Learning-augmented (Online) Algorithms with Societal Design Criteria

New techniques for teaching performance analysis

Mohammad Hajiesmaili

TeaPACS@SIGMETRICS, June 2023

SOLAR Lab
Sustainability, Optimization, Learning, and Algorithms Research

Manning College of Information & Computer Sciences

University of Massachusetts Amherst

This Talk is about Trade-offs!

ALGORITHMS WITH PREDICTIONS

SOCIETAL ALGORITHM DESIGN

- How these two are related …
- What are foundational new analysis techniques that could be part of teaching relevant courses?

Algorithms with Predictions

aka, learning-augmented algorithms, algorithms with ML advice

Black-box ML and good practical performance

Theoretical approach and worst-case analysis

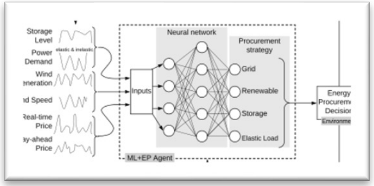

- Exploit data/history
- Build robust models to predict future
- Good on average performance
- Less robust to outliers

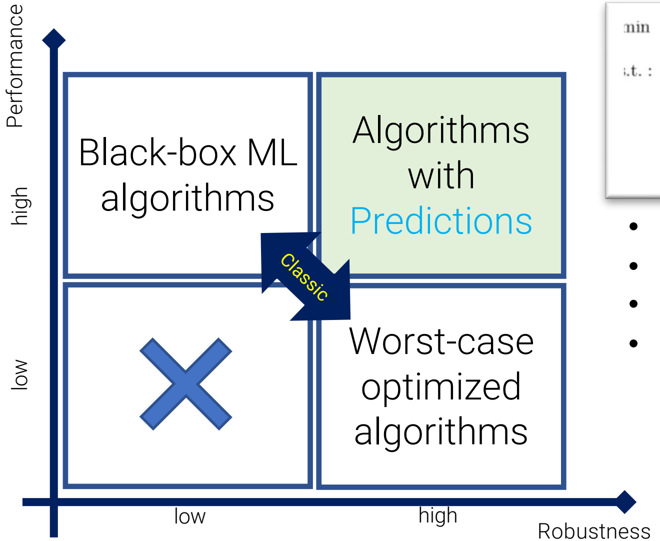

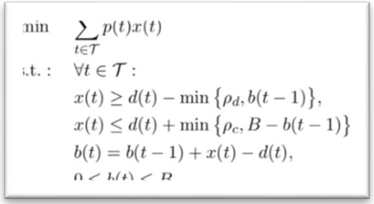

- Future is fully unknown
- Assume worst-possible future
- Pessimistic
- Hard to exploit patterns in data

An Emerging Topic

- Research
  - Sigmetrics Workshop on Learning-augmented Algorithms: Theory and Applications, 2023
  - Data-driven Decision Processes, Simon Institute, 2022
  - Workshop on Algorithms with Predictions, EPFL, 2022
  - Workshop on Algorithms with Predictions, STOC, 2022
  - Machine Learning for Algorithms, FODSI, 2021
  - Workshop on Algorithms with Predictions, STOC 2020
- Teaching
  - Algorithms with Predictions: UMass Amherst, 2022
  - Learning Augmented Algorithms, CWI, 2021
  - Machine Learning for Algorithm Design, Columbia University, 2020
  - Learning-augmented algorithms, MIT, 2019
- Paper Repository
  - https://algorithms-with-predictions.github.io/, 143 papers (June 15, 2023)

Societal Algorithm Design

AI-SDM

AI Institute for Societal Decision Making

Funded by NSF

Societal Criteria
- Fairness
- Carbon Awareness
- Privacy
- Safety

Motivating Applications
- Sustainability
- Healthcare
- Disaster response
- Law and Policy

SIMONS INSTITUTE for the Theory of Computing

Societal Considerations Applications

Program	Data-Driven Decision Processes
Date	Monday, Nov. 7 – Thursday, Nov. 10, 2022

Online Algorithms

- Inputs arrive one-by-one
- The algorithm makes an irrevocable decision
- A long list of classic online problems
  - Ski-rental
  - Paging/caching, k-server, metrical task systems
  - Bin packing
  - Knapsack problem
  - Matching
  - Job scheduling
  - Secretary problem, prophet inequalities
  - Set cover, online covering

Performance metrics

Online Algorithms

- Competitive ratio: $CR(ALG) = \max_{\omega \in \Omega} \frac{ALG(\omega)}{OPT(\omega)}$
- Multiplicative quantity
- Mainly in classic algorithm design

Online Learning

- Regret
- Additive quantity
- Mainly in theoretical ML

Algorithms with Predictions

Desiderata

Consistency: How well the algorithm performs when the prediction is good?

Near-offline performance if predictions are reasonably good

Graceful degradation of performance with quality of predictions

Robustness: How well the algorithm performs when the prediction is bad?

Near-online performance, even if predictions are very bad

Parsimony: Algorithms should use predictions if they are useful, otherwise, do not ask
Retain both consistency and robustness
Predictions can be expensive to obtain

Robustness vs. Consistency

[Mahdian-Nazerzadeh-Saberi'07, Lykouris-Vassilvitskii'18, Purohit-Svitkina-Kumar'18]

- Consistency is worst-case performance when prediction is accurate. An algorithm is η-consistent if

Good prediction!

$$\max_{\omega\in\Omega} \frac{ALG(\omega, p^\star(\omega))}{OPT(\omega)} \le \eta$$

- Robustness is the worst-case performance overall. An algorithm is γ-robust if

$$\max_{\omega\in\Omega, p} \frac{ALG(\omega, p)}{OPT(\omega)} \le \gamma$$

Any prediction!

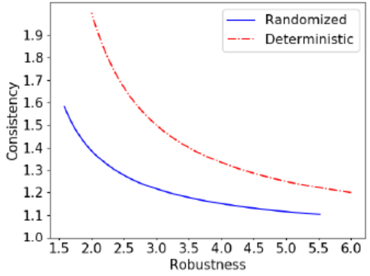

- Pareto-optimality: for given robustness, no online algorithms can achieve better consistency.

- M. Mahdian, H. Nazerzadeh, and A. Saberi. "Online optimization with uncertain information." *ACM Trans on Algorithms* 2012
- T. Lykouris and S. Vassilvitskii. "Competitive caching with machine learned advice." *Journal of the ACM*, 2021
- M. Purohit, Z. Svitkina, and R. Kumar. "Improving online algorithms via ML predictions." *NeurIPS* 2018
- B. Etienne, A. Maggiori, and O. Svensson. "The primal-dual method for learning augmented algorithms." *NeurIPS 2020*
- A. Wei and F. Zhang. "Optimal robustness-consistency trade-offs for learning-augmented online algorithms." *NeurIPS* 2020

Online Knapsack Problem (OKP)

- A knapsack with unit capacity
- N items arrive online
  - Item j has values v_j and weight w_j
- Problem formulation:

$$\max_{x_j} \sum_{j=1}^{n} v_j x_j \text{, s.t., } \sum_{j=1}^{n} w_j x_j \leq 1, \qquad x_j \in \{0,1\}$$

- Assumption:
  - The value density of items are bounded and known, i.e., $\frac{v_j}{w_i} \in [L, U], \forall j$
- **Closely related to**
  - One-way trading
  - K-search problem

Online Algorithms for OKP

[Zhou, Chakrabarty, Lukose '08]

- ZCL: Admit the new item only if its unit value is above a threshold.
- Threshold function: $\phi(z) = (Ue/L)^{zL}/e$
- Competitive ratio: $O(\log U/L)$

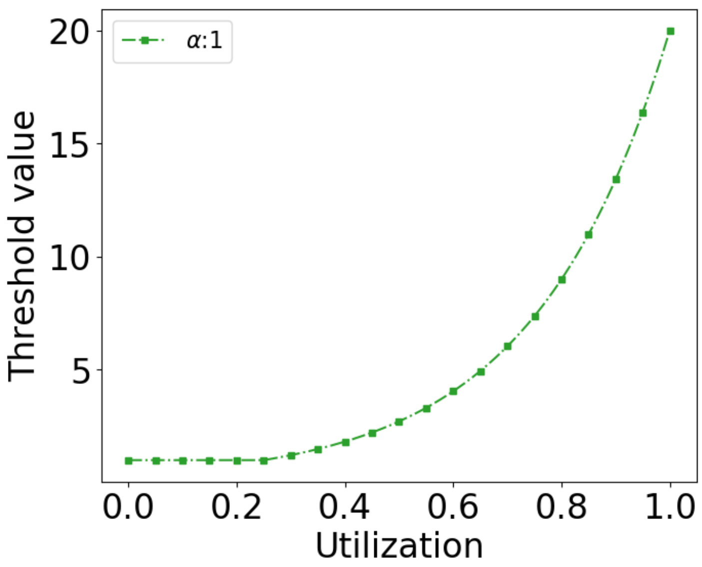

Extensive Literature on OKP

Later this week in FCRC

The Online Knapsack Problem with Departures
B. Sun, L. Yang, M. Hajiesmaili, J. Lui, A. Wierman, D. Towsley, and D. Tsang
ACM SIGMETRICS 2023, Tuesday June 20

Near-optimal Online Algorithms for Joint Pricing and Scheduling in EV Charging Networks
R. Bostandoost, B. Sun, C. Joe-Wong, and M. Hajiesmaili
ACM e-Energy 2023, Wednesday June 21 (Best paper candidate)

Design of Threshold Function

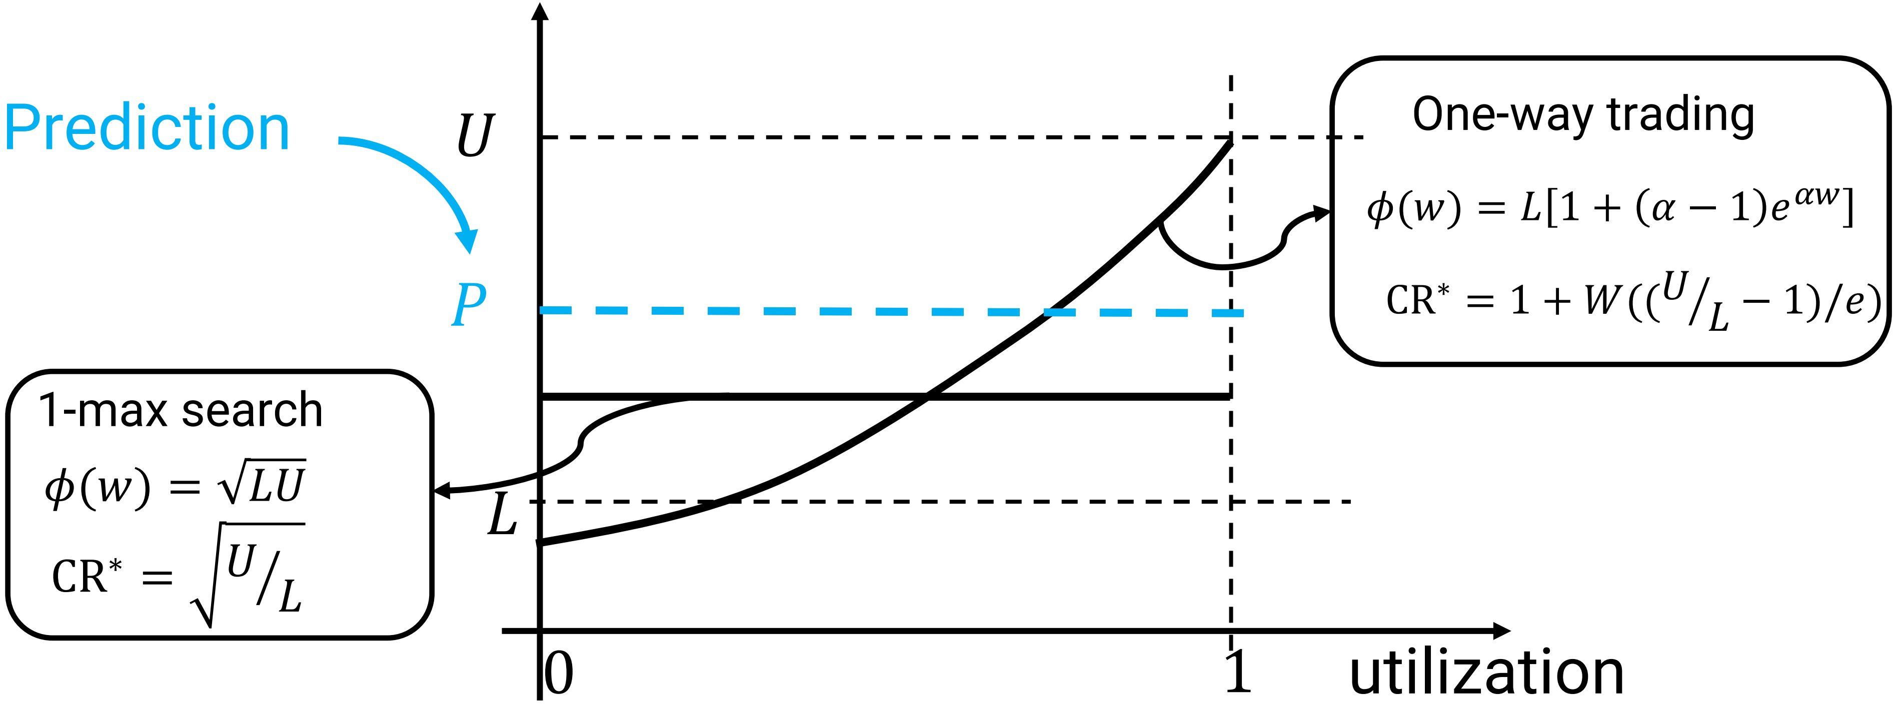

Challenge: given a prediction of the maximum value P, how to incorporate it into the threshold function to guarantee robustness and consistency?

An Intuitive Learning-augmented Algorithm

- Given $\lambda \in [0,1]$ and prediction P
  - **Reserve** $(1-\lambda)$ resource to trade at the price no less than P
  - **Run** α-competitive online algorithm with resource λ
- This algorithm is $(\frac{\alpha}{\lambda+(1-\lambda)\alpha})$-consistent and $(\frac{\alpha}{\lambda+(1-\lambda)\alpha/\theta})$-robust

- B. Sun, R. Lee, M. Hajiesmaili, A. Wierman, and D. Tsang. "Pareto-optimal learning-augmented algorithms for online conversion problems." *NeurIPS 2021*
- R. Lee, B. Sun, J. Lui, and M. Hajiesmaili. "Pareto-Optimal Learning-Augmented Algorithms for Online k-Search Problems." *arXiv preprint arXiv:2211.06567* 2022

Lower bound is about Trade-off!

- **Lower bound**: Any γ-robust online algorithm must have a consistency

$$\eta \geq \frac{\theta}{\frac{\theta}{\gamma} + (\theta - 1)(1 - \frac{1}{\gamma}\ln\frac{\theta - 1}{\gamma - 1})}$$

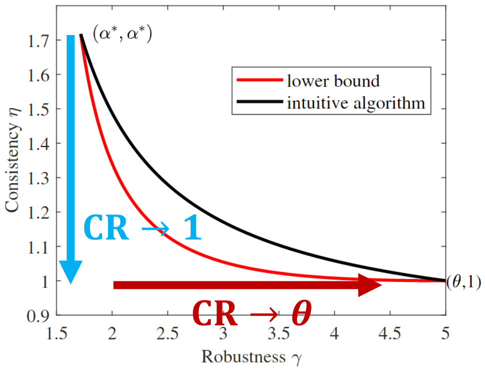

As consistency improves CR → 1 (best possible ratio), robustness degrades CR → θ (worst possible ratio)

- B. Sun, R. Lee, M. Hajiesmaili, A. Wierman, and D. Tsang. "Pareto-optimal learning-augmented algorithms for online conversion problems." *NeurIPS 2021*
- R. Lee, B. Sun, J. Lui, and M. Hajiesmaili. "Pareto-Optimal Learning-Augmented Algorithms for Online k-Search Problems." *arXiv preprint arXiv:2211.06567* 2022

Threshold Function for One-way Trading

- The threshold function consists of multiple segments
  - price region near prediction P, ensure consistency
  - price region away from prediction P, ensure robustness

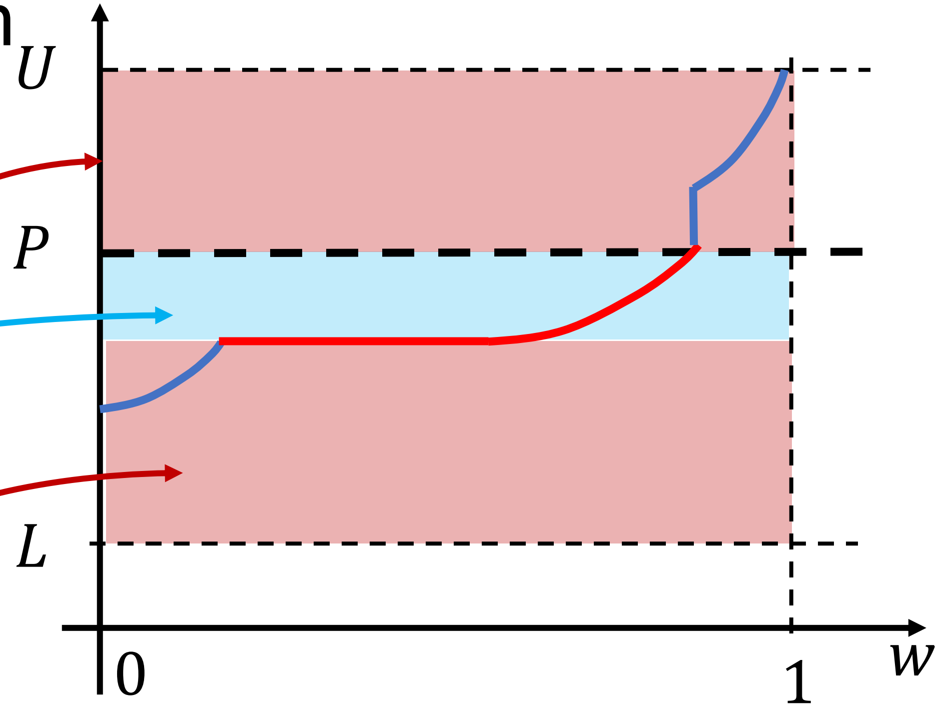

Pareto-optimal Trade-off

- What about the same analysis for the online knapsack problem?
  - Open problem!
  - We need a different prediction!
- There has been a recent progress for unit-size online knapsack problem

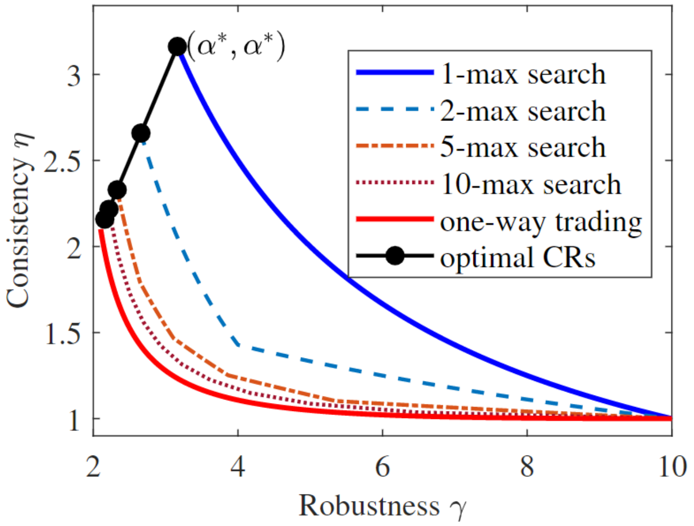

- B. Sun, R. Lee, M. Hajiesmaili, A. Wierman, and D. Tsang. "Pareto-optimal learning-augmented algorithms for online conversion problems." *NeurIPS 2021*
- R. Lee, B. Sun, J. Lui, and M. Hajiesmaili. "Pareto-Optimal Learning-Augmented Algorithms for Online k-Search Problems." *arXiv preprint arXiv:2211.06567* 2022
- S. Balseiro, C. Kroer, and R. Kumar. "Single-leg revenue management with advice." ACM *EC 2023*

Lesson Learned

	Consistency vs. robustness	Lower bounds are about trade-offs
	It is all about the predictions	Nontrivial predictions may help a lot more.
	Parametric analysis of performance metrics	Analyzing the degradation factor (the price of learning)

Fairness in Online Algorithms

- Regret settings
  - Full information feedback
    - Tue@SIGMETRICS: Online Fair Allocation with Perishable Resource
    - Thu@SIGMETRICS: Enabling Long-term Fairness in Dynamic Resource Allocation
    - Our work: No-regret Algorithms for Fair Resource Allocation
    - A lot more
  - Bandit feedback
    - Fair Exploration via Axiomatic Bargaining

Fairness in the Online Knapsack Problem

motivation

Time fairness in the context of the secretary problem. [Arsenis, Kleinberg, EC ’22]

The high-level idea is that we should treat an input based on its quality and not on its arrival time.

Time fairness in online knapsack?

Example: managing bids for cloud computing resources.

Want to accept or reject bids using a resource while getting enough profit — Simple and “fair” acceptance criteria might be desirable.

Lechowicz, Adam, Rik Sengupta, Bo Sun, Shahin Kamali, and Mohammad Hajiesmaili. "**Time Fairness in Online Knapsack Problems**." *arXiv preprint arXiv:2305.13293* (2023).

The ZCL Algorithm is not Time Fair

- ZCL: Admit the new item only if its unit value is above a threshold.

- Threshold function: $\emptyset(z) = \left(\frac{Ue}{L}\right)^{z}\frac{L}{e}$

- Competitive ratio: $O(\log {U}/{L})$

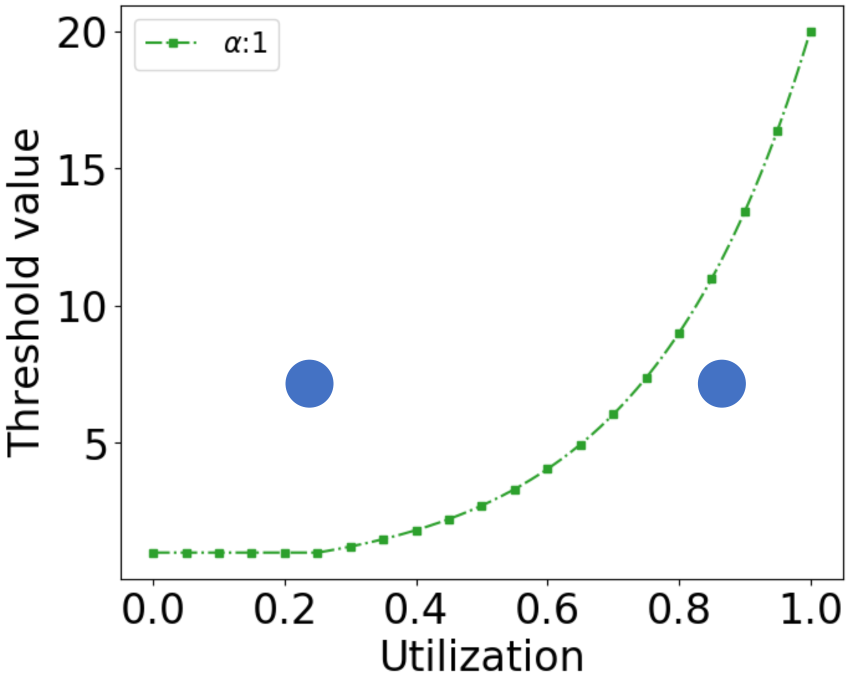

Time Fairness

- [Arsenis, Kleinberg '22] formalize Time-Independent Fairness.
- We propose α-Conditional Time-Independent Fairness (α-CTIF). ALG is α-CTIF if the following holds over all input sequences:

$$Pr\left[ALG\ accepts\ j-h\ item \mid \left(\frac{v_j}{w_j} = x\right) and\ (z_j + w_j \in \mathcal{A})\right] = p(x)$$

  - There exists some utilization interval $\mathcal{A} = [a, b]$s.t.$|b - a| = \alpha$, where an item with value density x is accepted with *constant probability* $p(x)$.

Achieving Fairness

without prediction ...

- An easy way to achieve 1-CTIF is to set a constant threshold ϕ such that any item is accepted iff its value is $\geq \phi$.
  - $^U/_L$-competitive … not interesting!

Does prediction help?

- **ECT:** An algorithm that balance between fairness vs. competitiveness

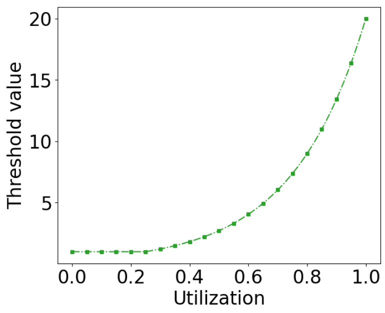

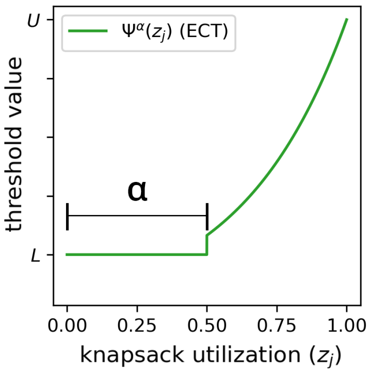

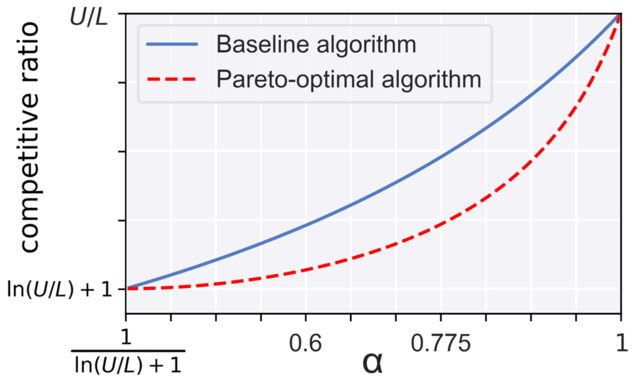

Achieving Fairness

use predictions to break the deterministic Pareto bound …

- What to predict matters!
  - We choose to predict the optimal threshold $d^\star$: the lowest value where accepting all items $\geq d^\star$ gives a constant approximation of OPT.
- LA-ECT
  - receives a prediction of $d^\star$
  - γ-CTIF
  - $\frac{1}{1-\gamma}(\log(U/L) + 1)$ -robust
  - $O\left(\frac{1}{\gamma}\right)$-consistent

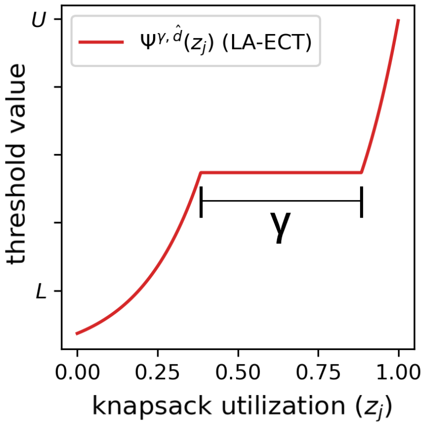

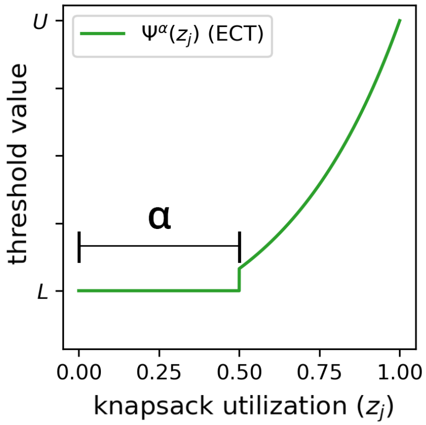

Numerical experiments

- Cloud traces from Google data centers
- Algorithms: ZCL, ECT, and LA-ECT

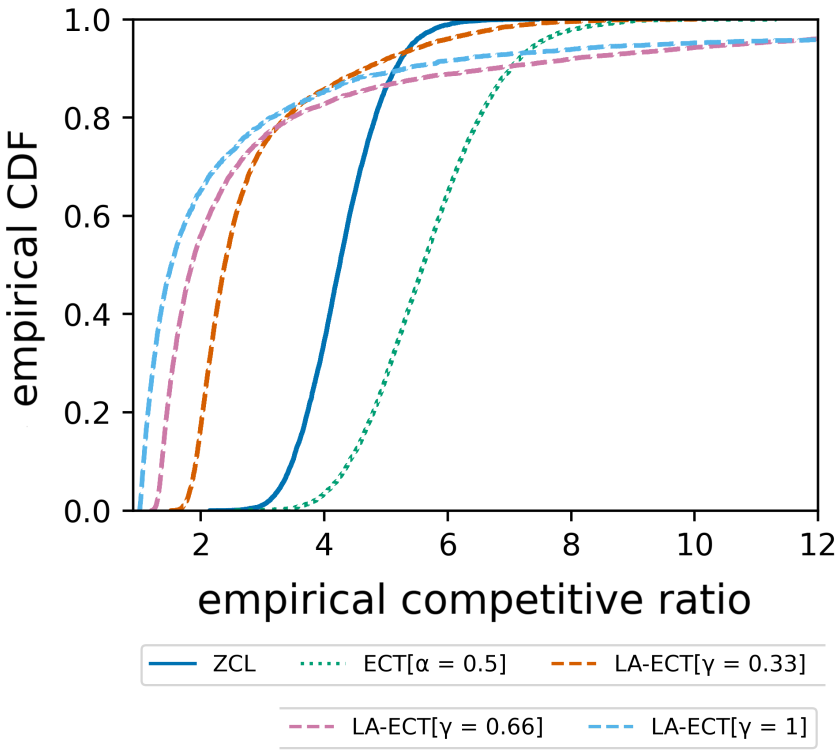

Carbon Intelligent Computing

The story is similar…

The War of the Efficiencies: Understanding the Tension between Carbon and Energy Optimization
W. Hanafy, R. Bostandoost, N. Bashir, D. Irwin, M. Hajiesmaili, P. Shenoy
HotCarbon 2023 – to appear

Online Pause and Resume Problem: Optimal Algorithms and Applications to Carbon-Aware Load Shifting
A. Lechowicz, N. Christianson, J. Zuo, N. Bashir, M. Hajiesmaili, A. Wierman, P. Shenoy
ArXiv:2303.17551

What is new in teaching?

- The performance analysis is inherently about trade-offs
  - Robustness vs. consistency
  - Fairness vs. competitiveness
- Meaningful societal algorithm design should be learning-augmented!
- Parametric analysis of algorithms
- Revisit the theoretical analysis of simple and practical algorithms
- Break the fundamental impossibility results

Thank you!

Mohammad Hajiesmaili

TeaPACS@SIGMETRICS, June 2023