

Ranking Policy Learning via Marketplace Expected Value Estimation From Observational Data

Ehsan Ebrahimzadeh¹, Nikhil Monga¹, Hang Gao¹, Alex Cozzi¹ and Abraham Bagherjeiran^{1,*†}

¹*eBay Search Ranking and Monetization*

Abstract

We develop a decision making framework to cast the problem of learning a ranking policy for search or recommendation engines in a two-sided e-Commerce marketplace as an expected reward optimization problem using observational data. As a value allocation mechanism, the ranking policy allocates retrieved items to designated slots to maximize the user utility from the slotted items at any given stage of the shopping journey. The objective of this allocation can in turn be defined with respect to the underlying probabilistic user browsing model as the expected number of interaction events on presented items matching the user intent, given the ranking context. Recognizing the effect of ranking as an intervention action to inform user interactions with slotted items and the corresponding economic value of interaction events for the marketplace, we formulate the expected reward of the marketplace as the *collective value* from all presented ranking actions. The key element in this formulation is the notion of *context value distribution*, which signifies not only the attribution of value to ranking interventions within a session but also the distribution of marketplace reward across user sessions. We build empirical estimates for the expected reward of the marketplace from observational data that account for the heterogeneity of economic value across session contexts as well as the distribution shifts in learning from observational user activity data. The ranking policy can then be trained by optimizing the empirical expected reward estimates via standard Bayesian inference techniques. We discuss the connections and distinctions between our proposed perspective and the standard supervised approach to learning to rank via empirical risk minimization with respect to standard information retrieval metrics. The specific focus of this paper is to highlight the significance of the empirical context value distribution in shaping the properties of the corresponding ranking policies by contrasting various empirical importance sampling distributions. We report empirical results from online randomized controlled experiments on a product search ranking task in a major e-commerce platform demonstrating the fundamental trade-offs governed by ranking policies trained on empirical reward estimates with respect to extreme choices of the context value distribution.

Keywords

Learning to Rank, Expected Reward Estimation, policy Learning, Two-Sided Marketplaces

1. Introduction

1.1. Motivation

Two-sided e-commerce marketplaces are intermediary economic platforms that connect buyers and sellers, usually providing a wide selection of products for the buyers from a diverse array of

Presented at the SURE workshop held in conjunction with the 18th ACM Conference on Recommender Systems (RecSys), 2024, in Bari, Italy.

✉ ebrahimzadeh@ebay.com (E. Ebrahimzadeh); nmonga@ebay.com (N. Monga); hanggao@ebay.com (H. Gao); acozzi@ebay.com (A. Cozzi); abagherjeiran@ebay.com (A. Bagherjeiran)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

sellers. The primary *buyer-focused* objective of the marketplace is to guide buyers through their search and discovery journeys to identify and purchase items that fulfill their *shopping intention*. Users' browsing and purchase journeys in the marketplace are impacted by an ecosystem of decision making systems, most notably via the ranking policies in various stages of their shopping journeys from discovery pages to the Search Engine Result Pages(SERP). An effective ranking policy aims to showcase a set of results that match the intent of the user in any given ranking context along the shopping journey with rewards realized as interaction events on the slotted items on the page. Collectively, user journeys are not equally likely to produce value for the marketplace, and the goal is to expand the set of successful user sessions, optimizing a suitable notion of long-term value for users and the marketplace. It is therefore essential that the ranking policy account for the utility of all stakeholders in this economic setting. In standard formulations of learning to rank in the information retrieval literature, however, there is usually no clear connection between the training objective for the ranking policy, the long-term value for the collective of the users and the key performance metrics of the marketplace. In this paper, focusing primarily on the search ranking policy invoked in response to users' search queries, we formulate the ranking policy learning as an optimization problem based on a (counterfactual) estimate of the marketplace reward from observational data.

1.2. Contributions and Related Work

Contribution 1.1. We propose a decision making framework establishing explicit connections between learning a ranking policy for a search/recommendation engines and building effective empirical estimates for a suitable notion of marketplace expected reward.

The problem of developing merit scores for ranking items, post a selection stage from a large pool of candidates, is widely studied in the context of recommendation systems[1], display advertising[2], sponsored search[3], and search ranking[4], where the sequential and hierarchical nature of user interaction events and the sparsity of success events[5, 1, 6, 7] in user journeys are taken into account. Value-aware policies in the context of advertising[3, 3], and economic recommender systems[8, 9, 1] account for business objectives, primarily through manipulations of the merit scores based on conversion likelihood estimates and the price of the candidate items to develop a point-wise notion of expected value for a given candidate item. In contrast, our approach is user focused in that the goal of the ranking policy is to optimize for the user utility in the sense of maximizing the expected number of engagements on desirable items at every stage of the search journey. An alternative formulation is to frame the search ranking policy learning as a multi-stakeholder multi-objective optimization problem[10, 11], with potentially conflicting objectives[12] that account either for business constraints[13] or group exposure constraints[14]. The notion of value for the marketplace is introduced in the ranking policy objective via an importance weighting distribution that signifies both the economic value and the likelihood of realizing some reward from an interaction event with an item that satisfies the user intent.

Contribution 1.2. We characterize the key elements in building effective (policy-dependent) expected reward estimates from observational data, controlling for (1) the heterogeneity of the session value distribution, (2) the contribution of interventions within a user journey via

the reward attribution scheme, and (3) the distribution shifts incurred by selection biases in observational data.

Reinforcement learning(RL) is a powerful framework to account for sequential interventions within the session by formulating the problem of recommending new items[15] or search ranking[16] as a Markov Decision Process (MDP). By expanding the planning horizon and adopting intermediary reward shaping techniques, RL-based approaches account for delayed rewards in the session, via suitable representations of the dynamic session context(state) in session trajectories. Recognizing the selection biases in the observed user behavior data, offline reinforcement learning techniques, including inverse propensity weighting[17], and actor-critic methods[18], are adopted to account for distribution shifts in learning from logged data. Similar counterfactual training techniques based on propensity weighting and potential outcome modeling are developed in the context of counterfactual learning to rank for search ranking problems[19, 20, 21]. There is, however, no clear account of the heterogeneity of marketplace reward across session trajectories, neither in the standard counterfactual supervised learning perspective nor in offline reinforcement learning approaches.

Contribution 1.3. We highlight the significance of the empirical session-context value distribution in building effective marketplace expected reward estimates by demonstrating fundamental performance trade-offs governed by the search ranking policies trained on extreme choices of the context value distribution via rigorous counterfactual evaluations as well as online randomized controlled experiments in a major e-commerce platform.

The definition of success events and the associated reward to the user events is flexible in our framework and is informed by the strategic choices of the marketplace. Specifically, an early-stage marketplace may focus on maximizing the collective number of engagements, while an acquisition-oriented marketplace targets the collective number of purchases, while a revenue-driven marketplace chooses to maximize the long-term gross merchandise value.

1.3. Notation

Here is a list of notation adopted throughout the paper. Sets and ordered sets(lists) are represented with upper-case calligraphic symbols; such as \mathcal{X} . Random quantities are shown in bold such as \mathbf{x} with realization x . The expected value of random variable \mathbf{x} is denoted by $\mathbb{E}[\mathbf{x}]$ and the conditional expectation of a random variable $\mathbf{z} = f(\mathbf{x}, \mathbf{y})$ given \mathbf{y} is denoted by $\mathbb{E}[\mathbf{z}|\mathbf{y}]$ or $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}(x)}[\mathbf{z}]$. For a function $f : \mathbb{X} \rightarrow \mathbb{R}$, the $|\mathcal{X}|$ dimensional array $[f(x)]_{x \in \mathcal{X}}$ is denoted by $f(\mathcal{X})$.

2. Problem Setup

2.1. Decision Making Framework

The marketplace is interested in maximizing the average total reward across all user session trajectories over a long time horizon

$$\frac{1}{T} \sum_{t \leq T} v_{s_t}, \quad (1)$$

where v_s is the economic value from a successfully served search session s . Our framework is flexible in the choice of the reward function and we discuss the fundamental trade-offs between multiple strategic marketplace long term reward choices, namely revenue-based, value per engagement and value per acquisition marketplaces. The reward from a session trajectory is assumed to be non-negative. Although our framework, can be extended to account for negative rewards, we ignore it in our formalization. We assume that the reward over search journeys is a stationary ergodic stochastic process. By invoking Birkhoff’s ergodic theorem[22], with probability 1, the long term temporal average is same as the expected reward, i.e.

$$\mathbb{E}_{\mathbf{s}, \mathbf{v}}[\mathbf{v}_{\mathbf{s}}], \quad (2)$$

where the expectation is taken with respect to the randomness in the session context and reward distribution.

While our formalization can naturally be extended to search journeys with complex goals, we focus on a typical e-commerce purchase decision making scenario of session trajectory a with a single product intent, ignoring sessions with multi-product purchase intent, as well as informational and navigational search sessions. We recognize that users’ decisions are impacted by multiple independently optimized decision making systems, but we are oblivious to potential interactions of the ranking policy with these systems, specifically to the closely related query understanding and candidate retrieval policies. We only focus on policy learning for search result pages with a single layer presentation semantic where the action of the ranking policy is the permutation/ranking of a largely homogeneous set of *comparable items* for a flat single-layered presentation of the results, ignoring the multiplicity of user’s search intent and diversity considerations for the result set.

We cast the problem into a Bayesian decision making framework with a user-focused perspective on the definition of success upon a ranking action. The ranking policy aims to increase the expected number of engagements on items that meet the user intent, and the reward is proportional to the likelihood of a success event(non-zero reward) from the user interactions on the search results page produced by the ranking policy. A crucial aspect of this framework is to account for distribution shifts in observational data, i.e. the distinction between the distribution of the logged search activity data that the policy is trained and the inference time distribution of user events.

2.2. Success From a Ranking Intervention

Given a search query q within a session context s , the ranking policy $\pi : \mathcal{D}_q \rightarrow \{1, \dots, N\}$ maps a candidate item d from the retrieved set \mathcal{D}_q to a slot $\pi(d)$. The notion of success with respect to a slotting $\pi(\mathcal{D}_q)$ of the items on the SERP q is defined based on the effectiveness of the policy in driving user interaction events(Click). Specifically, the objective of the ranking policy on a given ranked SERP is to increase the expected number of engagements on desirable items(suitably defined) $c_{\pi(\mathcal{D}_q)}$ given the session context upon issuing the query $s \prec q$; i.e.

$$\mathbb{E}[c_{\pi(\mathcal{D}_q)} | s \prec q], \quad (3)$$

where the expectation is with respect to the randomness in user preferences and browsing behaviors in the given query context and possibly the randomness in the ranking policy(if

stochastic). Note that $s_{\prec q}$ subsumes all the relevant contextual information upon issuing the query q ; including all the queries and the corresponding surfaced items, as well as the engaged items prior to the current query context. In the subsequent sections, we will make probabilistic assumptions on users' browsing and click behaviors and build effective policy dependent empirical estimators.

2.3. Success From a User Session

The notion of success with respect to a user session is defined based on interaction events on desirable items across all interventions by the marketplace within a user journey. Given the per query ranking objective $\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_q)} | s_{\prec q}]$ (i.e. the expected number of desirable engagements from the SERP given the session context upon issuing the query), the success from the overall user session is shaped by the distribution $\mathbb{P}(q|s)$ that signifies the contribution of the interactions on the ranked SERP q to the overall success of the user search session s .

$$\mathbb{E}_{\mathbf{q} \sim \mathbb{P}(q|s_{\prec q})} [\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_q)} | s_{\prec \mathbf{q}}]], \quad (4)$$

This distribution is usually referred to as the *success attribution distribution*, which is primarily studied in the context of online advertising[23]. The key difference is that in online advertising the unit of value attribution is item impression, while in this work we emphasize the attribution of value to the *ranked list* shown for the query given the prior context. This is also related to the *credit assignment problem* in reinforcement learning on how to attribute success to the intermediate actions of the agent.

2.4. Marketplace Expected Reward

The expected reward of the marketplace from the presented ranking is then shaped by the distribution of the value across search contexts, which signifies the economic value of the user-interaction events in the sessions for the marketplace. The random variable \mathbf{v}_s captures the strategic notion of the value of the session for the marketplace, which corresponds to the value of the interactions events on the item(s) that satisfy the user's intent. The expected reward of the marketplace can then be written as

$$\mathbb{E}_{\mathbf{s}, \mathbf{v}} [\mathbf{v}_s \mathbb{E}_{\mathbf{q} \sim \mathbb{P}(q|s_{\prec q})} [\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_q)} | \mathbf{s}_{\prec \mathbf{q}}]]], \quad (5)$$

We can also consider an alternative formulation where we assume that session value distribution $\mathbb{P}_v(s)$ subsumes both the *likelihood of the user session to lead to some reward* for the search engine as well as the *reward value* attributed to the user session s :

$$\mathbb{E}_{\mathbf{s} \sim \mathbb{P}_v(s)} [\mathbb{E}_{\mathbf{q} \sim \mathbb{P}(q|s_{\prec q})} [\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_q)} | \mathbf{s}_{\prec \mathbf{q}}]]]. \quad (6)$$

For a value-aware search engine with marketplace revenue objective, economic value is realized only in the event of a transaction as the success event from a search session and the reward is proportional to the price of the sold item(s). For a search engine that aims to optimize for the volume of transactions, it is more suitable to adopt a value per acquisition notion of reward oblivious to the price of the sold items. For a search engine with strategic goal of maximizing

user engagements for increased user retention and minimizing abandonment, it is more suitable to adopt a value per click notion of reward oblivious to the post click transaction events.

In the next section, we discuss empirical modeling techniques to build effective empirical reward estimates from observational data, which effectively frame the problem as a standard counterfactual empirical risk minimization with a value-aware context distribution.

3. Expected Reward Estimation from Observational Data

3.1. Estimating the Per Query Success

We are interested in maximizing $\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_q)} | s_{\prec q}]$, the expected number of *desirable engagements* across all the slots on the page, with a suitably parameterized ranking policy π . By hypothesizing an explanatory click model based on causal constructs that govern user browsing and engagement behaviors on search result pages, we can build effective likelihood models from which we can estimate the parameters of the ranking policy via maximum likelihood estimation using logged observational data. We instantiate this process with the simple widely adopted click models in information retrieval. Assuming a vanilla Sequential Browsing Model along with the standard Position-Dependent Examination Model, we can write the expected number of desirable engagements as

$$\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_q)} | s_{\prec q}] = \sum_{r=1}^N \mathbb{P}(\mathbf{c}_{\pi^{-1}(r)} = 1 | s_{\prec q}, \mathcal{D}_q^{\prec r}) \quad (7)$$

$$= \sum_{r=1}^N \mathbb{P}(\mathbf{c}_{\pi^{-1}(r)} = 1 | s_{\prec q}, r) \quad (8)$$

$$= \sum_{r=1}^N \mathbb{P}(\mathbf{o}_{\pi^{-1}(r)} = 1 | q) \times \mathbb{P}(\mathbf{R}_{\pi^{-1}(r)} = 1 | s_{\prec q}) \quad (9)$$

$$= \sum_{r=1}^N \mathbb{P}(\mathbf{o}_{\pi^{-1}(r)} = 1) \times \mathbb{P}(\mathbf{R}_{\pi^{-1}(r)} = 1 | s_{\prec q}) \quad (10)$$

where the first line follows from sequential browsing assumption with $\mathbf{c}_{\pi^{-1}(r)} = 1$ representing the click event on the item ranked at position r and $\mathcal{D}_q^{\prec r}$ representing the slotted items prior to the item ranked in position r . The second line follows from assuming that the user interaction event on a given slot is independent of the placed items in the previous slots. The third line follows from the standard examination-based click model that posits that a click event can be expressed as the intersection of a query specific examination event $\mathbf{o}_{\pi^{-1}(r)} = 1$ and a presentation-independent contextual relevance event $\mathbf{R}_{\pi^{-1}(r)} = 1$; and the last line follows from assuming a query-independent global rank discount function on the examination probabilities. The examination probabilities, a.k.a. propensity scores, are context-specific and can be estimated via explicit online interventions or from observational data. By considering a simple uni-variate model fit on the estimated propensities as a function of rank, one can build data-driven rank-discount functions to estimate users' examination effort. However, the standard approach is to

adopt vanilla log-based context-oblivious rank discounts as generic estimates of the examination probability $\hat{\mathbb{P}}(\mathbf{o}_{\pi^{-1}(r)} = 1)$, i.e.

$$\ell(r) = \frac{1}{\log(1+r)}.$$

Given an empirical estimate \hat{r} of Bayes contextual relevance probabilities $\mathbb{P}(\mathbf{R}_{\pi^{-1}(r)} = 1 | s \prec q)$, one can derive the standard discounted cumulative gain (DCG) estimate $\hat{\mathbb{E}}_{\text{DCG}}[\mathbf{c}_{\pi(\mathcal{D}_q)} | s \prec q]$ of the expected reward per query for the policy $\pi(\cdot)$ as

$$\ell(\pi(\mathcal{D}_q))^T \hat{r}(\mathcal{D}_q). \quad (11)$$

Upon building all the elements of the expected reward estimates, specifically the policy-dependent per query expected reward estimates, we can train the ranking policy by maximizing this empirical reward estimate, as elaborated in the next section.

3.2. Estimating the Session Expected Reward

3.2.1. In-session success attribution

Several techniques can be adopted to estimate the contribution of a ranked SERP q and the corresponding *observed or potential* interactions on that page to the overall success of the user search session s . A simple yet popular solution in the context of online advertising is to adopt an attribution distribution that assigns all the probability mass to the immediate query context preceding the post-click conversion event, which is referred to as Last Touch Attribution scheme. In contrast to this tight attribution scheme, one can assume a uniform distribution across all queries in the session in which the item with the attributed interaction event of interest was retrieved as a candidate item, oblivious to whether it was even impressed on the search result page. This approach is referred to as All Touch attribution scheme. Alternatively, One can assume a (Markovian) probabilistic graphical model on user's touch points within a session journey and infer a probabilistic multi-touch attribution distribution

$$\hat{\mathbb{P}}(q | s \prec q)$$

from observational data. Similarly, one can adopt an attention based sequence modeling approach and infer the contribution weights for interaction events along the user journey with a conversion prediction model. Lose attribution schemes, like the all touch attribution scheme, signify the powerful idea of *counterfactual training context generation* for ranking policy learning, where in contrast to predictive perspectives, the policy can collect reward from a ranking context where the item of interest was not observed by the user. As discussed in the empirical results section, such attribution schemes are particularly effective for capturing the user behavior in search sessions with longer feedback loops, e.g. sessions with high purchase value user intent.

3.2.2. Session Value estimation

In order to highlight the importance of the session value distribution,

$$\hat{\mathbb{P}}(s)$$

let us focus on a search engine with a value per acquisition objective. A straightforward empirical session value distribution is adopt a uniform distribution on sessions that lead to a transaction event. Such session value selection distribution leads to survivorship bias in training context selection in that traffic segments where transaction events are rare, e.g. user sessions with luxury intent, will be under-represented in training. A simple approach is to expand the definition of success events and estimate the likelihood of session success with a content-oblivious estimate based on the aggregate conversion likelihood of the *richest engagement event* attributed to the element(s) engaged. This perspective on building mixture distributions based on the richest post-click engagement event was shown to be effective in capturing potential conversions from browse-heavy user journeys[24].

For a revenue-focused marketplace, as discussed in Section 2.2, the value of a search session is proportional to the price of the item that matches the user intent. In the presence of an observed success event in the logged data, the purchase price of the item to which the success event is attributed is the realization of the session value; otherwise, in the absence of an interaction event, the value of the session has to be estimated from the content of the asked intent in user queries, or a Canonical set of actual or synthesized items that match the user intent.

3.3. Selection Bias Correction

One of the main challenges in learning from observational data is the distribution shift between the training data collected from the logging policies and the inference data distribution. We therefore have to introduce another set of techniques, e.g. importance weighting distributions, to account for this mismatch between the (population) expected reward in (6) and the estimated expected reward from the estimated quantities in the previous sections; that is,

$$\mathbb{E}_{\mathbf{s} \sim \hat{\mathbb{P}}(s)} [\mathbb{E}_{\mathbf{q} \sim \hat{\mathbb{P}}(q|s \prec q)} [\mathbb{E}[\mathbf{c}_{\pi(\mathcal{D}_{\mathbf{q}})} | \mathbf{s} \prec \mathbf{q}]]]. \quad (12)$$

An important source of distribution shift in observational search activity data is the the selection bias due to presentation of the items on the page and the sequential browsing of the users, implying that we only observe relevancy of the items to the user only in the event of an explicit user engagement and it is more likely to observe engagements on SERPs from higher ranking slots.

A key technique to account for this effect is to define a suitable notion of propensity, which is developed in the context of studying the effect of a treatment(an intervention) on a population by taking into account attributes of the treatment unit in the way the treatment is assigned. In the context of ranking, the treatment is defined in correspondence to the examination of a slotted item by the user, but the key difference with the standard applications of this concept is that the examination variable is not fully observable. An alternative approach based on potential outcome modeling, similar to actor-critic networks in the context of offline reinforcement learning, is proposed in [20], where distilled knowledge from a teacher model is used in the form of soft predicted relevance labels to account for unobserved user feedback to achieve variance reduction and improved generalization.

3.4. Variance Reduction and Generalization

Having discussed an array of importance weighting schemes to build empirical expected reward estimates, it is essential to develop variance reduction techniques to control the generalization behavior of expected reward estimators. For brevity of presentation, we briefly discuss the various reduction techniques adopted and ignore developing generalization bounds on the bias and variance of the estimation error of the proposed empirical reward estimation techniques.

3.4.1. Truncation and Bucketing

Clipping and truncated importance sampling techniques[25, 26] are popular techniques to control the variance and generalization behavior of inverse propensity weighting estimators when there is high variance in the estimated propensities. Since we combine multiple importance sampling techniques to account for selection bias, success likelihood, and context value distribution across highly heterogeneous user trajectories, we adopt this simple variance reduction technique off the shelf.

In building empirical session value distributions for a revenue focused marketplace reward, relying on the purchase price of the success items leads to a very high variance estimator, particularly in the presence of high heterogeneity in price intent across user trajectories. Instead, we can use a stratification technique by bucketing user sessions based on value buckets defined according to the empirical revenue distribution. Specifically, we can build a session value distribution based on the empirical revenue share of the bucket corresponding to the price of the purchased item.

3.4.2. Potential Outcome Modeling

One of the the primary challenges of counterfactual learning to rank from logged search activity data is that the relevancy of the items is observed only in the event of explicit user engagements. A popular idea in the context of contextual bandits and recommendation systems to circumvent the challenges in this partial information setting is to use predictive models for reward estimates as potential outcome models in conjunction with inverse propensity weighting[27, 28, 29]. There are a number of recent works in the context of unbiased response prediction that leverage and analyze the doubly robust technique[30, 31, 32, 21]. In [20], a generalized form of potential outcome modeling is proposed where the distilled knowledge from a relevance teacher is used in the form of soft predicted relevance labels to help the student with more effective list-wise comparisons, variance reduction, and improved generalization behavior. This is similar to the idea of actor-critic networks in the context of offline reinforcement learning[18], and augmentation policy in the context of contextual bandits[33]. Using knowledge distillation helps build training contexts from logged search contexts without user interaction events leveraging complex models. To simplify the discussions, we ignore discussing any details about the teacher models used in our experimental setup.

3.4.3. Stratification and Normalization

Effective stratification is a key technique in the context of importance weighting estimators, e.g. the context value binning idea discussed in sub-section 3.4.1 or training context stratification based on characteristics of logged training contexts[34]. We adopt Self-Normalizing propensity based estimators, recently analyzed in [35], where we use engagement ranks as yet another stratification dimension in our proposed estimators. Yet another standard variance reduction technique that we adopt to control the contribution of the search sessions with many success events in the observational data is to adopt normalization techniques; e.g. the standard Ideal cumulative gain normalization for the per query loss. We note that under this cumulative reward normalization technique, per item propensity weights should be reformulated as context weights.

Having equipped our empirical reward estimates with variance reduction techniques, from this point on, we can assume that the effect of all importance weighting schemes discussed so far are reflected in importance weights $\hat{v}_{q,s}$.

3.5. Optimization Objective for the Ranking Policy

We consider deterministic policies parameterized by a scoring function f , such that $\pi_f = \text{argSort}(f)$, oblivious to the representation of the items and the ranking context. An appealing approach, particularly in the context of online advertising and sponsored search, is to directly estimate the Bayes contextual relevance probabilities $\mathbb{P}(\mathbf{R}_d = 1 | s_{\prec q})$, or equivalently the counterfactual probability of click had the item been examined $\mathbb{P}(\mathbf{c}_d = 1 | s_{\prec q}, \text{do}(\mathbf{o}_d = 1))$ via a standard supervised predictive models, i.e.

$$\sum_{s,q} \hat{v}_{q,s} \sum_{d \in D_q} D(f(d) || \hat{r}_d), \quad (13)$$

where $\hat{v}_{q,s}$ is the empirical importance weights based on discussion in section 3.2, D is a distance measure, e.g. cross entropy, between the predicted distributions $f(d)$ and the properly debiased empirical label distribution \hat{r}_d . For estimating counterfactual probability of click that is *Contextually Well-Calibrated and Discriminative* for ranking, we need very complex models with rich feature representations, with careful data stratification and selection bias correction. Since absolute merit estimation is usually a harder problem than difference in merit estimation, we resort to alternative techniques for empirical expected reward optimization.

The standard alternative approach is to adopt the LambdaLoss framework[36] and optimize a pairwise upper bound on the (list-wise) empirical estimates for the expected number of engagements, $\ell_q(\pi_f, \hat{r})$, to circumvent the challenges of dealing with highly non-smooth rank-dependent policy function, which can be written as

$$\ell_q(\pi_f, \hat{r}) = \sum_{d,d' \in D_q} \Delta \hat{\mathbb{E}}_{\pi_f}(\text{swap}_{\hat{r}}(d, d')) \sigma(f(d) - f(d')), \quad (14)$$

where $\hat{\mathbb{E}}_{\pi_f}(\text{swap}_{\hat{r}}(d, d'))$ is the difference in the estimated expected number of engagements had the ranked slots of the item pairs (d, d') been swapped and $\sigma(\cdot)$ is some inverse link function,

e.g. softMax. The approximate surrogate objective, suitably weighted with the empirical reward estimates $\hat{v}_{q,s}$, expressed as

$$\sum_{s,q} \hat{v}_{q,s} \ell_q(\pi_f, \hat{r}) \quad (15)$$

can then be optimized using iterative optimization techniques, like Expectation-maximization; that is given an estimate $f^{(t)}$ at iteration t , in order to build $f^{(t+1)}$ from the gradient updates from the objective function, the difference in estimated objective $\hat{\mathbb{E}}_{\pi_{f^{(t)}}}$ from the swap operation is computed based on ranking order produced by $f^{(t)}$.

4. Evaluations and Discussions

In Section 3, we discussed essential elements of building empirical expected reward estimates for training effective search ranking policies. Since conducting thorough ablation studies for characterizing the effect of each element in building empirical expected reward estimates is not possible given the space constraints, we focus primarily on the rather under-explored element in the literature, which is the effect of *context value distribution* discussed in section 3.2 in shaping the properties and the generalization performance of the ranking policy.

We focus on a product search ranking scenario in a major E-commerce platform and evaluate candidate policies via online randomized control experiments, as well as rigorous counterfactual evaluations on user session data collected from the online traffic. Since all experiment are performed on proprietary data, we only report lifts compared to a simple clearly-specified baseline, with a focus on the relevant choices for controlling the *estimation error* with respect to the research question of interest, oblivious to the optimization framework, the feature representations, and the hypothesis class. Specifically, we only discuss the choice of the ranking objective and the relevant importance sampling and attribution techniques for building our estimators of interest, without discussing the details of the models.

4.1. Online Evaluation Framework

Since the main goal of the proposed decision making framework is to build search ranking policies that generalize with respect to a given notion of marketplace expected reward, we primarily evaluate the performance of the candidate policies in online randomized controlled experiments. Specifically, we adopt an experiment design and primary success metric defined with respect to lifts in cumulative reward in treated user sessions. This cumulative reward driven design is in contrast to the standard experiment design practices for incremental ranking changes, where the primary success metric is set to be the standard (immediate) ranking efficiency metrics that measure concentration of success events in Top slots, through simple attribution and aggregation schemes across search result pages. In fact, top slot engagement concentration metrics, e.g. per query DCG with respect to SERP interactions aggregated uniformly across all queries, which are usually tightly correlated with the marketplace reward, should only be treated as secondary metrics in the presence of a measurement of cumulative reward in online experiments. We do recognize, however, that DCG-type metrics are particularly

crucial for counterfactual off-policy evaluations, as approximations to the per query expected reward using logged data, because all we can do is to measure concentration of logged success events in top slots upon the shuffling action of the new target policy.

We establish the fundamental trade-offs between ranking policies trained on different empirical expected reward objectives primary based on session level cumulative reward metrics, including Number of Engagements, Number of Purchases, and Revenue, as measured in online AB tests. For metrics that attribute the observed effect to search events, we use a simple attribution schemes based on the immediate Search Result Page that precedes the user event of interest.

4.2. Training Objectives and Offline Evaluation Metrics

We adopt the standard supervised counterfactual training and evaluation framework based on logged search activity data collected from the online traffic of a major E-commerce platform. We are oblivious to the logging policy and collect datasets with importance sampling and reward attribution semantics based on the corresponding notions of expected reward of interest. Specifically, given a target notion of expected reward, the context value distribution remains the same for training and evaluation datasets. For candidate item selection per SERP, however, we sample three negative samples at random from impressed unengaged items within each training context, but keep all the candidate items to be re-ranked by the candidate ranker for the evaluation datasets.

For all empirical expected reward metrics, we use the same, suitably debiased and normalized, DCG approximation for the per query expected reward according to (11). Unless explicitly stated otherwise, we use the following vanilla empirical context value distribution for building expected reward estimates as training objectives and the counterfactual metrics.

Expected number of engagements $\hat{\mathbb{E}}[C]$: The session value distribution $\hat{\mathbb{P}}_C(s)$ is a uniform distribution across logged sessions with at least one click event. We consider a simple last touch attribution scheme $\hat{\mathbb{P}}_C(q|s \prec q)$ for the distribution of reward among queries within the session.

Expected number of purchases $\hat{\mathbb{E}}[P]$: The session value distribution $\hat{\mathbb{P}}_P(s)$ is uniform across logged sessions with at least one purchase event. We use a simple multi-touch attribution scheme $\hat{\mathbb{P}}_P(q|s \prec q)$ with uniform distribution across all queries in the converting session, where the purchased item appeared as a candidate.

Expected revenue $\hat{\mathbb{E}}[\text{Rev}]$: The session value distribution $\hat{\mathbb{P}}_R(s)$ is defined on the sessions with a transaction event according to the empirical revenue share of the bucket corresponding to the price of the purchased item. The same multi-touch attribution from above is adopted for this reward estimate as well.

To best highlight the heterogeneity of user behavior with respect to the underlying shopping intent and the associated fundamental trade-offs between different notions of marketplace expected reward, we also stratify our evaluations across traffic segments defined based on purchase price intent of the users, as realized in the price of the purchased item. The price intent bins are defined in such a way so that the empirical revenue distribution is roughly uniform across value buckets.

Online Metric m	$\Delta_{\mathcal{S}}(m_{\pi_{\mathcal{C}}}, m_{\pi_{\mathcal{P}}})$
Sessions With Any Engagement	>+3%
SERPs With Any Engagement	>+5%
Sessions With A Purchase Event	<-2%
Total Revenue	neutral

Table 1

Online AB test results contrasting policies with extreme session value distributions

4.3. Research Questions

4.3.1. Marketplace Reward Trade offs

The primary insight that we would like to highlight in our evaluations is the heterogeneity of users’ browsing and shopping intents, as reflected in different notions of marketplace reward from user sessions. These observations signify the crucial importance of the choice of the empirical session value distribution in shaping the properties and the generalization behavior of the search ranking policy.

We do this by contrasting the performance of ranking policies trained on expected reward estimates corresponding to extreme choices of the empirical session value distribution. Specifically, we compare a policy $\pi_{\mathcal{C}}$, corresponding to a scoring function $f_{\mathcal{C}}$, trained on a simple engagement-driven expected reward estimate based on the session value distribution $\hat{\mathbb{P}}_{\mathcal{C}}(s)$ against a policy $\pi_{\mathcal{P}}$, corresponding to a scoring function $f_{\mathcal{P}}$, trained on a simple acquisition-focused expected reward estimate based on the session value distribution $\hat{\mathbb{P}}_{\mathcal{P}}(s)$. We observe meaningfully different performance trade-offs between these extreme policies with respect to the primary notions of marketplace reward in an online randomized controlled experiment. Table 1 summarizes the key observations on the average effect size $\Delta_{\mathcal{S}}(m_{\pi_{\mathcal{C}}}, m_{\pi_{\mathcal{P}}})$ between engagement focused policy $\pi_{\mathcal{C}}$ and acquisition focused policy $\pi_{\mathcal{P}}$, with respect to different cumulative metrics m , over the global session traffic \mathcal{S} . The main takeaway from these observations is that the engagement focused policy $\pi_{\mathcal{C}}$, on the one hand, drives significantly higher share of search sessions with at least one click(> +3%), and on the other hand, leads to a significant drop in the share of search sessions with at least one purchase(< -2%).

It is interesting to note, however, that this drop is largely due to a significant loss in the number of bought items in search sessions with lower price intent, which usually take less exploration and browsing to identify and pinpoint the desirable item to purchase. Since the engagement-driven policy is more effective in driving success events with higher economic value in sessions that require more browsing effort, it can compensate for the revenue loss due to lower purchases in lower price intent segments, leading to an overall neutral effect size in total revenue.

In order to explore the fundamental trade-offs, highlighted in our online experiment, between different notions of marketplace expected reward across heterogeneous price intents in more depth, we build simple hybrid policies corresponding to a mixture of the engagement-based and acquisition based objectives. Specifically, we build a simple policy via a simple convex combination of the extreme polices

$$\pi_{\alpha} = \text{argSort}((1 - \alpha) f_{\mathcal{P}} + \alpha f_{\mathcal{C}}), \quad (16)$$

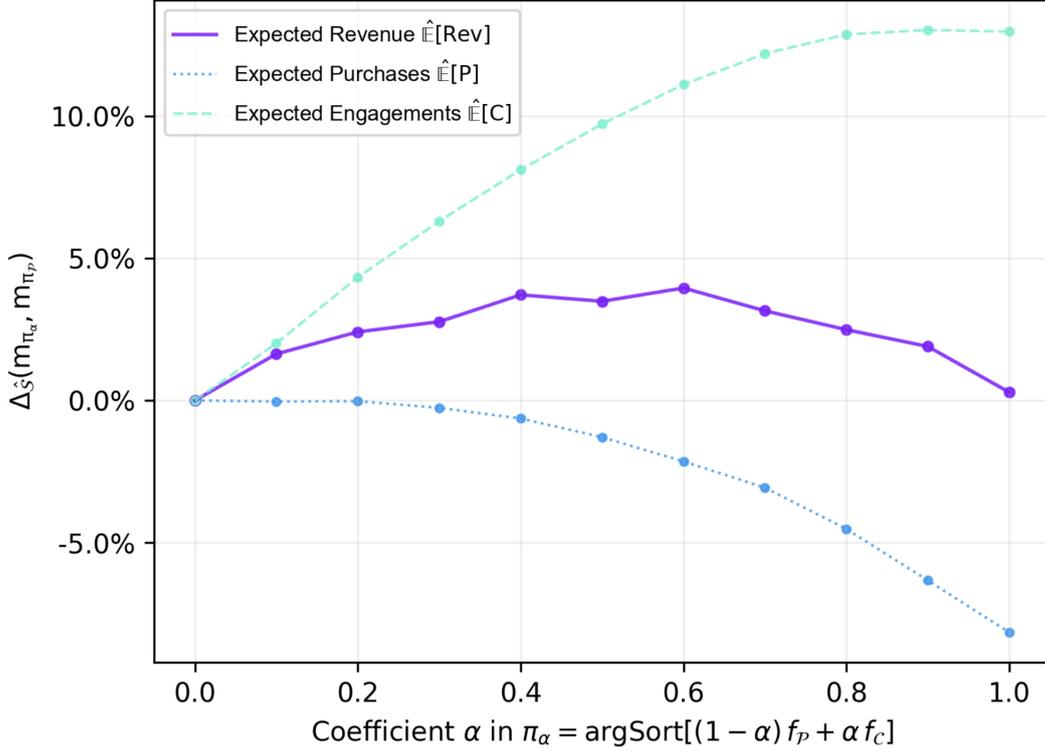


Figure 1: Counterfactual expected reward estimates $\Delta_{\hat{\mathcal{S}}}(m_{\pi_{\alpha}}, m_{\pi_{\mathcal{P}}})$ as a function of parameter α .

where π_{α} refers to the balanced ranking policy obtained via a linear combination of the scoring functions of the engagement focused policy, $\pi_{\mathcal{C}}$ and acquisition focused policy $\pi_{\mathcal{P}}$, for some $\alpha \in [0, 1]$. The parameterized policies π_{α} behave similarly to a policy trained on a corresponding mixture session value distribution $(1 - \alpha)\hat{\mathbb{P}}_{\mathcal{P}}(s) + \alpha\hat{\mathbb{P}}_{\mathcal{C}}(s)$.

Due to the scarcity of online experimentation traffic, we only conduct counterfactual off-policy evaluations for these parameterized policies. While our counterfactual estimates are largely aligned, at least directionally, with the measured effect sizes in online experiments, we point out that all counterfactual off-policy evaluations are fundamentally limited having access only to snapshots of the users' behavior in the logged sessions. In particular, if the logging policy is substantially different from the target policy to be evaluated, the offline evaluation metrics could be very biased.

Figure 1 highlights the essential trade offs between different expected reward estimates m from logged data $\hat{\mathcal{S}}$, with the acquisition focused policy as the baseline

$$\Delta_{\hat{\mathcal{S}}}(m_{\pi_{\alpha}}, m_{\pi_{\mathcal{P}}}). \quad (17)$$

Biasing the training objective heavily on one extreme, leads to significant drops in the estimated reward corresponding to the other extreme. As the contribution of the engagement-focused policy increases, by increasing $\alpha > 0$, we estimate higher expected number of engagements,

Price Intent Buckets p	Purchases	Engagements	Revenue
Low	-0.95%	-0.66%	-0.68%
Low-Moderate	+0.61%	-0.13%	-0.14%
Moderate	+0.53%	+0.44%	-0.41%
High	+4.10%	+0.94%	+1.35%
Very High	+3.89%	+0.74%	+2.11%

Table 2
Online AB Test Results contrasting policies trained on extreme success attribution schemes

with a saturation point of diminishing return, after which a sharp drop in the expected number of purchases is observed. Interestingly the estimated expected revenue is convex as a function of α , which we will discuss in our subsequent research focused on value-aware objectives.

Next, we explore the observed trade-offs in the global analysis above across heterogeneous segments \hat{S}_p corresponding to different price intent segments, where the attribution of a session to a value bucket is done with respect to the price of the purchased items. Figure 2 and 3 show the lift in estimated expected number of engagements $\hat{\mathbb{E}}[C]$ and estimated expected number of engagements $\hat{\mathbb{E}}[P]$, respectively, for the hybrid policy π_α across value segments \hat{S}_p with the acquisition focused policy as the baseline.

We clearly see that the extreme acquisition focused policy performs poorly in terms of the expected number of engagements, across all segments, with particularly larger effects sizes in high value price intents that require more exploration. We also observe that, as α increases, the lift in expected clicks $\Delta_{\hat{S}_p}(\hat{\mathbb{E}}_{\pi_\alpha}[C], \hat{\mathbb{E}}_{\pi_P}[C])$ increases, with a saturation point in lower price segments(which is in fact an inflation point for low price intent segments). On the contrary, biasing the policy towards the engagement-focused policy, by setting α close to 1, leads to a meaningful drop in the expected number of purchases, $\Delta_{\hat{S}_p}(\hat{\mathbb{E}}_{\pi_\alpha}[P], \hat{\mathbb{E}}_{\pi_P}[P])$, particularly in low value price segments, which constitute a high proportion of the overall number of purchases. An interesting observation, however, is that focusing more on an engagement-based objective is helpful for driving even higher expected number of purchases in higher price segments. We leave deeper dives on the observed trade offs for future work.

4.3.2. Tight Attribution of Purchase Events

In order to highlight the significance of the reward attribution scheme within a user session, we contrast the generalization performance of policies trained with respect to extreme choices of the query contribution distribution $\hat{\mathbb{P}}_{\mathcal{P}}(q|s_{<q})$. Specifically, we contrast the performance of a policy π_t trained on a session value distribution with a tight attribution of success events to search events, similar to the last touch scheme discussed earlier, to a policy π_l trained with respect to a loose multi-touch attribution of success events to search events, similar to the all touch scheme discussed earlier. While the overall cumulative rewards do not show sizable performance trade-offs between the two extreme policies, we highlight substantially different effect sizes across different purchase price intents. Table 2 clearly demonstrates that a loose attribution scheme for the empirical query contribution distribution helps with a significant improved generalization in higher price intent sessions, which tend to be more exploratory and

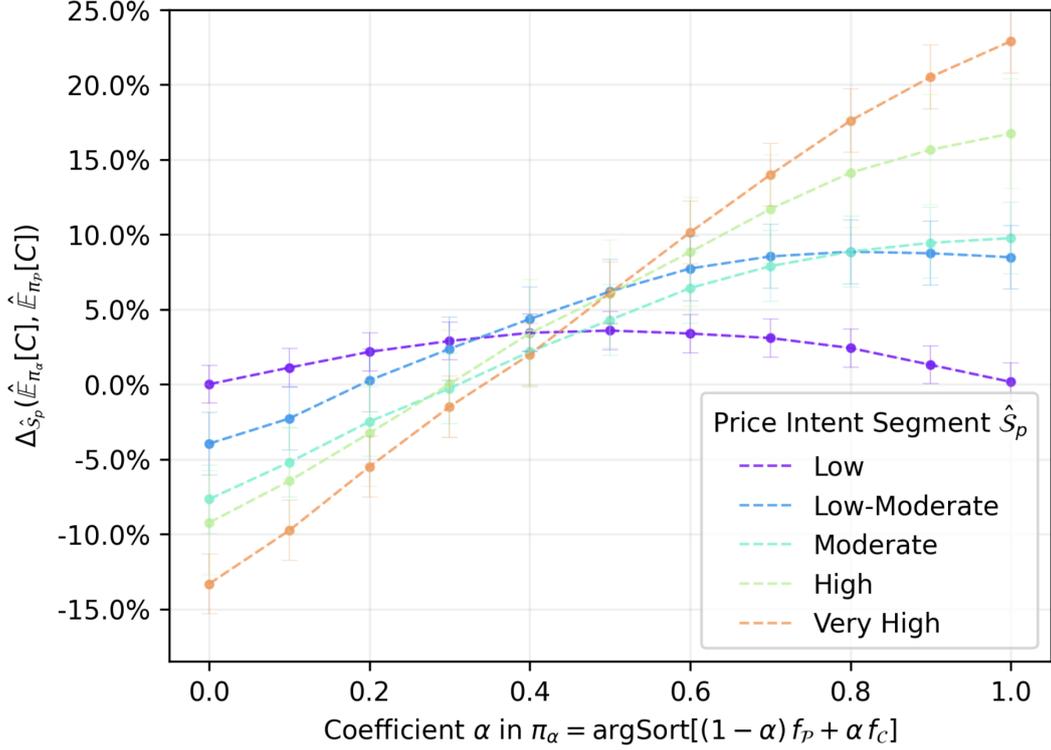


Figure 2: Counterfactual estimate for lift in the expected clicks $\Delta_{\hat{S}_p}(\hat{\mathbb{E}}_{\pi_\alpha}[C], \hat{\mathbb{E}}_{\pi_p}[C])$ across price segments as a function of α .

involve multiple ranking intervention touch points.

4.3.3. Purchase Price in Marketplace Reward

Finally, in order to highlight the significance of incorporating the purchase price in the session value distribution for a revenue focused marketplace reward, we highlight the results from an online AB test on a simple value-aware policy π_v in contrast to a value oblivious acquisition driven policy π_v . The primary difference between the two policies is the empirical session value distribution $\hat{\mathbb{P}}(s)$ in the corresponding expected reward estimate for the training objective, which depends also on the price of the sold item in the case of π_v , and all the other importance weighting distributions and per query reward estimates are the same. In Table 3, we clearly see a significant shift in the distribution of the accumulated revenue across price intent segments, which signifies the importance of taking into account the sparsity of purchase events from higher price intent sessions to avoid the selection bias incurred by the value oblivious policy, which is biased towards lower price intent segments, where sessions with purchase events are abundant.

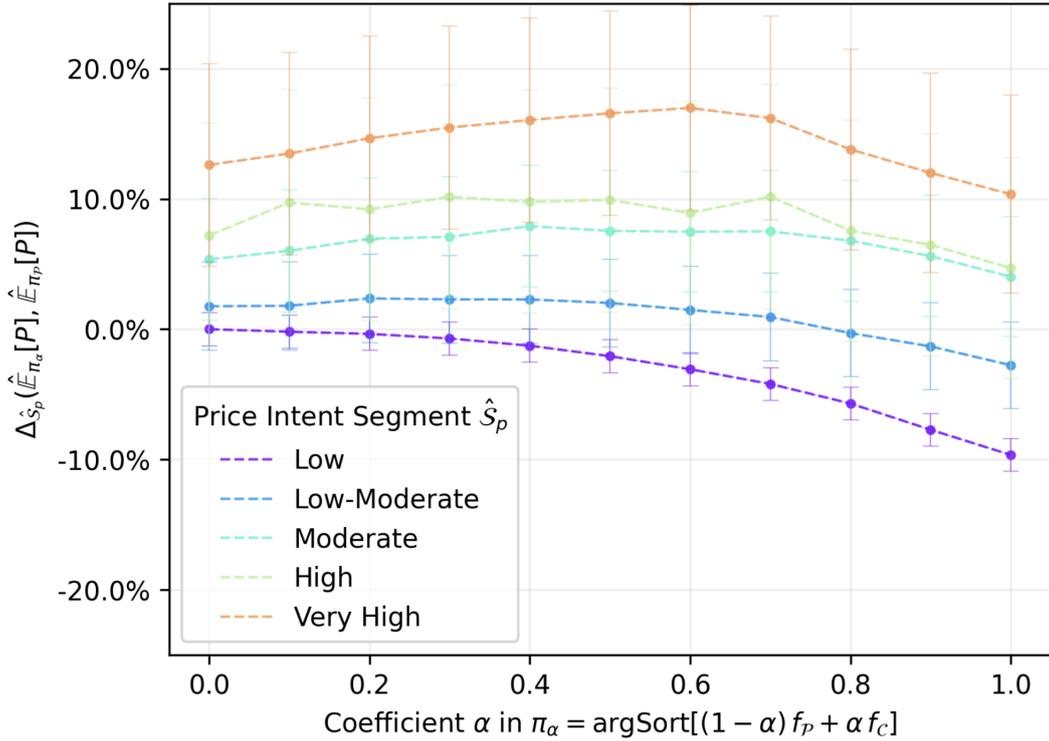


Figure 3: Counterfactual estimate for lift in the expected purchases $\Delta_{\hat{S}_p}(\hat{\mathbb{E}}_{\pi_\alpha}[P], \hat{\mathbb{E}}_{\pi_P}[P])$ across price segments as a function of α .

Price Intent Buckets p	$\Delta_{S_p}(\text{Rev}_{\pi_v}, \text{Rev}_{\pi_*})$
Low	-1%
Low-Moderate	-0.1%
Moderate	+0.5%
High	+1.3%
Very High	+1.9%

Table 3

Online AB test results contrasting a purchase value-aware policy against a value oblivious policy

5. Concluding Remarks

We established an explicit connection between the training objective for the search ranking policy and the key performance metrics of a two-sided commerce marketplace by building effective empirical estimates of the marketplace reward from observation data. Specifically, we highlighted the significance of the *search context value distribution* in building effective empirical estimates of the marketplace expected reward to inform the training and evaluation of the search ranking policy. We showcased empirical results from online randomized controlled experiments and counterfactual evaluations in a major e-commerce platform demonstrating the

fundamental trade-offs governed by extreme choices of the context value distribution.

References

- [1] X. Ma, L. Zhao, G. Huang, Z. Wang, Z. Hu, X. Zhu, K. Gai, Entire space multi-task model: An effective approach for estimating post-click conversion rate, in: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, 2018, pp. 1137–1140.
- [2] S. Chaudhuri, A. Bagherjeiran, J. Liu, Ranking and calibrating click-attributed purchases in performance display advertising, in: Proceedings of the ADKDD'17, 2017, pp. 1–6.
- [3] Y. Zhang, H. Dai, C. Xu, J. Feng, T. Wang, J. Bian, B. Wang, T.-Y. Liu, Sequential click prediction for sponsored search with recurrent neural networks, in: Proceedings of the AAAI conference on artificial intelligence, volume 28, 2014.
- [4] L. Wu, D. Hu, L. Hong, H. Liu, Turning clicks into purchases: Revenue optimization for product search in e-commerce, in: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, 2018, pp. 365–374.
- [5] H. Wang, T.-W. Chang, T. Liu, J. Huang, Z. Chen, C. Yu, R. Li, W. Chu, Escm2: entire space counterfactual multi-task model for post-click conversion rate estimation, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 363–372.
- [6] J. Jin, X. Chen, W. Zhang, Y. Chen, Z. Jiang, Z. Zhu, Z. Su, Y. Yu, Multi-scale user behavior network for entire space multi-task learning, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022, pp. 874–883.
- [7] C. O'Brien, K. S. Liu, J. Neufeld, R. Barreto, J. J. Hunt, An analysis of entire space multi-task models for post-click conversion prediction, in: Proceedings of the 15th ACM Conference on Recommender Systems, 2021, pp. 613–619.
- [8] A. De Biasio, N. Navarin, D. Jannach, Economic recommender systems—a systematic review, *Electronic Commerce Research and Applications* (2023) 101352.
- [9] A. De Biasio, A. Montagna, F. Aiolli, N. Navarin, A systematic review of value-aware recommender systems, *Expert Systems with Applications* (2023) 120131.
- [10] D. Mahapatra, C. Dong, Y. Chen, M. Momma, Multi-label learning to rank through multi-objective optimization, in: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2023, pp. 4605–4616.
- [11] D. Carmel, E. Haramaty, A. Lazerson, L. Lewin-Eytan, Multi-objective ranking optimization for product search using stochastic label aggregation, in: Proceedings of The Web Conference 2020, 2020, pp. 373–383.
- [12] M. Tsagkias, T. H. King, S. Kallumadi, V. Murdock, M. de Rijke, Challenges and research opportunities in ecommerce search and recommendations, in: *ACM Sigir Forum*, volume 54, ACM New York, NY, USA, 2021, pp. 1–23.
- [13] J. Tang, H. Gao, L. He, S. Katariya, Multi-objective learning to rank by model distillation, in: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024, pp. 5783–5792.
- [14] R. Mehrotra, J. McInerney, H. Bouchard, M. Lalmas, F. Diaz, Towards a fair marketplace:

- Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems, in: Proceedings of the 27th acm international conference on information and knowledge management, 2018, pp. 2243–2251.
- [15] G. Shani, D. Heckerman, R. I. Brafman, C. Boutilier, An mdp-based recommender system., *Journal of machine Learning research* 6 (2005).
 - [16] Y. Hu, Q. Da, A. Zeng, Y. Yu, Y. Xu, Reinforcement learning to rank in e-commerce search engine: Formalization, analysis, and application, in: Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining, 2018, pp. 368–377.
 - [17] M. Chen, A. Beutel, P. Covington, S. Jain, F. Belletti, E. H. Chi, Top-k off-policy correction for a reinforce recommender system, in: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, 2019, pp. 456–464.
 - [18] M. Chen, C. Xu, V. Gatto, D. Jain, A. Kumar, E. Chi, Off-policy actor-critic for recommender systems, in: Proceedings of the 16th ACM Conference on Recommender Systems, 2022, pp. 338–349.
 - [19] T. Joachims, A. Swaminathan, T. Schnabel, Unbiased learning-to-rank with biased feedback, in: Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, ACM, 2017, pp. 781–789.
 - [20] E. Ebrahimzadeh, A. Cozzi, A. Bagherjeiran, Counterfactual learning to rank via knowledge distillation, in: Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom’24), 2024.
 - [21] H. Oosterhuis, Doubly robust estimation for correcting position bias in click feedback for unbiased learning to rank, *ACM Transactions on Information Systems* 41 (2023) 1–33.
 - [22] R. Durrett, *Probability: theory and examples*, volume 49, Cambridge university press, 2019.
 - [23] W. Ji, X. Wang, D. Zhang, A probabilistic multi-touch attribution model for online advertising, in: Proceedings of the 25th acm international on conference on information and knowledge management, 2016, pp. 1373–1382.
 - [24] D. Seyler, E. Ebrahimzadeh, A. Cozzi, A. Bagherjeiran, Aligning ranking objectives with e-commerce search intent, in: Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom’23), 2023.
 - [25] L. Bottou, J. Peters, J. Quiñonero-Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, E. Snelson, Counterfactual reasoning and learning systems: The example of computational advertising., *Journal of Machine Learning Research* 14 (2013).
 - [26] E. L. Ionides, Truncated importance sampling, *Journal of Computational and Graphical Statistics* 17 (2008) 295–311.
 - [27] Y. Wang, D. Liang, L. Charlin, D. M. Blei, The deconfounded recommender: A causal inference approach to recommendation, *arXiv preprint arXiv:1808.06581* (2018).
 - [28] M. Dudík, J. Langford, L. Li, Doubly robust policy evaluation and learning, *arXiv preprint arXiv:1103.4601* (2011).
 - [29] T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, T. Joachims, Recommendations as treatments: Debiasing learning and evaluation, in: international conference on machine learning, PMLR, 2016, pp. 1670–1679.
 - [30] L. Zou, C. Hao, H. Cai, S. Wang, S. Cheng, Z. Cheng, W. Ye, S. Gu, D. Yin, Approximated doubly robust search relevance estimation, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022, pp. 3756–3765.

- [31] X. Wang, R. Zhang, Y. Sun, J. Qi, Doubly robust joint learning for recommendation on data missing not at random, in: International Conference on Machine Learning, PMLR, 2019, pp. 6638–6647.
- [32] Y. Saito, Doubly robust estimator for ranking metrics with post-click conversions, in: Proceedings of the 14th ACM Conference on Recommender Systems, 2020, pp. 92–100.
- [33] A. D. Tucker, T. Joachims, Variance-minimizing augmentation logging for counterfactual evaluation in contextual bandits, in: Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, 2023, pp. 967–975.
- [34] E. Ebrahimzadeh, A. Cozzi, A. Bagherjeiran, Intent-aware propensity estimation via click pattern stratification, in: Companion Proceedings of the ACM Web Conference 2023, 2023, pp. 751–755.
- [35] B. London, A. Buchholz, G. Di Benedetto, J. M. Lichtenberg, Y. Stein, T. Joachims, Self-normalized off-policy estimators for ranking (2023).
- [36] X. Wang, C. Li, N. Golbandi, M. Bendersky, M. Najork, The lambdaloss framework for ranking metric optimization, in: Proceedings of the 27th ACM international conference on information and knowledge management, 2018, pp. 1313–1322.