EDITOR RON VETTER
University of North Carolina Wilmington;
vetterr@uncw.edu

Learn to Play Maximum Revenue Auction

Xiaotie Deng, Peking University
Tao Xiao, Shanghai Jiao Tong University
Keyu Zhu, Georgia Institute of Technology

This installment of Computer's series highlighting the work published in IEEE Computer Society Journals comes from IEEE Transactions on Cloud Computing.

Auctions have recently become a key protocol to allocate resources and determine prices for services provided over the Internet, in cloud computing, and for the Internet of Things. Machine learning, in another direction, has naturally been utilized to learn the underlying value distributions of customers for better mechanism designs.

Our work[1] is motivated by a real-world scenario in which the task is to allocate cloud computing resources (for example, the Amazon Elastic Compute Cloud), and it is imperative to determine how much to charge for each service. Auctions are often repeated in cloud settings, and customers repeatedly play in the market. Learning their value distributions is expected in such a repeated game in which customers reveal more and more of their value distributions. From another perspective, Bayesian statistics heavily rely on a prior knowledge of data uncertainty to predict future events, and the optimization on decision variables affects future outcomes. This creates a scenario in an auction in which the auctioneer tries to achieve future optimality based on its learned probability distributions for the values of buyers, such as the case of the Myerson's auction, commonly referred to as the maximum revenue auction.

In the era of big data and cloud computing, the optimum auction in the Bayesian setting would be a game of two parties through data to form a conceptual two-stage process. The auctioneer collects bidding data to learn the prior distribution of the value distribution of buyers. Based on the learned prior information, the auctioneer's optimum auction extracts the maximum optimal revenue from the participating buyers.

Digital Object Identifier 10.1109/MC.2020.2981988
Date of current version: 4 June 2020

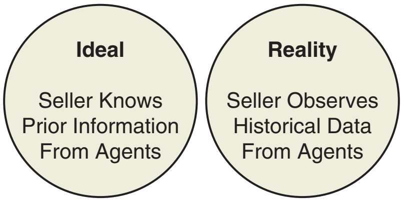

FIGURE 1. The traditional versus practical auction scenarios.

In another direction, the buyers can report their deviated value distributions to acquire a better utility function value in the equilibrium of the game of the seller and buyers. More specifically, buyers are assumed to have values over the resource, which follow certain probability distributions called the *prior*. Traditional Bayesian auction theory provides a good solution when the probability distributions are common knowledge.[2] In practice, however, this traditional theory faces a lot of challenges (Figure 1).

Statistical learning naturally becomes the most promising tool to learn the prior information, based on historical bids submitted by the customers. With the assumption that the bids are independent samples from the distributions, the optimal auction allows buyer value distributions to be learned. Hence, without specifying the auction, there will be no historical data collected. As long as the auction and learning algorithm are specified, individual agents will respond to the auction protocol and auctioneer's learning strategy and submit their strategic bids, which may or may not be the true prior information. Our work takes this central issue into consideration.

As the first article in this regime, we consider the learning mechanism to be exactly Myerson's auction, the revenue-optimal auction mentioned at the beginning. We consider practical scenarios in which value distributions from agents can be represented by parameters. For the manipulation performed by individual agents, we consider the simple and natural strategy space where the agents can manipulate over the parameters. Our work shows that it is possible to learn the revenue-optimal auction when individual priors are from power-law distribution family, while, for uniform and exponential distribution families, the learning task can be done when agents are from the same population.

> As long as the auction and learning algorithm are specified, individual agents will respond to the auction protocol and auctioneer's learning strategy and submit their strategic bids.

The more general sponsored search auctions were studied subsequently under our model. Revenue equivalence results were developed among generalized first price (GFP), generalized second price (GSP), and Vickrey–Clarke–Groves auctions in related hierarchical domains,[3] which also provides a revenue justification for the switch from GSP to GFP[4] for the repeated sponsored search auction.

ACKNOWLEDGMENT

This research was supported by the National Natural Science Foundation of China under grant 61632017.

REFERENCES

1. X. Deng, T. Xiao, and K. Zhu, "Learn to play maximum revenue auction," *IEEE Trans. Cloud Comput.*, vol. 7, no. 4, pp. 1057–1067, 2019. doi: 10.1109/TCC.2017.2712142.
2. R. B. Myerson, "Optimal auction design," *Math. Oper. Res.*, vol. 6, no. 1, pp. 58–73, 1981.
3. X. Deng, T. Lin, and T. Xiao, "Private data manipulation in optimal sponsored search auction," in *Proc. Web Conf. 2020*, Apr. 2020, Taipei, Taiwan, pp. 2676–2682.
4. R. Shields, "After months of preparation, Google set to roll out first-price auctions," *Adweek*, Sept. 5, 2019. [Online]. Available: https://www.adweek.com/programmatic/after-months-of-preparation-google-rolls-out-first-price-auctions/

XIAOTIE DENG is a chair professor of computer science at Peking University, China. Contact him at xiaotie@pku.edu.cn.

TAO XIAO is a Ph.D. student at Shanghai Jiao Tong University, China.

KEYU ZHU is a Ph.D. student at Georgia Tech, Atlanta.