Seller-side Outcome Fairness in Online Marketplaces

Zikun Ye[1], Reza Yousefi Maragheh[2], Lalitesh Morishetti[2], Shanu Vashishtha[2], Jason Cho[2], Kaushiki Nag[2], Sushant Kumar[2], Kannan Achan[2]

1 University of Washington,
2 Walmart Global Tech

NeurIPS
Dec 14, 2023

Outline

Outline

Background: Popularity bias

- Popularity bias: popular items are overly exposed in recommendations at the cost of less popular items that customers may find interesting. This is pervalent in online marketplaces.

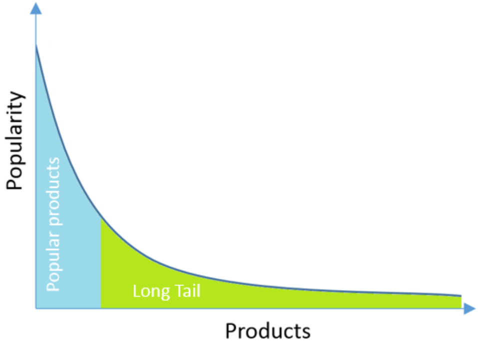

Background

- Drawbacks:
  - Lost of diversity of items.
  - Lost of seller loyalty: sellers may perceive unsatisfied promotions, sales on the marketplace.

Background

- Drawbacks:
  - Lost of diversity of items.
  - Lost of seller loyalty: sellers may perceive unsatisfied promotions, sales on the marketplace.
- Goals:
  - Incentivize sellers by: more exposures, sales, transparency and outcome fairness.
  - Walmart: more diversity without loss of revenue.

Paper's Question:

Design better and unified allocation rules (algorithms) to attack popularity bias issue and achieve outcome fairness.

Outline

Offline Allocation Model

- There are total T exposures/views, indexed by $t \in [T]$, and m sellers, indexed by $j \in [m]$. Seller j has K_j products to be sold.
- And for each view t:
  - Display K products.
  - With the side information, (predicted) conversion rate for each items c_{tjk}, $\forall j \in [m], k \in [K_j]$.
  - Expected revenue is $p_{jk}c_{tjk}$, $\forall j \in [m], k \in [K_j]$, where p_{jk} is unit revenue for the product.
-

$$
\begin{aligned}
\max_{x} \quad & \sum_{t=1}^{T}\sum_{j=1}^{m}\sum_{k=1}^{K_j} p_{jk}c_{tjk}x_{tjk} \\
\text{s.t.} \quad & \sum_{j=1}^{m}\sum_{k=1}^{K_j} x_{tjk} \le K, \ \forall t \in [T] \\
& x_{tjk} \in \{0,1\}, \ t \in [T], \ j \in [m], \ k \in [K_j]
\end{aligned} \tag{1}
$$

constraint: cardinality

Offline Regularized Allocation Model

- $$\begin{aligned} \max_{x} \quad & \sum_{t=1}^{T}\sum_{j=1}^{m}\sum_{k=1}^{K_j} p_{jk} c_{tjk} x_{tjk} + \textit{regularization term} \\ \text{s.t.} \quad & \sum_{j=1}^{m} x_{tjk} \le K, \ \forall t \in [T] \\ & x_{tjk} \in \{0,1\}, \ t \in [T], \ j \in [m], \ k \in [K_j] \end{aligned}$$

Challenges:

- Maximizing total revenue over all sellers can easily cause popularity bias
- Recommendation happens in real-time, most parameters are not known in advance or change over time. Moreover, it could be even biased due to inaccurate estimations from underlying machine learning models. Solved by bandit algorithms (exploration and exploitation)

Offline Regularized Allocation Model

$$\max_{x \in \mathcal{X}} \sum_{t=1}^{T} \sum_{j=1}^{m} \sum_{k=1}^{K_j} p_{jk} c_{tjk} x_{tjk} + r(\cdot)$$

Regularization function $r(a_1, a_2, ..., a_m)$ is a multi-dimensional concave function in aggregated outcomes (total exposures, clicks, purchases, revenue) at the seller level, $a_j,\ j \in [m]$ for sellers.

- Example 1 (No regularizer): $r(a) = 0$. Uncover most recommendation rules to maximize revenue.
- Example 2 (Above-Target Revenue $a_j = \sum_t \sum_{k=1}^{K_j} p_{jk} c_{tjk} x_{tjk}$): $r(a) = \sum_{j=1}^{m} \beta_j \min(a_j, \alpha_j)$ with the threshold $\alpha_j \geq 0$ and unit reward β_j. This regularizer can be used when the decision maker want all sellers have more than α_j revenue. β_j essentially captures the implicit non-monetary benefit by doing so.
- Example 3 (Max-min Revenue Fairness): $r(a) = \lambda \min_j(a_j)$. This regularizer imposes outcome fairness for sellers, i.e., no seller gets too-small revenue.

Outline

Primal Dual Models

Take Above-Target (Seller level) Exposures for illustration, i.e., we want to let each seller j has more than α_j exposures over her all products. Let $\mathcal{X} = \{x : x \geq 0, \sum_j \sum_k x_{tjk} \leq K\}$ be the constraint set.

$$\max_{x \in \mathcal{X}} \sum_{t=1}^{T} \sum_{j=1}^{m} \sum_{k=1}^{K_j} p_{jk} c_{tjk} x_{tjk} + \sum_{j=1}^{m} \beta_j \min\left(\sum_{t=1}^{T} \sum_{k=1}^{K_j} x_{tjk}, \alpha_j\right)$$

[Primal]

$$\min_{0 \leq \lambda_j \leq \beta_j} \sum_{t=1}^{T} \max_{x \in \mathcal{X}} \{(p_{jk} c_{tjk} + \lambda_j) x_{tjk}\} - \sum_{j=1}^{m} \alpha_j \lambda_j$$

[Dual]

Outline

Proposed Algorithm

Algorithm 1 Online Dual Gradient Descent with Learning

Initialize: decision variables $\lambda_j^1 = 0 \; \forall j \in [m]$, predicting machine learning model M^1 and its training method $\mathcal{L}$, step size η, learning rate γ_t, the total number of items H

for $t = 1$ to T **do**

  Step 1: Obverse the realization of contextual information s_{tjk} for prediction. Predict CVR $\hat{c}_{tjk}$ via the model $M^t(s_{tjk})$, for all sellers and their products.

  Step 2: Compute the fairness-aware rank score $\hat{f}_{tjk} = (p_{jk} + \lambda_j^t)\hat{c}_{tjk}$ for all sellers $j \in [m]$ and products $k \in [K_j]$. Then, find the empirical optimal product $\hat{l}$ with the highest score, i.e., $\hat{l} = argmax_{j\in[m],k\in[K_j]}(p_{jk} + \lambda_j^t)\hat{c}_{tjk}$.

  Step 3: (Exploitation-Exploration trade-off)

    With probability $\frac{1}{H+\gamma_t(\hat{f}_{t\hat{l}}-\hat{f}_{tjk})}$ to randomly display the other item jk. Otherwise, choose the empirical optimal item $\hat{l}$.

  Step 4: Observe outcomes and update dual variables via projected subgradient descent, $\lambda_j^{t+1} = \text{Proj}_{[0,\beta_j]}\big(\lambda_j^t - \eta(\sum_{k=1}^{K_j} \mathbb{I}_{\text{item } jk \text{ is purchased}} - \alpha_j/T)\big)$.

  Step 5: Update the prediction model $M^{t+1} = \mathcal{L}(M^t)$.

end for

Figure: Proposed Algorithm

Outline

Datassets

Done on two datasets:

1. Proprietary Dataset: 600 million sessions
2. Electronics Event History (EVS) Dataset: publicly available and includes 36,966 items from 999 brands and feedback data of 490,399 user sessions.

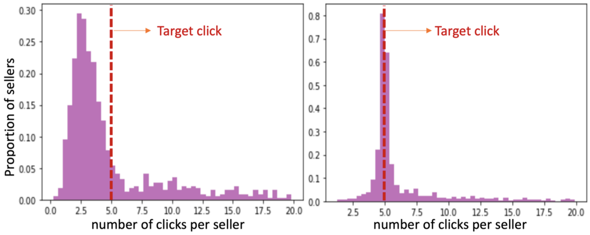

Figure: Shift in seller clicks after execution of our proposed algorithm

Results

Panel A: Relative Performance in Proprietary Dataset				
β/p	0.1	0.2	0.5	1.0
GMV (Revenue) - Relative change	0.13%	-0.21%	-0.46%	-0.84%
# of sellers with >5 clicks per T sessions	11%	17%	23%	28%

Panel B: Performance in EVS Dataset				
β/p	0.1	0.2	0.5	1.0
GMV (Revenue) - Relative change	0.00%	-0.04%	-3.71%	-8.73%
# of brands achieving the target - Alg1	29	32	36	35
Mean # of brands achieving the target - benchmark	16.4	16.4	16.4	16.4

Key Results:

- Significant right shift of total click distribution of sellers
- We can increase the fairness outcome and revenue at the same time.

Thank you for your attention!