Fixed Budget is No Harder Than Fixed Confidence in Best-Arm Identification up to Logarithmic Factors

Kapilan Balagopalan*, Yinan Li*, Yao Zhao, Tuan Ngo Nguyen, Anton Daitche, Houssam Nassif, Kwang-Sung Jun

Best Arm Identification (BAI): Fixed Confidence (FC) vs Fixed Budget (FB)

- Environment: K arms $\{1, \cdots, K\}$ with unknown mean rewards $\{\nu_1, \cdots \nu_K\}$. (i^* is the index of the unique best arm)

FB

- Initial budget B
- for each time step $t \leq B$ the learner samples an arm $I_t \in [K]$
- Observes reward $R_t \sim \nu_{I_t}$
- Learner outputs an estimated best arm $\hat{J}$ at the end of budget B

Goal: Minimize $\mathbb{P}(\hat{J} \neq i^*)$ (mis-identification probability)

FC

- Failure rate δ
- for each time step t the learner samples an arm $I_t \in [K]$
- Observes reward $R_t \sim \nu_{I_t}$
- **If the stopping condition is met** the learner outputs an estimated best arm $\hat{J}$ at the end of budget B. ($\tau = t$)

Certification: $\mathbb{P}(\hat{J} \neq i^*) \leq \delta$
Goal: Minimize τ (stopping time)

Sample complexity: Fixed Confidence (FC) vs Fixed Budget (FB)

- Although FC and FB are different configurations, we can still compare an FC algorithm and FB algorithm based on '**sample complexity**''

FB: Sample complexity

- The budget B necessary to such that $\mathbb{P}(\hat{J} \neq i^*) \leq \delta$
- Example: If an FB algorithm guarantees

$$\mathbb{P}(\hat{J} \neq i^*) \leq F \exp(-\frac{B}{H}) \textbf{ then}$$

$$B \geq H \log(F/\delta) \implies \mathbb{P}(\hat{J} \neq i^*) \leq \delta$$

FC: Sample complexity

- **Certification:** $\mathbb{P}(\hat{J} \neq i^*) \leq \delta$ is already there.
- Sample complexity = τ (Stopping time)

Research questions

- **Question 1:** Which configuration is easy to solve ? Is FB setting is harder/easier than FC setting ?
  - **Answer:** FB is no harder than FC up to logarithmic factors
- **Question 2:** Is it possible to achieve a reduction of one setting to the other ?
  - **Answer:** Yes

We answer the above research questions by constructing FC2FB algorithm that takes a black-box FC algorithm and transforms it into an FB an algorithm.

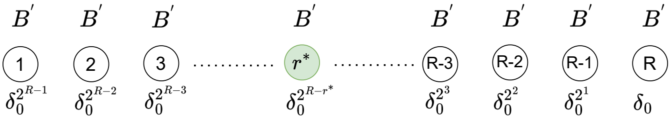

FC2FB: Algorithm

Algorithm

Algorithm 3 FC2FB

```
Input: Total budget B, algorithm A, base failure rate δ_0, Q ≤ B/2
Set R := ⌊log_2(B/Q)⌋, B' := ⌊B/R⌋
Ĵ ← any arbitrary arm
for each stage r = 1, 2, ..., R do
    L_r = 2^{R−r}
    Run the algorithm A(δ_0^{L_r}) with the budget limit of B'
    If the algorithm did not self-terminate, set Ĵ_r ← 0; otherwise, let Ĵ_r be the output arm from the algorithm.
    if Ĵ_r ≠ 0 then
        Ĵ ← Ĵ_r
        break
    end if
end for
Output: Ĵ
```

- Budget B is equally allocated to R stages with budget $B' = \frac{B}{R}$ each.
- For each stage $r \leq R$, FC algorithm $\mathcal{A}$ is initialized with doubly exponentially increasing failure rate $\delta = \delta_0^{2^{R-r}}$
- If $\mathcal{A}$ self-terminates at stage r (stopping condition met) then output $\hat{J}$ is set to be output of the $\mathcal{A}$ at stage r.
- If none of the instance self-terminates, output is set to be an arbitrary arm.

FC2FB: Main results

- If the input FC algorithm $\mathcal{A}$ satisfies a δ-stopping time,

$$T^*_\delta = A \ln(1/\delta) + C$$

- Then, FC2FB satisfies

Theorem (Correctness of FC2FB)

For a strong fixed-confidence algorithm $\mathcal{A}$, given a total budget $B \geq 2\left(A \ln\left(\frac{1}{\delta_0}\right) + (C+1)\right) \ln\left(2\frac{A}{Q} \ln\left(\frac{1}{\delta_0}\right) + \frac{2(C+1)}{Q}\right)$, $\delta_0 \leq 0.5$ and $Q \leq \frac{B}{2}$, algorithm FC2FB satisfies,

$$\mathbb{P}\left(\hat{J} \neq 1\right) \leq 3 \exp\left(-\frac{B}{\frac{4Q}{\ln(\frac{1}{\delta_0})} + 4A \log_2\left(\frac{B}{Q}\right)}\right).$$

FC2FB: FB is no harder than FC

- Our main theorem also implies that the sample complexity of FC2FB is,

$$O\left(A \ln(1/\delta) \cdot \ln(\ln(1/\delta)) + C\right) \ .$$

- This sample complexity is of same order that of the FC algorithm $\mathcal{A}$ ignoring the $\log \log(1/\delta)$
- Hence **FB is no harder than FC** up to logarithmic factors.
- We constructively proved that For all FC algorithms, there exist a FB algorithm with equal sample efficiency up to logarithmic factors using a generic framework regardless of problem domain. (Eg: K-armed bandits, Unimodal bandits, Linear bandits, Cascading bandits, etc.)

Contribution

- **Theory:** Research questions answered:
  - **Question 1:** Which configuration is easy to solve ? Is FB setting is harder/easier than FC setting ?
    - **Answer:** FB is no harder than FC up to logarithmic factors
  - **Question 2:** Is it possible to achieve a reduction of one setting to the other ?
    - **Answer:** Yes
- **Application:** Practical tool to construct an FB algorithm for any domains.
  - This tool can be used to construct FB algorithm from efficient FC algorithms.
  - We have illustrated such applications for linear, unimodal and cascading bandits.

Experiments

- We conduct experiments for K-armed bandits with known heterogeneous noise.
- In known heterogeneous noise case, we first construct an FC algorithm PE-KHN that has better sample efficiency compared to existing FB algorithms SH and SHVar.
- We then apply FC2FB on PE-KHN and show that the resulting FB algorithm performs better than SH and SHVar.

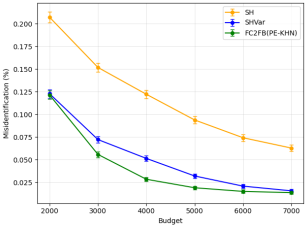

Probability of misidentifying the best arm in the Gaussian bandit as a function of the budget B, with $K = 32$ arms. Results are averaged over 5,000 trials.

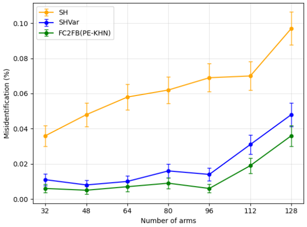

Probability of misidentifying the best arm in the Gaussian bandit as a function of the number of arms. The budget is fixed at 6000. Results are averaged over 5,000 trials.