On the Problem of Best Arm Retention

Houshuang Chen, Yuchen He and Chihao Zhang

Shanghai Jiao Tong University

2024.7

Best Arm Identification

Best Arm Identification

Input: n arms

Best Arm Identification

Input: n arms



Best Arm Identification

Input: n arms



In round t:

A_t: the player's choice

$r_t(A_t) \sim \mathcal{D}(\mu_{A_t})$

Best Arm Identification

Input: n arms



Output: one arm

In round t:

A_t: the player's choice

$r_t(A_t) \sim \mathcal{D}(\mu_{A_t})$

Best Arm Identification

Input: n arms



Output: one arm



In round t:

A_t: the player's choice

$r_t(A_t) \sim \mathscr{D}(\mu_{A_t})$

Best Arm Identification

Input: n arms



Output: one arm



In round t:

A_t: the player's choice

$r_t(A_t) \sim \mathscr{D}(\mu_{A_t})$

(ε, δ)-BAI: **identify** an ε-optimal arm with probability $1 - \delta$

Best Arm Identification

Input: n arms 

Output: one arm 

In round t:

A_t: the player's choice

$r_t(A_t) \sim \mathscr{D}(\mu_{A_t})$

(ε, δ)-BAI: **identify** an ε-optimal arm with probability $1 - \delta$

optimal arm:

$$a^* = \arg\max_{i \in [n]} \mu_i$$

ε-optimal arm:

$$\mu_{a^*} - \mu < \varepsilon$$

Best Arm Identification

Input: n arms 

Output: one arm 

In round t:

A_t: the player's choice

$r_t(A_t) \sim \mathcal{D}(\mu_{A_t})$

(ε, δ)-BAI: **identify** an ε-optimal arm with probability $1 - \delta$

Sample complexity bound: $\Theta\left(\frac{n}{\varepsilon^2}\log\frac{1}{\delta}\right)$

(see [EMM02, MT04])

optimal arm:

$a^* = \arg\max_{i \in [n]} \mu_i$

ε-optimal arm:

$\mu_{a^*} - \mu < \varepsilon$

Best Arm Retention

Best Arm Retention

Input: n arms

Best Arm Retention

Input: n arms



Best Arm Retention

Input: n arms



Output: a subset of m arms

Best Arm Retention

Input: n arms



Output: a subset of m arms



Best Arm Retention

Input: n arms 

Output: a subset of m arms 

(ε, δ)-BAR: **retain** an ε-optimal arm with probability $1 - \delta$

Best Arm Retention

Input: n arms 

Output: a subset of m arms 

(ε, δ)-BAR: **retain** an ε-optimal arm with probability $1 - \delta$

ε-BAR: **retain** an arm which is ε-optimal in expectation

Best Arm Retention

Input: n arms 

Output: a subset of m arms 

(ε, δ)-BAR: **retain** an ε-optimal arm with probability $1 - \delta$

ε-BAR: **retain** an arm which is ε-optimal in expectation

Target:

small **sample complexity** / small **regret**

Best Arm Retention

Input: n arms 

Output: a subset of m arms 

(ε, δ)-BAR: **retain** an ε-optimal arm with probability $1 - \delta$

ε-BAR: **retain** an arm which is ε-optimal in expectation

Target:

small **sample complexity** / small **regret**

Regret:

$$R(T) = \mathbf{E}\left[\sum_{t=1}^{T} r_t(a^*) - r_t(A_t)\right]$$

(ε, δ)-BAR

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly: 

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly: 

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly: 

2. Run an (ε, δ')-BAI algorithm on this subset:

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:



2. Run an (ε, δ')-BAI algorithm on this subset:



(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:



2. Run an (ε, δ')-BAI algorithm on this subset:



(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:



2. Run an (ε, δ')-BAI algorithm on this subset:

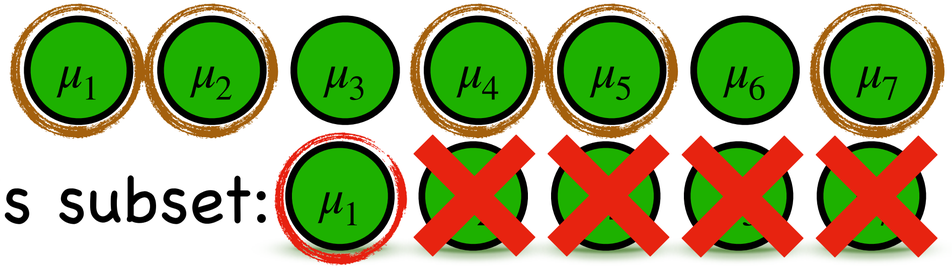


3. Discard the other $n - m$ arms and output:

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:
2. Run an (ε, δ')-BAI algorithm on this subset:
3. Discard the other $n - m$ arms and output:



(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly: 

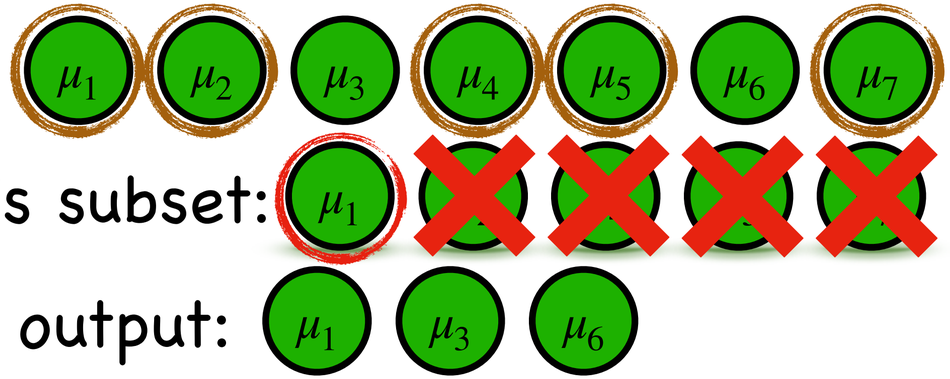
2. Run an (ε, δ')-BAI algorithm on this subset: 

3. Discard the other $n - m$ arms and output: 

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:
2. Run an (ε, δ')-BAI algorithm on this subset:
3. Discard the other $n - m$ arms and output:

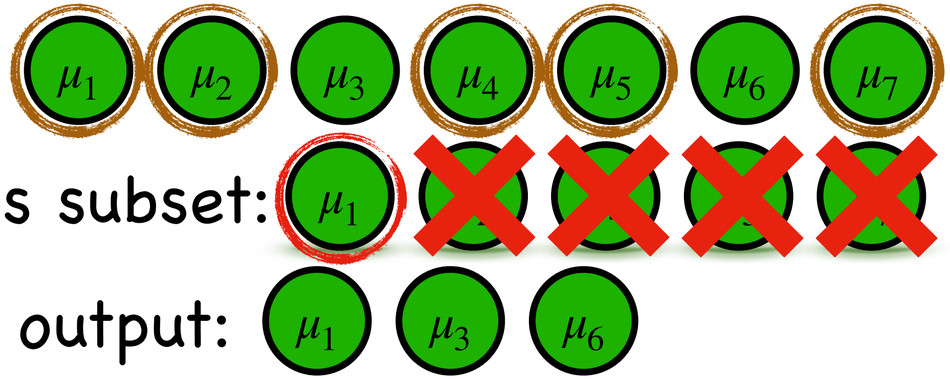


The optimal arm is chosen in step 1 with probability $\Theta\left(\frac{n-m}{n}\right)$

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:
2. Run an (ε, δ')-BAI algorithm on this subset:
3. Discard the other $n - m$ arms and output:



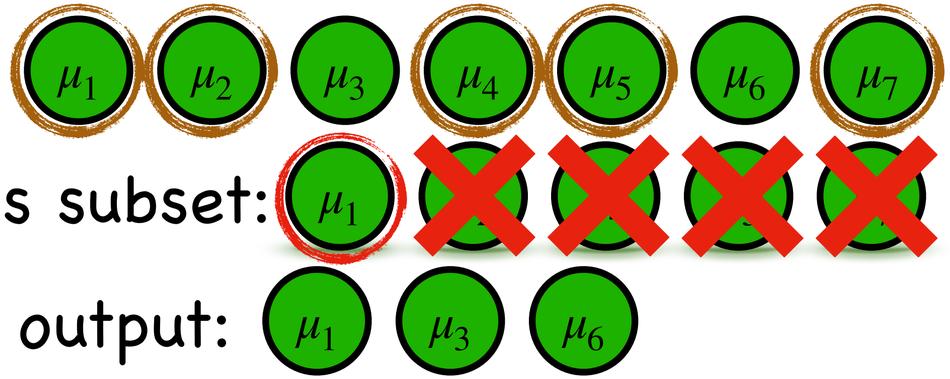
The optimal arm is chosen in step 1 with probability $\Theta\left(\frac{n-m}{n}\right)$

Choose $\delta' = \Theta\left(\frac{n\delta}{n-m}\right)$

(ε, δ)-BAR

(ε, δ)-BAR: retain an ε-optimal arm with probability $1 - \delta$

1. Choose $n - m + 1$ arms randomly:
2. Run an (ε, δ')-BAI algorithm on this subset:
3. Discard the other $n - m$ arms and output:



The optimal arm is chosen in step 1 with probability $\Theta\left(\frac{n-m}{n}\right)$

Choose $\delta' = \Theta\left(\frac{n\delta}{n-m}\right)$

Sample $\mathrm{O}\left(\frac{n-m}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$ times

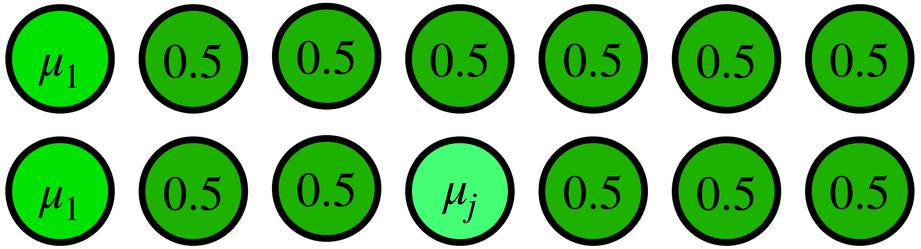
Why this is enough?

Why this is enough?

$\mathcal{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

Why this is enough?

$\mathcal{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

$\mathcal{H}_j$:

$\mu_j = 0.5 + 2\varepsilon$

Why this is enough?

$\mathcal{H}_1$: 

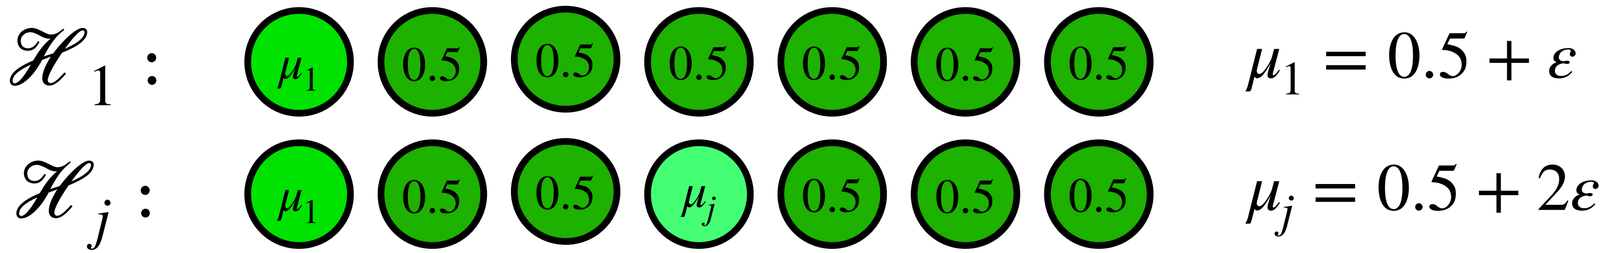
$\mu_1 = 0.5 + \varepsilon$

$\mathcal{H}_j$: 

$\mu_j = 0.5 + 2\varepsilon$

Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathcal{H}_1$

Why this is enough?



Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

Arm j_k should be retained on $\mathscr{H}_{j_k}$

Why this is enough?

$\mathscr{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

$\mathscr{H}_j$: 

$\mu_j = 0.5 + 2\varepsilon$

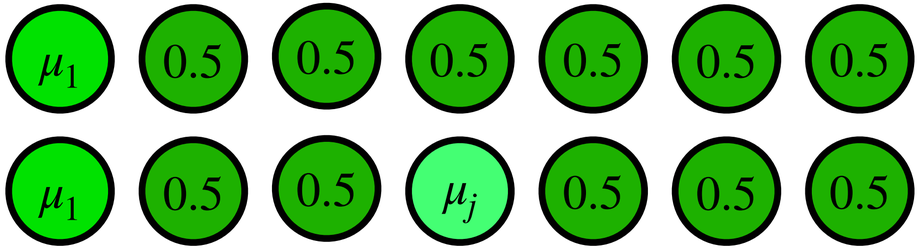
Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

Arm j_k should be retained on $\mathscr{H}_{j_k}$

$\Rightarrow$

Distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

Why this is enough?

$\mathscr{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

$\mathscr{H}_j$:

$\mu_j = 0.5 + 2\varepsilon$

Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

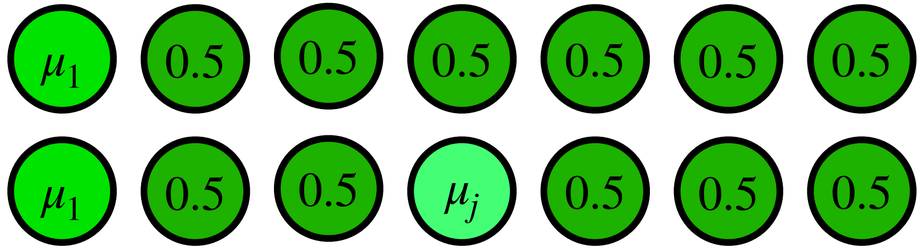
Arm j_k should be retained on $\mathscr{H}_{j_k}$

$\Rightarrow$

Distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

$\Omega\left(\frac{1}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$ samples

Why this is enough?

$\mathscr{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

$\mathscr{H}_j$:

$\mu_j = 0.5 + 2\varepsilon$

Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

Arm j_k should be retained on $\mathscr{H}_{j_k}$

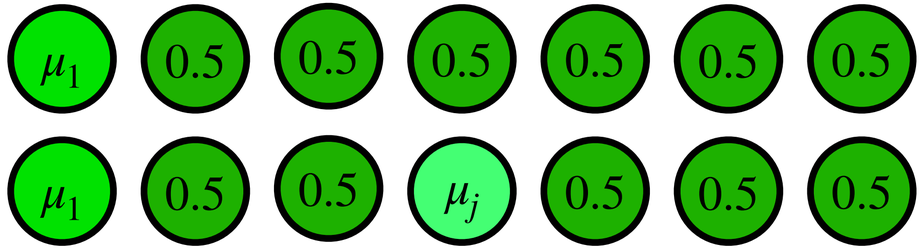
$\Rightarrow$

Distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

$\Omega\left(\frac{1}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$ samples

Technical difficulties:

Why this is enough?

$\mathscr{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

$\mathscr{H}_j$:

$\mu_j = 0.5 + 2\varepsilon$

Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

Arm j_k should be retained on $\mathscr{H}_{j_k}$

$\Rightarrow$

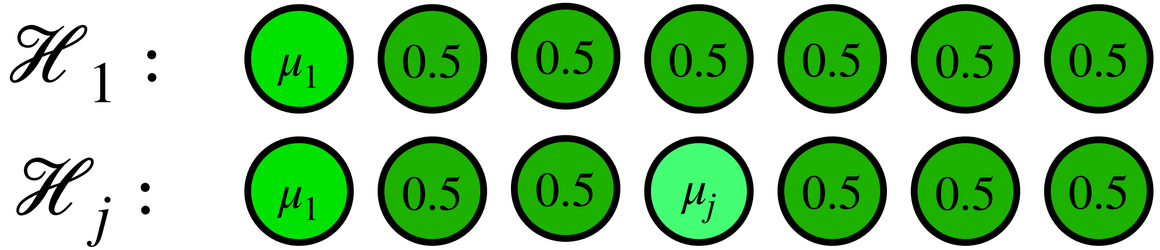
Distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

$\Omega\left(\frac{1}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$ samples

Technical difficulties:

- Success retention $\Leftrightarrow$ distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

Why this is enough?



$\mu_1 = 0.5 + \varepsilon$

$\mu_j = 0.5 + 2\varepsilon$

Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

Arm j_k should be retained on $\mathscr{H}_{j_k}$

$\Rightarrow$

Distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

$\Omega\left(\frac{1}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$ samples

Technical difficulties:

- Success retention $\Leftrightarrow$ distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$
- $j_1, j_2, \ldots, j_{n-m}$ are randomized

Why this is enough?

$\mathscr{H}_1$: 

$\mu_1 = 0.5 + \varepsilon$

$\mathscr{H}_j$: 

$\mu_j = 0.5 + 2\varepsilon$

Discard arms $j_1, j_2, \ldots, j_{n-m}$ on $\mathscr{H}_1$

Arm j_k should be retained on $\mathscr{H}_{j_k}$

$\Rightarrow$ Distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$

$\Omega\left(\frac{1}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$ samples

$$\Rightarrow \Omega\left(\frac{n-m}{\varepsilon^2}\log\frac{n-m}{n\delta}\right) \text{ samples in total}$$

Technical difficulties:

- Success retention $\Leftrightarrow$ distinguish $\mathscr{H}_{j_k}$ and $\mathscr{H}_1$
- $j_1, j_2, \ldots, j_{n-m}$ are randomized

Our Results

Problem	Target	Bound
(ε, δ)-BAR	Sample complexity	$\Theta\left(\frac{n-m}{\varepsilon^2}\log\frac{n-m}{n\delta}\right)$
ε-BAR	Sample complexity	$\Theta\left(\frac{(n-m)^3}{n^2\varepsilon^2}\right)$
ε-BAR	Regret	$\mathrm{O}\left(\frac{(n-m)^2}{n\varepsilon}\left(1+\sqrt{\frac{m}{n-m}}\right)\right), \Omega\left(\frac{(n-m)^2}{n\varepsilon}\right)$

Open problems

Open problems

What is the minimum regret needed for ε-BAR?

Open problems

What is the minimum regret needed for ε-BAR?

Is there a unified algorithm with **both** small **sample complexity** and small **regret**?

Open problems

What is the minimum regret needed for ε-BAR?

Is there a unified algorithm with **both** small **sample complexity** and small **regret**?

Is there a **trade-off** between sample complexity and regret?