Optimization and Uncertainty

Summer term 2023

Prof. Dr. Martin Hoefer
Conrad Schecker, Lisa Wilhelmi

GOETHE UNIVERSITÄT FRANKFURT AM MAIN

Institute of Computer Science
Algorithms and Complexity

Assignment 3

Issued: 02.05.2023
Due: 09.05.2023, **10:00h**

Exercise 3.1 Prophet *with multiple items* (3 + 3 + 1 = 7 points)

For the Prophet problem, let p_i be the probability that v_i is optimal, where $i = 1, ..., n$. Furthermore, suppose that τ_i is such that $\Pr[v_i \geq \tau_i] = p_i$, i.e., the p_i^{th} percentile for v_i. For simplicity, assume that such a threshold τ_i always exists. Define

$$\tilde{v}_i(p_i) := \mathbb{E}[v_i \mid v_i \geq \tau_i]$$

as the expected value of v_i given that it lies in the top p_i^{th} percentile. We consider the following strategy to pick one item: *When item i shows up, if no item has been chosen among* $1, ..., i-1$, *reject it with probability 1/2 outright, else accept it if* $v_i \geq \tau_i$.

a) Show that the algorithm achieves a value of at least $\frac{1}{4} \cdot \mathbb{E}[v_{\max}]$.

b) Suppose we are now allowed to choose up to k out of the n items. Thus, the sequence ends with acceptance of the k-th item or after the arrival of the n-th item. The quantity of interest is given by $\mathbb{E}[\text{sum of values of best } k \text{ items}]$. Show that the algorithm achieves at least $\frac{1}{4}$ of it. *Hint: Redefine p_i as the probability that v_i is among the top k values.*

c) If the distributions are finite, then a threshold τ_i as defined above might not exist. Can you adjust the algorithm to handle this case? Explain your answer.

Exercise 3.2 *Two-round IID Prophets* (3 + 4 = 7 points)

Consider two rounds of an IID Prophet problem. The finite distribution $\mathcal{D}$ has support $\{0, 1, n\}$, $n > 1$. The probabilities are $\Pr[v_1 = 1] = q$, $\Pr[v_1 = n] = (1-q)/n$, and $\Pr[v_1 = 0] = (1-1/n)(1-q)$. The same holds for v_2.

a) Let ALG^* be the optimal online algorithm for this instance. Determine $\mathbb{E}[v(ALG^*)]$ as a function of q and n.

b) Derive a lower bound on the competitive ratio as a function of q and n. What are (asymptotically) optimal choices for q and n when we want to maximize this lower bound? What is the resulting bound? Prove your answer.

Exercise 3.3 *Markov Decision Processes* (2 + 4 = 6 points)

Consider a stochastic decision problem similar to the one with the envelopes solved in the lecture. There are n boxes where box i contains a prize of 1 Euro with probability q_i and is empty otherwise. The game ends when a non-empty box is opened. At each point in time, one can also decide to stop playing. This means, the final prize is either 0 Euros or 1 Euro. Opening box i costs c_i Euros. The goal is to maximize the final prize minus the costs for opening.

a) Model this problem as a Markov decision process. In particular, give the state and action sets as well as transition probabilities and rewards.

b) Derive an optimal policy for this problem.

Contacts: {schecker,wilhelmi}@em.uni-frankfurt.de.