ADKDD Aug 2021

Relevance Constrained Re-ranking in Sponsored Listing Recommendations

Zhen Ge, Wei Zhou, Jesse Lute, Adam Ilardi

ebay

Background

Sponsored Listing

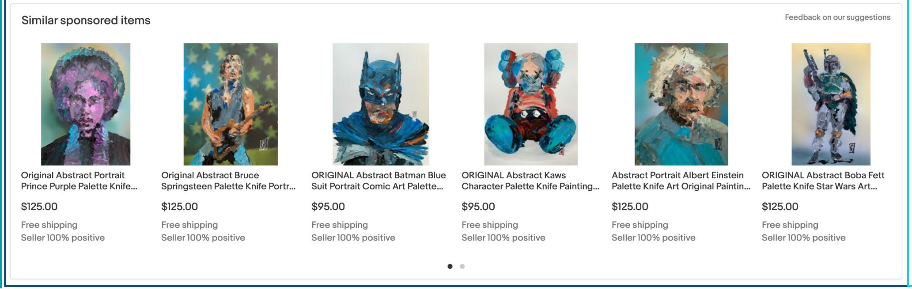

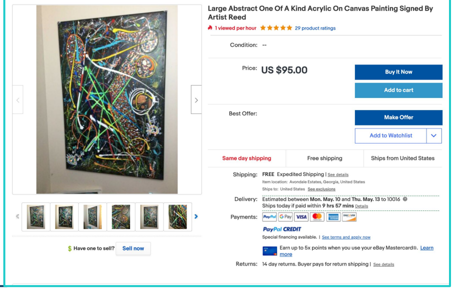

similar item recommendation

Seed item

features (e.g. title, price, item attributes, user preference)

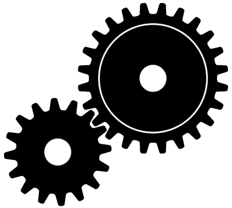

GBT model

purchase probability
p

$$score = \overbrace{p \cdot c}^{\text{Organic revenue}} + w \cdot \overbrace{p \cdot b}^{\text{Ad revenue}}$$

p : Purchase probability

c : Selling Cost

w : Ad Revenue Weight

b : Ad Bid Rate

Background

Revenue Relevance Trade-off

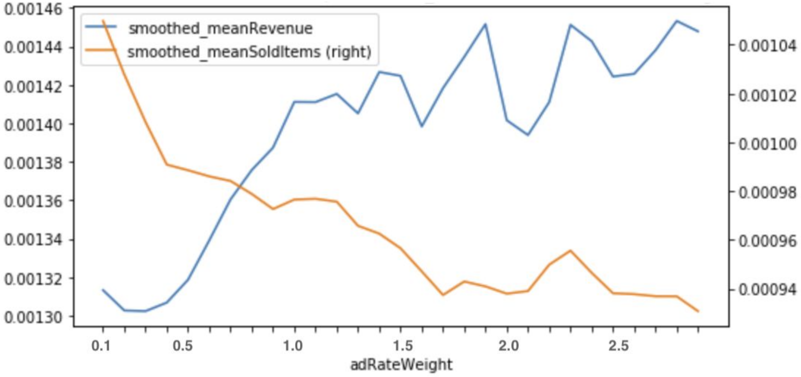

Figure 1. Average Ad revenue & Average Item Sold Count vs. Ad Revenue Weight

Many similar ones

Few similar ones

$$score = p \cdot c + \underline{w} \cdot p \cdot b \quad \rightarrow \quad score = p \cdot c + \underline{f(\delta)} \cdot p \cdot b$$

ebay

Dynamic Ad Revenue Weight Optimization (DARWO)

Kullback–Leibler divergence constraint

Purchase probability ranking

P $\quad item\_ranking = \{r_1 = 0.5, r_2 = 0.45, r_3 = 0.4, r_4 = 0.3, r_5 = 0.2, r_6 = 0.15\}$

r_i is the ***i*** th ranking position

Final ranking given different *w*

$$score = p \cdot c + f(\delta) \cdot p \cdot b$$

Q_{w1} $\quad item\_ranking_{w1} = \{r_1 = 0.45, r_2 = 0.5, r_3 = 0.3, r_4 = 0.4, r_5 = 0.15, r_6 = 0.1\}$

Q_{w2} $\quad item\_ranking_{w2} = \{r_1 = 0.1, r_2 = 0.5, r_3 = 0.45, r_4 = 0.3, r_5 = 0.4, r_6 = 0.2\}$

$$D_{kl}(P||Q_{w1}) = 0.021$$ ★

$$D_{kl}(P||Q_{w2}) = 0.258$$

Dynamic Ad Revenue Weight Optimization (DARWO)

Greedy Optimization

$$Q^* = \arg\max_{w \in W} f_{rev}(Q_w)$$

$$\text{s.t.} D_{kl}(P||Q_w) \leq \theta_{KL}$$

θ_{KL} is a constant. $f_{rev} : \mathbb{R}^k \rightarrow \mathbb{R}$ is the estimated ad revenue for ranking Q_w.

Distribution Q_w represents the ranked list X_w

$$X^w = [x_1^w, x_2^w, \ldots, x_k^w]$$

$$f_{rev}(Q_w) = f_{rev}(X^w)$$

$$= \sum_{r=1}^{k} a_r^w \cdot v_r$$

a_w is the ad revenue for the item at slot r. v_r is the unbiased click through rate for slot r.

Dynamic Ad Revenue Weight Optimization (DARWO)

Regression

What if we don't want to search for the w, can we estimate it?

$$D_{kl}(w) = f(\text{recall set statistics, context}, w)$$

Recall set statistics: recall set size, and the maximum, minimum, mean, median, and standard deviation of all recall set item's price, PTR score, and bid rate

Ordinary Least Squares Regression (OLS)

$$R^2 = 0.43$$

GBT Regression

$$\underline{D_{kl}(w)} = f(\text{recall set statistics, context}, \underline{w})$$

Experiments

Part - 1

DESIGN

- Control (production): Fixed ad revenue weight at 0.25 for all impressions.
- Treatment 1: OLS based dynamic ad revenue weight (DARWO) variant.
- Treatment 2: Fixed ad revenue weight at 1.75. 1.75 is selected because it's the median value of predicted ad revenue weight from the OLS DARWO variant.

RESULT

Table 1: OLS DARWO vs. Fixed ad revenue weight by marketplace

		Ad Revenue	Purchase Count
US	treatment 1	**+3.81%**[1]	**-4.05%**
	treatment 2	**+5.33%**	**-5.07%**
UK	treatment 1	**+6.89%**	**-4.11%**
	treatment 2	**+5.81%**	**-6.55%**
AU	treatment 1	**+7.10%**	-1.97%
	treatment 2	**+8.30%**	**-3.11%**
DE	treatment 1	**+6.44%**	**-3.44%**
	treatment 2	**+5.38%**	**-4.68%**

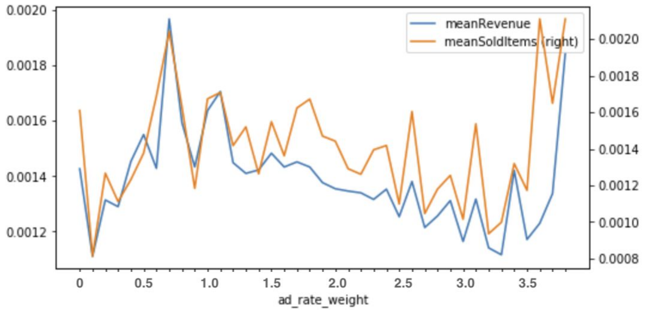

Figure 2. Average Ad revenue & Average Item Sold Count vs. Ad Revenue Weight (DARWO)

Ad revenue and purchase count have a pearson correlation of 0.724 $(p < 10^{-6})$

[1] bold numbers are significant with p<0.1

Experiments

Part - 2

DESIGN

- Control: OLS DARWO variant
- Treatment 1: GBT DARWO variant
- Treatment 2: Greedy optimization DARWO variant

RESULT

Table 2: GBT DARWO vs. Greedy DARWO by marketplace

		Ad Revenue	Purchase Count
US	treatment 1	**+7.72%**	**+5.45%**
	treatment 2	**+8.70%**	**+8.53%**
UK	treatment 1	**+7.00%**	0.90%
	treatment 2	**+4.13%**	**+4.91%**
AU	treatment 1	**+7.45%**	+2.49%
	treatment 2	**+8.50%**	**+6.78%**
DE	treatment 1	**+4.67%**	+1.33%
	treatment 2	+2.40%	**+3.57%**

Experiments

Summary

Compounded revenue and purchase changes (p<0.1):

GBT: Revenue lift: **12.6%**
Purchase lift: -1.8%

Greedy: Revenue lift: **11.0%**
Purchase lift: **2.5%**

Table 3: Offline Purchase Ranking Comparison: Production, GBT DARWO and Greedy DARWO

	Mean Reciprocal Rank	NDCG@6	NDCG@12
Production	0.508	0.567	0.615
GBT DARWO	0.480	0.544	0.593
Greedy DARWO	0.516	0.576	0.620

Conclusion

Effective

- KL divergence can be used as a quality measurement for a re-ranked list;
- Controlling global standard by adjusting local ranked list's relevance individually

Can be estimated

KL divergence can be estimated through local inventory based features.

Easy to implement

This ad hoc re-ranking stage is completely independent of the previous ranking or conversion stages.

Thank you!

Questions? zhge@ebay.com

Zhen Ge, Wei Zhou, Jesse Lute, Adam Ilardi