
Indexed Minimum Empirical Divergence for Unimodal Bandits

Hassan Saber

Université de Lille, Inria, CNRS, Centrale Lille
UMR 9189 – CRIStAL, F-59000 Lille, France
hassan.saber@inria.fr

Pierre Ménard

Otto von Guericke Universität Magdeburg
pierre.menard@ovgu.de

Odalric-Ambrym Maillard

Université de Lille, Inria, CNRS, Centrale Lille
UMR 9189 – CRIStAL, F-59000 Lille, France
odalric.maillard@inria.fr

Abstract

We consider a multi-armed bandit problem specified by a set of one-dimensional family exponential distributions endowed with a unimodal structure. We introduce IMED-UB, an algorithm that optimally exploits the unimodal-structure, by adapting to this setting the Indexed Minimum Empirical Divergence (IMED) algorithm introduced by Honda and Takemura [2015]. Owing to our proof technique, we are able to provide a concise finite-time analysis of the IMED-UB algorithm. Numerical experiments show that IMED-UB competes with the state-of-the-art algorithms.

1 Introduction

The multi-armed bandit problem is a popular framework to formalize sequential decision making problems. It was first introduced in the context of medical trials [Thompson, 1933, 1935] and later formalized by Robbins [1952]: A bandit is specified by a configuration, that is a set of unknown probability distributions, $\nu = (\nu_a)_{a \in \mathcal{A}}$ with means $(\mu_a)_{a \in \mathcal{A}}$. At each time $t \in \mathbb{N}$, the learner chooses an arm $a_t \in \mathcal{A}$, based only on the past, the learner then receives and observes a reward X_t , conditionally independent, sampled according to ν_{a_t} . The goal of the learner is to maximize the expected sum of rewards received over time (up to some unknown horizon T), or equivalently minimize the *regret* with respect to the algorithm constantly receiving the highest mean reward

$$R(\nu, T) = \mathbb{E}_\nu \left[\sum_{t=1}^T \mu^* - X_t \right] \text{ where } \mu^* = \max_{a \in \mathcal{A}} \mu_a.$$

Both means and distributions are *unknown*, which makes the problem non trivial, and the learner only knows that $\nu \in \mathcal{D}$ where \mathcal{D} is a given set of bandit configurations. This problem received increased attention in the middle of the 20th century, and the seminal paper Lai and Robbins [1985] established the first lower bound on the cumulative regret, showing that designing an algorithm that is optimal uniformly over a given set of configurations \mathcal{D} comes with a price. The study of the lower performance bounds in multi-armed bandits successfully lead to the development of asymptotically optimal algorithms for specific configuration sets, such as the KLUCB algorithm [Lai, 1987, Cappé et al., 2013, Maillard, 2018] for exponential families, or alternatively the DMED and IMED algorithms from Honda and Takemura [2011, 2015]. The lower bounds from Lai and Robbins [1985], later extended by Burnetas and Katehakis [1997] did not cover all possible configurations, and in particular *structured* configuration sets were not handled until Agrawal et al. [1989] and then Graves and Lai [1997] established generic lower bounds. Here, structure refers to the fact that pulling an arm

may reveals information that enables to refine estimation of other arms. Unfortunately, designing numerical efficient algorithms that are provably optimal remains a challenge for many structures.

Structured configurations. Motivated by the growing popularity of bandits in a number of industrial and societal application domains, the study of *structured configuration sets* has received increasing attention over the last few years: The linear bandit problem is one typical illustration Abbasi-Yadkori et al. [2011], Srinivas et al. [2010], Durand et al. [2017], for which the linear structure considerably modifies the achievable lower bound, see Lattimore and Szepesvari [2017]. The study of a *unimodal* structure naturally appears in many contexts, e.g. single-peak preference economics, voting theory or wireless communications, and has been first considered in Yu and Mannor [2011] from a bandit perspective, then in Combes and Proutiere [2014] and Trinh et al. [2020] providing an explicit lower bound together with an algorithm exploiting this specific structure. Other structures include Lipschitz bandits Magureanu et al. [2014], and we refer to the manuscript Magureanu [2018] for other examples, such as cascading bandits that are useful in the context of recommender systems. In Combes et al. [2017], a generic algorithm is introduced called OSSB (Optimal Structured Stochastic Bandit), stepping the path towards generic multi-armed bandit algorithms that are adaptive to a given structure. More recently in Degenne et al. [2020], the authors introduce an adaptation of the KLUCB strategy to handle structured multi-armed bandit problems.

Unimodal-structure. We assume a *unimodal* structure similar to that considered in Yu and Mannor [2011] and Combes and Proutiere [2014]. That is, there exists an undirected graph $G = (\mathcal{A}, E)$ whose vertices are arms \mathcal{A} , and whose edges E characterize a partial order among means $(\mu_a)_{a \in \mathcal{A}}$. This partial order is assumed unknown to the learner. We assume that there exists a unique optimal arm $a^* = \operatorname{argmax}_{a \in \mathcal{A}} \mu_a$ and that for all sub-optimal arm $a \neq a^*$, there exists a path $P_a = (a_1 = a, \dots, a_{\ell_a} = a^*) \in \mathcal{A}^{\ell_a}$ of length $\ell_a \geq 2$ such that for all $i \in [1, \ell_a - 1]$, $(a_i, a_{i+1}) \in E$ and $\mu_{a_i} < \mu_{a_{i+1}}$. Lastly, we assume that $\nu \subset \mathcal{P} := \{p(\mu), \mu \in \Theta\}$, where $p(\mu)$ is an exponential-family distribution probability with density $f(\cdot, \mu)$ with respect to some positive measure λ on \mathbb{R} and mean $\mu \in \Theta \subset \mathbb{R}$. \mathcal{P} is assumed to be known to the learner. Thus, for all $a \in \mathcal{A}$ we have $\nu_a = p(\mu_a)$. We denote by $\mathcal{D}_{(\mathcal{P}, G)}$ or simply \mathcal{D} the structured set of such unimodal-bandit distributions characterized by (\mathcal{P}, G) . In the following, we assume that \mathcal{P} is a set of one-dimensional exponential family distributions.

Contributions. In this paper, we provide novel regret minimization results related to the unimodal structure. We first revisit the Indexed Minimum Empirical Divergence (IMED) algorithm from Honda and Takemura [2015] introduced for unstructured multi-armed bandits, and adapt it to the unimodal-structured setting. We introduce in Section 3 the IMED-UB algorithm that is limited to the pulling of the current best arm or their no more than d nearest arms at each time step, with d the maximum degree of nodes in G . Being constructed from IMED, IMED-UB does not require any optimization procedure and does not separate exploration from exploitation rounds. IMED-UB appears to be a *local* algorithm. We prove in Theorem 6 that IMED-UB is asymptotically optimal. Furthermore, this novel algorithm competes with the state-of-the-art algorithms in practice. This is confirmed by numerical illustrations on synthetic data. We believe that the construction of this algorithm together with the proof techniques developed in this paper are of independent interest for the bandit community.

Notations. Let $\nu \in \mathcal{D}$. Let $\mu^* = \max_{a \in \mathcal{A}} \mu_a$ be the optimal mean and $a^* = \operatorname{argmax}_{a \in \mathcal{A}} \mu_a$ be the optimal arm of ν . We define for an arm $a \in \mathcal{A}$ its sub-optimality gap $\Delta_a = \mu^* - \mu_a$. Considering an horizon $T \geq 1$, thanks to the chain rule we can rewrite the regret as follows:

$$R(\nu, T) = \sum_{a \in \mathcal{A}} \Delta_a \mathbb{E}_\nu [N_a(T)], \quad (1)$$

where $N_a(t) = \sum_{s=1}^t \mathbb{I}_{\{a_s = a\}}$ is the number of pulls of arm a at time t .

2 Regret lower bound

In this subsection, we recall for completeness the known lower bound on the regret when we assume a unimodal structure. In order to obtain non trivial lower bound we consider algorithms that are *consistent* (aka uniformly-good).

Definition 1 (Consistent algorithm) An algorithm is consistent on \mathcal{D} if for all configuration $\nu \in \mathcal{D}$, for all sub-optimal arm a , for all $\alpha > 0$,

$$\lim_{T \rightarrow \infty} \mathbb{E}_\nu \left[\frac{N_a(T)}{T^\alpha} \right] = 0.$$

We can derive from the notion of consistency an asymptotic lower bound on the regret, see Combes and Proutiere [2014].

Proposition 2 (Lower bounds on the regret) Let us consider a consistent algorithm. Then, for all configuration $\nu \in \mathcal{D}$, it must be that

$$\liminf_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} \geq c(\nu) := \sum_{a \in \mathcal{V}_{a^*}} \frac{\Delta_a}{\text{KL}(\mu_a | \mu^*)},$$

where $\text{KL}(\mu | \mu') = \int_{\mathbb{R}} \log(f(x, \mu)/f(x, \mu')) f(x, \mu) \lambda(dx)$ denotes the Kullback-Leibler divergence between $\nu = p(\mu)$ and $\nu' = p(\mu')$, for $\mu, \mu' \in \Theta$.

Remark 3 The quantity $c(\nu)$ is a fully explicit function of ν (it does not require solving any optimization problem) for some set of distributions ν (see Remark 4). This useful property no longer holds in general for arbitrary structures. Also, it is noticeable that $c(\nu)$ does not involve all the sub-optimal arms but only the ones in \mathcal{V}_{a^*} . This indicates that sub-optimal arms outside \mathcal{V}_{a^*} are sampled $o(\log(T))$, which contrasts with the unstructured stochastic multi-armed bandits. See Combes and Proutiere [2014] for further insights.

Remark 4 For Bernoulli distributions, a possible setting is to assume $\lambda = \delta_0 + \delta_1$ (with δ_0, δ_1 Dirac measures), $\Theta = (0, 1)$ and for $\mu \in \Theta$, $f(\cdot, \mu) =: x \in \{0, 1\} \mapsto \mu^x (1 - \mu)^{1-x}$. Then for all $\mu, \mu' \in (0, 1)$, $\text{KL}(\mu | \mu') = \mu \log(\mu/\mu') + (1 - \mu) \log((1 - \mu)/(1 - \mu'))$. For Gaussian distributions (variance $\sigma^2 = 1$), we assume λ to be the Lebesgue measure, $\Theta = \mathbb{R}$, and for $\mu \in \mathbb{R}$, $f(\cdot, \mu) =: x \in \mathbb{R} \mapsto (\sqrt{2\pi})^{-1} e^{-(x - \mu)^2/2}$. Then for all $\mu, \mu' \in \mathbb{R}$, $\text{KL}(\mu | \mu') = (\mu' - \mu)^2/2$. For Exponential distributions, we assume λ to be the Lebesgue measure, $\Theta =]0; +\infty[$, and for $\mu > 0$, $f(\cdot, \mu) =: x > 0 \mapsto e^{-x/\mu}/\mu$. Then for all $\mu, \mu' > 0$, $\text{KL}(\mu | \mu') = \log(\mu'/\mu) + \mu/\mu' - 1$.

3 Optimal algorithm for unimodal-structured bandits

We present in this section a novel algorithm that matches the asymptotic lower bound of Proposition 2. This algorithm is inspired by the Indexed Minimum Empirical Divergence (IMED) proposed by Honda and Takemura [2011]. The general idea behind this algorithm is, following the intuition given by the lower bound, to narrow on the current best arm and its neighbourhood for pulling an arm at a given time step.

Notations. The empirical mean of the rewards from the arm a is denoted by $\hat{\mu}_a(t) = \sum_{s=1}^t \mathbb{I}_{\{a_s=a\}} X_s / N_a(t)$ if $N_a(t) > 0$, 0 otherwise. We also denote by $\hat{\mu}^*(t) = \max_{a \in \mathcal{A}} \hat{\mu}_a(t)$ and $\hat{\mathcal{A}}^*(t) = \arg \max_{a \in \mathcal{A}} \hat{\mu}_a(t)$ respectively the current best mean and the current set of optimal arms.

3.1 The IMED-UB algorithm.

We first pull each arm once. For all arm $a \in \mathcal{A}$ and time step $t \geq 1$ we introduce the IMED index

$$I_a(t) = N_a(t) \text{KL}(\hat{\mu}_a(t) | \hat{\mu}^*(t)) + \log(N_a(t)),$$

with the convention $0 \times \infty = 0$. This index can be seen as a transportation cost for moving a sub-optimal arm to an optimal one plus an exploration term: the logarithm of the number of pulls. When an optimal arm is considered, the transportation cost is null and there is only the exploration part. Note that, as stated in Honda and Takemura [2011], $I_a(t)$ is an index in the weaker sense since it cannot be determined only by samples from the arm a but also uses the empirical mean of the current optimal arm. We define IMED-UB (Indexed Minimum Empirical Divergence for Unimodal Bandits), described in Algorithm 1, to be the algorithm consisting of pulling an arm $a_t \in \{\hat{a}_t^*\} \cup \mathcal{V}_{\hat{a}_t^*}$ with

minimum index at each time step t , where is $\hat{a}_t^* \in \operatorname{argmin}_{\hat{a}^* \in \hat{\mathcal{A}}^*(t)} N_{\hat{a}^*}(t)$ is a current best arm. This is a natural algorithm since the lower bound on the regret given in Proposition 2 involves only the arms in \mathcal{V}_{a^*} , the neighbourhood of the arm a^* of maximal mean.

Algorithm 1 IMED-UB

Pull each arm once
for $t = |\mathcal{A}| \dots T - 1$ **do**
 Choose $\hat{a}_t^* \in \operatorname{argmin}_{\hat{a}^* \in \hat{\mathcal{A}}^*(t)} N_{\hat{a}^*}(t)$ (chosen arbitrarily)
 Pull $a_{t+1} \in \operatorname{argmin}_{a \in \{\hat{a}_t^*\} \cup \mathcal{V}_{\hat{a}_t^*}} I_a(t)$ (chosen arbitrarily)
end for

3.2 Asymptotic optimality of IMED-UB

In this section, we state the main theoretical result of this paper.

Theorem 5 (Upper bounds) *Let us consider a set of distributions $\nu \in \mathcal{D}$ and let a^* its optimal arm. Let \mathcal{V}_{a^*} be the sub-optimal arms in the neighbourhood of a^* . Then under the IMED-UB algorithm for all $0 < \varepsilon < \varepsilon_\nu$, for all horizon time $T \geq 1$, for all $a \in \mathcal{V}_{a^*}$,*

$$\mathbb{E}_\nu [N_a(T)] \leq \frac{1 + \alpha_\nu(\varepsilon)}{\operatorname{KL}(\mu_a | \mu_{a^*})} \log(T) + 2d C_\varepsilon \sqrt{\log(c_\varepsilon T)} + d(1 + c_{\varepsilon_\nu}^{-1}) + d(2d + 3) \frac{2\sigma_{\varepsilon_\nu}^2 e^{\varepsilon_\nu^2 / 2\sigma_\varepsilon^2}}{\varepsilon^2} + 1$$

and, for all $a \notin \{a^*\} \cup \mathcal{V}_{a^*}$,

$$\mathbb{E}_\nu [N_a(T)] \leq 2d C_\varepsilon \sqrt{\log(c_\varepsilon T)} + d(1 + c_{\varepsilon_\nu}^{-1}) + d(2d + 3) \frac{2\sigma_{\varepsilon_\nu}^2 e^{\varepsilon_\nu^2 / 2\sigma_\varepsilon^2}}{\varepsilon^2} + 1,$$

where d is the maximum degree of nodes in G , $\varepsilon_\nu = \min_{a \neq a'} |\mu_a - \mu_{a'}| / 2$,

$\sigma_\varepsilon^2 = \max_{a \in \mathcal{A}} \left\{ \mathbb{V}_{X \sim p(\mu^*)}(X) : \mu^* \in [\mu_a - \varepsilon, \mu_a] \right\}$ and $c_\varepsilon, C_\varepsilon > 0$ are the constants involved in Theorem 15. $\alpha_\nu(\cdot)$ is a non-negative function depending only on ν such that $\lim_{\varepsilon \rightarrow 0} \alpha_\nu(\varepsilon) = 0$ (see Section 4.1 for more details).

In particular one can note that the arms in the neighbourhood of the optimal one are pulled $\mathcal{O}(\log(T))$ times while the other sub-optimal arms are pulled $\mathcal{O}(\sqrt{\log(T)})$ of times under IMED-UB. This is coherent with the lower bound that only involves the neighbourhood of the best arm. More precisely, combining Theorem 5 and the chain rule (1) gives the asymptotic optimality of IMED-UB with respect to the lower bound of Proposition 2.

Corollary 6 (Asymptotic optimality) *With the same notations as in Theorem 5, then under the IMED-UB algorithm*

$$\limsup_{T \rightarrow \infty} \frac{R(\nu, T)}{\log(T)} \leq c(\nu) = \sum_{a \in \mathcal{V}_{a^*}} \frac{\Delta_a}{\operatorname{KL}(\mu_a | \mu_{a^*})}.$$

A finite time analysis of IMED-UB is provided in following Section 4.

4 IMED-UB finite time analysis

At a high level, the key interesting step of the proof is to realize that the considered algorithm implies empirical lower and empirical upper bounds on the numbers of pulls (see Lemma 7, Lemma 8). Then, based on concentration lemmas (see Section B), the algorithm-based empirical lower bounds ensure the reliability of the estimators of interest (Lemma 12). Interestingly, this makes use of arguments based on recent concentration of measure that enable to control the concentration without adding some $\log \log$ bonus (such a bonus was required for example in the initial analysis of the KL-UCB strategy from Cappé et al. [2013]). Then, combining the reliability of these estimators with the obtained algorithm-base empirical upper bounds, we obtain upper bounds on the average numbers of pulls (Theorem 5). The proof is concise to fit mostly in the next few pages.

4.1 Notations

Let us consider $\nu \in \mathcal{D}$ and let us denote by a^* its best arm. We recall that for all $a \in \mathcal{A}$, $\mathcal{V}_a = \{a' \in \mathcal{A} : (a, a') \in E\}$ is the neighbourhood of arm a in graph $G = (\mathcal{A}, E)$, and that

$$d = \max_{a \in \mathcal{A}} |\mathcal{V}_a|, \quad \varepsilon_\nu = \min_{a \neq a'} \frac{|\mu_a - \mu_{a'}|}{2}. \quad (2)$$

Then, there exists a function $\alpha_\nu(\cdot)$ such that for all $0 < \varepsilon < \varepsilon_\nu$, for all $a \neq a^*$,

$$\text{KL}(\mu_a + \varepsilon | \mu^* - \varepsilon) \leq (1 + \alpha_\nu(\varepsilon))^{-1} \text{KL}(\mu_a | \mu^*) \quad (3)$$

and $\lim_{\varepsilon \downarrow 0} \alpha_\nu(\varepsilon) = 0$. At each time step $t \geq 1$, \hat{a}_t^* is arbitrarily chosen in $\operatorname{argmin}_{a \in \hat{\mathcal{A}}^*(t)} N_a(t)$ where

$$\hat{\mathcal{A}}^*(t) = \operatorname{argmax}_{a \in \mathcal{A}} \hat{\mu}_a(t).$$

4.2 Algorithm-based empirical bounds

The IMED-UB algorithm implies inequalities between the indexes that can be rewritten as inequalities on the numbers of pulls. While lower bounds involving $\log(t)$ may be expected in view of the asymptotic regret bounds, we show lower bounds on the numbers of pulls involving instead $\log(N_{a_{t+1}}(t))$, the logarithm of the number of pulls of the current chosen arm. We also provide upper bounds on $N_{a_{t+1}}(t)$ involving $\log(t)$.

We believe that establishing these empirical lower and upper bounds is a key element of our proof technique, that is of independent interest and not *a priori* restricted to the unimodal structure.

Lemma 7 (Empirical lower bounds) *Under IMED-UB, at each step time $t \geq |\mathcal{A}|$, for all $a \in \mathcal{V}_{\hat{a}_t^*}$,*

$$\log(N_{a_{t+1}}(t)) \leq N_a(t) \text{KL}(\hat{\mu}_a(t) | \hat{\mu}^*(t)) + \log(N_a(t)) \quad (4)$$

and

$$N_{a_{t+1}}(t) \leq N_{\hat{a}_t^*}(t). \quad (5)$$

Proof For $a \in \mathcal{A}$, by definition, we have $I_a(t) = N_a(t) \text{KL}(\hat{\mu}_a(t) | \hat{\mu}^*(t)) + \log(N_a(t))$, hence

$$\log(N_a(t)) \leq I_a(t).$$

This implies, since the arm with minimum index is pulled, $\log(N_{a_{t+1}}(t)) \leq I_{a_{t+1}}(t) = \min_{a' \in \{\hat{a}_t^*\} \cup \mathcal{V}_{\hat{a}_t^*}} I_{a'}(t) \leq I_{\hat{a}_t^*}(t) = \log(N_{\hat{a}_t^*}(t))$. By taking the $\log^{-1}(\cdot)$, the last inequality allows us to conclude. \blacksquare

Lemma 8 (Empirical upper bounds) *Under IMED-UB at each step time $t \geq |\mathcal{A}|$,*

$$N_{a_{t+1}}(t) \text{KL}(\hat{\mu}_{a_{t+1}}(t) | \hat{\mu}^*(t)) \leq \log(t). \quad (6)$$

Proof As above, by construction we have

$$I_{a_{t+1}}(t) \leq I_{\hat{a}_t^*}(t).$$

It remains, to conclude, to note that

$$N_{a_{t+1}}(t) \text{KL}(\hat{\mu}_{a_{t+1}}(t) | \hat{\mu}^*(t)) \leq I_{a_{t+1}}(t),$$

and

$$I_{\hat{a}_t^*}(t) = \log(N_{\hat{a}_t^*}(t)) \leq \log(t). \quad \blacksquare$$

4.3 Non-reliable current means

For all arms $a, a' \in \mathcal{A}$ and for all accuracy $\varepsilon > 0$, let $\mathcal{E}_{a,a'}^+(\varepsilon)$ be the set of times where the current mean of arm a ε -deviates from above while arm a has more pulls than the current pulled arm a' ,

$$\mathcal{E}_{a,a'}^+(\varepsilon) := \{t \in \llbracket 1, T-1 \rrbracket : a_{t+1} = a', N_{a'}(t) \leq N_a(t), \widehat{\mu}_a(t) \geq \mu_a + \varepsilon\}. \quad (7)$$

We similarly define

$$\mathcal{E}_{a,a'}^-(\varepsilon) := \{t \in \llbracket 1, T-1 \rrbracket : a_{t+1} = a', N_{a'}(t) \leq N_a(t), \widehat{\mu}_a(t) \leq \mu_a - \varepsilon\}. \quad (8)$$

We also define

$$\mathcal{E}_{a,a'}(\varepsilon) = \mathcal{E}_{a,a'}^+(\varepsilon) \cup \mathcal{E}_{a,a'}^-(\varepsilon). \quad (9)$$

Definition 9 (KL-log deviation) For $\varepsilon > 0$, the couple of arms $(a, a') \in \mathcal{A}^2$ shows ε^- -KL-log deviation at time step $t \geq 1$ if the following conditions are satisfied

- (1) $a_{t+1} = a'$
- (2) $\widehat{\mu}_a(t) \leq \mu_a - \varepsilon$
- (3) $\log(N_{a'}(t)) \leq N_a(t) \text{KL}(\widehat{\mu}_a(t) | \mu_a - \varepsilon) + \log(N_a(t))$.

For all couple of arms $(a, a') \in \mathcal{A}^2$ and for all accuracy $\varepsilon > 0$, let $\mathcal{K}_{a,a'}^-(\varepsilon)$ be the set of times where couple of arms (a, a') shows ε^- -KL-log deviation, that is

$$\mathcal{K}_{a,a'}^-(\varepsilon) := \left\{ t \in \llbracket 1, T-1 \rrbracket : \begin{array}{l} (1) \quad a_{t+1} = a' \\ (2) \quad \widehat{\mu}_a(t) \leq \mu_a - \varepsilon \\ (3) \quad \log(N_{a'}(t)) \leq N_a(t) \text{KL}(\widehat{\mu}_a(t) | \mu_a - \varepsilon) + \log(N_a(t)) \end{array} \right\}. \quad (10)$$

We note that

$$\mathcal{E}_{a,a'}^-(\varepsilon) \subset \mathcal{K}_{a,a'}^-(\varepsilon).$$

We can now resort to concentration arguments in order to control the size of these sets, which yields the following upper bounds. We defer the proof to Appendix A.1.

Lemma 10 (Bounded subsets of times) For $\varepsilon > 0$, for $(a, a') \in \mathcal{A}^2$,

$$\begin{aligned} \mathbb{E}_\nu \left[\left| \mathcal{E}_{a,a'}^+(\varepsilon) \right| \right], \mathbb{E}_\nu \left[\left| \mathcal{E}_{a,a'}^-(\varepsilon) \right| \right] &\leq \frac{2\sigma_\varepsilon^2 e^{\varepsilon^2/2\sigma_\varepsilon^2}}{\varepsilon^2} \\ \mathbb{E}_\nu \left[\left| \mathcal{K}_{a,a'}^-(\varepsilon) \setminus \mathcal{E}_{a,a'}^-(\varepsilon) \right| \right] &\leq 1 + c_\varepsilon^{-1} + 2C_\varepsilon \sqrt{\log(c_\varepsilon T)}, \end{aligned}$$

where $\sigma_\varepsilon^2 = \max_{a \in \mathcal{A}} \left\{ \mathbb{V}_{X \sim p(\mu')} (X) : \mu' \in [\mu_a - \varepsilon, \mu_a] \right\}$, $c_\varepsilon, C_\varepsilon > 0$ are the constants involved in Theorem 15.

4.4 Non-reliable current best arm

For accuracy $\varepsilon > 0$, let $\mathcal{M}^*(\varepsilon)$ be the set of times $t \geq 1$ that do not belong to $\mathcal{E}_{\widehat{a}_t^*, a_{t+1}}^+(\varepsilon)$ and where some of the current best arm \widehat{a}_t^* differs from a^* ,

$$\mathcal{M}^*(\varepsilon) := \left\{ t \geq |\mathcal{A}| : \begin{array}{l} (1) \quad t \notin \mathcal{E}_{\widehat{a}_t^*, a_{t+1}}^+(\varepsilon) \\ (2) \quad \widehat{a}_t^* \neq a^* \end{array} \right\}. \quad (11)$$

Lemma 11 (Relation between subsets of times) Under IMED-UB, for all accuracy $0 < \varepsilon < \varepsilon_\nu = \min_{a \neq a'} |\mu_a - \mu_{a'}|/2$,

$$\mathcal{M}^*(\varepsilon) \subset \bigcup_{\substack{t \geq 1 \\ a \in \mathcal{V}_{\widehat{a}_t^*}}} \mathcal{K}_{a, a_{t+1}}^-(\varepsilon_\nu). \quad (12)$$

Proof Let us consider $t \in \mathcal{M}^*(\varepsilon)$. Since $\hat{a}_t^* \neq a^*$, there exists $a \in \mathcal{V}_{\hat{a}_t^*}$ such that

$$\mu_a > \mu_{\hat{a}^*}. \quad (13)$$

Then, since $\hat{a}_t^* \in \operatorname{argmax}_{a \in \mathcal{A}} \hat{\mu}_a(t)$, we have

$$\hat{\mu}_{\hat{a}_t^*}(t) = \hat{\mu}^*(t) \geq \hat{\mu}_a(t). \quad (14)$$

Since $t \in \mathcal{M}^*(\varepsilon)$, $t \notin \mathcal{E}_{\hat{a}_t^*, a_{t+1}}^+(\varepsilon)$. By considering empirical lower bounds (5) and Equation (7), we have

$$\mu_{\hat{a}_t^*} + \varepsilon \geq \hat{\mu}_{\hat{a}_t^*}(t). \quad (15)$$

By combining Equations (14) and (15), it comes

$$\mu_{\hat{a}_t^*} + \varepsilon \geq \hat{\mu}^*(t) \geq \hat{\mu}_a(t). \quad (16)$$

Since $\varepsilon < \varepsilon_\nu \leq |\mu_a - \mu_{\hat{a}_t^*}|/2$, Equation (13) and previous Equation (16) imply

$$\mu_a - \varepsilon_\nu > \hat{\mu}_{\hat{a}_t^*}(t) \geq \hat{\mu}_a(t). \quad (17)$$

Since $a \in \mathcal{V}_{\hat{a}_t^*}$, empirical lower bounds (4) imply

$$\log(N_{a_{t+1}}(t)) \leq N_a(t) \operatorname{KL}(\hat{\mu}_a(t) | \hat{\mu}^*(t)) + \log(N_a(t)). \quad (18)$$

The classical monotonic properties of $\operatorname{KL}(\cdot | \cdot)$ and Equation (17) imply

$$\begin{cases} \hat{\mu}_a(t) < \mu_a - \varepsilon_\nu \\ \operatorname{KL}(\hat{\mu}_a(t) | \hat{\mu}^*(t)) \leq \operatorname{KL}(\hat{\mu}_a(t) | \mu_a - \varepsilon_\nu). \end{cases} \quad (19)$$

Combining Equations (17) and (19), we get

$$\begin{cases} \hat{\mu}_a(t) < \mu_a - \varepsilon_\nu \\ \log(N_{a_{t+1}}(t)) \leq N_a(t) \operatorname{KL}(\hat{\mu}_a(t) | \mu_a - \varepsilon_\nu) + \log(N_a(t)), \end{cases} \quad (20)$$

which means $t \in \mathcal{K}_{a, a_{t+1}}^-(\varepsilon_\nu)$. ■

4.5 Reliable current means and current best arm

In this subsection, we characterize subsets of times where both the mean of current pulled arm and the optimal mean are well-estimated.

Let us consider for $0 < \varepsilon < \varepsilon_\nu$, for $a \neq a^*$,

$$\mathcal{U}_a(\varepsilon) = \{t \geq |\mathcal{A}| : a_{t+1} = a\} \cap \left(\bigcup_{t \geq 1} \mathcal{E}_{a_{t+1}, a_{t+1}}^+(\varepsilon) \cup \mathcal{E}_{\hat{a}_t^*, a_{t+1}}^-(\varepsilon) \cup \mathcal{E}_{\hat{a}_t^*, a_{t+1}}^+(\varepsilon) \cup \mathcal{M}^*(\varepsilon) \right). \quad (21)$$

Then, Lemma 11 implies

$$\mathcal{U}_a(\varepsilon) \subset \bigcup_{\substack{a' \in \{a\} \cup \mathcal{V}_a \\ a'' \in \mathcal{V}_{a'}}} \mathcal{E}_{a', a}^+(\varepsilon) \cup \mathcal{E}_{a'', a}^-(\varepsilon) \cup \mathcal{K}_{a'', a}^-(\varepsilon_\nu). \quad (22)$$

In particular, from Lemma 10 and previous Equation (22) we have

$$\begin{aligned} \mathbb{E}_\nu[\mathcal{U}_a(\varepsilon)] &\leq 2d(d+1) \frac{2\sigma_\varepsilon^2 e^{\varepsilon^2/2\sigma_\varepsilon^2}}{\varepsilon^2} + d \left(\frac{2\sigma_{\varepsilon_\nu}^2 e^{\varepsilon_\nu^2/2\sigma_{\varepsilon_\nu}^2}}{\varepsilon_\nu^2} + 1 + c_{\varepsilon_\nu}^{-1} + 2C_\varepsilon \sqrt{\log(c_\varepsilon T)} \right) \\ &\leq d(2d+3) \frac{2\sigma_{\varepsilon_\nu}^2 e^{\varepsilon_\nu^2/2\sigma_{\varepsilon_\nu}^2}}{\varepsilon^2} + d \left(1 + c_{\varepsilon_\nu}^{-1} + 2C_\varepsilon \sqrt{\log(c_\varepsilon T)} \right), \end{aligned} \quad (23)$$

where $d = \max_{a \in \mathcal{A}} |\mathcal{V}_a|$ is the maximum degree of nodes in graph \mathcal{G} .

Lemma 12 (Reliable current means) *Under IMED-UB, for all accuracy $0 < \varepsilon < \varepsilon_\nu = \min_{a \neq a^*} |\mu_a - \mu_{a^*}|/2$, for all sub-optimal arm $a \neq a^*$, for all time step $t \notin \mathcal{U}_a(\varepsilon)$, $t \geq |\mathcal{A}|$, such that $a_{t+1} = a$,*

$$\begin{cases} \hat{a}_t^* = a^* \\ \hat{\mu}^*(t) \geq \mu^* - \varepsilon \\ \hat{\mu}_a(t) \leq \mu_a + \varepsilon. \end{cases}$$

4.6 Upper bounds on the numbers of pulls of sub-optimal arms

In this subsection, we now combine the different results of the previous subsections to prove Theorem 5.

Proof [Proof of Theorem 5.] For $0 < \varepsilon < \varepsilon_\nu$, for $a \neq a^*$, let us consider $t \notin \mathcal{U}_a(\varepsilon)$, $t \geq |\mathcal{A}|$, such that $a_{t+1} = a$. From empirical upper bounds (6), we have

$$N_a(t) \text{KL}(\widehat{\mu}_a(t) | \widehat{\mu}^*(t)) \leq \log(t). \quad (24)$$

From Lemma 12 and Algorithm 1, we have $a \in \mathcal{V}_{a^*}$ and $\widehat{\mu}_a(t) \leq \mu_a + \varepsilon < \mu^* - \varepsilon \leq \widehat{\mu}^*(t)$. From classical monotonic properties of $\text{KL}(\cdot | \cdot)$ and Equation (3), we have $\text{KL}(\widehat{\mu}_a(t) | \widehat{\mu}^*(t)) \geq \text{KL}(\mu_a + \varepsilon | \mu^* - \varepsilon) \geq (1 + \alpha_\nu(\varepsilon))^{-1} \text{KL}(\mu_a | \mu^*)$. In view of Equation (24), this implies

$$\forall t \notin \mathcal{U}_a(\varepsilon), t \geq |\mathcal{A}|, \text{ such that } a_{t+1} = a, \quad \begin{cases} a \in \mathcal{V}_{a^*} \\ N_a(t) \leq \frac{(1 + \alpha_\nu(\varepsilon)) \log(t)}{\text{KL}(\mu_a | \mu^*)}. \end{cases} \quad (25)$$

For all arm $a \in \mathcal{A}$, for all time step $t \geq |\mathcal{A}|$, we denote by

$$\tau_a(t) = \max \{t' \in [|\mathcal{A}|; t] : a_{t'+1} = a \text{ and } t' \notin \mathcal{U}_a(\varepsilon)\} \quad (26)$$

the last time step before time step t that does not belong to $\mathcal{U}_a(\varepsilon)$ such that we pull arm a .

Then, from Equations (25) and (26) we have

$$\begin{aligned} \forall a \neq a^*, \forall t \geq 1, \quad N_a(t) &= N_a(|\mathcal{A}|) + \sum_{t' \geq |\mathcal{A}|}^{t-1} \mathbb{I}_{\{a_{t'+1} = a\}} \\ &\leq 1 + \sum_{t' \geq 1}^{t-1} \mathbb{I}_{\{a_{t'+1} = a, t' \in \mathcal{U}_a(\varepsilon)\}} + \sum_{t' \geq |\mathcal{A}|}^{t-1} \mathbb{I}_{\{a_{t'+1} = a, t' \notin \mathcal{U}_a(\varepsilon)\}} \\ &\leq 1 + |\mathcal{U}_a(\varepsilon)| + \sum_{t' \geq |\mathcal{A}|}^{t-1} \mathbb{I}_{\{a_{t'+1} = a, t' \notin \mathcal{U}_a(\varepsilon)\}} \\ &\leq 1 + |\mathcal{U}_a(\varepsilon)| + \mathbb{I}_{\{a \notin \mathcal{V}_{a^*}\}} \times 0 + \mathbb{I}_{\{a \in \mathcal{V}_{a^*}\}} \times N_a(\tau_a(t)) \\ &\leq 1 + |\mathcal{U}_a(\varepsilon)| + \mathbb{I}_{\{a \in \mathcal{V}_{a^*}\}} \frac{(1 + \alpha_\nu(\varepsilon)) \log(\tau_a(t))}{\text{KL}(\mu_a | \mu^*)} \\ &\leq 1 + |\mathcal{U}_a(\varepsilon)| + \mathbb{I}_{\{a \in \mathcal{V}_{a^*}\}} \frac{(1 + \alpha_\nu(\varepsilon)) \log(t)}{\text{KL}(\mu_a | \mu^*)}. \end{aligned}$$

This implies

$$\forall a \neq a^*, \forall t \geq 1, \quad N_a(t) \leq \begin{cases} \frac{(1 + \alpha_\nu(\varepsilon)) \log(t)}{\text{KL}(\mu_a | \mu^*)} + |\mathcal{U}_a(\varepsilon)| + 1 & \text{if } a \in \mathcal{V}_{a^*} \\ |\mathcal{U}_a(\varepsilon)| + 1 & \text{if } a \notin \mathcal{V}_{a^*}. \end{cases} \quad (27)$$

From Equation (23), averaging these inequalities allows us to conclude. ■

5 Numerical experiments

In this section, we compare empirically the following algorithms : OSUB, UTS [Combes and Proutiere, 2014, Trinh et al., 2020] and IMED-UB described in Algorithm 1. We illustrate how performs the IMED-UB algorithm under Bernoulli, Gaussian (variance $\sigma^2 = 0.25$) or Exponential distribution assumption. For the experiments we consider a graph \mathcal{G} with maximal degree $d = 2$ and the unimodal unimodal vectors of means $\mu = (0.05, 0.10, 0.15, 0.20, 0.25, 0.20, 0.15, 0.10, 0.05)$, and average regrets over 500 runs for each distribution family. Based on these experiments (Figure 1), it seems that IMED-UB competes with OSUB and UTS.

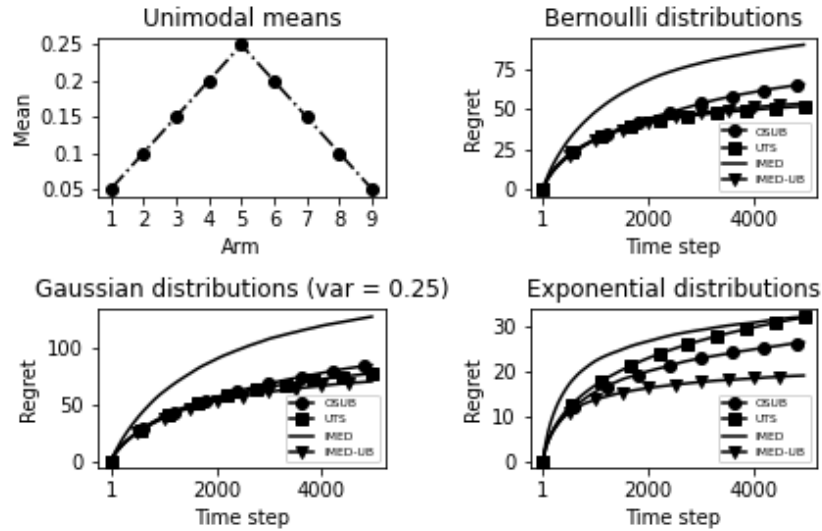


Figure 1: Cumulative regrets averaged over 500 runs.

Conclusion

In this paper, we have revisited the setup of unimodal multi-armed bandits: We introduced a novel variant based on the IMED algorithm. This algorithm does not separate exploration from exploitation rounds and is proven optimal for one-dimensional exponential family distributions. Remarkably, the IMED-UB algorithm do not require any optimization procedure, which can be interesting for practitioners. We also provided a novel proof algorithm, in which we make explicit empirical lower and upper bounds, before tackling the handling of bad events by specific concentration tools. This proof technique greatly simplifies and shorten the analysis of IMED-UB. Last, we provided numerical experiments that show the practical effectiveness of IMED-UB.

Acknowledgments

This work has been supported by the French Ministry of Higher Education and Research, Inria, the French Agence Nationale de la Recherche (ANR) under grant ANR-16-CE40-0002 (the BADASS project), the MEL, the I-Site ULNE regarding project R-PILOTE-19-004-APPRENF.

Pierre Ménard is supported by the SFI Sachsen-Anhalt for the project REBCI ZS/2019/10/102024 by the Investitionsbank Sachsen-Anhalt.

References

- Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- R. Agrawal, D. Teneketzis, and V. Anantharam. Asymptotically efficient adaptive allocation schemes for controlled iid processes: Finite parameter space. *IEEE Transactions on Automatic Control*, 34(3), 1989.
- A. N. Burnetas and M. N. Katehakis. Optimal adaptive policies for Markov decision processes. *Mathematics of Operations Research*, 22(1):222–255, 1997.
- O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, and G. Stoltz. Kullback–Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541, 2013.
- R. Combes and A. Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, 2014.
- R. Combes, S. Magureanu, and A. Proutiere. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2017.
- R. Degenne, H. Shao, and W. Koolen. Structure adaptive algorithms for stochastic bandits. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2443–2452. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/degenne20b.html>.
- A. Durand, O.-A. Maillard, and J. Pineau. Streaming kernel regression with provably adaptive mean, variance, and regularization. *arXiv preprint arXiv:1708.00768*, 2017.
- T. L. Graves and T. L. Lai. Asymptotically efficient adaptive choice of control laws in controlled markov chains. *SIAM journal on control and optimization*, 35(3):715–743, 1997.
- J. Honda and A. Takemura. An asymptotically optimal policy for finite support models in the multiarmed bandit problem. *Machine Learning*, 85(3):361–391, 2011.
- J. Honda and A. Takemura. Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. *Machine Learning*, 16:3721–3756, 2015.
- T. L. Lai. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, pages 1091–1114, 1987.
- T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- T. Lattimore and C. Szepesvári. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737, 2017.
- S. Magureanu. *Efficient Online Learning under Bandit Feedback*. PhD thesis, KTH Royal Institute of Technology, 2018.
- S. Magureanu, R. Combes, and A. Proutiere. Lipschitz bandits: Regret lower bounds and optimal algorithms. *Machine Learning*, 35:1–25, 2014.
- O.-A. Maillard. Boundary crossing probabilities for general exponential families. *Mathematical Methods of Statistics*, 27(1):1–31, 2018.
- H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58:527–535, 1952.
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pages 1015–1022. Omnipress, 2010.
- W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

- W. R. Thompson. On a criterion for the rejection of observations and the distribution of the ratio of deviation to sample standard deviation. *The Annals of Mathematical Statistics*, 6(4):214–219, 1935.
- C. Trinh, E. Kaufmann, C. Vernade, and R. Combes. Solving bernoulli rank-one bandits with unimodal thompson sampling. In *International Conference on Algorithmic Learning Theory*, 2020.
- J. Y. Yu and S. Mannor. Unimodal bandits. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 41–48, 2011.