

EECS598: Prediction and Learning: It's Only a Game

Fall 2013

## Lecture 19: UCB Algorithm and Adversarial Bandit Problem

Prof. Jacob Abernethy

Scribe: Hossein Keshavarz

**Announcements**

- There is no class on Wednesday, November 27
- Good job on project proposals so far!

**19.1 Review on stochastic multi-armed bandit problem**

Consider a gambler playing against a  $N$ -armed slot machine and sequentially pulls up arms to minimize the expected regret. Each arm,  $i \in \{1, \dots, N\}$  has a probability distribution  $\mathcal{D}_i$  supported on the closed interval  $[0, 1]$ . Let  $\mu_i$  be the expected value of the corresponding loss to arm  $i$  and define  $\tilde{I} = \operatorname{argmin}_{1 \leq j \leq N} \mu_j$  as the index of the optimal expected value among arms. Suppose that the gambler selects arm  $I_t$  at round  $t$ , then the expected regret of the gambler up to round  $T$  is defined as

$$\mathbb{E} - \text{regret} = \mathbb{E} \left( \sum_{t=1}^T (X_{I_t, t} - X_{\tilde{I}, t}) \right)$$

where  $X_{i,t} \sim \mathcal{D}_i$  denotes the loss associated with choosing arm  $i$  at round  $t$ . It's worthwhile to mention that the expected value is taken over the distributions  $\{\mathcal{D}_i\}_{i=1}^N$  and the randomness of choosing arms.

**19.2 Analysis of greedy algorithm**

This section is devoted to state and prove the theorem regarding the upper bound on the expected regret of the greedy algorithm. The greedy algorithm is introduced in the last lecture.

**Theorem 19.1.** *Suppose that there exists a positive scalar  $\Delta$  such that  $\mu_j - \mu_{\tilde{I}} \geq \Delta$  for any  $j \neq \tilde{I}$ . Then the expected regret of the greedy algorithm has the following upper bound*

$$\mathbb{E} - \text{regret} \leq 1 + \frac{2N \log(2NT)}{\Delta^2}$$

*Proof.* Let us to decompose the expected regret as sum of two terms in which the first term is the associated regret to the sampling phase and the second term corresponds to the exploitation phase.

$$\begin{aligned}
\mathbb{E} - \text{regret} &= \mathbb{E} \left( \sum_{t=1}^{mN} (X_{I_t,t} - X_{\tilde{I},t}) \right) + \mathbb{E} \left( \sum_{t=1+mN}^T (X_{I_t,t} - X_{\tilde{I},t}) \right) \stackrel{(a)}{\leq} Nm + \mathbb{E} \left( \sum_{t=1}^{mN} (X_{I_t,t} - X_{\tilde{I},t}) \right) \\
&\stackrel{(b)}{\leq} Nm + \mathbb{E} \left( \sum_{t=1}^{mN} \mathbb{I}_{I_t \neq \tilde{I}} \right) = Nm + \sum_{t=1+mN}^T \mathbb{P}(I_t \neq \tilde{I}) \leq Nm + T \max_{1+Nm \leq t \leq T} \mathbb{P}(I_t \neq \tilde{I})
\end{aligned} \tag{19.1}$$

Note that inequalities (a) and (b) are immediate consequence of the fact that  $X_{k,t} \in [0,1]$  for all  $1 \leq k \leq N$  and  $1 \leq t \leq T$ . Therefore, in order to complete the proof, we need to control the deterministic term  $\mathbb{P}(I_t \neq \tilde{I})$  uniformly from above. If  $I_t \neq \tilde{I}$  then  $\hat{\mu}_{I_t} \leq \hat{\mu}_{\tilde{I}}$ , hence using some straightforward algebraic manipulation we obtain

$$\Delta \leq \mu_{I_t} - \mu_{\tilde{I}} = (\mu_{I_t} - \hat{\mu}_{I_t}) + (\hat{\mu}_{I_t} - \hat{\mu}_{\tilde{I}}) + (\hat{\mu}_{\tilde{I}} - \mu_{\tilde{I}}) \leq |\mu_{I_t} - \hat{\mu}_{I_t}| + 0 + |\hat{\mu}_{\tilde{I}} - \mu_{\tilde{I}}| \leq 2 \max_{1 \leq j \leq N} |\mu_j - \hat{\mu}_j|$$

In other words,  $\mathbb{P}(I_t \neq \tilde{I}) \leq \mathbb{P}(\max_{1 \leq j \leq N} |\mu_j - \hat{\mu}_j| \geq \frac{\Delta}{2})$  for any  $t$ . Combination of union bound (inequality (a)) and Hoeffding's inequality (used to prove inequality (b)) give us the desired upper bound.

$$\begin{aligned}
\eta = \max_{1+Nm \leq t \leq T} \mathbb{P}(I_t \neq \tilde{I}) &\leq \mathbb{P} \left( \max_{1 \leq j \leq N} |\mu_j - \hat{\mu}_j| \geq \frac{\Delta}{2} \right) \stackrel{(a)}{\leq} N \max_{1 \leq j \leq N} \mathbb{P} \left( |\mu_j - \hat{\mu}_j| \geq \frac{\Delta}{2} \right) \\
&\stackrel{(b)}{\leq} 2N \exp \left( -\frac{m\Delta^2}{2} \right)
\end{aligned} \tag{19.2}$$

Letting the optional parameter  $m = \frac{2}{\Delta^2} \log \left( \frac{2N}{\eta} \right)$  and substituting inequality (19.2) into inequality (19.1) yield

$$\mathbb{E} - \text{regret} \leq Nm + T\eta = \frac{2N}{\Delta^2} \log \left( \frac{2N}{\eta} \right) + T\eta$$

Choosing  $\eta = \frac{1}{T}$  terminates the proof.  $\square$

Note that the reason that  $\Delta^2$  appeared in the denominator of the expected regret is that the inequality  $\mathbb{E} \left( \sum_{t=1}^{mN} (X_{I_t,t} - X_{\tilde{I},t}) \right) \leq Nm$  is not tight enough.

**Lemma 19.2** (Hoeffding's inequality). *Let  $Z_1, \dots, Z_n$  are independent and identically distributed random variables such that  $Z_1 \in [0, 1]$  almost sure and  $\mathbb{E}(Z_1) = \mu$ . Then,*

$$\mathbb{P} \left( \left| \frac{1}{n} \sum_{j=1}^n Z_j - \mu \right| \geq \epsilon \right) \leq 2 \exp(-2n\epsilon^2)$$

### 19.3 Upper Confidence Bound (UCB) algorithm

In order to introduce UCB algorithm let us to introduce some notation. Define,  $T_j(t) := \sum_{s=1}^t \mathbb{I}_{(I_s=j)}$  as the number of times that the gambler played the arm  $j$  up to time  $t$  and let the empirical expected regret of arm  $j$  as the following

$$\hat{\mu}_{j,t} = \frac{1}{|T_j(t)|} \sum_{s=1}^t X_{j,s} \mathbb{I}_{(I_s=j)} \quad (19.3)$$

<p><b>Algorithm 1.</b> UCB algorithm to solve multi armed bandit problem</p> <p><b>Initialization</b> Play each arm once</p> <p><b>For</b> <math>t = 1</math> to <math>T</math></p> <p>    <b>For</b> <math>j = 1</math> to <math>N</math></p> <p>        Update <math>\hat{\mu}_{j,t}</math> according to the identity (19.3)</p> <p>    end</p> <p>    Play the arm <math>I_t = \operatorname{argmin}_{1 \leq j \leq N} \left( \hat{\mu}_{j,t} - \sqrt{\frac{3 \log t}{ T_j(t) }} \right)</math></p> <p>end</p>
---

The arm selection criterion of UCB algorithm encourages to choose the rarely chosen arms. The next theorem characterizes the upper bound on the expected regret of UCB algorithm. The interested reader is referred to [1] for further details and proof of Theorem 19.3.

**Theorem 19.3.** Let  $\Delta_j = \mu_j - \mu_{\tilde{I}}$  such that  $\Delta_j > 0$  for any  $j \neq \tilde{I}$ . Then the expected regret of UCB algorithm can be upper bounded by the following inequality.

$$\mathbb{E} - \text{regret} \leq \mathcal{O} \left( \sum_{j \neq \tilde{I}} \frac{\log T}{\Delta_j} \right)$$

Theorem 19.3 shows that the expected regret of UCB algorithm is far better than greedy method for small  $\Delta_j$ 's. Strictly speaking, if  $\Delta_{\min} = \min_{j \neq \tilde{I}} \Delta_j$ , then

$$\mathbb{E} - \text{regret} \leq \mathcal{O} \left( \frac{N \log T}{\Delta_{\min}} \right)$$

### 19.4 Adversarial Bandit

Unlike the multi armed bandit setting where the loss of each action has a stationary distribution over time, in the adversarial bandit problem there is no statistical assumption about the form of the generating process of losses. In this new formulation, the associated regret to each arm is determined at each round by an adversary and the player only knows the reward of previously chosen actions. The only assumption about the loss vector is that,  $\ell^t \in [0, 1]^N$  for each round  $t$ . Since the adversary can assign low reward to the previously selected actions and high rewards to the unseen arms, hence, a deterministic policy of arm selection can not optimize the expected

regret function. Finally, we need to mention that the distribution of action  $I_t$  only depends on the loss of previous actions,  $\{\ell_{I_s}^s\}_{s=1}^{t-1}$ .

Due to the lack of full information about the associated losses of choosing arms, the player can't run exponential weighted algorithm to optimize the regret function. One doubtful solution is to estimate of loss vector,  $\ell^t$ , based on the observation of a single component  $\ell_{I_t}^t$ . Taking advantage of the conditional expectation property, we can show that if  $\tilde{\ell}^t$  is an unbiased estimator of  $\ell^t$ , i.e.  $\mathbb{E}(\tilde{\ell}^t | \mathcal{F}_{t-1}) = \ell^t$  where  $\mathcal{F}_{t-1}$  is the generated  $\sigma$ -field by observations up to round  $t-1$ , then the expected regrets with respect to  $\ell^t$  and  $\tilde{\ell}^t$  are the same.

$$\mathbb{E}\left(\sum_{t=1}^T \tilde{\ell}^t \cdot (p^t - p^*)\right) = \mathbb{E}\left\{\mathbb{E}\left(\sum_{t=1}^T \tilde{\ell}^t \cdot (p^t - p^*) | \mathcal{F}_{t-1}\right)\right\} = \mathbb{E}\left\{\sum_{t=1}^T \mathbb{E}(\tilde{\ell}^t | \mathcal{F}_{t-1}) \cdot (p^t - p^*)\right\} = \mathbb{E}\left(\sum_{t=1}^T \ell^t \cdot (p^t - p^*)\right)$$

#### 19.4.1 Unbiased estimation of $\ell^t$

We claim that the following procedure which is called exponential weighted algorithm with  $\epsilon$ -exploration generates an unbiased estimator of  $\ell^t$ .

1. With probability  $\epsilon$ , choose  $p^t = \frac{1}{N}\langle 1, \dots, 1 \rangle$  and select  $I_t \sim p^t$  (uniformly at random). Let  $\tilde{\ell}^t = \langle 0, \dots, 0, \frac{N\ell_{I_t}^t}{\epsilon}, 0, \dots, 0 \rangle$ .
2. With probability  $1-\epsilon$  choose  $p^t$  by exponential weighted algorithm on the loss vectors  $\{\tilde{\ell}^s\}_{s=1}^{t-1}$  and let  $\tilde{\ell}^t = 0$ .

*proof of claim.*

$$\mathbb{E}(\tilde{\ell}^t) = (1-\epsilon) \cdot 0 + \epsilon \sum_{j=1}^N \frac{1}{N} \langle 0, \dots, 0, \frac{N\ell_j^t}{\epsilon}, 0, \dots, 0 \rangle = \ell^t$$

□

Since  $\tilde{\ell}^t$  is an unbiased estimator of  $\ell^t$ , so at the first glance, it seems that the expected regret of above algorithm is exactly equal to the expected regret of EWA. However, a contingent reader notices that  $\tilde{\ell}^t$  is no longer in the closed cube  $[0, 1]^N$ . Recalling the proof of Theorem 3.2. in the lecture notes, one can easily show that

$$\mathbb{E}\text{-regret of EWA with } \epsilon \text{ exploration} \leq T\epsilon + (1-\epsilon) \left( \frac{\log N}{\eta} + \eta \sum_{t=1}^T \|\tilde{\ell}^t\|_\infty^2 \right) = \mathcal{O}\left(T\epsilon + \frac{\log N}{\eta} + \eta T \frac{N^2}{\epsilon^2}\right)$$

Now choosing the regularization parameters  $\epsilon^2 = \frac{N}{T} \sqrt{2T \log N}$  and  $\eta = \frac{\epsilon}{N} \sqrt{\frac{2 \log N}{T}}$ , the optimal upper bound is given by

$$\mathbb{E}\text{-regret of EWA with } \epsilon \text{ exploration} \leq \mathcal{O}\left(\sqrt{NT}^{\frac{3}{4}} (\log N)^{\frac{1}{4}}\right) \quad (19.4)$$

It's worthwhile to mention that although  $\max_{1 \leq t \leq T} \|\tilde{\ell}^t\|_\infty \leq \frac{N}{\epsilon}$ , but most of the times (with probability  $1 - \epsilon$ ), we have  $\tilde{\ell}^t = 0 \in [0, 1]$ . Hence, the proof can be slightly modified in a smart way to obtain the following upper bound.

$$\mathbb{E} - \text{regret of EWA with } \epsilon \text{ exploration} \leq \mathcal{O}\left(T\epsilon + \frac{\log N}{\eta} + \eta T \frac{N^2}{\epsilon}\right)$$

which leads to  $\mathbb{E} - \text{regret of EWA with } \epsilon \text{ exploration} \leq \mathcal{O}\left(\sqrt{NT}^{\frac{2}{3}}\right)$ .

**Question:** Is there any algorithm with the expected regret  $\mathcal{O}\left(\sqrt{NT}\right)$ ? Yes! EXP<sub>3</sub> algorithm

### 19.5 EXP<sub>3</sub> Algorithm

Let  $\tilde{L}^t$  be the cumulative loss up to round  $t$ .

<p><b>Algorithm 2.</b> EXP<sub>3</sub> algorithm to for adversarial multi-armed bandit problem</p> <p><b>Input</b> Regularization parameter <math>\eta</math></p> <p><b>Initialization</b> St initial value for <math>p^1</math></p> <p><b>For</b> <math>t = 1</math> to <math>T</math></p> <p>    <b>For</b> <math>j = 1</math> to <math>N</math></p> <p>        Sample <math>I_t</math> absed on the distribution <math>p^t</math></p> <p>        Observe <math>\ell_{I_t}^t</math></p> <p>        Let <math>\tilde{\ell}^t = \langle 0, \dots, 0, \frac{\ell_{I_t}^t}{p_{I_t}^t}, 0, \dots, 0 \rangle</math></p> <p>        Update <math>p^{t+1}</math> by <math>p_j^{t+1} = \frac{\exp(-\eta \tilde{L}_j^t)}{\sum_{j=1}^N \exp(-\eta \tilde{L}_j^t)}</math></p> <p>    <b>end</b></p>
---

The EXP<sub>3</sub> algorithm will be analyzed in the next class.

## References

- [1] P. Auer, N. Cesa-Bianchi and P. Fischer, "Finite-time analysis of the multi-armed bandit problem", *Machine learning* 47, no. 2-3 (2002): 235-256.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund and R. E. Schapire, "Gambling in a rigged casino: The adversarial multi-armed bandit problem", *In Foundations of Computer Science, 1995*. Proceedings., 36th Annual Symposium on, pp. 322-331. IEEE, 1995.