

Open Problem: Optimal Best Arm Identification with Fixed Budget

Chao Qin
Columbia University

CQ2199@COLUMBIA.EDU

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

Best arm identification or *pure exploration* problems have received much attention in the COLT community since [Bubeck et al. \(2009\)](#) and [Audibert et al. \(2010\)](#). For any bandit instance with a unique best arm, its asymptotic complexity in the so-called *fixed-confidence setting* has been completely characterized in [Garivier and Kaufmann \(2016\)](#) and [Chernoff \(1959\)](#), while little is known about the asymptotic complexity in its “dual” setting called *fixed-budget setting*. This note discusses the open problems and conjectures about the instance-dependent asymptotic complexity in the fixed-budget setting.

Keywords: multi-armed bandit, best arm identification, pure exploration, asymptotic complexities

1. Introduction and problem formulation

We consider the so-called best arm identification (BAI) or pure exploration problems where there is a finite number of arms. An experimenter can sequentially select arms to measure and observes independent noisy observations of their quality. The experimenter’s goal is to confidently identify a best arm through allocating measurement effort in an adaptive and intelligent manner. BAI problems have also been studied under different names for several decades, e.g., *ranking and selection* or *ordinal optimization* in the literature of statistics and operations research. The literature of machine learning mainly studies BAI problems in two settings. One is called *fixed-confidence setting* where the objective is minimizing the expected number of collected samples while guaranteeing the probability of incorrect decision after the stopping time less than a pre-specified level, and the other is called *fixed-budget setting* where the objective is minimizing the probability of incorrect decision after a given budget of samples is used up. For any bandit instance with a unique best arm, its asymptotic complexity in the fixed-confidence setting has been fully characterized. See for example, [Garivier and Kaufmann \(2016\)](#) and [Chernoff \(1959\)](#). Although both settings seem “dual” to each other, the instance-dependent asymptotic complexity in the fixed-budget setting is unclear for a very long time. This note briefly include the existing results in the fixed-confidence setting and discusses the open problems and conjectures about the instance-dependent asymptotic complexity in the fixed-budget setting.

We use bold letters to denote vectors. A bandit instance $\boldsymbol{\mu}$ consists of k unknown distributions or arms $\boldsymbol{\mu} = (\mu_1, \dots, \mu_k)$ with respective expectations $\boldsymbol{\theta} = (\theta_1, \dots, \theta_k)$. For the ease of exposition, we assume the bandit instance $\boldsymbol{\mu}$ has a unique best arm. Denote it by $I^*(\boldsymbol{\mu}) \triangleq \arg \max_{i \in [k]} \theta_i$ where $[k] \triangleq \{1, \dots, k\}$. The bandit instance $\boldsymbol{\mu}$ is *unknown* to an experimenter who wants to confidently identify the best arm $I^*(\boldsymbol{\mu})$ at the end of the experiment. At each time $t = 1, 2, \dots$, according to the information collected so far, she can choose an arm $I_t \in [k]$ to measure and then observes an independent noisy observation Y_{t, I_t} drawn from distribution μ_{I_t} .

2. Fixed-confidence setting and its known results

In the fixed-confidence setting, the experimenter can stop gathering samples at any time and returns an estimate of the identity of the best arm after that. The experimenter’s algorithm is then composed of three rules: a sampling rule that determines which arm to sample at each time, a stopping rule that decides whether to stop at each time, and a decision rule that at the stopping time τ , returns an estimate \hat{I}_τ of the identity of the best arm based on the first τ observations.

Let \mathcal{S} be the class of bandit instances with a unique best arm. [Garivier and Kaufmann \(2016\)](#) studies algorithms that guarantee a *uniformly* small probability of incorrect decision (at the stopping time) below a pre-specified level $\delta > 0$, in the sense that

$$\forall \boldsymbol{\mu} \in \mathcal{S}, \quad \mathbb{P}_{\boldsymbol{\mu}} \left(\hat{I}_{\tau_\delta} \neq I^*(\boldsymbol{\mu}) \right) \leq \delta \quad (1)$$

where τ_δ is an *almost surely finite* stopping time. The notation $\mathbb{P}_{\boldsymbol{\mu}}(\cdot)$ indicates that we are evaluating the probability of events when the observations from chosen arms are drawn under the bandit instance $\boldsymbol{\mu}$. In the learning theory literature, such algorithms are called δ -*Probably-Approximately-Correct* or δ -*PAC*. Among such algorithms, we would like to minimize the expected number of collected samples denoted by $\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]$. [Garivier and Kaufmann \(2016\)](#) shows that for any δ -PAC algorithm,

$$\forall \boldsymbol{\mu} \in \mathcal{S}, \quad \liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]}{\log(1/\delta)} \geq \Gamma_{\text{fc}}^*(\boldsymbol{\mu}) \quad (2)$$

where

$$\Gamma_{\text{fc}}^*(\boldsymbol{\mu}) = \left(\sup_{\boldsymbol{w} \in \Sigma_k} \inf_{\boldsymbol{\nu} \in \text{Alt}(\boldsymbol{\mu})} \sum_{i=1}^k w_i \text{KL}(\mu_i \parallel \nu_i) \right)^{-1}. \quad (3)$$

Here Σ_k is the probability simplex of dimension $k - 1$; $\text{Alt}(\boldsymbol{\mu}) \triangleq \{\boldsymbol{\nu} \in \mathcal{S} : I^*(\boldsymbol{\nu}) \neq I^*(\boldsymbol{\mu})\}$ is the set of bandit instances whose unique best arm is different from $\boldsymbol{\mu}$ ’s unique best arm; $\text{KL}(p \parallel q)$ denotes the Kullback-Leibler (KL) divergence between distributions p and q . The subscript *fc* in Γ_{fc}^* is the acronym of “fixed-confidence”. Besides the information-theoretic lower bound in Equation (2), [Garivier and Kaufmann \(2016\)](#) also proposes the so-called Track-and-Stop algorithms that are δ -PAC and can guarantee

$$\forall \boldsymbol{\mu} \in \mathcal{S}, \quad \limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]}{\log(1/\delta)} \leq \Gamma_{\text{fc}}^*(\boldsymbol{\mu}). \quad (4)$$

Since the lower and upper bounds in Equations (2) and (4) are the same, the function $\Gamma_{\text{fc}}^* : \mathcal{S} \rightarrow \mathbb{R}$ characterizes the asymptotic complexity in the fixed-confidence setting.

3. Fixed-budget setting and its open problems

In the fixed-budget setting, a budget of n samples is fixed and given. After collecting n samples, the experimenter needs to decide an estimate of the identity of the best arm denoted by \hat{I}_n . An algorithm is then only consists of a sampling rule and a decision rule. The experimenter’s objective in the fixed-budget setting is to minimize the probability of incorrect decision defined as

$$p_{\boldsymbol{\mu}, n} \triangleq \mathbb{P}_{\boldsymbol{\mu}} \left(\hat{I}_n \neq I^*(\boldsymbol{\mu}) \right).$$

This setting seems “dual” to the fixed-confidence setting in the sense that instead of minimizing the number of samples subject to a uniformly small probability of incorrect decision, here we minimize the probability of incorrect decision subject to a fixed budget of samples. However, little is known about the asymptotic complexity in the fixed-budget setting.

Open problem 1. The first and foremost open problem is whether there are a desirable algorithm class \mathcal{A} and a well-defined function $\Gamma_{\text{fb}}^* : \mathcal{S} \rightarrow \mathbb{R}$ such that for any algorithm in \mathcal{A} ,

$$\forall \mu \in \mathcal{S}, \quad \liminf_{n \rightarrow \infty} \frac{n}{\log(1/p_{\mu,n})} \geq \Gamma_{\text{fb}}^*(\mu)$$

and there is an algorithm that belongs to \mathcal{A} and guarantees

$$\forall \mu \in \mathcal{S}, \quad \limsup_{n \rightarrow \infty} \frac{n}{\log(1/p_{\mu,n})} \leq \Gamma_{\text{fb}}^*(\mu).$$

Here the subscript fb in Γ_{fb}^* is the acronym of “fixed-budget”.

Discussion on potential algorithm class. Kaufmann et al. (2016) studies the so-called *consistent* algorithms such that for any $\mu \in \mathcal{S}$, the probability of incorrect decision $p_{\mu,n}$ goes to zero when n increases to infinity. The class of consistent algorithms is relatively large, and we believe it might not be the right algorithm class for characterizing the asymptotic complexity in the fixed-budget setting. Note that in the fixed-confidence setting, the class of δ -PAC algorithms defined in Equation (1) is restrictive in the sense that it requires a uniformly small probability of incorrect decision (at the stopping time) for any bandit instance $\mu \in \mathcal{S}$. This restriction helps the analysis of the asymptotic complexity in the fixed-confidence setting. We believe it is necessary to come up with a natural but more restrictive algorithm class in the fixed-budget setting. For example, besides the convergence of the probability of incorrect decision to zero, we may also need to control the convergence rate of the algorithms in the class. One potential algorithm class contains all the algorithms that perform uniformly no worse than uniform sampling, i.e., for any algorithm in this class, it achieves a lower or the same value of $\limsup_{n \rightarrow \infty} \frac{n}{\log(1/p_{\mu,n})}$ for any bandit instance $\mu \in \mathcal{S}$. This leads to the following open problem.

Open problem 2. This open problem is whether there is an algorithm other than uniform sampling itself that performs uniformly no worse than uniform sampling in the fixed-budget setting.

Indeed in the fixed-confidence setting, one can show that for any bandit instance μ , $\Gamma_{\text{fc}}^*(\mu)$ is less than or equal to the value of $\limsup_{n \rightarrow \infty} \frac{\mathbb{E}_{\mu}[\tau_{\delta}]}{\log(1/\delta)}$ under uniform sampling. This implies those asymptotically optimal algorithms in the fixed-confidence setting perform uniformly no worse than uniform sampling. We tend to believe that those algorithms also have advantages over uniform sampling in the fixed-budget setting, but the answer to this open problem is unclear.

4. Conjectures

In this section, we state two existing conjectures in the literature. Unfortunately, neither of them is correct in general.

Conjecture 1. Since the fixed-budget and fixed-confidence settings are “dual” to each other, one conjecture is that $\Gamma_{\text{fb}}^* = \Gamma_{\text{fc}}^*$.

Conjecture 2. Another conjecture is that $\Gamma_{\text{fb}}^* = \Gamma_{\text{na}}^*$ where Γ_{na}^* defined later is the asymptotic complexity in a *non-adaptive* version of the fixed-budget setting studied in [Glynn and Juneja \(2004\)](#) (and the subscript na in Γ_{na}^* is the acronym of “non-adaptive”). They consider sampling rules that fix the probability vector \mathbf{w} of selecting k arms in each time and thus do not adapt to the observations from sequentially selected arms. They show that for any bandit instance $\boldsymbol{\mu} \in \mathcal{S}$,

$$\forall \mathbf{w} \in \Sigma_k, \quad \liminf_{n \rightarrow \infty} \frac{n}{\log(1/p_{\boldsymbol{\mu},n})} \geq \Gamma_{\text{na}}^*(\boldsymbol{\mu})$$

and

$$\exists \mathbf{w}^*(\boldsymbol{\mu}) \in \Sigma_k, \quad \limsup_{n \rightarrow \infty} \frac{n}{\log(1/p_{\boldsymbol{\mu},n})} \leq \Gamma_{\text{na}}^*(\boldsymbol{\mu})$$

where

$$\Gamma_{\text{na}}^*(\boldsymbol{\mu}) = \left(\sup_{\mathbf{w} \in \Sigma_k} \inf_{\nu \in \text{Alt}(\boldsymbol{\mu})} \sum_{i=1}^k w_i \text{KL}(\nu_i \parallel \mu_i) \right)^{-1}. \quad (5)$$

At the first glance, the complexity term $\Gamma_{\text{na}}^*(\boldsymbol{\mu})$ in Equation (5) looks the same as $\Gamma_{\text{fc}}^*(\boldsymbol{\mu})$ in Equation (3). Indeed they are different since KL divergence is not symmetrical in general, but for Gaussian distributions, $\Gamma_{\text{na}}^*(\boldsymbol{\mu}) = \Gamma_{\text{fc}}^*(\boldsymbol{\mu})$. Note that the optimal sampling vector $\mathbf{w}^*(\boldsymbol{\mu})$ depends on the knowledge of unknown bandit instance $\boldsymbol{\mu}$, so it is unknown a priori. Hence, the sampling rule that always fixes the optimal sampling vector $\mathbf{w}^*(\boldsymbol{\mu})$ for each bandit instance $\boldsymbol{\mu}$ is not a valid choice for the adaptive fixed-budget setting of our interest.

Neither conjecture is correct. The results in [Ariu et al. \(2021\)](#) imply neither conjecture is correct in general for Bernoulli bandits. Inspired by the construction in [Carpentier and Locatelli \(2016\)](#), [Ariu et al. \(2021\)](#) constructs a set of bandit instances with large number of arms and shows that neither conjecture can hold for all the instances. We believe that one can also show similar negative results for Gaussian bandits.

5. Known results for two-armed bandits

Though neither conjecture is correct in general, [Kaufmann et al. \(2016\)](#) shows that both conjectures hold for two-armed Gaussian bandits with known variances, i.e., $\Gamma_{\text{fb}}^*(\boldsymbol{\mu}) = \Gamma_{\text{fc}}^*(\boldsymbol{\mu}) = \Gamma_{\text{na}}^*(\boldsymbol{\mu})$ for any such bandit instance $\boldsymbol{\mu}$. It further proves that the optimal sampling rule is non-adaptive, which fixes the sampling vector $(\frac{\sigma_1}{\sigma_1 + \sigma_2}, \frac{\sigma_2}{\sigma_1 + \sigma_2})$ where σ_1 and σ_2 are the known variances of the two arms. Recently, [Kato et al. \(2022\)](#) shows that when the gap between the unknown means of the two arms goes to zero, even the variances are also unknown, the upper bound of the proposed algorithm matches the instance-dependent lower bound in [Kaufmann et al. \(2016\)](#). [Adusumilli \(2022\)](#) studies the diffusion regime of two-armed Gaussian bandits and proves that the same sampling vector $(\frac{\sigma_1}{\sigma_1 + \sigma_2}, \frac{\sigma_2}{\sigma_1 + \sigma_2})$ is also minimax optimal. However, for two-armed Bernoulli bandits, [Kaufmann et al. \(2016\)](#) shows that though the optimal sampling vector exists, it requires the knowledge of unknown means of the arms, which is unknown a priori. It is unclear whether there is an algorithm can achieve the asymptotic optimality without such a requirement.

Acknowledgments

We thank Kaito Ariu, Rémy Degenne, Sandeep Juneja, Masahiro Kato, Junpei Komiyama, Wouter M. Koolen, Pierre Ménard, Daniel Russo and Assaf Zeevi for fruitful discussions.

References

- Karun Adusumilli. Minimax policies for best arm identification with two arms, 2022. URL <https://arxiv.org/abs/2204.05527>.
- Kaito Ariu, Masahiro Kato, Junpei Komiyama, Kenichiro McAlinn, and Chao Qin. Policy choice and best arm identification: Asymptotic analysis of exploration sampling, 2021. URL <https://arxiv.org/abs/2109.08229>.
- Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In *23rd Annual Conference on Learning Theory*, pages 41–53, 2010.
- Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *International conference on Algorithmic learning theory*, pages 23–37. Springer, 2009.
- Alexandra Carpentier and Andrea Locatelli. Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *29th Annual Conference on Learning Theory*, pages 590–604, 2016.
- Herman Chernoff. Sequential design of experiments. *Annals of Mathematical Statistics*, 30(3): 755–770, 1959.
- Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *29th Annual Conference on Learning Theory*, pages 998–1027, 2016.
- Peter Glynn and Sandeep Juneja. A large deviations perspective on ordinal optimization. In *Simulation Conference, 2004. Proceedings of the 2004 Winter*, volume 1. IEEE, 2004.
- Masahiro Kato, Kaito Ariu, Masaaki Imaizumi, Masahiro Nomura, and Chao Qin. Best arm identification with a fixed budget under a small gap, 2022. URL <https://arxiv.org/abs/2201.04469>.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.