

Structured Learning in Sequential Selection Problems

In this talk I will describe two vignettes of sequential decision making problems arising in the OR literature. The first is the classical “house selling” problem (or optimal stopping): given a random sequence of independent observations revealed one at a time over some finite horizon of play, the objective is to design an algorithm that “stops” this sequence to maximize the expected value of the “stopped” observation. The second vignette is the worker assignment problem (or sequential stochastic assignment): arriving items (“jobs”) with independent stochastic attributes need to be sequentially matched to a pool of awaiting recipients (“workers”). Once each job/worker are matched they are no longer admissible for further assignment, and the objective is to maximize the expected cumulative value of the resulting matchings. Both problems date back more than 50 years and can be solved in a relatively straightforward manner using dynamic programming; both have been used as modeling constructs in numerous applications including recent online marketplaces. Somewhat surprisingly, neither has been studied extensively when the underlying information on the stochastic primitives is unknown a priori. While it is possible to analyze this incomplete information setting with generic multi-purpose learning theoretic methods, these tend to be quite inefficient when further “structure” is present in the problem and can be exploited to construct more customized algorithms. The main focus of this talk is to broadly describe possible learning theoretic formulations of the two vignettes, and illustrate some “design principles” for constructing near-optimal policies.