# **Online Learning**

1

Pierre Gaillard December 2. 2024

INRIA

These monographs are available online.

- Cesa-Bianchi and Lugosi 2006
- Shalev-Shwartz et al. 2012
- Hazan et al. 2016
- Lattimore and Szepesvári 2019

## Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

#### In classical supervised machine learning, the learner

- 1. observes training data with labels,
- 2. builds a program to minimize the training error
- 3. controls the error of new data if they are similar to the training data



 $\rightarrow$  Learning method  $\rightarrow$  Prediction on test data

In some applications, the environment may evolve over time and the data may be available sequentially.

**Spam detection**: can be seen as a game between spammer and spam filters. Each trying to fool the other one. The data is possibly adversarial.

Necessity to take a robust approach by learning as ones goes along from experiences as more aspects of the problem are observed.

This is the goal of sequential learning (or sequential learning).

**In sequential learning**, we do not have any training data. Data are acquired and treated on the fly. Feedbacks are received and algorithms updated step by step.



This field has received a lot of attention recently because of the possible applications coming from internet:

- ads to display,
- repeated auctions,

- spam detection,
- $\operatorname{experts}/\operatorname{algorithm}$  aggregation

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an action  $x_t \in \mathcal{K}$  (compact decision/parameter set);
- the environment chooses a loss function  $f_t: \mathcal{K} \to [0,1];$
- the player suffers loss  $f_t(x_t)$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K} \rightarrow full-information$  feedback
  - the loss of the chosen action only:  $f_t(x_t)$

 $\rightarrow$  bandit feedback.

Goal. Minimize the cumulative loss:

$$\widehat{\mathcal{L}}_{\mathcal{T}} \stackrel{\mathrm{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) \,.$$

# Example: Multi-armed bandit

(Thomson, 1933)

A simple stochastic model:

- K arms (actions: here price signals)
- Each arm k is associated an unknown probability distribution with mean  $\mu_k$



**Setting:** sequentially pick an arm  $k_t$  and get reward  $X_{k_t,t}$  with mean  $\mu_{k_t}$ 

Goal: maximize the expected cumulative reward

$$\mathbb{E}\left[\sum_{t=1}^{T} X_{k_t,t}\right]$$

Exploration vs Exploitation trade-off.

Maximize one's gains in casino? Hopeless ....



Historical motivation (Thomson, 1933): clinical trials, for each patient t in a clinical study

- choose a treatment  $k_t$
- observe response to the treatment  $X_{k_t,t}$

**Goal:** maximize the number of patient healed (or find the best treatment)

**Successful because of many applications coming from Internet**: recommender systems, online advertisements,...

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $x_t \in \mathcal{X}$  (optional step)
- the player chooses an action x<sub>t</sub> = k<sub>t</sub> ∈ K := {1,..., K} (compact decision/parameter set);
- the environment chooses a loss function  $f_t : \mathcal{K} \to [0, 1]$  (by sampling the arms);
- the player suffers loss  $f_t(x_t) = 1 X_{k_t,t}$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K} \rightarrow full-information$  feedback
  - the loss of the chosen action only:  $f_t(x_t) = X_{k_t,t}$
- $\rightarrow$  bandit feedback.

The goal of the player is to minimize his cumulative loss:

$$\widehat{L}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) \,.$$

There is some sequence of observations  $y_1, \ldots, y_T \in [0, 1]$  to be predicted step by step with the help of expert forecasts.

At each time step  $t \ge 1$ 

- the environment reveals experts forecasts  $c_t(k)$  for  $k=1,\ldots,K$
- the player chooses a weight vector  $p_t \in \Delta_K \stackrel{\text{def}}{=} \{ p \in [0, 1]^K : \sum_{k=1}^K p_k = 1 \}$ (here  $x_t$  is denoted  $p_t$  and  $\mathcal{K} = \Delta_K$ )
- the player forecasts  $\widehat{y}_t = \sum_{k=1}^{\mathcal{K}} p_t(k) c_t(k)$
- the environment reveals  $y_t \in [0,1]$  and the player suffers loss  $f_t(p_t) = f(\hat{y}_t, y_t)$  where  $f: [0,1]^2 \rightarrow [0,1]$  is a loss function.

Considering  $\mathcal{K} := \Delta_{\mathcal{K}}$  and  $x_t := p_t$ , we recover the general setting. The inputs correspond to the expert advice  $c_t(k)$  that are often revealed before the learner makes his decision  $p_t$ .

## **Example 2: Prediction with expert advice**

There is some sequence of observations  $y_1, \ldots, y_T \in [0, 1]$  to be predicted step by step with the help of expert forecasts.

At each time step  $t \ge 1$ 

- the environment reveals experts forecasts  $c_t(k)$  for  $k=1,\ldots,K$
- the player chooses a weight vector  $p_t \in \Delta_{\mathcal{K}} \stackrel{\text{def}}{=} \{ p \in [0, 1]^{\mathcal{K}} : \sum_{k=1}^{\mathcal{K}} p_k = 1 \}$ (here  $x_t$  is denoted  $p_t$  and  $\mathcal{K} = \Delta_{\mathcal{K}}$ )
- the player forecasts  $\widehat{y}_t = \sum_{k=1}^{K} p_t(k) c_t(k)$
- the environment reveals  $y_t \in [0,1]$  and the player suffers loss  $f_t(p_t) = f(\hat{y}_t, y_t)$  where  $f: [0,1]^2 \rightarrow [0,1]$  is a loss function.

 $f(\widehat{y}_t, y_t) = |\widehat{y}_t - y_t|/|y_t|$ 

- pinball loss.

Player's performance is then measured via a loss function  $f_t(p_t) = f(\hat{y}_t, y_t)$  which measures the distance between the prediction  $\hat{y}_t$  and the output  $y_t$ :

- squared loss  $f(\widehat{y}_t,y_t)=(\widehat{y}_t-y_t)^2$
- absolute loss  $f(\widehat{y}_t, y_t) = |\widehat{y}_t y_t|$
- absolute percentage of error

All these loss functions are convex, which will play an important role in the analysis.

Short term prediction (one day ahead) of the French electricity consumption



Important because electricity is hard to store.

Production



Demand

Short term prediction (one day ahead) of the French electricity consumption



Many experts (statisticians or data scientists) design prediction models:



Simultaneously, the French electricity market is evolving (electric cars,...)

Short term prediction (one day ahead) of the French electricity consumption



Many experts (statisticians or data scientists) design prediction models:



Simultaneously, the French electricity market is evolving (electric cars,...)

Short term prediction (one day ahead) of the French electricity consumption



Combine the predictions using adaptive methods:

Each day,

1. Assign a weight to each expert based on past performance

 $x_t = weight vector$ 

2. Predict the weighted average  $\hat{y}_t = \langle x_t, c_t \rangle$  and suffer loss

$$f_t(x_t) = \left(y_t - \widehat{y}_t\right)^2$$



If the environment chooses large losses  $f_t(x)$  for all decisions  $x \in \mathcal{K}$ , it is impossible for the player to ensure small cumulative loss.

 $\rightarrow$  Relative criterion: the regret of the player is the difference between the cumulative loss he incurred and that of the best fixed decision in hindsight.

## **Definition (Regret)**

The regret of the player with respect to a fixed parameter  $x^* \in \mathcal{K}$  after  $\mathcal{T}$  time steps is

$$\operatorname{Reg}_{T}(x^{*}) \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{T} f_{t}(x_{t}) - \sum_{t=1}^{T} f_{t}(x^{*}).$$

The regret (or uniform regret) is defined as  $\operatorname{Reg}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sup_{x^* \in \mathcal{K}} \operatorname{Reg}_{\mathcal{T}}(x^*)$ .

#### We have some approximation-estimation decomposition:

$$\sum_{t=1}^{T} f_t(x_t) = \underbrace{\inf_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t(x)}_{\text{Approximation error = how good the possible actions are.}} + \underbrace{\text{Reg}_{T}}_{\text{Sequential estimation error of the best action}}$$

#### We will focus on the regret in these lectures.

The goal of the player is to ensure a sublinear regret  $\text{Reg}_T = o(T)$  as  $T \to \infty$  and this for any possible sequence of losses  $f_1, \ldots, f_T$ .

 $\rightarrow$  the average performance of the player will approach on the long term the one of the best decision.

The losses  $f_t$  are unknown to the player beforehand and may be:

Adversarial setting (lessons 1, 2, and 3): No stochastic assumption on the process generating the losses f<sub>t</sub>. The latter are deterministic and may be chosen by some adversary. Typically, the problem can be seen as a game between the player who aims at optimizing with respect to x<sub>1</sub>,..., x<sub>T</sub> against an environment who aims at mazimizing with respect to loss<sub>t</sub>,..., loss<sub>T</sub> and x<sup>\*</sup>. Players's goal is to control the quantity:

 $\inf_{x_1} \sup_{f_1} \inf_{x_2} \sup_{f_2} \dots \inf_{x_T} \sup_{f_T} \sup_{x^* \in \mathcal{K}} \operatorname{Reg}_{\mathcal{T}}(x^*).$ 

- Stochastic setting (lessons 4, 5, and 6): the losses are generated by some stochastic process (e.g., i.i.d.). The regret bounds hold then in expectation or with high probability.

This may be caused by many phenomena, e.g. by

- some observation to be predicted if  $f_t(x) = f(x, y_t)$ . For instance, if the goal is to predict the evolution of the temperature  $y_1, \ldots, y_T$ , the latter changes over time and a prediction x is evaluated with  $f_t(x) = (x y_t)^2$ .
- noise: the environment is stochastic and the variation over time *t* models some noise effect.
- a changing environment. For instance, if the player is playing a game against some adversary that evolves and adapts to its strategy. A typical example is the case of spam detections. If the player tries to detect spams, while some spammers (the environment) try at the same time to fool the player with new spam strategies.

## **Exercise:** what about best $x_t^*$ at every round?

#### Regret

$$\operatorname{Reg}_{T} = \sum_{t=1}^{T} f_{t}(x_{t}) - \inf_{x^{*} \in \mathcal{K}} \sum_{t=1}^{T} f_{t}(x^{*})$$

Instead considering the regret with respect to a fixed  $x^* \in \mathcal{K}$ , one would be tempted to minimize the quantity

$$\operatorname{\mathsf{Reg}}_T^* \stackrel{\mathrm{\tiny def}}{=} \sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T \inf_{x \in \mathcal{K}} f_t(x)$$

where the infimum is inside the sum.

**Exercise:** Show that the environment can ensure  $\operatorname{Reg}_T^*$  to be linear in T by choosing properly the loss functions  $f_t$ .

Introduction: What is online learning?

Online Linear Optimization

Online Convex Optimization

Adversarial bandits

Stochastic bandits

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an action  $x_t \in \mathcal{K}$  (compact decision/parameter set);
- the environment chooses a loss function  $f_t : \mathcal{K} \to [0, 1]$ ;
- the player suffers loss  $f_t(x_t)$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K} \rightarrow \text{full-information feedback}$
- - the loss of the chosen action only:  $f_t(x_t) \rightarrow bandit$  feedback.

## **Goal.** Minimize the regret

$$\operatorname{Reg}_{T} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{T} f_{t}(x_{t}) - \inf_{x \in \mathcal{K}} \sum_{t=1}^{T} f_{t}(x)$$

We will start with the simple case where the decision set  $\mathcal{K}$  is the  $\mathcal{K}$ -dimensional simplex

$$\Delta_{K} \stackrel{ ext{def}}{=} \left\{ p \in \left[ 0,1 
ight]^{K} : \sum_{k=1}^{K} p_{k} = 1 
ight\}.$$
 (decision set)

Since the decisions  $x_t$  are probability distributions in  $\mathcal{K} = \Delta_{\mathcal{K}}$ , in this part we will denote them by  $p_t$  instead of  $x_t$ . We assume that the loss functions  $f_t$  are linear

$$\forall p \in \mathcal{K}, \qquad f_t(p) = \sum_{k=1}^{K} p(k)g_t(k) \in [-1, 1]$$
 (linear loss)

where  $g_t = (g_t(1), \ldots, g_t(K)) \in [-1, 1]^K$  is a loss vector chosen by the environment at round t.

At round t the player needs to choose a weight vector  $p_t \in \Delta_K$ .

#### How to choose the weights? The player should

- give more weight to actions that performed well in the past.
- not give all the weight to the current best action, otherwise it would not work (see Exercise next).

The exponentially weighted average forecaster (EWA) also called Hedge performs this trade-off by choosing a weight that decreases exponentially fast with the past errors.

Introduction: What is online learning?

Online Linear Optimization

The exponentially weighted average forecaster (EWA)

Application to prediction with expert advice

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

The exponentially weighted average forecaster

Parameter:  $\eta > 0$ Initialize:  $p_1 = \left(\frac{1}{K}, \dots, \frac{1}{K}\right)$ 

For  $t = 1, \ldots, T$ 

- select  $p_t$ ; incur loss  $f_t(p_t) = p_t^\top g_t$  and observe  $g_t \in [-1, 1]^K$ ;
- update for all  $k \in \{1, \dots, K\}$

$$p_{t+1}(k) = \frac{e^{-\eta \sum_{s=1}^{t} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t} g_s(j)}}$$

Consider the strategy, called "Follow The Leader" (FTL) that puts all the mass on the best action so far:

$$p_t \in \operatorname*{arg\,min}_{p \in \mathcal{K}} \sum_{s=1}^{t-1} f_s(p)$$
. (FTL)

#### **Exercise:**

- 1. Show that  $p_t(k) > 0$  implies that  $k \in \arg\min_j \sum_{s=1}^{t-1} g_s(j)$
- 2. Show that the regret of FTL might be linear: i.e., there exists a sequence  $g_1, \ldots, g_T \in [-1, 1]^K$  such that  $\text{Reg}_T \ge \Omega(T)$ .

## Solution

Consider the strategy, called "Follow The Leader" (FTL) that puts all the mass on the best action so far:

$$p_t \in \operatorname*{arg\,min}_{p \in \mathcal{K}} \sum_{s=1}^{t-1} f_s(p)$$
. (FTL)

#### Exercise:

1. Show that  $p_t(k) > 0$  implies that  $k \in \arg\min_j \sum_{s=1}^{t-1} g_s(j)$ 

#### Solution

Assume that there exists  $k \in [K]$  such that  $p_t(k) > 0$  and  $k \notin \arg \min_j \sum_{s=1}^{t-1} g_s(j)$ . Then, there exists  $k' \neq k$  such that  $\sum_{s=1}^{t-1} g_s(k') < \sum_{s=1}^{t-1} g_s(k)$ . Therefore,

$$\sum_{s=1}^{t-1} f_s(p_t) = \sum_{s=1}^{t-1} \sum_{j=1}^{K} p_s(j) g_s(j) = \sum_{s=1}^{t-1} \sum_{j \neq k} p_s(j) g_s(j) + p_s(k) \sum_{s=1}^{t-1} g_s(k)$$
$$> \sum_{s=1}^{t-1} \sum_{j \neq k} p_s(j) g_s(j) + p_s(k') \sum_{s=1}^{t-1} g_s(k) = \sum_{s=1}^{t-1} f_s(q_t),$$

where  $q_t(j) = p_t(j)$  if  $j \notin \{k, k'\}$  and  $q_t(k) = 0$  and  $q_t(k') = p_t(k') + q_t(k')$ . This yields a contradiction.

Consider the strategy, called "Follow The Leader" (FTL) that puts all the mass on the best action so far:

$$p_t \in \operatorname*{arg\,min}_{p \in \mathcal{K}} \sum_{s=1}^{t-1} f_s(p)$$
. (FTL)

#### Exercise:

- 1. Show that  $p_t(k) > 0$  implies that  $k \in \arg\min_j \sum_{s=1}^{t-1} g_s(j)$
- 2. Show that the regret of FTL might be linear: i.e., there exists a sequence  $g_1, \ldots, g_T \in [0, 1]^K$  such that  $\operatorname{Reg}_T \ge \Omega(T)$ .

#### Solution

It suffices to choose  $g_t(k) = 1$  if  $p_t(k) > 0$  and  $g_t(k) = 0$  otherwise. The cumulative loss of FTL is T while there exists an action with cumulative loss smaller then T/K.

## Regret guarantee for EWA

#### Theorem 1 (Regret bound for EWA)

Let  $T \ge 1$ . For all sequences of loss vectors  $g_1, \ldots, g_T \in [-1, 1]^K$ , EWA achieves the bound

$$\operatorname{Reg}_{T} \stackrel{\text{def}}{=} \sum_{t=1}^{T} f_{t}(p_{t}) - \min_{\rho \in \Delta_{K}} \sum_{t=1}^{T} f_{t}(\rho) \leqslant \eta \sum_{t=1}^{T} \sum_{k=1}^{K} p_{t}(k) g_{t}(k)^{2} + \frac{\log K}{\eta},$$
(1)

where we recall  $f_t : p \in \Delta_K \mapsto p^\top g_t$ .

Therefore, for the choice  $\eta = \sqrt{\frac{\log K}{T}}$ , EWA satisfies the regret bound  $\operatorname{Reg}_T \leq 2\sqrt{T \log K}$ .

This regret bound is optimal (see [1]).

**Exercise:** Generalize the above theorem when the losses  $g_1, \ldots, g_T \in [-B, B]^K$  for some B > 0.

[1] Cesa-Bianchi and Lugosi 2006.

## **Proof (Step 1 - Reformulation of the regret for linear losses)**

First, we remark that by definition of  $f_t: p \mapsto p \cdot g_t$  we have

R

$$eg_T \stackrel{\text{def}}{=} \sum_{t=1}^T f_t(p_t) - \min_{p \in \Delta_K} \sum_{t=1}^T f_t(p)$$
$$= \sum_{t=1}^T p_t \cdot g_t - \min_{p \in \Delta_K} \sum_{t=1}^T p \cdot g_t$$
$$= \sum_{t=1}^T p_t \cdot g_t - \min_{p \in \Delta_K} \sum_{k=1}^K \sum_{t=1}^T p(k)g_t(k)$$

Now, we can see that the minimum over  $p \in \Delta_{\mathcal{K}}$  is reached on a corner of the simplex. Therefore

$$\operatorname{Reg}_{T} = \sum_{t=1}^{T} p_{t} \cdot g_{t} - \min_{1 \leqslant k \leqslant K} \sum_{t=1}^{T} g_{t}(k)$$

# **Proof (Step 2 – Upper-bound of** $W_T$ )

We denote  $W_t(j) = e^{-\eta \sum_{s=1}^t g_t(j)}$  and  $W_t = \sum_{j=1}^K W_t(j)$ . The proof will consist in upper-bounding and lower-bounding  $W_T$ . We have

$$\begin{split} \mathcal{W}_{t} &= \sum_{j=1}^{K} W_{t-1}(j) e^{-\eta g_{t}(j)} & \leftarrow W_{t}^{(j)} = W_{t-1}(j) e^{-\eta g_{t}(j)} \\ &= W_{t-1} \sum_{j=1}^{K} \frac{W_{t-1}(j)}{W_{t-1}} e^{-\eta g_{t}(j)} \\ &= W_{t-1} \sum_{j=1}^{K} p_{t}(j) e^{-\eta g_{t}(j)} & \leftarrow p_{t}(j) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_{s}(j)}}{\sum_{k=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_{s}(k)}} = \frac{W_{t-1}(j)}{W_{t-1}} \\ &\leq W_{t-1} \sum_{j=1}^{K} p_{t}(j) (1 - \eta g_{t}(j) + \eta^{2} g_{t}(j)^{2}) & \leftarrow e^{x} \leq 1 + x + x^{2} \text{ for } x \leq 1 \\ &= W_{t-1} (1 - \eta p_{t} \cdot g_{t} + \eta^{2} p_{t} \cdot g_{t}^{2}), \end{split}$$

where we assumed in the inequality  $-\eta g_t(j) \leq 1$  and where we denote  $g_t = (g_t(1), \ldots, g_t(K))$ ,  $g_t^2 = (g_t(1)^2, \ldots, g_t(K)^2)$  and  $p_t = (p_t(1), \ldots, p_t(K))$ .

Now, using  $1 + x \leq e^x$ , we get:

$$W_t \leqslant W_{t-1}(1 - \eta p_t \cdot g_t + \eta^2 p_t \cdot g_t^2) \leqslant W_{t-1} \exp\left(-\eta p_t \cdot g_t + \eta^2 p_t \cdot g_t^2\right).$$

By induction on  $t = 1, \ldots, T$ , this yields using  $W_0 = K$ 

$$W_{T} \leqslant K \exp\left(-\eta \sum_{t=1}^{T} p_{t} \cdot g_{t} + \eta^{2} \sum_{t=1}^{T} p_{t} \cdot g_{t}^{2}\right).$$
(2)

## **Proof (Step 3 – Lower-bound of** $W_T$ **)**

On the other hand, upper-bounding the maximum with the sum,

$$\exp\left(-\eta\min_{j\in[K]}\sum_{t=1}^{T}g_t(j)\right)\leqslant\sum_{j=1}^{K}\exp\left(-\eta\sum_{t=1}^{T}g_t(j)\right)\leqslant W_{T}.$$

Combining the above inequality with Inequality (2) and taking the log, we get

$$-\eta \min_{j \in [K]} \sum_{t=1}^{T} g_t(j) \leqslant -\eta \sum_{t=1}^{T} p_t \cdot g_t + \eta^2 \sum_{t=1}^{T} p_t \cdot g_t^2 + \log K.$$
(3)

Dividing by  $\eta$  and reorganizing the terms proves the first inequality:

$$\operatorname{Reg}_{T} = \sum_{t=1}^{T} p_{t} \cdot g_{t} - \min_{1 \leq j \leq K} \sum_{t=1}^{T} g_{t}(j) \leq \eta \sum_{t=1}^{T} p_{t} \cdot g_{t}^{2} + \frac{\log K}{\eta}$$

Optimizing  $\eta$  and upper-bounding  $p_t \cdot g_t^2 \leqslant 1$  concludes the second inequality.

## Regret guarantee for EWA

#### Theorem 1 (Regret bound for EWA)

Let  $T \ge 1$ . For all sequences of loss vectors  $g_1, \ldots, g_T \in [-1, 1]^K$ , EWA achieves the bound

$$\operatorname{Reg}_{T} \stackrel{\text{def}}{=} \sum_{t=1}^{T} f_{t}(p_{t}) - \min_{\rho \in \Delta_{K}} \sum_{t=1}^{T} f_{t}(\rho) \leqslant \eta \sum_{t=1}^{T} \sum_{k=1}^{K} p_{t}(k) g_{t}(k)^{2} + \frac{\log K}{\eta},$$
(1)

where we recall  $f_t : p \in \Delta_K \mapsto p^\top g_t$ .

Therefore, for the choice  $\eta = \sqrt{\frac{\log K}{T}}$ , EWA satisfies the regret bound  $\operatorname{Reg}_T \leq 2\sqrt{T \log K}$ .

This regret bound is optimal (see [1]).

**Exercise:** Generalize the above theorem when the losses  $g_1, \ldots, g_T \in [-B, B]^K$  for some B > 0.

[1] Cesa-Bianchi and Lugosi 2006.
The previous algorithms EWA depends on a parameter  $\eta > 0$  that needs to be optimized according to K and T. For instance, for EWA using the value

$$\eta = \sqrt{\frac{\log K}{T}}$$

The bound of Theorem 1 is only valid for horizon T.

However, the learner might not know the time horizon in advance and one might want an algorithm with guarantees valid simultaneously for all  $T \ge 1$ .

We can avoid the assumption that T is known in advance, at the cost of a constant factor, by using the so-called doubling trick.

# Anytime algorithm: the doubling trick

Whenever we reach a time step t which is a power of 2, we restart the algorithm (forgetting all the information gained in the past) setting  $\eta$  to  $\sqrt{\log K/t}$ . Let us denote EWA-doubling this algorithm.

Theorem 2 (Anytime bound on the regret)

For all  $T \ge 1$ , the regret of EWA-doubling is then upper-bounded as:

 $\operatorname{Reg}_{T} \leqslant 7\sqrt{T \log K}$ .

The same trick can be used to turn most online algorithms into anytime algorithms (even in more general settings: bandits, general loss,...).

We can use the <u>doubling trick</u> whenever we have an algorithm with a regret of order  $\mathcal{O}(T^{\alpha})$  for some  $\alpha > 0$  with a known horizon T to turn it into an algorithm with a regret  $\mathcal{O}(T^{\alpha})$  for all  $T \ge 1$ .

### Proof

For simplicity we assume  $T = 2^{M+1} - 1$ . The regret of EWA-doubling is then upper-bounded as:

$$\operatorname{Reg}_{T} = \sum_{t=1}^{T} f_{t}(p_{t}) - \min_{p \in \Delta_{K}} \sum_{t=1}^{T} f_{t}(p)$$

$$\leqslant \sum_{t=1}^{T} f_{t}(p_{t}) - \sum_{m=0}^{M} \min_{p \in \Delta_{K}} \sum_{t=2^{m}}^{2^{m+1}-1} f_{t}(p)$$

$$= \sum_{m=0}^{M} \underbrace{\sum_{t=2^{m}}^{2^{m+1}-1} f_{t}(p_{t}) - \min_{p \in \Delta_{K}} \sum_{t=2^{m}}^{2^{m+1}-1} f_{t}(p)}_{R_{m}}$$

Now, we remark that each term  $R_m$  corresponds to the expected regret of an instance of EWA over the  $2^m$  rounds  $t = 2^m, \ldots, 2^{m+1} - 1$  and run with the optimal parameter  $\eta = \sqrt{\log K/2^m}$ . Therefore, using Theorem 1, we get  $R_m \leq 2\sqrt{2^m \log K}$ , which yields:

$$\operatorname{Reg}_{T} \leqslant \sum_{m=0}^{M} 2\sqrt{2^{m}\log K} \leqslant 2(1+\sqrt{2})\sqrt{2^{M+1}\log K} \leqslant 7\sqrt{T\log K}$$

Another solution is to use time-varying parameters  $\eta_t$  replacing T with the current value of t. The analysis is however less straightforward.

**Exercise**: Prove a regret bound for the time-varying choice  $\eta_t = \sqrt{\log K/t}$  in EWA.

Introduction: What is online learning?

### Online Linear Optimization

The exponentially weighted average forecaster (EWA)

Application to prediction with expert advice

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

## Reminder of the setting of prediction with expert advice

At each time step  $t \ge 1$ 

- the environment reveals experts forecasts  $c_t(k)$  for  $k = 1, \ldots, K$
- the player chooses a weight vector  $p_t \in \Delta_K \stackrel{\text{def}}{=} \{p \in [0,1]^K : \sum_{k=1}^K p_k = 1\}$ (here  $x_t$  is denoted  $p_t$  and  $\mathcal{K} = \Delta_K$ )
- the player forecasts  $\widehat{y}_t = \sum_{k=1}^{K} p_t(k) c_t(k)$
- the environment reveals  $y_t \in [0,1]$  and the player suffers loss  $f_t(p_t) = f(\hat{y}_t, y_t)$  where  $f: [0,1]^2 \rightarrow [0,1]$  is a loss function.

The goal is to minimize the regret with respect to the best expert

$$\operatorname{\mathsf{Reg}}_{T}^{\operatorname{expert}} \stackrel{\text{def}}{=} \sum_{t=1}^{T} f(\widehat{y}_{t}, y_{t}) - \min_{1 \leqslant k \leqslant K} \sum_{t=1}^{T} f(c_{t}(k), y_{t}),$$

where  $\hat{y}_t = p_t \cdot c_t$  are the prediction of the algorithm and  $y_t$  the observations to be predicted sequentially.

## Reminder of the setting of prediction with expert advice

At each time step  $t \ge 1$ 

- the environment reveals experts forecasts  $c_t(k)$  for  $k = 1, \ldots, K$
- the player chooses a weight vector  $p_t \in \Delta_K \stackrel{\text{def}}{=} \{ p \in [0,1]^K : \sum_{k=1}^K p_k = 1 \}$ (here  $x_t$  is denoted  $p_t$  and  $\mathcal{K} = \Delta_K$ )
- the player forecasts  $\widehat{y}_t = \sum_{k=1}^{K} p_t(k) c_t(k)$
- the environment reveals  $y_t \in [0,1]$  and the player suffers loss  $f_t(p_t) = f(\hat{y}_t, y_t)$  where  $f: [0,1]^2 \to [0,1]$  is a loss function.

Player's performance is then measured via a loss function  $f_t(p_t) = f(\hat{y}_t, y_t)$  which measures the distance between the prediction  $\hat{y}_t$  and the output  $y_t$ :

- squared loss  $f(\widehat{y}_t, y_t) = (\widehat{y}_t y_t)^2$
- absolute loss  $f(\widehat{y}_t, y_t) = |\widehat{y}_t y_t|$

- $f(\widehat{y}_t, y_t) = |\widehat{y}_t y_t|/|y_t|$
- pinball loss.

- absolute percentage of error

All these loss functions are convex, how can we apply our analysis for linear losses?

We state below a corrolary to Theorem 1 when the loss functions  $f(\cdot, \cdot)$  are convex in there first argument.

**Corollary 1 (Regret of EWA for prediction with expert advice and convex loss)** Let  $T \ge 1$ . Assume that the loss function  $f: (x, y) \in \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$  is convex and takes values in [-1, 1]. Then, EWA applied with the vector vectors  $g_t = (f(c_t(1), y_t), \dots, f(c_t(\mathcal{K}), y_t)) \in [-1, 1]^{\mathcal{K}}$  has a regret upper-bounded by

$$\mathsf{Reg}_{\mathcal{T}}^{\mathsf{expert}} \stackrel{\mathrm{def}}{=} \sum_{t=1}^{T} f(\widehat{y}_t, y_t) - \min_{1 \leqslant k \leqslant K} \sum_{t=1}^{T} f(c_t(k), y_t) \leqslant 2\sqrt{T \log K}$$

where  $\hat{y}_t = p_t \cdot c_t$  and were  $\eta > 0$  is well-tuned.

Therefore, the average error of the algorithm will converge to the average error of the best expert. This is the case for the square loss, the absolute loss or the absolute percentage of error. It suffices to remark that by convexity of  $f(\cdot, \cdot)$  in its first argument

$$\operatorname{Reg}_{T}^{\operatorname{expert}} = \sum_{t=1}^{T} f(p_{t} \cdot c_{t}, y_{t}) - \min_{1 \leq k \leq K} \sum_{t=1}^{T} f(c_{t}(k), y_{t})$$
$$\leq \sum_{t=1}^{T} p_{t} \cdot g_{t} - \min_{1 \leq k \leq K} \sum_{t=1}^{T} g_{t}(k) \stackrel{\text{def}}{=} \operatorname{Reg}_{T}.$$

The result is then obtained by Theorem 1.

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

Introduction: What is online learning?

**Online Linear Optimization** 

Online Convex Optimization

Online Mirror Descent and variants (EG, OGD)

Faster rate with curvature (OGD, EWA, ONS)

Adversarial bandits

Stochastic bandits

# Setting of an online learning problem/online convex optimization

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)

- the loss of the chosen action only:  $f_t(x_t)$ 

- the player chooses an action  $x_t \in \mathcal{K}$  (compact decision/parameter set);
- the environment chooses a loss function  $f_t : \mathcal{K} \to [0, 1];$
- the player suffers loss  $f_t(x_t)$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K}$ 
    - ightarrow full-information feedback
      - bandit feedback.

Goal. Minimize the regret

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) - \inf_{x \in \mathcal{K}} \sum_{t=1}^{\mathcal{T}} f_t(x)$$

**Previously:** we have seen an algorithm when  $f_t(x) = \langle x, g_t \rangle$  and  $\mathcal{K} = \Delta_{\mathcal{K}}$ . How to generalize for convex  $f_t$ ?

**Setting:** simplex decision set  $\mathcal{K} = \Delta_{\mathcal{K}}$ , convex and differentiable loss functions

**Assumptions and notations:** Actions are denoted by  $p_t$  (instead of  $x_t$ ). The losses are assumed to be convex and Lipschitz

 $\forall p_t \in \mathcal{K}, \qquad \left\| \nabla f_t(p_t) \right\|_{\infty} \leqslant G_{\infty}.$ 

We will see a simple trick, so-called the gradient trick that allows to extend the results we saw for linear losses to convex losses.

The resulted algorithm is called the Exponentiated Gradient forecaster (EG). It consists in playing EWA with the gradients  $g_t = \nabla f_t(p_t) \in [-G_{\infty}, G_{\infty}]^{\kappa}$  as loss vectors.



For  $g_t = \nabla f_t(x_t)$ , the linear loss  $\tilde{f}_t(x) = g_t^\top x$  satisfies for any  $x \in \mathcal{K}$  $f_t(x_t) - f_t(x) \leq g_t^\top (x_t - x) \leq \tilde{f}_t(x_t) - \tilde{f}_t(x)$ .

To prevent infinite regret, need finite  $|\tilde{f}_t(x)|$  and hence bounds on the dual norms of the domain and gradients

$$|\widetilde{f}_t(x)|\leqslant \|g_t\|_p\|x\|_q,\qquad rac{1}{p}+rac{1}{q}=1\,.$$

#### The Exponentiated Gradient forecaster (EG)

Parameter:  $\eta > 0$ Initialize:  $p_1 = \left(\frac{1}{K}, \dots, \frac{1}{K}\right)$ 

For  $t = 1, \ldots, T$ 

- select  $p_t$ ; incur loss  $f_t(p_t)$  and observe  $f_t : \mathcal{K} \to [0, 1]$ ;
- compute the gradient  $g_t = 
  abla f_t(p_t) \in [-G_\infty, G_\infty]^K$
- update for all  $k \in \{1, \dots, K\}$

$$p_{t+1}(k) = \frac{e^{-\eta \sum_{s=1}^{t} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t} g_s(j)}}$$

#### Theorem 3

Let  $T \ge 1$ . For all sequences of convex differentiable losses  $f_1, \ldots, f_T : \mathcal{K} \to \mathbb{R}$  with bounded gradient  $\max_{p \in \mathcal{K}} \|\nabla f_t(p)\|_{\infty} \leq G_{\infty}$ , EWA applied with  $g_t = \nabla f_t(p_t)$  achieves the regret bound

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\mathrm{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(p_t) - \min_{p \in \mathcal{K}} \sum_{t=1}^{\mathcal{T}} f_t(p) \leqslant \eta \, G^2_\infty \, \mathcal{T} + rac{\log \mathcal{K}}{\eta} \, .$$

Therefore, for the choice  $\eta = \frac{1}{G_{\infty}} \sqrt{\frac{\log K}{T}}$ , EWA satisfies the regret bound  $\operatorname{Reg}_{T} \leq 2G_{\infty}\sqrt{T\log K}$ .

(4)

## Proof

**1.** Apply the regret bound of EWA with  $g_t$  (see Theorem 1 of last class):

$$\sum_{t=1}^{T} p_t \cdot g_t - \min_{p \in \Delta_K} \sum_{t=1}^{T} p \cdot g_t \leqslant \eta \sum_{t=1}^{T} \sum_{k=1}^{K} p_t(k) g_t(k)^2 + \frac{\log K}{\eta}$$

Remark that the theorem also holds for loss vectors  $g_t \in [-G_{\infty}, G_{\infty}]^K$  as soon as  $\eta \leq 1/G_{\infty}$ . Upper-bounding  $g_t(j)^2 \leq \|\nabla f_t(p_t)\|_{\infty}^2 \leq G_{\infty}^2$ , substituting  $g_t = \nabla f_t(p_t)$ , this yields for all  $p \in \Delta_K$ 

$$\sum_{t=1}^{T} p_t \cdot \nabla f_t(p_t) - p \cdot \nabla f_t(p_t) \leqslant \eta T G_{\infty}^2 + \frac{\log K}{\eta}$$

2. Gradient inequality: by convexity of the losses

$$f_t(p_t) - f_t(p) \leqslant (p_t - p) \cdot 
abla f_t(p_t)$$

which yields

$$\sum_{t=1}^{T} f_t(p_t) - f_t(p) \leqslant \eta T G_{\infty}^2 + \frac{\log K}{\eta} \,.$$

**3. Optimize**  $\eta$ :  $\eta = \frac{1}{G_{\infty}} \sqrt{\frac{\log K}{T}}$ .

**Setting:** A sequence of observations  $y_1, \ldots, y_T \in [0, 1]$  is to be predicted with the help of K expert advice  $c_t(k) \in [0, 1]$  for  $1 \leq k \leq K$ . The learner predict  $\hat{y}_t = \sum_{k=1}^{K} p_t(k)c_t(k)$  and suffers a loss  $f(\hat{y}_t, y_t)$ .

If the loss function is convex and Lipschitz in its first argument, we can apply Theorem 3 with  $f_t : p \mapsto f(p \cdot c_t, y_t)$ .

For instance, with the absolute loss,  $\mathit{G}_{\infty}=1$  and EG satisfies:

$$\sum_{t=1}^{T} \left| \widehat{y}_t - y_t \right| - \min_{p \in \mathcal{K}} \sum_{t=1}^{T} \left| p \cdot c_t - y_t \right| \leqslant 2\sqrt{T \log K} \,.$$

Hence, on the long run we perform as good as the best convex combination of the experts which may outperform the best expert.

### Setting: convex differentiable Lipschitz loss function, convex and compact decision set ${\cal K}$

## **Online Gradient Descent (OGD)**

Parameter:  $\eta > 0$ Initialize:  $x_1 \in \mathcal{K}$  arbitrarily chosen

For  $t = 1, \ldots, T$ 

- select  $x_t$ ; incur loss  $f_t(x_t)$  and observe  $f_t:\mathcal{K} 
  ightarrow [0,1];$
- compute the gradient  $\nabla f_t(x_t)$
- update

$$x_{t+1} = \operatorname{Proj}_{\mathcal{K}}\left(x_t - \eta \nabla f_t(x_t)\right).$$

where  $\mathsf{Proj}_{\mathcal{K}}$  is the Euclidean projection onto  $\mathcal{K}.$ 

### **Online Gradient Descent**

$$x_{t+1} \leftarrow \mathsf{Proj}_{\mathcal{K}}\left(x_t - \eta \nabla f_t(x_t)\right)$$

#### Theorem 4 (Regret of OGD)

Let  $D, G, \eta > 0$ . Assume that  $\max_{x,x' \in \mathcal{K}} ||x - x'|| \leq D$  and. Then for any sequence  $f_1, \ldots, f_T$  of convex differentiable loss functions such that  $\max_{x \in \mathcal{K}} ||\nabla f_t(x)|| \leq G$ , the regret of OGD satisfies

$$\sum_{t=1}^{l} f_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{l} f_t(x) \leq \frac{D^2}{2\eta} + \frac{\eta}{2} G^2 T.$$

In particular, for  $\eta = \frac{D}{G\sqrt{T}}$ , we have  $\operatorname{Reg}_T \leqslant DG\sqrt{T}$ .

# Comparison of EG and OGD

Assume that  $\mathcal{K} = \Delta_{\mathcal{K}}$  is the simplex and the loss functions are sub-differentiable convex functions with  $\|\nabla f_t\|_{\infty} \leq G_{\infty}$ . Then both EG and OGD are possible algorithms (see Theorems 3 and 12).

We saw in Theorem 3 that EG has a regret bound  $\operatorname{Reg}_T \leq 2G_{\infty}\sqrt{T\log K}$ . In this case, for all  $p, p' \in \Delta_K$ 

$$\|p-p'\| = \sum_{k=1}^{K} (p(i)-p'(i))^2 \leq \sum_{i=1}^{K} |p(i)-p'(i)| \leq \sum_{i=1}^{K} p(i)+p'(i) = 2,$$

and  $\|\nabla f_t(p)\| \leq \sqrt{K} \|\nabla f_t(p)\|_{\infty} \leq \sqrt{K} G_{\infty}$ . Therefore, the regret of OGD is upper-bounded by  $R_t \leq G_{\infty} \sqrt{2KT}$ . Thus

$$\mathsf{EG:} \quad \mathsf{Reg}_{\mathcal{T}} \leqslant 2 \mathcal{G}_{\infty} \sqrt{\mathcal{T} \log K} \qquad \mathsf{and} \qquad \mathsf{OGD:} \quad \mathsf{Reg}_{\mathcal{T}} \leqslant \mathcal{G} \sqrt{2\mathcal{T}} \leqslant \mathcal{G}_{\infty} \sqrt{2\mathcal{KT}} \,.$$

The dependence on K of OGD is suboptimal in this case. This is solved by OMD, a generalization of both algorithms.

### **Online Gradient Descent**

$$x_{t+1} \leftarrow \mathsf{Proj}_{\mathcal{K}}\left(x_t - \eta \nabla f_t(x_t)\right)$$

#### Theorem 4 (Regret of OGD)

Let  $D, G, \eta > 0$ . Assume that  $\max_{x,x' \in \mathcal{K}} ||x - x'|| \leq D$  and. Then for any sequence  $f_1, \ldots, f_T$  of convex differentiable loss functions such that  $\max_{x \in \mathcal{K}} ||\nabla f_t(x)|| \leq G$ , the regret of OGD satisfies

$$\sum_{t=1}^{I} f_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{I} f_t(x) \leqslant \frac{D^2}{2\eta} + \frac{\eta}{2} G^2 T.$$

In particular, for  $\eta = \frac{D}{G\sqrt{T}}$ , we have  $\operatorname{Reg}_T \leq DG\sqrt{T}$ .

Recall the update of OGD:

$$\mathsf{OGD}: \quad x_{t+1} \leftarrow \mathsf{Proj}_{\mathcal{K}}\left(\underbrace{x_t - \eta \nabla f_t(x_t)}_{z_t}\right)$$

1. Upper-bound the regret with gradient inequality: by convexity

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) - f_t(x^*) \stackrel{\operatorname{\mathsf{Convexity}}}{\leqslant} \sum_{t=1}^{\mathcal{T}} \langle \nabla f_t(x_t), x_t - x^* \rangle$$

# Proof (Step 2)

2. Get a telescoping sum:

$$\begin{aligned} \|x_{t+1} - x^*\|^2 & \stackrel{\text{Projection}}{\leq} \|z_t - x^*\|^2 \\ &= \|x_t - \eta \nabla f_t(x_t) - x^*\|^2 \\ &= \|x_t - x^*\|^2 + \eta^2 \|\nabla f_t(x_t)\|^2 - 2\eta \langle \nabla f_t(x_t), x_t - x^* \rangle \end{aligned}$$

$$x_{t+1} \leftarrow \operatorname{Proj}_{\mathcal{K}}\left(\underbrace{x_t - \eta \nabla f_t(x_t)}_{z_t}\right)$$

Thus,

$$\langle 
abla f_t(x_t), x_t - x^* 
angle \leqslant rac{1}{2\eta} \left( \left\| x_t - x^* \right\|^2 - \left\| x_{t+1} - x^* \right\|^2 
ight) + rac{\eta}{2} \left\| 
abla f_t(x_t) \right\|^2$$

Summing over  $t = 1, \ldots, T$  and it telescopes

$$\operatorname{Reg}_{T} \leq \frac{1}{2\eta} \left( \left\| x_{1} - x^{*} \right\|^{2} - \left\| x_{\mathcal{I}+1} - x^{*} \right\|^{2} \right) + \frac{\eta}{2} G^{2} T$$
$$\leq \frac{D^{2}}{2\eta} + \frac{\eta G^{2} T}{2}$$

Exercise: Prove an upper-bound on the regret of OGD

- a) when  $\eta$  is calibrated with a doubling trick.
- b) when  $\eta$  is calibrated using a time-varying parameter  $\eta_t = D/(G\sqrt{t})$

**Exercise**: Prove an upper-bound on the regret of OGD with respect to any sequence of points  $x_1^*, \ldots, x_t^* \in \mathcal{K}$  such that  $\sum_{t=2}^T \|x_t^* - x_{t-1}^*\| \leq X$ 

$$\sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(x_t^*) \leqslant \ldots$$

### Generalization of OGD to better exploit the geometry of the decision space $\mathcal{K}$ .

OMD is the online counterpart of the Mirror Descent algorithm from convex optimization.

Gradient Descent updates are performed into a dual space defined by a function  $\psi : \mathcal{K} \to \mathbb{R}$  called mirror map. The goal is to transform the decision space into one more suited to Euclidean geometry.



#### **Definition (Bregman divergence)**

For any continuously differentiable convex function  $\psi,$  the Bregman divergence with respect to  $\psi$  is defined as

$$D_\psi(x,y)\leqslant\psi(x)-\psi(y)-
abla\psi(y)\cdot(x-y)\quadorall x,y\in\mathcal{K}$$
 .

It is the difference between the value of the regularization function at x and the value of its first order Taylor approximation.

$$\psi(x)$$

$$\psi(y)$$

$$\psi(y)$$

$$\psi(y) + \langle \nabla \psi(y), x - y \rangle$$

$$\psi(y) + \langle \nabla \psi(y), x - y \rangle$$

# Online Mirror Descent (OMD) – Lazy and agile versions

#### **Online Mirror Descent (OMD)**

Parameters:  $\eta > 0$ , mirror map  $\psi$ Initialize:  $x_1 \in \mathcal{K}$ ,  $z_1 = x_1$ 

For  $t = 1, \ldots, T$ 

- select  $x_t$ ; incur loss  $f_t(x_t)$  and observe  $f_t : \mathcal{K} \to [0, 1]$ ; compute the gradient  $\nabla f_t(x_t)$
- update  $z_t$  such that

$$\begin{aligned} \nabla \psi(z_{t+1}) &= \nabla \psi(x_t) - \eta \nabla f_t(x_t) &\leftarrow & \text{agile version} \\ \nabla \psi(z_{t+1}) &= \nabla \psi(z_t) - \eta \nabla f_t(x_t) &\leftarrow & \text{lazy version} \end{aligned}$$

- project according to the Bregman divergence

$$x_{t+1} \in \operatorname*{arg\,min}_{x \in \mathcal{K}} D_{\psi}(x, z_{t+1}).$$

**Main difference**: the lazy version performs all gradient steps in the dual space before computing the projection over  $\mathcal{K}$  to form predictions  $x_t$ , while the agile version alternates projections and gradient steps in the dual space.

$$\mathbf{OGD}: \quad x_{t+1} \leftarrow \operatorname{Proj}_{\mathcal{K}} \left( x_t - \eta \nabla f_t(x_t) \right) \qquad \mathbf{OMD}: \quad \frac{\nabla \psi(z_{t+1}) = \nabla \psi(x_t) - \eta \nabla f_t(x_t)}{x_{t+1} \in \arg \min_{x \in \mathcal{K}} D_{\psi}(x, z_{t+1})}$$

If  $\mathcal{K} \subset \mathbb{R}^d$ , we can choose  $\psi(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2$ .

Then

$$abla \psi(x) = x$$
 and  $D_\psi(x,y) = rac{1}{2} \|x-y\|^2$ .

Therefore, the update of OMD becomes  $z_{t+1} = x_t - \eta \nabla f_t(x_t)$  and  $x_{t+1} = \text{Proj}_{\mathcal{K}}(z_{t+1})$ .

We recover the online gradient descent algorithm.

## OMD (agile or lazy) with negative entropy = EG

$$\mathsf{EG}: \qquad \begin{aligned} g_t &= \nabla f_t(x_t) \\ x_{t+1}(k) &= \frac{x_t(k)e^{-\eta_{\mathcal{S}t}(k)}}{\sum_{j=1}^K x_t(j)e^{-\eta_{\mathcal{S}t}(j)}} \end{aligned} \qquad \mathsf{OMD}: \qquad \begin{aligned} \nabla \psi(z_{t+1}) &= \nabla \psi(x_t) - \eta \nabla f_t(x_t) \\ x_{t+1} &\in \operatorname{arg\,min}_{x \in \mathcal{K}} D_{\psi}(x, z_{t+1}) \end{aligned}$$

If  $\mathcal{K} = \Delta_{\mathcal{K}}$ . We consider the negative entropy  $\psi(x) = \langle x, \log x \rangle$  so that  $\nabla \psi(x) = 1 + \log x$ . The update of OMD is then

$$1 + \log(z_{t+1}(i)) = 1 + \log x_t(i) - \eta g_t(i),$$

where  $g_t = \nabla f_t(x_t) \in \mathbb{R}^K$ . This can be rewritten

$$z_{t+1}(i) = x_t(i)e^{-\eta g_t(i)}.$$

The projection to the simplex is a simple renormalization (exercise next), we thus recover EG.

# Exercice: Bregmann projection over the simplex is the renormalization

### Exercise

Let 
$$z \in \mathbb{R}^d_+$$
,  $\psi(x) = \langle x, \log x \rangle$  and  $x_* = \arg \min_{x \in \Delta_d} D_{\psi}(x, z)$ . Show that  
1. For any  $x \in \Delta_d$ ,  $D_{\psi}(x, z) = \langle x, \log \frac{x}{z} \rangle + ||z||_1 - 1 \ge -\log(||z||_1) + ||z||_1 - 1$   
2. For all  $k \in [d]$ ,  $x_* = z/||z||_1$ .

# Exercice: Bregmann projection over the simplex is the renormalization

### Exercise

Let 
$$z \in \mathbb{R}^d_+$$
,  $\psi(x) = \langle x, \log x \rangle$  and  $x_* = \arg \min_{x \in \Delta_d} D_{\psi}(x, z)$ . Show that  
1. For any  $x \in \Delta_d$ ,  $D_{\psi}(x, z) = \langle x, \log \frac{x}{z} \rangle + ||z||_1 - 1 \ge -\log(||z||_1) + ||z||_1 - 1$   
2. For all  $k \in [d]$ ,  $x_* = z/||z||_1$ .

### Solution:

#### 1.

$$\begin{split} D_{\psi}(x,z) &= \psi(x) - \psi(z) - \langle \nabla \psi(z), x - z \rangle = \langle x, \log x \rangle - \langle z, \log z \rangle - \langle 1 + \log z, x - z \rangle \\ &= \langle x, \log \frac{x}{z} \rangle + \|z\|_1 - 1 \stackrel{\text{Jensen}}{\geqslant} - \log\left(\left\langle x, \frac{z}{x} \right\rangle\right) + 1 - \|z\|_1 = -\log(\|z\|_1) + \|z\|_1 - 1 \end{split}$$

# Exercice: Bregmann projection over the simplex is the renormalization

### Exercise

Let 
$$z \in \mathbb{R}^d_+$$
,  $\psi(x) = \langle x, \log x \rangle$  and  $x_* = \arg \min_{x \in \Delta_d} D_{\psi}(x, z)$ . Show that  
1. For any  $x \in \Delta_d$ ,  $D_{\psi}(x, z) = \langle x, \log \frac{x}{z} \rangle + ||z||_1 - 1 \ge -\log(||z||_1) + ||z||_1 - 1$   
2. For all  $k \in [d]$ ,  $x_* = z/||z||_1$ .

## Solution:

#### 1.

$$\begin{split} D_{\psi}(x,z) &= \psi(x) - \psi(z) - \langle \nabla \psi(z), x - z \rangle = \langle x, \log x \rangle - \langle z, \log z \rangle - \langle 1 + \log z, x - z \rangle \\ &= \langle x, \log \frac{x}{z} \rangle + \|z\|_1 - 1 \stackrel{\text{Jensen}}{\geqslant} - \log\left(\left\langle x, \frac{z}{x} \right\rangle\right) + 1 - \|z\|_1 = -\log(\|z\|_1) + \|z\|_1 - 1 \end{split}$$

2.

$$D_{\psi}\left(rac{z}{\|z\|_{1}},z
ight) = \langle rac{z}{\|z\|_{1}}, -\log(\|z\|_{1})
angle + \|z\|_{1} - 1 = -\log(\|z\|_{1}) + \|z\|_{1} - 1$$

# Equivalent Formulation of OMD (agile version)

OMD (agile version) is equivalent to the following update

$$x_{t+1} = rgmin_{x \in \mathcal{K}} \left\{ \left\langle 
abla f_t(x_t), x - x_t 
ight
angle + rac{1}{\eta} D_\psi(x, x_t) 
ight\}$$

Indeed,

$$\begin{aligned} \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Big\{ \langle \nabla f_t(x_t), x - x_t \rangle + \frac{1}{\eta} D_{\psi}(x, x_t) \Big\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Big\{ \eta \langle \nabla f_t(x_t), x - \varkappa_t \rangle + \psi(x) - \psi(\varkappa_t) - \langle \nabla \psi(x_t), x - \varkappa_t \rangle \Big\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Big\{ \langle \eta \nabla f_t(x_t) - \nabla \psi(x_t), x \rangle + \psi(x) \Big\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Big\{ \langle -\nabla \psi(z_{t+1}), x \rangle + \psi(x) \Big\} \quad \text{since} \quad \nabla \psi(z_{t+1}) = \nabla \psi(x_t) - \eta \nabla f_t(x_t) \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Big\{ - \psi(z_{t+1}) - \langle \nabla \psi(z_{t+1}), x - z_{t+1} \rangle + \psi(x) \Big\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Big\{ D_{\psi}(x, z_{t+1}) \Big\} \end{aligned}$$

## Equivalent formulation of OMD (lazy version) – linearized FTRL

OMD (lazy version) is the linearized version of Follow The Regularized Leader (FTRL)

$$x_{t+1} = \operatorname*{arg\,min}_{x \in \mathcal{K}} \left\{ \sum_{s=1}^{t} \langle \nabla f_s(x_s), x - x_s \rangle + \frac{1}{\eta} D_{\psi}(x, x_1) \right\}$$

Indeed, for the lazy version  $\nabla \psi(z_{t+1}) = \nabla \psi(z_t) - \eta \nabla f_t(x_t) = \nabla \psi(z_1) - \sum_{s=1}^t \eta \nabla f_s(x_s)$ , and  $z_1 = x_1$ . Thus

$$\begin{aligned} \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \left\{ \sum_{s=1}^{t} \left\langle \nabla f_{s}(x_{s}), x - x_{s} \right\rangle + \frac{1}{\eta} D_{\psi}(x, x_{1}) \right\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \left\{ \eta \sum_{s=1}^{t} \left\langle \nabla f_{s}(x_{s}), x - \underbrace{x_{s}} \right\rangle + \psi(x) - \underbrace{\psi(x_{1})}_{} - \left\langle \nabla \psi(x_{1}), x - \underbrace{x_{t}} \right\rangle \right\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \left\{ \left\langle \sum_{s=1}^{t} \eta \nabla f_{s}(x_{s}) - \nabla \psi(x_{1}), x \right\rangle + \psi(x) \right\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \left\{ \left\langle -\nabla \psi(z_{t+1}), x \right\rangle + \psi(x) \right\} \\ &= \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \left\{ D_{\psi}(x, z_{t+1}) \right\} \end{aligned}$$

$$x_{t+1} = \underset{x \in \mathcal{K}}{\arg\min} \left\{ \begin{array}{ll} \left\langle \nabla f_t(x_t), x - x_t \right\rangle + \frac{1}{\eta} D_{\psi}(x, x_t) & \text{(agile)} \\ \sum_{s=1}^t \left\langle \nabla f_s(x_s), x - x_s \right\rangle + \frac{1}{\eta} D_{\psi}(x, x_1) & \text{(lazy)} \end{array} \right\}$$

#### Theorem 5 (Regret of OMD)

Let  $T \ge 1$ . Let  $\mathcal{K}$  be a compact and convex,  $x_1 \in \mathcal{K}$ , and  $\psi : \mathcal{K} \to \mathbb{R}$  be  $\mu$ -strongly convex mirror map with respect to some norm  $\|\cdot\|$ . Assume

$$D_{\psi}(x,x_1)^{1/2} \leqslant D$$
 and  $\|\nabla f_t(x_t)\|_* \leqslant G_*$ .

Then, for any convex  $f_1, \ldots, f_T : \mathcal{K} \to \mathbb{R}$ , the regret of OMD (agile or lazy version) with  $\eta = DG_*^{-1}\sqrt{2\mu/T}$  is upper bounded as

$$\operatorname{Reg}_{\mathcal{T}} \leqslant DG_* \sqrt{\frac{2T}{\mu}}$$

We exactly the regrets of OGD and EG since  $\psi : x \mapsto \frac{1}{2} ||x||_2$  and  $\psi : x \mapsto \langle x, \log x \rangle$  are resp. 1-strongly convex with  $\| \cdot \|_2$  and  $\| \cdot \|_1$  norms.
## **Proof (agile version)**

Let  $t \ge 1$  and  $x \in \mathcal{K}$ . Denote  $\Phi_t(x) = \langle \nabla f_t(x_t), x \rangle + \frac{1}{\eta} D_{\psi}(x, x_t)$ . From the optimality condition on  $x_{t+1}$ ,  $\langle \nabla \Phi_t(x_{t+1}), x_{t+1} - x \rangle \le 0$ , which entails  $\langle \nabla f_t(x_t), x_{t+1} - x \rangle \le \frac{1}{\eta} \langle \nabla \psi(x_{t+1}) - \nabla \psi(x_t), x - x_{t+1} \rangle$  and yields

$$egin{aligned} &f_t(x_t)-f_t(x)\leqslant \langle 
abla f_t(x_t),x_t-x
angle = \langle 
abla f_t(x_t),x_t-x_{t+1}
angle + \langle 
abla f_t(x_t),x_{t+1}-x
angle \ &\leqslant \langle 
abla f_t(x_t),x_t-x_{t+1}
angle + rac{1}{\eta}\langle 
abla \psi(x_{t+1})-
abla \psi(x_t),x-x_{t+1}
angle \ &= \langle 
abla f_t(x_t),x_t-x_{t+1}
angle + rac{1}{\eta}igg(D_\psi(x,x_t)-D_\psi(x,x_{t+1})-D_\psi(x_{t+1},x_t)igg) \ &\leqslant \langle 
abla f_t(x_t),x_t-x_{t+1}
angle + rac{1}{\eta}igg(D_\psi(x,x_t)-D_\psi(x,x_{t+1})-rac{\mu}{2}\|x_t-x_{t+1}\|^2igg)\,, \end{aligned}$$

where the last inequality because  $\psi$  is  $\mu$ -strongly convexity. Moreover,

$$\langle \nabla f_t(x_t), x_t - x_{t+1} \rangle \overset{\mathsf{Cauchy-Schwarz}}{\leqslant} \|x_t - x_{t+1}\| \left\| \nabla f_t(x_t) \right\|_* \overset{\mathsf{Young's}}{\leqslant} \frac{\mu}{2\eta} \|x_t - x_{t+1}\|^2 + \frac{\eta}{2\mu} \left\| \nabla f_t(x_t) \right\|_*^2$$

Combining the last two upper-bounds and summing over t = 1, ..., T entails the stated regret upper bound.

Left as exercise

$$\begin{array}{|c|c|c|c|c|} \hline x_{t+1} = \mathop{\arg\min}_{x \in \mathcal{K}} \left\{ \begin{array}{c} \langle \nabla f_t(x_t), x - x_t \rangle + \frac{1}{\eta} D_{\psi}(x, x_t) & (\mathsf{agile}) \\ \sum_{s=1}^t \langle \nabla f_s(x_s), x - x_s \rangle + \frac{1}{\eta} D_{\psi}(x, x_1) & (\mathsf{lazy}) \end{array} \right\} \xrightarrow{\mathcal{K}} \xrightarrow{v_{\psi}} \overbrace{\nabla \psi} \xrightarrow{v_{\psi}} \overbrace{\nabla \psi(x_t) - \eta_t \nabla f_t(x_t)} \xrightarrow{\nabla \psi} \xrightarrow{v_{\psi}} \overbrace{\nabla \psi(x_t) - \eta_t \nabla f_t(x_t)} \xrightarrow{v_{\psi}} \xrightarrow{v_{\psi}}$$

OMD provide efficient algorithms with generic  $O(\sqrt{T})$  regret bounds but suffer from some drawbacks

- $O(\sqrt{T})$  is sub-optimal for some losses with curvature (e.g., squared, logistic, or log loss) for which logarithmic regret may be achieved  $\rightarrow$  ONS
- needs bounded gradients which is not the case in some settings (e.g., portfolio selection).
- the best mirror map  $\psi$  should be specified by the user  $\rightarrow$  AdaGrad.

Introduction: What is online learning?

**Online Linear Optimization** 

**Online Convex Optimization** 

Online Mirror Descent and variants (EG, OGD)

Faster rate with curvature (OGD, EWA, ONS)

Adversarial bandits

Stochastic bandits

### Curvature

When losses have curvature, logarithmic regret may be achieved.

- Convexity:

$$f_t(x) \ge f_t(x_t) + \langle \nabla f_t(x_t), x - x_t \rangle$$

- Strong convexity:

$$f_t(x) \ge f_t(x_t) + \langle \nabla f_t(x_t), x - x_t \rangle + \frac{\gamma}{2} \|x - x_t\|^2$$

- Exp-concavity:  $x \mapsto \exp(-\eta f_t(x))$  is concave, which implies (exercise)

$$f_t(x) \ge f_t(x_t) + \langle 
abla f_t(x_t), x - x_t 
angle + rac{\gamma}{2} \langle 
abla f_t(x_t), x - x_t 
angle^2$$

for  $\gamma = \frac{1}{2} \min \left\{ \eta, \frac{1}{4G_*D} \right\}$ .



 $f_t(x)$ 

### Definition ( $\eta$ -exp-concavity)

For  $\eta \in \mathbb{R}$ , a function f is said to be  $\eta$ -exp-concave if  $x \mapsto e^{-\eta f(x)}$  is concave.

### **Properties:**

- Exp-concavity  $\Rightarrow$  convexity because  $-\log$  is convex and decreasing.
- Strong convexity + bounded domain and gradients  $\Rightarrow$  exp-concavity.
- $\eta$ -exp-concavity  $\Rightarrow \eta'$ -exp-concavity for  $0 \leqslant \eta' \leqslant \eta$ .

Many losses are exp-concave:

- strongly convex losses on bounded domain logistic loss
- squared loss is  $\frac{1}{2Y^2}\text{-expconcave on }[0,Y]$  relative entropy

**Exercise:** Prove the above the above facts.

## Exercise: Exp-concavity implies a quadratic lower-bound

Show that if the loss  $f_t$  is  $\eta$ -exp-concave with gradients bounded by G and a domain diameter bounded by D, then it can be lower-bounded by a quadratic approximation: for all  $x, x_t \in \mathcal{K}$ ,

$$f_t(x) \ge f_t(x_t) + \langle 
abla f_t(x_t), x - x_t 
angle + rac{\gamma}{2} \langle 
abla f_t(x_t), x - x_t 
angle^2$$

for  $\gamma = \frac{1}{2} \min \left\{ \eta, \frac{1}{4G_*D} \right\}$ .



#### Corollary 2 (Regret of EWA for prediction with expert advice and exp-concave loss)

In the setting of prediction with expert advice, if the loss functions  $f(\cdot, y_t)$  are  $\eta$ -exp-concave for all  $y_t$ , then EWA run with vectors  $g_t = (f(x_t(1), y_t), \dots, f(x_t(K), y_t)) \in \mathbb{R}^K$  with parameter  $\eta > 0$  satisfies

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}}^{\operatorname{expert}} \stackrel{\text{def}}{=} \sum_{t=1}^{\mathcal{T}} f(\widehat{y}_t, y_t) - \min_{1 \leqslant k \leqslant K} \sum_{t=1}^{\mathcal{T}} f(x_t(k), y_t) \leqslant \frac{\log K}{\eta} \,,$$

for all  $T \ge 1$ .

The worst-case regret does not increase with T but grows logarithmically in the dimension K.

# Proof (Step 1)

We define  $W_t(i) = e^{-\eta \sum_{s=1}^t g_s(i)}$  and  $W_t = \sum_{i=1}^N W_t(i)$ . We have  $W_t = \sum_{i=1}^N W_{t-1}(j) e^{-\eta g_t(j)}$  $\leftarrow W_t(j) = W_{t-1}(j)e^{-\eta g_t(j)}$  $= W_{t-1} \sum_{i=1}^{N} \frac{W_{t-1}(j)}{W_{t-1}} e^{-\eta g_t(j)}$  $= W_{t-1} \sum_{i=1}^{N} p_t(j) e^{-\eta g_t(j)}$  $\leftarrow p_t(j) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}{\sum_{k=1}^{N} e^{-\eta \sum_{s=1}^{t-1} g_s(k)}} = \frac{W_{t-1}(j)}{W_{t-1}}$  $\leq W_{t-1} \exp\left(-nf(p_t \cdot x_t, v_t)\right)$  $\leftarrow \eta$ -exp-concavity

Now, by induction on  $t = 1, \ldots, T$ , this yields using  $W_0 = K$ 

$$W_{\mathcal{T}} \leqslant K \exp\left(-\eta \sum_{t=1}^{\mathcal{T}} f(\widehat{y}_t, y_t)\right).$$
 (5)

On the other hand, upper-bounding the maximum with the sum,

$$\exp\left(-\eta\min_{j\in[K]}\sum_{t=1}^{T}g_t(j)\right) \leqslant \sum_{j=1}^{K}\exp\left(-\eta\sum_{t=1}^{T}g_t(j)\right) \leqslant W_{T}.$$

Combining the above inequality with Inequality (5) and taking the log concludes the proof.

Can we obtain a regret with respect to the best combination of experts

 $\min_{p} \sum_{t=1}^{T} f_t(p)$ 

instead of the regret with respect to the best fixed expert?

**Continuous EWA** 

$$p_t = \frac{\int_{\mathcal{K}} p e^{-\eta \sum_{s=1}^{t-1} f_s(p)} d\mu(p)}{\int_{\mathcal{K}} e^{-\eta \sum_{s=1}^{t-1} f_s(p)} d\mu(p)},$$

where  $\mu$  is the uniform (Lebesgue) measure on  $\mathcal{K} = \Delta_{\mathcal{K}}$ .

### Theorem 6 (Regret of continuous EWA)

Let  $T \ge 1$ . For all sequences of  $\eta$ -exp-concave losses  $f_1, \ldots, f_t$  the continuous EWA forecaster satisfies

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(p_t) - \inf_{\rho \in \mathcal{K}} \sum_{t=1}^{\mathcal{T}} f_t(\rho) \leqslant \frac{1 + (\mathcal{K} - 1) \log(\mathcal{T} + 1)}{\eta}$$

Nice theoretical result but hard to implement because of the integral.

In practice,  $p_t$  can be computed by using (1/T)-discretization grid of  $\mathcal{K}$  (bad complexity of order  $\mathcal{T}^{\mathcal{K}}$ !) or by using Monte-Carlo methods to approximate the integral.

The proof starts similarly to the one of Theorem 2. Let us denote  $W_t(p) = e^{-\eta \sum_{s=1}^t f_s(p)}$ ,  $W_t = \int_{\mathcal{K}} W_t(p) d\mu(p)$  and  $d\hat{\mu}_t(p) = W_t(p) d\mu(p)/W_t$ . Then,

$$\begin{split} \mathcal{N}_{\mathcal{T}} &= \int_{\mathcal{K}} e^{-\eta \sum_{t=1}^{T} f_{t}(p)} d\mu(p) \\ &= W_{\mathcal{T}-1} \int_{\mathcal{K}} \frac{W_{\mathcal{T}-1}(p)}{W_{\mathcal{T}-1}} e^{-\eta f_{\mathcal{T}}(p)} d\mu(p) \\ &= W_{\mathcal{T}-1} \int_{\mathcal{K}} e^{-\eta f_{\mathcal{T}}(p)} d\hat{\mu}_{\mathcal{T}-1}(p) & \leftarrow p_{\mathcal{T}} = \int_{\mathcal{K}} p d\hat{\mu}_{\mathcal{T}-1}(p) \\ &\leqslant W_{\mathcal{T}-1} \exp\left(-\eta f_{\mathcal{T}}(p_{\mathcal{T}})\right) & \leftarrow \eta\text{-exp-concavity} \\ &\leqslant \exp\left(-\eta \sum_{t=1}^{T} f_{t}(p_{t})\right), & \leftarrow \text{ induction} \end{split}$$

(6)

# **Proof (Step 2 – Lower bound of** $W_T$ **)**

For simplicity, we assume that  $f_t$  are continuous. Therefore the infimum is a minimum and let  $p^* \in \arg \min_{p \in \mathcal{K}} \sum_{t=1}^{T} f_t(p)$  and define

$$\mathcal{K}_arepsilon \stackrel{ ext{def}}{=} \Big\{ (1-arepsilon) p^* + arepsilon q, \quad q \in \mathcal{K} \Big\}, \qquad arepsilon \in (0,1)\,.$$

By exp-concavity of  $f_t$ , we have for all t and all  $p = (1 - \varepsilon)p^* + \varepsilon q$ 

$$e^{-\eta f_t(p)} \geqslant (1-arepsilon) e^{-\eta f_t(p^*)} + arepsilon e^{-\eta f_t(q)} \geqslant (1-arepsilon) e^{-\eta f_t(p^*)}$$

Therefore, for all  $p \in \mathcal{K}_{\varepsilon}$ 

$$e^{-\eta \sum_{t=1}^{T} f_t(p)} \geqslant (1-arepsilon)^T e^{-\eta \sum_{t=1}^{T} f_t(p^*)}$$

Integrating both parts over  $\mathcal{K}_{\varepsilon}$  and using  $\mu(\mathcal{K}_{\varepsilon}) = \varepsilon^{\mathcal{K}-1}\mu(\mathcal{K})$  (exercise) we get

$$W_{\mathcal{T}} \geq \int_{\mathcal{K}_{\varepsilon}} e^{-\eta \sum_{t=1}^{T} f_{t}(p)} d\mu(p) \geq \mu(\mathcal{K}) \varepsilon^{\mathcal{K}-1} (1-\varepsilon)^{\mathcal{T}} e^{-\eta \sum_{t=1}^{T} f_{t}(p^{*})}$$

Combining with (6), using  $W_0 = \mu(\mathcal{K})$ , taking the log and reorganizing the terms yields

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(p_t) - \sum_{t=1}^{\mathcal{T}} f_t(p^*) \leqslant \frac{(\mathcal{K}-1)\log\frac{1}{\varepsilon} + \mathcal{T}\log\frac{1}{1-\varepsilon}}{\eta}$$

Optimizing  $\varepsilon = 1/(T+1)$  concludes the proof since

$$T\lograc{1}{1-arepsilon} = T\log\left(1+rac{1}{T}
ight)\leqslant 1$$
 .

## Logarithmic regret for OGD under strong-convexity

Online Gradient Descent:

$$x_{t+1} \leftarrow \operatorname{Proj}_{\mathcal{K}} \left( x_t - \frac{\eta_t}{\eta_t} \nabla f_t(x_t) \right)$$

#### Theorem 7 (Regret of OGD under strong-convexity)

Let  $D, G, \gamma > 0$ . Assume that  $\max_{x,x' \in \mathcal{K}} ||x - x'|| \leq D$  and. Then for any sequence  $f_1, \ldots, f_T$  of  $\gamma$ -strongly convex differentiable loss functions such that  $\max_{x \in \mathcal{K}} ||\nabla f_t(x)|| \leq G$ , the regret of OGD with  $\eta_t = 1/(\gamma t)$  satisfies

$$\operatorname{Reg}_{T} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{T} f_{t}(x_{t}) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_{t}(x) \leqslant \frac{G^{2}}{2\gamma} (1 + \log T).$$

# Proof

1. Upper-bound the regret with strong convexity:

$$\mathsf{Reg}_{\mathcal{T}} \stackrel{\text{def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) - f_t(x^*) \stackrel{\mathsf{Strong Convexity}}{\leqslant} \sum_{t=1}^{\mathcal{T}} \langle \nabla f_t(x_t), x_t - x^* \rangle - \frac{\gamma}{2} \big\| x_t - x^* \big\|^2$$

2. Upper-bound the gradient term as for OGD analysis

$$\langle 
abla f_t(x_t), x_t - x^* 
angle \leqslant rac{1}{2\eta_t} \left( \|x_t - x^*\|^2 - \|x_{t+1} - x^*\|^2 
ight) + rac{\eta_t}{2} \|
abla f_t(x_t)\|^2$$

3. Substitute in the previous inequality and conclude

$$\begin{aligned} \mathsf{Reg}_{T} &\leq \sum_{t=1}^{T} \frac{1}{2\eta_{t}} \left( \left\| x_{t} - x^{*} \right\|^{2} - \left\| x_{t+1} - x^{*} \right\|^{2} \right) + \frac{\eta_{t} G^{2}}{2} - \frac{\gamma}{2} \left\| x_{t} - x^{*} \right\|^{2} \\ &= \frac{1}{2} \sum_{t=1}^{T} \underbrace{\left( \frac{1}{\eta_{t}} - \frac{1}{\eta_{t-1}} - \gamma \right)}_{t=1} \left\| x_{t} - x^{*} \right\|^{2} + \frac{G^{2}}{2} \sum_{t=1}^{T} \frac{1}{\gamma_{t}} \\ &\leq \frac{G^{2}}{\gamma} (1 + \log T) \end{aligned}$$

# **Online Newton Step (ONS)**

Idea: Include the quadratic upper-bound of the loss into FTRL:

$$x_t = \operatorname*{arg\,min}_{x \in \mathcal{K}} \left\{ \sum_{s=1}^{t-1} \langle \nabla f_s(x_s), x - x_s \rangle + \frac{\gamma}{2} \langle \nabla f_s(x_s), x - x_s \rangle^2 + \frac{\lambda}{2} \|x\|^2 \right\}$$

### Theorem 8 (Regret of ONS)

Let  $T \ge 1$ . Let  $\mathcal{K} \subset \mathbb{R}^d$  be a centered convex set with diameter D > 0. Then, for any *G*-Lipschitz,  $\eta$ -exp-concave losses  $f_1, \ldots, f_T : \mathcal{K} \to \mathbb{R}$ , the regret of ONS with  $\lambda > 0$  and  $\gamma \leqslant \frac{1}{2} \min \left\{ \eta, \frac{1}{4GD} \right\}$  satisfies

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \leqslant rac{\lambda}{2} D^2 + rac{d}{2\gamma} \log \left( 1 + rac{\gamma \, T G^2}{d\lambda} 
ight).$$

If the losses are  $\eta$ -exp-concave  $\rightarrow$  logarithmic regret.

**Complexity**: Ignoring the projection step (may be difficult), it requires the inversion of a  $d \times d$  matrices and could be done in  $O(d^3)$  operations per round.

# Proof (Step 1)

For any  $t \ge 1$ , define the matrix  $A_t = \gamma \sum_{s=1}^t \nabla f_t(x_t) \nabla f_t(x_t)^\top + \lambda I$  and set for any  $x \in \mathcal{K}$ 

$$\psi_t(x) = rac{\gamma}{2} \sum_{s=1}^t \langle 
abla f_s(x_s), x - x_s 
angle^2 + rac{\lambda}{2} \|x\|^2 \qquad ext{and} \qquad \Phi_t(x) = \sum_{s=1}^t \left\langle 
abla f_s(x_s), x - x_s 
ight
angle + \psi_t(x) \,.$$

Then,  $\psi_t$  and  $\Phi_t$  are 1-strongly convex with respect to the norm  $\|\cdot\|_{A_t}$ .

The ONS update is then

$$x_{t+1} = \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \left\{ \sum_{s=1}^{t} \left\langle \nabla f_s(x_s), x - x_s \right\rangle + \psi_t(x) \right\} = \underset{x \in \mathcal{K}}{\operatorname{arg\,min}} \Phi_t(x) + \psi_t(x) = \operatorname{arg\,min}_{x \in \mathcal{K}} \Phi_t(x) + \psi_$$

On the one hand, by Cauchy-Schwarz inequality

$$\sum_{t=1}^{T} \langle \nabla f_t(x_t), x_t - x_{t+1} \rangle \leqslant \frac{1}{2} \sum_{t=1}^{T} \left\| \nabla f_t(x_t) \right\|_{A_t^{-1}}^2 + \frac{1}{2} \sum_{t=1}^{T} \left\| x_t - x_{t+1} \right\|_{A_t}^2.$$
(7)

# Proof (Step 2)

On the other hand, fixing  $x \in \mathcal{K}$ , by definition of  $\Phi_t(x)$ 

$$\langle \nabla f_t(x_t), x_{t+1} - x_t \rangle = \Phi_t(x_{t+1}) - \Phi_{t-1}(x_{t+1}) - \frac{\gamma}{2} \langle f_t(x_t), x_{t+1} - x_t \rangle^2$$
(8)

Thus

$$\sum_{t=1}^{T} \langle \nabla f_t(x_t), x_{t+1} - x \rangle = -\Phi_{\mathcal{T}}(x) + \psi_t(x) + \sum_{t=1}^{T} \langle \nabla f_t(x_t), x_{t+1} - x_t \rangle \qquad \leftarrow \text{ by Definition of } \Phi_{\mathcal{T}}(x)$$

$$= -\Phi_{\tau}(x) + \psi_{\tau}(x) + \sum_{t=1}' \Phi_{t}(x_{t+1}) - \Phi_{t-1}(x_{t+1}) - \frac{\gamma}{2} \langle f_{t}(x_{t}), x_{t+1} - x_{t} \rangle^{2} \qquad \leftarrow \text{by (8)}$$

$$= \Phi_{\tau}(x_{\tau+1}) = \Phi_{\tau}(x) - \Phi_{0}(x_{1}) + \psi_{\tau}(x) + \sum_{t=1}^{T} \Phi_{t-1}(x_{t}) - \Phi_{t-1}(x_{t+1}) - \frac{\gamma}{2} \langle f_{t}(x_{t}), x_{t+1} - x_{t} \rangle^{2}$$

$$\leqslant \psi_{T}(x) - \frac{1}{2} \sum_{t=1}^{T} \left\| x_{t+1} - x_{t} \right\|_{A_{t-1}}^{2} - \frac{\gamma}{2} \langle f_{t}(x_{t}), x_{t+1} - x_{t} \rangle^{2} = \psi_{T}(x) - \frac{1}{2} \sum_{t=1}^{T} \left\| x_{t+1} - x_{t} \right\|_{A_{t}}^{2}.$$
(9)

where the last inequality is because  $\Phi_T(x_{T+1}) \leq \Phi_T(x)$  and by strong convexity of  $\Phi_{t-1}$  and by the optimality condition of  $x_t = \arg \min_{x \in \mathcal{K}} \Phi_{t-1}(x)$  we have  $\Phi_{t-1}(x_t) - \Phi_{t-1}(x) \leq -\frac{1}{2} ||x_t - x||^2_{A_{t-1}}$ .

# Proof (Step 3)

Combining (7) and (9) into the quadratic upper-bound yields

$$\sum_{t=1}^{T} f_t(x_t) - f_t(x) \overset{\text{Exp-concavity}}{\leqslant} \frac{\lambda}{2} \|x\|^2 - \psi_T(x) + \sum_{t=1}^{T} \langle \nabla f_t(x_t), x_t - x \rangle \leqslant \frac{\lambda}{2} \|x\|^2 + \frac{1}{2} \sum_{t=1}^{T} \left\| \nabla f_t(x_t) \right\|_{A_t^{-1}}^2.$$
(10)

Lemma 1 (Lemma 11.11 of Cesa-Bianchi and Lugosi 2006)

For any full rank matrix A and any vector  $x \in \mathbb{R}^d$ , then

$$x^{ op}(A + xx^{ op})^{-1}x = 1 - \det(A)/\det(A + xx^{ op})$$

Thus

$$\sum_{t=1}^T \left(1 - \frac{\det(A_{t-1})}{\det(A_t)}\right) \overset{1-u \leqslant -\log u}{\leqslant} \sum_{t=1}^T \log \frac{\det(A_t)}{\det(A_{t-1})} = \log \frac{\det(A_{\mathcal{T}+1})}{\det(A_0)} \leqslant d \log \left(1 + \frac{\gamma \, \mathcal{T} \mathcal{G}^2}{d\lambda}\right),$$

where the last line used that  $\|\nabla f_t(x_t)\| \leq G$ . Substituting into (10) concludes the proof.

Introduction: What is online learning?

Online Linear Optimization

Online Convex Optimization

Adversarial bandits

Stochastic bandits

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an action  $x_t \in \mathcal{K}$  (compact decision/parameter set);
- the environment chooses a loss function  $f_t:\mathcal{K} 
  ightarrow [0,1];$
- the player suffers loss  $f_t(x_t)$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K} \longrightarrow$  full-information feedback
  - the loss of the chosen action only:  $f_t(x_t) \rightarrow bandit$  feedback.

**Goal:** Minimize the regret

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) - \inf_{x \in \mathcal{K}} \sum_{t=1}^{\mathcal{T}} f_t(x)$$

### Previous results for full-information feedback

#### The Exponentially Weighted Average (EWA) forecaster

$$p_t(k) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}$$

(EWA)

achieves a cumulative regret  $\text{Reg}_T \lesssim \sqrt{T \log K}$  when the set of actions is the K-dimensional simplex and for linear losses  $f_t(p) = p^\top g_t$  with  $g_t \in [-1, 1]^K$ .

In particular, we saw the intermediate regret-bound if  $-\eta g_t(k)\leqslant 1$ 

$$\sum_{t=1}^T p_t \cdot g_t - \min_{1\leqslant j\leqslant K} \sum_{t=1}^T g_t(j) \leqslant \eta \sum_{t=1}^T \sum_{k=1}^K p_t(k)g_t(k)^2 + rac{\log K}{\eta} \,.$$

(\*)

Note that the loss vectors  $g_t$  may depend on past information  $p_1, g_1, \ldots, g_{t-1}, p_t$ .

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

The exponentially weighted average algorithm for bandits

Adversarial bandits with experts

OGD without Gradients

Stochastic bandits

# Adversarial multi-armed bandit and pseudo-regret

**Setting:**  $\mathcal{K} = \{1, \dots, K\}$ . At round *t*, the player chooses an action  $k_t \in \{1, \dots, K\}$  and suffers and observes the loss  $f_t(k_t) \in [0, 1]$  only.

**Regret** with respect to action  $k \in [K]$  by

$$\operatorname{\mathsf{Reg}}_{T}(k) \stackrel{\mathrm{\tiny def}}{=} \sum_{t=1}^{T} f_{t}(k_{t}) - \sum_{t=1}^{T} f_{t}(k).$$

Instead of minimizing the expected regret  $\mathbb{E}[\operatorname{Reg}_{T}] = \mathbb{E}[\max_{k} \operatorname{Reg}_{T}(k)]$ , we will start with an easier objective, the pseudo-regret.

### **Definition (Pseudo-regret)**

$$\bar{\operatorname{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \max_{k \in [K]} \mathbb{E} \left[ \operatorname{Reg}_{\mathcal{T}}(k) \right] = \max_{k \in [K]} \mathbb{E} \left[ \sum_{t=1}^{T} f_t(k_t) - \sum_{t=1}^{T} f_t(k) \right].$$
 (pseudo regret)

### **Oblivious vs adaptive adversary**

$$\bar{\mathsf{Reg}}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \max_{k \in [\mathcal{K}]} \mathbb{E}\big[\mathsf{Reg}_{\mathcal{T}}(k)\big] = \max_{k \in [\mathcal{K}]} \mathbb{E}\Big[\sum_{t=1}^{\mathcal{T}} f_t(k_t) - \sum_{t=1}^{\mathcal{T}} f_t(k)\Big]$$

The expectation is taken with respect to the randomness of the algorithm: the decisions  $k_t$  are random.

We can distinguish two types of adversaries:

- oblivious adversary: all the loss functions  $f_1, \ldots, f_t$  are chosen in advance before the game starts and do not depend on the past player decisions  $k_1, \ldots, k_T$ . In this case, the losses  $f_t(k)$  are determinist and there is thus equality:  $\bar{\text{Reg}}_T = \mathbb{E}[\text{Reg}_T]$ .
- adaptive adversary: the loss function  $f_t$  at round  $t \ge 1$  may depend on past information  $\sigma(k_1, \ldots, k_{t-1})$ . It is thus random. By Jensen's inequality  $\max_{k \in [K]} \mathbb{E}[\operatorname{Reg}_{\mathcal{T}}(k)] \le \mathbb{E}[\max_{k \in [K]} \operatorname{Reg}_{\mathcal{T}}(k)]$  and thus  $\operatorname{Reg}_{\mathcal{T}} \le \mathbb{E}[\operatorname{Reg}_{\mathcal{T}}]$ .

The Exponentially Weighted Average (EWA) forecaster

$$p_t(k) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}$$

**Question:** Can we use directly  $p_t(k)$  as defined by EWA with  $g_t = (f_t(1), \ldots, f_t(K))$  and sample  $k_t \sim p_t$  as we did for random EWA?

 $\Box$  Yes  $\Box$  No

#### The Exponentially Weighted Average (EWA) forecaster

$$p_t(k) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}$$

(EWA)

**Question:** Can we use directly  $p_t(k)$  as defined by EWA with  $g_t = (f_t(1), \ldots, f_t(K))$  and sample  $k_t \sim p_t$  as we did for random EWA?

**Answer:** No, since the player does not observe  $f_t(k)$  for  $k \neq k_t$  and cannot compute  $p_t$ .

### How to use EWA for bandits?

#### The Exponentially Weighted Average (EWA) forecaster

$$p_t(k) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}$$

٢	۱۸	1	٨	١
l	v	<i>v i</i>	4	J

**Question:** What about setting using  $f_t(k)$  if we observe it and 0 otherwise:

$$g_t(k) = \begin{cases} f_t(k) & \text{if } k = k_t & \leftarrow \text{i.e., decision } k \text{ is observed} \\ 0 & \text{otherwise} \end{cases}$$

#### $\Box$ Yes $\Box$ No

### How to use EWA for bandits?

The Exponentially Weighted Average (EWA) forecaster

$$p_t(k) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}$$

(EWA)

**Question:** What about setting using  $f_t(k)$  if we observe it and 0 otherwise:

$$g_t(k) = \begin{cases} f_t(k) & \text{if } k = k_t & \leftarrow \text{i.e., decision } k \text{ is observed} \\ 0 & \text{otherwise} \end{cases}$$
?

Answer: No, because this estimate would be biased:

$$\mathbb{E}_{k_t \sim p_t} \big[ g_t(k_t) \big] = p_t(k) f_t(k) \neq f_t(k) \,.$$

In other words, the actions that are less likely to be chosen by the algorithm (small weight  $p_t(k)$ ) are more likely to be unobserved and incur 0 loss. We need to correct this phenomenon.

#### The Exponentially Weighted Average (EWA) forecaster

$$p_t(k) = \frac{e^{-\eta \sum_{s=1}^{t-1} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t-1} g_s(j)}}$$

Therefore, we choose

$$g_t(k)=rac{f_t(k)}{p_t(k)}\mathbbm{1}\{k=k_t\}\,,$$

which leads to the algorithm EXP3 detailed below.

#### EXP3

Parameter:  $\eta > 0$ Initialize:  $p_1 = \left(\frac{1}{K}, \dots, \frac{1}{K}\right)$ 

For  $t = 1, \ldots, T$ 

- draw  $k_t \sim p_t$ ; incur loss  $f_t(k_t)$  and observe  $f_t(k_t) \in [0,1];$
- update for all  $k \in \{1, \dots, K\}$

$$p_{t+1}(k) = \frac{e^{-\eta \sum_{s=1}^{t} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t} g_s(j)}}, \quad \text{where } g_s(k) = \frac{f_s(k)}{p_s(k)} \mathbb{1}\{k = k_s\}$$

### Pseudo-Regret bound for EXP3

$$p_{t+1}(k) = \frac{e^{-\eta \sum_{s=1}^{t} g_s(k)}}{\sum_{j=1}^{K} e^{-\eta \sum_{s=1}^{t} g_s(j)}}, \quad \text{where } g_s(k) = \frac{f_s(k)}{p_s(k)} \mathbb{1}\{k = k_s\}$$
(EXP3)

### **Theorem 9**

Let  $T \ge 1$ . The pseudo-regret of EXP3 run with  $\eta = \sqrt{\frac{\log K}{KT}}$  is upper-bounded as:

$$\bar{\operatorname{Reg}}_{T} \stackrel{\text{def}}{=} \max_{k \in [K]} \mathbb{E} \left[ \sum_{t=1}^{T} f_{t}(k_{t}) - \sum_{t=1}^{T} f_{t}(k) \right] \leqslant 2\sqrt{KT \log K}$$

Applying EWA to the estimated losses  $g_t(j)$  that are completely observed and taking the expectation:

$$\mathbb{E}\left[\sum_{t=1}^{T} p_t \cdot g_t - \min_{j \in [K]} \sum_{t=1}^{T} g_t(j)\right] \leq \frac{\log K}{\eta} + \eta \sum_{t=1}^{T} \mathbb{E}\left[p_t \cdot g_t^2\right].$$
(\*)

The rest of the proof consists in computing the expectations:

 $\mathbb{E}[p_t \cdot g_t] = \mathbb{E}[f_t(k_t)], \qquad \mathbb{E}[g_t(j)] = \mathbb{E}[f_t(j)] \qquad \text{and} \qquad \mathbb{E}[p_t \cdot g_t^2] \leqslant K$ (11)

## Proof

Denote by  $\mathcal{F}_{t-1} \stackrel{\text{def}}{=} \sigma(p_1, f_1, k_1, \dots, k_{t-1}, p_t, f_t)$  the past information available at round t for the adversary (which cannot use the randomness of  $k_t$  but can use  $p_t$ ).

Note that  $f_t$  and  $p_t$  are  $\mathcal{F}_{t-1}$ -measurable by assumption.

1) Proof that  $\mathbb{E}[g_t(j)] = \mathbb{E}[f_t(j)]$ 

$$\forall j \in [\mathcal{K}] \qquad \mathbb{E}\Big[g_t(j)\Big|\mathcal{F}_{t-1}\Big] = \mathbb{E}\Big[\frac{f_t(j)}{p_t(j)}\mathbb{1}\{j=k_t\}\Big|\mathcal{F}_{t-1}\Big] = \sum_{k=1}^{\mathcal{K}} p_t(k)\frac{f_s(j)}{p_t(j)}\mathbb{1}\{j=k\} = f_t(j)$$

**2)** Proof that  $\mathbb{E}[p_t \cdot g_t] = \mathbb{E}[f_t(k_t)]$ 

$$\mathbb{E}\Big[p_t \cdot g_t\Big] = \mathbb{E}\Big[\sum_{j=1}^{K} p_t(j)g_t(j)\Big] = \mathbb{E}\Big[\sum_{j=1}^{K} p_t(j)\mathbb{E}\Big[g_t(j)\Big|\mathcal{F}_{t-1}\Big]\Big]$$
$$= \mathbb{E}\Big[\sum_{j=1}^{K} p_t(j)f_t(j)\Big] = \mathbb{E}\Big[\mathbb{E}\big[f_t(k_t)\big|\mathcal{F}_{t-1}\big]\Big] = \mathbb{E}\big[f_t(k_t)\big].$$
## Therefore, using

$$\mathbb{E}[p_t \cdot g_t] = \mathbb{E}[f_t(k_t)] \quad \text{and} \quad \mathbb{E}[g_t(j)] = \mathbb{E}[f_t(j)] \quad (12)$$

we have

$$\mathbb{E}\left[\sum_{t=1}^{T} p_t \cdot g_t - \min_{j \in [\mathcal{K}]} \sum_{t=1}^{T} g_t(j)\right] \ge \max_{j \in [\mathcal{K}]} \mathbb{E}\left[\sum_{t=1}^{T} p_t \cdot g_t - \sum_{t=1}^{T} g_t(j)\right]$$
$$= \max_{j \in [\mathcal{K}]} \mathbb{E}\left[\sum_{t=1}^{T} f_t(k_t) - \sum_{t=1}^{T} f_t(j)\right] = \bar{\operatorname{Reg}}_T.$$

## Proof

**3)** Proof that  $\mathbb{E}[p_t \cdot g_t^2] \leq K$ 

$$\mathbb{E}\left[p_t \cdot g_t^2\right] = \mathbb{E}\left[\sum_{j=1}^{K} p_t(j)g_t(j)^2\right] = \mathbb{E}\left[\sum_{j=1}^{K} p_t(j) \mathbb{E}\left[g_t(j)^2 \middle| \mathcal{F}_{t-1}\right]\right]$$
$$= \mathbb{E}\left[\sum_{j=1}^{K} \sum_{k=1}^{K} p_t(j)p_t(k) \left(\frac{f_t(j)}{p_t(j)}\mathbb{1}\{j=k\}\right)^2\right]$$
$$= \mathbb{E}\left[\sum_{j=1}^{K} \sum_{k=1}^{K} p_t(k) \frac{f_t(j)^2}{p_t(j)}\mathbb{1}\{j=k\}\right]$$
$$= \mathbb{E}\left[\sum_{j=1}^{K} f_t(j)^2\right] \leqslant K.$$

4) Conclusion. Substituting into Inequality (\*) yields

$$\bar{\operatorname{Reg}}_{\mathcal{T}} \leqslant rac{\log K}{\eta} + \eta K \mathcal{T}.$$

and optimizing  $\eta = \sqrt{KT/(\log K)}$  concludes.

The issue with the above regret bound is that it bounds the pseudo-regret and not the expected regret. This is because we have

$$\mathbb{E}\bigg[\min_{j}\sum_{t=1}^{T}g_{t}(j)\bigg] \leqslant \min_{j}\mathbb{E}\bigg[\sum_{t=1}^{T}g_{t}(j)\bigg] = \min_{j\in[K]}\mathbb{E}\bigg[\sum_{t=1}^{T}f_{t}(j)\bigg]$$

but not

$$\mathbb{E}\left[\min_{j}\sum_{t=1}^{T}g_{t}(j)\right] \nleq \mathbb{E}\left[\min_{j}\sum_{t=1}^{T}f_{t}(j)\right].$$
(13)

Hence, controlling the cumulative loss agains the best estimated action only controls the pseudo regret and not the true regret.

# EXP3.P

### EXP3.P

Parameters:  $\eta > 0, \beta \in (0, 1), \gamma \in (0, 1)$ Initialize:  $p_1 = (\frac{1}{K}, \dots, \frac{1}{K})$ 

For  $t = 1, \ldots, T$ 

- draw  $k_t \sim p_t$ ; receive reward  $r_t(k_t) = 1 f_t(k_t)$  and observe  $r_t(k_t) \in [0, 1]$ ;
- update for all  $k \in \{1, \dots, K\}$

$$p_{t+1}(k) = (1-\gamma) rac{e^{\eta \sum_{s=1}^t g_s(k)}}{\sum_{j=1}^K e^{\eta \sum_{s=1}^t g_s(j)}} + rac{\gamma}{K} \, ,$$

where  $g_s(k) = \frac{r_s(k)\mathbbm{1}\{k=k_s\}+\beta}{p_s(k)}$ .

The weights  $p_t(k)$  of EXP3.P are necessary larger than  $\gamma/K$  and thus  $|\eta g_t(j)| \leq 1$  as soon as  $\eta(1+\beta)K/\gamma \leq 1$ .

# Regret bound for Exp3.P

### Theorem 10

For well-chosen parameters  $\gamma \in (0,1)$ ,  $\beta \in (0,1)$  and  $\eta > 0$  satisfying  $\eta(1+\beta)K/\gamma \leq 1$ , for any  $\delta > 0$ , the EXP3.P algorithm achieves

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}}\leqslant 6\sqrt{\mathcal{T}\mathcal{K}\log\mathcal{K}} + \sqrt{rac{\mathcal{T}\mathcal{K}}{\log\mathcal{K}}\log(1/\delta)}\,.$$

with probability at least  $1 - \delta$ .

With the choice  $\delta = 1/T$  it yields

$$\mathbb{E}\big[\operatorname{\mathsf{Reg}}_{\mathcal{T}}\big] \leqslant 6\sqrt{\mathcal{T}\mathcal{K}\log\mathcal{K}} + \sqrt{\frac{\mathcal{T}\mathcal{K}}{\log\mathcal{K}}}\log(\mathcal{T}) + 1$$

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

The exponentially weighted average algorithm for bandits

Adversarial bandits with experts

OGD without Gradients

Stochastic bandits

# Setting of adversarial bandits with experts

#### Setting

At each time step  $t = 1, \ldots, T$ 

- N experts propose recommendations  $h_t(i) \in [K]$  for  $i \in [N]$
- the environment chooses a loss function  $f_t: \mathcal{K} \to [0,1];$
- the player chooses an action  $k_t \in [\mathcal{K}]$
- the player suffers loss  $f_t(k_t)$
- the player observes the loss of the chosen action only:  $f_t(k_t)$

Goal: compete with the best expert, i.e., minimize

$$\operatorname{Reg}_{T}^{\exp} \stackrel{\text{\tiny def}}{=} \max_{i=1,\ldots,N} \mathbb{E}\left[\sum_{t=1}^{T} f_{t}(k_{t}) - \sum_{t=1}^{T} f_{t}(h_{t}(i))\right]$$

with respect to the experts.

## By using EXP3 on the set of experts instead of the set of actions, we would get

 $\bar{\operatorname{Reg}}_{\mathcal{T}} \leqslant \sqrt{TN \log N}$ .

However it does not take into account the information on the reward of all experts that choose the same action  $h_t(i) = k_t$ .

# EXP4

#### EXP4

Parameter:  $\eta > 0$ 

For each round  $t = 1, \ldots, n$ 

- 1. Get expert advice  $h_t(1), \ldots, h_t(N) \in [K]$
- 2. Draw an expert  $i_t$  with probability distribution  $q_t \in \Delta_N$
- 3. Choose decision  $k_t = h_t(i_t)$
- 4. Compute the estimated loss for each decision

$$g_t(k) = \frac{f_t(k)}{p_t(k)} \mathbb{1}\{k = k_t\}$$

where  $p_t \stackrel{\text{def}}{=} \sum_{i=1}^{N} q_t(i) \delta_{f_t(i)} \in \Delta_{\mathcal{K}}.$ 

- 5. Compute the estimated loss of the experts component-wise  $g_t(h_t(i))$
- 6. Update the probability distribution over the experts component-wise

$$q_{t+1}(i) = \frac{\exp\left(-\eta \sum_{s=1}^{t} g_s(h_s(i))\right)}{\sum_{j=1}^{N} \exp\left(\eta \sum_{s=1}^{t} g_s(h_s(j))\right)}, \qquad \forall 1 \leqslant i \leqslant N$$

Initialize:  $q_1 = \left(\frac{1}{N}, \ldots, \frac{1}{N}\right)$ .

## Theorem 11

EXP4 with  $\eta = \sqrt{\log N/(KT)}$  satisfies  $\operatorname{Reg}_T^{exp} \leq 2\sqrt{TK \log N}$ .

Proof left as exercise.

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

The exponentially weighted average algorithm for bandits Adversarial bandits with experts

OGD without Gradients

Stochastic bandits

## Beyond finite set of actions?

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an action  $x_t \in \mathcal{K}$  (compact decision/parameter set);
- the environment chooses a loss function  $f_t : \mathcal{K} \to [0, 1];$
- the player suffers loss  $f_t(x_t)$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K} \longrightarrow$
  - the loss of the chosen action only:  $f_t(x_t)$

 $\begin{array}{rl} \rightarrow & \mbox{full-information feedback} \\ \rightarrow & \mbox{bandit feedback}. \end{array}$ 

The goal of the player is to minimize his cumulative loss:

$$\widehat{L}_T \stackrel{\text{\tiny def}}{=} \sum_{t=1}^T f_t(x_t) \,.$$

This lecture: we saw variants of EXP3 when  $\mathcal{K}$  is finite.

What if the losses  $f_t$  are convex but  $\mathcal{K}$  is any bounded convex set in  $\mathbb{R}^d$ ?

# **Online Gradient Descent**

In the full information setting (when gradient can be observed), we saw OGD algorithm:

$$x_{t+1} \leftarrow \mathsf{Proj}_{\mathcal{K}}\left(x_t - \eta \nabla f_t(x_t)\right)$$

#### Theorem 12 (Regret of OGD)

Let  $D, G, \eta > 0$ . Assume that  $\mathcal{K}$  has diameter bounded by D and the convex losses have sub-Gradients bounded by G in  $f_2$ -norm, the regret of OGD satisfies

$$\sum_{t=1}^{T} f_t(x_t) - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t(x) \leq DG\sqrt{T}.$$

How to adapt this algorithm to the bandit setting? That is, when only  $f_t(x_t)$  are observed and not  $\nabla f_t(x_t)$ ?

## **Point-wise gradient estimators**

$$x_{t+1} \leftarrow \operatorname{Proj}_{\mathcal{K}} \left( x_t - \eta \nabla f_t(x_t) \right)$$

Similarly to EXP3, the idea is to replace the gradient in OGD with unbiased estimators. That is try to find an observable random variable  $\hat{g}_t$  that satisfies

 $\mathbb{E}[\widehat{g}_t] \approx \nabla f_t(x_t)$ 

Example: one-dimensional gradient estimate

$$f'(x) = \lim_{\delta \to 0} \frac{f(x+\delta) - f(x-\delta)}{2\delta}$$

Thus we can define

$$\widehat{g}(x) = \begin{cases} \frac{f(x+\delta)}{\delta} & \text{with proba } \frac{1}{2} \\ -\frac{f(x-\delta)}{\delta} & \text{with proba } \frac{1}{2} \end{cases} \quad \text{which yields} \quad \mathbb{E}[\widehat{g}(x)] = \frac{f(x+\delta) - f(x-\delta)}{2\delta}$$

Thus in expectation, for small  $\delta$ ,  $\hat{g}(x)$  approximates f'(x).

## Point-wise gradient estimators: multi-dimensional case

We show here how the one-dimensional pointwise gradient estimator can be extended to the multi-dimensional case.

We define  $\hat{f}_t$  to be a smoothed version of the loss:

 $\widehat{f}_t(x) = \mathbb{E}_v \big[ f_t(x + \delta v) \big]$ 

where  $v \sim Unif(\mathbb{B})$ . If  $\delta$  is small,  $\hat{f}_t$  is a good approximation of  $f_t$ .

#### Lemma 2

Let  $\hat{f}_t(x) = \mathbb{E}[f_t(x + \delta v)]$  where  $v \sim Unif(\mathbb{B})$  be a smoothed version of the loss, then

$$\mathbb{E}_{u}\left[\frac{d}{\delta}f_{t}(x_{t}+\delta u)u\right]=\nabla\widehat{f}_{t}(x).$$

#### Proof.

Left as exercise. See Lem. 6.7, Hazan et al. 2016.

# **OGD** without Gradients

Similarly to EXP3, the idea is to replace the gradient in OGD with unbiased estimators.

### **OGD** without gradients

For  $t = 1, \ldots, T$ 

- Draw  $u_t \in \mathbb{S}$  uniformly at random in the unit sphere
- Set  $\widehat{x}_t = x_t + \delta u_t$  a random perturbation of the current point  $x_t$
- Play  $\hat{x}_t$
- Estimate the gradient in  $x_t$  with

$$\widehat{g}_t = rac{d}{\delta} f_t(\widehat{x}_t) u_t$$

- Update

$$x_{t+1} \leftarrow \mathsf{Proj}_{\mathcal{K}_{\delta}}\left(x_t - \eta \widehat{g}_t\right)$$

where  $\mathcal{K}_{\delta} = \{ x \in \mathcal{K} \quad \text{s.t} \quad x + \delta u \in \mathcal{K} \quad \forall u \in \mathbb{S} \}$ 

# Regret of OGD without gradients

### OGD without gradients:

$$x_{t+1} \leftarrow \operatorname{Proj}_{\mathcal{K}_{\delta}}(x_t - \eta \widehat{g}_t)$$
 where  $\widehat{g}_t = \frac{d}{\eta} f_t(\widehat{x}_t) u_t$  and  $\widehat{x}_t = x_t + \delta u_t$ 

#### Theorem 13

If the losses are in [-1,1] and G-Lipschitz, OGD without gradients with parameters  $\delta = \min\{D, (1/2)\sqrt{Dd/G}T^{-1/4}\}$  and  $\eta = D\delta/(dT^{1/2})$  satisfies the expected regret bound

$$\sum_{t=1}^{T} \mathbb{E} \left[ f_t(\widehat{x}_t) \right] - \min_{x \in \mathcal{K}} \sum_{t=1}^{T} f_t(x) \leq 2d\sqrt{T} + 2\sqrt{GDd} T^{3/4}$$

# Proof (Step 1)

#### Denote

$$x^* \in \operatorname*{arg\,min}_{x \in \mathcal{K}} \sum_{t=1}^T f_t(x) \quad \text{and} \quad x^*_{\delta} = \operatorname{Proj}_{\mathcal{K}_{\delta}}(x^*).$$

Then,

 $\left\|x^* - x^*_{\delta}\right\| \leqslant \delta$ 

Thus, if the losses are G-Lipschitz

$$\operatorname{Reg}_{T} := \sum_{t=1}^{T} \mathbb{E}[f_{t}(\widehat{x}_{t})] - \sum_{t=1}^{T} f_{t}(x^{*}) \leqslant \sum_{t=1}^{T} \mathbb{E}[f_{t}(\widehat{x}_{t})] - \sum_{t=1}^{T} f_{t}(x^{*}_{\delta})$$
$$\leqslant \sum_{t=1}^{T} \mathbb{E}[f_{t}(x_{t})] - \sum_{t=1}^{T} f_{t}(x^{*}_{\delta}) + \delta TG$$
$$\leqslant \sum_{t=1}^{T} \mathbb{E}[\widehat{f}_{t}(x_{t})] - \sum_{t=1}^{T} \widehat{f}_{t}(x^{*}_{\delta}) + 3\delta TG \qquad (*)$$

where  $\widehat{f}_t(x) = \mathbb{E}_v[f_t(x + \delta v)]$  with  $v \sim Unif(\mathbb{B})$  are the smoothed versions of the losses.

# Proof (Step 2)

Now, recall that the algorithm runs OGD with  $\hat{g}_t$  in place of the gradients:

 $x_{t+1} \leftarrow \operatorname{Proj}_{\mathcal{K}_{\delta}}(x_t - \eta \widehat{g}_t)$ 

Defining the pseudo-loss  $h_t(x) = \hat{f}_t(x) + (\hat{g}_t - \nabla \hat{f}_t(x_t))^\top x$ , we can see that

$$abla h_t(x_t) = 
abla \widehat{f}_t(x_t) + \widehat{g}_t - 
abla \widehat{f}_t(x_t) = \widehat{g}_t$$

Therefore, the algorithm actually runs OGD on the losses  $h_t$  and thus satisfies the OGD regret bound (see Lecture 2)

$$\sum_{t=1}^{T}h_t(x_t)-\sum_{t=1}^{T}h_t(x^*_\delta)\leqslant rac{D^2}{2\eta}+rac{\eta}{2}\sum_{t=1}^{T}\|\widehat{g}_t\|^2\,.$$

Furthermore, by construction of the gradient estimator, we have  $\mathbb{E}_{u_t}[\widehat{g}_t] = \nabla \widehat{f}_t(x_t)$ , which yields

$$\mathbb{E}_{u_t}[h_t(x_t)] = \widehat{f}_t(x_t) \quad \text{and} \quad \mathbb{E}_{u_t}[h_t(x_{\delta}^*)] = \widehat{f}_t(x_{\delta}^*)$$

Thus taking the expectation in the previous regret bound entails

$$\sum_{t=1}^{T} \mathbb{E}[\widehat{f}_t(x_t)] - \sum_{t=1}^{T} \widehat{f}_t(x_{\delta}^*) = \mathbb{E}\left[\sum_{t=1}^{T} h_t(x_t) - \sum_{t=1}^{T} h_t(x_{\delta}^*)\right] \leq \frac{D^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \mathbb{E}[\|\widehat{g}_t\|^2]$$
(\*\*)

# Proof (Step 3)

Combining the two bounds (\*) and (\*\*) that we have proved, we get

$$\operatorname{Reg}_{T} \leqslant \frac{D^{2}}{2\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \mathbb{E} \left[ \|\widehat{g}_{t}\|^{2} \right] + 3\delta TG$$

Then, since  $|f_t(x)| \leq 1$  for all  $x \in \mathcal{K}$ ,

$$\|\widehat{g}_t\|^2 = \left(\frac{d}{\delta}f_t(\widehat{x}_t)\right)^2 \leqslant \frac{d^2}{\delta^2}$$

This finally yields the regret

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} \leqslant \frac{D^2}{2\eta} + \frac{\eta d^2 T}{2\delta^2} + 3\delta TG \leqslant 2d\sqrt{T} + 2\sqrt{GDd} T^{3/4}$$

for the choices of  $\delta$  and  $\eta$ .

Convex bandits is still an active research area with many open problems.

The above regret bound of order  $O(T^{3/4})$  is suboptimal.

More complicated methods can achieve  $O(\sqrt{T})$  regret but with sub-optimal dependence on d and worst computational complexities.

More information can be found in Hazan et al. 2016.

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

The stochastic bandit problem

Consequences of the stochasticity

Exploration and exploitation

Optimism in face of uncertainty: UCB

Stochastic Linear Bandits

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an action  $x_t \in \mathcal{K}$  (compact decision/parameter set);
- the environment chooses a loss function  $f_t: \mathcal{K} \to [0,1];$
- the player suffers loss  $f_t(x_t)$  and observes
  - the losses of every actions:  $f_t(x)$  for all  $x \in \mathcal{K}$  –

- the loss of the chosen action only:  $f_t(x_t)$ 

The goal of the player is to minimize his cumulative loss:

$$\widehat{L}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \sum_{t=1}^{\mathcal{T}} f_t(x_t) \,.$$

- ightarrow full-information feedback
- $\rightarrow$  bandit feedback.

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an arm  $k_t \in \mathcal{K}$  (compact decision/parameter set, most often  $\{1, \ldots, \mathcal{K}\}$ );
- the player observes
  - the rewards of every arm:  $X_t^k \sim \nu_k$  for all  $k \in \mathcal{K} \to -$  full-information feedback
  - the reward of the chosen arm only:  $X_t^{k_t} \sim \nu_{k_t} \longrightarrow$  bandit feedback.

The goal of the player is to maximize their cumulative reward.

# **Regret?**

### We could use the definition of the regret from adversarial bandits:

Definition (Regret, attempt 1)

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} = \max_k \sum_{t=1}^{\mathcal{T}} X_t^k - \sum_{t=1}^{\mathcal{T}} X_t^{k_t} \,.$$

Let's see why we don't use that definition.

Notations and assumptions:

- The arm set is  $[K] = \{1, \dots, K\}.$
- $\mu^k = \mathbb{E}_{X \sim \nu_k}[X]$ , assumed finite for all arms k.
- $\mu^* = \max_{k \in [K]} \mu^k$ .

## The first notion of regret is inadequate

$$\operatorname{\mathsf{Reg}}_{\mathcal{T}} = \max_{k} \sum_{t=1}^{\mathcal{T}} X_{t}^{k} - \sum_{t=1}^{\mathcal{T}} X_{t}^{k_{t}}$$

 $u_k$  Bernoulli(1/2) for all  $k \in [K]$ .  $\mu^k = 1/2$  for all k.

All arms are the same  $\rightarrow$  there is no bad choice and no bad algorithm.

But:

$$\mathbb{E}\operatorname{Reg}_{T} = \mathbb{E}[\max_{k \in [K]} \sum_{t=1}^{T} X_{t}^{k}] - T/2$$
$$= \mathbb{E}[\max_{k \in [K]} \sum_{t=1}^{T} (X_{t}^{k} - 1/2)]$$
$$\approx \sqrt{T \log K}$$

(See any course/book/wikipedia article on symmetric random walks).

We want a regret notion that does not blow up with stochastic fluctuations.

The regret is defined as

$$\operatorname{Reg}_{T} = \max_{k} \sum_{t=1}^{T} \mu^{k} - \sum_{t=1}^{T} \mu^{k_{t}} = T \mu^{*} - \sum_{t=1}^{T} \mu^{k_{t}}$$

Recall that  $\mu^k = \mathbb{E}_{X \sim \nu_k}[X].$ 

Most often, we bound the expected regret  $\mathbb{E}[\operatorname{Reg}_{\mathcal{T}}]$ .

Note that the expectation here is over the random rewards and the randomness of the algorithm, if there is any.

Suppose that the set of arms is finite: [K].

Define the gap of arm  $k \in [K]$  by  $\Delta_k = \mu^* - \mu^k$ .

$$\operatorname{Reg}_{T} = T\mu^{*} - \sum_{t=1}^{T} \mu^{k_{t}} = \sum_{t=1}^{T} (\mu^{*} - \mu^{k_{t}}) = \sum_{t=1}^{T} \Delta_{k_{t}} = \sum_{k=1}^{K} N_{T}^{k} \Delta_{k} ,$$

where  $N_T^k = \sum_{t=1}^T \mathbb{I}\{k_t = k\}$  is the number of pulls of arm k up to time T.

Bounding the regret  $\Leftrightarrow$  bounding the number of pulls of bad arms

At each time step  $t = 1, \ldots, T$ 

- the player observes a context  $c_t \in \mathcal{X}$  (optional step)
- the player chooses an arm  $k_t \in \mathcal{K}$  (compact decision/parameter set, most often  $\{1, \ldots, K\}$ );
- the player observes the reward of the chosen arm only:  $X_t^{k_t} \sim \nu_{k_t}$  (independent of other rewards).

The goal of the player is to minimize their expected regret:  $\mathbb{E}[\operatorname{Reg}_T] = \sum_{k=1}^{K} \mathbb{E}[N_T^k] \Delta_k$ .

## Setting variants:

- Contextual bandit:  $X_t^{k_t} \sim \nu_{k_t}(c_t)$ , for a known context  $c_t$
- Linear bandit:  $\nu_{k_t} = \mathcal{N}(x^{ op} c_{k_t}, 1)$
- Structured bandit: the algorithm knows constraints on (μ<sup>k</sup>)<sub>k∈[K]</sub>, e.g. Lipschitz, linear, monotone...

## Goal variants: instead of minimizing the regret, we want to

- Minimize the simple regret: return an arm at time T, and minimize its expected gap.
- Identify the best arm: return an arm at time *T*, and minimize the probability that its not one of the best ones.

Relaxed assumptions: rewards not independent, distributions changing with time, etc.

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

Stochastic bandits

The stochastic bandit problem

Consequences of the stochasticity

Exploration and exploitation

Optimism in face of uncertainty: UCB

Stochastic Linear Bandits

Main idea: we can estimate the mean of the arms with the empirical mean.

Let  $(X_s)_{s\in\mathbb{N}}$  be iid random variables with  $\mathbb{E}[|X_1|] < \infty$  and expectated value  $\mathbb{E}[X_1] = \mu$ . Let  $\bar{X}_t = \sum_{s=1}^t X_s$ .

Theorem 14 (Strong law of large numbers)

$$ar{X}_t \stackrel{a.s.}{\longrightarrow} \mu$$
, that is  $\mathbb{P}(ar{X}_t o \mu) = 1$ .

**Theorem 15 (Central limit theorem)** If  $\mathbb{V}[X] = \sigma^2 < \infty$ , then  $\sqrt{t}(\bar{X}_t - \mu) \xrightarrow{d} \mathcal{N}(\mu, \sigma^2)$ .

#### Problem: those are asymptotic results.

Main question: if I have 15 samples of arm k, how reliable is my estimate for  $\mu^k$  ?

Our main tools are concentration inequalities: bounds on the probability that the empirical mean (or another statistic) is far from its expected value.

## Theorem 16 (Hoeffding's inequality)

If  $X_1, \ldots, X_t$  are independent random variables almost surely in [a, b] then for all  $\delta \in (0, 1)$  we have

$$\mathbb{P}\left(\sum_{s=1}^{t} X_s - \mathbb{E}\left[\sum_{s=1}^{t} X_s\right] \ge (b-a)\sqrt{\frac{t}{2}\log\frac{1}{\delta}}\right) \le \delta.$$

Equivalently, for all  $\varepsilon \geq 0$ ,

$$\mathbb{P}\left(\sum_{s=1}^{t} X_s - \mathbb{E}\left[\sum_{s=1}^{t} X_s\right] \ge \varepsilon\right) \le \exp\left(-\frac{2\varepsilon^2}{t(b-a)^2}\right)$$

# Proof

Proof under a sub-Gaussian assumption. Exercise: bounded implies sub-Gaussian.

Assumption: for all s,  $X_s$  is  $\sigma^2$ -sub-Gaussian, which means that for all  $\lambda \in \mathbb{R}$ ,

 $\mathbb{E}[e^{\lambda(X_s-\mu_s)}] \leq e^{\frac{1}{2}\sigma^2\lambda^2}.$ 

## Proof
In the analysis of bandit algorithms, we want to bound  $\hat{\mu}_t^k - \mu^k$ , where  $\hat{\mu}_t^k = \frac{1}{N_t^k} \sum_{s=1}^t X_s^{k_s} \mathbb{I}\{k_s = k\}.$ 

 $k_s$  is a random variable that depends on all previous rewards.

Issue:  $\hat{\mu}_t^k$  is a sum of a random number of random variables which are not independent.

- $\widehat{\mu}_t^k$  is not unbiased.
- $\widehat{\mu}_t^k$  is not a sum of a fixed number of independent random variables.
- Hoeffding's inequality does not apply.

How to avoid the difficulty: union bounds, or martingale arguments (see proofs later in the course).

Introduction: What is online learning?

Online Linear Optimization

Online Convex Optimization

Adversarial bandits

Stochastic bandits

The stochastic bandit problem

Consequences of the stochasticity

Exploration and exploitation

Optimism in face of uncertainty: UCB

Stochastic Linear Bandits

Goal: minimize  $\mathbb{E}[\operatorname{Reg}_T] = T\mu^* - \sum_{t=1}^T \mu^{k_t}$ .

Since the empirical mean of an arm concentrate around its expected value, can we simply pull the arm with highest empirical mean?

#### **Definition (Follow-The-Leader)**

The FTL algorithm first explores each arm once  $k_t = t$  for  $k \leq K$  and then pulls arm  $k_t = \arg \max_{k \in [K]} \hat{\mu}_{t-1}^k$  for all  $t \geq K + 1$ .

Full information: yes, FTL is optimal.

Bandit: answer is no, FTL does not work. It has linear expected regret in most settings.

### FTL still fails

Need to not only exploit, but also explore.

**Explore-Then-Commit** 

Parameter:  $m \ge 1$ .

- 1. Exploration
  - For rounds  $t = 1, \ldots, mK$  explore by drawing each arm m times.
  - Compute for each arm k its empirical mean of rewards obtained by pulling arm k m times

$$\widehat{\mu}_{mK}^{k} = \frac{1}{m} \sum_{s=1}^{Km} \mathbb{I}\{k_s = k\} X_s^k \,.$$

**2. Exploitation**: keep playing the best arm  $\arg \max_k \hat{\mu}_{mK}^k$  for the remaining rounds  $t = mK + 1, \dots, T$ .

#### Theorem 17 (Thm 6.1, Lattimore and Szepesvári 2019)

If all distributions are bounded in [0,1] and  $1\leqslant m\leqslant T/K$  then ETC has expected regret

$$\mathbb{E}[\operatorname{\mathsf{Reg}}_T] \leqslant m \sum_{k=1}^K \Delta_k + (T - mK) \sum_{k=1}^K \Delta_k \exp\left(-m\Delta_k^2\right).$$

- m too large  $\Rightarrow$  too much exploration, linear regret.
- m too small  $\Rightarrow$  too little exploration, linear regret.
- What *m* should we choose?

### Proof

### Proof

### Finding the right trade-off

Two arms bandit: arm 1 is the best arm, arm 2 has gap  $\Delta.$ 

ETC verifies

$$\mathbb{E}[\operatorname{\mathsf{Reg}}_{\mathcal{T}}] \leqslant m\Delta + (\mathcal{T} - 2m)\Delta e^{-m\Delta^2}.$$

**Theorem 18**  
If 
$$K = 2$$
 and  $m = \max\{1, \left\lceil \frac{\log(T\Delta^2)}{\Delta^2} \right\rceil\}$ , then
$$\mathbb{E}[\operatorname{Reg}_T] \leq \Delta + \frac{1 + \log(T\Delta^2)}{\Delta}.$$

This is a distribution dependent bound, meaning that it depends on the gap.

Issue with those bounds: meaningless if  $\Delta$  is small.

#### Worst case bound

#### ETC verifies

$$\mathbb{E}[\operatorname{\mathsf{Reg}}_{\mathcal{T}}] \leqslant m\Delta + (\mathcal{T} - 2m)\Delta e^{-m\Delta^2}$$
.

**Theorem 19**  
If 
$$K = 2$$
 and  $m = \max\left\{1, \left\lceil \frac{\log(T\Delta^2)}{\Delta^2} \right\rceil\right\}$ , then  
 $\mathbb{E}[\operatorname{Reg}_T] \le \min\left\{\Delta + \frac{1 + \log(T\Delta^2)}{\Delta}, \ T\Delta\right\} \lesssim \sqrt{T\log T}$ .

This is close to optimal: we can prove a lower bound of order  $\sqrt{T}$ .

Problems:

- m depends on  $\Delta$ , which is unknown.
- What can we do for K > 2?

Introduction: What is online learning?

Online Linear Optimization

**Online Convex Optimization** 

Adversarial bandits

#### Stochastic bandits

The stochastic bandit problem

Consequences of the stochasticity

Exploration and exploitation

Optimism in face of uncertainty: UCB

Stochastic Linear Bandits

The homework is available on my webpage:

http://pierre.gaillard.me/teaching.html

It is due by Jan. 10th 2025.

Upload your notebook using the form on my webpage:

http://pierre.gaillard.me/teaching/online\_learning\_uga.php

The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$



The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$





The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$





The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$





The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$





The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$



The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[ L_t^k, U_t^k \right] \, .$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$



The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[L_t^k, U_t^k\right].$$

It is **optimistic**, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$



The UCB (Upper-Confidence-Bound) algorithm is a very popular bandit algorithm that has several advantages over ETC:

- It does not rely on an initial exploration phase but explores on the fly as rewards are observed.
- Unlike ETC, it does not require knowledge of gaps and behaves well when there are more than two arms.

For each arm k, it builds a confidence interval on its expected reward based on past observation

$$I_t^k = \left[L_t^k, U_t^k\right].$$

It is optimistic, acting as if the best possible rewards are the real rewards:

$$k_t \in \underset{k \in \{1,...,K\}}{\operatorname{arg\,max}} U_t^k.$$



## **Confidence intervals**

How to design the upper confidence bounds?

 $\rightarrow$  concentration inequalities. Here Hoeffding's inequality.

#### Theorem 20 (Hoeffding's inequality)

If  $X_1, \ldots, X_t$  are independent random variables almost surely in [a, b] with same mean  $\mu$  then for all  $\delta \in (0, 1)$  we have

$$\mathbb{P}\left(rac{1}{t}\sum_{s=1}^t X_s - \mu \geq \sqrt{rac{(b-a)^2}{2t}\lograc{1}{\delta}}
ight) \leq \delta \ .$$

Careful: UCB is adaptive, hence  $\hat{\mu}_t$  is not exactly a sum of independent random variables. But we will make it work.

For rewards in [0, 1]:  $U_t^k = \widehat{\mu}_{t-1}^k + \sqrt{\frac{2\log t}{N_{t-1}^k}}$ 

**Initialization** For rounds t = 1, ..., K pull arm  $k_t = t$ . **For** t = K + 1, ..., T, choose  $k_t \in \underset{k \in [K]}{\operatorname{arg max}} \left\{ \widehat{\mu}_{t-1}^k + \sqrt{\frac{2 \log t}{N_{t-1}^k}} \right\},$ and get reward  $X_t^{k_t}$ .

### **Regret Bound**

#### Theorem 21

If the distributions  $\nu_k$  have supports all included in [0,1] then for all k such that  $\Delta_k > 0$ 

$$\mathbb{E}\big[N_T^k\big] \leqslant \frac{8\log T}{\Delta_k^2} + 2$$

In particular, this implies that the expected regret of UCB is upper-bounded as

$$\mathbb{E}[\operatorname{\mathsf{Reg}}_{T}] \leqslant 2K + \sum_{k:\Delta_{k}>0} \frac{8\log T}{\Delta_{k}} \,.$$

Remarks :

- we can also prove  $\mathbb{E}[\operatorname{Reg}_T] \lesssim \sqrt{KT \log(T)}$ . Close to the optimal  $O(\sqrt{KT})$ .
- Deals with multiple gaps, without any knowledge of the gaps, unlike ETC.
- Bounded can be replaced by sub-Gaussian.

Idea: if the means belong to the confidence intervals and the arms are pulled enough, the algorithm cannot pull a suboptimal arm.

We prove that if  $k_t = k \neq *$ , then one of these inequalities must be false:

$$\mu^{*} \leq \widehat{\mu}_{t-1}^{*} + \sqrt{\frac{2\log t}{N_{t-1}^{*}}} \qquad \leftarrow \mu^{*} \text{ smaller than UCB} \qquad (i)$$
$$\mu^{k} \geq \widehat{\mu}_{t-1}^{k} - \sqrt{\frac{2\log t}{N_{t-1}^{k}}} \qquad \leftarrow \mu_{k} \text{ larger than LCB} \qquad (ii)$$
$$N_{t-1}^{k} \geq \frac{8\log t}{\Delta_{k}^{2}} \qquad \leftarrow k \text{ played enough} \qquad (iii)$$

$$\mu^* \leq \widehat{\mu}_{t-1}^* + \sqrt{\frac{2\log t}{N_{t-1}^*}} \quad \text{and} \quad \mu^k \geq \widehat{\mu}_{t-1}^k - \sqrt{\frac{2\log t}{N_{t-1}^k}} \quad \text{and} \quad N_{t-1}^k \geq \frac{8\log t}{\Delta_k^2}$$

Prove that if k is pulled at t, then there is a contradiction.

#### **Proof 3: decomposition wrt events**

One of these is false:

$$\mu^* \leq \widehat{\mu}_{t-1}^* + \sqrt{\frac{2\log t}{N_{t-1}^*}} \quad ; \quad \mu^k \geq \widehat{\mu}_{t-1}^k - \sqrt{\frac{2\log t}{N_{t-1}^k}} \quad ; \quad N_{t-1}^k \geq \frac{8\log t}{\Delta_k^2}$$
  
Then:  $\mathbb{E}[N_T^k] \leq u + \sum_{t=u+1}^T \left( \mathbb{P}\{(i) \text{ is false}\} + \mathbb{P}\{(ii) \text{ is false}\} \right) \qquad \text{for } u = \left\lceil \frac{8\log T}{\Delta_k^2} \right\rceil.$ 

## **Proof 4: probability of the concentration event**

We show: 
$$\mathbb{P}(\mu^k < \widehat{\mu}_{t-1}^k - \sqrt{\frac{2\log t}{N_{t-1}^k}}) \le t^{-3}$$
.

For 
$$u = \frac{8\log T}{\Delta_k^2}$$
,  $\mathbb{E}\left[N_T^k\right] \leqslant u + \sum_{t=u+1}^T \left(\mathbb{P}\left\{\mu^* > \widehat{\mu}_{t-1}^* + \sqrt{\frac{2\log t}{N_{t-1}^*}}\right\} + \mathbb{P}\left\{\mu^k < \widehat{\mu}_{t-1}^k - \sqrt{\frac{2\log t}{N_{t-1}^k}}\right\}\right)$ .  
Each of these probabilities is smaller than  $t^{-3}$ .

$$\mathbb{E}[\operatorname{\mathsf{Reg}}_{\mathcal{T}}] \leq \frac{8\log \mathcal{T}}{\Delta_k^2} + 2\sum_{t=u+1}^{\mathcal{T}} \frac{1}{t^3} \leq \frac{8\log \mathcal{T}}{\Delta_k^2} + 2.$$

The bound of the regret then comes from  $\mathbb{E}[\operatorname{Reg}_{\mathcal{T}}] = \sum_{k} \mathbb{E}[N_{\mathcal{T}}^{k}] \Delta_{k}$ .

#### $\varepsilon$ -greedy

First choose a parameter  $\varepsilon \in (0, 1)$ , then at each round, select the arm with the highest empirical mean with probability  $\varepsilon$  (i.e., be greedy), and explore by playing a random arm with probability  $\varepsilon$ .

Works quite well in practice and is used in many application because of its simple implementation (in particular in reinforcement learning).

Choosing  $\varepsilon \approx K/(\Delta^2 T)$  yields to an upper-bound of order  $K \log T/\Delta^2$ . However it requires the knowledge of  $\Delta$ .

#### Thompson Sampling

Thomson sampling was the first algorithm proposed for bandits by Thomson in 1933. It assumes a uniform prior over the expected rewards  $\mu_i \in (0, 1)$ , then at each round  $t \ge 1$ , it

- computes  $\hat{\nu}_{k,t}$  the posterior distribution of the rewards of each arm k given the rewards observed so far;
- samples  $x_{k,t} \sim \widehat{\nu}_{k,t}$  independently;
- selects  $k_t \in \arg \max_{k \in \{1, \dots, K\}} x_{k,t}$ .

Thomson sampling has a similar upper-bound of order  $O(K \log T/\Delta)$  than the one achieved by UCB. Somewhat different proof techniques.

An advantage over UCB is the possibility of incorporating easily prior knowledge on the arms.

UCB proved easier to adapt to structured bandits (it can be hard to sample a posterior conditioned on structural information).

Introduction: What is online learning?

Online Linear Optimization

Online Convex Optimization

Adversarial bandits

#### Stochastic bandits

The stochastic bandit problem Consequences of the stochasticit Exploration and exploitation

Optimism in face of uncertainty: UCB

Stochastic Linear Bandits

#### Main motivation: use contexts.

Unknown parameter:  $\mu^* \in \mathbb{R}^d$ .

At each time step  $t = 1, \ldots, T$ 

- the environment chooses  $\mathcal{K}_t \subseteq \mathbb{R}^d$ , the decision set;
- the player chooses an action  $x_t \in \mathcal{K}_t$ ;
- given  $x_t$ , the environment draws the reward

$$X_t = x_t^\top \mu^* + \varepsilon_t$$

where  $\varepsilon_t$  is i.i.d. 1-subgaussian noise.  $(\forall \lambda > 0, \mathbb{E}[\exp(\lambda \varepsilon_t)] \leq \exp(\lambda^2/2))$ 

- the player only observes the feedback  $X_t$ .

The player wants to minimize its expected regret defined as

$$\mathbb{E}\operatorname{Reg}_{\mathcal{T}} \stackrel{\text{\tiny def}}{=} \mathbb{E}\left[\sum_{t=1}^{\mathcal{T}}\max_{\mathbf{x}\in\mathcal{K}_{t}}\mathbf{x}^{\top}\boldsymbol{\mu}^{*} - \sum_{t=1}^{\mathcal{T}}\mathbf{x}_{t}^{\top}\boldsymbol{\mu}^{*}\right]\,.$$

#### **E**xamples

- Finite-armed bandit: if  $\mathcal{K}_t = (e_1, \dots, e_d)$ , unit vectors in  $\mathbb{R}^d$  and  $\mu^* = (\mu_1, \dots, \mu_d)$ , we recover the setting of finite-armed bandit (with *d* arms).
- <u>Contextual linear bandit</u>: if  $c_t \in \mathcal{X}$  is a context observed by the player and the reward function  $\mu$  is of the form

$$\mu(x,x) = \psi(x,x)^{\top} \mu^*, \qquad \forall (x,x) \in [K] \times \mathcal{X},$$

for some unknown parameter  $\mu^* \in \mathbb{R}^d$  and <u>feature map</u>  $\psi : [K] \times \mathcal{X} \to \mathbb{R}^d$ .

- Combinatorial bandit:  $\mathcal{K}_t \subseteq \{0,1\}^d \to \text{combinatorial bandit problems. Example: decision set = possible paths in a graph, the vector <math>\mu^*$  assigns to each edge a reward corresponding to its cost and the goal is to find the smallest path with smallest cost.

Algorithm LinUCB - UCB for linear bandits.

- Build confidence region for the parameter:  $C_t$  such that  $\mu^* \in C_t$  with high probability.
- Build confidence bounds for the arm means:  $U_t^{\times} = \max_{\mu \in \mathcal{C}_t} x^{\top} \mu$ .
- Be optimistic: pull  $x_t = \arg \max_x U_t^x$ .

Main question: how do we get  $C_t$ ?

#### **Confidence region**

After time *t*, the algorithm observed:

$$X_1 = x_1^\top \mu^* + \varepsilon_1$$
$$X_2 = x_2^\top \mu^* + \varepsilon_2$$
$$\dots$$
$$X_t = x_t^\top \mu^* + \varepsilon_t$$

The unknown parameter we want to estimate is  $\mu^*$ .

Denoting by  $I_d$  the  $d \times d$  identity matrix and picking  $\lambda > 0$ , we can estimate  $\mu^*$  with regularized least square

$$\widehat{\mu}_t \stackrel{\text{def}}{=} \operatorname*{arg\,min}_{\mu \in \mathbb{R}^d} \left\{ \sum_{s=1}^t (X_s - x_s^\top \mu)^2 + \lambda \|\mu\|^2 \right\} = V_t^{-1} \sum_{s=1}^t x_s X_s \,,$$

where  $V_t \stackrel{\text{def}}{=} \lambda I_d + \sum_{s=1}^t x_s x_s^{\top}$ .
### Lemma 3

Let  $\delta \in (0,1)$ . Then, with probability at least  $1 - \delta$ , if  $\max_{x \in \mathcal{K}_t} \|x\|_2 \leqslant 1$ , for all  $t \ge 1$ 

$$\left\|\widehat{\mu}_t - \mu^*\right\|_{V_t} \leqslant \sqrt{\lambda} \|\mu^*\| + \sqrt{2\log(1/\delta) + d\log\left(1 + \frac{T}{\lambda}\right)} \stackrel{\text{\tiny def}}{=} \beta(\delta) \,,$$

where  $\|\mu\|_{V_t}^2 = \mu^\top V_t \mu$ .

Conclusion: with probability  $1 - \delta$ , for all  $t \ge 1$ ,

$$\mu^* \in C_t, \quad \text{where} \quad C_t \stackrel{\text{def}}{=} \left\{ x \in \mathbb{R}^d : \left\| \mu - \widehat{\mu}_{t-1} \right\|_{V_{t-1}} \leqslant \beta(\delta/T) \right\}.$$
(14)

#### Theorem 22

Let  $T \ge 1$  and  $\mu^* \in \mathbb{R}^d$ . Assume that for all  $x \in \bigcup_{t=1}^T \mathcal{K}_t$ ,  $|x^\top \mu^*| \le 1$ ,  $||\mu^*|| \le 1$  and  $||x|| \le 1$ , then LinUCB with  $C_t$  defined as above satisfies the regret bound

 $\mathbb{E}\operatorname{Reg}_{\mathcal{T}} \leqslant \Box_{\lambda} d\sqrt{\mathcal{T}} \log(\mathcal{T}),$ 

where  $\Box_{\lambda}$  is a constant that may depend on  $\lambda$ .

Remark:

-  $O(\sqrt{T})$ : the exponent does not depend on d.

With probability 1 - 1/T, for all  $t \ge 1$ ,  $\mu^* \in C_t$ , where  $C_t \stackrel{\text{def}}{=} \left\{ x \in \mathbb{R}^d : \left\| \mu - \widehat{\mu}_{t-1} \right\|_{V_{t-1}} \le \beta(1/T^2) \right\}$ .

## Proof 2

## Proof 3

LinUCB with  $C_t$  defined as above satisfies the regret bound

 $\mathbb{E}\operatorname{Reg}_{T} \lesssim d\sqrt{T}\log(T)$ ,

To prove it, we assumed the following lemma:

#### Lemma 4

Let  $\delta \in (0,1)$ . Then, with probability at least  $1 - \delta$ , if  $\max_{x \in \mathcal{K}_t} \|x\|_2 \leqslant 1$ , for all  $t \ge 1$ 

$$\left\|\widehat{\mu}_t - \mu^*\right\|_{V_t} \leqslant \sqrt{\lambda} \|\mu^*\| + \sqrt{2\log(1/\delta) + d\log\left(1 + rac{T}{\lambda}
ight)} \stackrel{ ext{def}}{=} eta(\delta)$$

where  $\|\mu\|_{V_t}^2 = \mu^{\top} V_t \mu$ .

Under additional assumptions, it is possible to improve the regret bound  $O(d\sqrt{T} \log T)$ .

- If the set of available actions at time t is fixed and finite; i.e.,  $x_t \in \mathcal{K}$  where  $|\mathcal{K}| = \mathcal{K}$ . Then, it is possible to achieve

$$\mathbb{E}\operatorname{Reg}_{T} \leqslant \Box \sqrt{Td \log(TK)},$$

which improves the previous bound by a factor  $\sqrt{d}/\log(K)$  and improves the classical bound of UCB  $O(\sqrt{TK \log T})$  by a factor  $K/\sqrt{d}$ .

- Another possible improvement when  $d \gg 1$  is to assume that  $\mu^*$  is  $m_0$ -sparse (i.e., most of its components are zero). Then under assumptions, one can get a regret of order  $\tilde{O}(\sqrt{dm_0T})$ .

# Thank you!

- Nicolo Cesa-Bianchi and Gábor Lugosi. <u>Prediction, learning, and games</u>. Cambridge university press, 2006.
- [2] Elad Hazan et al. **"Introduction to online convex optimization".** In: Foundations and Trends® in Optimization 2.3-4 (2016), pp. 157–325.
- [3] Tor Lattimore and Csaba Szepesvári. "Bandit algorithms". In: preprint (2019).
- [4] Shai Shalev-Shwartz et al. **"Online learning and online convex optimization".** In: Foundations and Trends® in Machine Learning 4.2 (2012), pp. 107–194.
- [5] Martin Zinkevich. "Online convex programming and generalized infinitesimal gradient ascent". In: Proceedings of the 20th International Conference on Machine Learning (ICML-03). 2003, pp. 928–936.