

Several structured thresholding bandit problem

Alexandra Carpentier, Institut für Mathematische Stochastik (Allemagne)

In this talk we will discuss the thresholding bandit problem, i.e. a sequential learning setting where the learner samples sequentially K unknown distributions for T times, and aims at outputting at the end the set of distributions whose means μ_k are above a threshold τ . We will study this problem under four structural assumptions, i.e. shape constraints: that the sequence of means is monotone, unimodal, concave, or unstructured (vanilla case). We will provide in each case minimax results on the performance of any strategies, as well as matching algorithms. This will highlight the fact that even more than in batch learning, structural assumptions have a huge impact in sequential learning.

This work is based on a joint work with James Cheshire, Pierre Menard, Andrea Locatelli and Maurilio Gutzeit.