



Department of Decision Sciences

Statistics Seminar

Uniformly Bounded Regret in the Multi-Secretary Problem

Alessandro Arlotto

Duke University

Thursday, 25th May 2017

12:30pm Room 3-E4-SR03 Via Roentgen 1 Milano

Abstract

In the well-known secretary problem, n positive numbers are sequentially presented to a decision maker who decides when to stop and select the most recent number. The goal is to maximize the probability of choosing the largest of the n numbers. In the k -choice (multi-secretary) variant, the decision maker is allowed to select k numbers and the goal is to maximize the aggregate value of the chosen elements. Assuming that numbers are i.i.d and drawn from a known distribution of finite support, we prove that the best regret---the gap between the optimal online policy and the offline upper bound---where all n numbers are made visible at time 0---is uniformly bounded in the horizon length n and the budget k . Our proof is constructive: we develop an adaptive Budget-Ratio(BR) policy that achieves this performance. The policy selects or skips values depending on where the budget ratio---the residual budget divided by the remaining horizon---stands relative to multiple thresholds that are simply computed from the distribution. Being adaptive is crucial. We also prove that the minimal regret among non-adaptive policies grows like the square root of n : the difference is the value of adaptiveness. (Joint work with Itai Gurvich.)