Optimization and Uncertainty

Summer term 2023

Prof. Dr. Martin Hoefer
Conrad Schecker, Lisa Wilhelmi

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

Assignment 10

Issued: 04.07.2023
Due: 11.07.2023, **10:00h**

Exercise 10.1 *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 should be played in the next round.

Exercise 10.2 *Regret of Explore-and-Exploit* (4 points)

Give a **formal and detailed** proof of a variant of Theorem 31:

For $k = \left(\frac{T}{n}\right)^{2/3} \cdot (\ln T)^{1/3}$*, the (expected) regret of the simple Explore-and-Exploit algorithm is at most* $\mathcal{O}\left(n^{1/3} \cdot T^{2/3} \cdot (\ln T)^{1/3}\right)$.

Hint: Use $\delta = \sqrt{\frac{2 \ln T}{k}}$.

Exercise 10.3 *ExpertClassification with k classes* (5 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 10.4 *No-regret property for Experts* (3 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* $L_{Alg}^{(T)} = T$ *and* $\min_i L_i^{(T)} \leq T/2$.

Exercise 10.5 *Adversary models in RWM* (2 + 2 + 2 = 6 points)

For the analysis of the Randomized Weighted Majority algorithm for EXPERTS, an adversary was considered that generated the cost $\ell^{(t)} := \ell^{(t)}_{\mathrm{RWM}}$ 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, $\ell^1, \ell^2, \dots, \ell^T$, are generated and fixed before round 1 and the first decision of the algorithm. Vector ℓ^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 ℓ^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 ℓ^t in round t.

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

The assignments and further information on the course are provided on our website:
https://algo.cs.uni-frankfurt.de/lehre/oau/sommer23/oau23.shtml

Contacts: {schecker,wilhelmi}@em.uni-frankfurt.de.