首届致远学术节
学生科研成果展示

Learn to Play Maximum Revenue Auction

Tao Xiao, Xiaotie Deng, Keyu Zhu
Advisor: Prof. Xiaotie Deng

Abstract

Auctions for allocating resources and determining prices have become widely applied for services over the Internet, Cloud Computing, and Internet of Things in recent years. Very often, such auctions are conducted multiple times. They may be expected to gradually reveal participants' true value distributions, with which, it eventually would result in a possibility to fully apply the celebrated Myerson's optimal auction to extract the maximum revenue, in comparison to all truthful protocols. There is however a subtlety in the above reasoning as we are facing a problem of exploration and exploitation, i.e., a task of learning the distribution and a task of applying the learned knowledge to revenue maximization. In this work, we make the first step effort to understand what economic settings would make this double task possible exactly or approximately. The question opens up greater challenges in the wider areas where auctions are conducted repeatedly with a possibility of improved revenue in the dynamic process, most interestingly in auctioning cloud resources.

Keyword: *Bayesian auction, equilibrium analysis*

Approaches, Objectives and Results

We are interested in learning against a collection of self-interest-seek bidders for auctioneer's maximum revenue in the Bayesian auction. The task is more challenging because of two combined tasks of statistical learning of bidders' value distributions and optimal revenue extraction of the learned information. The model is motivated by a realistic auction scenario in today's Internet application like cloud computing resources allocation: data of bidders may be stored in the cloud and auctions are conducted repeatedly. Our study combines a game theoretical analysis with statistical learning.

The auctioneer aims to design a market rule of pricing and allocation such that participating bidders eventually do no have any incentive to change their bidding strategies, i.e., arriving at a Nash equilibrium. A natural way of modeling repeated bidding behavior is to assume that bids are drawn from a strategic prior. Bidders submit information including their parameterized value distributions (not necessarily true ones) and bids that follow bidders' reported value distributions.

We restrict our discussion to a probability model for true values and bids in a class of distributions characterized by a privately known parameter for each bidder. Therefore, in repeated implementation of the auction, each bidder selects its bids from a value distribution from the class. The auctioneer then applies the Myerson's auction based on the reported bids and their bidding distributions.

Based on several important probability distribution classes of bidders' priors, we find, in equilibrium, scenarios of no-cheating and strategic cheating behavior in terms of reporting their value distributions of bids. In the former part of our models, the Myerson's auction is fully learnable while it is approximately learnable in the latter part .

For power-law distributions, we show that bidders' dominant strategy is to bid truthfully, namely, report their true prior distributions (to be learned). Therefore, this leads to the best scenario for the auctioneer to implement the Myerson's auction where the bidders will report their true distributions and true value estimates of the item simultaneously. Similar results are found in other value distributions.

个人信息：2014级，数学
邮　　箱：keyu.zhu@tum.de