Title: Secretary Problems and Online Auctions
Name: MohammadHossein Bateni
Affil./Addr. Google Inc., New York, NY, USA
Keywords: Secretary problem; Mechanism design; Online algorithm; Competitive analysis; Strategyproof; Truthfulness; Submodularity; Matroid constraint; Knapsack constraint
SumOriWork: 2004; Hajiaghayi, Kleinberg, Parkes
2008; Babaioff, Immorlica, Kempe, Kleinberg
2013; Bateni, Hajiaghayi, Zadimoghaddam

Secretary Problems and Online Auctions

MohammadHossein Bateni
Google Inc., New York, NY, USA

Years aud Authors of Summarized Original Work

2004; Hajiaghayi, Kleinberg, Parkes
2008; Babaioff, Immorlica, Kempe, Kleinberg
2013; Bateni, Hajiaghayi, Zadimoghaddam

Keywords

Secretary problem; Mechanism design; Online algorithm; Competitive analysis; Strategyproof; Truthfulness; Submodularity; Matroid constraint; Knapsack constraint

Problem Definition

The classic secretary problem, a prime example of stopping theory, has been studied extensively in the computer science literature. Consider the scenario where an employer is interested in hiring one secretary out of a pool of candidates. The difficulty is that, although the employer does not know the utility of a candidate before she is interviewed, the irrevocable hiring decision for each candidate has to be made right after the interview and prior to interviewing the subsequent candidates. The goal is nonetheless to pick the best candidate, or maximize the probability of achieving this.

Optimization angle

The above scenario is hopeless from an algorithmic point of view since an adversarial input makes it impossible to hire the best candidate. We can take either of two paths to make the problem tractable: restrict the set of utilities or the arrival order of candidates. The former path yields, for instance, the stochastic variant of the problem.

However, we follow the second idea here that leads to the classic secretary problem. The extra assumption, then, is that the candidates arrive in a random error; i.e., although each candidate may have an arbitrary adversarial utility, every permutation of the candidates is equally likely to be the arrival order.

A folklore solution to the problem, often attributed to [3], is to look into the first $\frac{1}{e}$ fraction of the candidates (called the "tuning set") without giving them any offers, and then hire the first candidate with utility more than every one in the tuning set. It is not difficult to show that this approach hires the best candidate with probability at least $\frac{1}{e}$. Indeed, it is known that this is the best possible performance.

There are two questions to be answered, once we extend the problem to multiple secretaries.

1. *What subsets of secretaries can be hired together?* The simplest answer is to allow at most k secretaries to be hired. Alternately, we can place (several) knapsack and/or matroid constraints on the feasible set. The former assigns a cost to each hire—say, the requested salary—that is to paid out of a given budget. The latter permits only those combinations that form an independent set according to a given matroid. It is easy to see that both generalize the cardinality constraint.
2. *How do we compute the utility of a set?* The utility of a set can be defined as the sum of the utilities of individual secretaries in the set. More generally, a submodular or subadditive function may be employed to describe the utility of a set.

We then attempt to hire a feasible set of secretaries of maximum expected utility.

Mechanism design angle

Mechanism design literature has looked at this problem from a slightly different angle. In this setting, the players that arrive in a random order declare a *bid*—i.e., how much they value the item being sold—and then the seller decides who should get the item (or items) and how much they should be charged. Such decisions are to be taken irrevocably as in the optimization problem discussed above.

The players can play strategically, though, by declaring higher or lower bids in order to increase their chances of winning the item or to reduce the price they pay. In addition, they may declare their arrival/departure time untruthfully to achieve a better result. We want to design a "truthful" auction that precludes such undesirable outcomes. Although we allow the player to declare any nonnegative bid (if it is in her favor), we do not let them state an arrival time that is earlier than their actual one. (Presence intervals may be overlapping and/or nested.) We say that a mechanism is value-strategyproof if no player can benefit from declaring a bid different from her real value. Similarly the mechanism is called time-strategyproof if there is no benefit in stating the arrival/departure times untruthfully. We look for mechanisms that are both time- and value-strategyproof.

Key Results

Optimization

Kleinberg [6] studies the multiple-choice generalization where the goal is to hire k candidates, whose total utility (defined as the sum of the individual utilities) is maximized. He presents a tight performance guarantee of $1 + \Theta(\frac{1}{\sqrt{k}})$ for the problem. In the case of $k = 1$, this is equivalent to the classic secretary problem. (The nontrivial direction follows from a construction where the utilities are hugely different.) Kleinberg's algorithm partitions the set of candidates into two (almost) equal pieces, recursively hires $\frac{k}{2}$ secretaries in the first, sets the threshold for the second piece by looking at the solution to the first piece, and picks as many as $\frac{k}{2}$ secretaries in the second piece who are better than threshold.

Babaioff et al. [1] look at the generalization where there is a restriction on the set of candidates that can be hired together; the restriction is in the form of a matroid. They present an $O(\log n)$ competitive ratio in this case along with improved bounds when the matroid has a special form. Their general matroid algorithm partitions the items into logarithmically many sets of almost equal utility, and focuses (randomly) on one such set, which reduces the problem into that of maximizing the cardinality of the solution (solved via the greedy method).

The case of submodular utilities is discussed in Bateni et al. [2]: several matroid or knapsack constraints can be placed on the set of feasible candidates, and the total utility of a set is computed by a submodular function of the participating candidates. They provide constant competitive ratios as long as a fixed number of knapsack constraints are present. When (a constant number of) matroid constraints are involved, too, their performance guarantees grow to $O(\log^2 k)$ where k is the rank of the matroid. They divide the input into different pieces where at most one secretary should be picked from each, not losing too much utility in the process. As a result, the submodular function collapses to an additive one within each piece (by taking the marginal values of secretaries with respect to the current solution). The classic algorithm is then used inside each piece. The main idea behind the matroid algorithm is that we only need to show that, whatever choices we have already committed to, there are enough options left that can appropriately augment the current solution. The argument goes by proving the existence of a magical solution with k' secretaries any of whose $\frac{k'}{2}$-size subsets has significant contribution (say, at least a $\frac{1}{\log k}$ fraction of the optimum) in the submodular function. Had we known k', a simple greedy algorithm would have sufficed to find a solution similar to the magical set. At the cost of another factor $O(\log k)$, we can guess k'.

Furthermore, Bateni et al. show that subadditive utility functions make the problem much more difficult. In particular, they provide matching $\Theta(\sqrt{k})$ competitive ratios.

Mechanism Design

The Dynkin's algorithm for the classic secretary problem can be readily turned into an auction: Set the price after observing the tuning set, and then sell to anyone with a higher bid. This mechanism is not truthful, though, since high-bid players spanning across the time threshold have an incentive to declare later arrival time (i.e., after the threshold); this way, they will win the item but do not set the price.

Nevertheless, Hajiaghayi et al. [5] show how one can modify the mechanism slightly to make it truthful: after the threshold, consider the option of selling the item to the agent with the highest bid so far—if she is still present—and charge her the second-highest bid so far. Their method achieves constant competitiveness for both efficiency and revenue. Their $1/e$ competitiveness for efficiency is best possible since it generalizes the the optimization problem; however, when comparing the revenue to that achieved by the Vickrey auction, their upper bound of $1/e^2$ for for competitiveness fares against a lower bound of $1/e$. (It is possible, they show, to modify the mechanism slightly to trade efficiency loss for revenue gain; for instance, simultaneous 4 competitiveness for both objectives is possible.)

The general idea for the transformation is to define a "tuning period" where the price is set for everyone. Then, not only a simple auction-like mechanism is employed in the "hiring phase" to obtain a strategyproof mechanism, but also extra care should be given to the "transition phase" (from tuning to hiring) so as not to incentivize untruthful declaration of arrival time for those whose presence spans the transition. The same approach can be applied to the multiple-choice secretary problem to obtain constant-factor comeptitve mechanisms (for efficiency and revenue), but this bound is far from the one achieved in the optimization setting by Kleinberg [6].

Open Problems

Though there has been some improvements on the matroid case, we still do not know which cases are hard and admit no constant-factor competitive ratio. For submodular utilities (and simple cardinality constraints), in particular, there is a gap between $(1 - \frac{1}{e})/(e + 1)$ algorithmic result [4] and the $1 - \frac{1}{e}$ (or $1 - \frac{1}{\sqrt{k}}$) target known for linear utilities.

Cross-References

Algorithmic Mechanism Design

Recommended Reading

1. Babaioff M, Immorlica N, Kempe D, Kleinberg R (2008) Online auctions and generalized secretary problems. SIGecom Exch 7(2):1–11, DOI http://doi.acm.org/10.1145/1399589.1399596
2. Bateni M, Hajiaghayi MT, Zadimoghaddam M (2013) Submodular secretary problem and extensions. ACM Transactions on Algorithms 9(4):32
3. Dynkin EB (1963) The optimum choice of the instant for stopping a markov process. Sov Math Dokl 4:627–629
4. Feldman M, Naor J, Schwartz R (2011) Improved competitive ratios for submodular secretary problems (extended abstract). In: APPROX, pp 218–229
5. Hajiaghayi MT, Kleinberg R, Parkes DC (2004) Adaptive limited-supply online auctions. In: EC, pp 71–80
6. Kleinberg R (2005) A multiple-choice secretary algorithm with applications to online auctions. In: SODA, pp 630–631