
Multiple Identifications in Multi-Armed Bandits

Sébastien Bubeck

SBUBECK@PRINCETON.EDU

Department of Operations Research and Financial Engineering, Princeton University

Tengyao Wang

TENGYAOW@PRINCETON.EDU

Department of Mathematics, Princeton University

Nitin Viswanathan

NVISWANA@PRINCETON.EDU

Department of Computer Science, Princeton University

Abstract

We study the problem of identifying the top m arms in a multi-armed bandit game. Our proposed solution relies on a new algorithm based on successive rejects of the seemingly bad arms, and successive accepts of the good ones. This algorithmic contribution allows to tackle other multiple identifications settings that were previously out of reach. In particular we show that this idea of successive accepts and rejects applies to the multi-bandit best arm identification problem.

1. Introduction

We are interested in the following situation: An agent faces K unknown distributions, and he is allowed to do n sequential evaluations of the form (i, X) where $i \in \{1, \dots, K\}$ is chosen by the agent and X is a random variable drawn from the i^{th} distribution and revealed to the agent. The goal of the agent after the n evaluations is to identify a subset of the distributions (or *arms* in the multi-armed bandit terminology) corresponding to some prespecified criterion. This setting was introduced in (Bubeck et al., 2009), where the goal was to identify the distribution with maximal mean. Note that in this formulation of the problem the evaluation budget n is fixed. Another possible formulation is the one of the PAC model studied in (Even-Dar et al., 2002; Mannor & Tsitsiklis, 2004) where an accuracy ε and a probability of correctness δ are prespecified, and one wants to minimize the number of evaluations to attain this prespecified accuracy

and probability of correctness. This latter formulation has a long history which goes back to the seminal work (Bechhofer, 1954).

In this paper we focus on the fixed budget setting of (Bubeck et al., 2009). For this fixed budget problem, (Audibert et al., 2010) proposed a new analysis and an optimal algorithm (up to a logarithmic factor). In particular this work introduced a notion of *best arm identification complexity*, and it was shown that this quantity, denoted H , characterizes the hardness of identifying the best distribution in a specific set of K distributions. Intuitively, it was shown that the number of evaluations n has to be $\Omega(H/\log K)$ to be able to find the best arm, and the algorithm SR (Successive Rejects) finds it with $\mathcal{O}(H \log^2 K)$ evaluations. Furthermore in the latter paper the authors also suggested the open problem of generalizing the analysis and algorithms to the identification of the m distributions with the top m means. Our main contribution is to solve this open problem. We suggest a non-trivial extension of the complexity H , denoted $H^{(m)}$, to the problem of identifying the top m distributions, and we introduce a new algorithm, called SAR (Successive Accepts and Rejects), that requires only $\tilde{\mathcal{O}}(H^{(m)})$ evaluations¹ to find the top m arms. We also propose a numerical comparison between SAR, SR and uniform sampling for the problem of finding the m top arms. Interestingly the experiments show that SR performs badly for $m > 1$, which shows that the tradeoffs involved in this generalized problem are fundamentally different from the ones for the single best arm identification.

Note that this problem and the associated complexity $H^{(m)}$ have been studied recently in the PAC model

Proceedings of the 30th International Conference on Machine Learning, Atlanta, Georgia, USA, 2013. JMLR: W&CP volume 28. Copyright 2013 by the author(s).

¹In the m -best arms identification problem we write $u_n = \tilde{\mathcal{O}}(v_n)$ when $u_n = \mathcal{O}(v_n)$ up to logarithmic factor in K

in (Kalyanakrishnan et al., 2012). The present paper is a concurrent and independent work with respect to the latter paper. In particular, the algorithm LUCB designed in (Kalyanakrishnan et al., 2012) is based on fundamentally different ideas (confidence intervals and UCB type algorithms) than the one developed in this paper.

As a by-product of our new analysis we are also able to solve an open problem of (Gabillon et al., 2011). In this paper the authors studied the setting where the agent faces M distinct best arm identification problems. A multi-bandit identification complexity was introduced, that we denote $H^{[M]}$. On the contrary to the setting of single best arm identification, here the algorithm proposed in (Gabillon et al., 2011) that needs of order of $H^{[M]}$ evaluations to find the best arm in each bandit requires to know the complexity $H^{[M]}$ to tune its parameters. Using our SAR machinery, we construct a parameter-free algorithm that identifies the best arm in each bandit with $\tilde{\mathcal{O}}(H^{[M]})$ evaluations².

Both the m -best arms identification and the multi-bandit best arm identification have numerous potential applications. We refer the interested reader to the previously cited papers for several examples.

2. Problem setup

We adopt the terminology of multi-armed bandits. The agent faces K arms and he has a budget of n evaluations (or *pulls*). To each arm $i \in \{1, \dots, K\}$ there is an associated probability distribution ν_i , supported³ on $[0, 1]$. These distributions are unknown to the agent. The sequential evaluations protocol goes as follows: at each round $t = 1, \dots, n$, the agent chooses an arm I_t , and observes a reward drawn from ν_{I_t} independently from the past given I_t . In the m -best arms identification problem, at the end of the n evaluations, the agent selects m arms denoted J_1, \dots, J_m . The objective of the agent is that the set $\{J_1, \dots, J_m\}$ corresponds to the set of arms with the m highest mean rewards.

Denote by μ_1, \dots, μ_K the mean of the arms. In the following we assume that $\mu_1 > \dots > \mu_K$. The ordering

²In the multi-bandit best arm identification problem we write $u_n = \tilde{\mathcal{O}}(v_n)$ when $u_n = \mathcal{O}(v_n)$ up to logarithmic factor in MK

³One can directly generalize the discussion to σ -subgaussian distributions.

assumption comes without loss of generality, and the assumption that the means are all distinct is made for sake of notation (the complexity measures are slightly different if there is an ambiguity for the top m means). We evaluate the performance of the agent's strategy by the probability of misidentification, that is

$$e_n = \mathbb{P}(\{J_1, \dots, J_m\} \neq \{1, \dots, m\}).$$

Finer measures of performance can be proposed, such as the simple regret $r_n = \sum_{i=1}^m (\mu_i - \mathbb{E}\mu_{J_i})$. However, as it was argued in (Audibert et al., 2010), for a first order analysis it is enough to focus on the quantity e_n .

In the (single) best arm identification, (Audibert et al., 2010) introduced the following complexity measures. Let $\Delta_i = \mu_1 - \mu_i$ for $i \neq 1$, $\Delta_1 = \mu_1 - \mu_2$,

$$H_1 = \sum_{i=1}^K \frac{1}{\Delta_i^2} \quad \text{and} \quad H_2 = \max_{i \in \{1, \dots, K\}} i \Delta_i^{-2}.$$

It is easy to see that these two complexity measures are equivalent up to a logarithmic factor since we have (see (Audibert et al., 2010))

$$H_2 \leq H_1 \leq \log(2K)H_2. \quad (1)$$

[Theorem 4, (Audibert et al., 2010)] shows that the complexity H_1 represents the hardness of the best arm identification problem. While the proof of the latter result is difficult and technical, it is intuitively obvious that H_1 is a lower bound on the number of evaluations necessary to identify the best arm. Indeed, merely to *check* if an arm has mean μ_i or μ^* , one needs to sample it of order of $1/\Delta_i^2$ times. The surprising fact is that of order of H_1 evaluations suffices to identify the best arm. Note that as far as upper bounds are concerned, the quantity H_2 proved to be a useful surrogate for H_1 to express the bounds on e_n .

For the m -best arms identification problem we define the following gaps and the associated complexity measures:

$$\begin{aligned} \Delta_i^{(m)} &= \begin{cases} \mu_i - \mu_{m+1} & \text{if } i \leq m \\ \mu_m - \mu_i & \text{if } i > m \end{cases}, \\ H_1^{(m)} &= \sum_{i=1}^K \frac{1}{(\Delta_i^{(m)})^2}, \\ H_2^{(m)} &= \max_{i \in \{1, \dots, K\}} i \left(\Delta_i^{(m)} \right)^{-2}, \end{aligned}$$

where the notation $(i) \in \{1, \dots, K\}$ is defined such that $\Delta_{(1)}^{(m)} \leq \dots \leq \Delta_{(K)}^{(m)}$. We conjecture that a similar lower bound to [Theorem 4, (Audibert et al., 2010)] with H_1 replaced by $H_1^{(m)}$ holds true for the m -best arms identification problem. Again such a statement is pretty much obvious. Indeed assume that we know the critical values μ_m and μ_{m+1} . Then for each arm we have an hypothesis testing problem which corresponds to deciding whether $\mu_i > \mu_{m+1}$ (i.e., the arm is in the m best) or $\mu_i < \mu_m$ (i.e., the arm is not in the m best). For this hypothesis testing problem to be feasible, it is easy to see that one needs of order of $1/\left(\Delta_i^{(m)}\right)^2$ evaluations of arm i . While this does not prove the lower bound, it strongly suggests that $\Omega\left(H_1^{(m)}\right)$ evaluations are necessary to identify the top m arms.

In this paper we shall focus on positive results. In particular we will prove an upper bound on e_n that gets small when $n = \tilde{O}\left(H_2^{(m)}\right)$ (recall that by (1), $\tilde{O}\left(H_2^{(m)}\right) = \tilde{O}\left(H_1^{(m)}\right)$). This result is derived in Section 3, where we introduce our key algorithmic contribution, the SAR (Successive Accepts and Rejects) algorithm. We also present experiments for this setting in Section 5.

In Section 4 we consider the framework of multi-bandit introduced in (Gabillon et al., 2011), where the agent faces M distinct best arm identification problems. For sake of notation we assume that each problem $m \in \{1, \dots, M\}$ has the same number of arms K . We also restrict our attention to the single best arm identification within each problem, but we could deal with m -best arms identification within each problem. We denote by $\nu_1(m), \dots, \nu_K(m)$ the unknown distributions of the arms in problem m . We define similarly all the relevant quantities for each problem, that is $\mu_1(m) > \dots > \mu_K(m), \Delta_1(m), \dots, \Delta_K(m), H_1(m)$ and $H_2(m)$. Finally we denote by (i, m) the arm i in problem m . In the multi-bandit best arm identification, the forecaster performs n sequential evaluations of the form $(I_t, m_t) \in \{1, \dots, K\} \times \{1, \dots, M\}$. At the end of the n evaluations, the agent selects one arm for each problem, denoted $(J_1, 1), \dots, (J_M, M)$. The objective of the agent is to find the arm with the highest mean reward in each problem, that is in this setting the probability of misidentification can be written as

$$e_n = \mathbb{P}(\exists m \in \{1, \dots, M\} : J_m \neq 1).$$

Following (Gabillon et al., 2011) we introduce the fol-

lowing complexity measure

$$H_1^{[M]} = \sum_{m=1}^M H_1(m).$$

Again we define a sort of weaker complexity measure by ordering the gaps. Let

$$\Delta_1^{[M]} \leq \Delta_2^{[M]} \leq \dots \leq \Delta_{MK}^{[M]}$$

be a rearrangement of $\{\Delta_i(m) : 1 \leq i \leq K, 1 \leq m \leq M\}$ in ascending order, and let

$$H_2^{[M]} = \max_{k \in \{1, \dots, MK\}} k \left(\Delta_k^{[M]}\right)^{-2}.$$

Similarly to the m -best arms identification problem we conjecture that the lower bound [Theorem 4, (Audibert et al., 2010)] with H_1 replaced by $H_1^{[M]}$ holds true for the multi-bandit best arm identification problem. Here we prove, using a variation of SAR, that $n = \tilde{O}\left(H_2^{[M]}\right)$ (recall that by (1), $\tilde{O}\left(H_2^{[M]}\right) = \tilde{O}\left(H_1^{[M]}\right)$) is enough to have a small probability of error e_n . This result is derived in Section 4. The improvement with respect to (Gabillon et al., 2011) is that our strategy is parameter-free, while the theoretical Gap-E introduced in (Gabillon et al., 2011) requires the knowledge of $H_1^{[M]}$ to tune its parameter. Moreover the analysis of SAR is much simpler than the one of Gap-E.

For each arm i and all time rounds $t \geq 1$, we denote by $T_i(t) = \sum_{s=1}^t \mathbb{1}_{I_s=i}$ the number of times arm i was pulled from rounds 1 to t , and by $X_{i,1}, X_{i,2}, \dots, X_{i,T_i,t}$ the sequence of associated rewards. Introduce $\hat{\mu}_{i,s} = \frac{1}{s} \sum_{t=1}^s X_{i,t}$ the empirical mean of arm i after s evaluations. Denote by $X_{i,s}(m)$ and $\hat{\mu}_{i,s}(m)$ the corresponding quantities in the multi-bandit problem.

3. m -best arms identification

In this section we describe and analyze a new algorithm, called SAR (Successive Accepts and Rejects), for the m -best arms identification problem, see Figure 1 for its precise description. The idea behind SAR is similar to the one for SR (Successive Rejects) that was designed for the (single) best arm identification problem, with the additional feature that SAR sometimes *accepts* an arm because it is confident enough that this arm is among the m top arms. Informally SAR proceeds as follows. First the algorithm divides the time (i.e., the n rounds) in $K - 1$ phases. At the end of each phase, the algorithm either accepts the

arm with the highest empirical mean or dismisses the arm with the lowest empirical mean, and in both cases the corresponding arm is deactivated. During the next phase, it pulls equally often each active arm. The key to decide whether to accept or reject during a certain phase k is to rely on estimates for the gaps $\Delta_i^{(m)}$. More precisely, assume that the algorithm has already accepted $m - m(k)$ arms $J_1, \dots, J_{m-m(k)}$, i.e. there is $m(k)$ arms left to find. Then, at the end of phase k , SAR computes for the $m(k)$ empirical best arms (among the active arms) the distance (in terms of empirical mean) to the $(m(k) + 1)^{th}$ empirical best arm among the active arms. On the other hand for the active arms that are not among the $m(k)$ empirical best arms, SAR computes the distance to the $m(k)^{th}$ empirical best arm. Finally SAR deactivates the arm i_k that maximizes these empirical distances. If i_k is currently the empirical best arm, then SAR accepts i_k and sets $m(k+1) = m(k) - 1$, $J_{m-m(k+1)} = i_k$, and otherwise it simply rejects i_k . The length of the phases are chosen similarly to what was done for the SR algorithm.

Theorem 1 *The probability of error of SAR in the m -best arms identification problem satisfies*

$$e_n \leq 2K^2 \exp\left(-\frac{n-K}{8\log(K)H_2^{(m)}}\right).$$

Proof Consider the event ξ defined by

$$\xi = \left\{ \forall i \in \{1, \dots, K\}, k \in \{1, \dots, K-1\}, \left| \frac{1}{n_k} \sum_{s=1}^{n_k} X_{i,s} - \mu_i \right| \leq \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \right\}.$$

By Hoeffding's Inequality and an union bound, the probability of the complementary event $\bar{\xi}$ can be bounded as follows

$$\begin{aligned} \mathbb{P}(\bar{\xi}) &\leq \sum_{i=1}^K \sum_{k=1}^{K-1} \mathbb{P}\left(\left| \frac{1}{n_k} \sum_{s=1}^{n_k} X_{i,s} - \mu_i \right| > \frac{1}{4} \Delta_{(K+1-k)}^{(m)}\right) \\ &\leq \sum_{i=1}^K \sum_{k=1}^{K-1} 2 \exp(-2n_k (\Delta_{(K+1-k)}^{(m)}/4)^2) \\ &\leq 2K^2 \exp\left(-\frac{n-K}{8\log(K)H_2^{(m)}}\right), \end{aligned}$$

Let $A_1 = \{1, \dots, K\}$, $m(1) = m$, $\overline{\log}(K) = \frac{1}{2} + \sum_{i=2}^K \frac{1}{i}$, $n_0 = 0$ and for $k \in \{1, \dots, K-1\}$,

$$n_k = \left\lceil \frac{1}{\overline{\log}(K)} \frac{n-K}{K+1-k} \right\rceil.$$

For each phase $k = 1, 2, \dots, K-1$:

- (1) For each active arm $i \in A_k$, select arm i for $n_k - n_{k-1}$ rounds.
- (2) Let $\sigma_k : \{1, \dots, K+1-k\} \rightarrow A_k$ be the bijection that orders the empirical means by $\hat{\mu}_{\sigma_k(1), n_k} \geq \hat{\mu}_{\sigma_k(2), n_k} \geq \dots \geq \hat{\mu}_{\sigma_k(K+1-k), n_k}$. For $1 \leq r \leq K+1-k$, define empirical gaps

$$\hat{\Delta}_{\sigma_k(r), n_k} = \begin{cases} \hat{\mu}_{\sigma_k(r), n_k} - \hat{\mu}_{\sigma_k(m(k)+1), n_k} & \text{if } r \leq m(k) \\ \hat{\mu}_{\sigma_k(m(k)), n_k} - \hat{\mu}_{\sigma_k(r), n_k} & \text{if } r \geq m(k) + 1 \end{cases}$$

- (3) Let $i_k \in \operatorname{argmax}_{i \in A_k} \hat{\Delta}_{i, n_k}$ (ties broken arbitrarily). Deactivate arm i_k , that is set $A_{k+1} = A_k \setminus \{i_k\}$.
- (4) If $\hat{\mu}_{i_k, n_k} > \hat{\mu}_{\sigma_k(m(k)+1), n_k}$ then arm i_k is accepted, that is set $m(k+1) = m(k) - 1$ and $J_{m-m(k+1)} = i_k$.

Output: The m accepted arms J_1, \dots, J_m .

Figure 1. SAR (Successive Accepts and Rejects) algorithm for m -best arms identification.

where the last inequality comes from the fact that

$$\begin{aligned} &n_k \left(\Delta_{(K+1-k)}^{(m)} \right)^2 \\ &\geq \frac{n-K}{\log(K)(K+1-k)} \left(\Delta_{(K+1-k)}^{(m)} \right)^{-2} \\ &\geq \frac{n-K}{\log(K)H_2^{(m)}}. \end{aligned}$$

Thus, it suffices to show that on the event ξ , the algorithm does not make any error. We prove this by induction on k . Let $k \geq 1$. Assume the algorithm makes no error in all previous $k-1$ stages, that is no bad arm $i > m$ has been accepted and no good arm $i \leq m$ has been rejected. Note that event ξ implies that at the end of stage k , all empirical means are within $\frac{1}{4} \Delta_{(K+1-k)}^{(m)}$ of the respective true means.

Let $A_k = \{a_1, \dots, a_{K+1-k}\}$ be the the set of active arms during phase k . We order the a_i 's such that $\mu_{a_1} > \mu_{a_2} > \dots > \mu_{a_{K+1-k}}$. To slightly lighten the

notation we denote $m' = m(k)$ for the number of arms that are left to find in phase k . The assumption that no error occurs in the first $k - 1$ stages implies that

$$a_1, a_2, \dots, a_{m'} \in \{1, \dots, m\},$$

and

$$a_{m'+1}, \dots, a_{K+1-k} \in \{m+1, \dots, K\}.$$

If an error is made at stage k , it can be one of the following two types:

1. The algorithm accepts a_j at stage k for some $j \geq m' + 1$.
2. The algorithm rejects a_j at stage k for some $j \leq m'$.

Again to slightly shorten the notation we denote $\sigma = \sigma_k$ for the bijection (from $\{1, \dots, K + 1 - k\}$ to A_k) such that $\hat{\mu}_{\sigma(1), n_k} \geq \hat{\mu}_{\sigma(2), n_k} \geq \dots \geq \hat{\mu}_{\sigma(K+1-k), n_k}$. Suppose Type 1 error occurs. Then $a_j = \sigma(1)$ since if the algorithm accepts, it must accept the empirical best arm. Furthermore we also have

$$\hat{\mu}_{a_j, n_k} - \hat{\mu}_{\sigma(m'+1), n_k} \geq \hat{\mu}_{\sigma(m'), n_k} - \hat{\mu}_{\sigma(K+1-k), n_k}, \quad (2)$$

since otherwise the algorithm would rather reject arm $\sigma(K + 1 - k)$. The condition $a_j = \sigma(1)$ and the event ξ implies that

$$\begin{aligned} \hat{\mu}_{a_j, n_k} &\geq \hat{\mu}_{a_1, n_k} \\ \Rightarrow \mu_{a_j} + \frac{1}{4} \Delta_{(K+1-k)}^{(m)} &\geq \mu_{a_1} - \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \\ \Rightarrow \Delta_{(K+1-k)}^{(m)} &> \frac{1}{2} \Delta_{(K+1-k)}^{(m)} \geq \mu_{a_1} - \mu_{a_j} \geq \mu_{a_1} - \mu_{m+1} \end{aligned}$$

We then look at the condition (2). In the event of ξ , for all $i \leq m'$, we have

$$\begin{aligned} \hat{\mu}_{a_i, n_k} &\geq \mu_{a_i} - \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \\ &\geq \mu_{a_{m'}} - \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \\ &\geq \mu_m - \frac{1}{4} \Delta_{(K+1-k)}^{(m)}. \end{aligned}$$

So there are $m + 1$ arms in A_k (namely $a_1, a_2, \dots, a_{m'}, a_j$) whose empirical means are at least $\mu_m - \frac{1}{4} \Delta_{(K+1-k)}^{(m)}$, which means $\hat{\mu}_{\sigma(m'+1), n_k} \geq \mu_m - \frac{1}{4} \Delta_{(K+1-k)}^{(m)}$. On the other hand, $\hat{\mu}_{\sigma(K+1-k), n_k} \leq \hat{\mu}_{a_{K+1-k}, n_k} \leq \mu_{a_{K+1-k}} + \frac{1}{4} \Delta_{(K+1-k)}^{(m)}$. Therefore, using those two observations and (2) we

deduce

$$\begin{aligned} &\left(\mu_{a_j} + \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \right) - \left(\mu_m - \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \right) \\ &\geq \left(\mu_m - \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \right) - \left(\mu_{a_{K+1-k}} + \frac{1}{4} \Delta_{(K+1-k)}^{(m)} \right) \\ &\Rightarrow \Delta_{(K+1-k)}^{(m)} \geq 2\mu_m - \mu_{a_j} - \mu_{a_{K+1-k}} > \mu_m - \mu_{a_{K+1-k}}. \end{aligned}$$

Thus so far we proved that if there is a Type 1 error, then

$$\Delta_{(K+1-k)}^{(m)} > \max(\mu_{a_1} - \mu_m, \mu_m - \mu_{a_{K+1-k}}).$$

But at stage k , only $k - 1$ arms have been accepted or rejected, thus $\Delta_{(K+1-k)}^{(m)} \leq \max(\mu_{a_1} - \mu_m, \mu_m - \mu_{a_{K+1-k}})$. By contradiction, we conclude that Type 1 error does not occur.

Suppose Type 2 error occurs. The reasoning is symmetric to Type 1. In fact, if we rephrase the problem as finding the $K - m$ worst arms instead of the m best arms, this is exactly the same as Type 1 error. Hence Type 2 error cannot occur as well. This completes the induction and consequently the proof of the theorem. \blacksquare

4. Multi-bandit best arm identification

In this section we use the idea of SAR for multi-bandit best arm identification. Here at the end of each phase we estimate the gaps $\Delta_i(m)$ within each problem, and we reject the arm with the largest such estimated gap. Moreover if a problem is left with only one active arm, then this arm is accepted and the problem is deactivated. The corresponding strategy is described precisely in Figure 2

Theorem 2 *The probability of error of SAR in the multi-bandit best arm identification problem satisfies*

$$e_n \leq 2M^2 K^2 \exp\left(-\frac{n - MK}{8 \log(MK) H_2^{[M]}}\right).$$

Proof Consider the event ξ defined by

$$\xi = \left\{ \forall 1 \leq i \leq K, 1 \leq m \leq M, 1 \leq k \leq MK - 1, \left| \frac{1}{n_k} \sum_{s=1}^{n_k} X_{i,s}(m) - \mu_i(m) \right| \leq \frac{1}{4} \Delta_{(MK+1-k)} \right\}.$$

Following the same reasoning as in the proof of Theorem 1, it suffices to show that in the event of ξ the

Let $A_1 = \{(1, 1), \dots, (K, M)\}$, $\overline{\log}(MK) = \frac{1}{2} + \sum_{i=2}^{MK} \frac{1}{i}$, $n_0 = 0$ and for $k \in \{1, \dots, MK - 1\}$,

$$n_k = \left\lceil \frac{1}{\overline{\log}(MK)} \frac{n - MK}{MK + 1 - k} \right\rceil.$$

For each phase $k = 1, 2, \dots, MK - 1$:

- (1) For each active pair (arm, problem) $(i, m) \in A_k$, select arm i in problem m for $n_k - n_{k-1}$ rounds.
- (2) Let $h_k(m)$ be the arm with the highest empirical mean $\widehat{\mu}_{i, n_k}(m)$ among the active arms in the active problem m (that is such that $(i, m) \in A_k$).
- (3) If there is a problem m such that $h_k(m)$ is the last active arm in problem m , then deactivate both the arm and the problem, and accept the arm. That is, set $A_{k+1} = A_k \setminus \{(h_k(m), m)\}$ and $J_m = h_k(m)$. Otherwise proceed to step (4).

(4) Let

$$(i_k, m_k) \in \operatorname{argmax}_{(i, m) \in A_k} (\widehat{\mu}_{h_k(m), n_k}(m) - \widehat{\mu}_{i, n_k}(m)).$$

(Ties broken arbitrarily.) Deactivate arm i_k in problem m_k , that is set $A_{k+1} = A_k \setminus \{(i_k, m_k)\}$.

Output: The M accepted arms $(J_1, 1), \dots, (J_M, M)$ (where the last accepted arm is defined by the unique element of A_{MK}).

Figure 2. SAR (Successive Accepts and Rejects) algorithm for the multi-bandit best arm identification.

algorithm makes no error. We do this by induction on the phase k of the algorithm. Let $k \geq 1$. Assume the algorithm makes no error in all previous $k - 1$ stages. Then at phase k , for each active problem m , the arm $(1, m)$ is still active. Moreover, as only $k - 1$ arms have been deactivated, one clearly has

$$\max_{(i, m) \in A_k} (\mu_1(m) - \mu_i(m)) \geq \Delta_{(MK+1-k)}.$$

Suppose the above maximum is achieved for the arm (i^*, m^*) , so we have

$$\mu_1(m^*) - \mu_{i^*}(m^*) \geq \Delta_{(MK+1-k)}. \quad (3)$$

Assume now that the algorithm makes an error at the end of phase k , i.e. some arm $(1, m)$ is deactivated and it was not the last active arm in problem m . For this to happen, we necessarily have for some $j \in \{2, \dots, K\}$ (e.g., $j = h_k(m)$),

$$\widehat{\mu}_{j, n_k}(m) - \widehat{\mu}_{1, n_k}(m) \geq \widehat{\mu}_{1, n_k}(m^*) - \widehat{\mu}_{i^*, n_k}(m^*). \quad (4)$$

Clearly on the event ξ one has

$$\begin{aligned} & \widehat{\mu}_{j, n_k}(m) - \widehat{\mu}_{1, n_k}(m) \\ &= \widehat{\mu}_{j, n_k}(m) - \mu_j(m) \\ & \quad + \mu_j(m) - \mu_1(m) + \mu_1(m) - \widehat{\mu}_{1, n_k}(m) \\ &< \frac{1}{2} \Delta_{(MK+1-k)}. \end{aligned}$$

On the other hand, using (3) and ξ , one has

$$\begin{aligned} & \widehat{\mu}_{1, n_k}(m^*) - \widehat{\mu}_{i^*, n_k}(m^*) \\ &= \widehat{\mu}_{1, n_k}(m^*) - \mu_1(m^*) \\ & \quad + \mu_1(m^*) - \mu_{i^*}(m^*) + \mu_{i^*}(m^*) - \widehat{\mu}_{i^*, n_k}(m^*) \\ &\geq \frac{1}{2} \Delta_{(MK+1-k)}. \end{aligned}$$

Therefore, $\widehat{\mu}_{1, n_k}(m^*) - \widehat{\mu}_{i^*, n_k}(m^*) > \widehat{\mu}_{j, n_k}(m) - \widehat{\mu}_{1, n_k}(m)$, contradicting (4). This completes the induction and the proof. \blacksquare

5. Experiments

In this section we revisit the simple experiments of (Audibert et al., 2010) in the setting of multiple identifications. Since our objective is simply to illustrate our theoretical analysis we focus on the m -best arms identification problem, but similar numerical simulations could be conducted in the multi-bandit setting and compared to the results of (Gabillon et al., 2011).

We compare our proposed strategy SAR to three competitors: The uniform sampling strategy that divides evenly the allocation budget n between the K arms, and then return the m arms with the highest empirical mean (see (Bubeck et al., 2011) for a discussion of this strategy in the single best arm identification). The SR strategy is the plain Successive Rejects strategy of (Audibert et al., 2010) which was designed to find the (single) best arm. We slightly improve it for m -best identification by running only $K - m - 1$ phases (while still using the full budget n) and then returning the last m surviving arms. Finally we consider the extension of UCB-E to the m -best arms identification problem, which is based on a similar idea than the extension Gap-E of (Gabillon et al., 2011) for the multi-bandit best arm identification, see Figure 3 for the details. Note that this last algorithm requires to know the complexity $H_1^{(m)}$. One could propose an adaptive version, using ideas described in (Audibert et al., 2010), but for sake of simplicity we restrict our attention to the non-adaptive algorithm.

In our experiments we consider only Bernoulli distributions, and the optimal arm always has parameter

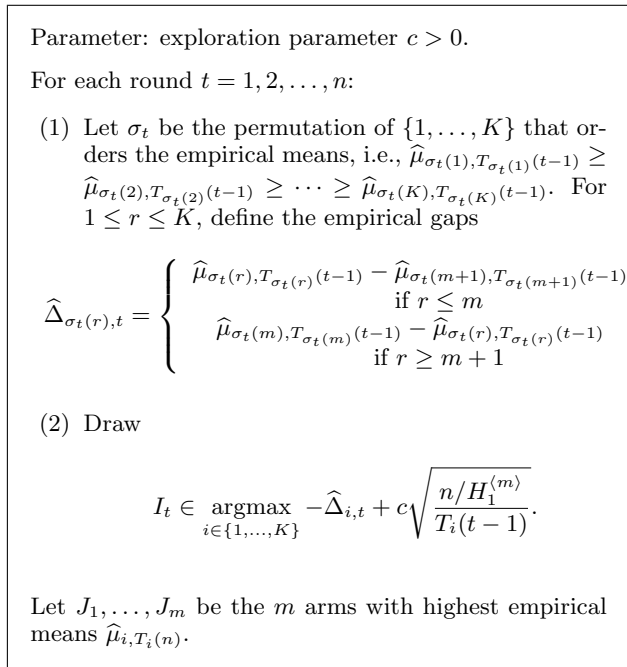


Figure 3. Gap-E algorithm for the m -best arms identification problem.

1/2. Each experiment corresponds to a different situation for the gaps, they are either clustered in few groups, or distributed according to an arithmetic or geometric progression. For each experiment we plot the probability of misidentification for each strategy, varying m between 2 and $K-1$. The allocation budget for each experiment is chosen to be roughly equal to $\max_{1 \leq m \leq K-1} H_1^{(m)}$. We report our results in Figure 4. The parameters for the experiments are as follows:

- Experiment 1: One group of bad arms, $K = 20$, $\mu_{2:20} = 0.4$ (meaning for any $j \in \{2, \dots, 20\}$, $\mu_j = 0.4$)
- Experiment 2: Two groups of bad arms, $K = 20$, $\mu_{2:6} = 0.42$, $\mu_{7:20} = 0.38$.
- Experiment 3: Geometric progression, $K = 4$, $\mu_i = 0.5 - (0.37)^i$, $i \in \{2, 3, 4\}$.
- Experiment 4: 6 arms divided in three groups, $K = 6$, $\mu_2 = 0.42$, $\mu_{3:4} = 0.4$, $\mu_{5:6} = 0.35$.
- Experiment 5: Arithmetic progression, $K = 15$, $\mu_i = 0.5 - 0.025i$, $i \in \{2, \dots, 15\}$.
- Experiment 6: Three groups of bad arms, $K = 30$, $\mu_{2:6} = 0.45$, $\mu_{7:20} = 0.43$, $\mu_{21:30} = 0.38$.

It is interesting to note that SR performs badly for m -best arms identification when $m > 1$, as it has even worse performances than the naive uniform sampling in many cases. This shows that the tradeoffs involved in finding the single best arm and finding the top m arms are fundamentally different. As expected SAR always outperforms uniform sampling, and Gap-E has slightly better performances than SAR (but Gap-E requires an extra information to tune its parameter, and the adaptive version comes with no provable guarantee).

References

- Audibert, J.-Y., Bubeck, S., and Munos, R. Best arm identification in multi-armed bandits. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, 2010.
- Bechhofer, R. E. A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Annals of Mathematical Statistics*, 25:16–39, 1954.
- Bubeck, S., Munos, R., and Stoltz, G. Pure exploration in multi-armed bandits problems. In *Proceedings of the 20th International Conference on Algorithmic Learning Theory (ALT)*, 2009.
- Bubeck, S., Munos, R., and Stoltz, G. Pure exploration in finitely-armed and continuously-armed bandits. *Theoretical Computer Science*, 412:1832–1852, 2011.
- Even-Dar, E., Mannor, S., and Mansour, Y. Pac bounds for multi-armed bandit and markov decision processes. In *Proceedings of the Fifteenth Annual Conference on Computational Learning Theory (COLT)*, 2002.
- Gabillon, V., Ghavamzadeh, M., Lazaric, A., and Bubeck, S. Multi-bandit best arm identification. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. Pac subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012.
- Mannor, S. and Tsitsiklis, J. N. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5:623–648, 2004.

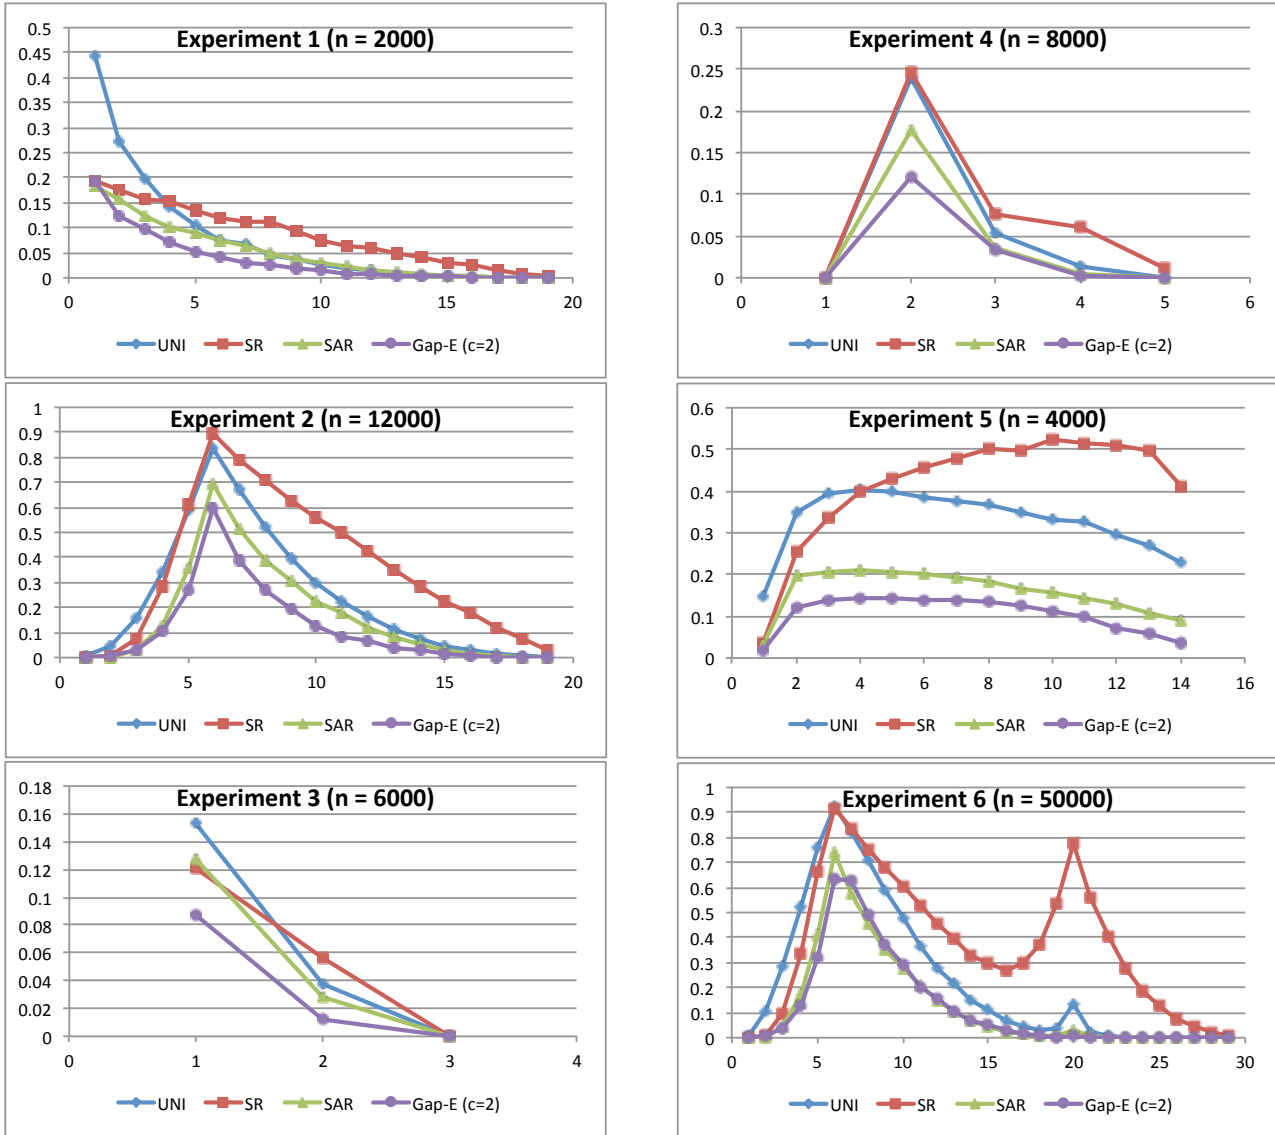


Figure 4. Numerical simulations for the m -best arms identification problem. We chose $c = 2$ (exploration parameter) for the Gap-E algorithm in all experiments.