- If q has enough margin, Perceptron terminates after $O\left(\frac{1}{\gamma^2}\right)$ rounds.
- Else, resample

4. From known to unknown C.

Subsampling and ***global eliminations*** of $\log T$ layers.

→ Extension of the multi-layering race technique of [Lykouris, Mirrokni, Paes Leme, STOC18] for continuous action spaces.