



Good arm identification via bandit feedback

Hideaki Kano^{1,2} · Junya Honda^{1,2} · Kentaro Sakamaki¹ · Kentaro Matsuura³ · Atsuyoshi Nakamura⁴ · Masashi Sugiyama^{1,2}

Received: 19 April 2018 / Accepted: 9 January 2019 / Published online: 6 March 2019
© The Author(s) 2019

Abstract

We consider a novel stochastic multi-armed bandit problem called *good arm identification* (GAI), where a good arm is defined as an arm with expected reward greater than or equal to a given threshold. GAI is a pure-exploration problem in which a single agent repeats a process of outputting an arm as soon as it is identified as a good one before confirming the other arms are actually not good. The objective of GAI is to minimize the number of samples for each process. We find that GAI faces a new kind of dilemma, the *exploration-exploitation dilemma of confidence*, which is different from the best arm identification. As a result, an efficient design of algorithms for GAI is quite different from that for the best arm identification. We derive a lower bound on the sample complexity of GAI that is tight up to the logarithmic factor $O(\log \frac{1}{\delta})$ for acceptance error rate δ . We also develop an algorithm whose sample complexity almost matches the lower bound. We also confirm experimentally that our proposed algorithm outperforms naive algorithms in synthetic settings based on a conventional bandit problem and clinical trial researches for rheumatoid arthritis.

Keywords Thresholding bandits · Multi-armed bandits · Reinforcement learning · Machine learning

1 Introduction

The stochastic multi-armed bandit (MAB) problem is one of the most fundamental problems for sequential decision-making under uncertainty (Sutton and Barto 1998). It is regarded as a subfield of reinforcement learning in which an agent aims to acquire a policy to select the best-rewarding action via trial and error. In the stochastic MAB problem, a single agent repeatedly plays K slot machines called *arms*, where an arm generates a stochastic reward when pulled. At each round t , the agent pulls arm $i \in [K] = \{1, 2, \dots, K\}$ and then observes an i.i.d. reward $X_i(t)$ from distribution ν_i with expectation $\mu_i \in [0, 1]$.

Editor: Yung-Kyun Noh.

✉ Hideaki Kano
kanohideaki@outlook.jp

Extended author information available on the last page of the article

One of the most classic MAB formulations is the *cumulative regret minimization* (Lai and Robbins 1985; Auer et al. 2002), where the agent tries to maximize the cumulative reward over the fixed number of trials. In this setting, the agent faces the *exploration-exploitation dilemma of reward*, where the exploration means that the agent pulls seemingly suboptimal arms to discover the arm whose expected reward is largest, and the exploitation indicates that the agent pulls the currently best arm to increase the cumulative reward. The related frameworks can be widely applied to various real-world problems such as clinical trials (Grieve and Krams 2005; Genovese et al. 2013; Choy et al. 2013; Curtis et al. 2015; Liu et al. 2017) and personalized recommendations (Tang et al. 2015).

Another classic branch of the MAB problem is the *best arm identification* (Kaufmann et al. 2016; Kalyanakrishnan et al. 2012), which is a pure-exploration problem that the agent tries to identify the best arm $a^* = \arg \max_{i \in \{1, 2, \dots, K\}} \mu_i$. So far, the conceptual idea of the best arm identification has also been successfully applied to many kinds of real-world problems (Koenig and Law 1985; Schmidt et al. 2006; Zhou et al. 2014; Jun et al. 2016). Recently, the *thresholding bandit problem* was proposed (Locatelli et al. 2016) as a variant of pure-exploration MAB formulations. In the thresholding bandit problem, the agent tries to correctly partition all the K arms into good arms and bad arms, where a good arm is defined as an arm whose expected reward is greater than or equal to a given threshold $\xi > 0$, and a bad arm is defined as an arm whose expected reward is lower than the threshold ξ . However, in practice, neither correctly partitioning all the K arms nor exactly identifying the very best arm is always needed; rather, finding some of reasonably good arms as fast as possible is often more useful.

Take a problem of personalized recommendations for example. The objective is to increase our profit by sending direct emails recommending personalized items. In this problem, timely recommendation is a key, because the best sellers in the past are not necessarily the best sellers in the future. Now, there arise three troubles if this problem is formulated as the best arm identification or the thresholding bandit problem. First, an inflation of exploration costs could break out when the purchase probabilities of the multiple best sellers are much close with each other. Although this trouble can be partly relaxed by the ϵ -best arm identification (Even-Dar et al. 2006), in which an arm with expectation greater than or equal to $\max_{i \in [K]} \mu_i - \epsilon$ is also acceptable, the tolerance parameter ϵ has to be set very conservatively. Second, recommending even the best sellers is not a good idea if the “best” purchase probability is too small considering the advertising costs. Third, it needlessly increases exploration costs to partition all items into good (or profitable) items and bad (or not profitable) items, if it is enough to find only some good items to increase our profit. For the above reasons, the formulation of the personalized recommendation problem as the best arm identification or the thresholding bandit problem is not necessarily effective.

Similar troubles also occur in clinical trials for finding drugs (Kim et al. 2011) or for finding appropriate doses of a drug (Grieve and Krams 2005; Genovese et al. 2013; Choy et al. 2013; Curtis et al. 2015; Liu et al. 2017), where the number of patients is extremely limited. In such a case, it is vitally important to find some drugs or doses with satisfactory effect as fast as possible rather than either to classify all drugs or doses into satisfactory ones and others or to identify the exactly best ones.

In this paper, we propose a new bandit framework named *good arm identification* (GAI), where a good arm is defined as an arm whose expected reward is greater than or equal to a given threshold. We formulate GAI as a pure-exploration problem in the *fixed confidence* setting, which is often considered in conventional pure-exploration problems. In the fixed confidence setting, an acceptance error rate δ is fixed in advance, and we minimize the number of pulling arms needed to assure the correctness of the output with probability greater than

or equal to $1 - \delta$. In GAI, a single agent repeats a process of outputting an arm as soon as the agent identifies it as a good one with error probability at most δ . If it is found that there remain no good arms, then the agent stops working. Although the agent does not face the exploration-exploitation dilemma of reward since GAI is a pure-exploration problem, the agent suffers from a new kind of dilemma, that is the *exploration-exploitation dilemma of confidence*, where the exploration means that the agent pulls other arms than the currently best one in order to discover the arm that the agent can identify as a good one with the least arm-pulls, and the exploitation indicates that the agent pulls the currently best arm to increase the confidence on the goodness.

To address the dilemma of confidence, we propose a Hybrid algorithm for the Dilemma of Confidence (HDoC). The sampling strategy of HDoC is based on the upper confidence bound (UCB) algorithm for the cumulative regret minimization (Auer et al. 2002), and the identification rule (that is, the criterion to output an arm as a good one) of HDoC is based on the lower confidence bound (LCB) for the best arm identification (Kalyanakrishnan et al. 2012). In addition, we show that a lower bound on the sample complexity for GAI is $\Omega(\lambda \log \frac{1}{\delta})$, and HDoC can find λ good arms within $O(\lambda \log \frac{1}{\delta} + (K - \lambda) \log \log \frac{1}{\delta})$ samples. This result suggests that HDoC is superior to naive algorithms based on conventional pure-exploration problems, because they require $O(K \log \frac{1}{\delta})$ samples.

For the personalized recommendation problem, the GAI approach is more appropriate, because the agent can quickly identify good items since the agent only focuses on finding good items rather than identifying the best item (as in the best arm identification) and bad items (as in the thresholding bandit). Certainly, there exists a possibility that the recommended item does not possess the best purchase probabilities. However, that does not necessarily matter when customers' interests and item repositories undergo frequent changes, because identifying the exactly best item requires too many samples, and thus we cannot do that in practice. In addition, thanks to the absolute comparison, not the relative comparison, the inflation of exploration costs does not break out even if the purchase probabilities are close to each other, and then the agent can refrain from recommending items when the purchase probabilities are too small.

Our contributions can be summarized as four folds. First, we formulate a novel pure-exploration problem called GAI and find there is a new kind of dilemma, that is, the exploration-exploitation dilemma of confidence. Second, we derive a lower bound for GAI in the *fixed confidence* setting. Third, we propose the HDoC algorithm and show that an upper bound on the sample complexity of HDoC almost matches the lower bound. Fourth, we experimentally demonstrate that HDoC outperforms two naive algorithms derived from other pure exploration problems in synthetic settings based on the thresholding bandit problem (Locatelli et al. 2016) and the clinical trial researches for rheumatoid arthritis (Genovese et al. 2013; Choy et al. 2013; Curtis et al. 2015; Liu et al. 2017).

2 Good arm identification

In this section, we first formulate GAI as a pure exploration problem in the fixed confidence setting. Next, we derive a lower bound on the sample complexity for GAI. We give the notation list in Table 1.

Table 1 Notation list

K	Number of arms
m	Number of good arms (unknown)
\hat{m}	Number of arms that the agent outputs before outputting \perp (NULL)
δ	Acceptance error rate
ξ	Threshold determining whether arms are good or not
$a(t)$	Pulled arm at round t
μ_i	Expected reward of arm i (unknown)
$\hat{\mu}_i(t)$	Empirical mean of the rewards of arm i by the end of round t
$\hat{\mu}_{i,n}$	Empirical mean of the rewards when arm i has been pulled n times
$N_i(t)$	Number of samples of arm i which has been pulled by the end of round t
τ_λ	Round that agent identifies λ good arms
τ_{stop}	Round that agent outputs \perp (NULL)
	$\tilde{\mu}_i(t) = \hat{\mu}_i(t) + \sqrt{\frac{\log t}{2N_i(t)}}$ $\bar{\mu}_i(t) = \hat{\mu}_i(t) + \sqrt{\frac{\log(4KN_i^2(t)/\delta)}{2N_i^2(t)}}$ $\underline{\mu}_i(t) = \hat{\mu}_i(t) - \sqrt{\frac{\log(4KN_i^2(t)/\delta)}{2N_i^2(t)}}$ $\Delta_i = \mu_i - \xi $ $\Delta_{i,j} = \mu_i - \mu_j$ $\Delta = \min \left\{ \min_{i \in [K]} \Delta_i, \min_{\lambda \in [K-1]} \Delta_{\lambda, \lambda+1/2} \right\}$ $n_i = \frac{1}{(\Delta_i - \epsilon)^2} \log \left(\frac{4\sqrt{K}/\delta}{(\Delta_i - \epsilon)^2} \log \frac{5\sqrt{K}/\delta}{(\Delta_i - \epsilon)^2} \right)$

2.1 Problem formulation

Let K be the number of arms, $\xi \in (0, 1)$ be a threshold and $\delta > 0$ be an acceptance error rate. Each arm $i \in [K] = \{1, 2, \dots, K\}$ is associated with Bernoulli distribution v_i with mean μ_i . The parameters $\{\mu_i\}_{i=1}^K$ are unknown to the agent. We define a good arm as an arm whose expected reward is greater than or equal to threshold ξ . The number of good arms is denoted by m which is unknown to the agent and, without loss of generality, we assume an indexing of the arms such that

$$\mu_1 \geq \mu_2 \geq \dots \geq \mu_m \geq \xi \geq \mu_{m+1} \geq \dots \geq \mu_K.$$

The agent is naturally unaware of this indexing. At each round t , the agent pulls an arm $a(t) \in [K]$ and receives an i.i.d. reward drawn from distribution $v_{a(t)}$. The agent outputs an arm when it is identified as a good one. The agent repeats this process until there remain no good arms, where the stopping time is denoted by τ_{stop} . To be more precise, the agent outputs $\hat{a}_1, \hat{a}_2, \dots, \hat{a}_{\hat{m}}$ as good arms (which are different from each other) at rounds $\tau_1, \tau_2, \dots, \tau_{\hat{m}}$, respectively, where \hat{m} is the number of arms that the agent outputs as good ones. The agent stops working after outputting \perp (NULL) at round τ_{stop} when the agent finds that there remain no good arms. If all arms are identified as good ones, then the agent stops after outputting \hat{a}_K and \perp together at the same round. For $\lambda > \hat{m}$ we define $\tau_\lambda = \tau_{\text{stop}}$. Now, we introduce the definitions of (λ, δ) -PAC (Probably Approximately Correct) and δ -PAC.

Definition 1 ((λ, δ) -PAC) An algorithm satisfying the following conditions is called (λ, δ) -PAC: if there are at least λ good arms then $\mathbb{P}[\{\hat{m} < \lambda\} \cup \bigcup_{i \in \{\hat{a}_1, \hat{a}_2, \dots, \hat{a}_\lambda\}} \{\mu_i < \xi\}] \leq \delta$ and if there are less than λ good arms then $\mathbb{P}[\hat{m} \geq \lambda] \leq \delta$.

Definition 2 (δ -PAC) An algorithm is called δ -PAC if the algorithm is (λ, δ) -PAC for all $\lambda \in [K]$.

The agent aims to minimize $\{\tau_1, \tau_2, \dots, \tau_{\text{stop}}\}$ simultaneously by a δ -PAC algorithm. On the other hand, the minimization of τ_{stop} corresponds to the thresholding bandit if we consider the fixed confidence setting.

As we can easily see from these definitions, the condition for a (λ, δ) -PAC algorithm is weaker than that for a δ -PAC algorithm. Thus, there is a possibility that we can construct a good algorithm to minimize τ_λ by using a (λ, δ) -PAC algorithm rather than a δ -PAC algorithm if a specific value of λ is considered. Nevertheless, we will show that the lower bound on τ_λ for (λ, δ) -PAC algorithms can be achieved by a δ -PAC algorithm without knowledge of λ .

2.2 Lower bound on the sample complexity

We give a lower bound on the sample complexity for GAI. This proof is given in Sect. 5.

Theorem 1 Under any (λ, δ) -PAC algorithm, if there are $m \geq \lambda$ good arms, then

$$\mathbb{E}[\tau_\lambda] \geq \left(\sum_{i=1}^{\lambda} \frac{1}{d(\mu_i, \xi)} \log \frac{1}{2\delta} \right) - \frac{m}{d(\mu_\lambda, \xi)}, \quad (1)$$

where $d(x, y) = x \log(x/y) + (1-x) \log((1-x)/(1-y))$ is the binary relative entropy, with convention that $d(0, 0) = d(1, 1) = 0$.

This lower bound on the sample complexity for GAI is given in terms of top- λ expectations $\{\mu_i\}_{i=1}^{\lambda}$. In the next section we confirm that this lower bound is tight up to the logarithmic factor $O(\log \frac{1}{\delta})$.

3 Algorithms

In this section, we first consider naive algorithms based on other pure-exploration problems. Next, we propose an algorithm for GAI and bound its sample complexity from above. Pseudo codes of all the algorithms are described in Algorithm 1. These algorithms can be decomposed into two components: a sampling strategy and an identification criterion. A sampling strategy is a policy to decide which arm the agent pulls. An identification criterion is a policy for the agent to decide whether arms are good or bad. All the algorithms adopt the same identification criterion of Lines 5–11 in Algorithm 1, which is based on the Lower Confidence Bound (LCB) for the best arm identification Kalyanakrishnan et al. (2012). See Remark 3 at the end of Sect. 3.2 for other choices of identification criteria.

3.1 Naive algorithms

We consider two naive algorithms: the Lower and Upper Confidence Bounds algorithm for GAI (LUCB-G), which is based on the LUCB algorithm for the best arm identification

(Kalyanakrishnan et al. 2012) and the Anytime Parameter-free Thresholding algorithm for GAI (APT-G), which is based on the APT algorithm for the thresholding bandit problem (Locatelli et al. 2016). In both algorithms, the sampling strategy is the same as the original algorithms. These algorithms sample all arms at the same order $O(\log \frac{1}{\delta})$.

3.2 Proposed algorithm

We propose a Hybrid algorithm for the Dilemma of Confidence (HDoC). The sampling strategy of HDoC is based on the UCB score of the cumulative regret minimization (Auer et al. 2002). As we will see later, the algorithm stops within $t = O(\log \frac{1}{\delta})$ rounds with high probability. Thus, the second term of the UCB score of HDoC in (2) is $O(\sqrt{\frac{\log \log(1/\delta)}{N_i(t)}}$, whereas that of LUCB-G in (3) is $O(\sqrt{\frac{\log(1/\delta)}{N_i(t)}}$. Therefore, the HDoC algorithm pulls the currently best arm more frequently than LUCB-G, which means that HDoC puts more emphasis on exploitation than exploration.

Algorithm 1 HDoC / LUCB-G / APT-G

- 1: **Input:** a threshold ξ , an acceptance error rate δ and a set of arms $\mathcal{A} \leftarrow [K]$.
- 2: Pull each arm once.
- 3: **repeat**
- 4: HDoC: Pull arm $\hat{a}^* = \arg \max_{i \in \mathcal{A}} \tilde{\mu}_i(t)$ for

$$\tilde{\mu}_i(t) = \hat{\mu}_i(t) + \sqrt{\frac{\log t}{2N_i(t)}}. \tag{2}$$

LUCB-G: Pull arm $\hat{a}^* = \arg \max_{i \in \mathcal{A}} \bar{\mu}_i(t)$ for

$$\bar{\mu}_i(t) = \hat{\mu}_i(t) + \sqrt{\frac{\log(4KN_i^2(t)/\delta)}{2N_i(t)}}. \tag{3}$$

APT-G: Pull arm $\hat{a}^* = \arg \min_{i \in \mathcal{A}} \beta_i(t)$ for

$$\beta_i(t) = \sqrt{N_i(t)} |\xi - \hat{\mu}_i(t)|.$$

- 5: **if** $\mu_{\hat{a}^*}(t) = \hat{\mu}_{\hat{a}^*}(t) - \sqrt{\frac{\log(4KN_{\hat{a}^*}^2(t)/\delta)}{2N_{\hat{a}^*}(t)}} \geq \xi$ **then**
 - 6: Output \hat{a}^* as a good arm.
 - 7: Delete \hat{a}^* from \mathcal{A} .
 - 8: **end if**
 - 9: **if** $\bar{\mu}_{\hat{a}^*} = \hat{\mu}_{\hat{a}^*}(t) + \sqrt{\frac{\log(4KN_{\hat{a}^*}^2(t)/\delta)}{2N_{\hat{a}^*}(t)}} < \xi$ **then**
 - 10: Delete \hat{a}^* from \mathcal{A} .
 - 11: **end if**
 - 12: **until** $\bar{\mu}_i < \xi, \forall i \in \mathcal{A}$.
-

The correctness of the output of the HDoC algorithm can be verified by the following theorem, whose proof is given in Appendix A.

Theorem 2 *The HDoC algorithm is δ -PAC.*

This theorem means that the HDoC algorithm outputs a bad arm with probability at most δ .

Next we give an upper bound on the sample complexity of HDoC. We bound the sample complexity in terms of $\Delta_i = |\mu_i - \xi|$ and $\Delta_{i,j} = \mu_i - \mu_j$.

Theorem 3 *Assume that $\Delta_{\lambda,\lambda+1} > 0$. Then, for any $\lambda \leq m$ and $\epsilon < \min\{\min_{i \in [K]} \Delta_i, \Delta_{\lambda,\lambda+1}/2\}$,*

$$\begin{aligned} \mathbb{E}[\tau_\lambda] &\leq \sum_{i \in [\lambda]} n_i + \sum_{i \in [K] \setminus [\lambda]} \left(\frac{\log(K \max_{j \in [K]} n_j)}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \delta n_i \right) + \frac{K^{2 - \frac{\epsilon^2}{(\min_{i \in [K]} \Delta_i - \epsilon)^2}}}{2\epsilon^2} \\ &\quad + \frac{K \left(5 + \log \frac{1}{2\epsilon^2} \right)}{4\epsilon^2}, \\ \mathbb{E}[\tau_{\text{stop}}] &\leq \sum_{i \in [K]} n_i + \frac{K}{2\epsilon^2}, \end{aligned}$$

where

$$n_i = \frac{1}{(\Delta_i - \epsilon)^2} \log \left(\frac{4\sqrt{K/\delta}}{(\Delta_i - \epsilon)^2} \log \frac{5\sqrt{K/\delta}}{(\Delta_i - \epsilon)^2} \right).$$

We prove this theorem in Appendix B. The following corollary is straightforward from this theorem.

Corollary 1 *Let $\Delta = \min\{\min_{i \in [K]} \Delta_i, \min_{\lambda \in [K-1]} \Delta_{\lambda,\lambda+1}/2\}$. Then, for any $\lambda \leq m$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_\lambda]}{\log(1/\delta)} \leq \sum_{i \in [\lambda]} \frac{1}{2\Delta_i^2}, \tag{4}$$

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_{\text{stop}}]}{\log(1/\delta)} \leq \sum_{i \in [K]} \frac{1}{2\Delta_i^2}, \tag{5}$$

$$\mathbb{E}[\tau_\lambda] = O \left(\frac{\lambda \log \frac{1}{\delta} + (K - \lambda) \log \log \frac{1}{\delta} + K \log \frac{K}{\Delta}}{\Delta^2} \right), \tag{6}$$

$$\mathbb{E}[\tau_{\text{stop}}] = O \left(\frac{K \log(1/\delta) + K \log(K/\Delta)}{\Delta^2} \right). \tag{7}$$

Proof Since

$$\limsup_{\delta \rightarrow 0} \frac{n_i}{\log(1/\delta)} = \frac{1}{2(\Delta_i - \epsilon)^2},$$

we obtain (4) and (5) by letting $\epsilon \downarrow 0$. We obtain (6) and (7) by letting $\epsilon = \Delta/2$ in Theorem 3. □

Note that $d(\mu_i, \xi) \geq 2(\mu_i - \xi)^2 = 2\Delta_i^2$ from Pinsker’s inequality and its coefficient two cannot be improved. Thus we see that the upper bound in (4) in Corollary 1 is almost optimal in view of the lower bound in Theorem 1 for sufficiently small δ . The authors believe that the coefficient $2\Delta_i^2$ can be improved to $d(\mu_i, \xi)$ by the techniques in the KL-UCB (Kullback-Leibler-UCB) algorithm (Cappé et al. 2012) and the Thompson sampling algorithm (Agrawal

and Goyal 2012), although we use the sampling strategy based on the UCB algorithm (Auer et al. 2002) for simplicity of the analysis. Eq. (6) means that the sample complexity of $\mathbb{E}[\tau_\lambda]$ scales with $O(\lambda \log \frac{1}{\delta} + (K - \lambda) \log \log \frac{1}{\delta})$ for moderately small δ , which is contrasted with the sample complexity $O(K \log \frac{1}{\delta})$ for the best arm identification (Kaufmann et al. 2016). Furthermore, we see from (5) and (7) that the HDoC algorithm reproduces the optimal sample complexity for the thresholding bandits (Locatelli et al. 2016).

Remark 1 We can easily extend GAI in a Bernoulli setting to GAI in a Gaussian setting with known variance σ^2 . In the proofs of Theorems 2 and 3, we used the assumption of the Bernoulli reward only in Hoeffding’s inequality expressed as

$$\mathbb{P}[\hat{\mu}_{i,n} \leq \mu_i - \epsilon] \leq e^{-2n\epsilon^2},$$

where $\hat{\mu}_{i,n}$ is the empirical mean of the rewards when arm i has been pulled n times. When each reward follows a Gaussian distribution with variance σ^2 , the distribution of the empirical mean is evaluated as

$$\mathbb{P}[\hat{\mu}_{i,n} \leq \mu_i - \epsilon] \leq e^{-\frac{n\epsilon^2}{2\sigma^2}}$$

by Cramér’s inequality. By this replacement the score of HDoC becomes $\tilde{\mu}_i(t) = \hat{\mu}_i(t) + \sqrt{\frac{2\sigma^2 \log t}{N_i(t)}}$, the score of LUCB-G becomes $\bar{\mu}_i(t) = \hat{\mu}_i(t) + \sqrt{\frac{2\sigma^2 \log(4KN_i^2(t)/\delta)}{N_i(t)}}$ and the score for identifying good arms becomes $\underline{\mu}_i(t) = \hat{\mu}_i(t) - \sqrt{\frac{2\sigma^2 \log(4KN_i^2(t)/\delta)}{N_i(t)}}$ in a Gaussian setting given variance σ^2 , while the score of APT-G in a Gaussian setting is the same as the score of APT-G in a Bernoulli setting.

Remark 2 Theorem 2 and the evaluation of τ_{stop} in Theorem 3 do not depend on the sampling strategy and only use the fact that the identification criterion is given by Lines 5–11 in Algorithm 1. Thus, these results still hold even if we use the LUCB-G and APT-G algorithms.

Remark 3 The evaluation of the error probability is based on the union bound over all rounds $t \in \mathbb{N}$, and the identification criterion in Lines 5–11 in Algorithm 1 is designed for this evaluation. The use of the union bound does not worsen the asymptotic analysis for $\delta \rightarrow 0$ and we use this identification criterion to obtain a simple sample complexity bound. On the other hand, it is known that the empirical performance can be considerably improved by, for example, the bound based on the law of iterated logarithm in Jamieson et al. (2014) that can avoid the union bound. We can also use an identification criterion based on such a bound to improve empirical performance but this does not affect the result of relative comparison since we use the same identification criterion between algorithms with different sampling strategies.

3.3 Gap between lower and upper bounds

As we can see from Theorem 3 and its proof, an arm $i > \lambda$ (that is, an arm other than top- λ ones) is pulled roughly $O(\frac{\log \log(1/\delta)}{\Delta_{\lambda,i}})$ times until HDoC outputs λ good arms. On the other hand, the lower bound in Theorem 1 only considers $O(\log \frac{1}{\delta})$ term and does not depend on arms $i > \lambda$. Therefore, in the case where $(K - \lambda)$ is very large compared to $\frac{1}{\delta}$ (more specifically, in the case of $K - \lambda = \Omega\left(\frac{\log(1/\delta)}{\log \log(1/\delta)}\right)$), there still exists a gap between the lower

bound in (1) and the upper bound in (6). Furthermore, the bound in (6) becomes meaningless when $\Delta_{\lambda, \lambda+1} \approx 0$. In fact, the $O(\log \log \frac{1}{\delta})$ term for small $\Delta_{\lambda, \lambda+1}$ is not negligible in some cases as we will see experimentally in Sect. 4.

To fill this gap, it is necessary to consider the following difference between the cumulative regret minimization and GAI. Let us consider the case of pulling two good arms with the same expected rewards. In the cumulative regret minimization, which of these two arms is pulled makes no difference in the reward and, for example, it suffices for pulling these two arms alternately. On the other hand in GAI, the agent should output one of these good arms as fast as possible; hence, it is desirable to pull one of these equivalent arms with a biased frequency. However, the bias in the numbers of samples between seemingly equivalent arms increases the risk to miss an actually better arm and this dilemma becomes a specific difficulty in GAI. The proposed algorithm, HDoC, is not designed to cope with this difficulty and, improving $O(\log \log \frac{1}{\delta})$ term from this viewpoint is important future work.

4 Numerical experiments

In this section we experimentally compare the performance of HDoC with that of LUCB-G and APT-G. In all experiments, each arm is pulled five times as burn-in and the results are the averages over 1,000 independent runs.

4.1 Threshold settings

We consider three settings named Threshold 1–3, which are based on Experiment 1–2 in Locatelli et al. (2016) and Experiment 4 in Mukherjee et al. (2017).

Threshold 1 (Three group setting) Ten Bernoulli arms with mean $\mu_{1:3} = 0.1$, $\mu_{4:7} = 0.35 + 0.1 \cdot (0 : 3)$ and $\mu_{8:10} = 0.9$, and threshold $\xi = 0.5$, where $(i : j)$ denotes $\{i, i + 1, i + 2, \dots, j - 1, j\}$.

Threshold 2 (Arithmetically progressive setting) Six Bernoulli arms with mean $\mu_{1:6} = 0.1 \cdot (1 : 6)$ and threshold $\xi = 0.35$.

Threshold 3 (Close-to-threshold setting) Ten Bernoulli arms with mean $\mu_{1:3} = 0.55$ and $\mu_{4:10} = 0.45$ and threshold $\xi = 0.5$.

4.2 Medical settings

We also consider two medical settings of dose-finding in clinical trials as GAI. In general, the dose of a drug is quite important. Although high doses are usually more effective than low doses, low doses can be effective than high doses because high doses often cause bad side effects. Therefore, it is desirable to list various doses of a drug with satisfactory effect, which can be formulated as GAI. We considered two instances of the dose-finding problem based on Genovese et al. (2013) and Liu et al. (2017) as Medical 1–2, respectively, specified as follows. In both settings, the threshold ξ corresponds to the satisfactory effect.

Medical 1 (Dose-finding of secukinumab for rheumatoid arthritis with satisfactory effect) Five Bernoulli arms with mean $\mu_1 = 0.36$, $\mu_2 = 0.34$, $\mu_3 = 0.469$, $\mu_4 = 0.465$, $\mu_5 = 0.537$, and threshold $\xi = 0.5$.

Here, $\mu_1, \mu_2, \dots, \mu_5$ represent placebo, secukinumab 25mg, 75mg, 150mg and 300mg, respectively. The expected reward indicates American College of Rheumatology 20% Response (ACR20) at week 16 given in (Genovese et al. 2013, Table 2).

Table 2 Averages and standard deviations of arm-pulls over 1000 independent runs in Threshold 1–3 and Medical 1–2 for $\delta = 0.05$

Thre. 1	τ_1	τ_2	τ_3	τ_4	τ_5	τ_{stop}	
HDoC	114.0 ± 21.8	146.7 ± 22.6	186.8 ± 34.4	778.7 ± 741.1	5629.2 ± 1759.6	10264.1 ± 2121.1	
LUCB-G	134.4 ± 26.6	167.3 ± 27.5	197.0 ± 31.2	798.0 ± 246.5	5702.9 ± 1589.9	10258.2 ± 2054.9	
APT-G	6067.5 ± 1789.3	6282.5 ± 1810.9	6473.4 ± 1820.9	8254.0 ± 1909.5	10161.3 ± 2062.0	10243.0 ± 2062.0	
Thre. 2	τ_1	τ_2	τ_3	τ_{stop}			
HDoC	202.2 ± 106.7	825.8 ± 1048.7	5237.2 ± 1614.8	10001.2 ± 2051.5			
LUCB-G	259.5 ± 120.4	763.7 ± 260.4	5566.6 ± 1575.5	9961.8 ± 1957.7			
APT-G	6891.3 ± 1776.4	7990.7 ± 1839.1	9971.4 ± 1976.5	10048.9 ± 1976.5			
Thre. 3	τ_1	τ_2	τ_3	τ_{stop}			
HDoC	7081.3 ± 2808.4	10955.9 ± 2954.4	18063.0 ± 9252.0	46136.6 ± 4699.4			
LUCB-G	10333.0 ± 3330.0	14183.6 ± 3186.7	17162.8 ± 2997.7	46059.1 ± 4740.4			
APT-G	44326.0 ± 4708.1	45212.9 ± 4727.2	45676.7 ± 4727.8	45852.7 ± 4727.8			
Med. 1	τ_1	τ_{stop}	Med. 2	τ_1	τ_2	τ_3	τ_{stop}
HDoC	10170.1 ± 4276.6	31897.1 ± 5791.0	HDoC	111.8 ± 67.9	267.6 ± 125.1	–	–
LUCB-G	10524.9 ± 3189.7	31827.9 ± 5786.0	LUCB-G	110.7 ± 66.2	262.3 ± 117.9	–	–
APT-G	30847.1 ± 5724.9	31571.7 ± 5829.4	APT-G	–	–	–	–

Bold numbers indicate that the results have statistical significance

Table 3 Averages and standard deviations of arm-pulls over 1000 independent runs in Threshold 1–3 and Medical 1–2 for $\delta = 0.005$

Thre. 1	τ_1	τ_2	τ_3	τ_4	τ_5	τ_{stop}	
HDoC	130.5 ± 28.6	168.3 ± 25.9	209.6 ± 37.1	864.1 ± 518.1	6183.0 ± 1956.0	11464.8 ± 2238.4	
LUCB-G	164.6 ± 31.4	200.4 ± 31.1	235.0 ± 34.0	926.6 ± 257.3	6357.5 ± 1696.9	11323.9 ± 2202.6	
APT-G	7300.6 ± 1625.6	7471.0 ± 1642.3	7647.7 ± 1661.0	9433.1 ± 1902.8	11344.9 ± 2292.9	11417.3 ± 2294.5	
Thre. 2	τ_1	τ_2	τ_3	τ_{stop}			
HDoC	231.3 ± 122.3	897.4 ± 1012.0	5783.7 ± 1564.5	11197.2 ± 2129.6			
LUCB-G	301.1 ± 127.2	875.4 ± 265.9	6163.5 ± 1612.2	11099.1 ± 2054.0			
APT-G	8070.7 ± 1616.8	9159.3 ± 1726.7	11081.3 ± 2067.4	11173.4 ± 2058.2			
Thre. 3	τ_1	τ_2	τ_3	τ_{stop}			
HDoC	9721.1 ± 3396.3	14019.9 ± 3154.4	18711.2 ± 7128.7	50937.3 ± 5008.2			
LUCB-G	11926.5 ± 3488.7	15943.1 ± 3145.2	18980.1 ± 3072.1	50700.7 ± 4803.3			
APT-G	49632.9 ± 5010.9	50458.2 ± 4925.2	50841.6 ± 4910.3	50989.3 ± 4905.0			
Med. 1	τ_1	τ_{stop}	Med. 2	τ_1	τ_2	τ_3	τ_{stop}
HDoC	11274.7 ± 4844.1	34840.2 ± 5942.5	HDoC	134.5 ± 71.1	310.1 ± 126.4	–	–
LUCB-G	11739.9 ± 3339.7	34595.5 ± 5890.0	LUCB-G	135.5 ± 72.5	315.0 ± 132.6	–	–
APT-G	34516.4 ± 5994.6	35189.0 ± 6055.4	APT-G	–	–	–	–

Bold numbers indicate that the results have statistical significance
 Symbol “–” denotes that the agent does not output arm $\lambda \in [K]$ or \perp (NULL) within 100,000 arm-pulls in almost all the runs

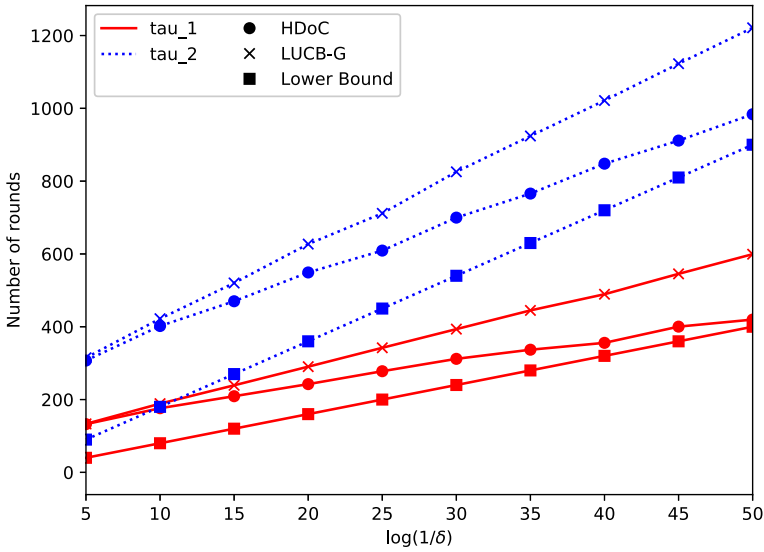


Fig. 1 Number-of-round plots of HDoC, LUCB-G and the lower bound for $\log \frac{1}{\delta} = 5, 10, \dots, 50$ in Medical 2

Medical 2 (Dose-finding of GSK654321 for rheumatoid arthritis with satisfactory effect)
 Seven Gaussian arms with mean $\mu_1 = 0.5, \mu_2 = 0.7, \mu_3 = 1.6, \mu_4 = 1.8, \mu_5 = 1.2, \mu_6 = 1.0$ and $\mu_7 = 0.6$ with variance $\sigma^2 = 1.44$ and threshold $\xi = 1.2$.

Here, $\mu_1, \mu_2, \dots, \mu_7$ represent the positive effect¹ of placebo, the dose of GSK654321 0.03, 0.3, 10, 20 and 30 mg/kg, respectively, where GSK654321 (Liu et al. 2017) is a developing drug with nonlinear dose-response, which is based on the real drug GSK315234 (Choy et al. 2013). The expected reward indicates change from the baseline in Δ Disease Activity Score 28 (DAS28) given in (Liu et al. 2017, Profile 4). The threshold $\xi = 1.2$ is based on Curtis et al. (2015).

4.3 Results

First we compare HDoC, LUCB-G and APT-G for acceptance error rates $\delta = 0.05, 0.005$. Tables 2 and 3 show the averages and standard deviations of $\tau_1, \tau_2, \dots, \tau_\lambda$ and τ_{stop} for these algorithms. In most settings, HDoC outperforms LUCB-G and APT-G. In particular, the number of samples required for APT-G is very large compared to those required for HDoC or LUCB-G, and the stopping times of all algorithms are close as discussed in Remark 2. The results verify that HDoC addresses GAI more efficiently than LUCB-G or APT-G.

In Medical 2, we can easily see $\tau_3, \tau_{\text{stop}} = +\infty$ with high probability since the expected reward μ_5 is equal to the threshold ξ . Moreover, APT-G fails to work completely, since it prefers to pull an arm whose expected reward is closest to the threshold ξ and selects the arm with mean μ_5 almost all the times. In fact, Tables 2, 3 show that APT-G cannot identify even one good arm within 100,000 arm-pulls whereas HDoC and LUCB-G can identify some good arms reasonably even in such a case.

¹ The original values (smaller than zero) in Liu et al. (2017) represent the negative effect and we inverted the sign to denote the positive effect.

As shown in Tables 2, 3, the performance of HDoC is almost the same as that of LUCB-G in Medical 2, where the expectations of the arms are very close to each other, taking the variance σ^2 into consideration. Figure 1 shows the result of an experiment to investigate the behavior of HDoC and LUCB-G for Medical 2 in more detail, where τ_1, τ_2 are plotted for (possibly unrealistically) small δ . Here “Lower bound” in the figure is the asymptotic lower bound $\sum_{i=1}^{\lambda} \frac{2\sigma^2 \log(1/\delta)}{\Delta_i^2}$ of τ_λ for normal distributions (see Theorem 1 and Remark 1). Since the result of HDoC asymptotically approaches to the lower bound, the $O(\log \frac{1}{\delta})$ term of the sample complexity of HDoC is almost optimal, and the results show that the effect of $O(\log \log \frac{1}{\delta})$ term is not negligible for practical acceptance error rates such as $\delta = 0.05$ and 0.005.

5 Proof of Theorem 1

In this section, we prove Theorem 1 based on the following proposition on the expected number of samples to distinguish two sets of reward distributions.

Proposition 1 (Lemma 1 in Kaufmann et al. 2016) *Let v and v' be two bandit models with K arms such that for all i , the distributions v_i and v'_i are mutually absolutely continuous. For any almost-surely finite stopping time σ and event \mathcal{E} ,*

$$\sum_{i=1}^K \mathbb{E}[N_i(\sigma)] \text{KL}(v_i, v'_i) \geq d(\mathbb{P}_v[\mathcal{E}], \mathbb{P}_{v'}[\mathcal{E}]),$$

where $\text{KL}(v_i, v_j)$ is the Kullback-Leibler divergence between distributions v_i and v_j , and $d(x, y) = x \log(x/y) + (1 - x) \log((1 - x)/(1 - y))$ is the binary relative entropy, with convention that $d(0, 0) = d(1, 1) = 0$.

Standard proofs on the best arm identification problems set \mathcal{E} as an event such that $\mathbb{P}[\mathcal{E}] \geq 1 - \delta$ under any δ -PAC algorithm. On the other hand, we leave $\mathbb{P}[\mathcal{E}]$ to range from 0 to 1 and establish a lower bound as a minimization problem over $\mathbb{P}[\mathcal{E}]$.

Proof of Theorem 1 Fix $j \in [m]$ and consider a set of Bernoulli distributions $\{v'_i\}$ with expectations $\{\mu'_i\}$ given by

$$\mu'_i = \begin{cases} \xi - \epsilon, & \text{if } i = j, \\ \mu_i, & \text{if } i \in [K] \setminus \{j\}. \end{cases}$$

Let $\mathcal{E}_j = \{j \in \{\hat{\alpha}_i\}_{i=1}^{\min\{\lambda, \hat{m}\}}\}$ and $p_j = \mathbb{P} \left[j \in \{\hat{\alpha}_i\}_{i=1}^{\min\{\lambda, \hat{m}\}} \right]$ under $\{v_i\}$. Since j is not a good arm under $\{v'_i\}$, we obtain from Proposition 1 that

$$\begin{aligned} \mathbb{E}[N_j]d_j &\geq d(p_j, \min\{\delta, p_j\}) \\ &= \max \left\{ p_j \log \frac{1}{\min\{\delta, p_j\}} - h(p_j) + (1 - p_j) \log \frac{1}{1 - \min\{\delta, p_j\}}, 0 \right\} \\ &\geq \max \left\{ p_j \log \frac{1}{\min\{\delta, p_j\}} - \log 2, 0 \right\} \\ &\geq \max \left\{ p_j \log \frac{1}{\delta} - \log 2, 0 \right\}, \end{aligned}$$

where we set $d_i = d(\mu_i, \xi - \epsilon)$ and $h(p) = -p \log p - (1 - p) \log(1 - p) \leq \log 2$ is the binary entropy function.

Here note that

$$\begin{aligned} \sum_{i=1}^m p_i &= \mathbb{E}_\nu \left[\left| [m] \cap \{\hat{a}_i\}_{i=1}^{\min\{\lambda, \hat{m}\}} \right| \right] \\ &\geq \lambda \mathbb{P}_\nu[\{\{\hat{a}_i\}_{i=1}^{\min\{\lambda, \hat{m}\}} \subset [m]\}, \hat{m} \geq \lambda] \geq \lambda(1 - \delta) \end{aligned}$$

under any (λ, δ) -PAC algorithm. Thus we have

$$\sum_{i=1}^K \mathbb{E}[N_i] \geq \sum_{i=1}^m \mathbb{E}[N_i] \geq C^*,$$

where C^* is the optimal value of the optimization problem

$$\begin{aligned} (P_1) \quad \text{minimize} \quad & \sum_{i=1}^m \frac{1}{d_i} \max \left\{ p_i \log \frac{1}{\delta} - \log 2, 0 \right\}, \quad \text{subject to} \quad \sum_{i=1}^m p_i \geq \lambda(1 - \delta), \\ & 0 \leq p_i \leq 1, \quad \forall i \in [m], \end{aligned}$$

which is equivalent to the linear programming problem

$$\begin{aligned} (P_2) \quad \text{minimize} \quad & \sum_{i=1}^m \frac{x_i}{d_i}, \quad \text{subject to} \quad \sum_{i=1}^m p_i \geq \lambda(1 - \delta), \\ & x_i \geq p_i \log \frac{1}{\delta} - \log 2, \quad \forall i \in [m], \\ & 0 \leq p_i \leq 1, \quad x_i \geq 0, \quad \forall i \in [m]. \end{aligned}$$

The dual problem of (P_2) is given by

$$\begin{aligned} (P'_2) \quad \text{maximize} \quad & \lambda(1 - \delta)\alpha - (\log 2) \sum_{i=1}^m \beta_i - \sum_{i=1}^m \gamma_i \\ \text{subject to} \quad & \beta_i \leq \frac{1}{d_i}, \quad \forall i \in [m], \\ & \alpha - \beta_i \log \frac{1}{\delta} - \gamma_i \leq 0, \quad \forall i \in [m], \\ & \alpha, \beta_i, \gamma_i \geq 0, \quad \forall i \in [m]. \end{aligned}$$

Here consider the feasible solution of (P'_2) given by

$$\alpha = \frac{1}{d_\lambda} \log \frac{1}{\delta}, \quad \beta_i = \begin{cases} \frac{1}{d_i}, & i \leq \lambda, \\ \frac{1}{d_\lambda}, & i > \lambda, \end{cases} \quad \gamma_i = \begin{cases} \left(\frac{1}{d_\lambda} - \frac{1}{d_i} \right) \log \frac{1}{\delta}, & i \leq \lambda, \\ 0, & i > \lambda, \end{cases}$$

which attains the objective function

$$\begin{aligned} & \frac{\lambda(1 - \delta)}{d_\lambda} \log \frac{1}{\delta} - (\log 2) \left(\sum_{i \leq \lambda} \frac{1}{d_i} + \frac{m - \lambda}{d_\lambda} \right) - \sum_{i \leq \lambda} \left(\frac{1}{d_\lambda} - \frac{1}{d_i} \right) \log \frac{1}{\delta} \\ & = \sum_{i \leq \lambda} \left(\frac{1}{d_i} \log \frac{1}{\delta} - \frac{\log 2}{d_i} \right) - \frac{\lambda \delta}{d_\lambda} \log \frac{1}{\delta} - \frac{(m - \lambda) \log 2}{d_\lambda} \end{aligned}$$

$$\begin{aligned} &\geq \sum_{i \leq \lambda} \frac{1}{d_i} \log \frac{1}{2\delta} - \frac{\lambda}{d_\lambda} - \frac{(m - \lambda) \log 2}{d_\lambda} \quad \text{by } \sup_{0 < \delta \leq 1} \delta \log(1/\delta) = 1/e < 1 \\ &\geq \sum_{i \leq \lambda} \frac{1}{d_i} \log \frac{1}{2\delta} - \frac{m}{d_\lambda}. \end{aligned}$$

Since the objective function of a feasible solution for the dual problem (P₂') of (P₂) is always smaller than the optimal value C* of (P₂), we have

$$\begin{aligned} \sum_{i=1}^m \mathbb{E}[N_i] &\geq C^* \\ &\geq \sum_{i \leq \lambda} \frac{1}{d_i} \log \frac{1}{2\delta} - \frac{m}{d_\lambda} \\ &= \sum_{i \leq \lambda} \frac{1}{d(\mu_i, \xi - \epsilon)} \log \frac{1}{2\delta} - \frac{m}{d(\mu_\lambda, \xi - \epsilon)}. \end{aligned}$$

We complete the proof by letting $\epsilon \downarrow 0$. □

6 Conclusion

In this paper, we considered and discussed a new multi-armed bandit problem called good arm identification (GAI). The objective of GAI is to minimize not only the total number of samples to identify all good arms but also the number of samples until identifying λ good arms for each $\lambda = 1, 2, \dots$, where a good arm is an arm whose expected reward is greater than or equal to threshold ξ . Even though GAI, which is a pure-exploration problem, does not face the exploration-exploitation dilemma of reward, GAI encounters a new kind of dilemma: the exploration-exploitation dilemma of confidence. We derived a lower bound on the sample complexity of GAI, developed an efficient algorithm, HDoC, and then we theoretically showed the sample complexity of HDoC almost matches the lower bound. We also experimentally demonstrated that HDoC outperforms algorithms based on other pure-exploration problems in the three settings based on the thresholding bandit and two settings based on the dose-finding problem in the clinical trials.

Acknowledgements JH acknowledges support by KAKENHI 16H00881. MS acknowledges support by KAKENHI 17H00757. This work was partially supported by JST CREST Grant Number JPMJCR1662, Japan.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

A Proof of Theorem 2

In this appendix we prove Theorem 2 based on the following lemma.

Lemma 1

$$\mathbb{P} \left[\bigcup_{n \in \mathbb{N}} \{\bar{\mu}_{i,n} < \xi\} \right] \leq \frac{\delta}{K}, \quad \text{for any } i \in [m],$$

$$\mathbb{P} \left[\bigcup_{n \in \mathbb{N}} \{ \underline{\mu}_{i,n} \geq \xi \} \right] \leq \frac{\delta}{K}, \quad \text{for any } i \in [K] \setminus [m].$$

Proof For any $i \in [m]$,

$$\begin{aligned} \mathbb{P} \left[\bigcup_{n \in \mathbb{N}} \{ \bar{\mu}_{i,n} < \xi \} \right] &\leq \sum_{n \in \mathbb{N}} \mathbb{P} [\bar{\mu}_{i,n} < \xi] \\ &\leq \sum_{n \in \mathbb{N}} \mathbb{P} [\bar{\mu}_{i,n} < \mu_i] && \text{by } \mu_i \geq \xi \text{ for } i \in [m] \\ &\leq \sum_{n \in \mathbb{N}} e^{-2n \left(\sqrt{\frac{\log(4Kn^2/\delta)}{2n}} \right)^2} && \text{by Hoeffding's inequality} \\ &= \sum_{n \in \mathbb{N}} \frac{\delta}{4Kn^2} \\ &= \frac{\pi^2 \delta}{24K} \quad \text{by } \sum_{n \in \mathbb{N}} \frac{1}{n^2} = \frac{\pi^2}{6} \\ &\leq \frac{\delta}{K}. \end{aligned}$$

For any $i \in [K] \setminus [m]$, the same argument holds. □

Proof of Theorem 2 We show that HDoC is (λ, δ) -PAC for arbitrary $\lambda \in [K]$.

First we consider the case that there are more than or equal to λ good arms and show

$$\mathbb{P} \left[\{ \hat{m} < \lambda \} \cup \bigcup_{i \in \{ \hat{a}_1, \hat{a}_2, \dots, \hat{a}_\lambda \}} \{ \mu_i < \xi \} \right] \leq \delta. \tag{8}$$

Since we are now considering the case $m \geq \lambda$, the event $\{ \hat{m} < \lambda \}$ implies that at least one good arm $j \in [m]$ is regarded as a bad arm, that is, $\{ \underline{\mu}_{j,n} \leq \xi \}$ occurs for some $j \in [m]$ and $n \in \mathbb{N}$. Thus we have

$$\begin{aligned} \mathbb{P} [\hat{m} < \lambda] &\leq \sum_{j \in [m]} \mathbb{P} \left[\bigcup_{n \in \mathbb{N}} \{ \bar{\mu}_{j,n} < \xi \} \right] \\ &\leq \sum_{j \in [m]} \frac{\delta}{K} \quad \text{by Lemma 1} \\ &\leq \frac{m\delta}{K}. \end{aligned} \tag{9}$$

On the other hand, since the event $\bigcup_{i \in \{ \hat{a}_1, \hat{a}_2, \dots, \hat{a}_\lambda \}} \{ \mu_i < \xi \}$ implies that $j \in \{ \hat{a}_i \}_{i=1}^\lambda$ for some bad arm $j \in [K] \setminus [m]$, we have

$$\begin{aligned} \mathbb{P} \left[\bigcup_{i \in \{ \hat{a}_1, \hat{a}_2, \dots, \hat{a}_\lambda \}} \{ \mu_i < \xi \} \right] &\leq \sum_{j \in [K] \setminus [m]} \mathbb{P} [j \in \{ \hat{a}_i \}_{i=1}^\lambda] \\ &\leq \sum_{j \in [K] \setminus [m]} \mathbb{P} \left[\bigcup_{n \in \mathbb{N}} \{ \underline{\mu}_{j,n} \geq \xi \} \right] \end{aligned}$$

$$\leq \frac{(K - m)\delta}{K} \tag{10}$$

in the same way as (9). We obtain (8) by putting (9) and (10) together.

Next we consider the case that the number of good arms m is less than λ and show

$$\mathbb{P}[\hat{m} \geq \lambda] \leq \delta. \tag{11}$$

Since there are at most $m < \lambda$ good arms, the event $\{\hat{m} \geq \lambda\}$ implies that $j \in \{\hat{a}_i\}_{i=1}^\lambda$ for some $j \in [K] \setminus [m]$. Thus, in the same way as (10) we have

$$\begin{aligned} \mathbb{P}[\hat{m} \geq \lambda] &\leq \sum_{j \in [K] \setminus [m]} \mathbb{P}[j \in \{\hat{a}_i\}_{i=1}^\lambda] \\ &\leq \frac{(K - m)\delta}{K} \\ &\leq \delta, \end{aligned}$$

which proves (11). □

B Proof of Theorem 3

In this appendix, we prove Theorem 3 based on the following lemmas, and we define $T = K \max_{i \in [K]} \lfloor n_i + 2 \rfloor$.

Lemma 2 *If $n \geq n_i$ then*

$$\mathbb{P}[\mu_{i,n} \leq \xi] \leq e^{-2n\epsilon^2}, \quad \forall i \in [m], \tag{12}$$

$$\mathbb{P}[\bar{\mu}_{i,n} \geq \xi] \leq e^{-2n\epsilon^2}, \quad \forall i \in [K] \setminus [m]. \tag{13}$$

Proof We only show (12) for $i \in [m]$. Eq. (13) for $i \in [K] \setminus [m]$ is exactly the same. From Hoeffding’s inequality it suffices to show that for $n \geq n_i$

$$\sqrt{\frac{\log(4Kn^2/\delta)}{2n}} \leq \mu_i - \xi - \epsilon = \Delta_i - \epsilon.$$

We write $c = (\Delta_i - \epsilon)^2 \leq 1$ in the following for notational simplicity. Then we can express $n \geq n_i$ as

$$n = \frac{1}{c} \log \frac{4t\sqrt{K/\delta}}{c}$$

for some $t > \log \frac{5\sqrt{K/\delta}}{c} > \log(5\sqrt{2}) > 1$. Then

$$\begin{aligned} \sqrt{\frac{\log(4Kn^2/\delta)}{2n}} \leq \Delta_i - \epsilon &\Leftrightarrow \log(4Kn^2/\delta) \leq 2cn \\ &\Leftrightarrow \log \frac{4K \left(\log \frac{4t\sqrt{K/\delta}}{c} \right)^2}{c^2\delta} \leq \log \frac{16t^2K}{c^2\delta} \\ &\Leftrightarrow \log \frac{4t\sqrt{K/\delta}}{c} \leq 2t \end{aligned}$$

$$\begin{aligned} &\Leftrightarrow \log \frac{4\sqrt{K/\delta}}{c} + t - 1 \leq 2t \quad \text{by } \log x \leq x - 1 \\ &\Leftrightarrow \log \frac{4\sqrt{K/\delta}}{ec} \leq t. \end{aligned} \tag{14}$$

We obtain the lemma since $t > \log \frac{5\sqrt{K/\delta}}{c}$ satisfies (14). □

Lemma 3

$$\begin{aligned} \mathbb{E} \left[\sum_{n=1}^{\infty} \mathbb{1}[\underline{\mu}_{i,n} \leq \xi] \right] &\leq n_i + \frac{1}{2\epsilon^2}, \quad \forall i \in [m], \\ \mathbb{E} \left[\sum_{n=1}^{\infty} \mathbb{1}[\bar{\mu}_{i,n} \geq \xi] \right] &\leq n_i + \frac{1}{2\epsilon^2}, \quad \forall i \in [K] \setminus [m]. \end{aligned}$$

Proof If arm $i \in [m]$ then

$$\begin{aligned} \mathbb{E} \left[\sum_{n=1}^{\infty} \mathbb{1}[\underline{\mu}_{i,n} \leq \xi] \right] &\leq \mathbb{E} \left[\sum_{n=1}^{n_i} 1 + \sum_{n=n_i+1}^{\infty} \mathbb{1}[\underline{\mu}_{i,n} \leq \xi] \right] \\ &\leq n_i + \sum_{n=1}^{\infty} \mathbb{P}[\underline{\mu}_{i,n} < \xi] \\ &\leq n_i + \sum_{n=1}^{\infty} e^{-2n\epsilon^2} \\ &\leq n_i + \frac{1}{e^{2\epsilon^2} - 1} \\ &\leq n_i + \frac{1}{2\epsilon^2}. \end{aligned}$$

If arm $i \in [K] \setminus [m]$ then

$$\begin{aligned} \mathbb{E} \left[\sum_{n=1}^{\infty} \mathbb{1}[\bar{\mu}_{i,n} \geq \xi] \right] &\leq \mathbb{E} \left[\sum_{n=1}^{n_i} 1 + \sum_{n=n_i+1}^{\infty} \mathbb{1}[\bar{\mu}_{i,n} \geq \xi] \right] \\ &\leq n_i + \sum_{n=1}^{\infty} \mathbb{P}[\bar{\mu}_{i,n} \geq \xi] \\ &\leq n_i + \sum_{n=1}^{\infty} e^{-2n\epsilon^2} \\ &\leq n_i + \frac{1}{e^{2\epsilon^2} - 1} \\ &\leq n_i + \frac{1}{2\epsilon^2}. \end{aligned}$$

□

Lemma 4

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}[a(t) \in [\lambda]] + \sum_{t=1}^T \mathbb{1}[a(t) \notin [\lambda], \tilde{\mu}^*(t) \geq \mu_{\lambda} - \epsilon] \right]$$

$$\leq \sum_{i \in [\lambda]} n_i + \sum_{i \in [K] \setminus [\lambda]} \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \frac{K}{2\epsilon^2}. \tag{15}$$

Proof For the first term of (15) we have

$$\begin{aligned} \sum_{t=1}^{\infty} \mathbb{1}[a(t) \in [\lambda]] &\leq \sum_{i \in [\lambda]} \sum_{t=1}^{\infty} \mathbb{1}[a(t) = i] \\ &= \sum_{i \in [\lambda]} \sum_{t=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{1}[a(t) = i, N_i(t) = n]. \end{aligned} \tag{16}$$

Since the event $\{a(t) = i, N_i(t) = n\}$ occurs for at most one $t \in \mathbb{N}$ we have

$$\begin{aligned} \sum_{t=1}^{\infty} \sum_{n=1}^{\infty} \mathbb{1}[a(t) = i, N_i(t) = n] &\leq \sum_{n=1}^{\infty} \mathbb{1}\left[\bigcup_{t=1}^{\infty} \{a(t) = i, N_i(t) = n\}\right] \\ &\leq \sum_{n=1}^{\infty} \mathbb{1}[\mu_{i,n} \leq \xi]. \end{aligned} \tag{17}$$

By combining (16) and (17) with Lemma 3 we obtain

$$\mathbb{E}\left[\sum_{t=1}^{\infty} \mathbb{1}[a(t) \in [\lambda]]\right] \leq \sum_{i \in [\lambda]} \left(n_i + \frac{1}{2\epsilon^2}\right). \tag{18}$$

Next we consider the second term of (15). By using the same argument as (16) we obtain for $i \notin [\lambda]$ that

$$\begin{aligned} &\sum_{t=1}^T \mathbb{1}[a(t) = i, \tilde{\mu}^*(t) \geq \mu_\lambda - \epsilon] \\ &\leq \sum_{n=1}^T \mathbb{1}\left[\bigcup_{t=1}^T \{a(t) = i, \tilde{\mu}^*(t) \geq \mu_\lambda - \epsilon, N_i(t) = n\}\right] \\ &\leq \sum_{n=1}^T \mathbb{1}\left[\bigcup_{t=1}^T \{\tilde{\mu}^*(t) = \tilde{\mu}_i(t) \geq \mu_\lambda - \epsilon, N_i(t) = n\}\right] \\ &\leq \sum_{n=1}^T \mathbb{1}\left[\hat{\mu}_{i,n} + \sqrt{\frac{\log T}{2n}} \geq \mu_\lambda - \epsilon\right] \\ &\leq \sum_{n=1}^T \mathbb{1}\left[\hat{\mu}_{i,n} + \sqrt{\frac{\log T}{2n}} \geq \mu_i + \Delta_{\lambda,i} - \epsilon\right] \\ &\leq \sum_{n=1}^{\frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2}} 1 + \sum_{n=\frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + 1}^T \mathbb{1}\left[\hat{\mu}_{i,n} + \sqrt{\frac{\log T}{2 \cdot \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2}}} \geq \mu_i + \Delta_{\lambda,i} - \epsilon\right] \\ &= \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \sum_{n=\frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + 1}^T \mathbb{1}[\hat{\mu}_{i,n} \geq \mu_i + \epsilon]. \end{aligned}$$

By taking the expectation we have

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[a(t) = i, \tilde{\mu}^*(t) \geq \mu_\lambda - \epsilon] \right] \leq \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \sum_{n=1}^{\infty} \mathbb{P}[\hat{\mu}_{i,n} \geq \mu_i + \epsilon]$$

$$\leq \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \sum_{n=1}^{\infty} e^{-2n\epsilon^2} \tag{19}$$

$$= \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \frac{1}{e^{2\epsilon^2} - 1}$$

$$\leq \frac{\log T}{2(\Delta_{\lambda,i} - 2\epsilon)^2} + \frac{1}{2\epsilon^2}, \tag{20}$$

where (19) follows from Hoeffding’s inequality. We complete the proof by combining (18) with (20). \square

Lemma 5

$$\mathbb{E} \left[\sum_{t=T+1}^{\infty} \mathbb{1}[t \leq \tau_\lambda] \right] \leq \frac{K^{2 - \frac{\epsilon^2}{(\min_{i \in [K]} \Delta_i - \epsilon)^2}}}{2\epsilon^2}.$$

Proof Note that at the t -th round some arm is pulled at least $\lceil (t - 1)/K \rceil$ times. Furthermore, $N_i(t) \geq \lceil (t - 1)/K \rceil$ implies that the arm i is still in $\mathcal{A}(t)$ when the arm i is pulled $\lceil (t - 1)/K \rceil - 1$ times. Thus we have

$$\sum_{t=T+1}^{\infty} \mathbb{1}[t \leq \tau_\lambda] \leq \sum_{i=1}^m \sum_{t=T+1}^{\infty} \mathbb{1}[N_i(t) \geq \lceil (t - 1)/K \rceil, t \leq \tau_\lambda]$$

$$\leq \sum_{i=1}^m \sum_{t=T+1}^{\infty} \mathbb{1}[\mu_{i, \lceil (t-1)/K \rceil - 1} \leq \xi] + \sum_{i=m+1}^K \sum_{t=T+1}^{\infty} \mathbb{1}[\bar{\mu}_{i, \lceil (t-1)/K \rceil - 1} \geq \xi]. \tag{21}$$

From the definition of $T = K \max_{i \in [K]} \lfloor n_i + 2 \rfloor$, we have $\lceil (t - 1)/K \rceil - 1 \geq n_i$ for all $i \in [K]$. Thus, the expectation of (21) is bounded by Lemma 2 as

$$\mathbb{E} \left[\sum_{t=T+1}^{\infty} \mathbb{1}[t \leq \tau_\lambda] \right] \leq K \sum_{t=T+1}^{\infty} e^{-2\epsilon^2(\lceil (t-1)/K \rceil - 1)}$$

$$\leq K \sum_{t=T+1}^{\infty} e^{-2\epsilon^2((t-1)/K - 1)}$$

$$= \frac{K e^{-2\epsilon^2((T-1)/K - 1)}}{e^{2\epsilon^2/K} - 1} \text{ by } \sum_{i=1}^{\infty} e^{-ai} = \frac{1}{e^a - 1}, \text{ for } a > 0$$

$$\leq \frac{K e^{-2\epsilon^2(\max_i n_i - 2)}}{e^{2\epsilon^2/K} - 1} \text{ by } T = K \max_{i \in [K]} \lfloor n_i + 2 \rfloor$$

$$\leq \frac{K e^{-\frac{2\epsilon^2}{(\min_{i \in [K]} \Delta_i - \epsilon)^2} (\log \sqrt{K})}}{e^{2\epsilon^2/K} - 1}$$

$$\begin{aligned}
 &= \frac{K^{1 - \frac{\epsilon^2 (\log \sqrt{K})}{(\min_{i \in [K]} \Delta_i - \epsilon)^2}}}{e^{2\epsilon^2/K} - 1} \\
 &\leq \frac{K^{2 - \frac{\epsilon^2}{(\min_{i \in [K]} \Delta_i - \epsilon)^2}}}{2\epsilon^2}, \tag{22}
 \end{aligned}$$

where (22) follows from

$$\begin{aligned}
 n_i &\geq \frac{1}{(\Delta_i - \epsilon)^2} \log \left(4\sqrt{K} \log 5\sqrt{K} \right) \\
 &\geq \frac{\log \sqrt{K}}{(\Delta_i - \epsilon)^2} + \log \left(4 \log 5\sqrt{2} \right) \\
 &\geq \frac{\log \sqrt{K}}{(\Delta_i - \epsilon)^2} + 2.05.
 \end{aligned}$$

□

Lemma 6

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, \tilde{\mu}^*(t) < \mu_\lambda - \epsilon] \right] \leq \frac{3K}{4\epsilon^2} + \frac{K \log \frac{1}{2\epsilon^2}}{4\epsilon^2} + \delta \sum_{i \in [K] \setminus [\lambda]} n_i$$

Proof The summation is decomposed into

$$\begin{aligned}
 &\sum_{t=1}^{\infty} \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, \tilde{\mu}^*(t) < \mu_\lambda - \epsilon] \\
 &\leq \sum_{t=1}^T \mathbb{1}[\tilde{\mu}^*(t) < \mu^* - \epsilon, [\lambda] \cap \mathcal{A}(t) \neq \emptyset] + \sum_{t=1}^T \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, [\lambda] \cap \mathcal{A}(t) = \emptyset], \tag{23}
 \end{aligned}$$

where $\mathcal{A}(t) = \{i \in [K] : \underline{\mu}_i(t) < \xi \leq \bar{\mu}_i(t)\}$. From definition $\tilde{\mu}^*(t) = \max_{i \in \mathcal{A}(t)} \tilde{\mu}_i(t)$ the first term of (23) is evaluated as

$$\begin{aligned}
 &\sum_{t=1}^T \mathbb{1}[\tilde{\mu}^*(t) < \mu_\lambda - \epsilon, [\lambda] \cap \mathcal{A}(t) \neq \emptyset] \\
 &\leq \sum_{i \in [\lambda]} \sum_{t=1}^T \mathbb{1}[\tilde{\mu}^*(t) < \mu_\lambda - \epsilon, i \in \mathcal{A}(t)] \\
 &\leq \sum_{i=1}^{\lambda} \sum_{t=1}^T \mathbb{1}[\tilde{\mu}_i(t) < \mu_\lambda - \epsilon] \\
 &\leq \sum_{i=1}^{\lambda} \sum_{n=1}^T \sum_{t=1}^T \mathbb{1}[\tilde{\mu}_i(t) < \mu_\lambda - \epsilon, N_i(t) = n] \\
 &\leq \sum_{i=1}^{\lambda} \sum_{n=1}^T \sum_{t=1}^T \mathbb{1} \left[\hat{\mu}_i(t) + \sqrt{\frac{\log t}{2n}} < \mu_\lambda - \epsilon, N_i(t) = n \right]
 \end{aligned}$$

$$\begin{aligned} &\leq \sum_{i=1}^{\lambda} \sum_{n=1}^T \sum_{t=1}^T \mathbb{1}[t < e^{2n(\hat{\mu}_{i,n} - \mu_{\lambda} + \epsilon)^2}, \hat{\mu}_{i,n} < \mu_{\lambda} - \epsilon] \\ &\leq \sum_{i=1}^{\lambda} \sum_{n=1}^T e^{2n(\hat{\mu}_{i,n} - \mu_{\lambda} + \epsilon)^2} \mathbb{1}[\hat{\mu}_{i,n} < \mu_{\lambda} - \epsilon]. \end{aligned} \tag{24}$$

Let $P_{i,n}(x) = \mathbb{P}[\hat{\mu}_{i,n} < x]$. Then the expectation of the inner summation of (24) is bounded by

$$\begin{aligned} &\sum_{n=1}^T \mathbb{E}[e^{2n(\hat{\mu}_{i,n} - \mu_{\lambda} + \epsilon)^2} \mathbb{1}[\hat{\mu}_{i,n} < \mu_{\lambda} - \epsilon]] \\ &\leq \sum_{n=1}^T \int_0^{\mu_{\lambda} - \epsilon} e^{2n(x - \mu_{\lambda} + \epsilon)^2} dP_{i,n}(x) \\ &= \sum_{n=1}^T \left([e^{2n(x - \mu_{\lambda} + \epsilon)^2} P_{i,n}(x)]_0^{\mu_{\lambda} - \epsilon} - 4n \int_0^{\mu_{\lambda} - \epsilon} (x - \mu_{\lambda} + \epsilon) e^{2n(x - \mu_{\lambda} + \epsilon)^2} P_{i,n}(x) dx \right) \\ &\hspace{15em} \text{by integration by parts} \\ &\leq \sum_{n=1}^T \left(e^{-2n\epsilon^2} - 4n \int_0^{\mu_{\lambda} - \epsilon} (x - \mu_{\lambda} + \epsilon) e^{2n(x - \mu_{\lambda} + \epsilon)^2} e^{-2n(x - \mu_{\lambda})^2} dx \right) \\ &\hspace{15em} \text{by Hoeffding's inequality} \\ &= \sum_{n=1}^T \left(e^{-2n\epsilon^2} + 4ne^{-2n\epsilon^2} \int_0^{\mu_{\lambda} - \epsilon} (\mu_{\lambda} - \epsilon - x) e^{4n\epsilon(x - \mu_{\lambda} + \epsilon)} dx \right) \\ &\leq \sum_{n=1}^{\infty} e^{-2n\epsilon^2} \left(1 + \frac{1}{4n\epsilon^2} \right) \\ &= \frac{1}{e^{2\epsilon^2} - 1} + \frac{-\log(1 - e^{-2\epsilon^2})}{4\epsilon^2} \text{ by } -\log(1 - x) = \sum_{n=1}^{\infty} \frac{x^n}{n} \\ &= \frac{1}{e^{2\epsilon^2} - 1} + \frac{2\epsilon^2 + \log\left(\frac{1}{e^{2\epsilon^2} - 1}\right)}{4\epsilon^2} \\ &\leq \frac{1}{2\epsilon^2} + \frac{1}{2} + \frac{\log \frac{1}{2\epsilon^2}}{4\epsilon^2} \\ &\leq \frac{5}{8\epsilon^2} + \frac{\log \frac{1}{2\epsilon^2}}{4\epsilon^2} \text{ by } \epsilon < \frac{1}{2}. \end{aligned} \tag{25}$$

Combining (24) with (25) we obtain

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\tilde{\mu}^*(t) < \mu^* - \epsilon, [\lambda] \cap \mathcal{A}(t) \neq \emptyset] \right] \leq \frac{5K}{8\epsilon^2} + \frac{K \log \frac{1}{2\epsilon^2}}{4\epsilon^2}. \tag{26}$$

Next we bound the second term of (23). Note that $\{t \leq \tau_{\lambda}, [\lambda] \cap \mathcal{A}(t) = \emptyset\}$ implies that $\bar{\mu}_j(t') \leq \xi$ occurred for some $j \in [\lambda]$ and $t' < t$. Thus we have

$$\begin{aligned}
 & \sum_{t=1}^T \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, [\lambda] \cap \mathcal{A}(t) = \emptyset] \\
 & \leq \sum_{i \in [K] \setminus [\lambda]} \sum_{t=1}^T \mathbb{1} \left[a(t) = i, \bigcup_{j \in [\lambda]} \bigcup_{t' < t} \{\bar{\mu}_j(t') \leq \xi\} \right] \\
 & \leq \sum_{i \in [K] \setminus [\lambda]} \sum_{t=1}^T \mathbb{1}[a(t) = i] \left(\sum_{j \in [\lambda]} \mathbb{1} \left[\bigcup_t \{\bar{\mu}_j(t) \leq \xi\} \right] \right) \\
 & \leq \sum_{i \in [K] \setminus [\lambda]} \sum_{n=1}^T \mathbb{1}[\bar{\mu}_{i,n} \geq \xi] \left(\sum_{j \in [\lambda]} \mathbb{1} \left[\bigcup_t \{\bar{\mu}_j(t) \leq \xi\} \right] \right), \tag{27}
 \end{aligned}$$

where we used the same argument as (17) in (27). We can bound the expectation of (27) by Lemma 1 and 3 as

$$\begin{aligned}
 \sum_{t=1}^T \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, [\lambda] \cap \mathcal{A}(t) = \emptyset] & \leq \frac{\lambda \delta}{K} \left(\sum_{i \in [K] \setminus [\lambda]} n_i + \frac{K - \lambda}{2\epsilon^2} \right) \\
 & \leq \delta \sum_{i \in [K] \setminus [\lambda]} n_i + \frac{K}{8\epsilon^2} \text{ by } \lambda(K - \lambda) \leq K^2/4. \tag{28}
 \end{aligned}$$

We obtain the lemma by putting (23), (26) and (28) together. □

Proof of Theorem 3 The stopping time is decomposed into

$$\begin{aligned}
 \tau_\lambda & = \sum_{t=1}^\infty \mathbb{1}[a(t) \in [\lambda], t \leq \tau_\lambda] + \sum_{t=1}^\infty \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda] \\
 & \leq \sum_{t=1}^\infty \mathbb{1}[a(t) \in [\lambda]] \\
 & \quad + \sum_{t=1}^\infty \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, \tilde{\mu}^*(t) \geq \mu_\lambda - \epsilon] \sum_{t=1}^\infty \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, \tilde{\mu}^*(t) < \mu_\lambda - \epsilon] \\
 & \leq \sum_{t=1}^\infty \mathbb{1}[a(t) \in [\lambda]] \\
 & \quad + \sum_{t=1}^T \mathbb{1}[a(t) \notin [\lambda], \tilde{\mu}^*(t) \geq \mu_\lambda - \epsilon] + \sum_{t=T+1}^\infty \mathbb{1}[t \leq \tau_\lambda] \\
 & \quad + \sum_{t=1}^\infty \mathbb{1}[a(t) \notin [\lambda], t \leq \tau_\lambda, \tilde{\mu}^*(t) < \mu_\lambda - \epsilon].
 \end{aligned}$$

Lemmas 4–6 bound the expectation of these terms, which complete the proof. □

References

- Agrawal, S., & Goyal, N. (2012). Analysis of thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th annual conference on learning theory* (vol. 23, pp. 39.1–39.26).
- Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2), 235–256.
- Cappé, O., Garivier, A., Maillard, O. A., Munos, R., & Stoltz, G. (2012). Kullback-leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 41(3), 1516–1541.
- Choy, E., Bendit, M., McAleer, D., Liu, F., Feeney, M., Brett, S., et al. (2013). Safety, tolerability, pharmacokinetics and pharmacodynamics of an anti- oncostatin m monoclonal antibody in rheumatoid arthritis: Results from phase ii randomized, placebo-controlled trials. *Arthritis Research & Therapy*, 15(5), R132.
- Curtis, J., Yang, S., Chen, L., Pope, J., Keystone, E., Haraoui, B., et al. (2015). Determining the minimally important difference in the clinical disease activity index for improvement and worsening in early rheumatoid arthritis patients. *Arthritis Care & Research*, 67(10), 1345–1353.
- Even-Dar, E., Mannor, S., & Mansour, Y. (2006). Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7, 1079–1105.
- Genovese, M. C., Durez, P., Richards, H. B., Supronik, J., Dokoupilova, E., Mazurov, V., et al. (2013). Efficacy and safety of secukinumab in patients with rheumatoid arthritis: a phase ii, dose-finding, double-blind, randomised, placebo controlled study. *Annals of the Rheumatic Diseases*, 72(6), 863–869.
- Grieve, A. P., & Krams, M. (2005). ASTIN: A bayesian adaptive dose-response trial in acute stroke. *Clinical Trials*, 2(4), 340–351.
- Jamieson, K., Malloy, M., Nowak, R., & Bubeck, S. (2014). lil' ucb : An optimal exploration algorithm for multi-armed bandits. In *Proceedings of The 27th conference on learning theory* (vol. 35, pp. 423–439).
- Jun, K. S., Jamieson, K., Nowak, R., & Zhu, X. (2016). Top arm identification in multi-armed bandits with batch arm pulls. In *Proceedings of the 19th international conference on artificial intelligence and statistics* (pp. 139–148).
- Kalyanakrishnan, S., Tewari, A., Auer, P., & Stone, P. (2012). PAC subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th international conference on machine learning* (pp. 655–662).
- Kaufmann, E., Cappé, O., & Garivier, A. (2016). On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1), 1–42.
- Kim, E. S., Herbst, R. S., Wistuba, I. I., Lee, J. J., Blumenschein, G. R., Tsao, A., et al. (2011). The BATTLE trial: Personalizing therapy for lung cancer. *Cancer Discovery*, 1(1), 44–53.
- Koenig, L. W., & Law, A. M. (1985). A procedure for selecting a subset of size m containing the l best of k independent normal populations, with applications to simulation. *Communications in Statistics—Simulation and Computation*, 14(3), 719–734.
- Lai, T., & Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1), 4–22.
- Liu, F., Walters, S. J., & Julious, S. A. (2017). Design considerations and analysis planning of a phase 2a proof of concept study in rheumatoid arthritis in the presence of possible non-monotonicity. *BMC Medical Research Methodology*, 17(1), 149.
- Locatelli, A., Gutzeit, M., & Carpentier, A. (2016). An optimal algorithm for the thresholding bandit problem. In *Proceedings of the 33rd international conference on machine learning* (pp. 1690–1698).
- Mukherjee, S., Naveen, K. P., Sudarsanam, N., & Ravindran, B. (2017). Thresholding bandits with augmented UCB. In *Proceedings of the 26th international joint conference on artificial intelligence* (pp. 2515–2521).
- Schmidt, C., Branke, J., & Chick, S. E. (2006). Integrating techniques from statistical ranking into evolutionary algorithms. In *Applications of evolutionary computing* (pp. 752–763). Springer: Heidelberg.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (1st ed.). Cambridge: MIT Press.
- Tang, L., Jiang, Y., Li, L., Zeng, C., & Li, T. (2015). Personalized recommendation via parameter-free contextual bandits. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval* (pp. 323–332).
- Zhou, Y., Chen, X., & Li, J. (2014). Optimal pac multiple arm identification with applications to crowdsourcing. In *Proceedings of the 31st international conference on machine learning* (pp. 217–225).

Affiliations

Hideaki Kano^{1,2} · Junya Honda^{1,2} · Kentaro Sakamaki¹ · Kentaro Matsuura³ ·
Atsuyoshi Nakamura⁴ · Masashi Sugiyama^{1,2}

Junya Honda
jhonda@k.u-tokyo.ac.jp

Kentaro Sakamaki
sakamaki@m.u-tokyo.ac.jp

Kentaro Matsuura
kmatsuu@its.jnj.com

Atsuyoshi Nakamura
atsu@main.ist.hokudai.ac.jp

Masashi Sugiyama
sugi@k.u-tokyo.ac.jp

¹ University of Tokyo, Tokyo, Japan

² RIKEN, Tokyo, Japan

³ Johnson & Johnson, Tokyo, Japan

⁴ Hokkaido University, Sapporo, Japan