MaxGap Bandit:Adaptive Sampling for Approximate Ranking

Sumeet Katariya, Ardhendu Tripathy, Robert Nowak

University of Wisconsin-Madison

Problem

Motivation

1) Surveys

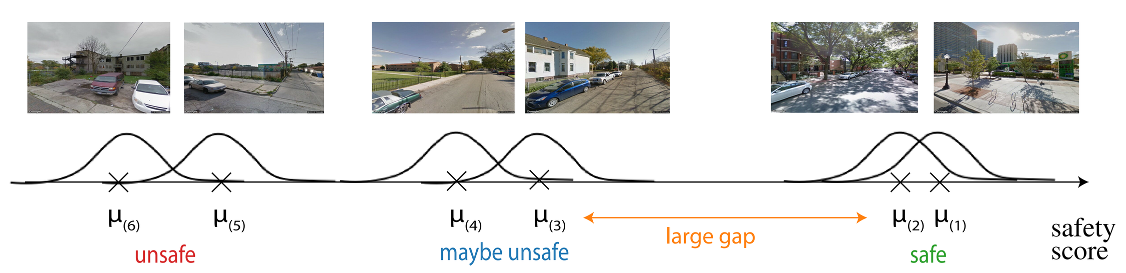

2) Grading

3) Identify Outliers

<u>Formulation</u>
Adaptively sample to partially order distributions according to their means (identify large gaps).

The simplest problem for 2 clusters is: MaxGap Bandit

MaxGap Bandit: Cluster using largest gap

- K arms with means $\mu_1, \ldots, \mu_K$.
- Arm with the largest gap

$$m = \arg\max_{i \in [K-1]} (\mu_{(i)} - \mu_{(i+1)})$$

- Defines two clusters

$$C_1 = \{(1), \ldots, (m)\} \quad \text{and} \quad C_2 = [K] \setminus C_1$$

- Let algorithm $\mathcal{A}$ stop after $T(\mathcal{A})$ samples and return clusters $\hat{C}_1, \hat{C}_2 = [K] \setminus \hat{C}_1$

Objective: $\min_{\mathcal{A}} T(\mathcal{A})$
subject to $\mathbb{P}(\hat{C}_1 \neq C_1) \leq 1 - \delta$

Existing Solutions don't work

- Best Arm Identification on gaps:
1) Only K out of $\binom{K}{2}$ gaps matter
2) arm rewards not independent
- Adaptive Sorting

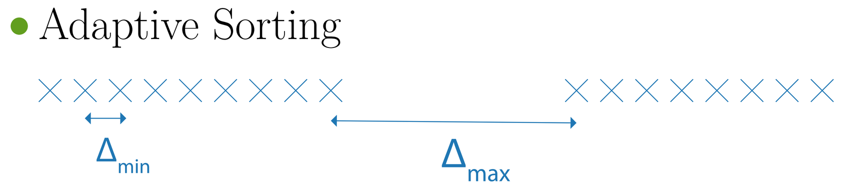

Requires $O(K/\Delta^2_{\min})$ samples (compared to $O(K/\Delta^2_{\max})$ needed in this work)

Solution

Algorithms

- <u>Elimination</u>: Play active arms whose gap upper bound is higher than the gap lower bound
- <u>UCB</u>: Play all (there will be multiple) arms with the highest gap UCB
- <u>Top2UCB</u>: Play arms with the top-two highest gap UCBs

How to compute confidence intervals on gaps from confidence intervals on means?

Gap Upper Bounds: MIP

Computing upper bounds requires solving a MIP

$$U\Delta^r_a = \max_{\{b \in [K]: b \neq a\}} \max_{\mu_1, \ldots, \mu_K} (\mu_b - \mu_a)$$

$$1)\ l_i \leq \mu_i \leq r_i \,\forall i \in [K]$$
$$2)\ \forall i \notin \{a, b\}, \mu_i \notin (\mu_a, \mu_b)$$

$$\boxed{\begin{aligned} &\mu_i \leq \mu_a + M(1 - y_i) \\ &\mu_i \geq \mu_b - M y_i \\ &y_i \in \{0, 1\} \end{aligned}}$$

Gap Upper Bounds: Efficient $O(K^2)$ Algorithm

Max-gap if μ_a known

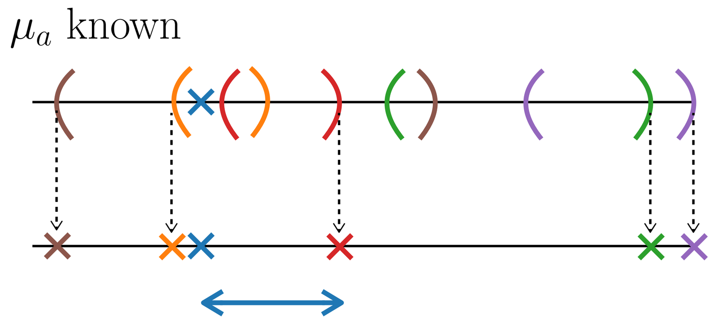

Max-gap if μ_a unknown

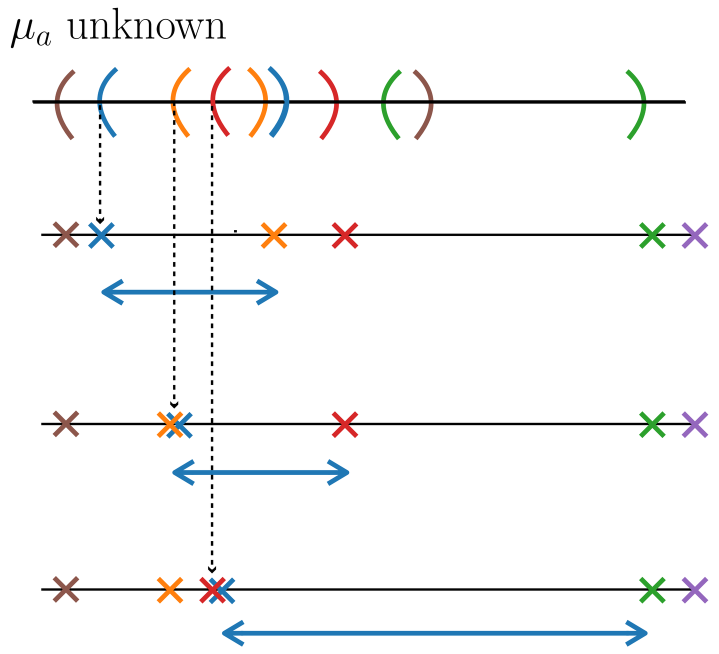

Gap Lower Bounds

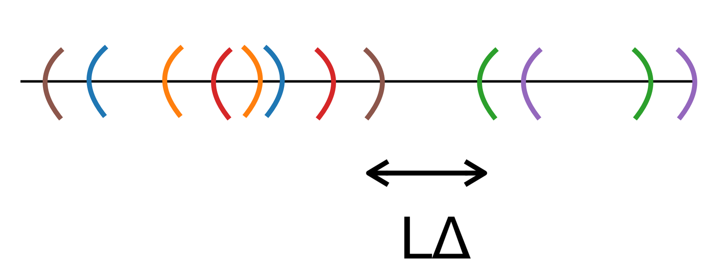

Analysis

Sample Complexity

$$H = \sum_{i \neq \{m, m+1\}} \frac{1}{\gamma_i^2}$$

where

$$\gamma_i = \min\{\gamma_i^r, \gamma_i^l\}$$

$$\gamma_i^r = \max_{j: \Delta_{i,j} > 0} \min\{\Delta_{i,j}, \Delta_{\max} - \Delta_{i,j}\}$$

$$\gamma_i^l = \max_{j: \Delta_{i,j} < 0} \min\{\Delta_{j,i}, \Delta_{\max} - \Delta_{j,i}\}$$

Distribution i can be eliminated quickly if there is another distribution j that has a moderately large gap from i (so that this gap can be easily detected), but not too large so that gap is easy to distinguish from $\Delta_{\max}$.

MaxGap Bandit $\neq$ Best-arm Identification on Gaps

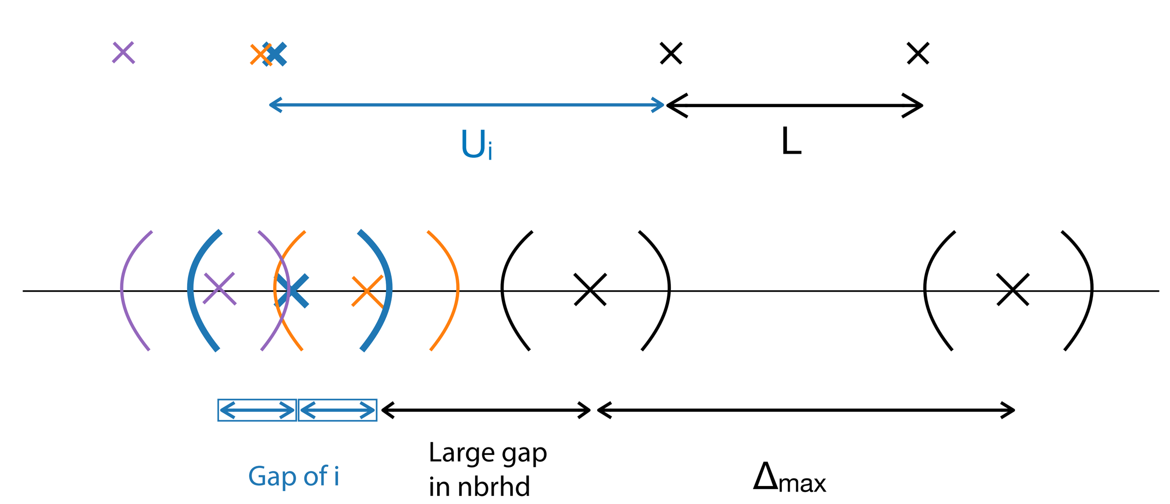

> An arm may have a small gap, but if there is a large gap in its neighborhood, it cannot be eliminated quickly.

Minimax Lower Bound

For the above problem, we prove a lower bound that matches the upper bound.

Experiments

Empirical vs Theoretical Stopping Time

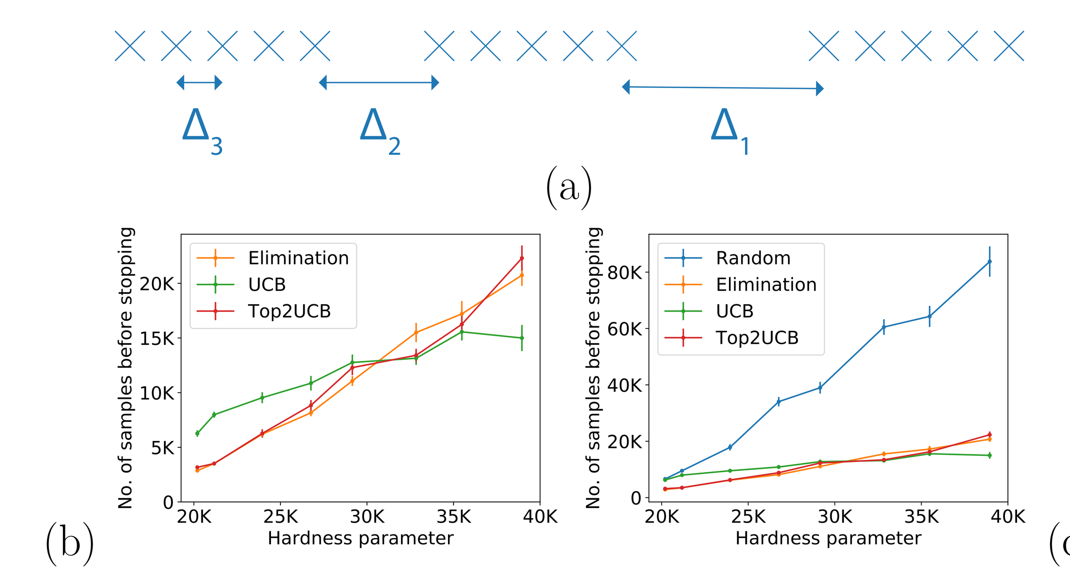

Adaptive Gains

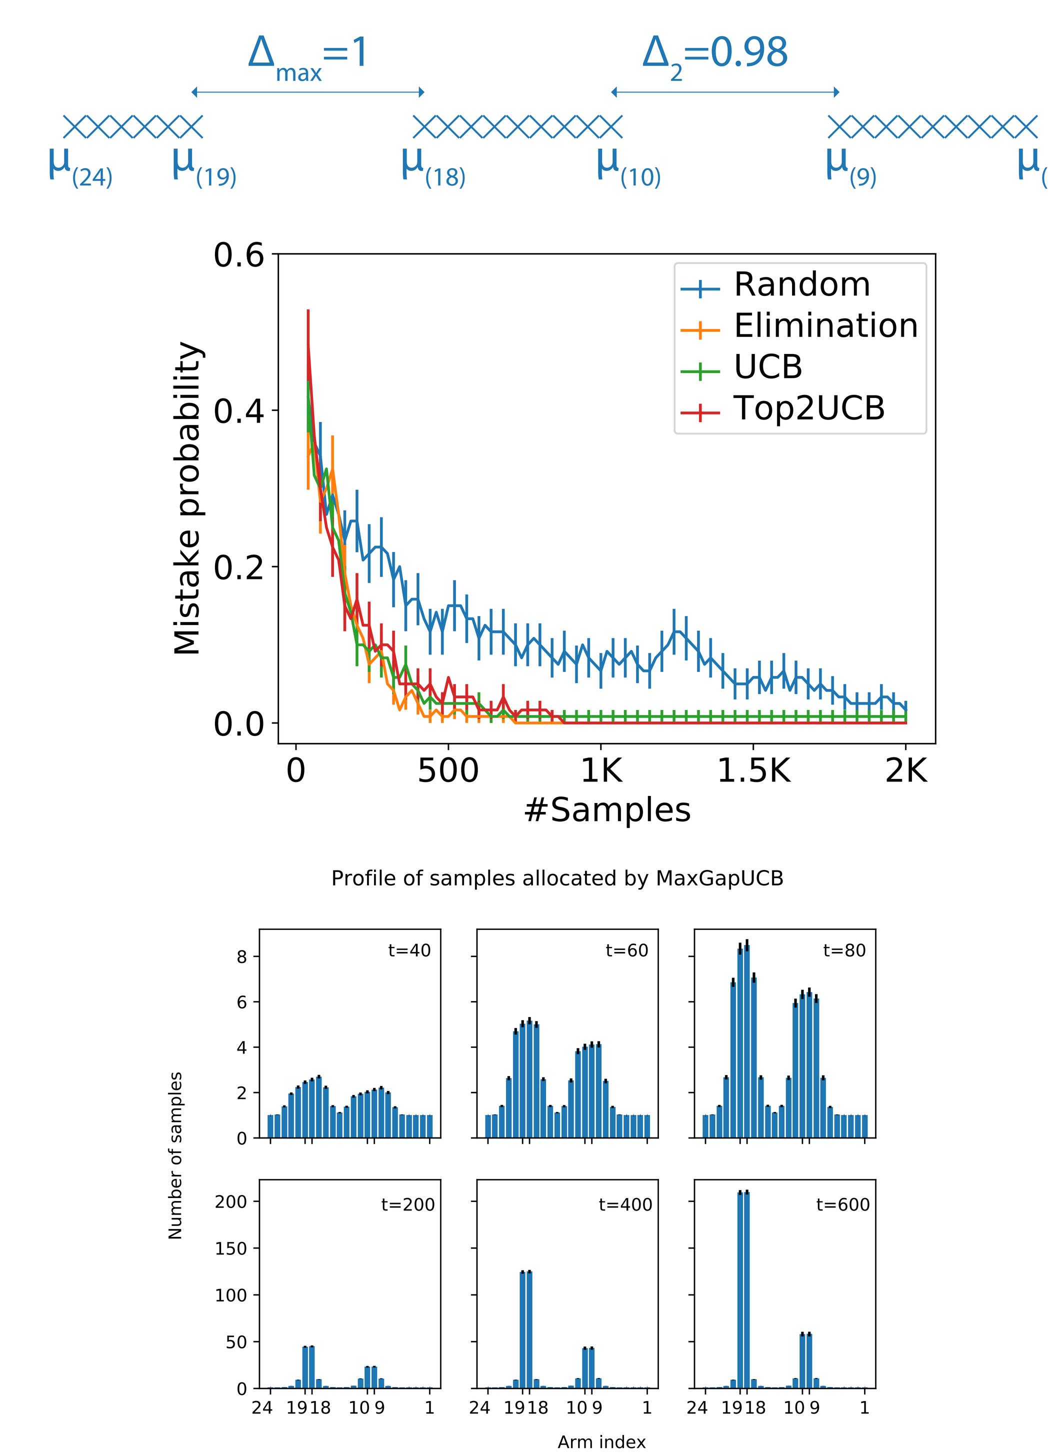

Streetscore Dataset