Online Mechanism Design with Predictions

ERIC BALKANSKI, Columbia University, USA
VASILIS GKATZELIS, Drexel University, USA
XIZHI TAN, Drexel University, USA
CHERLIN ZHU, Columbia University, USA

In the framework of “algorithms with predictions,” algorithms are augmented with a machine-learned prediction and the goal is to obtain improved guarantees when the prediction is correct (*consistency*) while simultaneously guaranteeing some worst-case bounds even if the prediction is arbitrarily wrong (*robustness*). The majority of the work on this framework has focused on online algorithms with predictions regarding future input. A subsequent line of work has focused on mechanism design, where the prediction is regarding the private information of strategic agents. In this paper, we initiate the study of online mechanism design with predictions, which combines the challenges of online algorithms with predictions and mechanism design with predictions. We consider the problem of designing a revenue-maximizing auction to sell a single item to strategic bidders who arrive and depart over time. We study the learning-augmented version of this problem, where the auction designer is given a prediction regarding the maximum value over all agents. Our main result is a strategyproof mechanism whose revenue guarantees are α-consistent with respect to the highest value and $(1-\alpha^2)/4$-robust with respect to the second-highest value, for $\alpha \in [0, 1]$. We show that this trade-off is optimal within a broad family of auctions.

1 SUMMARY OF CONTRIBUTIONS

A well-established shortcoming of worst-case analysis is that it often leads to overly pessimistic conclusions. On the other hand, any non-trivial performance guarantee that can be established through worst-case analysis is very robust since it holds no matter what the input may be. In an attempt to overcome the limitations of worst-case analysis without compromising its robustness, the recently proposed framework of “algorithms with predictions” allows algorithms to be augmented with a machine-learned prediction that they can use as a guide. Crucially, this prediction may be highly inaccurate, so depending too heavily on it can lead to very poor performance in the worst case. Therefore, the goal in this framework is to use such a prediction so that a strong performance can be guaranteed whenever the prediction is accurate (known as the *consistency* guarantee) while simultaneously maintaining non-trivial worst-case guarantees even if the prediction is inaccurate (known as the *robustness* guarantee).

Since this framework was introduced, a surge of work has utilized it toward a refined analysis of algorithms, data structures, and mechanisms. The vast majority of this work has focused on the design and analysis of online algorithms. An even more recent line of work has successfully adapted this framework for the design and analysis of mechanisms interacting with strategic bidders. One of the canonical problems in mechanism design is the design of auctions for selling goods to a group of strategic bidders, aiming to maximize the revenue. The main obstacle in achieving this goal is the fact that the amount that each bidder is willing to pay is private information that the designer needs to carefully elicit. Learning-augmented mechanisms are therefore enhanced with predictions regarding the value of this private information, which can potentially alleviate this obstacle.

In this work, we initiate the study of online mechanism design with predictions, bringing together the two lines of work on online algorithms with predictions and mechanism design with predictions. Specifically, we consider the problem of revenue maximization from selling a good to strategic

Authors’ addresses: Eric Balkanski, eb3224@columbia.edu, Columbia University, New York, NY, USA; Vasilis Gkatzelis, gk4tzelis@gmail.com, Drexel University, Philadelphia, PA, USA; Xizhi Tan, Drexel University, Philadelphia, PA, USA, xizhi@drexel.edu; Cherlin Zhu, Columbia University, New York, NY, USA, cz2740@columbia.edu.

bidders that arrive and depart over time. This problem combines the challenges of both lines of work since the designer needs to carefully elicit the unknown, private, value of each bidder, while also not knowing (and being unable to elicit) the values of the bidders who have not yet arrived. In fact, designing an auction for such dynamic settings can be more demanding because, apart from the combined information limitations that the designer faces, the bidders may not only strategically misreport their value for the good(s) being sold, but also strategically misrepresent their arrival and departure times.

Our main result is a strategyproof mechanism whose revenue guarantees are α-consistent with respect to the highest value and $(1-\alpha^2)/4$-robust with respect to the second-highest value, for $\alpha \in [0,1]$. To achieve strategyproofness, our auction must very carefully determine the time at which the item is allocated, the bidder who receives the item, and the item's price in order to handle bidders who might be active during multiple phases, which is the main technical (and novel) challenge in our setting.

We also show that this trade-off is optimal within a broad and natural family of auctions, meaning that any α-consistent mechanism in that family has robustness at most $(1-\alpha^2)/4$. An interesting fact about the online mechanism design problem that we study is that proving impossible results for revenue maximization is significantly more demanding than proving impossibility results for the closely related secretary problem. First, the performance of the auction depends not only on who gets the good, but also at what price. For example, the auction could potentially offer the good to the highest value bidder for a price greater than the second-highest value, leading to revenue greater than the benchmark. Second, the price that the auction offers to a bidder can be an arbitrary function of the values observed among previous bidders, instead of just their relative ordering, giving rise to a very rich design space and making impossibility arguments quite demanding.

This work was published at The Twenty-Fifth ACM Conference on Economics and Computation [1].

REFERENCES

[1] Eric Balkanski, Vasilis Gkatzelis, Xizhi Tan, and Cherlin Zhu. 2024. Online mechanism design with predictions. In *25th ACM Conference on Economics and Computation.*