# CPSC 532D — 7. ONLINE LEARNING

Bingshan Hu with Danica J. Sutherland University of British Columbia, Vancouver Fall 2024

In batch (offline) learning, usually, we have a learning phase and a prediction phase. Learner first receives a batch of i.i.d. training samples and then uses these data to learn a hypothesis, e.g., an ERM (learning phase). The learned hypothesis will be used for predicting the labels of future samples (prediction phase). In contrast, in online learning, there is no separation between the learning phase and the prediction phase. We blend these two phases together by specifying a learning protocol that regulates all parties participating in the learning.

# 7.1 ONLINE BINARY CLASSIFICATION IN REALIZABLE SETTING

Still, we have an instance space  $\mathcal{X}$ , a label space  $\mathcal{Y} = \{-1, 1\}$ , a hypothesis class  $\mathcal{H} : \mathcal{X} \to \mathcal{Y}$ , and the 0 - 1 loss function  $\ell_{0-1}(y, \hat{y}) = \mathbf{1} \{y \neq \hat{y}\}$ . We play this game sequentially with the following learning protocol.

In each round  $t = 1, 2, \ldots, T$ ,

- 1. Nature (Adversary/Environment) selects  $x_t \in \mathcal{X}$  and reveals it to Learner;
- 2. Learner chooses a hypothesis  $f_t \in \mathcal{H}$  and predicts  $\hat{y}_t = f_t(x_t) \in \{-1, 1\}$ ;
- 3. Nature plays label  $y_t$  and <u>reveals</u> it to Learner ;
- 4. Learner <u>obtains</u> a data sample  $(x_t, y_t)$  and <u>suffers</u> loss  $\ell(y_t, \hat{y}_t) = \mathbf{1} \{y_t \neq \hat{y}_t\}$ .

If  $\ell(y_t, \hat{y}_t) = \mathbf{1} \{y_t \neq \hat{y}_t\} = 1$ , we say Learner makes a mistake in round *t*, as its predicted label is not correct. The goal of Learner is to make as few mistakes as possible. Intuitively, to choose a good predictor  $f_t$  in round *t*, Learner should use all the data samples attained in previous rounds  $S_{t-1} := ((x_1, y_1), (x_2, y_2), \dots, (x_{t-1}, y_{t-1}))$  and even the input  $x_t$  in round *t*.

The data sequence  $((x_1, y_1), (x_2, y_2), \dots, (x_T, y_T))$  does not need to be iid, which is quite different from the assumption usually batch learning needs.

Since Adversary can decide  $y_t$  based on  $\hat{y}_t$ , it is hopeless to learn for some some hypothesis class  $\mathcal{H}$ , e.g., two constant functions  $x \mapsto +1$  and  $x \mapsto -1$  are in  $\mathcal{H}$ . Because Adversary can make Learner unhappy in each round by declaring  $y_t = -\hat{y}_t$ .

So, we need to put some constraints for Adversary.

We get started with a simple learning problem setup. It is about the aforementioned binary classification problem under the assumption of *realizability*, i.e., the true label function  $f^* \in \mathcal{H}$ , i.e.,  $y_t = f^*(x_t)$  for all  $t \in [T]$ , and Learner knows  $\mathcal{H}$  and the fact that  $f^* \in \mathcal{H}$ . <sup>1</sup> With this restriction on how Adversary generates the data sequence,

Chapter 21.1 of [SSBD14].

This is a picture for this.

No statistical assumption.

With the assumption of realizability, setting  $y_t = -\hat{y}_t$  for all  $t \in [T]$  may not always be possible.

For more, visit https://cs.ubc.ca/~dsuth/532D/24w1/.

<sup>&</sup>lt;sup>1</sup>Nature is regulated to play a label function that is consistent with the history data  $S_{t-1}$ .

Learner should make as few mistakes as possible. We are interested in how many mistakes Learner makes after playing this sequential game after T rounds.

**Hypothesis class**  $\mathcal{H}$  is finite. Since we know the true label function  $f^* \in \mathcal{H}$ , i.e., there is a perfect hypothesis incurring zero loss in  $\mathcal{H}$ , it is quite natural to eliminate any hypothesis in  $\mathcal{H}$  that has made a mistake after observing  $(x_t, y_t)$  at the end of each round t. This intuition gives us an idea, called *version space*, to develop learning algorithms. The version space at the end of round t is defined as  $\mathcal{H}_t = \{f \in \mathcal{H} : f(x_s) = y_s, \forall s \in [t]\}$ . Then, we can maintain a sequence of version spaces  $\mathcal{H} = \mathcal{H}_0 \supseteq \mathcal{H}_1 \supseteq \ldots \supseteq \mathcal{H}_{t-1} \supseteq \mathcal{H}_t \supseteq \ldots \supseteq \{f^*\}$ . Hypotheses that have made mistakes will be kicked out from the version space at the end of each round t. We eliminate hypotheses from  $\mathcal{H}$  during the learning once we are confident they are not good almost surely!

Let's see Consistent Algorithm first.

Initially, we set  $\mathcal{H}_0 = \mathcal{H}$ . In each round  $t = 1, 2, \ldots, T$ , 1. Learner receives  $x_t$ ; 2. Learner chooses  $f_t \in \mathcal{H}_{t-1}$  arbitrarily and predicts  $\hat{y}_t = f_t(x_t)$ ; 3. Nature reveals true label  $y_t = f^*(x_t)$  and Learner updates  $\mathcal{H}_t =$  $\{f \in \mathcal{H}_{t-1} : f(x_t) = y_t\}.$ 

Note that the number of the learning horizon T. For any length of T, the number of mistakes is at most  $|\mathcal{H}| - 1$ .

Note that the number of mistakes does not depend on Now, we exam how many mistakes (the total loss  $\sum_{t=1}^{T} \ell(y_t, \hat{y}_t) = \sum_{t=1}^{T} \ell(y_t, f_t(x_t))$ ) that the learning horizon T. For Consistent Algorithm makes by the end of round T. It is not hard to see that the number of mistakes is at most  $|\mathcal{H}| - 1$ , as each mistake forces Learner to eliminate at least one hypothesis from the version space at the end of that round.

f(x) = o(g(x)) is equivalent One may ask is it possible to reduce the number of mistakes from  $|\mathcal{H}|$  to  $o(|\mathcal{H}|)$ ? to  $\lim_{x\to\infty} \frac{f(x)}{g(x)} = 0$  if g(x) > 0. The answer is yes!

> Actually, we can do better if we do not pick  $f_t$  in an arbitrary way! We introduce another idea, called *majority vote*. Since some hypotheses in  $\mathcal{H}_{t-1}$  predict +1 while some predict -1, we simply count which side has more supporters. The side of more supporters wins! Combining majority vote with version space, we have a much nicer algorithm, called Halving Algorithm, which reduces the number of mistakes from  $\mathcal{O}(|\mathcal{H}|)$  to log  $|\mathcal{H}|$ .

Initially, we set  $\mathcal{H}_0 = \mathcal{H}$ . In round t = 1, 2, ..., T, 1. Learner receives  $x_t$ ; 2. Learner predicts  $\hat{y}_t \in \arg \max |\{f \in \mathcal{H}_{t-1} : f(x_t) = y\}|$ ;  $y \in \{-1, 1\}$ 3. Nature reveals true label  $y_t = f^*(x_t)$  and Learner updates  $\mathcal{H}_t =$  $\{f \in \mathcal{H}_{t-1} : f(x_t) = y_t\}.$ 

**THEOREM 7.1.** Halving Algorithm makes at most  $\log_2(|\mathcal{H}|)$  mistakes.

*Proof.* According to the majority vote, the version space is halved on each mistake. If Learner makes a mistake in round t, we have  $|\mathcal{H}_t| \leq \frac{1}{2} |\mathcal{H}_{t-1}|$ . Let M be the total

number of mistakes. We have

$$1 \le |\mathcal{H}_{T}| \le |\mathcal{H}_{0}| 2^{-M} = |\mathcal{H}| 2^{-M}$$

which yields  $M \leq \log_2(|\mathcal{H}|)$ . Note that  $f^*$  is always in the version spaces. So, we have  $1 \leq |\mathcal{H}_{\mathrm{T}}|$ .

**GENERAL HYPOTHESIS CLASS INCLUDING**  $\mathcal{H}$  **IS INFINITE.** Still under the assumption of realizability, now, let us consider a hypothesis class  $\mathcal{H}$  that may be infinite in size. For some hypothesis class  $\mathcal{H}$ , Learner is able to have a strategy that guarantees it makes a finite number of mistakes. For other hypothesis classes, Nature can force Learner to make infinitely many mistakes. To gain an understanding of which hypothesis class falls in the former category and which falls in the latter, we need to have new concepts, something like VC dimension to measure the richness of a hypothesis class. We aim at characterizing online learnability. In particular, we target the following question: What is the optimal online classification learning algorithm for a given hypothesis class  $\mathcal{H}$ ?

Littlestone dimension characterizes online learnability. A hypothesis  $\mathcal{H}$  is online learnable if it has a finite Littlestone dimension denoted as  $Ldim(\mathcal{H})$ . The idea of Littlestone dimension is to view online learning as a 2-player sequential game between Learner and Adversary. The job of Adversary is to force Learner to make mistakes while preserving realizability.

How does Adversary choose  $x_t$  to force Learner to make the maximum number of mistakes, while ensuring realizability? It is easy as Adversary can always choose  $y_t = -\hat{y}_t$  for the first Ldim( $\mathcal{H}$ ) rounds in the sequential game.

The strategy for Adversary can be formally described by using a complete binary tree. Each node of the tree is associated with an instance  $x_t \in \mathcal{X}$ . If Learner predicts  $\hat{y}_t = +1$ , Adversary will declare the prediction is wrong, i.e.,  $y_t = -\hat{y}_t = -1$  and will traverse to the right child of the current node. If Learner predicts  $\hat{y}_t = -1$ , Adversary will set  $y_t = -\hat{y}_t = +1$  and will traverse to the left child of the current node.

To introduce the definition of Litterstone dimension, let us give the definition of an  $\mathcal{H}$  shattered tree first.

Add a picture of the binary tree.

The depth of the tree is defined as the number of to a leaf, i.e., the number of

**DEFINITION 7.2.** ( $\mathcal{H}$  shattered tree.) A shattered tree of depth *d* is a sequence of *edges in a path from the root* inputs  $v_1, v_2, \ldots, v_{2^d-1} \in \mathcal{X}$  such that for every root-to-leaf path,  $\exists f^* \in \mathcal{H}$  such that  $l_{avers in the tree.}$ all labels along the path are achieved.

**DEFINITION 7.3.** (Littlestone dimension.) The Littlestone dimension of hypothesis class  $\mathcal{H}$ , Ldim( $\mathcal{H}$ ), is the maximal integer d such that there exists a shattered tree of depth *d* that  $\mathcal{H}$  shatters. If there is no such largest *d*, then the Litterstone dimension is infinite.

**THEOREM 7.4.** Any algorithm makes at least  $Ldim(\mathcal{H})$  mistakes.

*Proof.* Let  $d_* = \text{Ldim}(\mathcal{H})$ . Since  $\mathcal{H}$  has Littlestone dimension  $d_*$ , we know there exists an  $\mathcal{H}$  shattered tree with depth  $d_*$  that is shattered by  $\mathcal{H}$ . Adversary will "walk" on the tree for the first  $d_*$  rounds. In each round  $t \in [d_*]$ , Adversary sets  $y_t = -\hat{y}_t$ . Since the walk continues for  $d_*$  rounds, Learner makes  $d_*$  mistakes. Now, we check the assumption of realizability. Since  $\mathcal{H}$  shatters the tree, there  $\exists f_* \in \mathcal{H}$  such that all the labels along the root-to-leaf path selected by Adversary can be realized.  Now, we show a learning algorithm, **Standard Optimal Algorithm (SOA)**, makes at most  $Ldim(\mathcal{H})$  mistakes. The algorithm is similar to **Halving Algorithm**. We partition the version space by the end of round t - 1 into two sub-version spaces. Let  $\mathcal{H}_{t-1}^{(-1)} := \{f \in \mathcal{H}_{t-1} : f(x_t) = -1\}$  and  $\mathcal{H}_{t-1}^{(+1)} := \{f \in \mathcal{H}_{t-1} : f(x_t) = +1\}$ .

Initially, we set  $\mathcal{H}_0 = \mathcal{H}$ . In round t = 1, 2, ..., T, 1. Learner receives  $x_t$ ; 2. Learner predicts  $\hat{y}_t \in \underset{y \in \{-1,1\}}{\operatorname{arg max}} \operatorname{Ldim}\left(\mathcal{H}_{t-1}^{(y)}\right)$ ; 3. Nature reveals true label  $y_t = f^*(x_t)$  and Learner updates  $\mathcal{H}_t = \{f \in \mathcal{H}_{t-1} : f(x_t) = y_t\}$ .

**THEOREM 7.5.** SOA makes at most  $Ldim(\mathcal{H})$  mistakes.

*Proof.* We claim that whenever Learner makes a mistake in some round *t*, we have  $Ldim(\mathcal{H}_t) \leq Ldim(\mathcal{H}_{t-1}) - 1$  at the end of that round. Without loss of generality, we assume Learner predicts -1, but the true label is +1 in that round.

We use contradiction to complete the proof of the claim. Suppose  $Ldim(\mathcal{H}_t) = Ldim(\mathcal{H}_{t-1})$ .

We have the following inequalities.

- 1.  $\operatorname{Ldim}(\mathcal{H}_{t-1}^{(-1)}) \leq \operatorname{Ldim}(\mathcal{H}_{t-1})$ ;
- 2.  $\operatorname{Ldim}(\mathcal{H}_t) = \operatorname{Ldim}(\mathcal{H}_{t-1}^{(+1)})$ ;
- 3.  $\operatorname{Ldim}(\mathcal{H}_{t-1}^{(-1)}) \ge \operatorname{Ldim}(\mathcal{H}_{t-1}^{(+1)})$ .

From inequalities 2 and 3 above, we have  $\text{Ldim}(\mathcal{H}_{t-1}^{(-1)}) \ge \text{Ldim}(\mathcal{H}_t) = \text{Ldim}(\mathcal{H}_{t-1})$ , where the equality uses the assumption that  $\text{Ldim}(\mathcal{H}_t) = \text{Ldim}(\mathcal{H}_{t-1})$ . Combining with inequality 1, we have

$$\operatorname{Ldim}(\mathcal{H}_{t-1}^{(-1)}) = \operatorname{Ldim}(\mathcal{H}_{t-1}) = \operatorname{Ldim}(\mathcal{H}_t) = \operatorname{Ldim}(\mathcal{H}_{t-1}^{(+1)}).$$

Since  $\text{Ldim}(\mathcal{H}_{t-1}^{(-1)}) = \text{Ldim}(\mathcal{H}_{t-1}^{(+1)})$ , we can construct a shattered tree of depth  $\text{Ldim}(\mathcal{H}_{t-1}) + 1$  that can be shattered by  $\mathcal{H}_{t-1}$ , which yields a contradiction.

**COROLLARY 7.6.** Let  $\mathcal{H}$  be any hypothesis class. Then, SOA makes exactly  $Ldim(\mathcal{H})$  mistakes.

*Proof.* From Theorem 7.4, we know SOA makes at least  $Ldim(\mathcal{H})$  mistakes. From Theorem 7.5, we know SOA makes at most  $Ldim(\mathcal{H})$  mistakes. Combining these two results concludes the proof.

VC dimension vs Littlestone Dimension.

**EXAMPLE 1.** If  $\mathcal{H}$  is a finite hypothesis class, we have  $Ldim(\mathcal{H}) \leq \log_2(|\mathcal{H})|$ . why?

**EXAMPLE 2.** Let  $\mathcal{X} = [0, 1]$  and  $\mathcal{H} = \{x \mapsto \mathbf{1} \{x < a\} : a \in [0, 1]\}$  be the class of thresh-*There is a picture for this.* olds on the interval [0, 1]. Then, we have VCdim $(\mathcal{H}) = 1$ , but Ldim $(\mathcal{H}) = \infty$ . **THEOREM 7.7.** For any hypothesis class  $\mathcal{H}$ , we have  $VCdim(\mathcal{H}) \leq Ldim(\mathcal{H})$ . Further, the gap can be arbitrarily large.

*Proof.* We first prove that  $VCdim(\mathcal{H}) \leq Ldim(\mathcal{H})$ . Suppose  $VCdim(\mathcal{H}) = d$  and let  $\{x_1, x_2, \ldots, x_d\}$  be a shattered set by  $\mathcal{H}$ . Now, we construct a complete binary tree of inputs  $v_1, v_2, \ldots, v_{2^d-1}$ , where all nodes at depth *i* are set to be  $x_i$ .

Now, the definition of a shattered set clearly implies that we constructed a valid shattered tree of depth *d* and conclude that  $VCdim(\mathcal{H}) \leq Ldim(\mathcal{H})$ .

The class of threshold functions on the unit interval has VC dimension of 1, whereas its Littlestone dimension is infinite.  $\hfill \Box$ 

There is a picture for this.

## 7.2 DECISION-THEORETIC ONLINE LEARNING AND EXPONENTIAL WEIGHTS (HEDGE)

Practically, we should not assume realizability always holds, i.e., the true labels  $y_t = f^*(x_t), \forall t \in [T]$ , are generated by using  $f^*$ . Similar to agnostic setting in supervised learning, now, we can compare with the best predictor in  $\mathcal{H}$ . In online learning, we use the notion of regret, defined as

$$\mathcal{R}(\mathbf{T}) := \sup_{((x_1, y_1), \dots, (x_T, y_T))} \left\{ \sum_{t=1}^{T} \ell(y_t, f_t(x_t)) - \min_{f \in \mathcal{H}} \sum_{t=1}^{T} \ell(y_t, f(x_t)) \right\} , \qquad (7.1)$$

to measure the performance gap between  $\sum_{t=1}^{T} \ell(y_t, f_t(x_t))$ , the total loss Learner has

made, and  $\min_{f \in \mathcal{H}} \sum_{t=1}^{T} \ell(y_t, f(x_t))$ , the total loss of the best predictor in hindsight. If a learning algorithm achieves an o(T) regret bound, we say it is a *no-regret* algorithm. The intuition is if you play this game long enough, you can compete with the best predictor in hindsight. Note that o(T) regret implies  $\lim_{T \to \infty} \frac{o(T)}{T} = 0$ .

To control the regret at Learner's side, generally, we have two key principles.

1. Randomization. Deterministic Learner fails for some learning problems. Let hypothesis class  $\mathcal{H} = \{f_+ : x \to +1, f_- : x \to -1\}$  only contain two constant functions. Adversary can always give Learner  $y_t = -\hat{y}_t$  to force Learner to suffer loss. So, Learner suffers in total T loss. Now, we investigate the total

loss in hindsight of the best predictor in  $\mathcal{H}$ , that is,  $\min_{f \in \mathcal{H}} \sum_{t=1}^{T} \ell(y_t, f(x_t)) \leq \frac{1}{2}$ .

$$\left(\sum_{t=1}^{1} \ell(y_t, f_{-1}(x_t)) + \sum_{t=1}^{1} \ell(y_t, f_{+}(x_t))\right) = 0.5T$$
, which yields the regret is at least 0.5T.

Since Learner always fails if revealing the predicted label  $\hat{y}_t$ , to make the problem interesting, we need to give some power to Learner. Now, we change the learning protocol a bit by allowing Learner to reveal only a probability distribution over  $\{-1, 1\}$  instead of revealing the predicted label  $\hat{y}_t$  itself.

2. *Exploitation.* We want to track the empirical performance of each predictor, but we do not eliminate any of them during the learning, as some of them may not perform well at the beginning of the learning, but later turns out to the best one.

Since we do not plan to eliminate predictors during the learning, we can maintain a (*data-dependent*) *distribution* over all the predictors in  $\mathcal{H}$  taking account of each predictor's total loss suffered so far. For some predictor performing poorly, we put a small mass on it. We eliminate some predictor in a soft way!

**DECISION-THEORETIC ONLINE LEARNING (DTOL).** Since Learner is allowed to reveal a probability distribution over the labels and revealing a distribution over  $\{-1, 1\}$  can be translated to revealing a distribution over  $\mathcal{H}$ , now, we can re-formulate the online learning problem slightly and the modification will be, generally, useful for bandit problems (and even reinforcement learning problems). We do not use training samples at all. We get rid of the input space, label space, and hypothesis class. Instead, we have a fixed set of K actions, denoted by [K]. The learning protocol

Here, exploitation refers to utilizing the already learned information, i.e., the total loss of each predictor observed so far.

If the data sequence are i.i.d. according to some fixed but unknown distribution, we can eliminate predictors.

You can view each action in [K] as a hypothesis in H. is modified as follows.

In each round  $t = 1, 2, \ldots, T$ ,

- 1. Learner plays a probability distribution  $p_t \in [0, 1]^K$  over all actions ;
- 2. Nature plays loss vector  $\ell_t = (\ell_{1,t}, \ell_{2,t}, \dots, \ell_{K,t})$ ;
- 3. Learner observes  $\ell_t$  and suffer a loss  $\langle p_t, \ell_t \rangle = \sum_i p_{j,t} \ell_{j,t}$ .

**Remarks.** (1) A probability distribution  $p_t = (p_{1,t}, p_{2,t}, ..., p_{K,t})$  is a vector with all  $p_{j,t} \ge 0$  and  $\sum_j p_{j,t} = 1$ . (2) In Step 2, Nature can give  $\ell_t$  based on all the past information and even  $p_t$ .

For DTOL, we adapt the notion of regret shown in (7.1) to

$$\mathcal{R}(\mathbf{T}) := \sup_{\ell_1, \ell_2, \dots, \ell_{\mathbf{T}}} \left\{ \sum_{t=1}^{\mathbf{T}} \langle p_t, \ell_t \rangle - \min_{j \in [\mathbf{K}]} \sum_{t=1}^{\mathbf{T}} \ell_{j,t} \right\} \quad . \tag{7.2}$$

**EXPONENTIAL WEIGHTS (HEDGE).** There is a no-regret algorithm for DTOL. It has many names such as "Exponential Weights" and "Hedge". The idea of Hedge is to maintain weights over actions, and the weight of an action decays exponentially in the total loss incurred by that action over all previous rounds.

Input: learning rate  $\eta \in (0, 1]$ . Initialize:  $w_{j,0} = 1$  for all  $j \in [K]$ . For t = 1, 2, ..., T, 1. Set  $p_{j,t} = \frac{w_{j,t-1}}{\sum_{j'} w_{j',t-1}}$  for all  $j \in [K]$ ; 2. Observe loss vector  $\ell_t$  for Nature ; 3. Suffer loss  $\langle p_t, \ell_t \rangle = \sum_j p_{j,t} \ell_{j,t}$ ; 4. Update  $w_{j,t} = w_{j,t-1} \cdot e^{-\eta \ell_{j,t}}$  for all  $j \in [K]$ .

**THEOREM 7.8.** For any sequence of loss vectors  $(\ell_1, \ell_2, ..., \ell_T) \in ([0, 1]^K)^T$ , for any  $\eta \in (0, 1]$ , we have

$$\sum_{t=1}^{\mathrm{T}} \langle p_t, \ell_t \rangle - \min_{j \in [\mathrm{K}]} \sum_{t=1}^{\mathrm{T}} \ell_{j,t} \leq \frac{\eta}{2} \sum_{t=1}^{\mathrm{T}} \left\langle p_t, \ell_t^2 \right\rangle + \frac{\log \mathrm{K}}{\eta}$$

With  $\eta = \sqrt{\frac{2 \log K}{T}}$ , we have  $\mathcal{R}(T) \le \sqrt{2T \log K}$ .

*Proof.* of Theorem 7.8: Fix a sequence of loss vectors  $(\ell_1, \ell_2, \dots, \ell_T)$ . Let  $Z_t = \sum_{j=1}^{K} w_{j,t} = \sum_{j=1}^{K} w_{j,t-1} \cdot e^{-\eta \cdot \ell_{j,t}} = \sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t} \ell_{j,s}}$  be the total weight at the end of round t. Note that  $Z_0 = K$ . We have  $\log Z_T = \sum_{t=1}^{T} \log \frac{Z_t}{Z_{t-1}} + \log Z_0$ .

Sometimes, Learner chooses  $p_t \in \{e_1, e_2, \dots, e_K\}.$ 

Now, we construct the following upper bound to upper bound  $\log Z_T$ . We have

$$\log \frac{Z_t}{Z_{t-1}} = \log \frac{\sum_{j=1}^{K} w_{j,t-1} \cdot e^{-\eta \cdot \ell_{j,t}}}{\sum_{j'=1}^{K} w_{j',t-1}}$$

$$= \log \left( \sum_{j=1}^{K} \frac{w_{j,t-1}}{\sum_{j'=1}^{K} w_{j',t-1}} \cdot e^{-\eta \cdot \ell_{j,t}} \right)$$

$$= \log \left( \sum_{j=1}^{K} p_{j,t} \cdot e^{-\eta \cdot \ell_{j,t}} \right)$$

$$\leq^{(a)} \log \left( \sum_{j=1}^{K} p_{j,t} \cdot \left( 1 - \eta \cdot \ell_{j,t} + \frac{\eta^2 \cdot \ell_{j,t}^2}{2} \right) \right)$$

$$= \log \left( 1 - \eta \langle p_t, \ell_t \rangle + \frac{\eta^2}{2} \langle p_t, \ell_t^2 \rangle \right)$$

$$\leq^{(b)} -\eta \langle p_t, \ell_t \rangle + \frac{\eta^2}{2} \langle p_t, \ell_t^2 \rangle,$$

where step (a) uses  $e^{-x} \le 1 - x + x^2/2$  and step (b) uses  $\log(1 + x) \le x$ . We also have a lower bound on  $\log Z_T$ , which is

$$\log Z_{\mathrm{T}} = \log \left( \sum_{j=1}^{\mathrm{K}} e^{-\eta \sum_{t=1}^{\mathrm{T}} \ell_{j,t}} \right)$$
  

$$\geq \log \left( \max_{j \in [\mathrm{K}]} \left\{ e^{-\eta \sum_{t=1}^{\mathrm{T}} \ell_{j,t}} \right\} \right)$$
  

$$= \max_{j \in [\mathrm{K}]} \left\{ -\eta \sum_{t=1}^{\mathrm{T}} \ell_{j,t} \right\}.$$
(7.3)

Now, we have

$$\sum_{t=1}^{T} -\eta \langle p_t, \ell_t \rangle + \frac{\eta^2}{2} \langle p_t, \ell_t^2 \rangle \ge \sum_{t=1}^{T} \log \frac{Z_t}{Z_{t-1}} = \log Z_T - \log Z_0 \ge \max_{j \in [K]} \left\{ -\eta \sum_{t=1}^{T} \ell_{j,t} \right\} - \log K$$

$$\Rightarrow \sum_{t=1}^{T} \langle p_t, \ell_t \rangle - \min_{j \in [K]} \left\{ \sum_{t=1}^{T} \ell_{j,t} \right\} \le \frac{\eta}{2} \langle p_t, \ell_t^2 \rangle + \frac{\log K}{\eta} \quad .$$

$$(7.4)$$

**Remark.** Note that in order to achieve the  $\sqrt{2T \log K}$  regret bound, **Hedge** needs to input the learning rate  $\eta = \sqrt{\frac{2 \log K}{T}}$ , depending on the learning horizon T. Later, we will show how to use doubling-trick to get rid of it!

**Hedge with Doubling-Trick.** The regret bound  $\mathcal{R}(T) = \mathcal{O}(\sqrt{T \log K})$  in Theorem 7.8 relies on inputting the learning rate  $\eta = \sqrt{\frac{2 \log K}{T}}$  which relies on the knowledge of the time horizon T.

One may be curious to know is it possible to have a learning algorithm that does not need to know T in advance, but still preserving the same regret bound?

The answer is yes!!!

We introduce a useful idea, called *doubling-trick*, for solving online learning prob-Geometric doubling vs lems. Using (Geometric) doubling-trick, we can still achieve an  $\mathcal{O}(\sqrt{T \log K})$ ) regret

exponential doubling can be found here.

8

bound. The idea is to run the algorithm in epochs of lengths  $2^0, 2^1, \ldots, 2^r, \ldots$  until stopping.

At the beginning of each epoch  $r \ge 0$ , we set the learning rate  $\eta_r = \sqrt{\frac{2 \log K}{2^r}}$  based on  $2^r$  the length of the current epoch. At the end of epoch r, we reset the algorithm, that is, we forget all the stuff we have learned. Progressing in this way, we will run the algorithm for epochs r = 0, 1, ..., d, where  $d = \lceil \log_2(T+1) - 1 \rceil = \mathcal{O}(\log T)$ .

**THEOREM 7.9.** The regret of Hedge with doubling-trick is  $\mathcal{O}(\sqrt{T \log K})$ .

*Proof.* For any  $r \ge 0$ , from (7.8), we have

$$\sum_{t=2^{r+1}}^{2^{r+1}} \langle p_t, \ell_t \rangle - \min_{j_r \in [K]} \sum_{t=2^{r+1}}^{2^{r+1}} \ell_{j_r, t} \le \frac{\eta_r}{2} \sum_{t=2^{r+1}}^{2^{r+1}} \langle p_t, \ell_t^2 \rangle + \frac{\log K}{\eta_r} \le \mathcal{O}\left(\sqrt{2^r \log K}\right) .$$
(7.5)

The regret is

$$\begin{aligned} \mathcal{R}(\mathbf{T}) &= \sup_{\ell_{1},\ell_{2},...,\ell_{\mathrm{T}}} \left\{ \sum_{t=1}^{\mathrm{T}} \langle p_{t},\ell_{t} \rangle - \min_{j \in [\mathrm{K}]} \sum_{t=1}^{\mathrm{T}} \ell_{j,t} \right\} \\ &= \sup_{\ell_{1},\ell_{2},...,\ell_{\mathrm{T}}} \left\{ \sum_{r \geq 0} \sum_{t=2^{r+1}}^{2^{r+1}} \langle p_{t},\ell_{t} \rangle - \min_{j \in [\mathrm{K}]} \sum_{r \geq 0} \sum_{t=2^{r+1}}^{2^{r+1}} \ell_{j,t} \right\} \\ &\leq \sup_{\ell_{1},\ell_{2},...,\ell_{\mathrm{T}}} \left\{ \sum_{r \geq 0} \sum_{t=2^{r+1}}^{2^{r+1}} \langle p_{t},\ell_{t} \rangle - \sum_{r \geq 0} \min_{j_{r} \in [\mathrm{K}]} \sum_{t=2^{r+1}}^{2^{r+1}} \ell_{j_{r},t} \right\} \\ &= \sup_{\ell_{1},\ell_{2},...,\ell_{\mathrm{T}}} \left\{ \sum_{r \geq 0} \left( \sum_{t=2^{r+1}}^{2^{r+1}} \langle p_{t},\ell_{t} \rangle - \min_{j_{r} \in [\mathrm{K}]} \sum_{t=2^{r+1}}^{2^{r+1}} \ell_{j_{r},t} \right) \right\} \\ &= \sup_{\ell_{1},\ell_{2},...,\ell_{\mathrm{T}}} \left\{ \sum_{r \geq 0} \mathcal{O}(\sqrt{2^{r}\log \mathrm{K}}) \right\} \\ &= \sum_{r \geq 0} \mathcal{O}(\sqrt{2^{r}\log \mathrm{K}}) \\ &= \mathcal{O}\left( \sqrt{\mathrm{T}\log \mathrm{K}} \right) \quad . \end{aligned}$$

The last epoch may not have a full length, but we can add  $\vec{0}$ 's to make the last epoch have a full length

(7.6)

#### 7.3 BANDITS

Last time we have talked about Hedge/DTOL learning. We usually say it is a full information game as Learner is able to observe each individual loss in  $\ell_t = (\ell_{1,t}, \ell_{2,t}, \dots, \ell_{K,t})$ . Starting from this lecture, we will talk about multi-armed bandit (MAB) problems, where only some entry in the loss vector  $\ell_t = (\ell_{1,t}, \ell_{2,t}, \dots, \ell_{K,t})$  is revealed in each round.

Actions and arms are **LEARNING PROTOCOL.** We have a fixed arm set [K].

interchangeable. You can view an arm as a hypothesis.

You can view Step 1 and Step 2 occur simultaneously.

In Step 3, the remaining losses other than  $\ell_{J_t,t}$  are still hidden to Learner, which is the key difference between DTOL and bandits. In each round  $t = 1, 2, \ldots, T$ ,

- 1. Adversary/Environment selects a loss vector  $\ell_t = (\ell_{1,t}, \ell_{2,t}, \dots, \ell_{K,t})$  that is hidden to Learner ;
- 2. Learner pulls an arm  $J_t \in [K]$ ;
- 3. Learner suffers/observes loss  $\ell_{J_t,t}$  associated with the pulled arm  $J_t$ .

The goal of Learner to pull a sequence of arms  $(J_1, J_2, ..., J_T)$  to minimize the total loss by the end of round T.

Regret is defined as

$$\mathcal{R}(\mathbf{T}) := \sum_{t=1}^{\mathbf{T}} \ell_{\mathbf{J}_{t},t} - \min_{j \in [\mathbf{K}]} \sum_{t=1}^{\mathbf{T}} \ell_{j,t} \quad ,$$
(7.7)

which is a random variable as  $J_1, J_2, \ldots, J_T$  and all loss vectors are random.

Based on how the loss vectors  $(\ell_1, \ell_2, \dots, \ell_T)$  are generated, we have

- 1. Adversary bandits: no distributional assumption is made.
- 2. Stochastic bandits: all  $\ell_t$  are i.i.d. over time according to a fixed but unknown probability distribution.

Learner needs to make a good balance between

- 1. *Exploitation*: Learner needs to pull arms that have smaller losses, as the goal is to minimize the total loss.
- 2. *Exploration*: Learner needs to pull arms that have not been observed too often to gain information.

## 7.3.1 Adversarial bandits.

In adversarial bandits, the loss vectors can be generated adversarially. We use *pseudo regret* to measure performance of the algorithm used by Learner, defined as

$$\overline{\mathcal{R}}(\mathbf{T}) := \mathbb{E}\left[\sum_{t=1}^{\mathbf{T}} \ell_{\mathbf{J}_{t},t}\right] - \min_{j \in [\mathbf{K}]} \mathbb{E}\left[\sum_{t=1}^{\mathbf{T}} \ell_{j,t}\right] \quad ,$$
(7.8)

Sometimes people call it where the expectation is taken over  $(J_1, J_2, ..., J_T)$  and all loss vectors over T rounds. expected regret as it has an expectation. I am following

the notation of Bubeck and EXP3 (EXPONENTIAL WEIGHTS FOR EXPLORATION AND EXPLOITATION). It is quite Cesa-Bianchi [BC12]. Similar to Hedge, but in EXP3, only the weight associated with the pulled arm will be updated at the end of each round t, as only the loss for that arm is revealed.

Importance sampling We construct an estimator  $\tilde{\ell}_t = (\tilde{\ell}_{1,t}, \tilde{\ell}_{2,t}, \dots, \tilde{\ell}_{K,t})$  with each entry  $\tilde{\ell}_{j,t} = \ell_{j,t} \frac{1\{J_t=j\}}{p_{i,t}}$ .

It is not hard to see that for all arms not played in round t, we set  $\tilde{\ell}_{i,t} = 0$ . For the pulled arm in round t, we set  $\tilde{\ell}_{J_{t},t} = \frac{\ell_{J_{t},t}}{p_{J_{t},t}}$  instead of  $\ell_{J_{t},t}$  by making it more "important"! Actually, the constructed estimator  $\tilde{\ell}_{j,t}$  is an unbiased estimator of  $\ell_{j,t}$ , as we have

$$\mathbb{E}_{\mathbf{J}_t \sim p_t} \left[ \tilde{\ell}_{j,t} \right] = \mathbb{E}_{\mathbf{J}_t \sim p_t} \left[ \ell_{j,t} \frac{\mathbf{1}\{\mathbf{J}_t = j\}}{p_{j,t}} \right] = \ell_{j,t} \quad .$$
(7.9)

Input: learning rate  $\eta \in (0, 1]$ . Initialize:  $w_{j,0} = 1$  for all  $j \in [K]$ . For t = 1, 2, ..., T,

- 1. Adversary/Environment selects a loss vector  $\ell_t = (\ell_{1,t}, \ell_{2,t}, \dots, \ell_{K,t})$  that is hidden to Learner ;
- 2. Learner computes  $p_{j,t} = \frac{w_{j,t-1}}{\sum\limits_{i'} w_{j',t-1}}$  for all  $j \in [K]$ ;
- 3. Learner plays arm  $J_t \in [K]$  according to  $p_t = (p_{1,t}, p_{2,t}, \dots, p_{K,t})$ ; 4. Learner computes loss estimates  $\tilde{\ell}_{j,t} = \frac{\ell_{j,t}}{p_{j,t}} \mathbf{1} \{J_t = j\}$  for all  $j \in [K]$ ;
- 5. Update  $w_{j,t} = w_{j,t-1} \cdot e^{-\eta \tilde{\ell}_{j,t}}$  for all  $j \in [K]$ .

**THEOREM 7.10.** Assume that all loss vectors are bounded with [0, 1] support. If **EXP3** is run with learning rate  $\eta = \sqrt{\frac{\log K}{KT}}$ , the pseudo regret is at most  $\sqrt{2TK\log K}$ .

Proof.

$$\begin{aligned} \overline{\mathcal{R}}(\mathbf{T}) &= \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \ell_{\mathbf{J}_{t},t}\right] - \min_{j \in [\mathrm{K}]} \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \ell_{j,t}\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \left\langle p_{t}, \tilde{\ell}_{t} \right\rangle\right] - \min_{j \in [\mathrm{K}]} \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{E}_{\mathbf{J}_{t} \sim p_{t}}\left[\tilde{\ell}_{j,t}\right]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \left\langle p_{t}, \tilde{\ell}_{t} \right\rangle\right] - \min_{j \in [\mathrm{K}]} \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \tilde{\ell}_{j,t}\right] \\ &\leq \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \left\langle p_{t}, \tilde{\ell}_{t} \right\rangle\right] - \mathbb{E}\left[\min_{j \in [\mathrm{K}]} \sum_{t=1}^{\mathrm{T}} \tilde{\ell}_{j,t}\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \left\langle p_{t}, \tilde{\ell}_{t} \right\rangle - \min_{j \in [\mathrm{K}]} \sum_{t=1}^{\mathrm{T}} \tilde{\ell}_{j,t}\right] \\ &\leq (a) \mathbb{E}\left[\frac{\eta}{2} \sum_{t=1}^{\mathrm{T}} \left\langle p_{t}, \tilde{\ell}_{t}^{2} \right\rangle + \frac{\log \mathrm{K}}{\eta}\right] \\ &= (b) \mathbb{E}\left[\frac{\eta}{2} \sum_{t=1}^{\mathrm{T}} \sum_{j \in [\mathrm{K}]} p_{j,t} \frac{\ell_{j,t}^{2}}{p_{j,t}}\right] + \frac{\log \mathrm{K}}{\eta} \end{aligned}$$
(7.10)

Tuning  $\eta = \sqrt{\frac{\log K}{KT}}$  gives the stated regret bound.

Step (a) uses Theorem 7.8 in previous lecture. Note that from (7.9), we know  $\mathbb{E}_{J_t \sim p_t} | \tilde{\ell}_{j,t} | \in [0, 1]$ . So, it is safe to use Theorem 7.8 directly.

Step (b) uses 
$$\mathbb{E}_{J_t \sim p_t} \left[ \tilde{\ell}_{j,t}^2 \right] = \mathbb{E}_{J_t \sim p_t} \left[ \ell_{j,t}^2 \frac{1\{J_t=j\}}{p_{j,t}^2} \right] = \frac{\ell_{j,t}^2}{p_{j,t}^2}$$
.

**REMARK.** Note that inputting  $\eta = \sqrt{\frac{\log K}{KT}}$  into EXP3 means that it is not an anytime learning algorithm. To make it anytime, you can set the learning rate  $\eta_t = \sqrt{\frac{\log K}{tK}}$  in each round *t*. The regret analysis is much more complicated (refer to Theorem 3.1 in [BC12]).

## 7.3.2 Stochastic bandits

In stochastic bandits, we have a fixed arm set [K] and each arm  $j \in [K]$  is associated with a reward distribution  $v_j$ . We can use  $\Theta_K := (v_1, v_2, ..., v_K)$  to specify a K-armed stochastic bandit problem instance.

In each round  $t = 1, 2, \ldots, T$ ,

- 1. Environment generate a reward vector  $X_t = (X_{1,t}, X_{2,t}, ..., X_{K,t})$  with each  $X_{j,t} \sim v_j$ . This reward vector is hidden to Learner ;
- 2. Learner pulls an arm  $J_t \in [K]$ ;
- 3. Learner obtains/observes  $X_{J_t,t}$ , the reward of the pulled arm  $J_t$ .

We still use pseudo regret to measure performance of Alg, defined as

$$\overline{\mathcal{R}}(\mathbf{Alg}; \Theta_{\mathrm{K}}; \mathrm{T}) := \max_{j \in [\mathrm{K}]} \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathrm{X}_{j,t}\right] - \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathrm{X}_{\mathrm{J}_{t},t}\right] , \qquad (7.11)$$

where the randomness is taken over  $J_1, J_2, ..., J_T$  and all reward vectors. Note that since we have statistical assumptions on reward vectors, different  $\Theta_K$  may give different regret. That is also to say, for a fixed algorithm, when working over different problem instances, the regret could be different.

The goal of Learner is to pull arms sequentially to minimize regret.

Let  $\mu_j = \mathbb{E}_{X_j \sim v_j}[X_j]$  denote the mean reward of arm *j*. Without loss of generality, we assume the first arm is the optimal one, that is,  $\mu_1 > \mu_j$  for all  $j \neq 1$ .

For any  $j \neq 1$ , let  $\Delta_j := \mu_1 - \mu_j$  denote the mean reward gap. We also call it the sub-optimality gap between the optimal arm 1 and the sub-optimal arm *j*. Let  $\Delta_1 = 0$ .

Now, we can rewrite  $\overline{\mathcal{R}}(T)$  as

$$\overline{\mathcal{R}}(\mathbf{T}) = \max_{j \in [\mathbf{K}]} \mathbb{E}\left[\sum_{t=1}^{T} \mathbf{X}_{j,t}\right] - \mathbb{E}\left[\sum_{t=1}^{T} \mathbf{X}_{J_{t},t}\right]$$

$$= \mathbf{T} \cdot \boldsymbol{\mu}_{1} - \sum_{t=1}^{T} \mathbb{E}\left[\boldsymbol{\mu}_{J_{t}}\right]$$

$$= \sum_{t=1}^{T} \mathbb{E}\left[\boldsymbol{\mu}_{1} - \boldsymbol{\mu}_{J_{t}}\right]$$

$$= \sum_{t=1}^{T} \mathbb{E}\left[\Delta_{J_{t}}\right]$$

$$= \sum_{t=1}^{T} \sum_{j \in [\mathbf{K}]} \mathbb{E}\left[\mathbf{1}\left\{J_{t} = j\right\}\right] \cdot \Delta_{j}$$

$$= \sum_{j \in [\mathbf{K}]: \Delta_{j} > 0} \mathbb{E}\left[\sum_{t=1}^{T} \mathbf{1}\left\{J_{t} = j\right\}\right] \cdot \Delta_{j} \quad .$$
(7.12)

Let  $n_{j,t-1} := \sum_{s=1}^{t-1} \mathbf{1} \{J_s = j\}$  denote the total number of pulls for arm j by the end of round t - 1. Then,  $\mathbb{E}[n_{j,T}]$  is the expected number of pulls of arm j by the end of learning and  $\Delta_j$  is the singe round performance loss when pulling a sub-optimal arm j.

From the last step in (7.12), it is not hard to see, to minimize the regret, it is important to control the number of pulls of sub-optimal arms. But we have no idea which arms are sub-optimal. So, we have to pull each arm a certain amount of times in order to learn whether they are sub-optimal or not confidently. We need information!

Here, I should mention exploitation-vs-exploration.

**UPPER CONFIDENCE BOUND (UCB).** It is inspired by the principle of being optimistic in the face of uncertainty. Usually, all UCB-based algorithms are optimistic learning algorithms and follow a template to decompose regret. Also, they can be justified by concentration inequalities, e.g., Hoeffdin's inequality.

Recall  $\Theta_K := (v_1, v_2, \dots, v_K)$  is the bandit instance we are interested in. Actually, for developing regret minimization algorithm, we are interested in  $(\mu_1, \mu_2, \dots, \mu_K)$ . Note that only the mean reward gaps appear in the regret.

Recall  $n_{j,t-1} = \sum_{s=1}^{t-1} \mathbf{1} \{J_s = j\}$  is the number of pulls of arm *j* by the end of round t - 1.

Now, let  $\hat{\mu}_{j,n_{j,t-1}} := \frac{1}{n_{j,t-1}} \sum_{s=1}^{t-1} X_{j,s} \mathbf{1} \{J_s = j\}$  be the empirical mean of arm j by the end of round t - 1, i.e., the average of  $n_{j,t-1}$  iid random variables according to  $v_j$ .

Now, we can construct an empirical model  $\hat{\Theta}_t = (\hat{\mu}_{1,n_{1,t-1}}, \hat{\mu}_{2,n_{2,t-1}}, \dots, \hat{\mu}_{K,n_{K,t-1}}).$ 

Assume all reward distributions have a [0, 1] support. The idea of UCB1<sup>2</sup>, an algorithm in UCB family, is to construct an optimistic model  $\overline{\Theta}_t = (\bar{\mu}_{1,t}, \bar{\mu}_{2,t}, \dots, \bar{\mu}_{K,t})$  with each  $j \in [K]$ 

exploitation-vs-exploration  
Hoeffding's inequality  
upper confidence bound  
w.t.p. 
$$\bar{\mu}_{1,t} \ge \mu_1$$

$$\bar{\mu}_{j,t} = \hat{\mu}_{j,n_{j,t-1}} + \sqrt{\frac{2\ln(t)}{n_{j,t-1}}} \quad .$$
(7.13)

In each round  $t = 1, 2, \ldots, T$ ,

- 1. Environment generate a reward vector  $X_t = (X_{1,t}, X_{2,t}, ..., X_{K,t})$  with each  $X_{j,t} \sim v_j$ . This reward vector is hidden to Learner ;
- 2. Learner constructs the upper confidence bound  $\bar{\mu}_{j,t} = \hat{\mu}_{j,n_{j,t-1}} + \sqrt{\frac{2 \ln(t)}{n_{j,t-1}}}$  for all  $j \in [K]$ ;
- 3. Learner pulls the arm with the highest upper confidence bound, i.e.,  $J_t \in \underset{j \in [K]}{\operatorname{arg \,max}} \bar{\mu}_{j,t}$ ;
- 4. Learner observes  $X_{I_{t,t}}$ , the reward of the pulled arm  $J_t$ ;
- 5. Learner updates  $n_{J_t,t} = n_{J_t,t-1} + 1$  and the empirical mean  $\hat{\mu}_{J_t,n_{1,t}}$ .

<sup>&</sup>lt;sup>2</sup>UCB1 works for all Sub-Gaussian reward distributions.

**THEOREM 7.11.** If all reward distributions in  $\Theta_K$  have a [0, 1] support, we have

$$\overline{\mathcal{R}}(UCB1;\Theta_{\mathrm{K}};\mathrm{T}) \leq \sum_{j \in [\mathrm{K}]:\Delta_j > 0} \frac{8 \ln \mathrm{T}}{\Delta_j} + Constant$$

*Proof.* Fix a sub-optimal arm *j*, we upper bound  $\mathbb{E}[n_{j,T}]$ .

You can view  $L_j$  as the Let  $L_j := \frac{\Box \cdot \ln T}{\Delta_i^2}$ , where  $\Box$  is a constant that will be tuned later. We have

amount of observations needed to conclude that this arm is not the optimal one.

$$\mathbb{E}\left[n_{j,\mathrm{T}}\right] = \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j\right\}\right] \\
= \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, n_{j,t-1} \leq \mathrm{L}_{j}\right\}\right] + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, n_{j,t-1} > \mathrm{L}_{j}\right\}\right] \\
= \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, n_{j,t} > n_{j,t-1}, n_{j,t-1} \leq \mathrm{L}_{j}\right\}\right] + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, n_{j,t} > n_{j,t-1}, n_{j,t-1} > \mathrm{L}_{j}\right\}\right] \\
\leq \mathbb{L}_{j} + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, \bar{\mu}_{j,t} \geq \bar{\mu}_{1,t}, n_{j,t-1} > \mathrm{L}_{j}\right\}\right] \\
\leq \mathbb{L}_{j} + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, \bar{\mu}_{j,t} \geq \bar{\mu}_{1,t}, n_{j,t-1} > \mathrm{L}_{j}\right\}\right] \\
\leq \mathbb{L}_{j} + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{J_{t}=j, \bar{\mu}_{j,t} \geq \mu_{1}, n_{j,t-1} > \mathrm{L}_{j,t}\right\}\right] + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{\bar{\mu}_{1,t} \leq \mu_{1}, n_{j,t-1} > \mathrm{L}_{j}\right\}\right] \\
\leq \mathbb{L}_{j} + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{\bar{\mu}_{j,t} \geq \mu_{j} + \Delta_{j}, n_{j,t-1} > \mathrm{L}_{j}\right\}\right] + \mathbb{E}\left[\sum_{t=1}^{\mathrm{T}} \mathbf{1}\left\{\bar{\mu}_{1,t} \leq \mu_{1}\right\}\right] . \tag{7.14}$$

Now, we set  $L_j = \frac{8 \ln(T)}{\Delta_j^2}$ . Then, we have  $\Delta_j = \sqrt{\frac{8 \ln T}{L_j}}$ .

$$\mathbb{E}\left[\sum_{t=1}^{T} \mathbf{1}\left\{\bar{\mu}_{j,t} \ge \mu_{j} + \Delta_{j}, n_{j,t-1} > L_{j}\right\}\right] \\ = \mathbb{E}\left[\sum_{t=1}^{T} \mathbf{1}\left\{\hat{\mu}_{j,n_{j,t-1}} + \sqrt{\frac{2\ln(t)}{n_{j,t-1}}} \ge \mu_{j} + \sqrt{\frac{8\ln T}{L_{j}}}, n_{j,t-1} > L_{j}\right\}\right] \\ \le \mathbb{E}\left[\sum_{t=1}^{T} \mathbf{1}\left\{\hat{\mu}_{j,n_{j,t-1}} + \sqrt{\frac{2\ln(T)}{n_{j,t-1}}} \ge \mu_{j} + \sqrt{\frac{8\ln T}{n_{j,t-1}}}, n_{j,t-1} > L_{j}\right\}\right] \\ \le \mathbb{E}\left[\sum_{t=1}^{T} \mathbf{1}\left\{\hat{\mu}_{j,n_{j,t-1}} \ge \mu_{j} + \sqrt{\frac{2\ln T}{n_{j,t-1}}}, n_{j,t-1} > L_{j}\right\}\right] \\ \le \sum_{t=1}^{T} \sum_{h=L_{j}}^{t-1} \mathbb{E}\left[\mathbf{1}\left\{\hat{\mu}_{j,h} \ge \mu_{j} + \sqrt{\frac{2\ln T}{h}}\right\}\right] \\ \le T^{2} \cdot e^{-2\cdot 2\ln T} \\ = \mathcal{O}(1) \quad .$$

Similarly, we have 
$$\mathbb{E}\left[\sum_{t=1}^{T} \mathbf{1}\left\{\bar{\mu}_{1,t} \leq \mu_{1}\right\}\right] = \mathcal{O}(1).$$

If problem-dependent Now, we have  $\mathbb{E}\left[n_{j,T}\right] \leq L_j + \mathcal{O}(1) = \frac{8 \ln T}{\Delta_j^2} + \mathcal{O}(1)$ , which gives

 $\overline{\mathcal{R}}(\mathbf{UCB1};\Theta_{\mathbf{K}};\mathbf{T}) = \sum_{\substack{j \in [\mathbf{K}]:\Delta_j > 0}} \mathbb{E}\left[n_{j,\mathbf{T}}\right] \cdot \Delta_j$  $\leq \sum_{\substack{j \in [\mathbf{K}]:\Delta_j > 0}} \frac{8 \ln \mathbf{T}}{\Delta_j} + \mathbf{Constant} \quad .$ (7.16)

parameters, e.g., all  $\Delta_j$ , appear in the regret bound, we say it is a problem-dependent regret bound. **WORST-CASE REGRET BOUND FOR UCB1.** Let us consider a 2-armed bandit problem, where the first arm is the optimal one and the second arm has a mean reward gap  $\Delta = \frac{1}{T}$ . Clearly, according to Theorem 7.11, we have  $\overline{\mathcal{R}}(\mathbf{UCB1}; \Theta_{\mathrm{K}}; \mathrm{T}) = 8\mathrm{T} \ln \mathrm{T}$ , which is even worse than T. Does it mean UCB1 fails this learning task?

To answer this question, we are motivated to study the worst-case regret bound, defined as

$$\sup_{\Theta_{K}\in\Pi^{K}} \overline{\mathcal{R}}(\mathbf{UCB1};\Theta_{K};T) , \qquad (7.17)$$

where  $\Pi$  is a set of distributions with a [0, 1] support.

**THEOREM 7.12.** We have

$$\sup_{\Theta_{K}\in\Pi^{K}} \overline{\mathcal{R}}(UCB1;\Theta_{K};T) \le \mathcal{O}(\sqrt{KT\ln T}) \quad .$$
(7.18)

*Proof.* Fix  $\Theta_{K}$  and set  $\Delta := \sqrt{\frac{K \ln T}{T}}$ . We have

$$\overline{\mathcal{R}}(\mathbf{UCB1};\Theta_{\mathrm{K}};\mathrm{T}) = \sum_{t=1}^{\mathrm{T}} \sum_{j \in [\mathrm{K}]} \mathbb{E} \left[ \mathbf{1} \left\{ J_{t} = j \right\} \right] \cdot \Delta_{j} \\
= \sum_{t=1}^{\mathrm{T}} \sum_{j \in [\mathrm{K}]:\Delta_{j} \leq \Delta} \mathbb{E} \left[ \mathbf{1} \left\{ J_{t} = j \right\} \right] \cdot \Delta_{j} + \sum_{t=1}^{\mathrm{T}} \sum_{j \in [\mathrm{K}]:\Delta_{j} > \Delta} \mathbb{E} \left[ \mathbf{1} \left\{ J_{t} = j \right\} \right] \cdot \Delta_{j} \\
\leq \mathrm{T} \cdot \Delta + \sum_{t=1}^{\mathrm{T}} \sum_{j \in [\mathrm{K}]:\Delta_{j} > \Delta} \mathbb{E} \left[ \mathbf{1} \left\{ J_{t} = j \right\} \right] \cdot \Delta_{j} \\
\leq \mathrm{T} \cdot \Delta + \sum_{j \in [\mathrm{K}]:\Delta_{j} > \Delta} \left( \frac{8 \ln \mathrm{T}}{\Delta_{j}} + \mathcal{O}(1) \right) \\
\leq \mathrm{T} \cdot \Delta + \sum_{j \in [\mathrm{K}]:\Delta_{j} > \Delta} \left( \frac{8 \ln \mathrm{T}}{\Delta} + \mathcal{O}(1) \right) \\
\leq \mathrm{T} \cdot \Delta + \mathrm{K} \frac{8 \ln \mathrm{T}}{\Delta} + \mathcal{O}(\mathrm{K}) \\
= \mathcal{O}(\sqrt{\mathrm{KT} \ln \mathrm{T}}) \quad .$$
(7.19)

One may ask whether UCB1 is optimal in the worst case sense or not. The definition of minimax optimality, a joint of property between a family of algorithms and distributions, can answer this question. We skip the proof here and only show the conclusion: UCB1 is minimax optimal up to an extra  $\sqrt{\ln T}$  factor.

**ARM ELIMINATION ALGORITHM.** As UCB1 only has an  $\mathcal{O}(\sqrt{\text{KT} \ln T})$  worst-case regret bound, now, we show an algorithm that enjoys an  $\mathcal{O}(\sqrt{\text{KT} \ln K})$  worst-case regret bound, which is slightly better than UCB1.

Suppose we have a special K-armed bandit problem with one arm having a mean reward  $\mu_*$  and all the remaining arms having the same mean reward  $\mu_* - \Delta$ , i.e., the mean reward gap for any sub-optimal arm is  $\Delta$ . *Learner knows*  $\Delta$  *and*  $\mu_*$  *but does not know which arm is the optimal one.* To solve this special bandit problem, we can use the following **Arm Elimination** algorithm [AO10]:

Input: [K],  $\mu_*$ ,  $\Delta$ , and T.

- 1. Pull each arm  $n = \frac{2 \ln(T\Delta^2)}{\Delta^2}$  times and compute the empirical mean  $\hat{\mu}_{j,n}$  of each arm  $j \in [K]$ ;
- 2. Commit to the arm with the highest empirical mean until the end of the learning, i.e., pull arm  $J = \arg \max_{j \in [K]} \hat{\mu}_{j,n}$  for the remaining T Kn rounds.

**THEOREM 7.13.** Arm Elimination enjoys an  $\mathcal{O}\left(\frac{K\ln(T\Delta^2)}{\Delta} + \frac{K}{\Delta}\right)$  problem-dependent regret bound. It also enjoys an  $\mathcal{O}(\sqrt{KT\ln K})$  worst-case regret bound.

*Proof.* Let  $i_*$  denote the index of the optimal arm. Without loss of generality, we assume it is unique. We first upper bound the probability that the committed arm is not the optimal one. We have

$$\mathbb{P}\left\{\mathbf{J}\neq i_{*}\right\} \leq \mathbb{P}\left\{\max_{j\in[K]\setminus\{i_{*}\}}\hat{\mu}_{j,n}\geq \hat{\mu}_{i_{*},n}\right\} \leq \sum_{j\in[K]\setminus\{i_{*}\}}\mathbb{P}\left\{\hat{\mu}_{j,n}\geq \hat{\mu}_{i_{*},n}\right\} \leq (K-1)\cdot 2e^{-n\frac{\Delta^{2}}{2}} \quad .$$

$$(7.20)$$

The problem-dependent regret  $\overline{\mathcal{R}}(T)$  is

$$\underbrace{(K-1) \cdot n \cdot \Delta}_{\text{regret in Step 1}} + \underbrace{\mathbb{P} \{J \neq i_*\} \cdot (T-Kn) \cdot \Delta}_{\text{regret in Step 1}} \\ \leq K \cdot \frac{2\ln(T\Delta^2)}{\Delta^2} \cdot \Delta + K \cdot 2e^{-n\frac{\Delta^2}{2}} \cdot T \cdot \Delta \\ = 2K \frac{\ln(T\Delta^2)}{\Delta} + K \cdot \frac{2}{T\Delta^2} \cdot T \cdot \Delta \\ = \mathcal{O}\left(\frac{K\ln(T\Delta^2)}{\Delta} + \frac{K}{\Delta}\right) .$$
(7.21)

Let  $\tilde{\Delta} := \frac{e\sqrt{K}}{\sqrt{T}}$ . If  $\Delta \leq \tilde{\Delta}$ , we have the regret is at most  $T \cdot \tilde{\Delta} = e\sqrt{KT}$ . If  $\Delta > \tilde{\Delta}$ , we have  $\mathcal{O}\left(\frac{K\ln(T\Delta^2)}{\Delta} + \frac{K}{\Delta}\right) \leq \mathcal{O}\left(\frac{K\ln(T\tilde{\Delta}^2)}{\tilde{\Delta}} + \frac{K}{\tilde{\Delta}}\right) = \mathcal{O}(\sqrt{KT\ln K})$ , where the inequality uses the fact that  $f(x) = \frac{\ln(Tx^2)}{x}$  is a decreasing function when  $Tx^2 \geq e^2$ .

**ARM ELIMINATION ALGORITHM WITH DOUBLING-TRICK.** Since we cannot assume we know  $\Delta$  in advance and all sub-optimal arms have the same  $\Delta$ , we cannot use **Arm Elimination** directly for solving practical learning problems. A good thing is we can introduce **doubling-trick** into **Arm Elimination** to make it work by estimating  $\Delta_j$ .

Input: [K] and T. Initialization: Set  $\hat{\Delta}_1 = 0.5$  and  $B_1 = [K]$ . For epochs r = 1, 2, ... up to log T, 1. For each arm  $j \in B_r$ , pull it until the total number of pulls hits  $n_r = \frac{2\ln(KT\hat{\Delta}_r^2)}{\hat{\Delta}_r^2}$ ; 2. All arms  $i \in B_r$  such that  $\hat{\mu}_{i,n_r} + \sqrt{\frac{\log(KT\hat{\Delta}_r^2)}{2 \cdot n_r}} \geq \max_{\substack{j \in B_r} \hat{\mu}_{j,n_r}} - \sqrt{\frac{\log(KT\hat{\Delta}_r^2)}{2 \cdot n_r}}$ . upper confidence bound of arm i lower confidence bound of  $j_r^* \in \arg \max_{j \in B_r} \hat{\mu}_{j,n_r}$  (7.22) will be kept in  $B_{r+1}$ . Set  $\hat{\Delta}_{r+1} = \frac{\hat{\Delta}_r}{2} = 0.5^{r+1}$ . If  $|B_{r+1}| = 1$ , commit to that arm until the end of the learning.

**THEOREM 7.14.** Arm Elimination with Doubling-Trick has a  $\sum_{j \in [K]: \Delta_j > 0} \mathcal{O}\left(\frac{\ln(T\Delta_j^2)}{\Delta_j} + \frac{\ln K}{\Delta_j}\right)$  problem-dependent regret bound and an  $\mathcal{O}(\sqrt{KT \ln K})$  worst-case regret bound.

*Proof.* Let  $i_*$  denote the index of the unique optimal arm. Fix a sub-optimal arm j. Let  $r_j = \left[ \log \left( \frac{1}{\Delta_i} \right) \right]$ . Then, we have  $0.5\Delta_j \le 0.5^{r_j} = \hat{\Delta}_{r_j} \le \Delta_j$ .

We claim that the probability that this arm j is kept in  $B_{r_j+1}$  is very low. Formally, we have

$$\mathbb{P}\left\{j \in B_{r_{j}+1}\right\} = \mathbb{P}\left\{j \in B_{r_{j}}, \hat{\mu}_{j,n_{r_{j}}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}} \ge \max_{j \in B_{r_{j}}} \hat{\mu}_{j,n_{r_{j}}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}}\right\} \\ \le \mathbb{P}\left\{\hat{\mu}_{j,n_{r_{j}}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}} \ge \max_{j \in B_{r_{j}}} \hat{\mu}_{j,n_{r_{j}}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}}, i_{*} \in B_{r_{j}}\right\} \\ + \mathbb{P}\left\{\hat{\mu}_{j,n_{r_{j}}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}} \ge \max_{j \in B_{r_{j}}} \hat{\mu}_{j,n_{r_{j}}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}}, i_{*} \notin B_{r_{j}}\right\} \\ \le \mathbb{P}\left\{\hat{\mu}_{j,n_{r_{j}}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}} \ge \hat{\mu}_{i_{*},n_{r_{j}}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{2 \cdot n_{r_{j}}}}\right\}$$

$$(7.23)$$

 $\begin{array}{l} \text{UCB analysis, Hoeffding's inequality} \\ + \quad \mathbb{P}\left\{i^* \notin B_2\right\} + \mathbb{P}\left\{i^* \in B_2, \, i^* \notin B_3\right\} + \ldots + \mathbb{P}\left\{i^* \in B_2, \ldots, \, i^* \in B_{r_j-1}, \, i^* \notin B_{r_j}\right\} \\ \leq \quad \frac{2}{\text{KT}\hat{\Delta}_{r_j}^2} + \sum_{r=1}^{r_j-1} \frac{2}{\text{T}\hat{\Delta}_r^2} \\ \leq \quad \sum_{r=1}^{r_j} \frac{2}{\text{T}\hat{\Delta}_r^2} \\ = \quad \sum_{r=1}^{r_j} \frac{2}{\text{T}\cdot 0.5^{2r}} \\ = \quad \mathcal{O}\left(\frac{1}{\text{T}\cdot 0.5^{2r_j}}\right) \\ = \quad \mathcal{O}\left(\frac{1}{\text{T}\cdot \Delta_i^2}\right) \quad . \end{array}$ 

It is also fine to set  $\hat{\Delta}_0 = 0$ and start from r = 0.

You can view  $B_r$  as a version space in batch learning.

This is a picture for this.

The total regret from this sub-optimal arm j is at most

$$\underbrace{n_{r_{j}} \cdot \Delta_{j}}_{\text{regret until the end of epoch } r_{j}} + \underbrace{\mathbf{T} \cdot \mathbb{P}\left\{j \in \mathbf{B}_{r_{j}+1}\right\} \cdot \Delta_{j}}_{\text{regret nutil the end of epoch } r_{j}} \text{ regret for the remaining rounds} \\
\leq \frac{2\ln(\mathrm{KT}\hat{\Delta}_{r_{j}}^{2})}{\hat{\Delta}_{r_{j}}^{2}} \cdot \Delta_{j} + \mathbf{T} \cdot \mathcal{O}\left(\frac{1}{\mathrm{T} \cdot \Delta_{j}^{2}}\right) \cdot \Delta_{j} \tag{7.24} \\
\leq \frac{2\ln(\mathrm{KT}\Delta_{j}^{2})}{0.25\Delta_{j}^{2}} \cdot \Delta_{j} + \mathrm{T} \cdot \mathcal{O}\left(\frac{1}{\mathrm{T} \cdot \Delta_{j}^{2}}\right) \cdot \Delta_{j} \\
= \mathcal{O}\left(\frac{\ln(\mathrm{T}\Delta_{j}^{2})}{\Delta_{j}} + \frac{\ln \mathrm{K}}{\Delta_{j}}\right) .$$

Summing over all the sub-optimal arms, we have the problem-dependent regret is at most (1 - (T + 2))

$$\overline{\mathcal{R}}(\mathbf{T}) = \sum_{j \in [K]: \Delta_j > 0} \mathcal{O}\left(\frac{\ln(\mathbf{T}\Delta_j^2)}{\Delta_j} + \frac{\ln K}{\Delta_j}\right) \quad .$$
(7.25)

Let  $\tilde{\Delta} := \frac{e\sqrt{K \ln K}}{\sqrt{T}}$ . We have

$$\begin{aligned} \overline{\mathcal{R}}(\mathrm{T}) &\leq \mathrm{T} \cdot \widetilde{\Delta} + \sum_{j \in [\mathrm{K}]: \Delta_j > \widetilde{\Delta}} \mathcal{O}\left(\frac{\ln(\mathrm{T}\Delta_j^2)}{\Delta_j} + \frac{\ln \mathrm{K}}{\Delta_j}\right) \\ &\leq e\sqrt{\mathrm{K}\mathrm{T}\ln\mathrm{K}} + \sum_{j \in [\mathrm{K}]: \Delta_j > \widetilde{\Delta}} \mathcal{O}\left(\frac{\ln(\mathrm{T}\widetilde{\Delta}^2)}{\widetilde{\Delta}} + \frac{\ln \mathrm{K}}{\widetilde{\Delta}}\right) \\ &\leq e\sqrt{\mathrm{K}\mathrm{T}\ln\mathrm{K}} + \mathcal{O}\left(\frac{\mathrm{K}\ln(\mathrm{T}\widetilde{\Delta}^2)}{\widetilde{\Delta}} + \frac{\mathrm{K}\ln\mathrm{K}}{\widetilde{\Delta}}\right) \\ &\leq e\sqrt{\mathrm{K}\mathrm{T}\ln\mathrm{K}} + \mathcal{O}\left(\frac{\mathrm{K}\ln(e^2\mathrm{K}\ln\mathrm{K})}{\frac{e\sqrt{\mathrm{K}\ln\mathrm{K}}}{\sqrt{\mathrm{T}}}} + \frac{\mathrm{K}\ln\mathrm{K}}{\frac{e\sqrt{\mathrm{K}\ln\mathrm{K}}}{\sqrt{\mathrm{T}}}}\right) \\ &= \mathcal{O}(\sqrt{\mathrm{K}\mathrm{T}\ln\mathrm{K}}) \quad . \end{aligned}$$
(7.26)

We have

$$\begin{split} \mathbb{P}\left\{i^{*} \notin B_{2}\right\} &= \mathbb{P}\left\{\hat{\mu}_{i_{*},n_{1}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2 \cdot n_{1}}} < \max_{j \in B_{1} \setminus \{i_{*}\}} \hat{\mu}_{j,n_{1}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2 \cdot n_{1}}}\right\} \\ &\leq \sum_{j \in B_{1} \setminus \{i_{*}\}} \mathbb{P}\left\{\hat{\mu}_{i_{*},n_{1}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2 \cdot n_{1}}} < \hat{\mu}_{j,n_{1}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2 \cdot n_{1}}}\right\} \\ &\leq \sum_{j \in B_{1} \setminus \{i_{*}\}} \left(\mathbb{P}\left\{\hat{\mu}_{i_{*},n_{1}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2 \cdot n_{1}}} \le \mu_{1}\right\} + \mathbb{P}\left\{\hat{\mu}_{j,n_{1}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2 \cdot n_{1}}} \ge \mu_{j}\right\}\right) \\ &\leq \frac{2}{\mathrm{T}\hat{\Delta}_{1}^{2}} \quad . \end{split}$$
(7.27)

$$\mathbb{P}\left\{i^{*} \in B_{2}, i^{*} \notin B_{3}\right\}$$

$$= \mathbb{P}\left\{\hat{\mu}_{i_{*},n_{1}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2\cdot n_{1}}} \ge \max_{j \in B_{1}} \hat{\mu}_{j,n_{1}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{1}^{2})}{2\cdot n_{1}}}, \hat{\mu}_{i_{*},n_{2}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{2}^{2})}{2\cdot n_{2}}} < \max_{j \in B_{2} \setminus \{i_{*}\}} \hat{\mu}_{j,n_{2}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{2}^{2})}{2\cdot n_{2}}}\right\}$$

$$\le \mathbb{P}\left\{\hat{\mu}_{i_{*},n_{2}} + \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{2}^{2})}{2\cdot n_{2}}} < \max_{j \in B_{2} \setminus \{i_{*}\}} \hat{\mu}_{j,n_{2}} - \sqrt{\frac{\log(\mathrm{KT}\hat{\Delta}_{2}^{2})}{2\cdot n_{2}}}\right\}$$

$$\le \frac{2}{\mathrm{T}\hat{\Delta}_{2}^{2}} \quad .$$

$$(7.28)$$

Similarly, we have

$$\mathbb{P}\left\{i^* \in \mathcal{B}_2, \dots, i^* \in \mathcal{B}_{r_j-1}, i^* \notin \mathcal{B}_{r_j}\right\} \leq \frac{2}{T\hat{\Delta}_{r_j-1}^2} \quad . \qquad \Box$$

# REFERENCES

- [AO10] Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica* 61.1-2 (2010), pages 55–65.
- [BC12] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends*® *in Machine Learning* 5.1 (2012), pages 1–122. arXiv: 1204.5721.
- [SSBD14] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.