Sequential Learning: Homework #1

What I care about. I care about well-written proofs: with sufficient details, with calculations worked out and leading to pleasant and readable bounds. I favor quality of the writing over the quantity of questions answered. I give bonus points for elegant solutions.

Formats of your submission, deadline. Please send your solutions in a sequential manner, one exercise after the other. Wait for my OK to send a new solution, as I may request you to re-work a solution badly written. I may take 1 or 2 business days to get back to you, please take this into account when trying to abide by the deadline.

I expect to receive PDF files, with answers either handwritten and neatly scanned (as I do for my weekly lecture notes) or typed in LaTeX.

The PDF file must be named `YourName-HW1-ExN.pdf`, where `YourName` is to be replaced by your family name, and `N` by the exercise number. E.g., my submission for the first exercise would be named `Stoltz-HW1-Ex1.pdf`.

Deadline is Friday, March 4, at 6pm. This is a strict deadline: submitting after this deadline will negatively impact your grade, with the impact depending on the delay.

Beware: Typos. Most likely the statement comes with typos. This is part of the job. Try to correct them on your own!

Exercise 1: The polynomially weighted average forecaster

We consider the "vanilla" setting of linear losses, with $N \geqslant 2$ components: 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 $(\ell_{j,t})_{1 \leqslant j \leqslant N}$;
- The choices are publicly revealed.

The statistician aims to control the regret

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

We will actually denote by

$$R_{i,T} = \sum_{t=1}^{T} \sum_{j=1}^{N} p_{j,t}\, \ell_{j,t} - \sum_{t=1}^{T} \ell_{i,t}$$

the regret associated with the component $i \in \{1, \ldots, N\}$. We also denote by $u_+ = \max\{u, 0\}$ the non-negative part of a real number u, and write $\boldsymbol{u}_+$ the vector based on $\boldsymbol{u} = (u_1, \ldots, u_N) \in \mathbb{R}^N$ with components $(u_j)_+$.

Strategy: The statistician considers the following strategy, with hyperparameter $p \geqslant 2$: for $t \geqslant 1$,

$$p_{j,t} = \frac{(R_{j,t-1})_+^{p-1}}{\sum_{k=1}^{N} (R_{k,t-1})_+^{p-1}} \qquad \text{if} \quad \sum_{k=1}^{N} (R_{k,t-1})_+^{p-1} > 0$$

and $p_{j,t} = 1/N$ otherwise (this is in particular the case when $t = 1$).

Analysis in the case $p = 2$ (only requires Lecture #1)

We consider the special case $p = 2$ to have a smooth start. We introduce the instantaneous regret vectors: for all $t \geqslant 1$,

$$\boldsymbol{r}_t = (r_{i,t})_{1 \leqslant i \leqslant N} = \left(\sum_{j=1}^{N} p_{j,t}\, \ell_{j,t} - \ell_{i,t} \right)_{1 \leqslant i \leqslant N}$$

We then define the cumulative regret vector $\boldsymbol{R}_T = \boldsymbol{r}_1 + \ldots + \boldsymbol{r}_T$.

1. Explain why $(u + v)_+ \leqslant \bigl| u_+ + v \bigr|$ for all real numbers $(u, v) \in \mathbb{R}^2$ and why we therefore have

$$\bigl\| (\boldsymbol{R}_t)_+ \bigr\| \leqslant \bigl\| (\boldsymbol{R}_{t-1})_+ + \boldsymbol{r}_t \bigr\|$$

2. Show that

$$\bigl\| (\boldsymbol{R}_{t-1})_+ + \boldsymbol{r}_t \bigr\|^2 = \bigl\| (\boldsymbol{R}_{t-1})_+ \bigr\|^2 + \|\boldsymbol{r}_t\|^2$$

3. Provide a regret bound for the algorithm considered, say, for losses $\ell_{j,t}$ all lying in some $[m, M]$ range; provide a closed-form regret bound only depending on m, M, T and N.

4. Does the algorithm need to know m, M and T? Are the dependencies in T and N optimal?

Analysis for $p > 2$ (optional; only requires Lecture #1)

This part of the exercise is optional. If you send a solution, it must be a complete solution covering all questions 5–8 (not just questions 7–8).

The general analysis of this strategy relies on a function Φ defined as: for all $\boldsymbol{u} = (u_1, \ldots, u_N) \in \mathbb{R}^N$,

$$\Phi(\boldsymbol{u}) = \left(\sum_{i=1}^{N} (u_i^+)^p \right)^{2/p}$$

5. Show that for all $t \geqslant 2$, there exists $\xi_t \in \mathbb{R}^N$ such that

$$\Phi(\boldsymbol{R}_t) \leqslant \Phi(\boldsymbol{R}_{t-1}) + \frac{1}{2} \sum_{i,j=1}^{N} \partial^2_{ij} \Phi(\xi_t)\, r_{i,t}\, r_{j,t}$$

You may use that Φ is C^2–regular on a subset of $\mathbb{R}^N$ to determine. Carefully explain how you handle the cases where this C^2–regularity cannot be directly exploited.

6. Prove the bound

$$\sum_{i,j=1}^{N} \partial^2_{ij} \Phi(\xi_t)\, r_{i,t}\, r_{j,t} \leqslant 2(p-1) \|\boldsymbol{r}_t\|_p^2$$

You may do so by using that $\psi(x) = x^{2/p}$ is concave (thus $\psi'' \leqslant 0$) and by introducing $f(x) = x_+^p$ for the sake of more concise and more abstract calculations; Hölder's inequality may be useful as well.

7. Conclude to a $(M-m)\sqrt{(p-1)N^{2/p}T}$ regret bound.

8. Propose a good value of p (simple enough and readable, not necessarily some optimal value) so that the upper bound obtained is optimal as far as its dependencies in T and N are concerned.