Near-Optimal Bayesian Online Assortment of Reusable Resources

Yiding Feng (MSR) Rad Niazadeh (Booth) Amin Saberi (Stanford)

Microsoft Research

CHICAGO BOOTH

Stanford University

Personalized Online Assortment

- Users $[T]$ & *reusable* resources $[n]$
- Type τ of a user encodes
  - r_i^τ: rental fee from each resource i
  - $G_i^\tau = \mathbb{P}[d_i^\tau \leq x]$: rental duration CDF for each resource i
  - $\phi_i^\tau(S) = \mathbb{P}[x \text{ is chose in assortment } S]$

  [Assumption I] all choice models satisfy weak substitution

  [Assumption II]] oracle access to an optimal offline assortment solver $\mathrm{argmax}_S \sum \alpha_i \phi_i^\tau(S)$
- Each resource $i \in [n]$ has capacity c_i; let $c_{\min} = \min_i c_i$
- Users arrive at platform in an *online* fashion
  - when user $t \in [T]$ arrives, platform observes type $\tau \sim F_t$
  - platform selects assortment $S_t \subseteq [n]$ based on current resource availability
  - user t selects item $i \sim \phi_i^\tau(S)$, rental duration $d_i \sim G_i^\tau$
  - platform collects rental fee r_i^τ

Competitive Ratio & Revenue Benchmark

- **Goal**: given $F_1, \ldots, F_T$, design online (randomized) assortment algorithm that maximizes total collected rental fees
- **Competitive ratio**:

$$\min_{\text{problem instance}} \frac{\mathbb{E}[\text{ALG}]}{\mathbb{E}[\text{Benchmark}]}$$

- **Revenue benchmark**: any upper bound on the expected revenue of best online policy

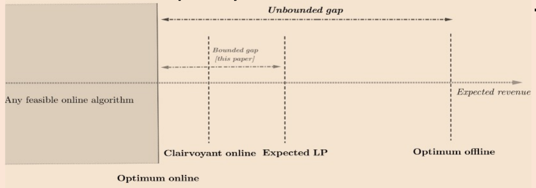

Expected LP

knows realized type sequence,

not knows realized choices/durations

inventory feasibility in expectation

Main Result

Main Theorem.

A near-optimal algorithm that (i) runs in polynomial time; and (ii) obtains the following competitive ratio w.r.t. Expected LP

$$\max\left(\frac{1}{2}, 1 - \varepsilon^*(c_{\min})\right), \qquad \varepsilon^*(c_{\min}) = \mathcal{O}\left(\sqrt{\frac{\log(c_{\min})}{c_{\min}}}\right)$$

- no stationary assumption, benchmark is strongest
- information theoretically optimal (up to log factors)

General Four-step Recipe

- Pre-processing:
  - solve Expected LP to obtain $\{y^*_{S,t,\tau}\}$

At each time $t \in [T]$:

- Simulation: <lossless rounding>
  - given realized type τ, sample assortment $\tilde{S} \sim \{y^*_{S,t,\tau}\}_{S \subseteq [n]}$
- Discarding: <handle ex post inventory feasibility>
  - for each $i \in \tilde{S}$, send "suggestion query" to $Discarding^{(i)}$
  - $Discarding^{(i)}$ is procedure (possibly adaptive) that decides whether to discard resource $i \in \tilde{S}$ or not
- Post-processing: <handle technical difficulties below>
  - given un-discarded resources $\hat{S}$, pick (possibly randomized) subset $\bar{S} \subseteq \hat{S}$ to help "global operation" of the algorithm

Technical Difficulties

- Weak substitution: discarding resourece i increases resource j's selecting probability
- Reusable resource: *positive* correlation of resource i across different time periods

Sub-assortment Sampling in Post-processing

- Input: assortment $\hat{S}$ and targeted probabilities $\{p_i\}_{i \in \hat{S}}$ s.t.

$$p_i \leq \phi_i^\tau(\hat{S})$$

- Output: sample sub-assortment $\bar{S} \subseteq \hat{S}$ matching targeted probability, i.e.,

$$\mathbb{E}_S[\phi_i^\tau(\bar{S})] = p_i$$

- Implications: handle difficulties due to weak substitution and resource reusability

Resource-wise Discarding Procedure

- Uniform discarding towards $1 - \varepsilon^*(c_{\min})$
  - discard with probability $\varepsilon^*(c_{\min})$
- DP-based discarding towards ½
  - calculate per-unit revenue-to-go with replenishment

$$\mathcal{V}_{i,t} = h(\{y^*_{S,t,\tau}\})$$

  - discard if $r_i^\tau < \mathcal{V}_{i,t+1} - \sum g_i^\tau(d) \cdot \mathcal{V}_{i,t+d}$

Hybrid Algorithms: Best-of-both-world

- **Question**: how to select discarding procedure for a given instance?
- Hybrid-1: for each resource i
  - compute $\mathcal{R}_i = \frac{c_i \mathcal{V}_{i,1}}{\text{"Optimal revenue of resource } i \text{ in Expected LP"}}$
  - select DP-based discarding procedure if $\mathcal{R}_i \geq 1 - \varepsilon^*(c_i)$
- Hybrid-2: for each user t
  - use Mento Carlo to estimate expected future revenue from each discarding procedure
  - select discarding procedure with higher estimation

SSRN preprint: 3714338