

Towards Quality Ad Selection: A Model-based Approach to Performance Filtering

Mandar S. Chaudhary, Javad Nejati, Mahmuda Rahman, Gajanan Adalinge and Abraham Bagherjeiran

eBay Inc.

Abstract

Major e-commerce platforms display advertisements (ads) on search results page through a three-phase approach: retrieval, selection, and ranking. The effectiveness of ad selection algorithms directly impacts the quality of candidates available for ranking on the search results page, thereby heavily influencing buyer engagement and advertiser performance. Ad selection algorithms need to efficiently prune the set of retrieved ads due to limited latency and resource capacity for the ranking stage. In order to pass better quality ads to the ranking stage, we propose a two-stage ad selection algorithm which filters ads that degrade buyer experience on search engine results page (SERP). Our algorithm is based on an ad performance filter which presents a novel approach for identifying and filtering low-performing ads. First, we formulate eligibility criteria to select ads with sufficient exposure on SERPs, and second, we leverage this criteria to identify ads with low buyer engagement. We demonstrate the efficacy of our approach by conducting online experiments using the A/B test framework of a major e-commerce platform. Results show that the proposed two-stage ad performance filter significantly improves Click-Through-Rate, and it highlights the impact of developing a well designed ad selection filter to enhance buyer and advertiser experiences.

Keywords

Sponsored products quality, Low-performing ads filter, label generation,

1. Introduction

In the rapidly evolving landscape of e-commerce, the effectiveness of sponsored search results is paramount to both advertisers and platform providers. Sponsored ads serve as a crucial revenue stream for e-commerce companies, while simultaneously influencing consumer purchasing decisions. These results, often displayed prominently on search engine results page (SERP), are a significant revenue driver for e-commerce platforms and also a critical tool for advertisers to reach potential customers. The quality of these sponsored listings impacts buyer satisfaction and engagement, making it imperative for platforms to ensure that the ads presented are relevant and engaging to the user.

Major e-commerce platforms follow a well-known multi-stage approach shown in Figure 1 to display ads on SERP [1, 2, 3]. This approach primarily consists of the ad retrieval, selection, and ranking stages that aim to maximize buyer engagement as well as the total value of ads displayed on the SERP. The retrieval phase matches the buyer query with the sponsored ads keywords specified by the advertiser and usually fetches millions of ad listings. The ad selection algorithm should efficiently trim this ad space and maintain quality listings for the final ranking stage. However, achieving this balance poses a complex challenge. On the one hand, e-commerce platforms must prioritize buyer experience by showcasing high-quality ads while on the other hand advertisers are keen on maximizing the exposure of their listings to enhance brand visibility and drive sales. This duality necessitates a sophisticated approach to ad selection, where the interests of both users and advertisers are harmonized. Additionally, this approach must also address the cold-start problem since advertisers daily onboard thousands of new ad listings in the sponsored ads program.

The goal of ad selection algorithms is to provide sufficient opportunities for each ad listing to surface on SERP, and once they have accumulated sufficient exposure, efficiently filter the ads that are

ECOM'25: SIGIR Workshop on eCommerce, Jul 17, 2025, Padua, Italy

✉ manchaudhary@ebay.com (M. S. Chaudhary); jnejati@ebay.com (J. Nejati); mahrahman@ebay.com (M. Rahman); gadalinge@ebay.com (G. Adalinge); abagherjeiran@ebay.com (A. Bagherjeiran)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Figure 1: Multi-stage sponsored ads funnel comprising of three primary stages: retrieval, selecting, and ranking.

unlikely to receive buyer engagement. It is very difficult to quantify these *sufficient opportunities* for each ad listing because the true performance of an ad is not known until it has surfaced on SERP to accumulate a few hundred or even thousands of impressions. Furthermore, allowing low-quality ads on premium placements on SERP can dilute the overall quality of search results, leading to suboptimal user experiences and reduced revenue potential. Low-quality ads, characterized by poor engagement metrics such as low click-through rates (CTR), can undermine the effectiveness of sponsored search results.

In this work, we introduce an innovative ad performance filter designed to filter low-performing ads that negatively impact buyer experience. The contributions of our proposed work can be summarized as follows,

- Construct an ad cohort to group similar ads, enabling comparative analysis of ad performance within each cohort based on buyer engagement metrics.
- We formulate a novel approach to quantify the *sufficient opportunities* provided to each ad in a cohort, and identify eligible ads with measurable performance. We build on the definition of eligible ads to propose our novel approach for identifying and filtering low-performing ads in each cohort.
- We demonstrate the efficacy of our approach by performing online experiments using A/B test framework of a major e-commerce platform, and results indicate a significant improvement in CTR by filtering ad impressions that do not receive clicks.

The rest of the paper is organized as follows. We present a survey of state-of-art prior works in Section 2 followed by details of the proposed method in Section 3, and experimental results in Section 4. In the end, we present discussions and future work in Section 5.

2. Related Work

Improving ad selection with ad quality filters is not a well-studied problem, since most of the prior work is primarily based on three different approaches: optimizing query-keyword matching [4, 5, 6, 7], early stage ranking with multi-task framework [8, 9, 10], and reinforcement learning to improve ad selection policy [1, 11, 12]. Although our proposed work does not directly fit into these approaches, it has some overlap with multi-task frameworks that learn ad quality signals to improve buyer engagement.

Prior works based on multi-task learning frameworks have incorporated different kinds of ad quality signals to select the best set of ads for the final ranking stage [3, 13, 14]. These signals are based on explicit buyer feedback or derived from buyer feedback using feature engineering to measure ad quality. For instance in [14] a multi-task learning framework is developed to learn two different quality signals that measure the ad dismiss rate and post-ad-click experience respectively. These signals contribute towards the improvement of the final CTR prediction task. Another prior work in [3] learned two ad quality events namely cross-out rate, which measures the number of times a user explicitly does not want to see an ad, and survey assessment, which records user ads rating with higher rating indicating




Figure 2: The ad quality filter pipeline operates in two stages. In the first stage, a Beta-Binomial Bayesian model serves as a guardrail to filter out low-quality ads across all cohorts. Ads that pass this initial filter proceed to the second stage, where they are evaluated for eligibility using the ad eligibility filter. Eligible ads are then further assessed using the low-performing ad filter. Ads deemed ineligible in the second stage may still be allowed to appear on the search results page (SERP), provided they have satisfied the criteria of the first-stage filter.

better ad quality. Both the prior works used explicit buyer feedback to create the ad quality labels. On the other hand, the work in [13] focused on predicting the time interval for an ad to be discontinued for user exposure where the labels for time interval were generated heuristically from the user feedback data.

The aforementioned prior works deliver promising results with the state-of-art deep learning based multi-task methods. While effective, these methods are difficult to scale in ad selection systems with limited latency and resource model capacity of the final ranking stage. We posit that similar to these works, the valuable buyer engagement feedback from the impressed ads can be harnessed to develop meaningful ad quality signals for improving the selection process with simple and efficient approaches. To the best of our knowledge, this is the first work to propose a novel light-weight two-stage filter with an intuitive definition to quantify ads performance from low buyer engagement.

3. Method

In this section we describe in detail our two-stage ad performance filter with the Beta Binomial Bayesian model as the first stage and the low-performance ad model as the second stage. Figure 2 shows the high-level approach of the two-stage filter pipeline. The first stage filter guardrails buyer experience from ads which do not meet the minimum quality criteria while the second stage filters identify ads with measurable performance on SERP and filters poor quality ads. As part of the second stage filter we describe our approach for defining an ad cohort and two label generation methods. The first label generation method develops an ad eligibility criteria to identify ads with sufficient impressions on SERP. These ads are considered to have received sufficient exposure and buyer interaction to reliably estimate their performance. The second label generation approach determines whether an eligible listing is low-performing based on its performance in their respective cohort.

3.1. Leveraging Beta-binomial Bayesian Model

We utilize the well-known Beta-Binomial Bayesian model to estimate posterior mean scores of Click-Through-Rate and Purchase-Through-Rate respectively for each ad listing [15]. These models have been widely used for measuring item’s performance by smoothing their quality scores with priors to address the cold-start problem [16, 17, 18]. The posterior mean for computing Click-Through-Rate of an ad listing a_i from exposure of impressions and clicks on country k can be estimated from the following beta distribution,

$$\text{Beta} \left(\frac{\zeta_{ik} + \alpha_{CTR}^k}{\phi_{ik} - \zeta_{ik} + \beta_{CTR}^k} \right) \quad (1)$$

where α_{CTR}^k and β_{CTR}^k are the priors of the beta distribution for country k , ζ_{ik} is the total click count and ϕ_{ik} is the total impression count for ad listing a_i . Similarly, we also compute the Purchase-Through-Rate (PTR) score of a_i by estimating beta binomial priors using the total sale count, θ_{ik} , and total impression count, ϕ_{ik} . The priors are useful for providing smoothed scores for new ad listings that do not have any exposure on the SERP thereby addressing the cold-start problem. Finally, the smoothed CTR and PTR scores of an ad listing are estimated as the mean of their corresponding beta distributions,

$$\gamma_{ik} = \frac{\zeta_{ik} + \alpha_{CTR}^k}{\phi_{ik} + \alpha_{CTR}^k + \beta_{CTR}^k} \quad (2)$$

$$\omega_{ik} = \frac{\theta_{ik} + \alpha_{PTR}^k}{\phi_{ik} + \alpha_{PTR}^k + \beta_{PTR}^k} \quad (3)$$

The beta binomial priors are estimated for each country using the well-known Method-of-Moments (MoM) and Maximum Likelihood Estimation (MLE) approaches. The priors are initialized using the estimates produced by MoM and refined iteratively using MLE until the difference in prior values between successive iterations is less than $1e-3$. We apply a threshold each on the smoothed CTR and PTR scores to filter listings that do not meet the minimum quality bar for CTR and PTR. All listings which pass the first-round are passed onto the next steps.

3.2. Identifying Ad Cohort

Our approach groups similar ad listings into cohorts to accurately measure ad performance. This is important because ads in different cohorts can display vastly different performance metrics. For example, an ad for search query *iphone case* might have a higher CTR than one for *sectional couch*, it could still be considered low-performing when compared to other ads within its own category. Therefore, we evaluate ad eligibility and click-through-rate performance within the context of their specific cohorts, allowing for more meaningful comparisons.

Major e-commerce platforms provide an option for sellers to list an item under a suitable business vertical and category based on the item's functionality. It creates a natural grouping of similar items and provides a clear organization of the inventory. Informally, an ad cohort is defined as a group of ads with sufficient listing count that exhibit a certain level of similarity in terms of semantics and functionality.

We consider each combination of country and listing category as a cohort to group similar ads. This combination can be extended to include different levels of granularity such as price, item aspects and query. However, higher granularity increases the data sparsity, which results in fewer ads in each cohort. In this work, we define an ad cohort by the $(country, category)$ combination as it provides a reasonable trade-off between data sparsity and ad group similarity. Next, we measure the click-through-rate performance of each ad cohort to establish the ad eligibility criteria. Specifically, we calculate the impressions and clicks for each cohort using a rank-discounted exponential decayed function.

Formally, we denote $S = \{s_i \mid i = 1, \dots, n\}$ as the set of n distinct cohorts, and the CTR score for each cohort is stored in $C = \{c_i \mid i = 1, \dots, n\}$.

$$c_i = \frac{click_i^t}{imp_i^t} \quad (4)$$

$$imp_i^t = \sum_{p,q \in T_i} (imp_i^p * \lambda^{\delta_{p,q}} + imp_i^q) \quad (5)$$

$$click_i^t = \sum_{r,s \in T_C} (click_i^r * \lambda^{\delta_{r,s}} + click_i^s) \quad (6)$$

where:

- $click_i^t$ and imp_i^t refer to the decayed click count and decayed impression count with rank discount at current timestamp t observed across T_I and T_C series of timestamps respectively, and λ is the decay factor. For simplicity, we will use the notation imp_i and $click_i$ for the decayed click count and impression count respectively.
- The timestamps $p, q \in T_I$ represent the consecutive timestamps when the listing received impressions and $\delta_{p,q}$ represents the time elapsed between p and q timestamps. Similarly, $r, s \in T_C$ where $r < s$ represents the timestamps with consecutive clicks.

3.3. Defining Ad Eligibility

Once the ad cohort is established, we assess an ad's eligibility to be categorized as low performing by evaluating its visibility on search engine results pages (SERPs) which is defined in terms of its impression count. The impression of each ad shown on SERP belongs to one of the cohorts $s_k \in S$, and we determine an ad is eligible if it has accumulated sufficient impression count under the given cohort s_k . Let $\mathcal{A} = \{a_{ij} \mid i = 1, \dots, m; j = 1, \dots, n\}$ denote the set of ads such that i^{th} ad listing, a_{ij} , has received at least one impression under cohort s_j . Correspondingly, $\mathcal{I} = \{imp_{ij} \mid i = 1, \dots, m; j = 1, \dots, n\}$ contains the rank-discounted exponential decayed impression count of the ad listings in their cohorts. Finally, the criteria to determine an ad is eligible under the given cohort is as follows,

$$imp_{ij} > \tau_j \quad (7)$$

$$\tau_j = \frac{1}{c_j} \quad (8)$$

where τ_j is the inverse of c_j and measures the average number of impressions per click for cohort s_j . The cohort's click-through-rate c_j score is used to estimate a threshold for the number of impressions each ad should be provided before their performance can be reliably judged. Ads with fewer than τ_j impressions are not considered eligible as they have not received sufficient exposure to the buyer on SERP. Such ads also referred to as non-eligible ads are allowed to pass through the second-stage filter and have the opportunity to surface as impression as long as they pass the first-round filter. The non-eligible group of ads comprises newly listed ads with no historical data, as well as existing listings that have not received any recent exposure on the SERP within a defined time window. As a result, our label generation process inherently handles the cold-start problem by treating new ad listings as not eligible for filtering by our two-stage filter. We define the time window and additional experiment details in Section 4.1.

Intuitively, the ad eligibility criterion determines τ_j opportunities in the form of impressions for each ad in a given cohort before their click-through-rate performance can be considered. Note that the same item can belong to more than one cohort and it can be labeled as eligible and low-performing in one cohort but not in another or it can be eligible and low-performing in all cohorts.

3.4. Detecting Low-performing Ads

In this step, we consider all eligible ads and determine if they are low-performing based on their quality scores. Consider the set of ads, $\tilde{\mathcal{A}} \subseteq \mathcal{A}$, which have received sufficient impression count with respect to their cohort and the set $\tilde{\mathcal{P}} \subseteq \mathcal{P}$ denote their quality scores. An ad is considered low-performing in a cohort if its quality score is lower than a threshold that is computed from the quality score distribution of the ads in the cohort. For a given cohort s_j , we calculate the lower l^{th} percentile of the quality score distribution from all eligible ads and set it as the threshold for identifying low-performing ads. An ad with quality score lower than this threshold is labeled as low-performing ad and the rest are labeled as not-low-performing. Below we present the equation for calculating this threshold \tilde{p}_j from the quality scores of all eligible items, \tilde{p}_j , in cohort s_j .

$$\tilde{p}_j = \{p_{kj} \mid p_{kj} \in \tilde{\mathcal{P}} \iff imp_{kj} > \tau_j\} \quad (9)$$

$$\tilde{p}_{lj} = \text{Percentile}(\tilde{p}_j, l) \quad (10)$$

Finally, we formulate the conditions for identifying an ad a_{ij} as a low-performing ad as follows,

$$(imp_{ij} > \tau_j) \text{ and } (p_{ij} < \tilde{p}_{lj}) \quad (11)$$

An ad in the s_j cohort with at least τ_j impression count and a quality score lower than \tilde{p}_{lj} will be labeled as an eligible and low-performing ad, since it has received sufficient exposure to the buyer on SERPs, and with buyer feedback incorporated into its quality score it has been observed to be among the worst performing listings in its cohort.

3.5. Model Training

We train two classification models to predict ad eligibility and low-performing ad respectively. The response variable in the ad eligibility prediction model is binary valued where a value 1 indicates the ad is eligible and 0 indicates an ineligible ad. Similarly, for low-performing ad model, the target variable is also binary valued with a value of 1 indicating a low-performing ad and 0 indicating the ad is not-low-performing. The predictor set for both models included a combination of content-based and historical features. We trained both classification models using the XGBoost algorithm [19] and logistic loss function by varying the number of trees in the model in the range of [1, 50].

The eligible ad model was trained by adding sample weights to the loss function. The sample weights were set to the rank-discounted decayed impression count of each ad listing thereby penalizing the model if it incorrectly predicts ad listings with high impression count. No such sample weights were applied for training the low-performance ad model.

$$\text{Loss}(y, \hat{y}) = -\frac{1}{N} \sum_{i=1}^N w_i [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (12)$$

where $w_i = \sum_k imp_{ik}$ is the rank weighted decayed impression count of the i^{th} ad listing.

The output of the ad eligibility prediction model is used to determine whether an ad should be further examined for low performance or whether it should be allowed more opportunities. As shown in Figure 2 if an ad has a low probability of being eligible it will have an opportunity to appear as impression on SERP whereas an ad with high probability of being eligible will receive another prediction score from the low-performing ad model, and the ad will be filtered if its probability score of being low-performing is higher than a threshold.

4. Experiments

In this section, we evaluate our proposed two-stage filter by describing the offline experiment setup and demonstrate the effectiveness with experiments on real-world traffic using A/B test framework.

4.1. Offline experiments

We sampled logs of sponsored ad listings on SERP of a major e-commerce platform over a period of three months. The dataset comprised of 3.5 billion impressions and each ad listing was labeled as eligible and low-performing based on the look-back period of one month. The training and validation datasets were generated by splitting the data by time where all ad listings with impressions before timestamp t were included in the training set and those after timestamp t were used for validation set.

We evaluate the effectiveness of the eligible ad model in distinguishing between eligible and non-eligible ads. The eligible category encompasses both low-performing and not-low-performing ads. Our analysis focuses on the change in item filter rate for low-performing ads, not-low-performing ads, and non-eligible ads as we vary the thresholds used by the eligible ad filter and the low-performance ad filter.




Figure 3: A comparison of the item filter rate change for different kinds of ads after passing the eligible ad model under three different threshold values. Figure 3a shows the item filter rate without using the prediction scores from the eligible ad model. Figures 3b-3c show the low-performing ad model applied on the ads predicted as eligible by applying thresholds of 0.3 and 0.6 respectively.

Table 1

Online comparison of metrics between control and treatment groups. Stat-sig results are highlighted in bold.

| Metrics | Lift (%) |
|------------------------|---------------|
| Click-Through-Rate | +0.51% |
| Sale-Through-Rate | +1.03% |
| Total Impression Count | -0.55% |
| Total Click Count | -0.05% |
| Total Sale Count | +0.47% |

In Figure 3a, we calculate item filter rates for the three types of ads subject to a threshold of zero to pass the eligible ad filter and apply only the low-performance ad filter. Simulations with varying thresholds of low-performance prediction score show that a significant portion of non-eligible ads get filtered alongside low-performing ads. For instance, with a threshold of 0.6, around 40% of all non-eligible ads get filtered which would lead to great dissatisfaction among advertisers. The eligible ad model prevents this by passing only eligible ads to the next stage. This is evident from Figures 3b-3c where ads with a eligible score greater than 0.3 and 0.6 respectively are applied the low-performance filter. However, this comes at the cost of also reducing the total fraction of low-performing ads that can be filtered from 100% to less than 50%. Therefore, carefully tuning threshold for eligible ad model is a trade-off between precision and recall as higher threshold presents more precise results but lowers the fraction of low-performing ads that can be filtered.

As briefly discussed in Section 3.3, our ad eligibility label generation mechanism mitigates the cold-start problem. We found that the ad eligibility model only predicted an ad as eligible for receiving a low-performing ad prediction score once it had accumulated, on average, at least 25 impressions. This behavior underscores the model’s effectiveness in handling the cold-start challenge.

4.2. Online A/B Test Results

We perform an online experiment for two weeks using the A/B test framework to evaluate the effectiveness of the two-stage low-performance ad filter across four different channels including, desktop, mobile web, iOS, and Android. The A/B test traffic was equally distributed between the control and treatment groups with users randomly assigned to each group. The results of the A/B test experiment are presented in Table 1.

The experimental results demonstrate the effectiveness of the two-stage filter in reducing the number of ad impressions which do not receive a click on SERP. Particularly, we observe a statistically significant reduction of -0.55% in the total ads impression count without affecting the total click count while slightly

positive trending total sale count of +0.47%. As a result there is a significant increase of +0.51% in Click-Through-Rate and an increase of +1.03% in Sale-Through-Rate with confidence interval of $[-0.1\%, +2.16\%]$. The statistical significance was measured by a two-sided t-test with a p-value of 0.05. We also observed the fraction of impressions from low-performing ads dropped by -19.6% compared to the Control group. By reducing the number of impressions without losing clicks or sales, the two-stage ad performance filter was able to improve the buyer experience.

The promising results support the design of developing an efficient two-stage filter that does not require substantial infrastructure investment as deep learning methods. These findings validate the approach of creating an innovative ad selection filter that emphasizes creating meaningful ad quality signals to improve sponsored search buyer experience.

5. Discussions and Future Work

In this work, we developed a simple and intuitive novel approach for identifying ads with poor click-through-rate on SERP. The proposed two-stage filter quantifies the opportunities for each ad and presents two label generation strategies for classifiers to learn the patterns of eligible and low-performing ads. The approach is evaluated on real-world traffic with an A/B test to illustrate the efficacy of the two-stage filter by filtering impressions that do not lead to a click.

As part of our future work, we plan to improve this approach in a few different ways. The proposed label generation strategies do not take advantage of the entire inventory as the ad quality signals are measured only for the impressed ad listings. To address this drawback of selection bias, we plan to improve the approach by generating pseudo-labels for non-impressed ads so they can be included in generating ad quality signals as well as model training. For instance, pseudo-labels for non-impressed inventory can be obtained from the final CTR ranker. We also plan to refine the approach for grouping ads based on their ad cohort by including additional information such as embedding similarity scores of ads. The embeddings can be generated by including several additional signals such as seller id, price, image and aspects. Lastly, we plan to develop a similar model for ads with low conversion rates to further improve buyer and advertiser experience.

References

- [1] Y. Lian, Z. Chen, X. Pei, S. Li, Y. Wang, Y. Qiu, Z. Zhang, Z. Tao, L. Yuan, H. Guan, et al., Optimizing ad pruning of sponsored search with reinforcement learning, in: Companion Proceedings of the Web Conference 2021, 2021, pp. 123–127.
- [2] X. Ma, P. Wang, H. Zhao, S. Liu, C. Zhao, W. Lin, K.-C. Lee, J. Xu, B. Zheng, Towards a better tradeoff between effectiveness and efficiency in pre-ranking: A learnable feature selection based approach, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2021, pp. 2036–2040.
- [3] X. Wang, Q. Jin, S. Huang, M. Zhang, X. Liu, Z. Zhao, Y. Chen, Z. Zhang, J. Yang, E. Wen, et al., Towards the better ranking consistency: A multi-task learning framework for early stage ads ranking, arXiv preprint arXiv:2307.11096 (2023).
- [4] Q. Cui, F.-S. Bai, B. Gao, T.-Y. Liu, Global optimization for advertisement selection in sponsored search, *Journal of Computer Science and Technology* 30 (2015) 295–310.
- [5] H. Wang, Y. Liang, L. Fu, G.-R. Xue, Y. Yu, Efficient query expansion for advertisement search, in: Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval, 2009, pp. 51–58.
- [6] A. Broder, P. Ciccolo, E. Gabrilovich, V. Josifovski, D. Metzler, L. Riedel, J. Yuan, Online expansion of rare queries for sponsored search, in: Proceedings of the 18th international conference on World wide web, 2009, pp. 511–520.
- [7] Y. Choi, M. Fontoura, E. Gabrilovich, V. Josifovski, M. Mediano, B. Pang, Using landing pages for

- sponsored search ad selection, in: Proceedings of the 19th international conference on World wide web, 2010, pp. 251–260.
- [8] J. Ma, Z. Zhao, X. Yi, J. Chen, L. Hong, E. H. Chi, Modeling task relationships in multi-task learning with multi-gate mixture-of-experts, in: Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining, 2018, pp. 1930–1939.
 - [9] H. Tang, J. Liu, M. Zhao, X. Gong, Progressive layered extraction (ple): A novel multi-task learning (mtl) model for personalized recommendations, in: Proceedings of the 14th ACM conference on recommender systems, 2020, pp. 269–278.
 - [10] Z. Zhao, L. Hong, L. Wei, J. Chen, A. Nath, S. Andrews, A. Kumthekar, M. Sathiamoorthy, X. Yi, E. Chi, Recommending what video to watch next: a multitask ranking system, in: Proceedings of the 13th ACM conference on recommender systems, 2019, pp. 43–51.
 - [11] S. Han, R. Lakritz, H. Wu, Augmented two-stage bandit framework: Practical approaches for improved online ad selection (2024).
 - [12] Q. Shi, F. Xiao, D. Pickard, I. Chen, L. Chen, Deep neural network with linucb: A contextual bandit approach for personalized recommendation, in: Companion Proceedings of the ACM Web Conference 2023, 2023, pp. 778–782.
 - [13] S. Kitada, H. Iyatomi, Y. Seki, Ad creative discontinuation prediction with multi-modal multi-task neural survival networks, *Applied Sciences* 12 (2022). URL: <https://www.mdpi.com/2076-3417/12/7/3594>. doi:10.3390/app12073594.
 - [14] N. Ma, M. Ispir, Y. Li, Y. Yang, Z. Chen, D. Z. Cheng, L. Nie, K. Barman, An online multi-task learning framework for google feed ads auction models, in: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022, pp. 3477–3485.
 - [15] G. Casella, An introduction to empirical bayes data analysis, *The American Statistician* 39 (1985) 83–87.
 - [16] C. Han, P. Castells, P. Gupta, X. Xu, V. Salaka, Addressing cold start in product search via empirical bayes, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022, pp. 3141–3151.
 - [17] Q. Liu, A. Singh, J. Liu, C. Mu, Z. Yan, J. Pedersen, Long or short or both? an exploration on lookback time windows of behavioral features in product search ranking, in: Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom’24), 2024.
 - [18] P. Gupta, T. Dreossi, J. Bakus, Y.-H. Lin, V. Salaka, Treating cold start in product search by priors, in: Companion Proceedings of the Web Conference 2020, 2020, pp. 77–78.
 - [19] C. T. G. C. XGBoost, et al., A scalable tree boosting system, in: Proc 22nd ACM SIGKDD Int Conf Knowl Discov Data Min, 2016, pp. 785–794.