Thomas Kesselheim
Rebecca Lehming

January 15, 2025

Algorithms and Uncertainty

Winter Term 2024/25

Tutorial Session - Week 13

Exercise 1:
Show that the Entropical regularizer is $\frac{1}{\eta}$-strongly convex with respect to the ℓ_1-norm in the experts setting.

Hint: Pinsker's inequality can be useful here. In particular, you might want to use it in the following special form: For $\mathbf{p}, \mathbf{q} \in [0,1]^d$ with $\sum_{i=1}^d p_i = \sum_{i=1}^d q_i = 1$ we have $\sum_{i=1}^d p_i \ln\left(\frac{p_i}{q_i}\right) \geq \frac{1}{2} \|\mathbf{p} - \mathbf{q}\|_1^2$.

Exercise 2:
We consider Online Quadratic Optimization. That is, in each round, the function f_t has the form $f_t(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}_t\|_2^2$ for some vector $\mathbf{y}_t$. Show that the regret of Follow-the-Leader is bounded by $4 \cdot (\max_t \|\mathbf{y}_t\|)^2 \cdot (\log T + 1)$.

Hint: Use Lemma 23.3 to bound the regret. In addition, note that for the k-th harmonic number H_k, we have $H_k \leq \log k + 1$.