
Qualitative Multi-Armed Bandits: A Quantile-Based Approach

Balázs Szörényi^{1,5,6}
Róbert Busa-Fekete²
Paul Weng^{3,4}
Eyke Hüllermeier²

SZORENYI@INF.U-SZEGED.HU
BUSAROBI@UPB.DE
PAWENG@CMU.EDU
EYKE@UPB.DE

¹INRIA Lille - Nord Europe, Sequel project, 40 avenue Halley, 59650 Villeneuve d'Ascq, France

²Department of Computer Science, University of Paderborn, Warburger Str. 100, 33098 Paderborn, Germany

³SYSU-CMU Joint Institute of Engineering, 132 East Waihuan Road, Guangzhou, 510006, China

⁴SYSU-CMU Shunde International Joint Research Institute, 9 Eastern Nanguo Road, Shunde, 528300, China

⁵MTA-SZTE Research Group on Artificial Intelligence, Tisza Lajos krt. 103., H-6720 Szeged, Hungary

⁶Department of Electrical Engineering, The Technion - Israel Institute of Technology, Haifa, Israel 32000

Abstract

We formalize and study the multi-armed bandit (MAB) problem in a generalized stochastic setting, in which rewards are not assumed to be numerical. Instead, rewards are measured on a qualitative scale that allows for comparison but invalidates arithmetic operations such as averaging. Correspondingly, instead of characterizing an arm in terms of the mean of the underlying distribution, we opt for using a quantile of that distribution as a representative value. We address the problem of quantile-based online learning both for the case of a finite (pure exploration) and infinite time horizon (cumulative regret minimization). For both cases, we propose suitable algorithms and analyze their properties. These properties are also illustrated by means of first experimental studies.

1. Introduction

The multi-armed bandit (MAB) problem (or simply bandit problem) refers to an iterative decision making problem in which an agent repeatedly chooses among K options, metaphorically corresponding to pulling one of K arms of a bandit machine. In each round, the agent receives a random reward that depends on the arm being selected. The agent's goal is to optimize an evaluation metric, e.g., the *error rate* (expected percentage of playing a suboptimal arm) or the *cumulative regret* (difference between the sum of rewards obtained and the (expected) rewards that could have been obtained by selecting the best arm in each round).

Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 2015. JMLR: W&CP volume 37. Copyright 2015 by the author(s).

In the *stochastic* multi-armed bandit setup, the distributions can vary with the arms but do not change with time. To achieve the desired goal, the agent has to tackle the classical exploration/exploitation dilemma: It has to properly balance the pulling of arms that were found to yield high rewards in earlier rounds and the selection of arms that have not yet been tested often enough (Auer et al., 2002; Cesa-Bianchi & Lugosi, 2006; Lai & Robbins, 1985).

MAB algorithms have not only been studied quite thoroughly from a theoretical point of view but have also been used in many real applications, such as medical treatment allocation design (Kuleshov & Precup, 2014), feature selection (Gaudel & Sebag, 2010; Busa-Fekete & Kégl, 2010) and crowdsourced labeling (Zhou et al., 2014). In many practical applications, however, numerical rewards are not provided in a natural way. Consider the example of clinical trials for testing pain medication. Here, the patients are asked to value their pain on a scale such as `no pain—mild—moderate—severe`, which is of qualitative nature. Computing averages on ordinal scales of that kind is clearly invalid and may lead to disputable conclusions.

In this paper, we therefore propose a setting in which the arm distributions are defined over a complete totally ordered set (L, \preceq) ; the corresponding online learning framework will be introduced formally in Section 3, after reviewing related work in Section 2. The quality of an arm is expressed in terms of a τ -quantile of the arm's distribution over L . Thus, arms are compared in terms of their τ -quantiles, and an arm is considered to be τ -optimal if its τ -quantile coincides with the highest τ -quantile of all arms.

We consider two quantile-based learning frameworks that we refer to as the *finite* and the *infinite horizon* cases, respectively. The finite horizon case (Section 4) is formalized in the PAC framework: the goal of the learner is to find a

τ -optimal arm with probability at least $1 - \delta$. As opposed to this, the infinite horizon case (Section 5) is formalized as a regret minimization problem, in which the regret depends on τ and the quantile functions of the arms. The difficulty of both setups stems from the fact that, when for all τ -optimal arms, the probability of getting qualitative rewards lower or equal to the optimal τ -quantile x^* is close or equal to τ , it is hard (or impossible) to guess x^* , which is essential to decide whether an arm is optimal or not.

2. Related Work

Pure exploration algorithms for the stochastic bandit problem sample the arms a certain number of times (not necessarily known in advance) and then output a recommendation, such as the best arm or the m best arms (Bubeck et al., 2009; Even-Dar et al., 2002; Bubeck et al., 2013; Gabillon et al., 2011; Cappé et al., 2013; Kalyanakrishnan et al., 2012). Since our quantile-based learning task in the finite horizon case is formulated in a PAC setting, it can be viewed as a pure exploration strategy, too. Yet, we do not assume that *absolute numerical* feedback can be generated for individual arms; instead, our feedback is of *qualitative* nature. Therefore, since averaging rewards is no longer meaningful, the preference relation over the arms is defined based on τ -quantiles instead of mean values of the underlying distribution on rewards.

Yu & Nikolova (2013) introduce a pure exploration setting where, instead of the means, the goodness value or payoff of the arms is defined based on some notion of risk, such as the *value-at-risk* (Schachter, 1997), a famous risk measure used in finance, which is a particular case of quantiles. Their setup is similar to the best arm identification problem (Bubeck et al., 2009), where the goal of the learner is to control the so-called simple regret, which is the difference between the payoff of the best arm and the expected payoff obtained by its recommendation. The algorithm proposed by Yu & Nikolova (2013) is based on their result concerning the concentration property of the estimators of various risk measures—these properties are preconditioned on the assumption that the density functions of arms are continuously differentiable, and their derivatives are bounded from above. The proposed algorithm is computationally demanding since it solves a non-linear constrained and integer-valued optimization task in each round; moreover, their results regarding the performance of the algorithm assume that the densities are bounded away from zero everywhere. In addition to these limitations on the reward distributions, our learning setup also differs from theirs in that we assume a PAC setting with finite horizon, where the goal is to find a τ -optimal arm with high probability. Thus, since the error of the learner is controlled, the algorithms are evaluated in terms of their sample complexity

(the number of samples taken prior to termination).

In the infinite case, the most commonly optimized property is the regret with respect to the maximum mean reward (Bubeck & Cesa-Bianchi, 2012). Nevertheless, alternative targets have already been considered, too, which led to interesting formal tasks. In a recent study, Carpentier & Valko (2014) formulate the regret in terms of the extreme value of the arm distributions. The goal here is to optimize the maximal regret observed. To this end, the learner intends to identify the most “abnormal” arm having the heaviest tail distribution, since the rewards received on a heavy-tailed arm are likely to deviate the most from its mean with highest probability. The learner is evaluated in terms of so-called extreme regret, which is the most extreme value found and compared to the most extreme value possible. The authors devise an algorithm, called EXTREMEHUNTER, based on the optimism in the face of uncertainty principle, which can achieve logarithmic expected regret in this setting. Sani et al. (2012) consider a MAB setting with a regret notion based on the principle of risk-aversion, where the risk is defined based on mean-variance risk. More specifically, there is a trade-off parameter that controls the influence of the mean and variance of the arm distributions. Thus, pulling an arm with high mean might result in a high regret if the variance of the rewards is high. The worst case regret of the proposed algorithm is $\mathcal{O}(KT^{2/3})$, and it is not clear whether it can be improved. As it is known that the worst case regret for the standard mean-based regret is $\mathcal{O}(\sqrt{KT})$, which is achieved by the MOSS algorithm by Audibert & Bubeck (2010), the optimization of regret based on risk aversion is conjectured to be a more complex problem.

In the preference-based bandit setup (Busa-Fekete & Hüllermeier, 2014), also known as duelling bandits (Yue et al., 2012), feedback about arms is not received in terms of absolute numerical rewards either. Instead, the learner is only allowed to compare the arms in a pairwise manner and receives binary feedback informing about the winner of the comparison. From this point of view, the feedback about the arms is even weaker than in our qualitative setting. Moreover, the notion of optimality of an arm can be defined in various ways in the preference-based setup. For example, a commonly used notion is that of a *Condorcet winner*, for which the probability of winning in a pairwise comparison is larger than $1/2$ against each other arm (Yue & Joachims, 2011; Zoghi et al., 2014).

3. Qualitative Multi-Armed Bandits

Formally, a standard *value-based* multi-armed or K -armed bandit problem is specified by real-valued random variables X_1, \dots, X_K associated, respectively, with K arms (that we simply identify by the numbers $1, \dots, K$). In each

time step t , the online learner selects one or more arms (depending on the specific problem) and obtains a random sample of the corresponding distributions. These samples, which are called rewards, are assumed to be independent of all previous actions and rewards. The goal of the learner can be defined in different ways, such as maximizing the sum of rewards over time (Lai & Robbins, 1985; Auer et al., 2002) or identifying, with high probability, an arm the expected value of which differs by at most ϵ from the highest one (Even-Dar et al., 2002; Kalyanakrishnan et al., 2012).

In the *qualitative* multi-armed bandit (QMAB) problem, the rewards are not necessarily real-valued. Instead, the arms X_1, \dots, X_K are random variables over a complete totally ordered set¹ (L, \preceq) . Accordingly, when arm k is played in the t -th round, it yields a qualitative payoff $x \in L$. Independence is again assumed to hold across rounds. We will also use the reverse order \succeq over L and the associated asymmetric (hence irreflexive) relations \prec and \succ of \preceq and \succeq , respectively. For simplicity, we shall assume that L is a subset of the real numbers and \preceq denotes the ordinary ranking over the reals. However, we shall not make use of the nominal values of the elements in L , only of their ordering.

3.1. Empirical CDF and Quantile Function

Let F^X denote the cumulative distribution function (CDF) of a random variable X . The *quantile function* $Q^X : [0, 1] \rightarrow L$ of X is defined as

$$Q^X(\tau) = \inf \{x \in L : \tau \leq F^X(x)\} .$$

We extend the domain of this function to the whole real line by defining $Q^X(\tau) = \inf L$ for $\tau < 0$ and $Q^X(\tau) = \sup L$ for $\tau > 1$.

As already mentioned, our aim is to compare arms in terms of their τ -quantiles, where $\tau \in [0, 1]$ is a (user-specified) parameter of the problem—the concrete learning tasks will be detailed in Sections 4 and 5, respectively. As will be seen, the highest τ -quantile of the arms, $x^* = \max_{1 \leq k \leq K} Q^{X_k}(\tau)$, will play a central role in both cases. The difficulty of our quantile-based approach is due to the fact that x^* is unknown and, moreover, hard to guess.²

Denote the j -th sample of arm k by $X_{k,j}$. The empirical estimate of the CDF (or empirical CDF) of X_k based on $X_{k,1}, \dots, X_{k,m}$ is

$$\widehat{F}_m^{X_k}(x) = \frac{1}{m} \sum_{j=1}^m \mathbb{I} \{X_{k,j} \preceq x\} ,$$

¹A totally ordered set is *complete* if every subset of it that has an upper bound also has a least upper bound.

²If x^* were known, the problem could be simplified to a standard value-based MAB with reward 1 in case the qualitative reward is at least as good as x^* and 0 otherwise.

where $\mathbb{I} \{\cdot\}$ is the indicator function. Denoting by $T_t(k)$ the number of times arm k has been pulled up to time t , the empirical CDF of X_k in round t is $\widehat{F}_{T_t(k)}^{X_k}(x)$.

The empirical estimator for the quantile function of arm k is based on the empirical distribution function:

$$\widehat{Q}_m^{X_k}(\tau) = \inf \left\{ x \in L : \tau \leq \widehat{F}_m^{X_k}(x) \right\}$$

The accuracy of these empirical estimates can be quantified using a concentration result of Dvoretzky et al. (1956), which upper-bounds the tail distribution of the deviation of the empirical cumulative distribution function in supremum norm. Its improved version (Massart, 1990) (having optimal constants) can be formulated in our case as follows:³

$$\mathbf{P} \left(\|F^{X_k} - \widehat{F}_m^{X_k}\|_\infty > c \right) \leq 2 \exp(-2mc^2) , \quad (1)$$

where $\|\cdot\|_\infty$ denotes the supremum norm.

For the sake of conciseness, we introduce an auxiliary function that determines the size of the confidence intervals:

$$c_m(\delta) = \sqrt{\frac{1}{2m} \log \frac{\pi^2 m^2}{3\delta}} \quad (2)$$

Proposition 1. *Fix some $1 \leq k \leq K$ and $\delta \in (0, 1)$. The following holds with probability at least $1 - \delta$: For all $m \geq 1$ and for every $0 \leq \tau \leq 1$,*

$$\widehat{Q}_m^{X_k}(\tau - c_m(\delta)) \preceq Q^{X_k}(\tau) \preceq \widehat{Q}_m^{X_k}(\tau + c_m(\delta)) \quad (3)$$

Proof. To simplify notations, denote F^{X_k} by F and $\widehat{F}_m^{X_k}$ by \widehat{F}_m . Combining the bound (1) with the uniform bound and the Basel problem one obtains that, with probability at least $(1 - \delta)$, $\|F - \widehat{F}_m\|_\infty \leq c_m(\delta)$ for all $m > 0$. In addition, $\|F - \widehat{F}_m\|_\infty \leq c_m(\delta)$ implies

$$\begin{aligned} Q(\tau) &= \inf \{x \in L : \tau \leq F(x)\} \\ &\preceq \inf \left\{ x \in L : \tau \leq \widehat{F}_m(x) - c_m(\delta) \right\} \\ &= \widehat{Q}_m(\tau + c_m(\delta)) \end{aligned}$$

and

$$\begin{aligned} \widehat{Q}_m(\tau - c_m(\delta)) &= \inf \left\{ x \in L : \tau \leq \widehat{F}_m(x) + c_m(\delta) \right\} \\ &\preceq \inf \{x \in L : \tau \leq F(x)\} \\ &= Q(\tau) \end{aligned}$$

□

³Each analysis in this paper also goes through using the Chernoff-Hoeffding bounds, essentially without any modification, at the cost of having slightly worse multiplicative constants (see Appendix C).

4. Finite Horizon: A PAC Algorithm

In this section, we consider the task of determining a “best” arm. In accordance with the goal highlighted in the introduction, the optimality of an arm is defined as follows.

Definition 1. An arm k is said to be τ -optimal if $Q^{X_k}(\tau) = x^*$, where

$$x^* = \max_{1 \leq k' \leq K} Q^{X_{k'}}(\tau). \quad (4)$$

Throughout this section, let k^* denote the index of a τ -optimal arm. Requiring the learner to output such an arm might be hard or even impossible to achieve in cases where the probability of getting qualitative rewards lower or equal to the optimal τ -quantile is close or equal to τ for all τ -optimal arms. Therefore, in the spirit of the PAC bandit setting introduced by Even-Dar et al. (2002), we are going to tolerate some approximation error.

Definition 2. An arm k is said to be (ϵ, τ) -optimal iff $Q^{X_k}(\tau + \epsilon) \geq x^*$. Put in words, a slight “negative” perturbation on the distribution of an (ϵ, τ) -optimal arm yields a τ -quantile that is higher or equal to x^* .

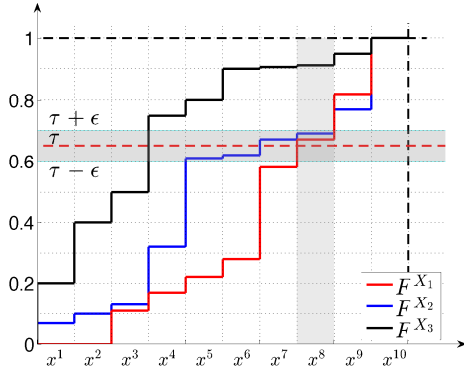


Figure 1. A qualitative MAB setup with three arms. The CDFs of the arms are plotted. The rewards come from $L = \{x^1, \dots, x^{10}\}$, where $x^1 \prec x^2 \prec \dots \prec x^{10}$.

Example 1. To illustrate the notion of (ϵ, τ) -optimality, consider the following simple qualitative MAB problem with three arms and parameters $\tau = 0.65$, $\epsilon = 0.05$. Each arm is a random variable over ten possible qualitative rewards $x^1 \prec x^2 \prec \dots \prec x^{10}$. Figure 1 depicts their cumulative distributions F^{X_1} , F^{X_2} and F^{X_3} . The τ -quantiles of arms 1, 2 and 3 are x^8 , x^7 and x^4 , respectively. The first arm (plotted in red) is τ -optimal, whence $k^* = 1$ and $x^* = x^8$. The second arm (plotted in blue) is not τ -optimal, since $F^{X_2}(x^7) > \tau$; yet, it is (ϵ, τ) -optimal since $F^{X_2}(x^7) - \epsilon < \tau$. The third arm (plotted in black) is not (ϵ, τ) -optimal, since the 0.7-quantile of X_3 is still given by $x^4 \prec x^8$.

In practice, there may be several τ -optimal and several

(ϵ, τ) -optimal arms. The goal of the learner is to identify one of them reliably.

Definition 3. An online learning algorithm is called (ϵ, τ, δ) -quantile learner if it outputs an (ϵ, τ) -optimal arm with probability at least $1 - \delta$.

4.1. The QPAC Algorithm

In Algorithm 1, we present our QPAC (Qualitative Probably Approximately Correct) algorithm, an adaptive elimination strategy inspired by Even-Dar et al. (2002). The algorithm computes lower and upper bounds x_t^- and x_t^+ of the optimal τ -quantile and exploits that (3) holds with high probability, which has several important consequences (as will be shown later in the analysis). One such consequence is that

$$x_t^- \preceq \widehat{Q}_t^{X_k}(\tau + \epsilon + c_t(\frac{\delta}{K})) \quad (5)$$

for all (ϵ, τ) -optimal arms k , and thus every arm h with $\widehat{Q}_t^{X_h}(\tau + \epsilon + c_t(\frac{\delta}{K})) \prec x_t^-$ can be eliminated. Another important consequence is that an arm k is necessarily (ϵ, τ) -optimal if

$$x_t^+ \preceq \widehat{Q}_t^{X_k}(\tau + \epsilon - c_t(\frac{\delta}{K})) \quad (6)$$

The rest will be detailed in the analysis.

Algorithm 1 QPAC(δ, ϵ, τ)

- 1: Set $\mathcal{A} = \{1, \dots, K\}$ ▷ Active arms
 - 2: $t = 1$
 - 3: **while** $\mathcal{A} \neq \emptyset$ **do**
 - 4: **for** $k \in \mathcal{A}$ **do**
 - 5: Pull arm k and observe $X_{k,t}$
 - 6: $x_t^+ = \max_{k \in \mathcal{A}} \widehat{Q}_t^{X_k}(\tau + c_t(\frac{\delta}{K}))$
 - 7: $x_t^- = \max_{k \in \mathcal{A}} \widehat{Q}_t^{X_k}(\tau - c_t(\frac{\delta}{K}))$
 - 8: **for** $k \in \mathcal{A}$ **do**
 - 9: **if** $\widehat{Q}_t^{X_k}(\tau + \epsilon + c_t(\frac{\delta}{K})) \prec x_t^-$ **then**
 - 10: $\mathcal{A} = \mathcal{A} \setminus \{k\}$ ▷ Discard k based on (5)
 - 11: **if** $x_t^+ \preceq \widehat{Q}_t^{X_k}(\tau + \epsilon - c_t(\frac{\delta}{K}))$ **then**
 - 12: $\widehat{k} = k$ ▷ Select k according to (6)
 - 13: **BREAK**
 - 14: $t = t + 1$
 - 15: **return** \widehat{k}
-

Let us illustrate the algorithm on Example 1.

Example 2. (Example 1 continued) The non- (ϵ, τ) -optimal arm 3 cannot be eliminated by QPAC unless $x_t^- \succ x^4$. This happens when $\widehat{Q}_t^{X_k}(\tau - c_m(\frac{\delta}{K})) \succ x^4$ for arm 1 or arm 2 (see line 7 of Algorithm 1). Therefore, x_t^- needs to be high enough to eliminate a non- (ϵ, τ) -optimal arm. Moreover, for eliminating the third arm (see line 10), we need $\widehat{Q}_t^{X_3}(\tau + \epsilon + c_t(\frac{\delta}{K})) \preceq x^4$ to hold, i.e., $\widehat{F}_t^{X_3}(x^4) - c_t(\frac{\delta}{K}) -$

$\epsilon > \tau$. Therefore, for eliminating a non- (ϵ, τ) -optimal arm, the estimate of its CDF needs to be tight enough as well. The selection mechanism of QPAC is based on a very similar argument as the one described above for elimination.

4.2. Analysis

The sample complexity of the qualitative PAC setting is very similar to the one of the value-based setting (see [Even-Dar et al. \(2002\)](#)). Before discussing the result in more detail, some further notation needs to be introduced.

First of all, denote the set of (ϵ, τ) -optimal arms by $\mathcal{K}_{\epsilon, \tau}$, and define

$$\Delta_k^\epsilon = \sup \left\{ \Delta \in [0, 1] \mid Q^{X_k}(\tau + \epsilon + \Delta) \prec \max_{h \in \mathcal{K}_{\epsilon, \tau}} Q^{X_h}(\tau - \Delta) \right\}$$

for $k = 1, \dots, K$, where $\sup \emptyset = 0$ by definition. Finally, let \vee denote the max operator.

Theorem 1. *Assume that algorithm QPAC is run with parameters (ϵ, δ, τ) on a problem with K arms X_1, \dots, X_K . Then, with probability at least $1 - \delta$, QPAC outputs an (ϵ, τ) -optimal arm after drawing*

$$\mathcal{O} \left(\sum_{k=1}^K \frac{1}{(\epsilon \vee \Delta_k^\epsilon)^2} \log \frac{K}{(\epsilon \vee \Delta_k^\epsilon) \cdot \delta} \right)$$

samples. Thus, QPAC is an (ϵ, τ, δ) -quantile learner.

The proof of Theorem 1 is deferred to Appendix A.

Remark 1. *Note that the sample complexity of QPAC depends on the number of arms, K , but not on the number of different rewards (i.e., size of L).*

Remark 2. *Lower bound on sample complexity for value-based PAC bandits had already been investigated by [Mannor & Tsitsiklis \(2004\)](#). A similar lower bound analysis also applies to the qualitative setting resulting in a lower bound of the form $\Omega(\sum_{k=1}^K \frac{1}{(\epsilon \vee \Delta_k^\epsilon)^2} \log \frac{1}{\delta})$ (see Appendix A.1). Therefore this lower bound shows that the sample complexity of the QPAC algorithm given in Theorem 1 is optimal up to a logarithmic factor.*

5. Infinite Horizon

In this section, we analyze the infinite horizon setting, where the goal of the online learner is normally defined as minimizing the cumulative regret of its actions in the course of time. First of all, this of course presupposes an appropriate definition of the notion of *regret*. Preferably, in order to allow for a simple accumulation, regret should be defined in a quantitative way.

In the standard value-based setting, the regret of choosing an arm is typically defined in terms of the difference $x^* - x$

between the reward observed, x , and the reward that would have been obtained (in expectation) by choosing the best arm, namely the arm with the highest expected reward x^* . In our setting, a quantification of regret in terms of differences is no longer possible, however, since arithmetic operations are not defined on our qualitative scale L . Instead, as explained before, we are only allowed to compare outcomes in terms of “better” or “worse”. This leads us quite naturally to a binary regret function $\text{regret}(x, y) = \mathbb{I}\{x \in G\} - \mathbb{I}\{y \in G\}$ for obtaining reward y instead of x , where G is the subset of outcomes in L considered to be “good”. Accordingly, the (expected) immediate regret of choosing arm X_k is

$$\begin{aligned} \rho_k &= \max_{k'=1, \dots, K} \mathbf{E}[\text{regret}(X'_k, X_k)] \\ &= \max_{k'=1, \dots, K} \mathbf{P}[X_{k'} \in G] - \mathbf{P}[X_k \in G] \end{aligned} \quad (7)$$

Now, the above definition of regret raises another question: What is the set G of good outcomes? In our setting, a natural answer is to consider an outcome x as good if $x \succeq x^*$, i.e., if x is at least as good as the optimal τ -quantile (4). However, in conjunction with the sharp discontinuity of the regret function (7), this definition renders regret minimization a truly hard problem. In fact, as shown by the following example, no algorithm with sublinear distribution independent regret guarantees exists. A more formal explanation of the linear worst case regret is deferred to the supplementary material (see Appendix B.2). Our worst case analysis is based on the fact that bandit instances given in Example 3 (a) and (b) are hard to distinguish.

Example 3. *Consider the qualitative MAB settings illustrated in Figure 2. In case (a), it is clear that x^3 should be considered as the only “good” reward, and thus $x^* = x^3$ and $G = \{x^3\}$. The second arm thus never returns good rewards, whereas the first arm does with probability $1 - \tau + \delta$. Therefore, the regret of arm 2 is $1 - \tau + \delta$. On the other hand, in case (c) both x^2 and x^3 should be considered good, so $x^* = x^2$ and $G = \{x^2, x^3\}$. Thus, while arm 2 returns a good reward consistently, the first arm is doing so only with probability $1 - \tau - \delta$. The regret of the first arm is $\tau + \delta$.*

As long as one cannot distinguish between cases (a) and (c) with high enough probability, the choice of which one to optimize for (which is crucial, as arm 2 has at least constant regret in (a), and the same holds for arm 1 in (c)) will remain random. (The problem of the learner, therefore, is to find out whether $\mathbf{P}[X_1 = x^1] \geq \tau$ or $\mathbf{P}[X_1 = x^1] < \tau$ for the first arm.) Thus, until that point is reached, any learner necessarily incurs linear regret in at least one of these examples. Additionally, to distinguish between the examples is getting more and more difficult as δ approaches 0 (for formal results see Appendix B.1 and B.2). This suggests that one cannot hope for sublinear worst case regret bounds.

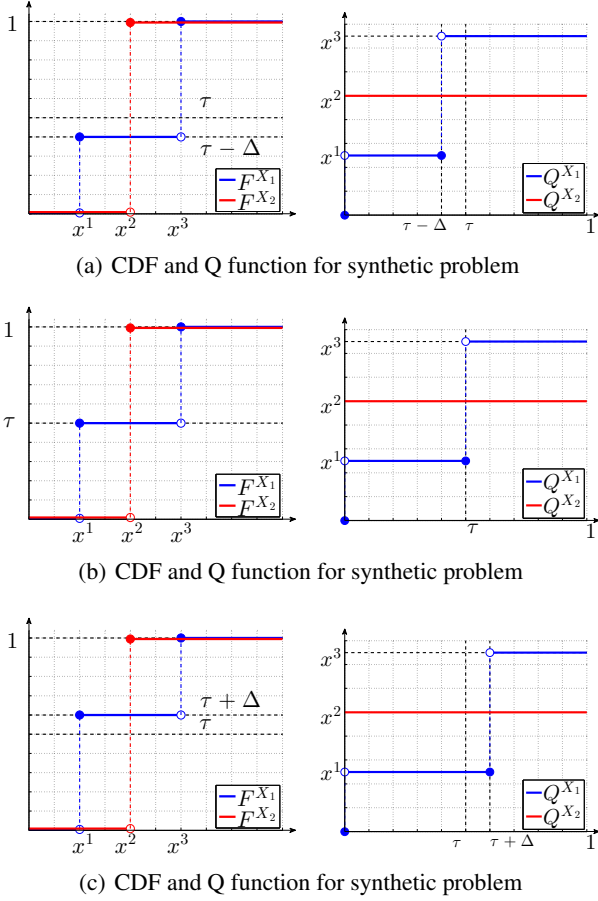


Figure 2. Synthetic qualitative MAB tasks with two arms.

Example 4 (Examples in Figure 2 continued). Now, consider cases (a) and (b), and the τ -quantile x^* . In case (b), $x^* = x^2$, thus $\mathbf{P}[X_1 \geq x^*] = 1 - \tau$ and $\mathbf{P}[X_2 \geq x^*] = 1$ while in case (a) (see Example 3), $\mathbf{P}[X_1 \geq x^*] = 1 - \tau + \delta$ and $\mathbf{P}[X_2 \geq x^*] = 0$. However, in order to distinguish the two cases, the learner needs to pull arm 1, leading to some non-negligible regret. This regret in (b), however, cannot be explained by any natural parameter (like the difference of the means in the quantitative case).

In order to avoid the problem in Example 4, we propose a slightly different definition of the set G of good outcomes. To this end, let

$$x^*(\tau') = \max_{k=1, \dots, K} Q^{X_k}(\tau')$$

for $\tau' \in [0, 1]$ (thus $x^* = x^*(\tau)$), and define

$$G = L_\tau = \{x \in L : x \geq x^*(\tau') \text{ for some } \tau' > \tau\}.$$

Correspondingly, the best arm with the minimal expected regret is defined by

$$k^* = \operatorname{argmax}_{1 \leq k \leq K} \mathbf{P}[X_k \in L_\tau],$$

Algorithm 2 QUCB(τ)

- 1: **for** rounds $t = 1, \dots, K$ **do**
- 2: set $k_t = t$ and $T(k_t) = 1$
- 3: pull arm k_t and observe sample $X_{k_t,1}$
- 4: **for** rounds $t = K + 1, K + 2, \dots$ **do**
- 5: $\hat{x}_t = \sup_{k=1, \dots, K} \hat{Q}_{T(k)}^{X_k}(\tau + c(t, T(k)))$
- 6: $k_t := \operatorname{argmin}_{k=1, \dots, K} (\hat{p}_{T(k)}^{X_k}(\hat{x}_t) - c(t, T(k)))$
- 7: set $T(k_t) = T(k_t) + 1$
- 8: pull arm k_t , and observe sample $X_{k_t, T(k_t)}$

the (expected) immediate regret of arm k is $\rho_k = \mathbf{P}[X_{k^*} \in L_\tau] - \mathbf{P}[X_k \in L_\tau]$, and $R_t = t\mathbf{P}[X_{k^*} \in L_\tau] - \mathbf{E}[\sum_{t'=1}^t \mathbb{I}\{X_{k_{t'}} \in L_\tau\}]$ is the expected cumulative regret, where $k_{t'}$ is the index of the arm chosen in round t' . In our example, this approach renders the first arm optimal in both (a) and (b), since in both cases $L_\tau = x^3$. Note also that in case of a “clear” separation (i.e., when $x^*(\tau) = x^*(\tau + \epsilon)$ for some $\epsilon > 0$) this regret is equivalent to the one based on $\mathbf{P}[X_k \geq x^*]$.

5.1. Algorithm QUCB

In Algorithm 2, we present the pseudocode of our QUCB (which is short for Qualitative Upper Confidence Bound) algorithm. In each round t , it uses an estimate $\hat{Q}_{T_i(k)}^{X_k}$ of the τ -quantiles and pulls the arm which maximizes this estimate. The confidence term used is $c(t, s) = \sqrt{2 \frac{\ln(t-1)}{s}}$. (The algorithm also needs to break ties, which is carried out in a similar fashion, but using the estimates of the p functions described later.) As a result, the accuracy of the estimate for the most promising arm will be increased. Suboptimal arms will be revealed as such as soon as the accuracy of the corresponding estimate is high enough—mathematically, this can be guaranteed thanks to the right-continuity of the quantile functions.

For selecting the most promising arms between those with $\mathbf{P}[X_k \notin L_\tau] < \tau$, the algorithm additionally keeps track of an estimate of another function, $p^{X_k}(x) = \mathbf{P}[X_k < x]$, using $\hat{p}_m^{X_k}(x) = \frac{1}{m} \sum_{j=1}^m \mathbb{I}\{X_{k,j} < x\}$. (Thus $\hat{F}_m^{X_k}(x) = \hat{p}_m^{X_k}(x)$ for $x \in L \setminus \{X_{k,1}, \dots, X_{k,m}\}$.)

In order to state the regret bound of QUCB, some further notation needs to be introduced. But first we need a technical lemma (see Appendix B for the proof).

Lemma 1. *If $\mathbf{P}[X_k \notin L_\tau] < \tau$ for some $1 \leq k \leq K$ then $(\inf L_\tau) \in L_\tau$, $\tau > \min_{k'} \mathbf{P}[X_{k'} < \inf L_\tau]$ and $\tau < \min_{k'} \mathbf{P}[X_{k'} \leq \inf L_\tau]$. Also, $Q^{X_k}(\tau) = x^*$.*

For arm k with $\mathbf{P}[X_k \notin L_\tau] > \tau$, define $\Delta_k = \mathbf{P}[X_k \notin L_\tau] - \tau$. Now, consider some arm k with $\mathbf{P}[X_k \notin L_\tau] \leq \tau$. In case $\mathbf{P}[X_k \notin L_\tau] \leq \tau$, X_k is also optimal, it is thus only interesting to upper bound $T_k(t)$ in case $\rho_k = \mathbf{P}[X_k \notin$

$L_\tau] - \mathbf{P}[X_{k^*} \notin L_\tau] > 0$. In that case, $\mathbf{P}[X_{k^*} \notin L_\tau] < \tau$ and Lemma 1 applies. Therefore, $\Delta_0 = \min_{k'} \mathbf{P}[X_{k'} \leq \inf L_\tau] - \tau > 0$. Based on these, for X_k satisfying $\mathbf{P}[X_k \notin L_\tau] \leq \tau$, define $\Delta_k = \min(\rho_k, \Delta_0)$. Then, we have (see Appendix B for the proof):

Theorem 2. *The expected cumulative regret of QUCB in round t is $R_t = \mathcal{O}\left(\sum_{k: \Delta_k > 0} \frac{\rho_k}{(\Delta_k)^2} \log t\right)$.*

For regret lower bounds see Appendix B.1 and B.2.

6. Experiments

6.1. Finite Horizon

The goal of the first experiment is to assess the impact of the parameters τ and ϵ on the sample complexity, that is, the number of samples taken by the algorithm prior to termination. We generated random bandit instances for which rewards are taken from a totally ordered discrete set $\{x^1, \dots, x^{10}\}$. In other words, the arm distributions are multinomial distributions. The parameters of the distributions are drawn uniformly at random from $(0, 1)$ and proportionally scaled so as to sum up to one. The sample complexities for various values of parameters ϵ and τ are shown in Figure 3. As can be seen, the smaller ϵ , the higher the sample complexity—thus, our algorithm scales gracefully with the approximation error allowed. The second observation is that the parameter τ has only a weak influence on the sample complexity. This can be explained by the fact that our confidence intervals are derived for the empirical CDF as a whole, without being tailored for a specific quantile.

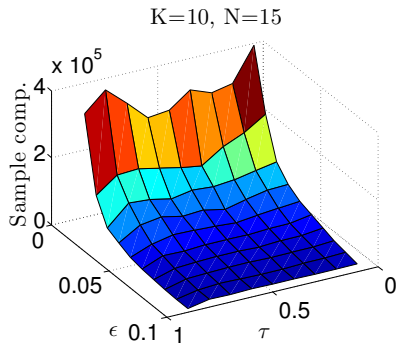


Figure 3. The sample complexity of QPAC for $K = 10$ and different values of the parameter τ and ϵ . The arm distributions are categorical distributions. The results are averaged over 100 repetitions. The confidence parameter δ was set to 0.05 for each run; accordingly, the average accuracy was significantly above $1 - \delta = 0.95$ in each case.

In the second experiment, we compare the performance of the QPAC algorithm with a standard PAC bandit algorithm on bandit problems for which the (ϵ, τ) -optimal arm coin-

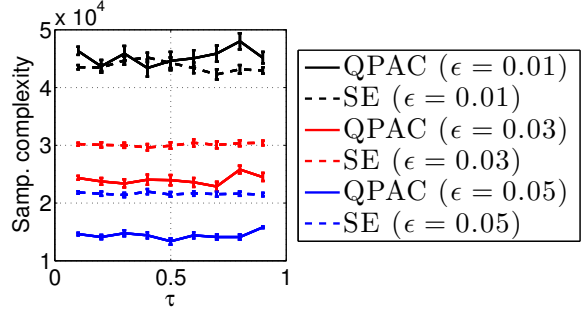


Figure 4. The sample complexity of SE and QPAC for the NHS problem with various parameter setting. The number K of arms was set to 15. The results are averaged over 100 repetitions. The confidence parameter δ was set to 0.05 for each run; accordingly, the average accuracy was significantly above $1 - \delta = 0.95$ in each case.

cides with the ϵ -best arm in terms of means, thereby assuring that both learners are seeking the same arm. As a baseline, we run the SUCCESSIVEELIMINATION learner (SE) by Even-Dar et al. (2002), which is an (ϵ, δ) -PAC learner (i.e., it returns an ϵ -optimal arm with probability at least $1 - \delta$). To guarantee a fair comparison, the confidence interval defined in (2) is used in our implementation of SE, which differs only in constants from the one of Even-Dar et al. (2002). We tested the algorithm on the “Needle in the Haystack” (NHS) problem, which consists of arms obeying Bernoulli distribution with parameter p —except for one of them, the target, which has a slightly higher parameter $p + p'$. Note that for $\tau = (1 - p) - p'/2$, we have $Q^{X_i}(\tau) = 1$ for the single τ -optimal arm, and $Q^{X_{i'}}(\tau) = 0$ otherwise. We run the experiments with $\tau = 0.1, \dots, 0.9$ and $p' = 0.1$; correspondingly, p was set to 0.85, 0.75, \dots , 0.05, respectively. The approximation error ϵ was set to 0.01, 0.03 and 0.05. As can be seen from the results in Figure 4, QPAC slightly outperforms SE in terms of sample complexity for smaller values of ϵ . This can be explained by the fact that, although both algorithms are using similar elimination strategies, the statistics they use are of different nature.

6.2. Infinite Horizon

In this section, we evaluate the QUCB algorithm on several numerical test cases. As a baseline, we run the standard UCB algorithm (Auer et al., 2002), which maximizes the sum of rewards. In each case, we plot the quantile-based cumulative regret and the average accuracy of the algorithms versus the number of iterations. By definition, the accuracy of an algorithm is 1, if it selects an arm from \mathcal{K}^* , and 0 otherwise. We run the algorithms on the following bandit problems:

1. Bandit instances defined in Figure 2(a) and 2(b). The

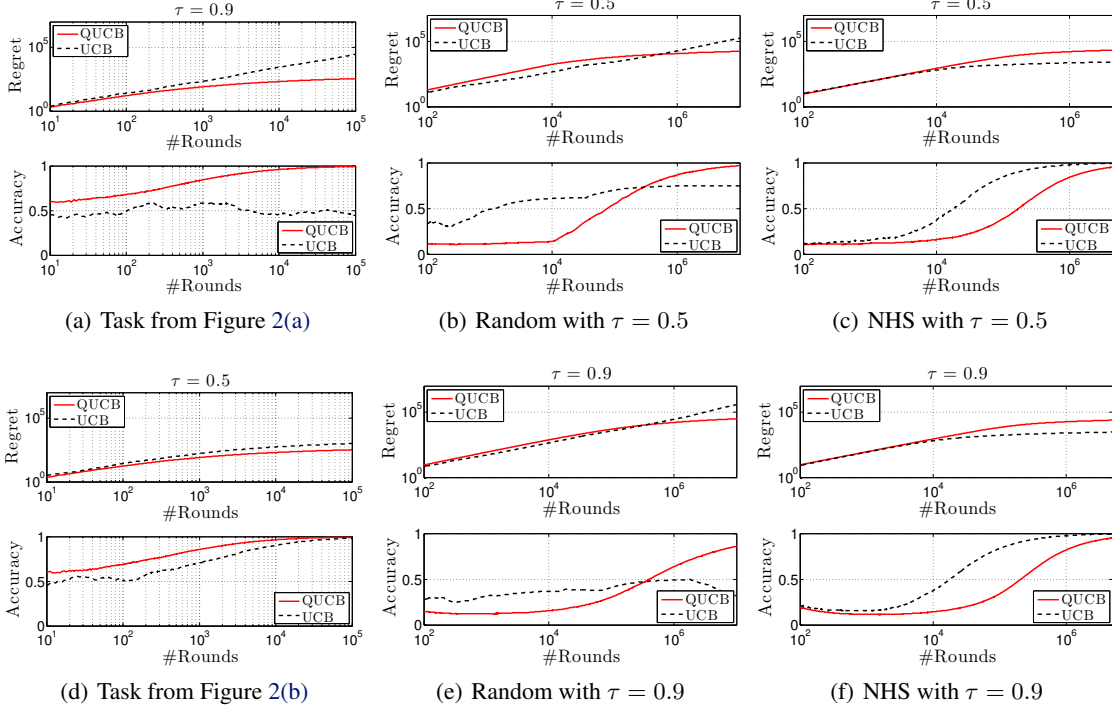


Figure 5. Cumulative regret and accuracy for various test cases.

results are shown in Figure 5(a) and 5(d), respectively. We set $x^1 = 1$, $x^2 = 2$ and $x^3 = 3$ in the case of UCB. The parameter δ was set to 0.1.

2. Multinomial arm distributions as described in the previous section, with parameters drawn uniformly at random. For the quantiles, we used $\tau \in \{0.5, 0.9\}$. The results are shown in Figure 5(b) and 5(e), respectively.
3. NHS problem with parameters $p = 0.45, p' = 0.1, \tau = 0.5$ and $p = 0.85, p' = 0.1, \tau = 0.9$. The results are shown in Figure 5(c) and 5(f).

In the first test case described in Figure 2(a), the mean of both arm distributions is $1/2$. Therefore, since UCB cannot distinguish the arms, its accuracy is fluctuating around $1/2$. As opposed to this, QUCB is able to identify the optimal arm. In the second test case defined in Figure 2(b), the best option is the second arm, both in terms of mean and τ -quantile. Accordingly, QUCB and UCB are both able to identify the optimal arm. On multinomial arm distributions, QUCB significantly outperforms the UCB algorithm. This can be explained by the fact that the median ($\tau = 1/2$) does not necessarily coincide with the mean—the higher τ , the more different the goals of the learners will actually become. As expected, the performance of both algorithms is on par in the case of the NHS problem.

7. Conclusion

We have investigated the setting of quantile-based online bandit learning in the qualitative case, that is, when rewards are coming from a complete totally ordered set but are not necessarily numerical. We introduced and analyzed a PAC algorithm in the finite horizon setting. Moreover, for the infinite horizon setting, we proposed an algorithm the (distribution-dependent) expected regret of which is growing logarithmically with time.

We have showed that sublinear regret in the qualitative setting is not achievable in the worst case (without using properties of the underlying distributions) in general. Since the standard reward expectation maximization problem has a known lower-bound of $\Omega(1/\sqrt{T})$ (Audibert & Bubeck, 2010) and the risk-aversion setup has $\Omega(T^{2/3})$ (Sani et al., 2012), therefore our worst case result implies that minimizing the quantile-based regret in the qualitative setting is intrinsically more difficult than the standard value-based and the risk-aversion bandit problems.

Acknowledgments

This work was supported by the European Community’s 7th Framework Programme (FP7/2007-2013) under grant agreement n° 270327 (project CompLACS) and by the German Research Foundation under grant HU 1284/8-1.

References

- Audibert, Jean-Yves and Bubeck, Sébastien. Regret bounds and minimax policies under partial monitoring. *J. Mach. Learn. Res.*, 11:2785–2836, 2010.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- Bubeck, S. and Cesa-Bianchi, N. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- Bubeck, S., Munos, R., and Stoltz, G. Pure exploration in multi-armed bandits problems. In *Proceedings of the 20th ALT*, ALT’09, pp. 23–37, 2009.
- Bubeck, S., Wang, T., and Viswanathan, N. Multiple identifications in multi-armed bandits. In *Proceedings of The 30th ICML*, pp. 258–265, 2013.
- Busa-Fekete, R. and Hüllermeier, E. A survey of preference-based online learning with bandit algorithms. In *Algorithmic Learning Theory (ALT)*, volume 8776, pp. 18–39, 2014.
- Busa-Fekete, R. and Kégl, B. Fast boosting using adversarial bandits. In *International Conference on Machine Learning (ICML)*, volume 27, pp. 143–150, Haifa, Israel, 2010. ACM.
- Cappé, O., Garivier, A., Maillard, O.-A., Munos, R., and Stoltz, G. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3): 1516–1541, 2013.
- Carpentier, A. and Valko, M. Extreme bandits. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems*, pp. 1089–1097, 2014.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning and Games*. Cambridge university press, 2006.
- Dvoretzky, A., Kiefer, J., and Wolfowitz, J. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, 27(3):642–669, 1956.
- Even-Dar, E., Mannor, S., and Mansour, Y. PAC bounds for multi-armed bandit and markov decision processes. In *Proceedings of the 15th Conference on Learning Theory (COLT)*, pp. 255–270, 2002.
- Gabillon, V., Ghavamzadeh, M., Lazaric, A., and Bubeck, S. Multi-bandit best arm identification. In *Advances in NIPS 24*, pp. 2222–2230, 2011.
- Gaudel, Romaric and Sebag, Michèle. Feature selection as a one-player game. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pp. 359–366, 2010.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. Pac subset selection in stochastic multi-armed bandits. In *Proceedings of the Twenty-ninth International Conference on Machine Learning (ICML 2012)*, pp. 655–662, 2012.
- Kuleshov, Volodymyr and Precup, Doina. Algorithms for multi-armed bandit problems. *CoRR*, abs/1402.6028, 2014.
- Lai, T.L. and Robbins, H. Asymptotically efficient allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- Mannor, Shie and Tsitsiklis, John N. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5:623–648, 2004.
- Massart, P. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The Annals of Probability*, 18(3): 1269–1283, 1990.
- Sani, Amir, Lazaric, Alessandro, and Munos, Rémi. Risk-aversion in multi-armed bandits. In *26th Annual Conference on Neural Information Processing Systems (NIPS)*, pp. 3284–3292, 2012.
- Schachter, B. An irreverent guide to Value-at-Risk. *Financial Engineering News*, 1, 1997.
- Yu, Jia Yuan and Nikolova, Evdokia. Sample complexity of risk-averse bandit-arm selection. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pp. 2576–2582. AAAI Press, 2013.
- Yue, Y., Broder, J., Kleinberg, R., and Joachims, T. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.
- Yue, Yisong and Joachims, Thorsten. Beat the mean bandit. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 241–248, 2011.
- Zhou, Yuan, Chen, Xi, and Li, Jian. Optimal pac multiple arm identification with applications to crowdsourcing. *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 217–225, 2014.
- Zoghi, M., Whiteson, S., Munos, R., and Rijke, M. Relative upper confidence bound for the k-armed dueling bandit problem. In *Proceedings of the Thirty-First International Conference on Machine Learning (ICML)*, pp. 10–18, 2014.