Applied Learning-Augmented Algorithms with Heterogeneous Predictors

JESSICA MAGHAKIAN, Stony Brook University, USA
RUSSELL LEE, UMass Amherst, USA
MOHAMMAD HAJIESMAILI, UMass Amherst, USA
JIAN LI, Binghamton University-SUNY, USA
RAMESH SITARAMAN, UMass Amherst and Akamai Tech, USA
ZHENHUA LIU, Stony Brook University, USA

For many application domains, the integration of machine learning (ML) models into decision making is hindered by the poor explainability and theoretical guarantees of black box models. Although the emerging area of learning-augmented algorithms offers a way to leverage ML while enjoying worst-case guarantees, existing work usually assumes access to only one predictor. We demonstrate how to more effectively utilize historical datasets and application domain knowledge by intentionally using predictors of *different* quantities. By leveraging the heterogeneity in our predictors, we are able to achieve improved performance, explainability and computational efficiency over predictor-agnostic methods. Theoretical results are supplemented by large-scale empirical evaluations with production data demonstrating the success of our methods on optimization problems occurring in large distributed computing systems.

ACM Reference Format:
Jessica Maghakian, Russell Lee, Mohammad Hajiesmaili, Jian Li, Ramesh Sitaraman, and Zhenhua Liu. 2023. Applied Learning-Augmented Algorithms with Heterogeneous Predictors. In . ACM, New York, NY, USA, 2 pages. https://doi.org/XXXXXXX.XXXXXXX

1 SUMMARY OF CONTRIBUTIONS

This poster is based on results of the authors, recently published at the International Conference on Machine Learning (ICML '23) [3].

A large gap still remains between the current learning-augmented algorithms framework and practice. For example, the learning-augmented algorithms literature overwhelming considers access to one prediction. Defaulting to the classical online assumption of no information about the future when one prediction is inaccurate neglects the vast domain knowledge of practitioners. Most applications typically have many different mathematical and computational models to forecast quantities of interest that can greatly improve algorithm performance.

Several recent papers [1, 2] consider the setting of online algorithms with multiple predictions. However these works exclusively consider multiple predictions of the *same* quantity. For example, Dinitz et al. consider a portfolio of predictions generated by different ML models that cover the hyperparameter space. Although the predictors are able to specialize to different scenarios, there is no way to know a priori which predictor will do well on the following problem instance, thus requiring exploration to find the right choice.

SIGMETRICS-LATA '23, June 03–05, 2018, Orlando, Florida

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

In our ICML paper, we proposed an alternative approach of incorporating predictions of *different* quantities. The first predictor we introduced was a *parameter predictor* that learns the correct value of a tunable parameter of an online algorithm. The second predictor was an *input predictor* that predicts the unknown future inputs of the online algorithm in the form of short look-ahead windows.

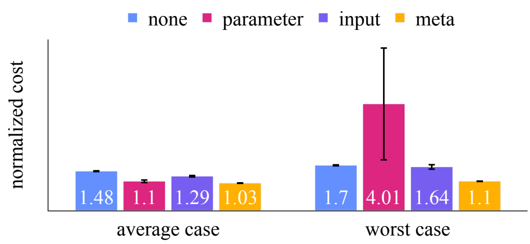

Fig. 1. Each predictor type has a different performance profile for average case (low prediction error) and worst case (high prediction error) settings. A simple meta-algorithm can learn to use the correct predictor for a given instance without the need for exploration.

Our Contributions. We presented the first work on online algorithms with heterogeneous predictors and the first work showing multiple noisy predictors outperform one in the worst-case. Our approaches have better explainability and computational efficiency over predictor-agnostic methods. We demonstrate that our algorithms can tackle real-world challenges such as COVID-19 related distributional shift in large distributed computing systems.

REFERENCES

[1] Keerti Anand, Rong Ge, Amit Kumar, and Debmalya Panigrahi. 2022. Online Algorithms with Multiple Predictions. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 582–598. https://proceedings.mlr.press/v162/anand22a.html

[2] Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. 2022. Algorithms with Prediction Portfolios. *arXiv preprint arXiv:2210.12438* (2022).

[3] Jessica Maghakian, Russell Lee, Mohammad Hajiesmaili, Jian Li, Ramesh Sitaraman, and Zhenhua Liu. 2023. Applied Online Algorithms with Heterogeneous Predictors. In *International Conference on Machine Learning*. PMLR, 582–598.