Online Stochastic Max-Weight Bipartite Matching: Beyond Prophet Inequalities

CHRISTOS PAPADIMITRIOU, Columbia University
TRISTAN POLLNER, Stanford University
AMIN SABERI, Stanford University
DAVID WAJC, Stanford University

CCS Concepts: • **Theory of computation** → **Online algorithms**.

Additional Key Words and Phrases: online matching, prophet inequalities, online optimum

ACM Reference Format:
Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. 2021. Online Stochastic Max-Weight Bipartite Matching: Beyond Prophet Inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3465456.3467613

Full version: Can be found at https://arxiv.org/abs/2102.10261.

The rich literature on online Bayesian selection problems has long focused on so-called prophet inequalities, which compare the gain of an online algorithm to that of a "prophet" who knows the future. An equally-natural, though significantly less well-studied benchmark is the optimum *online* algorithm, which may be omnipotent (i.e., computationally-unbounded), but not omniscient. What is the computational complexity of the optimum online? How well can a polynomial-time algorithm approximate it?

Problem Statement. The input is a (known) bipartite graph, with nodes of one side, termed offline nodes, present at time zero. At each time $t \geq 1$, online node t arrives with some (known) probability p_t, yielding a (known) value of v_t if it is matched. These matching choices must be made immediately and irrevocably, with the objective to maximize the overall gain. (*Note:* In the full paper, we consider an extension where online nodes sample weights to offline nodes from arbitrary distributions.)

The optimal offline algorithm (the prophet) can be approximated with a factor of 1/2 (see, e.g., [3]). This ratio is tight, as the problem generalizes the original single-item prophet inequality.

We study the optimal *online* algorithm: it is well-defined, and computable via an exponential-sized dynamic program. How well can be approximate this algorithm in polynomial time?

Hardness. We show that it is PSPACE-hard to approximate this problem within some constant $\alpha < 1$. To the best of our knowledge, this is the first hardness result for the computation of the optimal online algorithm in an online Bayesian selection setting. Our result builds on a sequence of reductions from the (PSPACE-hard) approximation of the stochastic SAT problem, due to Condon et

EC '21, July 18–23, 2021, Budapest, Hungary

ACM ISBN 978-1-4503-8554-1/21/07.
https://doi.org/10.1145/3465456.3467613

al. [1]. In particular, our proof proceeds via the use of expander graphs for the hardness of bounded-degree stochastic SAT, and reducing this problem to our online Bayesian matching problem.

Algorithm. On the flip side, we present a polynomial-time algorithm which approximates optimal online within a factor of 0.51—beating the best-possible prophet inequality.

At the core of our algorithm is a linear program with a constraint that naturally separates online from offline algorithms (previously observed in [5]). This LP's variables $\{x_{i,t}\}$ intuitively upper bound the probability edge (i, t) is matched by the optimal online algorithm. Our algorithm solves this LP and uses its solution to match each edge (i, t) with probability at least $0.51 \cdot x_{i,t}$, yielding our approximation ratio, by linearity of expectation. Broadly, for each arriving online node t, we pick a neighbor i with probability based on x, and if i is free, match (i, t) with some correcting probability, as in [2].

Unlike the algorithm of [2], we allow some online nodes t to make *two* choices of an offline neighbor, to increase the marginal probabilities for certain edges. Here, we rely crucially on the LP constraint separating online from offline algorithms, and a careful analysis of correlations between offline nodes' matched statuses. From this, we obtain our 0.51 approximation of the optimal online algorithm, beating the best-possible ratio of 1/2 of any prophet inequality.

Conclusions and Open Questions. It is natural to further study the efficient approximability of our problem. We suspect that much better approximation guarantees are achievable. A related interesting question is to obtain better approximation for the widely-studied special case of online nodes drawn from some i.i.d distribution.

More broadly, one might ask how well one can approximate the optimal online algorithm for online Bayesian selection problems under the numerous constraints studied in the literature, including matroid and matroid intersections, knapsack constraints, etc. For which of these problems is the online optimum easy to compute? Which admit a PTAS? Which admit constant approximations? Which are hard to approximate? We are hopeful that the ideas developed here will prove useful when exploring this promising research direction.

Follow-up work: Since this paper was posted to Arxiv, the last two authors have extended this paper's algorithm to obtain improved algorithms for the (seemingly unrelated) online edge coloring problem [4]. Their ideas can be used to improve our approximation ratio from 0.51 to 0.526.

ACKNOWLEDGMENTS

Research supported in part by NSF Awards CCF1763970, CCF191070, CCF1812919, ONR award N000141912550, and a gift from Cisco Research. We thank the authors of [5] for drawing our attention to their work, and the anonymous EC reviewers for useful comments.

REFERENCES

[1] Anne Condon, Joan Feigenbaum, Carsten Lund, and Peter Shor. 1997. Random debaters and the hardness of approximating stochastic functions. *SIAM Journal on Computing (SICOMP)* 26, 2 (1997), 369–400.

[2] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. 2020. Online Stochastic Max-Weight Matching: prophet inequality for vertex and edge arrival models. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 769–787.

[3] Michal Feldman, Nick Gravin, and Brendan Lucier. 2014. Combinatorial auctions via posted prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 123–135.

[4] Amin Saberi and David Wajc. 2021. The Greedy Algorithm is not Optimal for Online Edge Coloring. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*.

[5] Alfredo Torrico and Alejandro Toriello. 2017. Dynamic Relaxations for Online Bipartite Matching. *arXiv preprint arXiv:1709.01557* (2017).