Exercise 2: Adversarial sparse losses

The aim of this exercise is to study what happens when both a non-negativity and a sparsity assumptions are made on the vectors of losses picked by the opponent.

More formally, we consider the setting of linear losses, with N components, where at most s components are positive while the other components are null. The parameter $s \in \{1, \ldots, N\}$ is fixed throughout the game but is unknown to the statistician. The online protocol is the following.

Protocol: For all rounds $t = 1, 2, \ldots$,
- The statistician picks a convex combination $(p_{j,t})_{1 \leqslant j \leqslant N}$ while the environment simultaneously picks a loss vector $(\ell_{j,t})_{1 \leqslant j \leqslant N} \in [0, +\infty)^N$, with at most s non-null components;
- The choices are publicly revealed.

The statistician aims to control the regret

$$R_T = \sum_{t=1}^{T} \sum_{j=1}^{N} p_{j,t}\, \ell_{j,t} - \min_{1 \leqslant i \leqslant N} \sum_{t=1}^{T} \ell_{i,t}\,.$$

The question is:
What is the optimal order of magnitude of the regret under the non-negativity and sparsity assumptions?

Lower bound on the regret (only requires Lecture #1)

Consider the joint distribution over $\{0,1\}^N$ defined as the law of a random vector $\boldsymbol{L} = (L_1, \ldots, L_N)$ drawn in two steps. First, we pick s components uniformly at random among $\{1, \ldots, N\}$; we call them $K_1, \ldots, K_s$. Then, the components not picked ($k \neq K_j$ for all j) are associated with zero losses, $L_k = 0$. The losses L_k for picked components $K_1, \ldots, K_s$ are drawn according to a Bernoulli distribution with parameter $1/2$. The loss vector $\boldsymbol{L} \in [0,1]^N$ thus generated is indeed s–sparse and non-negative.

We fix an algorithm for the statistician, consider an i.i.d. sequence $\boldsymbol{L}_1, \boldsymbol{L}_2, \ldots$ of random vectors thus generated, and study the corresponding regret

$$R_T = \sum_{t=1}^{T} \sum_{j=1}^{N} p_{j,t}\, L_{j,t} - \min_{1 \leqslant i \leqslant N} \sum_{t=1}^{T} L_{i,t}\,.$$

1. Show that the expectation of the regret can be written as

$$\mathbb{E}\left[\frac{R_T}{\sqrt{T}}\right] = \mathbb{E}\left[\max_{1 \leqslant i \leqslant N} \frac{1}{\sqrt{T}} \sum_{t=1}^{T} X_t^{(i)}\right]$$

where the $\bigl(X_t^{(1)}, \ldots, X_t^{(N)}\bigr)$ are i.i.d. centered random vectors taking values in $[-1,1]^N$, with covariance matrix denoted by Γ.

please give a closed-form definition of the $X_t^{(i)}$ based on the $L_{i,t}$, and also compute Γ.

2. Explain why

$$\mathbb{E}\left[\max_{1 \leqslant i \leqslant N} \frac{1}{\sqrt{T}} \sum_{t=1}^{T} X_t^{(i)}\right] \longrightarrow \mathbb{E}\Bigl[\max_{1 \leqslant i \leqslant N} Z_i\Bigr]$$

where $(Z_1, \ldots, Z_N)$ follows the normal distribution $\mathcal{N}(\mathbf{0}, \Gamma)$, i.e., the centered normal distribution with covariance matrix Γ.

3. Consider the Gaussian random vector $(W_1, \ldots, W_N)$ with i.i.d. components W_i with distribution $\mathcal{N}\bigl(0, \mathrm{Var}(Z_1)\bigr)$. Show that Slepian's lemma (stated at the bottom of next page) is applicable and that it entails

$$\mathbb{E}\Bigl[\max_{1 \leqslant i \leqslant N} Z_i\Bigr] \geqslant \mathbb{E}\Bigl[\max_{1 \leqslant i \leqslant N} W_i\Bigr]$$

4. Conclude to an asymptotic lower bound of the order of $\sqrt{(T s \ln N)/N}$; state it carefully and rigorously.

Upper bound on the regret (requires Lecture #2)

5. Recall first how, under the non-negativity assumption, i.e., assuming that the losses $\ell_{j,t}$ all lie in $[0, M]$, we could prove the bound

$$R_T \leqslant 13M \ln N + 2\sqrt{M \min_{j=1,\ldots,N} \sum_{t=1}^{T} \ell_{j,t} \ln N}\,,$$

referred to as an "improvement for small cumulative losses."

More precisely, recall the algorithm at hand and the sketch of its performance bound above. (Answer in a about 10–15 lines only.)

6. Deduce a $13M \ln N + 2M\sqrt{(Ts \ln N)/N}$ bound on the regret of this algorithm under the sparsity assumption.

Does the algorithm need to know s to ensure this bound? Explain and comment.

Slepian's lemma (1962): Let $(Z_1, \ldots, Z_N)$ and $(W_1, \ldots, W_N)$ be two centered Gaussian random vectors in $\mathbb{R}^N$. If

$$\forall i \in \{1, \ldots, N\}^2, \qquad \mathbb{E}\big[Z_i^2\big] = \mathbb{E}\big[W_i^2\big]$$

and

$$\forall (i,j) \in \{1, \ldots, N\}^2, \qquad i \neq j \quad \Rightarrow \quad \mathbb{E}\big[Z_i Z_j\big] \leqslant \mathbb{E}\big[W_i W_j\big]\,,$$

then for all $t \in \mathbb{R}$,

$$\mathbb{P}\Big\{\max_{1 \leqslant i \leqslant N} Z_i > t\Big\} \geqslant \mathbb{P}\Big\{\max_{1 \leqslant i \leqslant N} W_i > t\Big\}\,.$$