

Problem Set 3

For the Exercise Sessions on October 1 and 8

Last name	First name	SCIPER Nr	Points

Problem 1: Time-Varying Bandits

Let ν be a time varying environment with K arms where all the arms except arm i has distribution $\mathcal{N}(0, 1)$ and the i -th arm has distribution $\mathcal{N}(\Delta_t, 1)$. Note that the distributions changes with time – hence the name “time-varying bandits.” Let π be our policy, where we assume that the policy does not depend on time. Our time horizon is T .

- (a) Let $\Delta_t = \frac{1}{t^p}$, where $p \in (0, 1)$. Show that for every policy the regret is upper bounded by cT^{1-p} , where c is a constant. (HINT: Integrals are simpler than sums.)
- (b) Consider any policy whose regret scales as $o(T^{1-p})$. I.e.,

$$\lim_{T \rightarrow \infty} \frac{R_T(\nu, \pi)}{T^{1-p}} = 0$$

Show that for such a policy we must have

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T P_\nu(A_t \neq i) = 0$$

- (c) Now suppose that $p \in (\frac{1}{2}, 1)$. Let ν' be an environment which has the same distributions except at arm $i' \neq i$. At arm i' , ν' gives reward distributed from $\mathcal{N}(2\Delta_t, 1)$. Show that

$$\sup_{T \in \mathbb{N}} D(P_\nu(A_1, X_1, \dots, A_{T-1}, X_{T-1}) || P_{\nu'}(A_1, X_1, \dots, A_{T-1}, X_{T-1})) < \infty$$

- (d) Show that, if,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T P_\nu(A_t = i') = 0$$

then,

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T P_{\nu'}(A_t = i') < 1$$

(Hint: If $d_2(0||p)$ is a finite number, can p be very close to 1?)

- (e) Show that, if for $p \in (1/2, 1)$, $R(\nu, \pi) = o(T^{1-p})$, then

$$R_T(\nu', \pi) \neq o(T^{1-p})$$

- (f) Conclude that, if we are in an adversarial setting where the adversary is allowed to choose a time varying environment, no matter how long we play, \sqrt{T} is the best regret we can ever hope to achieve. That is, show that, for any policy π , for any $\alpha < 1/2$ there exists a time varying environment ν s.t.

$$R_T(\nu, \pi) \neq o(T^\alpha)$$

Problem 2: Bandits

In the course we mentioned *contextual bandits*. E.g., imagine that you suspect that the rewards that you get from the various arms depend on the day of the week (more generally, a particular feature of the input, aka the “context”). In this case it makes sense to run a separate bandit algorithm for each of the possible contexts. As we discussed, the downside of this approach is that we now have less data for each of the bandit algorithms and hence it will take us longer to learn.

Let c be the context, where $c \in \mathcal{C}$. We run a separate bandit algorithm for each of the $|\mathcal{C}|$ possible contexts and the total number of steps we take is n . We further assume that the number of arms is K and the K does not depend on the context. We use the same bandit algorithm for each of the contexts and we assume that this bandit algorithm has an expected regret of $O(\sqrt{Kn \log(n)})$ (when run over n time steps for a fixed context) in the stochastic setting.

Consider the following setup. For each context there is a different stochastic bandit with K arms. At each time a context is sampled independently at random, the player sees the context, chooses an arm, and receives a reward according to the distribution P_i^c (reward distribution for context c and arm i).

Define the expected regret for this setup as the difference between the expected reward that we could have gotten if for each context we would have chosen the arm with the largest mean minus the expected reward we get using our scheme.

1. Show that the expected regret at time n is $O(\sqrt{Kn|\mathcal{C}|\ln(n)})$.

Discussion: Compared to the case of a single context note that you “only” pay a factor of $\sqrt{|\mathcal{C}|}$. If the number of contexts is small, this might be acceptable. Note further, that this bound is valid, regardless how often each context appears. *Hint:* Start by assuming that each context $c \in \mathcal{C}$ appears n_c times.

2. Assume now that $\mathcal{C} = [0, 1)$, i.e., \mathcal{C} is very large, and in fact uncountable. In this case we cannot use the strategy above. What would you do in such a case? Under what circumstances will such a scheme likely work? If we assume a fixed number of arms and a very large time horizon, how would you expect the regret to scale?

Problem 3: Thompson Sampling with Bernoulli Losses

This problem deals with a Bayesian approach to multi-arm bandits. Although we will not pursue this facet in the current problem, the Bayesian approach is useful since within this framework it is relatively easy to incorporate prior information into the algorithm.

Assume that we have K bandits, and that bandit k outputs a $\{0, 1\}$ -valued Bernoulli random variable with parameter $\theta_k \in [0, 1]$. Let π be the uniform prior on $[0, 1]^K$, i.e., the uniform prior on the set of all parameters $\theta = (\theta_1, \dots, \theta_K)$. Let

$$T_k^1(t) = |\{\tau \leq t : A_\tau = k; Y_\tau = 1\}|,$$

$$T_k^0(t) = |\{\tau \leq t : A_\tau = k; Y_\tau = 0\}|.$$

In words, $T_k^1(t)$ is the number of times up to and including time t that we have chosen action k and the output of arm k was 1 and similarly $T_k^0(t)$ is the number of times up to and including time t that we have chosen action k and the output of the arm k was 0.

The goal is to find the arm with the highest parameter, i.e., the goal is to determine

$$k^* = \operatorname{argmax}_k \theta_k.$$

In the Bayesian approach we proceed as follows. At time t :

1. Compute for each arm k the distribution $p(\theta_k(t) | T_k^1(t-1), T_k^0(t-1))$.
2. Generate samples of these parameters according to their distributions.
3. Pick the arm j with the largest sample.
4. Observe the output of the j -th arm, call it $Y_j(t)$, and update the counters T_j^1 and T_j^0 accordingly.

Show that this algorithm “works” in the sense that eventually it will pick the best arm. More precisely, show the following two claims.

1. Show that $p(\theta_k(t) | T_k^1(t-1), T_k^0(t-1))$ is a Beta distributed and determine α and β .
2. Show that as t tends to infinity the probability that we choose the correct arm tends to 1. [HINT: To simplify your life, you can assume that for every arm k , $T_k^1(t-1) + T_k^0(t-1) \xrightarrow{t \rightarrow \infty} \infty$.]

NOTE: Recall that the density of the Beta distribution on $[0, 1]$ with parameters α and β is equal to

$$f(x; \alpha, \beta) = \text{constant } x^{\alpha-1} (1-x)^{\beta-1}.$$

Further, the expected value of $f(x; \alpha, \beta)$ is $\frac{\alpha}{\alpha+\beta}$ and its variance is $\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$.

Problem 4: Bandits with Infinitely Many Arms

In the course we considered bandits with a finite number of K arms. In this problem we will see that the same ideas apply if we have infinitely many arms as long as there is some additional structure.

Assume that there is an unknown unit-norm vector $\theta \in \mathbb{R}^d$. For every unit-norm vector $u \in \mathbb{R}^d$, there is a bandit. It gives the reward $X_u = \langle u, \theta \rangle + Z_u$, where Z_u is a zero-mean unit-variance Gaussian that is independent over time and independent with respect to different bandits. The nature of the reward is known to the player.

Find a policy, i.e., a strategy of what bandit to probe at any given point in time given a specific history, that has a sublinear regret as time tends to infinity. You can assume that you know the horizon, i.e., we are looking for fixed-horizon policies.