Single-pass Streaming Lower Bounds for Multi-armed Bandits Exploration with Instance-sensitive Sample Complexity

Chen Wang

Rutgers University

Joint work with Sepehr Assadi

- Motivations and Definitions

Best Arm Identification Problem

Best Arm Identification Problem

- Given n arms with unknown rewards.

Best Arm Identification Problem

- Given n arms with unknown rewards.
- The rewards follows Sub-Gaussian distributions, e.g. Bernoulli distribution.

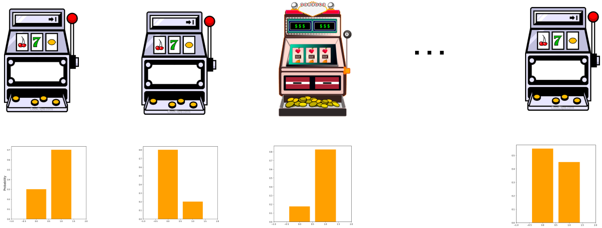

Best Arm Identification Problem

- Given n arms with unknown rewards.
- The rewards follows Sub-Gaussian distributions, e.g. Bernoulli distribution.

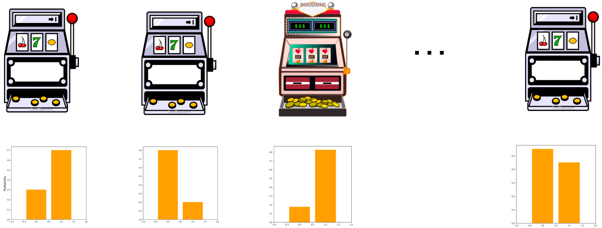

- Goal: find arm* — the *best* arm with highest *mean* of the distribution.

Best Arm Identification Problem

- Given n arms with unknown rewards.
- The rewards follows Sub-Gaussian distributions, e.g. Bernoulli distribution.

- Parameters $\Delta_{[i]}$: the gap between the best and the i-th best arms.

- Parameters may or may not be known.

Best Arm Identification Problem

Best Arm Identification Problem

- Natural Strategy: Pull each arm multiple times and record the empirical rewards.

Best Arm Identification Problem

- Natural Strategy: Pull each arm multiple times and record the empirical rewards.

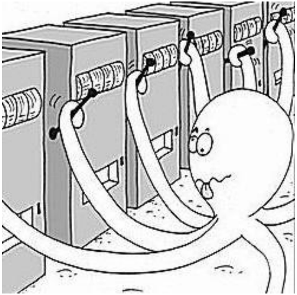

Best Arm Identification Problem

- ❑ Natural Strategy: Pull each arm multiple times and record the empirical rewards.

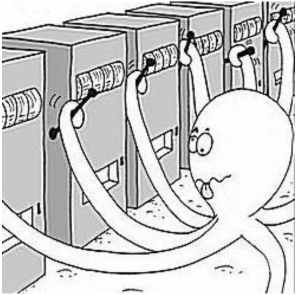

- ❑ Goal: Finding the best arm:
  - ➢ High enough (constant) probability
  - ➢ Minimum number of arm pulls (sample complexity).

Classical Results

Worst-case (w.r.t. $\Delta_{[2]}$):

- $\Theta(\frac{n}{\Delta_{[2]}^2})$ sample complexity.
- Upper bound — Median Elimination (Even-Dar et al., [COLT'02]).
- Asymptotically matching lower bound (Mannor and Tsitsiklis, [JMLR'04]).

Classical Results

Instance-sensitive (w.r.t. $\{\Delta_{[i]}\}_{i=2}^{n}$):

- $O(H_2 := \sum_{i=2}^{n} \frac{1}{\Delta_{[i]}^2} \log\log(\frac{1}{\Delta_i}))$ sample complexity.
- Upper bound algorithms (Karnin et al., [ICML'13]; Jamieson et al., [COLT'14]).

Classical Results

- Another interesting variate: instance-sensitive sample complexity.

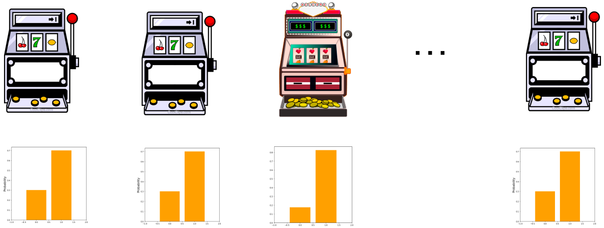

$$\Theta(\frac{n}{\Delta_{[2]}^2}) \approx H_2$$

- Another interesting variate: instance-sensitive sample complexity.

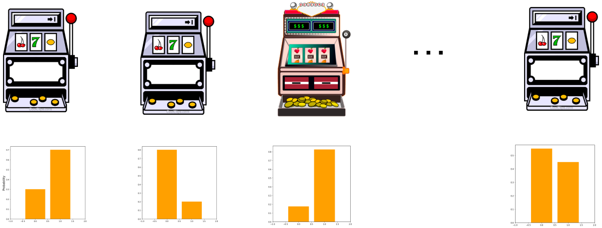

$$\Theta(\frac{n}{\Delta_{[2]}^2}) >> H_2$$

A twist on perspective: Memory

- So far most of the studies focused on smaller number of arm pulls (sample complexity).

A twist on perspective: Memory

- So far most of the studies focused on smaller number of arm pulls (sample complexity).

- But is it practical to store all the arms in massive datasets?

A twist on perspective: Memory

- So far most of the studies focused on smaller number of arm pulls (sample complexity).

- But is it practical to store all the arms in massive datasets?

- In large-scale applications, space becomes important.

Streaming Multi-armed Bandits Problem

Streaming Multi-armed Bandits Problem

The Streaming Model (Assadi and Wang, [STOC'20])

- A stream of arms
- Algorithm on-the-fly read arms to memory
- Algorithm can pull the arriving arm and store it
- Algorithm can pull a stored arm
- Algorithm can discard a stored arm

➢ Not Stored or Discarded -> cannot be retrieved anymore (lost forever)

Streaming Multi-armed Bandits

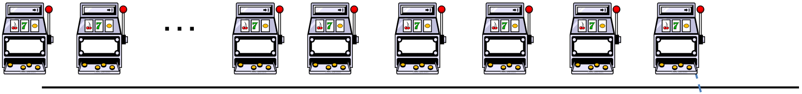

t=1, read/pull arm 1

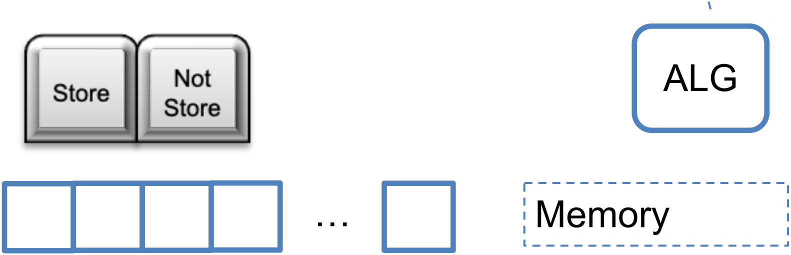

Streaming Multi-armed Bandits

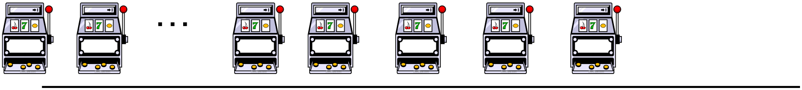

Store Arm 1

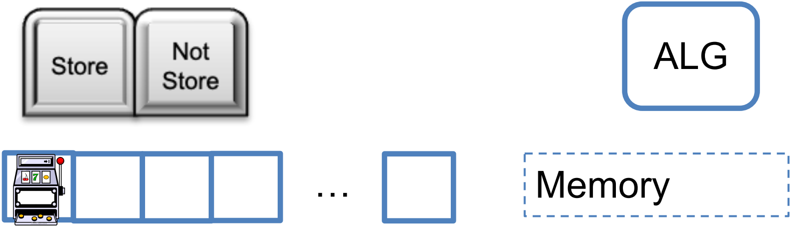

Streaming Multi-armed Bandits

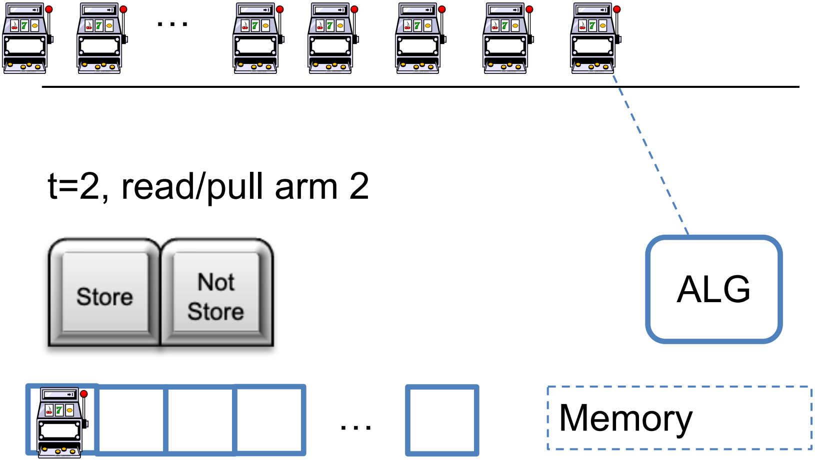

Streaming Multi-armed Bandits

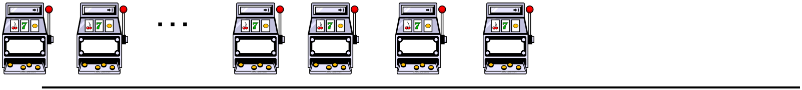

Store Arm 2

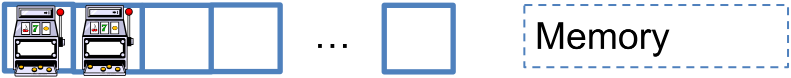

Streaming Multi-armed Bandits

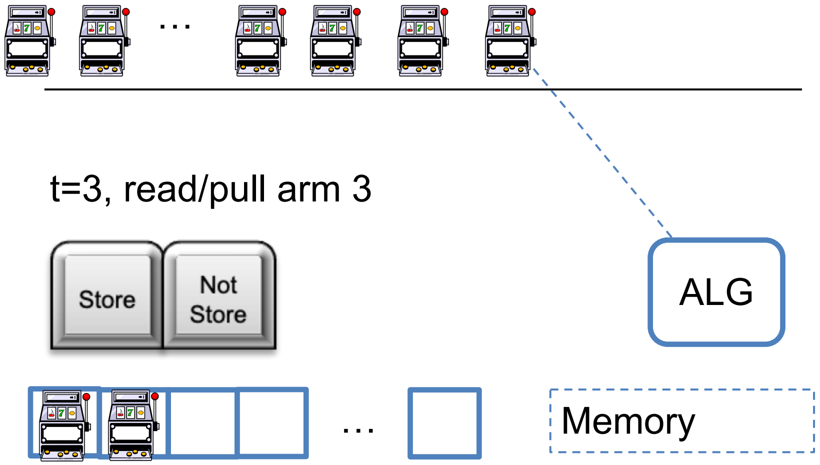

Streaming Multi-armed Bandits

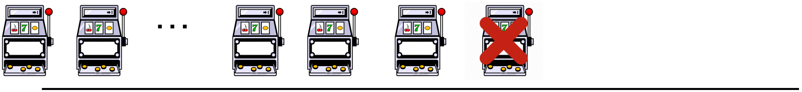

Not Store Arm 3
(lost forever)

ALG

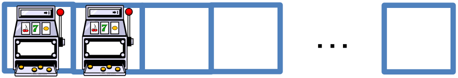

Memory

Streaming Multi-armed Bandits

Continue till the last arm

Streaming Multi-armed Bandits

Continue till the last arm

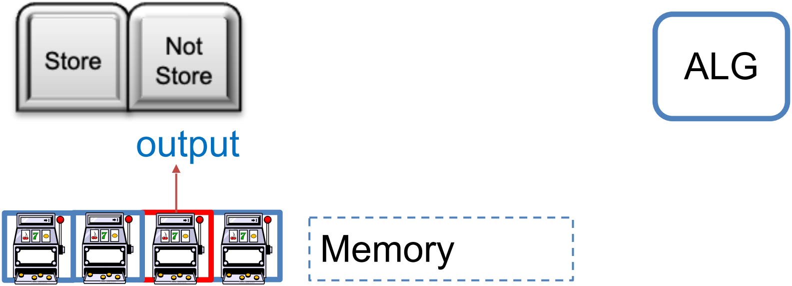

Memory Complexity: the *maximum* number of arms we ever stored.

Continue till the last arm

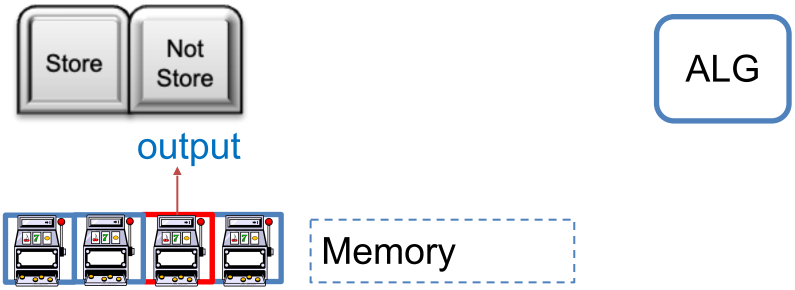

Streaming MABs Exploration: Worst-case Complexity

Streaming MABs Exploration: Worst-case Complexity

- (Assadi and Wang, [STOC'20]) There exists an algorithm that given parameter $\Delta_{[2]}$, finds the best arm with high constant probability and:

- $O(\frac{n}{\Delta_{[2]}^2})$ Sample Complexity
- *A Single arm* Space Complexity

Streaming MABs Exploration: Worst-case Complexity

- (Assadi and Wang, [STOC'20]) There exists an algorithm that given parameter $\Delta_{[2]}$, finds the best arm with high constant probability and:

- $O(\frac{n}{\Delta_{[2]}^2})$ Sample Complexity
- *A Single arm* Space Complexity

Same memory as we need to *output* the best arm.

Beyond Worst-case Streaming Sample Complexity

- Open questions following Assadi and Wang, [STOC'20]:

- What if parameter $\Delta_{[2]}$ is not given?

- What about the instance-sensitive exploration?

$$O(H_2 := \sum_{i=2}^{n} \frac{1}{\Delta_{[i]}^2} \log\log(\frac{1}{\Delta_i}))$$

Beyond Worst-case Streaming Sample Complexity

- Open questions following Assadi and Wang, [STOC'20]:

- What if parameter $\Delta_{[2]}$ is not given?

- What about the instance-sensitive exploration?

$$O(H_2 := \sum_{i=2}^{n} \frac{1}{\Delta_{[i]}^2} \log\log(\frac{1}{\Delta_i}))$$

Answers to the above questions in this work.

Main Results: Streaming Lower Bounds

First Lower Bound — Role of Parameter $\Delta_{[2]}$

First Lower Bound — Role of Parameter $\Delta_{[2]}$

- A very strong lower bound for streaming MABs algorithms without the knowledge of $\Delta_{[2]}$.

First Lower Bound — Role of Parameter $\Delta_{[2]}$

- A very strong lower bound for streaming MABs algorithms without the knowledge of $\Delta_{[2]}$.

- If without the knowledge of $\Delta_{[2]}$, and no other knowledge, any streaming algorithm to find the best arm:

- Either has a memory of $\Theta(n)$ arms.
- Or the sample complexity is $\omega(\frac{n}{f(\Delta_{[i]})})$ for every non-zero f.

First Lower Bound — Role of Parameter $\Delta_{[2]}$

- Without the knowledge of $\Delta_{[2]}$, sample complexity is unbounded as a function of $\Delta_{[2]}$.

First Lower Bound — Role of Parameter $\Delta_{[2]}$

- Without the knowledge of $\Delta_{[2]}$, sample complexity is unbounded as a function of $\Delta_{[2]}$.

- Sharp contrast to the upper bound by Assadi and Wang, [STOC'20].

First Lower Bound — Role of Parameter $\Delta_{[2]}$

- Without the knowledge of $\Delta_{[2]}$, sample complexity is unbounded as a function of $\Delta_{[2]}$.

- Sharp contrast to the upper bound by Assadi and Wang, [STOC'20].

- Sharp contrast to the upper bounds in offline algorithms — Karnin et al., [ICML'13]; Jamieson et al., [COLT'14] do not assume a known $\Delta_{[2]}$!

- In a single-pass stream, we proved lower bounds of $\omega(H_2)$ arm pulls for any streaming algorithm with $o(n)$ arm memory if:
  - Only the parameter $\Delta_{[2]}$ is given.
  - If the parameter $\Delta_{[2]}$ is given, and the arms arrive in a random order.
  - If all the gap parameters $\{\Delta_{[i]}\}_{i=2}^{n}$ are given.

A Streaming Algorithm with Strong Assumptions

- Finally, we devise a single-pass streaming algorithm when all assumptions hold.

- There exists an algorithm that given parameter $\Delta_{[2]}$, a random order of the stream, and the value of H_2, finds the best arm with high constant probability and:

- $O(H_2 + \text{poly}(\frac{\log(n)}{\Delta_{[2]}}))$ Sample Complexity

- *A Single arm* Space Complexity

- Our algorithm hinges on new ideas of budgeting the stored arm with number of arm pulls.

- The additive $O(\text{poly}(\frac{\log(n)}{\Delta_{[2]}}))$ term is negligible in large-scale application context ($n \sim 10^6$ and $\Delta_{[2]} \sim 0.1$ suffice).

Open Problems

❑ Multi-pass bounds for MABs:

- What if we are allowed to make multiple passes over the arms?
- Jin et al. [ICML'21]: $O(\log(1/\Delta_{[2]}))$ passes, $O(H_2)$ sample complexity.
- Without the knowledge of $\Delta_{[2]}$.
- What are the tight pass upper and lower bounds?

Open Problems

- Multi-pass bounds for MABs:
  - What if we are allowed to make multiple passes over the arms?
  - Jin et al. [ICML'21]: $O(\log(1/\Delta_{[2]}))$ passes, $O(H_2)$ sample complexity.
  - Without the knowledge of $\Delta_{[2]}$.
  - What are the tight pass upper and lower bounds?

Thank You!