Thomas Kesselheim

December 11, 2018
Due: lecture on December 18, 2018

Algorithms and Uncertainty

Winter Semester 2018/19

Exercise Set 7

Exercise 1: (4 Points)
Prove Theorem 15.1: Show that the expected regret of the basic explore-exploit algorithm is at most $O(nT^{\frac{2}{3}}\sqrt{\ln(nT)})$.

Hint: Use Hoeffding's inequality and the union bound to upper-bound the probability that there is an i for which $\hat{\mu}_i > \mu_i + \epsilon$ for a suitably chosen ϵ.

Exercise 2: (4 Points)
We consider a generalization of the algorithm *Weighted Majority* for classifiers with k different classes. (The case covered in the lecture, binary classification, is $k = 2$.) In each step, the algorithm chooses a class, which is recommended by the largest number of classifiers (so the class has a plurality).
Show that this algorithm makes at most $(2 + \eta) \min m_i + 2 \ln n/\eta$ errors, where m_i is the number of errors of classifier i.

Exercise 3: (4 Points)
Consider the modified update rule for Multiplicative Weights that sets $w_i^{(t+1)} = w_i^{(t)} \cdot (1 - \ell_i^{(t)} \eta)$. Show that Theorem 16.3 still holds.

Exercise 4: (4 Points)
Show that every no-regret algorithm in the experts setting has to be randomized. Consider the case $n = 2$ and for every deterministic algorithm construct a sequence such that $L_{\text{Alg}}^{(T)} = T$ and $\min_i L_i^{(T)} \leq \frac{T}{2}$.

Exercise 5: (4 Points)
Use Yao's principle to show that for every (randomized) algorithm there is a sequence $\ell^{(1)}, \ldots, \ell^{(T)}$ such that $L_{\text{Alg}}^{(T)} \geq \left(1 - \frac{1}{n}\right) T$ but $\sum_{t=1}^{T} \min_i \ell_i^{(t)} = 0$. This shows that the order of sum and maximum in the regret definition are important.

Bonus:
Use the one-sided version of Hoeffding's inequality to show a regret bound for UCB1 of $\sum_{i \neq i^*} \frac{2 \ln T}{\Delta_i} + 2\Delta_i$. The one-sided version of Hoeffding's inequality is as follows: Let $Z_1, \ldots, Z_N$ be independent random variables such that $a_i \leq Z_i \leq b_i$ with probability 1. Let $\bar{Z} = \frac{1}{N} \sum_{i=1}^{N} Z_i$ be their average. Then for all $\gamma \geq 0$

$$\Pr\left[\bar{Z} - \mathbf{E}[\bar{Z}] \geq \gamma\right] \leq \exp\left(-\frac{2N^2\gamma^2}{\sum_{i=1}^{N}(b_i - a_i)^2}\right) .$$