

Contextual Stochastic Bandits with Budget Constraints and Fairness Application

mercredi 3 avril 2024 15:20 (50 minutes)

We review the setting and fundamental results of contextual stochastic bandits, where at each round some vector-valued context x_t is observed and K actions are available, each action providing a stochastic reward with expectation given by some (partially unknown) function of x_t and a . The aim is to maximize the cumulative rewards obtained, or equivalently, to minimize the regret. This requires maintaining a good balance between the estimation (a.k.a., exploration) of the function and the exploitation of the estimates built. The literature also considers additional budget constraints (leading to so-called contextual bandits with knapsacks): actions now provide rewards but also costs. The literature also illustrated that costs may model fairness constraints. We will review these two lines of work and briefly describe our own contribution in this respect, related to a more direct strategy, able to handle \sqrt{T} cost constraints over T rounds, which is exactly what is needed for fairness applications. The recent results discussed at the end of the talk will be based on the joint work by Evgenii Chzhen, Christophe Giraud, Zhen Li, and Gilles Stoltz, Small total-cost constraints in contextual bandits with knapsacks, with application to fairness, Neurips, 2023.

Orateur: STOLTZ, Gilles (CNRS, LMO, Univ. Paris-Saclay)