Optimal robustness-consistency tradeoffs for learning-augmented metrical task systems

NICOLAS CHRISTIANSON, California Institute of Technology
JUNXUAN SHEN, California Institute of Technology
ADAM WIERMAN, California Institute of Technology

ACM Reference Format:
Nicolas Christianson, Junxuan Shen, and Adam Wierman. 2023. Optimal robustness-consistency tradeoffs for learning-augmented metrical task systems. In . ACM, New York, NY, USA, 1 page. https://doi.org/XXXXXXX.XXXXXXX

ABSTRACT

We discuss our work recently published in [1] on the design of learning-augmented algorithms for the metrical task systems (MTS) problem that achieve optimal tradeoffs between robustness (*i.e.*, worst-case competitiveness) and consistency (*i.e.*, performance in comparison to that of a black-box *advice* algorithm, such as a machine-learned algorithm trained on historical problem instance data). Our algorithm, which we call "Distance-Adaptive Robust weight Transport" (DART), is a randomized algorithm that probabilistically follows the decisions of either a black-box advice algorithm or a traditional competitive/robust algorithm for MTS. Specifically, at each time of the MTS instance, it updates the probabilities with which it follows the advice or robust algorithms according to their relative performance and the distance between the two decisions, coupling consecutive probability distributions with the optimal transportation coupling. We obtain the following upper bound on the robustness-consistency tradeoff obtained by DART.

THEOREM 0.1. *For any $\epsilon > 0$, the algorithm DART achieves $(1+\epsilon)$-consistency and $2^{O(1/\epsilon)}C$-robustness, where C is the competitive ratio of the provided robust algorithm.*

While this exponential tradeoff between robustness and consistency appears at first glance to be overly pessimistic, we further establish a lower bound demonstrating that in general, such a tradeoff is necessary.

THEOREM 0.2. *Any algorithm for MTS that is $(1+\epsilon)$-consistent must have robustness at least $2^{\Omega(1/\epsilon)}$.*

However, in certain special cases of the MTS problem, such as in the case of the k-server problem, DART achieves a tighter tradeoff between robustness and consistency without any modifications to the algorithm.

THEOREM 0.3 (INFORMAL). *When applied to the k-server problem, for any $\epsilon > 0$, DART obtains $(1+\epsilon)$-consistency and $O(\frac{k}{\epsilon})$-robustness.*

REFERENCES

[1] CHRISTIANSON, N., SHEN, J., AND WIERMAN, A. Optimal robustness-consistency tradeoffs for learning-augmented metrical task systems. In *International Conference on Artificial Intelligence and Statistics* (2023), PMLR, pp. 9377–9399.