

# Learning Modular Safe Policies in the Bandit Setting with Application to Adaptive Clinical Trials

Hossein Aboutaleb<sup>\*1</sup>, Doina Precup<sup>1</sup>, and Tibor Schuster<sup>2</sup>

<sup>1</sup>Department of Computer Science, McGill University. Mila Quebec AI Institute

<sup>2</sup>Department of Family Medicine, McGill University

## Abstract

The stochastic multi-armed bandit problem is a well-known model for studying the exploration-exploitation trade-off. It has significant possible applications in adaptive clinical trials, which allow for dynamic changes in the treatment allocation probabilities of patients. However, most bandit learning algorithms are designed with the goal of minimizing the expected regret. While this approach is useful in many areas, in clinical trials, it can be sensitive to outlier data, especially when the sample size is small. In this paper, we define and study a new robustness criterion for bandit problems. Specifically, we consider optimizing a function of the distribution of returns as a regret measure. This provides practitioners more flexibility to define an appropriate regret measure. The learning algorithm we propose to solve this type of problem is a modification of the BESA algorithm [Baransi *et al.*, 2014], which considers a more general version of regret. We present a regret bound for our approach and evaluate it empirically both on synthetic problems as well as on a dataset from the clinical trial literature. Our approach compares favorably to a suite of standard bandit algorithms. Finally, we provide a web application where users can create their desired synthetic bandit environment and compare the performance of different bandit algorithms online.

## Introduction

The multi-armed bandit is a standard model for researchers to investigate the exploration-exploitation trade-off, see e.g [Baransi *et al.*, 2014; Auer *et al.*, 2002; Sani *et al.*, 2012a; Chapelle and Li, 2011; Sutton and Barto, 1998]. One of the main advantage of multi-armed bandit problems is its simplicity that allows for a higher level of theoretical studies.

The multi-armed bandit problem consists of a set of arms, each of which generates a stochastic reward from a fixed but unknown distribution associated to it. Consider a series of multiple arm pulls (or steps)  $t = 1, \dots, T$  and selecting a specific arm  $a \in \mathcal{A}$  at each step i.e.  $a(t) = a_t$ . The standard goal

in the multi-armed bandit setting is to find the arm  $\star$  which has the maximum expected reward  $\mu_\star$  (or equivalently, minimum expected regret). The expected regret after  $T$  steps  $R_T$  is defined as the sum of the expected difference between the mean reward under  $\{a_t\}$  and the reward expected under the optimal arm  $\star$ :

$$R_T = \mathbb{E} \left[ \sum_{t=1}^T (\mu_\star - \mu_{a_t}) \right]$$

While this objective is very popular, there are practical applications, for example in medical research and AI safety [Garcia and Fernández, 2015] where maximizing expected value is not sufficient, and it would be better to have an algorithm sensitive also to the variability of the outcomes of a given arm. For example, consider multi-arm clinical trials where the objective is to find the most promising treatment among a pool of available treatments. Due to heterogeneity in patients' treatment responses, considering only the expected mean may not be of interest [Austin, 2011]. Specifically, as the mean is usually sensitive to outliers and does not provide information about the dispersion of individual responses, the expected reward has only limited value in achieving a clinical trial's objective. Due to these problems, previous contributions like [Sani *et al.*, 2012a] try to include the variance of rewards in the regret definition and develop algorithms to solve this slightly enhanced problem. While these modified approaches try to consider variability in the response of arms, they induce new problems due to the fact that the variance is not necessarily a good measure of variability for a distribution. This is because the variance equally penalizes responses that are above or below the mean response. Other articles like [Galichet *et al.*, 2013] try to use the conditional value at risk to define a better regret definition. Though the conditional value at risk may address the problem we faced with including variance, it may not reflect the amount of variability we could observe for a distribution over its entire domain. All in all, the consistency of treatments among patients is essential, with the ideal treatment usually defined as the one which has a high positive response rate while showing low variability in response among patients. Thus, the idea of consistency and safety seems to some extent subjective and problem dependent. As a result, it might be necessary to develop an algorithm which can work with an arbitrary definition of consistency for a distribution.

<sup>\*</sup>hossein.aboutalebi@mail.mcgill.ca

This kind of system design which allows the separation of different parts of a system (here regret function and learning algorithm) has already been explored in modular programming. In modular programming, we emphasize on splitting the entire system into independent modules which at the end, the composite of these modules builds our system. This design trick is necessary when we are dealing with the change of customer demands and we require our system to adapt with the new demands. Here, we follow the same paradigm by making regret definition independent of the learning algorithm. As a result, we allow more flexibility in defining the regret function which is capable of incorporating problem specific demands.

Finally, we achieve the aforementioned goals by extending one of the recent algorithms in the bandit literature called BESA (Best Empirical Sampled Average) [Baransi *et al.*, 2014]. One of the main advantage of BESA compared to other existing bandit algorithms is that it does not involve many hyper-parameters. This is especially useful when one does not have any prior knowledge or has insufficient prior knowledge about the different arms in the beginning. Also, this feature makes it easier to introduce modular design by using McDiarmid’s Lemma [El-Yaniv and Pechyony, 2009].

**Key contributions:** We provide a modular definition of regret called safety-aware regret which allows higher flexibility in defining the risk for multi-armed bandit problems. We propose a new algorithm called BESA+ which solves this category of problems. We show the upper-bounds of its safety-aware regret for two-armed and multi-armed bandits. For the experiment parts, we compare our model with some of the notable earlier research works and show that BESA+ has a satisfying performance. For the last experiment, we depict the performance of our algorithm on a real clinical dataset and illustrate that it is capable of solving the problem with user-defined safety-aware regret. Finally, for the first time as far as we know, we provide a web application which allows users to create their own custom environment and compare our algorithm with other works.

## Background and Notation

We consider the standard bandit setting with action (arm) set  $\mathcal{A}$ , where each action  $a \in \mathcal{A}$  is characterized by a reward distribution  $\varphi_a$ . The distribution for action  $a$  has mean  $\mu_a$  and variance  $\sigma_a^2$ . Let  $X_{a,i} \sim \varphi_a$  denote the  $i$ -th reward sampled from the distribution of action  $a$ . All actions and samples are independent. The bandit problem is described as an iterative game where, on each step (round)  $t$ , the player (an algorithm) selects action (arm)  $a_t$  and observes sample  $X_{a_t, N_{a_t, t}}$ , where  $N_{a, t} = \sum_{s=1}^t \mathbb{I}\{a_s = a\}$  denotes the number of samples observed for action  $a$  up to time  $t$  (inclusively). A policy is a distribution over  $\mathcal{A}$ . In general, stochastic distributions are necessary during the learning stage, in order to identify the best arm. We discuss the exact notion of “best” below.

We define  $I_S(m, j)$  as the set obtained by sub-sampling without replacement  $j$  elements from the set  $S$  of size  $m$ . Let  $\mathcal{X}_{a, t}$  denote the history of observations (records) obtained from action (arm)  $a$  up to time  $t$  (inclusively), such that  $|\mathcal{X}_{a, t}| = N_{a, t}$ . The notation  $\mathcal{X}_{a, t}(\mathcal{I})$  indicates the set of sub-

samples from  $\mathcal{X}_{a, t}$ , where sub-sample  $\mathcal{I} \subset \{1, 2, \dots, N_{a, t}\}$ .

The multi-armed bandit was first presented in the seminal work of Robbins [Robbins, 1985]. It has been shown that under certain conditions [Burnetas and Katehakis, 1996; Lai and Robbins, 1985], a policy can have logarithmic cumulative regret:

$$\liminf_{t \rightarrow \infty} \frac{\mathfrak{R}_t}{\log(t)} \geq \sum_{a: \mu_a < \mu_*} \frac{\mu_* - \mu_a}{K_{\text{inf}}(r_a; r_*)}$$

where  $K_{\text{inf}}(r_a; r_*)$  is the Kullback-Leibler divergence between the reward distributions of the respective arms. Policies for which this bound holds are called *admissible*.

Several algorithms have been shown to produce admissible policies, including UCB1 [Auer *et al.*, 2002], Thompson sampling [Chapelle and Li, 2011; Agrawal and Goyal, 2013] and BESA [Baransi *et al.*, 2014]. However, theoretical bounds are not always matched by empirical results. For example, it has been shown in [Kuleshov and Precup, 2014] that two algorithms which do not produce admissible policies,  $\epsilon$ -greedy and Boltzmann exploration [Sutton and Barto, 1998], behave better than UCB1 on certain problems. Both BESA and Thompson sampling were shown to have comparable performance with Softmax and  $\epsilon$ -greedy.

While the expected regret is a natural and popular measure of performance which allows the development of theoretical results, recently, some papers have explored other definitions for regret. For example, [Sani *et al.*, 2012b] consider a linear combination of variance and mean as the definition of regret for a learning algorithm  $A$ :

$$\widehat{MV}_t(A) = \widehat{\sigma}_t^2(A) - \widehat{\rho}\widehat{\mu}_t(A) \quad (1)$$

where  $\widehat{\mu}_t$  is the estimate of the average of observed rewards up to time step  $t$  and  $\widehat{\sigma}_t^2$  is a biased estimate of the variance of rewards up to time step  $t$ . The regret is then defined as:

$$\mathfrak{R}_t(A) = \widehat{MV}_t(A) - \widehat{MV}_{*, t}(A),$$

where  $*$  is the optimal arm. According to [Maillard, 2013], however, this definition is going to penalize the algorithm if it switches between optimal arms. Instead, in [Maillard, 2013], the authors devise a new definition of regret which controls the lower tail of the reward distribution. However, the algorithm to solve the corresponding objective function seems time-consuming, and the optimization to be performed may be intricate. Finally, in [Galichet *et al.*, 2013], the authors use the notion of conditional value at risk in order to define the regret.

## Measure of regret

Unlike previous works, we now give a formal definition of class of functions which can be used as a separate module inside our learning algorithm module to measure the regret. We call these class of functions “safety value functions”.

In the following section, we try to formally define these functions. Assume we have  $k$  arms ( $|\mathcal{A}| = k$ ) with reward distributions  $\varphi_1, \varphi_2, \dots, \varphi_k$ .

**Definition 0.1. safety value function:** Let  $\mathcal{D}$  denotes the set of all possible reward distributions for a given interval. The

safety value function  $v : \mathcal{D} \rightarrow \mathcal{R}$  provides a score for a given distribution.

The optimal arm  $\star$  under this value function is defined as

$$\star \in \arg \max_{a \in \mathcal{A}} (v(\varphi_a)) \quad (2)$$

The regret corresponding to the safety value function up to time  $T$  is defined as:

$$\mathfrak{R}_{T,v} = \mathbb{E} \left[ \sum_{t=1}^T (v(\varphi_\star) - v(\varphi_{a_t})) \right] \quad (3)$$

We call (3), *safety-aware regret*.

When the context is clear, we usually drop the subscript  $v$  and use only  $\mathfrak{R}_T$  for the ease of notation.

**Definition 0.2. Well-behaved safety value function:** Given a reward distribution  $\varphi_a$  over the interval  $[0, 1]$ , a safety value function  $v$  for this distribution is called well-behaved if there exists an unbiased estimator  $\hat{v}$  of  $v$  such that for any set of observation  $\{x_1, x_2, \dots, x_n\}$  sampled from  $\varphi_a$ , and for some constant  $\gamma$  we have:

$$\sup_{\hat{x}_i} |\hat{v}(x_1, \dots, x_i, \dots, x_n) - \hat{v}(x_1, \dots, \hat{x}_i, \dots, x_n)| < \frac{\gamma}{n} \quad (4)$$

If (4) holds for any reward distribution  $\varphi$  over the interval  $[0, 1]$ , we call the safety value function  $v$ , a well-behaved safety value function.

**Example 1:** For a given arm  $a$  which has reward distribution limited to interval  $[0, 1]$ , consider the safety value function  $\mu_a - \rho\sigma_a^2$  which measures the balance between the mean and the variance of the reward distribution of arm  $a$ .  $\rho$  is a hyper-parameter constant for adjusting the balance between variance and the mean. This is a well-behaved safety function if we use the following estimator for computing empirical mean and variance:

$$\hat{\mu}_{a,t} = \frac{1}{N_{a,t}} \sum_{i=1}^{N_{a,t}} r_{a,i} \quad (5)$$

$$\hat{\sigma}_{a,t}^2 = \frac{1}{N_{a,t} - 1} \sum_{i=1}^{N_{a,t}} (r_{a,i} - \hat{\mu}_{a,t})^2 \quad (6)$$

where  $r_{a,i}$  is the  $i$ th reward obtained from pulling arm  $a$ . It should be clear that the unbiased estimator  $\hat{\mu}_{a,t} - \rho\hat{\sigma}_{a,t}^2$  satisfies (4).  $\square$

Other types of well-behaved safety function can be defined as a function of standard deviation or conditional value at risk similar to the previous example. In the next section, we are going to develop an algorithm which can optimize the safety-aware regret.

## Proposed Algorithm

In order to optimize the safety-aware regret, we build on the BESA algorithm, which we will now briefly review. As discussed in [Baransi *et al.*, 2014], BESA is a non-parametric

(without hyperparameter) approach for finding the optimal arm according to the expected mean regret criterion. Consider a two-armed bandit with actions  $a$  and  $\star$ , where  $\mu_\star > \mu_a$ , and assume that  $N_{a,t} < N_{\star,t}$  at time step  $t$ . In order to select the next arm for time step  $t + 1$ , BESA first sub-samples  $s_\star = I_\star(N_{\star,t}, N_{a,t})$  from the observation history (records) of the arm  $\star$  and similarly sub-sample  $s_a = I_a(N_{a,t}, N_{\star,t}) = \mathcal{X}_{a,t}$  from the records of arm  $a$ . If  $\hat{\mu}_{s_a} > \hat{\mu}_{s_\star}$ , BESA chooses arm  $a$ , otherwise it chooses arm  $\star$ .

The main reason behind the sub-sampling is that it gives a similar opportunity to both arms. Consequently, the effect of having a small sample size, which may cause bias in the estimates diminishes. When there are more than two arms, BESA runs a tournament algorithm on the arms [Baransi *et al.*, 2014].

Finally, it is worth mentioning that the proof of the regret bound of BESA uses a non-trivial lemma for which authors did not provide any formal proof. In this paper, we will avoid using this lemma to prove the soundness of our proposed algorithm for a more general regret family. Also, we extend the proof for the multi-armed case which was not provided in the [Baransi *et al.*, 2014].

We are now ready to outline our proposed approach, which we call BESA+. As in [Baransi *et al.*, 2014], we focus on the two-arm bandit. For more than two arms, a tournament can be set up in our case as well.

---

### Algorithm BESA+ two action case

---

**Input:** Safety aware value function  $v$  and its estimate  $\hat{v}$

**Parameters:** current time step  $t$ , actions  $a$  and  $b$ . Initially  $N_{a,0} = 0, N_{b,0} = 0$

- 1: **if**  $N_{a,t-1} = 0 \vee N_{a,t-1} < \log(t)$  **then**
  - 2:  $a_t = a$
  - 3: **else if**  $N_{b,t-1} = 0 \vee N_{b,t-1} < \log(t)$  **then**
  - 4:  $a_t = b$
  - 5: **else**
  - 6:  $n_{t-1} = \min\{N_{a,t-1}, N_{b,t-1}\}$
  - 7:  $\mathcal{I}_{a,t-1} \leftarrow I_a(N_{a,t-1}, n_{t-1})$
  - 8:  $\mathcal{I}_{b,t-1} \leftarrow I_b(N_{b,t-1}, n_{t-1})$
  - 9: Calculate  $\tilde{v}_{a,t} = \hat{v}(\mathcal{X}_{a,t-1}(\mathcal{I}_{a,t-1}))$  and  $\tilde{v}_{b,t} = \hat{v}(\mathcal{X}_{b,t-1}(\mathcal{I}_{b,t-1}))$
  - 10:  $a_t = \arg \max_{i \in \{a,b\}} \tilde{v}_{i,t}$  (break ties by choosing arm with fewer tries)
  - 11: **end if**
  - 12: **return**  $a_t$
- 

If there is a strong belief that one arm should be better than the other then instead of using factor  $\log(t)$  in Algorithm BESA+, one can use  $\alpha \log(t)$  factor (where  $0 < \alpha < 1$  and is constant) to reduce the final regret.

The first major difference between BESA+ and BESA is the use of the safety-aware value function instead of the simple regret. A second important change is that BESA+ selects the arm which has been tried less up to time step  $t$  if the arm has been chosen less than  $\log(t)$  times up to  $t$ . Essentially, this change in the algorithm is negligible in terms of establishing the total expected regret, as we cannot achieve any better bound than  $\log(T)$  which is shown in Robbins' lemma [Lai

and Robbins, 1985]. This tweak also turns out to be vital in proving that the expected regret of the BESA+ algorithm is bounded by  $\log(T)$  (a result which we present shortly).

To better understand why this modification is necessary, consider a two arms scenario. The first arm gives a deterministic reward of  $r \in [0, 0.5)$  and the second arm has a uniform distribution in the interval  $[0, 1]$  with the expected reward of 0.5. If we are only interested in the expected reward ( $\mu$ ), the algorithm should ultimately favor the second arm. On the other hand, there exists a probability of  $r$  that the BESA algorithm is going to constantly choose the first arm if the second arm gives a value less than  $r$  on its first pull. In contrast, BESA+ evades this problem by letting the second arm be selected enough times such that it eventually becomes distinguishable from the first arm.

We are now ready to state the main theoretical result of our proposed algorithm.

**Theorem 0.1.** *Let  $v$  be a well-behaved safety value function. Assume  $\mathcal{A} = \{a, \star\}$  be a two-armed bandit with bounded rewards  $\in [0, 1]$ , and the value gap  $\Delta = v_\star - v_a$ . Given the value  $\gamma$ , the expected safety-aware regret of the Algorithm BESA+ up to time  $T$  is upper bounded as follows:*

$$\mathfrak{R}_T \leq \zeta_{\Delta, \gamma} \log(T) + \theta_{\Delta, \gamma} \quad (7)$$

where in (7),  $\zeta_{\Delta, \gamma}, \theta_{\Delta, \gamma}$  are constants which are dependent on the value of  $\gamma, \Delta$ .

*Proof.* Due to the page limit, we could not include all the proof. Here, we just provide a short overview of the proof. The proof mainly consists of two parts. The first part of our proof is similar to [Baransi *et al.*, 2014] but instead we have used McDiarmid's Lemma [El-Yaniv and Pechyony, 2009] [Tolstikhin, 2017]. For the second part of the proof, unlike [Baransi *et al.*, 2014], we have avoided using the unproven lemma in their work and instead tried to compute the upper bound directly by exploiting the *log* trick in our algorithm (this trick has been further elaborated in the first experiment). Interested reader can visit [here](#) to see the full proof. ■

**Theorem 0.2.** *Let  $v$  be a well-behaved safety value function. Assume  $\mathcal{A} = \{a_1, \dots, a_{k-1}, \star\}$  be a  $k$ -armed bandit with bounded rewards  $\in [0, 1]$ . Without loss of generality, consider the optimal arm is  $\star$  and the value gap for arm  $a, \star$  is  $\Delta_a = v_\star - v_a$ . Also consider  $\Delta_{max} = \max_{a \in \mathcal{A}} \Delta_a$ . Given the value  $\gamma$ , the expected safety-aware regret of the Algorithm BESA+ up to time  $T$  is upper bounded as follows:*

$$\mathfrak{R}_T \leq \frac{\Delta_{max} \lceil \log k \rceil}{\Delta_{\hat{a}}} [\zeta_{\Delta_{\hat{a}}, \gamma} \log(T) + \theta_{\Delta_{\hat{a}}, \gamma}] + k \Delta_{max} n \quad (8)$$

where in (8),  $\zeta, \theta$  are constants which are dependent on the value of  $\gamma, \Delta$ . Moreover,  $\hat{a}$  is defined:

$$\hat{a} = \arg \max_{a \in \mathcal{A}} \zeta_{\Delta_a, \gamma} \log(T) + \theta_{\Delta_a, \gamma}$$

for  $T \geq n$ .

*Proof.* We know that the arm  $\star$  has to play at most  $\lceil \log k \rceil$  matches (games) in order to win the round. If it losses any of

these  $\lceil \log k \rceil$  games, we know that at that round we will see a regret. This regret should be less than or equal to  $\Delta_{max}$ .

In the following, We use notation  $\mathbf{1}_{-a_\star, i}$  to denote the indicator for the event of  $a_\star$  losing the  $i$ th match ( $1 \leq i \leq \lceil \log k \rceil$ ).

$$\begin{aligned} \mathfrak{R}_T &= \sum_{t=1}^T \sum_{i=1}^k \Delta_{a_i} \mathbb{E}[\mathbf{1}_{a_t=a_i}] \\ &\leq \sum_{t=1}^T \sum_{i=1}^{\lceil \log k \rceil} \Delta_{max} \mathbb{E}[\mathbf{1}_{-a_\star, i}] \\ &\leq \sum_{t=1}^T \sum_{i=1}^{\lceil \log k \rceil} \Delta_{max} \max_{i'} \{\mathbb{E}[\mathbf{1}_{-a_\star, i'}]\} \\ &\leq \sum_{i=1}^{\lceil \log k \rceil} \Delta_{max} \sum_{t=1}^T \max_{i'} \{\mathbb{E}[\mathbf{1}_{-a_\star, i'}]\} \\ &\leq \frac{\Delta_{max} \lceil \log k \rceil}{\Delta_{\hat{a}}} \sum_{t=n}^T \Delta_{\hat{a}} \mathbb{E}[\mathbf{1}_{-a_\star, \hat{a}}] + k \Delta_{max} n \\ &\leq \frac{\Delta_{max} \lceil \log k \rceil}{\Delta_{\hat{a}}} [\zeta_{\Delta_{\hat{a}}, \gamma} \log(T) + \theta_{\Delta_{\hat{a}}, \gamma}] + k \Delta_{max} n \end{aligned} \quad (9)$$

■

## Empirical results

### Empirical comparison of BESA and BESA+

As discussed in the previous section, BESA+ has some advantages over BESA. We illustrate the example we discussed in the previous section through the results in Figures 1-3, for  $r \in \{0.2, 0.3, 0.4\}$ . Each experiment has been repeated 200 times. Note that while BESA has an almost a linear regret behavior, BESA+ can learn the optimal arm within the given time horizon and its expected accumulated regret is upper bounded by a log function. It is also easy to notice that BESA+ has a faster convergence rate compared with BESA. As  $r$  gets closer to 0.5, the problem becomes harder. This phenomenon is a direct illustration of our theoretical result.




Figure 1: Result of accumulated expected regret for  $r = 0.4$

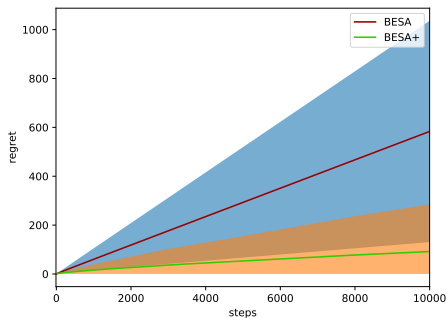


Figure 2: Result of accumulated expected regret for  $r = 0.3$




Figure 3: Result of accumulated expected regret for  $r = 0.2$

### Conditional value at risk safety value function

As discussed in [Galichet *et al.*, 2013], in some situations, we need to limit the exploration of risky arms. Examples include financial investment where investors may tend to choose risk-averse kind of strategy. Using conditional value at risk as a risk measure is one of the approaches to achieve this goal. Informally, conditional value at risk level  $\alpha$  is defined as the expected values of the quantiles of reward distribution where the probability of the occurrence of values inside this quantile is less than or equal to  $\alpha$ . More formally:

$$CVaR_\alpha = \mathbb{E}[X|X < v_\alpha] \quad (10)$$

where in (10),  $v_\alpha = \arg \max_\beta \{\mathbb{P}(X < \beta) \leq \alpha\}$ . To estimate (10), we have used the estimation measure introduced by [Chen, 2007]. This estimation is also employed in [Galichet *et al.*, 2013] work to derive their MARAB algorithm. Here, we have used this estimation for the Conditional value at risk safety value function which is the regret measure for this problem. Our environment consists of 20 arms where each arm reward distribution is the truncated Gaussian mixture consisting of four Gaussian distribution with equal probability. The reward of arms are restricted to the interval  $[0, 1]$ . To make the environment more complex, the mean and standard deviation of arms are sampled uniformly from the interval  $[0, 1]$  and  $[0.5, 1]$  respectively. The experiments are carried out for  $\alpha = 10\%$ . For MARAB algorithm, we have used grid search and set the value  $C = 1$ . The figures 4, 5 depict the results of the run for ten experiments. It is noticeable that in both figures BESA+ has a lower variance in experiments.




Figure 4: Accumulated regret figure. The safety value function here is conditional value at risk.




Figure 5: Percentage of optimal arm play figure. The safety value function here is conditional value at risk.

### Mean-variance safety value function

Next, we evaluated the performance of BESA+ with the regret definition provided by [Sani *et al.*, 2012a]. Here, we used the same 20 arms Gaussian mixture environment described in the previous section. We evaluated the experiments with  $\rho = 1$  which is the trade off factor between variance and the mean. The results of this experiment is depicted in figures 6, 7. The hyper-parameters used here for algorithms MV-LCB and ExpExp are based on what [Sani *et al.*, 2012a] suggests using. Again, we can see that BESA+ has a relatively small variance over 10 experiments.

### Real Clinical Trial Dataset

Finally, we examined the performance of BESA+ against other methods (BESA, UCB1, Thompson sampling, MV-LCB, and ExpExp) based on a real clinical dataset. This dataset includes the survival times of patients who were suffering from lung cancer [Ripley *et al.*, 2013]. Two different kinds of treatments (standard treatment and test treatment) were applied to them and the results are based on the number of days the patient survived after receiving one of the treatments. For the purpose of illustration and simplicity, we assumed non-informative censoring and equal follow-up times in both treatment groups. As the experiment has already been conducted, to apply bandit algorithms, each time a treatment is selected by a bandit algorithm, we sampled uniformly from the recorded results of the patients whom received that selected treatment and used the survival time as the reward signal. Figure 8 shows the distribution of treatment 1 and 2. We

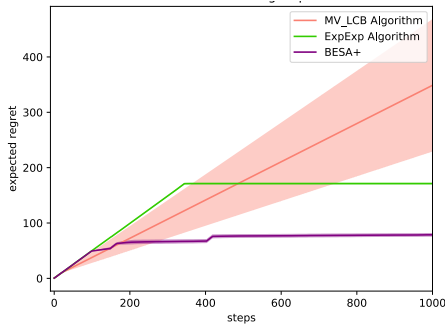


Figure 6: Accumulated regret figure. The safety value function here is mean-variance.

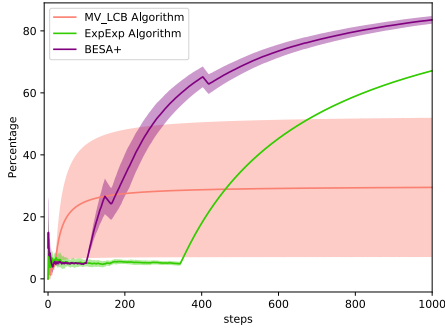


Figure 7: Percentage of optimal arm play figure. The safety value function here is mean-variance.

categorized the survival time into ten categories (category 1 showing the minimum survival time). It is interesting to notice that while treatment 2 has a higher mean than treatment 1 due to the effect of outliers, it has a higher level of variance compared to treatment 1. From figure 8 it is easy to deduce that treatment 1 has a more consistent behavior than treatment 2 and a higher number of patients who received treatment 2 died early. That is why treatment 1 may be preferred over treatment 2 if we use the safety value function described in Example 1. In this regard, by setting  $\rho = 1$ , treatment 1 has less expected mean-variance regret than treatment 2, and it should be ultimately favored by the learning algorithm. Figure 9 illustrates the performance of different bandit algorithms. It is easy to notice that BESA+ has relatively better performance than all the other ones.

### Web Application Simulator

As discussed earlier, for this project, we have developed a web application simulator for bandit problem where users can create their customized environment and run experiments online. Usually, research works provide limited experiments to testify their method. We tried to overcome this problem by developing this web application where the user can select number of arms and change their reward distribution. Then the web application will send the input to the web-server and show the results to the user by providing regret figures and additional figures describing the way algorithms have chosen arms over time. This software can be used as a benchmark for future bandit research and it is open sourced for future

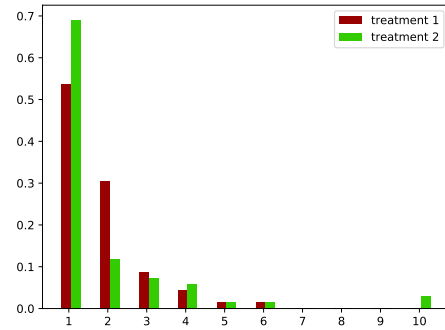


Figure 8: Distribution graph




Figure 9: Accumulated consistency-aware regret

extension. The link is [here](#)

### Conclusion and future work

In this paper, we developed a modular safety-aware regret definition which can be used to define the function of interest as a safety measure. We also modified the BESA algorithm and equipped it with new features to solve modular safety-aware regret bandit problems. We then computed the asymptotic regret of BESA+ and showed that it can perform like an admissible policy if the safety value function satisfies a mild assumption. Finally, we depicted the performance of BESA+ on the regret definition of previous works and showed that it can have better performance in most cases.

It is still interesting to investigate whether we can find better bounds for BESA+ algorithm with modular safety-aware regret definition. Another interesting path would be to research if we can define similar safety-aware regret definition for broader reinforcement learning problems including MDP environments.

### Acknowledgment

We would like to thank Audrey Durand for her comments and insight on this project. We also thank department of family medicine of McGill University and CIHR for their generous support during this project.

### References

[Agrawal and Goyal, 2013] Shipra Agrawal and Navin Goyal. Further optimal regret bounds for thompson

- sampling. In *Artificial Intelligence and Statistics*, pages 99–107, 2013.
- [Auer *et al.*, 2002] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [Austin, 2011] Peter C Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*, 46(3):399–424, 2011.
- [Baransi *et al.*, 2014] Akram Baransi, Odalric-Ambrym Maillard, and Shie Mannor. Sub-sampling for multi-armed bandits. In *ECML-KDD*, pages 115–131, 2014.
- [Burnetas and Katehakis, 1996] Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- [Chapelle and Li, 2011] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257, 2011.
- [Chen, 2007] Song Xi Chen. Nonparametric estimation of expected shortfall. *Journal of financial econometrics*, 6(1):87–107, 2007.
- [El-Yaniv and Pechyony, 2009] Ran El-Yaniv and Dmitry Pechyony. Transductive rademacher complexity and its applications. *Journal of Artificial Intelligence Research*, 35(1):193, 2009.
- [Galichet *et al.*, 2013] Nicolas Galichet, Michele Sebag, and Olivier Teytaud. Exploration vs exploitation vs safety: Risk-aware multi-armed bandits. In *Asian Conference on Machine Learning*, pages 245–260, 2013.
- [Garcia and Fernández, 2015] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [Kuleshov and Precup, 2014] Volodymyr Kuleshov and Doina Precup. Algorithms for multi-armed bandit problems. *arXiv preprint arXiv:1402.6028*, 2014.
- [Lai and Robbins, 1985] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [Maillard, 2013] Odalric-Ambrym Maillard. Robust risk-averse stochastic multi-armed bandits. In *ICML*, pages 218–233, 2013.
- [Ripley *et al.*, 2013] Brian Ripley, Bill Venables, Douglas M Bates, Kurt Hornik, Albrecht Gebhardt, David Firth, and Maintainer Brian Ripley. Package ‘mass’. *Cran R*, 2013.
- [Robbins, 1985] Herbert Robbins. Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pages 169–177. Springer, 1985.
- [Sani *et al.*, 2012a] Amir Sani, Alessandro Lazaric, and Rémi Munos. Risk-aversion in multi-armed bandits. In *Advances in Neural Information Processing Systems*, pages 3275–3283, 2012.
- [Sani *et al.*, 2012b] Amir Sani, Alessandro Lazaric, and Rémi Munos. Risk-aversion in multi-armed bandits. In *NIPS*, pages 3275–3283, 2012.
- [Sutton and Barto, 1998] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [Tolstikhin, 2017] IO Tolstikhin. Concentration inequalities for samples without replacement. *Theory of Probability & Its Applications*, 61(3):462–481, 2017.