Optimization and Uncertainty

Summer term 2021

Prof. Dr. Martin Hoefer
Tim Koglin, Lisa Wilhelmi

GOETHE UNIVERSITÄT FRANKFURT AM MAIN
Institute of Computer Science
Algorithms and Complexity

Assignment 9

Issued: 29.06.2021
Due: 06.07.2021, **10:00h**

Exercise 9.1 *Confidence bound* (2 points)

After 9 iterations of the UCB1 algorithm applied on a 4-armed bandit problem with $T = 12$ assume $P_1^{(9)} = 2, P_2^{(9)} = 4, P_3^{(9)} = 2, P_4^{(9)} = 1$ and $S_1 = 1.10, S_2 = 2.52, S_3 = 1.75, S_4 = 0.40$.

Compute which arm will be played in the next round.

Exercise 9.2 *Upper bound for UCB1* (4 points)

When the distributions of the n arms have similar expected rewards, the upper bound for UCB1 shown in lecture is very poor. Show that an upper bound on the expected regret, which is independent of the expected rewards, is given by $5\sqrt{n\,T\ln T} + 4n$, where T is the total number of rounds.
Hint: Divide the arms in two types according to whether their expected reward deviates more or less than $\varepsilon = \sqrt{\frac{n \ln T}{T}}$ *from the expected reward of the best arm* μ_{i^*}.

Exercise 9.3 ExpertClassification *with k classes* (4 points)

Consider a generalization of Weighted Majority for ExpertClassification (Algorithm 16 in the notes) to n classifiers with $k \in \mathbb{N}$ classes (the case covered in the lecture, binary classification, is $k = 2$) and T rounds in total: In each step, the algorithm chooses the class which is recommended by the largest number of classifiers.

Show that the number of mistakes made by the generalized algorithm is at most $(2+2\eta) \cdot \min_i M_i^T + 2\ln(n)/\eta$, where M_i^T is the total number of mistakes by classifier i and $\eta \in (0, 1/2]$ is the learning rate.

Exercise 9.4 *No-regret property for* Experts (2 points)

Show that every no-regret algorithm in the Experts problem needs to be randomized.
Hint: Consider the case $n = 2$ *and costs* $\ell_i^{(t)} = 1$ *whenever classifier* i *makes a mistake in round* t *and* $\ell_i^{(t)} = 0$ *otherwise. For every deterministic algorithm, construct a sequence of* T *rounds such that* $L_A^{(T)} = T$ *and* $\min_i L_i^{(T)} \le T/2$.

Exercise 9.5 *Adversary models in RWM* (2 + 2 + 2 points)

For the analysis of the Randomized Weighted Majority algorithm for Experts (Algorithm 17), an adversary was considered that generated the cost vector $\ell^{(t)} = (\ell_i^{(t)})_{i \in [n]}$ of the n experts in any round $t = 1, ..., T$. This implies that the analyzed algorithm meets the no-regret property, even if the costs are generated in a different (non adversarial) way. However, there are different variants for the adversarial model regarding the knowledge and the power of the adversary. Consider the following three cases.

a) **Oblivious Adversary:** All cost vectors of the experts, $\ell^{(1)}, \ell^{(2)}, \ldots, \ell^{(T)}$, are generated and fixed before round 1 and the first decision of the algorithm. Vector $\ell^{(t)}$ is only presented to the algorithm in round t.

b) **Adaptive Online Adversary:** In every round t, the adversary knows the probability distribution of the algorithm for choosing an expert. The choice of $\ell^{(t)}$ is based on this knowledge.

c) **Adaptive Offline Adversary:** In every round t, the adversary knows the expert that is chosen (after a random draw according to the probability distribution) by the algorithm. Based on that, the adversary chooses a cost vector $\ell^{(t)}$ in round t.

For all three models, argue whether or not there exists an algorithm with no-regret guarantee for the Experts problem.