Mechanism Design with Predictions: An Annotated Reading List

ERIC BALKANSKI
Columbia University
and
VASILIS GKATZELIS
Drexel University
and
XIZHI TAN
Drexel University

A surge of recent work has focused on analyzing the performance of algorithms guided by predictions, aiming to enhance their worst-case performance guarantees with improved guarantees when the predictions are accurate. This "learning-augmented" framework was recently also extended to mechanism design settings involving strategic agents and we provide an overview of these results.

Categories and Subject Descriptors: B.6.3 [**Theory of computing**]: Algorithmic mechanism design

General Terms: Mechanism Design, Beyond Worst-Case Analysis

Additional Key Words and Phrases: Consistency, Robustness

For more than half a century, the dominant approach for the mathematical analysis of algorithms in computer science has been worst-case analysis. While worst-case analysis provides a useful signal regarding the robustness of an algorithm, it can be overly pessimistic and it often leads to uninformative bounds or impossibility results that may not reflect real-world obstacles. Meanwhile, advances in machine learning have led to very practical algorithms, most of which do not provide any non-trivial worst-case performance guarantees. Motivated by the tension between worst-case analysis and machine learning, a surge of recent work aims to design robust algorithms guided by machine-learned predictions. The goal of this literature on "algorithms with predictions" is to *simultaneously* provide two types of guarantees: "robustness" (which corresponds to the classic worst-case guarantees, even if the predictions are arbitrarily bad) and "consistency" (i.e., the performance guarantees when the predictions are accurate). This "learning-augmented framework" has been used successfully in a variety of settings, e.g., toward a refined analysis of competitive ratios in online algorithms and running times in traditional algorithms.

A very recent line of work has deployed this learning augmented framework in settings involving strategic agents. In such settings, the designer often faces information limitations, e.g., the participating agents' may have private information which they can strategically misreport, which limits the designer's ability to reach

Authors' addresses: eb3224@columbia.edu, gkatz@drexel.edu, xizhi@drexel.edu

desired outcomes. Mechanism design has proposed solutions to this problem, but their worst-case guarantees are often underwhelming from a practical perspective. Could we design learning-augmented mechanisms that combine "robustness' with strong "consistency" guarantees? Below are some initial works in this direction.

(1) Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. "Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location". In: *EC '22: The 23rd ACM Conference on Economics and Computation.* ACM, 2022, pp. 497–528

This work initiated the line of research on learning-augmented mechanism design and showcased the power of predictions in strategic settings, focusing on the canonical problem of strategic facility location. For both the egalitarian social cost and utilitarian cost, this paper provided truthful mechanisms enhanced with a prediction regarding the optimal facility location. These mechanisms achieve either the optimal consistency with the best-possible robustness (i.e., the best of both worlds) or the optimal trade-off between the two notions.

(2) Chenyang Xu and Pinyan Lu. "Mechanism Design with Predictions". In: *IJCAI '22: The 31st International Joint Conference on Artificial Intelligence,* 2022, pp. 571–577

This work concurrently initiated mechanism design with predictions. Rather than focusing on a single problem, it sampled a variety of different mechanism design problems, including auction design, frugality, scheduling, and facility location. In each of these settings, the results are truthful mechanisms that utilize predictions to achieve consistency guarantees that are better than the best-known worst-case performance guarantees, while simultaneously maintaining non-trivial robustness guarantees.

(3) Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. "Improved Price of Anarchy via Predictions". In: *EC '22: The 23rd ACM Conference on Economics and Computation.* ACM, 2022, pp. 529–557

While most papers on learning-augmented mechanism design focus on centralized mechanisms, this work studies a decentralized setting where the mechanism has limited information and can affect the agents' decisions only indirectly. It proposes cost-sharing protocols for classic job scheduling and network creation games which use predictions regarding the missing information and induce better Nash equilibria and improved price of anarchy bounds.

(4) Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. "Strategyproof Scheduling with Predictions". In: *ITCS '23: The 14th Innovations in Theoretical Computer Science Conference,* vol. 251. 2023, 11:1–11:22

This work focused on the celebrated problem of makespan minimization in strategic scheduling introduced by one of the first papers in AGT. It was con-

jectured, and very recently validated, that the best deterministic mechanism cannot achieve an approximation better than n. In this work, the authors provided a polynomial time mechanism, enhanced with predictions, that is 6-consistent and $2n$-robust, thus achieving asymptotically the best of both worlds (asymptotically optimal consistency and robustness).

(5) Gabriel Istrate and Cosmin Bonchis. "Mechanism Design With Predictions for Obnoxious Facility Location". In: *CoRR* abs/2212.09521 (2022)

This work considers the obnoxious facility location problem where the agents wish to be far from the facility instead of close to it. For segments, squares, circles, and trees, the authors provide truthful mechanisms augmented with predictions and bounded their consistency and robustness. The trade-offs obtained are shown to be optimal in one dimension.

(6) Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. "Bicriteria Multidimensional Mechanism Design with Side Information". In: *CoRR* abs/2302.14234 (2023)

This work focuses on multidimensional mechanism design. Rather than focusing on any specific setting, it proposes a general meta-mechanism that incorporates different types of side information to achieve both high social welfare and high revenue. The approach is versatile and can accommodate various sources of side information.

(7) Andres Muñoz Medina and Sergei Vassilvitskii. "Revenue Optimization with Approximate Bid Predictions". In: *NIPS '17: Advances in Neural Information Processing Systems 30.* 2017, pp. 1858–1866

This closely related work appeared before the papers initiating the line of work on algorithms with predictions. It focuses on finding good reserve prices in advertising auctions and proposes a method to reduce reserve price optimization to a standard setting of prediction under squared loss. They used a predictor to define a clustering of the data and compute the empirically maximizing reserve price for each group. The reduction directly ties the revenue gained by the algorithm to the prediction error, but without bounded robustness guarantees.

(8) Michael Mitzenmacher and Sergei Vassilvitskii. "Algorithms with Predictions". In: *Commun. ACM* 65.7 (2022), pp. 33–35

A survey of some of the initial results on algorithms with predictions.

(9) Alexander Lindermayr and Nicole Megow. *Algorithms with Predictions.* URL: https://algorithms-with-predictions.github.io/

This frequently updated website keeps track and categorizes papers in the area of algorithms with predictions. It allows easy search of papers by performance measure and/or type of problem.