



Workshop 2019
September 17-18, 2019, Amsterdam

Rémi Degenne (CWI)
6PAC Associate team

Title: « A game perspective on bandit exploration »

Abstract:

Bandit exploration is the process of discovering information about a problem by sampling distributions. Given several probability distributions, an algorithm seeks the answer to a query (e.g. what distribution has the highest mean or whether one of them has mean above a threshold). To obtain that information it sequentially requests samples, then stops and returns an answer based on those observations.

We are interested in algorithms which return the correct answer with high probability while using the lowest amount of samples possible. We present a new family of algorithms for bandit exploration which sees it as a game between the algorithm requesting samples and an adversary trying to make the problem as difficult as possible by altering the distributions. For a wide class of bandit settings, these algorithms asymptotically use the lowest number of samples possible, while also having finite time guarantees and being computationally cheaper than other asymptotically optimal alternatives.

We finally discuss how that game point of view extends beyond the exploration setting, e.g. to regret minimization.