From Producer Success to Retention: a New Role of Search and Recommendation Systems on Marketplaces

Viet Ha-Thuc
Facebook Inc.
Menlo Park, CA, USA
vhathuc@fb.com

Matthew Wood
Facebook Inc.
Menlo Park, CA, USA
woodmd@fb.com

Yunli Liu
Facebook Inc.
Menlo Park, CA, USA
yunliliu@fb.com

Jagadeesan Sundaresan
Facebook Inc.
Menlo Park, CA, USA
jagadeesan@fb.com

ABSTRACT

In online marketplaces, an increasing number of producers depend on search and recommendation systems to connect them with consumers to make a living. In this talk, we discuss how these systems will need to evolve from the traditional formulations by incorporating the producer value into their objectives. Jointly optimizing the ranking functions behind these systems on both consumer and producer values is a new direction and raises many technical challenges. To overcome these, we lay out an end-to-end solution and present the results of applying this solution on Facebook Marketplace.

ACM Reference Format:
Viet Ha-Thuc, Matthew Wood, Yunli Liu, and Jagadeesan Sundaresan. 2021. From Producer Success to Retention: a New Role of Search and Recommendation Systems on Marketplaces. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21), July 11–15, 2021, Virtual Event, Canada.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3404835.3464924

1 INTRODUCTION

With the rise of the creator and gig economies and the popularity of online marketplaces fostering them in recent years, media (from short-form entertaining videos to yoga lessons to political blogs), products (from food to apparels to electronics) and services (such as, ride hail, food delivery and rental) are created, sold and provided by a large number of creators, sellers and providers. At the heart of these marketplaces, search engines and recommendation systems are the primary distribution channels connecting consumers and producers [1]. These systems not only help the consumers discover relevant and high-quality content, products or services but also play a critical role in the producer's success (or failure) and long-term retention (i.e., continuing to produce on these marketplaces).

[1]In this proposal, the term *producer* generally refers to creators, sellers and service providers while *consumer* broadly indicates users on the other side of the marketplaces, like viewers, buyers and guests. When discussing Facebook Marketplace specifically, we use *sellers* and *buyers*.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

ACM ISBN 978-1-4503-8037-9/21/07.
https://doi.org/10.1145/3404835.3464924

Given the important role in the producer side, search and recommendation systems need to evolve from the standard formulation of retrieving the most relevant documents matching a query and finding the items maximizing the likelihood that a consumer takes an action. In addition, these systems should also take into account the utility based on the producers' experience and future behavior.

Optimizing the ranking functions (rankers) behind these systems on such objectives, however, is a new direction and raises many technical challenges. It is difficult to measure the causal effect of ranking changes on producers. If we run a standard seller-side A/B test that exposes to a small percentage of producers, what we observe in the test would be significantly different from when the treatment is launched to all producers. Another problem is how to model the effect of up-ranking or down-ranking an individual item on the producer's experience and long-term behaviors.

In this proposal, we describe our approach to evolving our ranking systems on Facebook Marketplace, an online marketplace where buyers and sellers connect to buy and sell commercial products. Deviating from the typical approach of optimizing for consumer-centric metrics like click-through-rate or purchase rate, our rankers are jointly optimized for both consumer- and producer-side objectives. In particular, we describe an end-to-end solution from conducting A/B testing on the producer-side to incorporating producer-side success in ranking objectives and how the success will translate into retention in the long term. While we use Facebook Marketplace as an illustrative example, the solution is generalizable to other two-sided marketplaces in the industry.

2 A/B TESTING ON PRODUCER SIDE

To demonstrate the challenge of producer-side A/B testing, let us consider a simple test in which the items from the *new* producers in the test group are boosted. Because of the boost, these new producers get a better experience compared to the ones in the control group. However, when the boosting is ramped up to 100% (i.e., to all new producers), the items in the original test group will be ranked lower because significantly more items are getting boosted now. Thus, the impact when experimented at 1% is artificially inflated, and the producer-side A/B test result is incorrect.

The root cause of this is a violation of "stable unit treatment value assumption" (SUTVA) since the experience of the producers in either group also depends on the ranking model applied to the other producers outside the group. Given this, we develop a novel producer-side A/B testing framework based on a counterfactual property: the items in the treatment are ranked at where they would be if the treatment is ramped to 100% of the producers. Similarly, the items in the control are ranked where they would be if the status quo is applied to all. Thus, the difference between treatment and control is independent of what applies to the rest, i.e., satisfies SUTVA. The readers can find details in our previous work [2].

3 OPTIMIZE FOR PRODUCER SUCCESS

Commercial ranking systems are usually trained from historical engagement data. As a result, they tend to concentrate on the items from a small set of producers and do not give a fair distribution of the traffic to the majority of other producers. Thus, these producers with less distribution usually do not have the success they expect [2]. Furthermore, this could lead to increasing the inequality among user groups [1]. While the multi-armed bandit approach balancing exploit and explore is a well known technique to alleviate the cold-start problem [3], this was designed to optimize for consumer-centric objectives in the long run, rather than explicitly considering the producer-side objectives, e.g., producer success and fairness.

To optimize for seller success on Facebook Marketplace, we consider the objective of the number of unique sellers getting at least one call-to-action (*cta*), e.g., a message or an offer from a buyer, on a given day. The rationale of this objective is to distribute the success to as many sellers as possible. Specifically, given a ranking by the first-stage ranker purely optimized on the buyer side, we up-rank N items in top-K of the ranking whose sellers have not received any cta in the *past day* (i.e., unsuccessful) and having the highest *item values*. The values are the weighted sum of buyer value (the first-stage ranking score) and seller value as in Equation 1.

$$item\_value = buyer\_value + \lambda * seller\_value \tag{1}$$

$$\begin{aligned} seller\_value &= p(cta\_1d|imp) - p(cta\_1d|no\_imp) \\ &= 1 - p(no\_cta|imp) * p(no\_cta\_1d|no\_imp) \\ &\quad - p(cta\_1d|no\_imp) \\ &= 1 - [1 - p(cta|imp)] * p(no\_cta\_1d|no\_imp) \\ &\quad - [1 - p(no\_cta\_1d|no\_imp)] \\ &= p(cta|imp) * p(no\_cta\_1d|no\_imp) \end{aligned} \tag{2}$$

The seller value is the incremental change to the probability that the seller will get a *cta* in the *next day* when receiving an impression (*imp*) due to the boost in the current session. As mathematically derived in Equation 2, this incremental value turns out to be the product of the probability that the item gets a cta in the current session given the boosted impression and the probability that the seller will not succeed beyond this session (in the next 24 hours). Note that in the extreme cases when the item is non-relevant or low-quality, i.e., $p(cta|imp) \approx 0$, or the seller will certainly succeed anyway, i.e., $p(no\_cta\_1d|no\_imp) \approx 0$, the incremental value is zero, which intuitively makes sense. The incremental value formulation can also be generalized at various levels on the seller side. For instance, we can consider the incremental value of boosting on the *item success* defined as having a minimum number of cta or clicks or we can shorten or extend the time window over which success is measured. Another objective is the number of *seller groups* having cta rate above a threshold. Including this in the item value model would increase the number of groups (e.g., by different attributes like gender or race as suggested in [1]) achieving the minimum bar of success, thus reduce the gap between these groups.

4 FROM SUCCESS TO RETENTION

To empirically evaluate the optimization strategy, we run both buyer- and seller-side A/B tests. The buyer-side test is simply a standard A/B test which measures the short-term impact on buyer-side experience. The seller-side test is run on the counterfactual framework described in Section 2. Only the items of the sellers in the test group could get up-ranked, but they are positioned as if the reranking strategy in Section 3 was applied to all items. At the same time, the items of the sellers in the control group are ranked at the positions as if no reranking was applied at all. The empirical result from the seller-side A/B test shows a significant improvement of 10% on seller success (i.e., the number of unique sellers having at least one cta on a given day) while the buyer-side test reveals neutral results on short-term buyer experience.

Distributing success to more unique sellers on a marketplace would clearly improve sellers' experience and be an important objective by itself. However, we also want to examine the impact of the seller experience on seller retention and the long-term growth of the marketplace. To measure these, we run a seller-side back-test over more than six months as of this writing. In the backtest, we observe that the seller success improvement translates to a significant increase in seller retention and the number of posted listings. Comparing the number of transactions received by the sellers in the two groups, the test group is receiving an increasing number of transactions over time relative to the control group after initially having a similar number of transactions. The increase in transactions is interesting because by incorporating the seller-side objective, we actually deviate from the original rankings purely optimized for transactions. The reason behind this is two fold: (i) the increase in seller retention and listings (ii) the exploration effect of increasing the distribution for items and sellers that normally rarely get exposed.

5 CONCLUSIONS

With the rise of the creator and gig economies, search and recommendation systems play an increasingly important role for producers and the society at large. In this proposal, we lay out a vision in which these systems will evolve by incorporating the impact on the producer side into the objective function. We also present our end-to-end solution to the technical challenges emerging with this new direction as well as the initial results on Facebook Marketplace.

REFERENCES

[1] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *Proceedings of the 25th ACM SIGKDD*. 2221–2231.

[2] Viet Ha-Thuc, Avishek Dutta, Ren Mao, Matthew Wood, and Yunli Liu. 2020. A Counterfactual Framework for Seller-Side A/B Testing on Marketplaces. In *Proceedings of the 43rd International ACM SIGIR*. ACM, 2288–2296.

[3] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 661–670.