Pareto-Optimal Learning-Augmented Algorithms for Online Conversion Problems

Bo Sun, Russell Lee, Mohammad Hajiesmaili, Adam Wierman, Danny H.K. Tsang

Learning-Augmented Online Algorithms

- **Online algorithm** makes online decisions when inputs are revealed piece-by-piece and future input pieces are unknown
  - Optimize the worst-case **competitive ratio** of offline algorithm that knows the entire input in hindsight and online algorithm

$$\mathrm{CR} = \max_I \frac{\mathrm{OPT}(I)}{\mathrm{ALG}(I)}$$

- **Learning-augmented algorithm** leverages **machined-learned predictions** to design online algorithms
  - Characterize the competitive ratio as a function of the prediction error ε

$$\mathrm{CR}(\varepsilon) = \max_I \frac{\mathrm{OPT}(I)}{\mathrm{ALG}(I,\varepsilon)}$$

  - **Consistency:** the competitive ratio when the prediction is accurate, i.e., $\eta = \mathrm{CR}(0)$
  - **Robustness:** the worst-case competitive ratio regardless of the prediction error, i.e., $\gamma = \max_{\varepsilon \geq 0} \mathrm{CR}(\varepsilon)$

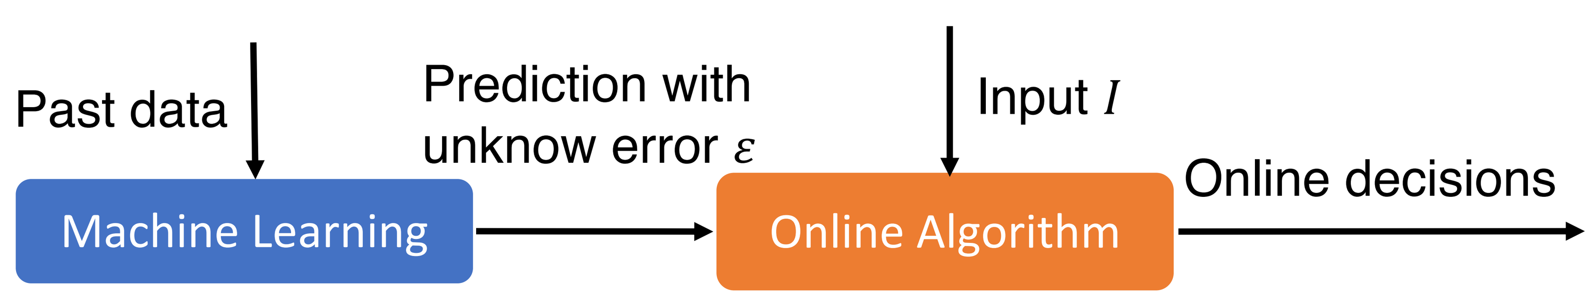

Learning-augmented online algorithms

- **Existing works**
  - defining consistency and robustness [Lykouris, ICML'18]
  - trade-offs between two metrics [Purohit, NeurIPS'18]
  - Pareto-optimal algorithms for the ski-rental problem [Bamas, NeurIPS'20][Wei, NeurIPS'20]
  - online algorithms with bounded robustness and consistency for other problems, e.g., secretary and online matching [Antoniadis, NeurIPS'20], metrical task systems [Antoniadis, ICML'20]

- **Contributions of this work**
  - Designing **Pareto-optimal algorithms** for **online conversion problems**, which achieve optimal trade-offs between consistency and robustness
  - Proposing a novel method of deriving the **lower bound** on the consistency-robustness trade-off, which may be of use beyond online conversion problems

Online Conversion Problem

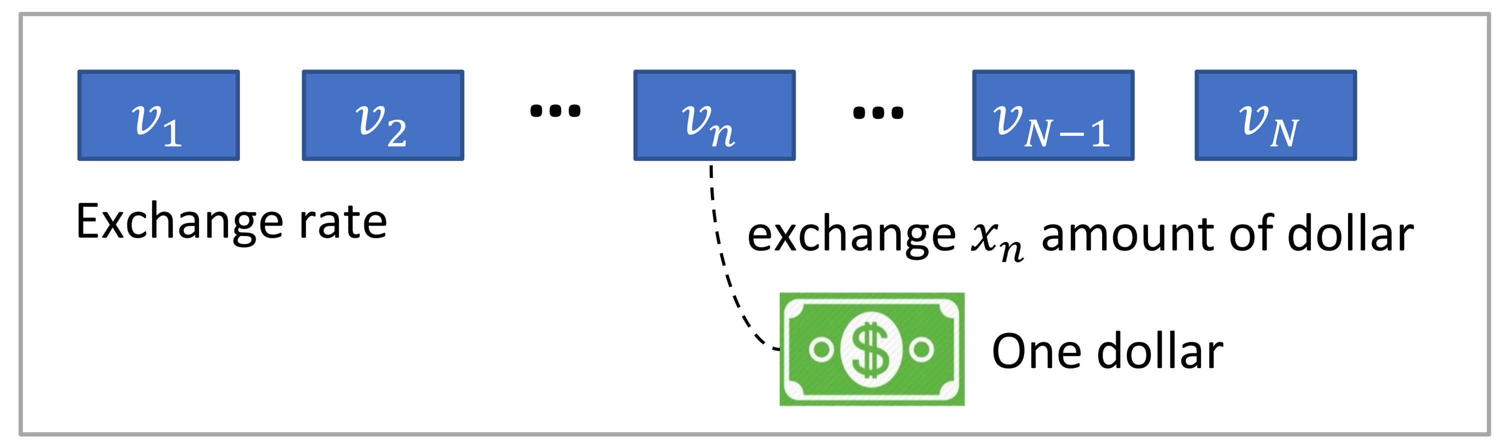

- **Online conversion problem** aims to exchange one asset (e.g., one dollar) to another asset (e.g., Yen) over time-varying exchange rates
  - For $n = 1, \dots, N$
    - Observe an exchange rate v_n
    - Decide the amount of dollar to trade, $x_n \in \mathcal{X}_n$, and obtain return $x_n v_n$
  - Goal: maximize the total return $\sum_{n\in[N]} x_n v_n$
  - Two settings:
    - **1-max search** with integral decisions, i.e., $\mathcal{X}_n = \{0,1\}$
    - **one-way trading** with fractional decisions, i.e., $\mathcal{X}_n = [0,1]$
  - Assumption: exchange rate is bounded $v_n \in [L, U]$, with $\theta = U/L$

Unified Online Threshold-Based Algorithm (OTA)

- **Inputs:**
  - Threshold function $\boldsymbol{\phi}(\boldsymbol{w})$:[0,1]→$[L,U]$
  - Initial utilization (cumulated traded dollar) $w^{(0)} \leftarrow 0$
- **For $n = 1, \dots, N$, do**
  - Observe current utilization $w^{(n-1)}$ and exchange rate v_n
  - Determine online decision x_n^* by solving a pseudo-problem

$$\max_{x_n \in \mathcal{X}_n} v_n x_n - \int_{w^{(n-1)}}^{w^{(n-1)}+x_n} \phi(u)du$$

return — pseudo-cost

  - Update $w^{(n)} \leftarrow w^{(n-1)} + x_n^*$

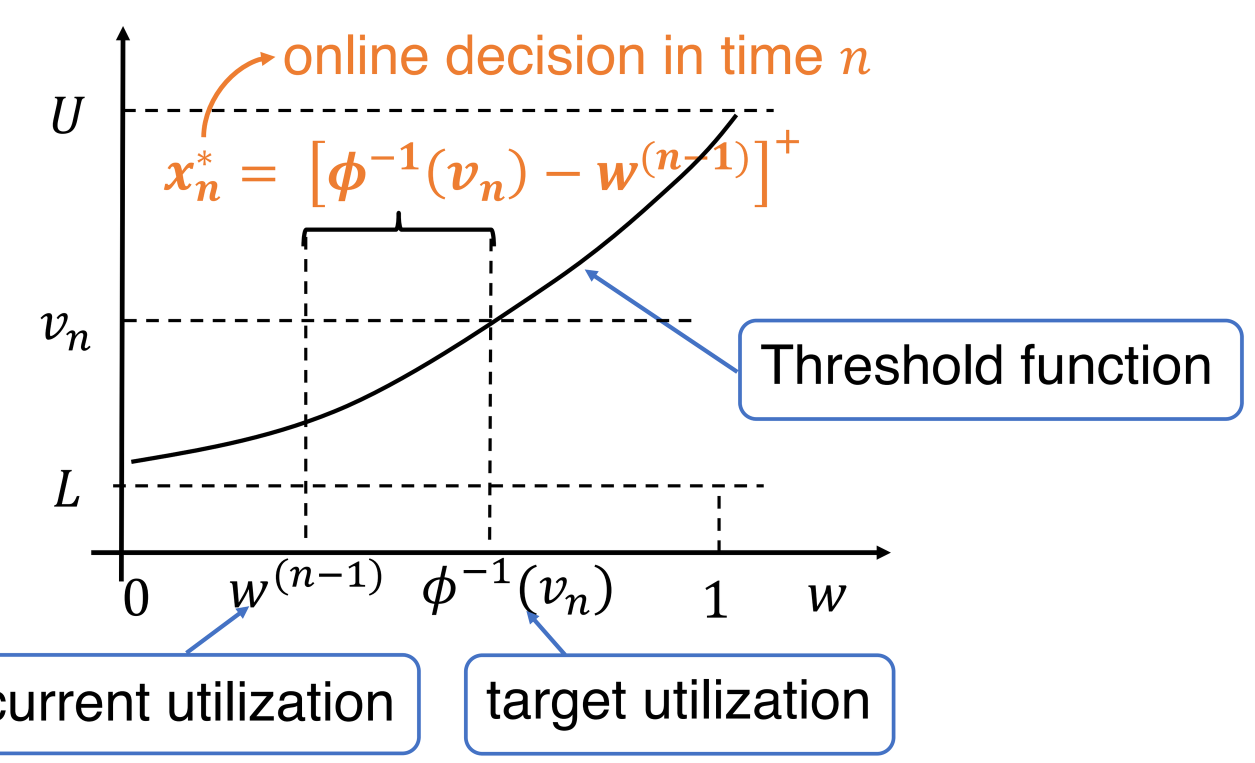

Pareto-Optimal Learning-Augmented OTA

- **Threshold function for the worst-case:**
  - 1-max-search: $\phi(w) = \sqrt{LU}$ achieves optimal $\mathrm{CR}^* = \sqrt{\theta}$
  - One-way trading: $\phi(w) = L[1 + (\alpha - 1)e^{\alpha w}]$ achieves optimal $\mathrm{CR}^* = \alpha = 1 + W((\theta - 1)/e)$, where $W(\cdot)$ is Lambert function
- **Threshold function for robustness and consistency,** given a prediction of the **maximum exchange rate $\boldsymbol{P}$**:
  - **Sufficient condition**: if the function consists of multiple segments and each segment satisfies **a set of differential equations**, OTA is with bounded consistency and robustness
  - **Pareto-Optimality**: for a given robustness, no online algorithms can achieve a better consistency than learning-augmented OTA

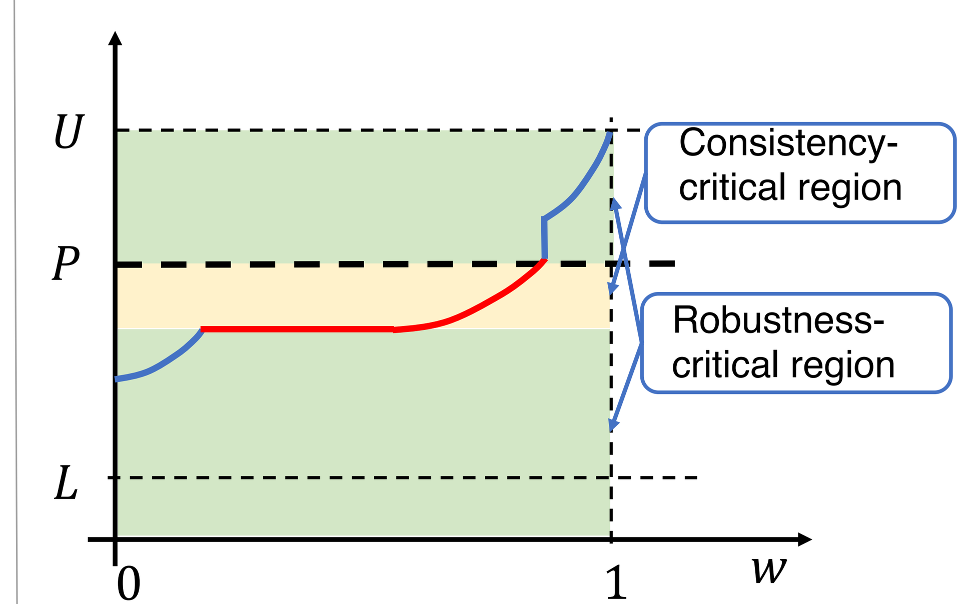

Threshold function for learning-augmented OTA

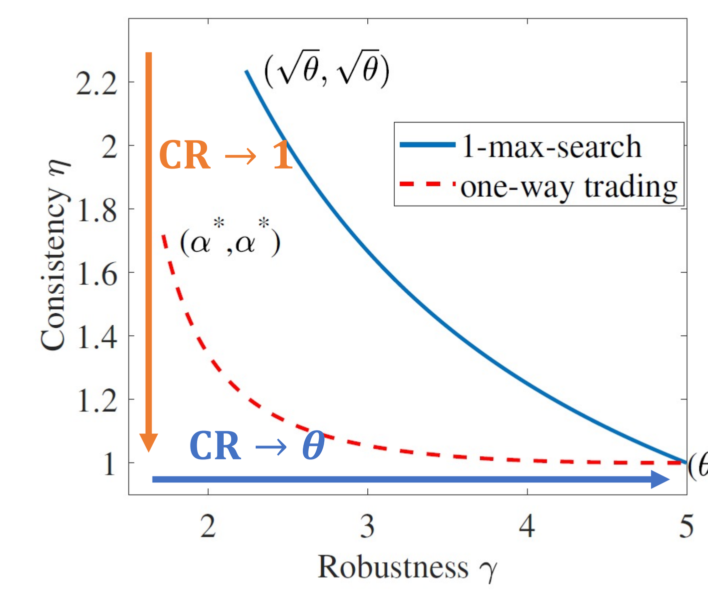

Pareto-boundaries for online conversion problems

Experiments on Bitcoin conversion

- **Setup:**
  - Historical BTC prices (2015-2020) divided into 250 instances
  - For each instance, select an OTA indexed by $\lambda \in [0,1]$, sell one BTC fraction-by-fraction, and evaluate its empirical ratio
  - Compared algorithms
    - ALG(λ^w): worst-case optimized online algorithm without prediction
    - ALG(λ^{off}): an algorithm that selects the best λ offline
    - ALG(λ^{alf}): an algorithm that selects λ using online learning algorithms
- **Learning-augmented OTA** outperforms pure online algorithms for varying **prediction errors** and **exchange rate crashes**