Online Ad Allocation with Predictions

Michelle Zhou[1,2], Alina Ene[2]

Lynbrook High, 1280 Johnson Ave, San Jose, CA 95129[1];
Boston University, Boston, MA 02215[2]

INTRODUCTION

- In the rapidly evolving digital landscape, **online advertising** plays a pivotal role in promoting products and services to the target audience
  - **Display Ads** and the **generalized assignment problem (GAP)** are two widely studied online packing problems in this domain
    - Both involve the immediate allocation of ad impressions to budget-constrained advertisers as impressions arrive in real-time
    - Advertisers value users differently based on search queries or demographic data
- Traditional worst-case algorithms are optimal in theory but might act overly cautious in practice due to the predictable nature of real-world data
- We implement an algorithm [1] for both problems that incorporates **machine-learned predictions** to improve performance beyond the worst case

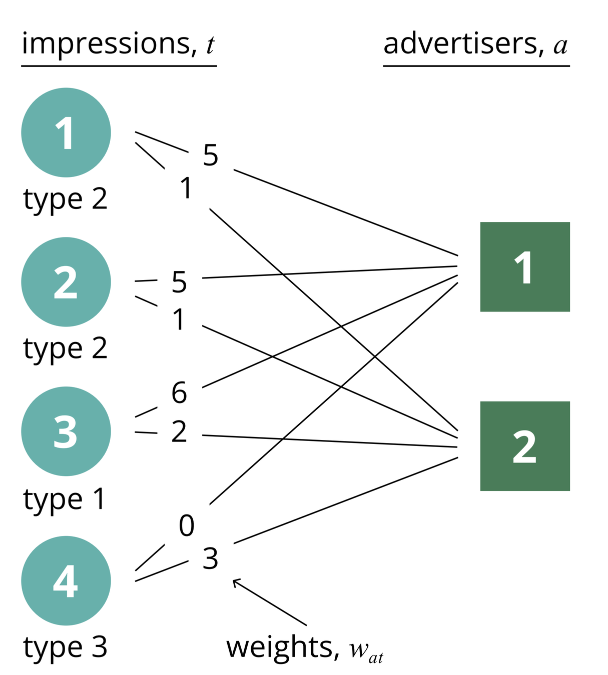

Fig. 1 A small-scale example of valuations. Advertisers value impressions of the same type (a group of homogeneous users) identically. Values might be a prediction of click-through probability or any valuation from the advertiser.

RESULTS

- - - Random Mixture: Run the worst-case algorithm for some parameter $q \in [0, 1]$ and follow the prediction exactly for $1 - q$

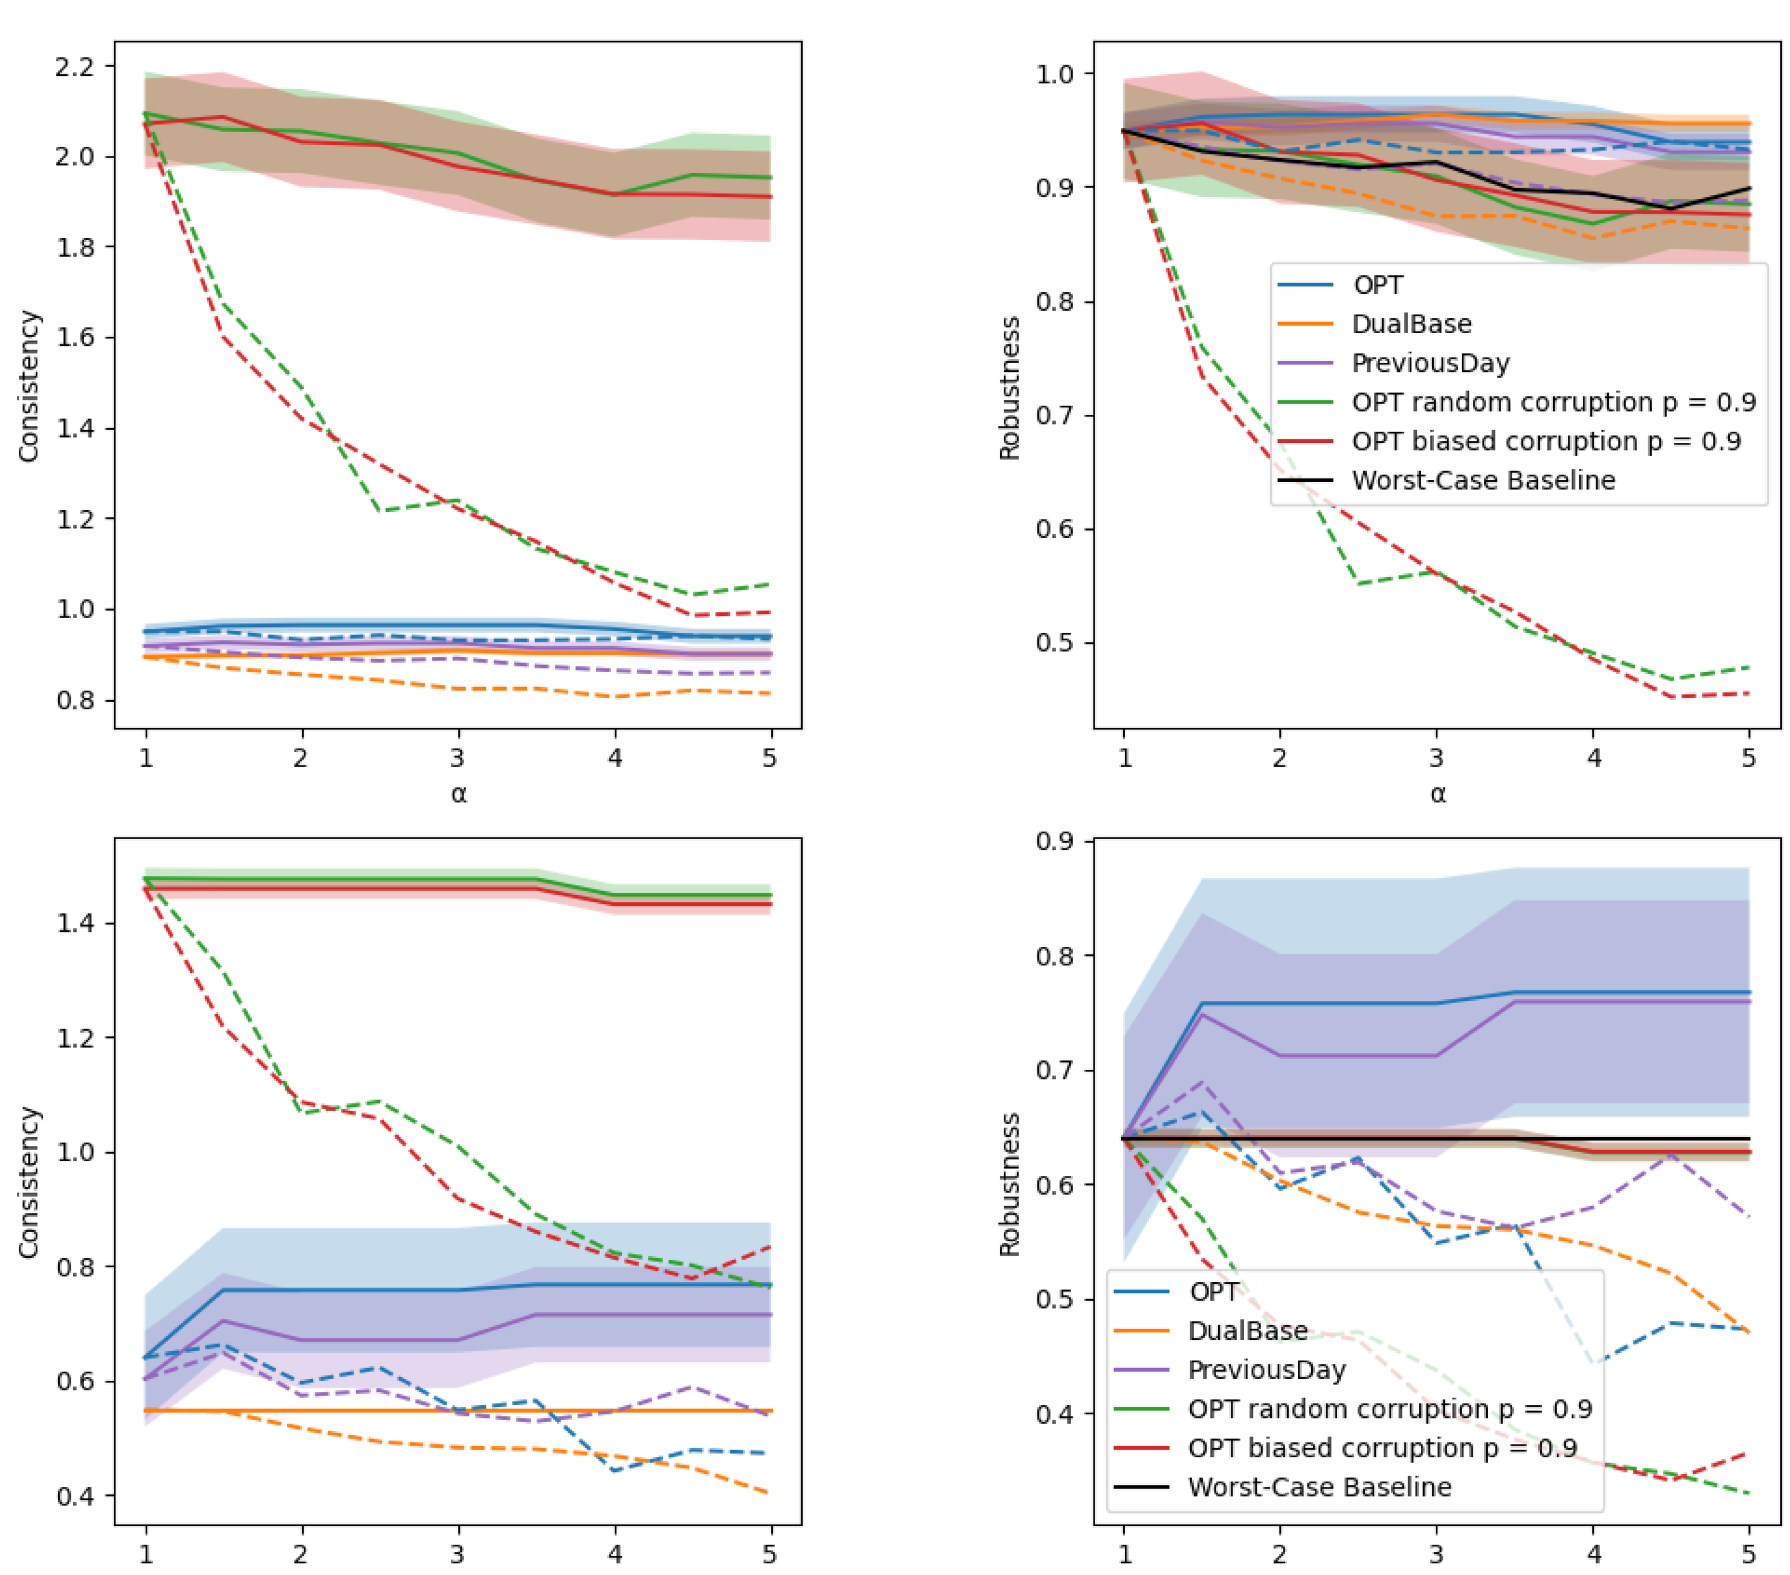

Fig. 3 Two illustrations of consistency (left) and robustness (right) for varying α. We run the algorithms 5 times and report average for both algorithms and the standard deviation only for our algorithm, to avoid clutter. $q = 1 / \alpha$.

- A synthetic instance of 10 advertisers, budgets $B_a \in [5, 20]$, and 100 impressions of 10 types
- Algorithm 1 bypasses the strong lower bound on the worst-case competitive ratio
- Algorithm 1 outperforms the random-mixture algorithm
- The **exponential averaging** thresholds (top) outperform the **minimum** thresholds (bottom)

DISCUSSION/CONCLUSIONS

- We implement a learning-augmented algorithm for Display Ads and GAP with free disposal that improves the performance beyond the worst-case

Future Work

- Improve the consistency of the performance
  - Some predictors do not always perform better than the worst-case
- Test different LP solvers
  - 20 advertisers and 400 impressions of 20 types takes CVXOPT ~12 seconds
- Test on real-world data

METHODS

Algorithm 1

- Input: Robustness-consistency trade-off parameter $\alpha \in [1, \infty)$, advertiser budgets, $B_a \in \mathbb{N}$
- Define the constants $B := \min_a B_a$, $e_B := (1 + 1 / B)^B$, and $\alpha_B := B\,(e_B^{\alpha/B} - 1)$
- For each advertiser, a, initialize $\beta_a \leftarrow 0$

for all arriving impressions, t → $a_{(\mathrm{PRD})} = \mathrm{PRD}(t)$, $a_{(\mathrm{EXP})} = \mathrm{argmax}_a\{w_{at} - \beta_a\}$ → $a = \mathrm{argmax}\{\alpha_B\,(w_{a(PRD)t} - \beta_{a(PRD)}),\ w_{a(EXP)t} - \beta_{a(EXP)}\}$ → allocate t to a and remove a's least valuable impression → update thresholds, β_a end for

- Incorporates predictions in the primal-dual algorithm with free disposal [2]
- A larger value of α means that we trust the prediction more

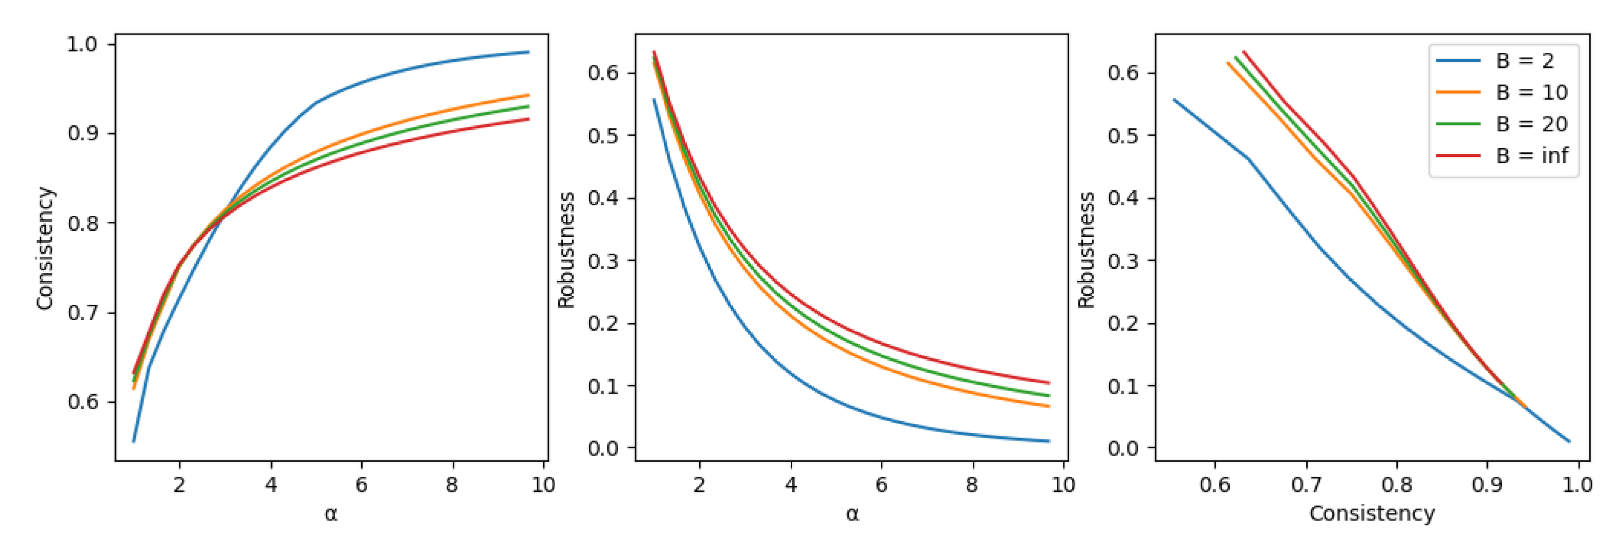

Fig. 2 An illustration of the robustness-consistency trade-off of Algorithm 1 for various values of α and budgets B.

Predictors

- **Optimum Solution (OPT)**: We solve the primal linear program (LP)
  - Random corruption: We reallocate to randomly chosen advertisers
  - Biased corruption: We reallocate according to a random permutation
- **Dual Base**: We solve the dual LP on a fraction of the impressions to obtain $\{\beta_a\}_a$, then allocate impressions to advertisers that maximize the discounted gain, $w_{at} - \beta_a$
- **Previous Day**: We solve the dual LP on impressions from the previous day to obtain $\{\beta_a\}_a$, then allocate impressions to advertisers that maximize the discounted gain, $w_{at} - \beta_a$

Synthetic Instances

- We generate advertisers' valuations for impression types by sampling from an **exponential distribution**
  - We sample the same number of impressions from each type, then sort them by ascending order of display times from a **Gaussian distribution**
- We equip advertisers with fixed budgets

Testing

- We use two important measures for the performance of the algorithm
  - The **robustness** ALG/OPT indicates how well the algorithm performs against the optimum solution
  - The **consistency** ALG/PRD measures how close the algorithm gets to the prediction's objective value
- We test a variety of implementations on a variety of synthetic instances
  - We test three algorithms, five predictors, and two thresholds
  - We test on varying advertisers, budgets, impressions, etc.

REFERENCES

[1] Fabian Spaeh and Aline Ene. (2023). Online ad allocation with predictions. URL https://doi.org/10.48550/arXiv.2302.01827.
[2] Jon Feldman, Nitish Korula, Vahab S. Mirrokni, S. Muthukrishnan, and Martin Pál. (2009). Online ad assignment with free disposal. In *Workshop of Internet Economics (WINE)*, pp. 374–385.

ACKNOWLEDGEMENTS

I would like to thank Professor Alina Ene for guiding me through this amazing research opportunity with incredible patience and kindness. A special thanks to the Boston University RISE program and my family for making this experience possible.