Exercise 5: Online linear regression based on picking convex weights

We consider the problem of online linear regression, with the following learning protocol: for all rounds $t = 1, 2, \ldots$,

- features $\boldsymbol{x}_t = (x_{1,t}, \ldots, x_{N,t}) \in \mathbb{R}^N$ are chosen by Nature and observed by the statistician;
- the statistician and Nature simultaneously pick linear weights $\boldsymbol{u}_t = (u_{1,t}, \ldots, u_{N,t}) \in \mathbb{R}^N$ and an outcome $y_t \in \mathbb{R}$, respectively;
- both players observe $\boldsymbol{u}_t$ and y_t, and the statistician suffers the prediction error

$$\left(y_t - \sum_{j=1}^{N} u_{j,t} x_{j,t} \right)^2 .$$

The statistician wants to control the regret

$$R_T\big(\mathcal{U}_U\big) = \left(y_t - \sum_{j=1}^{N} u_{j,t} x_{j,t} \right)^2 - \min_{\boldsymbol{v} \in \mathcal{U}_U} \left(y_t - \sum_{j=1}^{N} v_j x_{j,t} \right)^2$$

against weight vectors $\boldsymbol{v}$ in the $\mathbb{L}^1$–ball with radius $U > 0$:

$$\mathcal{U}_U = \big\{ \boldsymbol{v} \in \mathbb{R}^N : \|\boldsymbol{v}\|_1 \leqslant U \big\}.$$

A simple strategy based on convex weights (requires Lecture #3)

With each $\boldsymbol{x} = (x_1, \ldots, x_N) \in \mathbb{R}^N$, we associate $f(\boldsymbol{x}) \in \mathbb{R}^{2N}$ defined as

$$f(\boldsymbol{x}) = \big(Ux_1, \ldots, Ux_N, -Ux_1, \ldots, -Ux_N \big) .$$

1. Show that each linear weight vector $\boldsymbol{u} \in \mathcal{U}_U$ may be associated with a convex weight vector $g(\boldsymbol{u})$ of size $2N$ such that

   $$\forall \boldsymbol{x} \in \mathbb{R}^N, \qquad \boldsymbol{u} \cdot \boldsymbol{x} = g(\boldsymbol{u}) \cdot f(\boldsymbol{x}) ,$$

   where we use the short-hand notation $\cdot$ to denote inner products; e.g.,

   $$\boldsymbol{u} \cdot \boldsymbol{x} = \sum_{j=1}^{N} u_j x_j .$$

2. Construct a strategy (based on a black-box application of a suitable sub-strategy) with sublinear regret $R_T\big(\mathcal{U}_U\big)$. State its regret bound.

A note: other strategies

There exist other, direct, strategies (e.g., ridge regression) that lead to regret bounds in the setting of online linear regression. Some such strategies even have uniform regret guarantees against all vectors $\boldsymbol{v} \in \mathbb{R}^N$. But it would be too long to study them... so, let's end this first homework here!