

REWARDRANK: OPTIMIZING TRUE LEARNING-TO-RANK UTILITY

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 Traditional ranking systems optimize offline proxy objectives that rely on oversimplified assumptions about user behavior, often neglecting factors such as position bias and item diversity. Consequently, these models fail to improve true counterfactual utilities such as click-through rate or purchase probability, when evaluated in online A/B tests. We introduce RewardRank, a data-driven learning-to-rank (LTR) framework for counterfactual utility maximization. RewardRank first learns a reward model that predicts the utility of any ranking directly from logged user interactions, and then trains a ranker to maximize this reward using a differentiable soft permutation operator. To enable rigorous and reproducible evaluation, we further propose two benchmark suites: (i) Parametric Oracle Evaluation (PO-Eval), which employs an open-source click model as a counterfactual oracle on the Baidu-ULTR dataset, and (ii) LLM-as-User Evaluation (LAU-Eval), which simulates realistic user behavior via large language models on the Amazon-KDD-Cup dataset. RewardRank achieves the highest counterfactual utility across both benchmarks and demonstrates that optimizing classical metrics such as NDCG is sub-optimal for maximizing true user utility. Finally, using real user feedback from the Baidu-ULTR dataset, RewardRank establishes a new state of the art in offline relevance performance. Overall, our results show that learning-to-rank can be reformulated as direct optimization of counterfactual utility, achieved in a purely data-driven manner without relying on explicit modeling assumptions such as position bias.

1 INTRODUCTION

034 The goal of any ranking system is to model human decision-making in a way that maximizes user engagement and utility. However, real-world user behavior is shaped by subtle, context-dependent cognitive biases that traditional ranking losses fail to capture. Engagement often drops when users are presented with redundant or overly similar items, whereas introducing diversity or strategically positioning items can significantly enhance interest. For example, the decoy effect—where the presence of a less-attractive item increases preference for a similar alternative—has been observed in search interactions and shown to meaningfully influence user choices (Wang et al., 2025a). Other well-documented biases include position bias (Chen et al., 2024; Hager et al., 2024; Zou et al., 2022), brand bias (Li et al., 2025), and similarity aversion (Tversky & Simonson, 2004). In online advertising, the goal is often to maximize the probability that a user clicks on any item in the list, rather than just the top-ranked one. If data shows that users tend to click on the second position, it may be optimal to place the most engaging ad there to improve overall performance. Likewise, in recommendation scenarios, users may prefer a diverse mix of product styles or brands over a cluster of nearly identical, albeit highly relevant, items. Traditional ranking losses, which emphasize relevance at individual positions (typically the top), are ill-suited for modeling such list-level behaviors (Figure 1). They overlook the fact that user utility depends not just on which items are shown, but how they are arranged, highlighting the limitations of handcrafted objectives in capturing the interactive and comparative nature of real user decision-making.

051 A natural way to model user behavior is by learning preferences over full permutations of items within a query group (i.e., a query and its associated items). The ideal objective is to identify 052 and rank those permutations that are most likely to drive user engagement, which can be 053 formulated as a likelihood maximization problem: maximizing the probability of observing high-

engagement permutations while minimizing that of unengaged ones. However, the combinatorial explosion of the permutation space quickly renders this approach intractable; for instance, ranking 10 items results in $10!$ (over 3.6 million) possible arrangements. To address this, recent approaches adopt a utility-based framework (Feng et al., 2021; Shi et al., 2023; Xi et al., 2024; Ren et al., 2024; Wang et al., 2025b), where a utility model is trained to score permutations based on user preferences, and a ranker is subsequently optimized to generate item orders that maximize the predicted utility. While this framework reduces the combinatorial burden, it introduces two key challenges. First is the classic exploration-exploitation dilemma: the ranker must leverage known high-utility arrangements while also exploring novel permutations that may yield higher engagement. Second is utility model misspecification (akin to reward misspecification Clark & Amodei (2016); Coste et al. (2023)): if the learned reward model fails to accurately reflect true user preferences, the ranker may be misled, resulting in poor exploration and degraded overall performance.

In operational ranking systems, user interactions are logged for only a small fraction of the total permutation space. For example, in Figure 1, only 3 out of the 120 possible arrangements of 5 items for the query "laptop bag" are observed, constituting the factual space. These observed interactions define the *factual/observed* space, whereas the vast majority of unexposed, yet potentially high-utility, permutations form the *counterfactual/unobserved* space. The optimal arrangement that maximizes user engagement may exist anywhere within the full permutation space of 120 arrangements. One of the major challenges in counterfactual ranking lies in reliably evaluating unobserved permutations. Even if the full permutation space is modeled, evaluating ranking strategies under counterfactual settings remains challenging due to the lack of explicit supervision (Agarwal et al., 2019; Gupta et al., 2024b;a; Buchholz et al., 2024). For instance, in Figure 1, 117 out of 120 possible arrangements remain unobserved, making their evaluation inherently counterfactual. Existing approaches, such as offline A/B testing, inverse propensity scoring, or other debiasing techniques, are often costly, statistically unstable, or difficult to scale, making counterfactual evaluation a central bottleneck in listwise utility optimization.

To address these challenges, we propose REWARDRANK, a counterfactual utility maximization framework that models user behavior over full item permutations. Rather than scoring items in isolation, we learn a permutation-aware utility function that captures user preferences at the list level. To enable differentiable optimization over permutations, we employ the *SoftSort* operator (Prillo & Eisenschlos, 2020) to construct soft item embeddings, allowing end-to-end training of the ranking model with respect to utility gradients. To mitigate the effects of reward model misspecification—where the learned utility may diverge from actual user preferences—we introduce a correction term in the ranker’s training objective that improves robustness during optimization. For evaluation, we present two scalable, fully automated protocols that assess counterfactual performance without requiring human labels. *Parametric Oracle Evaluation (PO-Eval)* uses a pretrained, position-aware oracle to provide soft supervision and serve as a proxy for user behavior. *LLM-As-User Evaluation (LAU-Eval)* leverages large language models to simulate user preferences and assess ranking quality across unobserved permutations. Together, these methods enable efficient benchmarking of counterfactual ranking strategies and help align learned rankings with actual or simulated user utility.

Our key contributions can be summarized as:



Figure 1: Counterfactual ranking with true learning-to-rank utility. Three arrangements for the query "laptop bag", with item relevance/rating scores (0–5*). The top row ranks purely by relevance but suffers from similarity aversion due to identical color and style, lowering engagement. The middle row improves diversity but surfaces low-relevance items early, which may deter clicks. The bottom row balances diversity and relevance, placing distinct yet relevant items in top positions, leading to higher predicted utility and user engagement. (Figures are generated by GPT-4o)

- 108 • We introduce REWARDRANK, a framework for counterfactual utility maximization that
109 learns a permutation-aware reward model, capturing human list-level preferences and behav-
110 ioral biases without any explicit modeling assumptions such as position bias.
- 111 • We enable end-to-end ranking optimization using differentiable soft permutation operators,
112 and incorporating a per-item auxiliary loss along with a misspecified reward correction term
113 to aid counterfactual space exploration.
- 114 • We propose two large-scale automated evaluation protocols: *PO-Eval* (parametric oracle) and
115 *LAU-Eval* (LLM-as-user), and construct reproducible testbeds for scalable counterfactual
116 ranking evaluation. Experiments on these testbeds reveal that optimizing standard offline
117 ranking metrics such as NDCG do not reliably maximize true user utility.
- 118 • In both proposed counterfactual testbeds, REWARDRANK consistently achieves the highest
119 learning-to-rank (LTR) utility compared to existing and widely adopted ranking methods.
120 When trained on real click signals from an industry-scale dataset, REWARDRANK further
121 establishes a new state-of-the-art in relevance performance.

123 2 RELATED WORK

124 **Traditional ranking methods.** Traditional learning-to-rank (LTR) methods are typically categorized
125 into three classes: point-wise, pair-wise, and list-wise approaches. Point-wise methods treat ranking
126 as a regression or classification problem by independently assigning relevance scores to each item
127 (Burges et al., 2005a;b). While computationally efficient, they neglect interactions among items in the
128 ranked list. Pair-wise approaches, including RankSVM (Joachims, 2002), RankBoost (Freund et al.,
129 2003), and LambdaMART (Burges, 2006; Wu et al., 2010), aim to learn relative preferences between
130 item pairs, improving over point-wise methods but still failing to capture full list-level dependencies.
131 In contrast, list-wise methods optimize objectives over the entire ranking, such as NDCG (Cao et al.,
132 2007; Xia et al., 2008), offering better alignment with evaluation metrics.

133 Recent large-scale datasets such as Baidu-ULTR (Zou et al., 2022) have enabled realistic benchmarking
134 of ranking algorithms under user-interaction-driven settings, facilitating systematic studies on
135 position bias, distribution shift, and counterfactual evaluation in LTR (Hager et al., 2024). Building
136 on these advances, modern approaches have expanded beyond purely supervised objectives toward
137 *data-driven* and *representation-rich* formulations. Pretraining-based LTR models leverage large
138 language or multimodal corpora to learn transferable ranking priors (Hou et al., 2024), while latent
139 cross-encoding methods (Luo et al., 2022) and set-aware transformers Qin et al. (2021) jointly embed
140 queries and items to capture fine-grained contextual dependencies.

141 **Counterfactual Learning-to-Rank.** Prior work in counterfactual learning-to-rank (CLTR) primarily
142 addresses position bias in implicit feedback using methods such as inverse propensity scoring
143 (IPS) (Joachims et al., 2017) and doubly robust estimation (Oosterhuis, 2023). Extensions include
144 modeling trust bias (Agarwal et al., 2019) and jointly correcting for both position and trust biases (Var-
145 dasbi et al., 2020). Recent approaches explore policy optimization via proximal updates (Gupta et al.,
146 2024b) and extend this to trust-aware CLTR through proximal ranking objectives (Gupta et al., 2024a).
147 While effective, these methods often focus narrowly on position bias or make strong assumptions,
148 underscoring the need for broader utility-driven ranking frameworks, as pursued in this work.

149 **Utility-oriented counterfactual reranking.** Reranking methods enhance an initial ranked list by
150 applying a secondary model to better optimize downstream objectives such as user utility, fairness, or
151 diversity (Xi et al., 2024; Wang et al., 2025b). Recent work in counterfactual ranking predominantly
152 follows a two-stage framework consisting of a generator and an evaluator (Xi et al., 2024; Shi
153 et al., 2023; Ren et al., 2024; Wang et al., 2025b). For example, URCC (Xi et al., 2024) trains a
154 set-aware utility model and employs a context-sensitive pairwise LambdaLoss to guide the ranker.
155 NLGR (Wang et al., 2025b) leverages neighboring lists within a generator-evaluator setup for utility
156 optimization. PRS (Feng et al., 2021) adopts beam search to generate candidate permutations and
157 evaluates them using a permutation-wise scoring model, while PIER (Shi et al., 2023) uses SimHash
158 to select top-K candidates from the full permutation space efficiently.

159 Reranking approaches rely on a strong base ranker trained on logged data and typically explore
160 counterfactuals around its initial permutations Xi et al. (2024); Wang et al. (2025b), which constrains
161 exploration and limits discovery of globally optimal rankings. Importantly, these methods do not learn

162 an explicit reward model; instead, they assume a predefined metric such as NDCG to serve as the
 163 counterfactual reward (Joachims et al., 2017; Agarwal et al., 2019). While effective in certain settings,
 164 this reliance on a fixed evaluation metric can hinder adaptability to more general or task-specific
 165 reward signals.

166 **Differential approximation to ranking.** A key challenge in learning-to-rank is the mismatch
 167 between evaluation metrics (e.g., NDCG, MAP) and surrogate loss functions amenable to gradient-
 168 based optimization, due to the non-differentiable nature of sorting operations. To address this, prior
 169 work has either proposed smooth approximations to the rank function (e.g., ApproxNDCG (Qin et al.,
 170 2010)) or introduced differentiable approximations to argsort using soft permutation matrices (Grover
 171 et al., 2019; Prillo & Eisenschlos, 2020); for instance, PiRank (Swezey et al., 2021) and NeuralNDCG
 172 (Pobrotyn & Białobrzeski, 2021) utilize NeuralSort as a temperature-controlled surrogate. Another
 173 line of work leverages the Plackett–Luce distribution to model ranking policies in a differentiable
 174 manner (Oosterhuis, 2021). Methods like PG-RANK (Gao et al., 2023) use policy gradients to
 175 optimize the expected reward over the Plackett–Luce distribution based on REINFORCE, while
 176 ListNet (Cao et al., 2007) and ListMLE (Xia et al., 2008) employ the Plackett–Luce framework to
 177 derive smooth list-wise objectives.

178 Please refer to Appendix Section A.5 for additional related work.

180 3 LEARNING-TO-RANK PROBLEM: UTILITY MAXIMIZATION VS SORTING

182 A data sample of an LTR problem is a *query group* (QG) consisting of a query, q , and a set of L items,
 183 $\{x_\ell\}_{\ell=1}^L$, where L may vary. The *query* may represent, for example, a search string, a user profile, or
 184 other contextual information like device type and page layout. The *items* are candidate entities like
 185 webpages, songs, or products retrieved by an upstream system. We assume that the QGs are drawn
 186 i.i.d. from a distribution \mathcal{P} , i.e. $(q, \{x_\ell\}) \sim \mathcal{P}$. When a user is presented with a ranking/arrangement
 187 (permutation), $\pi : [L] \rightarrow [L]$, of the items of a QG, i.e. $(x_{\pi(1)}, \dots, x_{\pi(L)})$, they interact with the
 188 ranked items, yielding a stochastic utility $U(q, \{x_\ell\}, \pi) \in \mathbb{R}$, which is a hidden function of the QG
 189 and the ranking. In typical internet systems, the utility can represent outcomes such as whether a user
 190 clicks or purchases any item, or continuous measures such as the total minutes of media consumed.
 191 Our objective is to learn a ranking policy, f , mapping the QGs to permutations, that maximizes the
 192 expected utility return, i.e.

$$193 f^* = \underset{f}{\operatorname{argmax}} \mathbb{E}_{(q, \{x_\ell\}) \sim \mathcal{P}} [U(q, \{x_\ell\}, \pi = f(q, \{x_\ell\})] \quad (1)$$

194 Based on the choice of the utility, this objective corresponds to business metrics like click-through
 195 rate, units sold, or streamed minutes. The main challenge here is that the hidden stochastic utility
 196 function U is not directly observable. Instead we are given a training dataset, \mathcal{D} , consisting of N
 197 QGs (indexed by i) and their observed utility $\{u_i\}$ under some logged rankings $\{\pi_i\}$, i.e. $\mathcal{D} =$
 198 $\{(q_i, \{x_{i,\ell}\}_{\ell=1}^{L_i}, \pi_i, u_i)\}_{i \in [N]}$. We assume that similar hold-out test and validation dataset are also
 199 available. This setting can be viewed as an offline one-step reinforcement-learning problem in which
 200 the state space is comprised of all possible QGs in the support of \mathcal{P} , the action space is comprised of
 201 all item permutations, and the reward is the observed utility.

202 In practice, most QGs are unique, so we observe only one out of $L!$ possible rankings for each.
 203 Consequently, even if we propose a better alternative ranking for a given QG, the *counterfactual*
 204 utility it would have obtained remains unknown. To address this, traditional LTR algorithms optimize
 205 heuristic offline ranking metrics like Normalized Discounted Cumulative Gain (NDCG) (Järvelin &
 206 Kekäläinen, 2002; Burges, 2006), averaged over a test set. When a user interacts with a ranked QG,
 207 we also obtain per-item feedback signals $\{y_\ell \geq 0\}$ (e.g. whether an item was clicked or purchase,
 208 or how many minutes it was streamed). Usually, the overall QG-level utility u is some function of
 209 these per-item signals. Then, the NDCG of any new ranking $\hat{\pi}$, on a QG with feedbacks $\{y_\ell\}$ can be
 210 defined as

$$211 \text{NDCG}(\hat{\pi}, \{y_\ell\}) \triangleq \frac{\text{DCG}(\hat{\pi}, \{y_\ell\})}{\text{DCG}(\pi^*, \{y_\ell\})} \in [0, 1], \text{ where } \text{DCG}(r, \{y_\ell\}) \triangleq \sum_{\ell=1}^L \frac{2^{y_\ell} - 1}{\log_2(1 + r^{-1}(\ell))}. \quad (2)$$

212 DCG assigns a gain $2^{y_\ell} - 1$ for the item x_ℓ in a test QG, but its contribution to the metric is discounted
 213 by its position $r^{-1}(\ell)$ under the ranking r . Thus, NDCG is maximized when the items are ranked

216 in the descending order of their feedback values, i.e., under the optimal ranking π^* . Traditional
 217 LTR methods (Burges, 2006; Swezey et al., 2021), aim to maximize NDCG by optimizing various
 218 continuous relaxations of it. This heuristic of learning to move items with higher feedback signal
 219 to the top of list have been highly successful, potentially because (i) items with positive feedback
 220 are usually relevant, and (ii) users tend to focus their attention on the top of the list. However, such
 221 offline metrics are now well-known to be sub-optimal as they do not perfectly align with the true
 222 (hidden) utility (1) we aim to maximize (Wang et al., 2023; Jeunen et al., 2024). A key advantage of
 223 REWARDRANK over traditional LTR methods is its ability to *leverage data without click/purchase*
 224 *labels*. Whereas standard pipelines often discard sessions with no purchases/clicks (or treat them as
 225 uninformative negatives), our approach can still extract signal from these interactions via its utility
 226 modeling and preference estimation. This aligns with recent evidence that leveraging unlabeled or
 227 weakly labeled interaction data—e.g., through pretraining or preference modeling—improves ranking
 228 quality (Hou et al., 2024). In the next section, we introduce REWARDRANK, a data-driven ranking
 229 framework that directly maximizes the true LTR utility without relying on heuristics or specific
 230 user-behavior assumptions.

231 4 REWARDRANK: DATA-DRIVEN LTR UTILITY MAXIMIZATION

232 In this section, we present the REWARDRANK framework, which aims to maximize the true (hidden)
 233 LTR utility defined in (2). At a high level, REWARDRANK proceeds in two stages. First, using the
 234 logged training data \mathcal{D} , it learns a reward model that predicts the counterfactual utility for any QG
 235 and permutation. Then it trains a ranker using the reward model’s predictions as supervision, so as to
 236 maximize the expected counterfactual LTR utility of the ranker’s item arrangement (ranking) policy.
 237

238 4.1 STAGE 1: LEARNING THE UTILITY USING A REWARD MODEL

239 Let $g(q, \{x_\ell\}, \pi; \phi)$ denote the reward model (parameterized with ϕ) to predict the scalar utility
 240 for the QG $(q, \{x_\ell\})$ and a ranking π . It is trained solely on the logged query groups, rankings
 241 and observed utilities in the training dataset \mathcal{D} . When the utility $U \in \{0, 1\}$ is a binary random
 242 variable (e.g. click, purchase), we train $g \in [0, 1]$ by minimizing the average binary cross-entropy
 243 loss between the observed u_i and the predicted utilities $\hat{u}_i(\phi) := g(q_i, \{x_{i\ell}\}, \pi_i; \phi)$ over all $i \in [N]$:
 244

$$245 \min_{\phi} \left[\text{RewardLoss}(\phi) \triangleq -\frac{1}{N} \sum_{i=1}^N [u_i \log(\hat{u}_i(\phi)) + (1 - u_i) \log(1 - \hat{u}_i(\phi))] \right]. \quad (3)$$

246 When the true utility is a continuous random variable (e.g. minutes a song is streamed) we can
 247 use regression losses such mean squared error (MSE) $\min_{\phi} (1/N) \sum_{i=1}^N \|u_i - \hat{u}_i(\phi)\|^2$. In our
 248 experiments, the reward model is instantiated with a transformer encoder, $\text{Enc}_{\text{reward}}$, due to its ability to
 249 model functions over sequences (ranked list of items). Before passing a QG into $\text{Enc}_{\text{reward}}$ each query-item
 250 pair $[q, x_\ell]$ is embedded using a text encoder to create the token embedding \mathbf{e}_ℓ . Next, the ranking
 251 of the items π is encoded through position encodings $\{\mathbf{p}_k\}$. Then the position encoded tokens are
 252 passed to the transformer $\text{Enc}_{\text{reward}}$. Finally, the predicted utility is computed as the sigmoid of a linear
 253 function of the $[\text{CLS}]$ token output. This can be succinctly represented as:

$$254 g(q, \{x_\ell\}_{\ell=1}^L, \pi; \phi) = \sigma \left[\mathbf{v}^\top \text{Enc}_{\text{reward}}^{[\text{CLS}]} \left(\{\mathbf{e}_{\pi(k)} + \mathbf{p}_k\}_{k=1}^L \right) \right], \quad (4)$$

255 where $\mathbf{e}_{\pi(k)}$ is the token embedding of the query and the k -th ranked item. During the second stage
 256 of training the ranker, we freeze the reward model parameters ϕ .

257 **Auxiliary per-item predictor:** Typically the observed utility u is a byproduct of user’s interaction
 258 with the items. So, we hypothesize that predicting the per-item feedback signals $\{y_\ell\}$ as an auxiliary
 259 task would improve the overall quality of the LTR utility prediction. Thus we include an auxiliary
 260 prediction head on the item tokens’ outputs to predict the feedback signal observed at each ranked
 261 position $k \in [L]$. When $y_\ell \in \{0, 1\}$ is binary, the predictions can be instantiated as

$$262 \hat{y}_k(\phi) \triangleq \sigma \left[\tilde{\mathbf{v}}^\top \text{Enc}_{\text{reward}}^{(k)} \left(\{\mathbf{e}_{\pi(k)} + \mathbf{p}_k\}_{k=1}^L \right) \right], \quad \forall k \in [L], \quad (5)$$

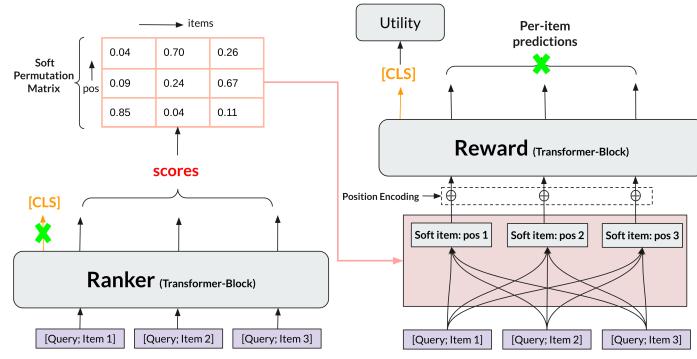


Figure 2: **REWARD RANK**. A ranker scores the items in a query group. These scores are used to compute soft item embeddings via a soft permutation matrix. Position encoded soft item embeddings are passed into a reward to estimate its utility. Finally, the ranker is optimized to maximize the predicted utility.

where $\text{Enc}_{\text{reward}}^{(k)}$ is the output token at the k -th position and σ is the sigmoid function. We can learn $\hat{y}_k(\phi)$ alongside $\hat{u}(\phi)$ by adding the average cross-entropy loss between \hat{y}_k and $y_{\pi(k)}$,

$$\text{ItemLoss}(\phi) \triangleq -\frac{1}{\sum_i L_i} \sum_{i=1}^N \sum_{k=1}^{L_i} [y_{i\pi(k)} \log(\hat{y}_{ik}(\phi)) + (1 - y_{i\pi(k)}) \log(1 - \hat{y}_{ik}(\phi))]. \quad (6)$$

as an additional regularizer to RewardLoss (3). Note that during the training of the ranker in the next stage, these auxiliary predictions can be discarded. Our ablation in Section 5.1 shows that the per-item loss provides a moderate boost in performance. We also apply the per-item loss to query groups (QGs) with no purchases (i.e., no positive labels). This enables us to exploit otherwise discarded sessions and stabilize learning in sparse-feedback regimes by providing item-level signals even when list-level purchase supervision is absent.

4.2 STAGE 2: RANKER REWARD MAXIMIZATION THROUGH SOFT SORTING

Typically, rankers are modeled as scoring functions that assign a score to each item in a QG. Then the items are ordered in the descending order of their scores to obtain the final ranking. We follow the same pattern and define $f(q, \{x_\ell\}; \theta)$ as a scoring-based ranker which maps a QG $(q, \{x_\ell\})$ to a set of item scores $\{s_\ell\}$. Following our reward model design, we instantiate f using the same transformer backbone architecture. Since QG has an unordered set of items, we do not use position encoding. Finally the score is computed as the linear function of the output item tokens, i.e.

$$s_\ell \triangleq f_\ell(q, \{x_\ell\}; \theta) \triangleq \mathbf{w}^\top \text{Enc}_{\text{ranker}}^{(\ell)} \left(\{\mathbf{e}_\ell\}_{\ell=1}^L \right), \quad \forall \ell \in [L], \quad (7)$$

where $\text{Enc}_{\text{ranker}}^{(\ell)}$ is the output token of the ℓ -th item. Our goal is to optimize the effective ranking $\hat{\pi}$ induced by these scores so that it maximizes the expected counterfactual utility, which is a hidden from us. This is where the reward model comes in handy, as it helps us predict the counterfactual utility as $\hat{u} := g(q, \{x_\ell\}, \hat{\pi})$. However, since sorting (of the scores) is a discontinuous operation, it is challenging to optimize the scores to maximize the reward. To enable an end-to-end optimization of the scorer f , we resort to a continuous relaxation of the sorting operation.

Soft Permutation via SoftSort. *SoftSort* (Prillo & Eisenschlos, 2020) is a continuous relaxation of sorting operation. It defines a *unimodal row-stochastic* matrix (Swezey et al., 2021) as the *soft* permutation matrix $\hat{\Pi}^{(\tau)} \in [0, 1]^{L \times L}$. Row k of this matrix corresponds to a probability distribution of the k -the ranked item over the set of all items. Formally, we define

$$\hat{\Pi}_{k,\ell}^{(\tau)} \triangleq \frac{\exp\left(-\frac{1}{\tau} |s_\ell - s_{\hat{\pi}(k)}|\right)}{\sum_{\ell'=1}^L \exp\left(-\frac{1}{\tau} |s_{\ell'} - s_{\hat{\pi}(k)}|\right)}, \quad \forall k, \ell \in [L], \quad (8)$$

324 where τ is a temperature parameter and $\widehat{\pi}(k)$ is the k -th ranked items when (hard) sorting by the
 325 scores $\{s_\ell\}$. $\widehat{\Pi}^{(\tau)}$ is a continuous function of the scores $\{s_\ell\}$ and when $\tau \rightarrow 0$, $\widehat{\Pi}^{(\tau)}$ tends to the
 326 binary hard-permutation matrix $\widehat{\Pi}$, where
 327

$$328 \lim_{\tau \rightarrow 0} \widehat{\Pi}_{k,\ell}^{(\tau)} = \widehat{\Pi}_{k,\ell} \triangleq \mathbb{I}\{\widehat{\pi}(k) = \ell\}, \quad \forall k, \ell \in [L], \quad (9)$$

331 assuming the scores are unique. Using this soft permutation matrix, we can compute a soft item
 332 embedding $\widehat{e}_k^{(\tau)}$ at position k as the following convex combination of the true item embeddings
 333

$$334 \widehat{e}_k^{(\tau)} \triangleq \sum_{\ell \in [L]} \widehat{\Pi}_{k,\ell}^{(\tau)} \mathbf{e}_\ell. \quad (10)$$

337 It is easy to verify that $\widehat{e}_k^{(\tau)} \rightarrow \mathbf{e}_{\widehat{\pi}(k)}$ when $\tau \rightarrow 0$. Note that there are alternate soft permutation
 338 matrices like NeuralSort (Grover et al., 2019), but we adopt SoftSort for its simplicity and state of
 339 the art performance (Prillo & Eisenschlos, 2020). We then compute a soft reward for these soft item
 340 embeddings using

$$341 \widehat{g}(\theta) \triangleq g(q, \{x_\ell\}, \widehat{\Pi}^{(\tau)}) \triangleq \sigma \left[\mathbf{v}^\top \text{Enc}_{\text{reward}}^{[\text{CLS}]} \left(\left\{ \widehat{e}_k^{(\tau)} + \mathbf{p}_k \right\}_{k=1}^L \right) \right], \quad (11)$$

344 This allows us to compute an approximate predicted reward (11) as a continuous function over the
 345 ranker scores $\{s_\ell\}$ through the SoftSort matrix. Finally, we optimize the parameters of scorer f to
 346 maximize the average approximate reward over the training set in an end-to-end manner:

$$348 \min_{\theta} \left[\text{RankerLoss}(\theta) \triangleq -\frac{1}{N} \sum_{i=1}^N \widehat{g}_i(\theta) \right]. \quad (12)$$

351 Even though REWARDRANK is maximizing the predicted utility of the soft ranking, we hypothesize
 352 that it generalizes well and produces rankings with higher expected counterfactual utility than prior
 353 LTR methods.

354 An alternative to *soft-permutation matrices* is the Plackett–Luce (PL) model, which offers efficient,
 355 closed-form gradients for ranking. However, counterfactual learning with PL requires Monte Carlo
 356 sampling, leading to high-variance estimates in large action spaces. While variance reduction
 357 helps (Gao et al., 2023), unbiased learning fundamentally depends on stochastic logging, which is
 358 incompatible with real-world deterministic rankers designed for stability and trust. Soft permutation
 359 relaxations like SoftSort (Prillo & Eisenschlos, 2020) approximate permutations in continuous space,
 360 enabling gradient-based optimization without sampling. Though computationally more expensive,
 361 they reduce variance and support end-to-end utility maximization. We pair SoftSort with a learned
 362 reward model that generalizes over logged data, enabling scalable training under deterministic logs.
 363 This approach trades unbiasedness for stability and practicality in real-world ranking systems.

364 **Mitigating reward misspecification.** One challenge of reward modeling the hidden counterfactual
 365 utility is model misspecification, i.e. a gap between the predicted and the true utilities. A misspecified
 366 reward can misguide the ranker into wrong ranking policies Coste et al. (2023); Clark & Amodei
 367 (2016). To mitigate this issue we propose a sample reweighting scheme which modifies the ranker
 368 loss as

$$370 \text{RankerLoss}^{(\lambda)}(\theta) \triangleq -\frac{1}{N} \sum_{i=1}^N w_i \cdot \widehat{g}_i(\theta), \text{ where } w_i = 1 - \lambda |u_i - \widehat{u}_i| \in [0, 1] \text{ and } \lambda \in [0, 1], \quad \forall i. \quad (13)$$

374 Above loss is a pessimistic upperbound to $\text{RankerLoss}(\theta)$ (12). This reward down-weighting scheme
 375 is motivated by a conjecture that when the observed utility u_i for the i -th training QG and the
 376 corresponding prediction $\widehat{u}_i(\theta)$ are different, the utility prediction on new ranking of this QG would
 377 also be less reliable. Through an ablation in Section 5.1 we show that reward misspecification
 correction slightly improves the REWARDRANK performance.

378 **5 EXPERIMENTAL RESULTS**

380 **Datasets.** Public large-scale datasets for learning-to-rank (LTR), especially in counterfactual settings,
 381 are scarce. To the best of our knowledge, we propose the first reproducible testbeds for counterfactual
 382 ranking evaluation. We utilize two existing large-scale LTR datasets: Baidu-ULTR (Hager et al.,
 383 2024; Zou et al., 2022) and Amazon KDD-Cup (Reddy et al., 2022), to construct these testbenches,
 384 enabling rigorous evaluation of permutation-aware ranking policies. Baidu-ULTR contains 1.8M
 385 query groups (11.7M query-document pairs) and 590K validation/test sessions. Amazon KDD-Cup
 386 comprises 130K queries and 2.6M annotated query-product pairs with rich textual metadata. We
 387 generate 400K training and 50K validation/test query groups by sampling permutations of products
 388 per query. See Appendix B.1 for further details.

389 **Implementation Details and Baselines.** Our reward models and rankers are based on a transformer
 390 architecture with 12 layers, 768 hidden dimensions, 12 attention heads, and roughly 110M parameters.
 391 We set $\tau = 0.5$ and $\lambda = 0.7$ for all REWARDRANK experiments, based on tuning over a held-out
 392 set. Ablations with varying values and further implementation details are provided in Appendix B.2.
 393 For comparison, we implement two utility-based counterfactual ranking methods: URCC (Xi et al.,
 394 2024), which uses a LambdaLoss-based pairwise objective, and PG-rank (Gao et al., 2023), which
 395 applies Plackett–Luce modeling with policy gradients. Our variants, URCC* and PG-rank*, replace
 396 their offline metric utility (e.g. NDCG) with our transformer-based reward model for improved
 397 counterfactual performance. Additionally, we train standard LTR baselines: ListNet (Cao et al.,
 398 2007), ListMLE (Xia et al., 2008), LambdaRank (Wang et al., 2018), and PiRank (Swezey et al.,
 399 2021), all using the same transformer architecture for fair comparison across supervision methods.

400 **5.1 LARGE-SCALE REPRODUCIBLE TESTBENCHES FOR COUNTERFACTUAL LTR**

401 To enable reproducible evaluation of ranking policies without online A/B testing, we introduce
 402 two complementary testbeds: PO-Eval, which leverages a parametric click model, and LAU-Eval,
 403 which simulates human-like shopping behavior via LLM reasoning. Together, they enable holistic,
 404 counterfactual assessment of ranking algorithms under both statistical and behavioral lenses.

406 **Parametric Oracle Evaluation (PO-Eval).** To simulate a click-based counterfactual recom-
 407 mendation setting, we build a testbed from the Baidu-ULTR dataset (Hager et al., 2024), employing a
 408 *pretrained parametric IPS model* as the oracle for supervision. This model estimates the click proba-
 409 bility at position ℓ as $P(C) = P(E_\ell) \cdot P(R_{q,i})$, where $P(E_\ell)$ is the position-dependent examination
 410 probability and $P(R_{q,i})$ is the click probability given examination. We use this oracle to sample binary
 411 clicks for training and later reuse it for counterfactual evaluation of new ranking policies. For each
 412 ranked query group (QG), we compute the expected utility as the probability of at least one click and
 413 the observed utility as a binary indicator of at least one sampled click. This setup provides a realistic
 414 and repeatable framework for evaluating how well learned rankers align with user behavior modeled
 415 by the IPS-oracle. See Appendix A.1 for details on the parametric model and the derivation of utility
 416 metrics.

417 Table 7 reports counterfactual evaluation results using PO-Eval, where we leverage a pre-trained
 418 parametric IPS-Oracle to simulate user clicks and assess ranking quality. The IPS-based utility
 419 $Pr(\#Clicks \geq 1)$ captures the expected probability of at least one click per ranked list, while
 420 $NDCG_{click}$ measures how high are the originally clicked items in the test dataset ranked. The
 421 *Upper-Bound* is computed by ranking items in descending order of $P(R)$, which maximizes utility
 422 due to the rearrangement inequality (Day, 1972) (see Appendix A). Traditional LTR baselines
 423 (ListNet, ListMLE, LambdaRank, PiRank), trained with per-item IPS-sampled clicks, achieve strong
 424 offline/surrogate metrics under Eq. 2 (e.g., $NDCG_{click}$) but fail to capture the true user utility in
 425 Eq. 1 (e.g., $Pr(\#Clicks \geq 1)$). URCC* yields the lowest performance, as it relies heavily on a strong
 426 *pretrained* ranker to initialize its search; without such initialization, its effectiveness diminishes
 427 significantly (see Appendix D). In particular, URCC* explores only the neighborhood of the current
 428 permutations via pairwise position swaps, which (i) induces quadratic complexity and (ii) leads
 429 to *pessimistic* exploration that can miss superior rankings outside this local region. In contrast,
 430 REWARDRANK does not require any pretrained ranker and performs counterfactual optimization
 431 directly, enabling broader exploration beyond the data rankings from logged data. For PG-Rank*,
 we observe that increasing the number of Monte Carlo samples ($MC = 1, 5, 10$) reduces variance in
 its estimates, which improves performance, albeit at the cost of longer training time (see Appendix

432
 433 **Table 1: Counterfactual and surrogate evaluation across two settings.** The table compares ranking
 434 methods under (i) **PO-Eval** and (ii) **LAU-Eval**. For each setting we report a *counterfactual* metric:
 435 $\Pr(\#\text{Clicks} \geq 1)$ for PO-Eval and $\Pr(\#\text{Purchases} \geq 1)$ for LAU-Eval, reflecting the probability of
 436 at least one positive user action; and an *offline/surrogate* metric: $\text{NDCG}_{\text{click}}$ and $\text{NDCG}_{\text{purchase}}$,
 437 respectively, computed from logged labels (Eq. 2). While most baselines achieve high surrogate
 438 scores, these gains do not consistently translate into higher counterfactual utility (Eq. 1). Notably,
 439 URCC* and PG-rank* attain competitive NDCG yet underperform on the counterfactual metric,
 440 whereas REWARDRANK delivers the highest purchase rate in LAU-Eval while remaining competitive
 441 on surrogate metrics. The `policy_in_data` row reflects the original logged ordering.
 442

443 Method	444 PO-Eval		444 LAU-Eval	
	445 <i>Counterfactual</i> (✓) $\Pr(\#\text{Clicks} \geq 1)$	446 <i>Offline</i> (✗) $\text{NDCG}_{\text{click}}$	445 <i>Counterfactual</i> (✓) $\Pr(\#\text{Purchases} \geq 1)$	446 <i>Offline</i> (✗) $\text{NDCG}_{\text{purchase}}$
Upper-Bound	0.553 ± 0.0007	—	—	—
Policy in data	0.475 ± 0.0004	0.211 ± 0.0003	0.497 ± 0.009	0.496 ± 0.009
ListNet (Cao et al., 2007)	0.523 ± 0.0007	0.376 ± 0.0002	0.521 ± 0.009	0.405 ± 0.009
ListMLE (Xia et al., 2008)	0.522 ± 0.0007	0.377 ± 0.0002	0.522 ± 0.008	0.402 ± 0.008
LambdaRank (Wang et al., 2018)	0.524 ± 0.0007	0.378 ± 0.0002	0.523 ± 0.009	0.406 ± 0.009
PiRank (Swezey et al., 2021)	0.525 ± 0.0007	0.378 ± 0.0002	0.528 ± 0.007	0.408 ± 0.009
URCC* (Xi et al., 2024)	0.462 ± 0.0005	0.315 ± 0.0004	0.471 ± 0.008	0.401 ± 0.007
PG-rank* (Gao et al., 2023)	0.501 ± 0.0005	0.327 ± 0.0002	0.489 ± 0.007	0.402 ± 0.008
REWARDRANK	0.536 ± 0.0007	0.370 ± 0.0002	0.561 ± 0.008	0.401 ± 0.007

453
 454 Section A.3 for details). In contrast REWARDRANK attains the highest utility under IPS-Oracle,
 455 despite slightly lower $\text{NDCG}_{\text{click}}$ than some baselines. This reflects a key distinction: proxy metrics,
 456 such as NDCG (Eqn. 2), may not fully align with the true user utility (Eqn. 1). By directly optimizing
 457 counterfactual reward, REWARDRANK better aligns with behavioral objectives beyond conventional
 458 ranking accuracy.
 459

460 **LLM-based User Simulation (LAU-Eval).** While PO-Eval captures position bias via IPS-Oracle
 461 supervision, it does not account for broader behavioral patterns such as brand bias, similarity aversion,
 462 or irrelevance bias. To complement PO-Eval and more fully assess human-centered ranking behavior,
 463 we introduce the LAU-Eval framework. In this setup, a large language model (LLM) is prompted
 464 to simulate user shopping behavior given a query and its associated product list from the Amazon
 465 KDD-Cup dataset. The prompt incorporates behavioral factors such as position bias, brand bias,
 466 irrelevance bias, and color bias (full details are provided in Appendix C.2). The LLM generates a
 467 binary purchase decision $D(\text{purchase}) \in \{0, 1\}$, which serves as the reward signal for training a
 468 reward model and optimizing rankers. For evaluation, the same prompt is used: each ranker’s ranked
 469 item list is assessed by the LLM, and performance is reported as the average purchase decision rate on
 470 a held-out test set. For LTR methods that do not rely on reward modeling, we instead use the per-item
 471 binary LLM-purchase decision as the training signal. Higher values indicate stronger alignment with
 472 human-centered behavioral criteria. Refer to the Appendix Section C.2 for implementation details.
 473

474 Under *LAU-Eval*, which measures binary purchase decisions made by the LLM, we observe clear
 475 differences across methods. The `policy_in_data` baseline (original item order) attains an
 476 average purchase rate of 0.497. Classical listwise approaches—ListNet, ListMLE, LambdaRank,
 477 and PiRank—yield only modest gains on the true utility $\Pr(\#\text{Purchase} \geq 1)$, reaching 0.500–0.513,
 478 while achieving very high scores on the offline/surrogate utility ($\text{NDCG}_{\text{purchase}}$). These LTR methods
 479 largely succeed by moving the purchased item to the top, which inflates surrogate metrics but does
 480 not faithfully capture true preferences under the LLM-Oracle, such as brand or color bias among
 481 the items, and therefore does not consistently increase purchases. This underscores the need to
 482 optimize *counterfactual utility* as the primary metric for modeling human ranking behavior. We
 483 also observe a clear mismatch between surrogate and counterfactual objectives for counterfactual
 484 baselines: both PG-rank* and URCC* attain a strong $\text{NDCG}_{\text{purchase}}$ (formulated by Eqn 2), yet *both*
 485 methods yield lower values on the counterfactual metric (purchase rate as formulated by Eqn 1).
 486 This indicates that optimizing the ranking-aware surrogate alone can overfit to list reshuffling (e.g.,
 487 moving a known purchased item to the top) without improving the actual decision outcome measured
 488 by $\Pr(\#\text{Purchase} \geq 1)$.
 489

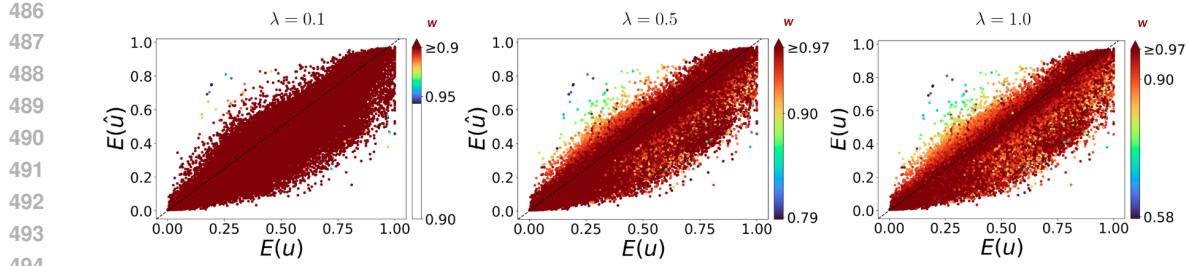


Figure 3: **Reward misspecification correction on PO-Eval.** Each point represents a ranked list with true utility (u : estimated by IPS-Oracle) and predicted utility (\hat{u} : estimated by utility model) from the ranker. Colors indicate ($w = 1 - \lambda |u_{\text{logged}} - \hat{u}_{\text{logged}}|$), showing how increasing (λ) down-weights overconfident or misaligned samples to emphasize well-calibrated predictions.

Ablations. We ablate the per-item regularizer and the two parameters of REWARDRANK: the SoftSort temperature τ , which controls the sharpness of the permutation approximation (8), and the misspecification correction strength λ , which down-weights rewards on QGs with high prediction error (13). Removing the auxiliary item-level reward loss (Eqn 6) decreased the final expected counterfactual utility of the learned ranker. This indicates that learning to predict the per-item feedback enhances the reward model’s generalization and hence improves downstream ranking performance. As shown in Figure 3, increasing λ progressively reduces the influence of unreliable reward estimates by lowering their instance weights, leading to more stable learning. As λ increases, the influence of low-confidence predictions (lower w) diminishes, effectively down-weighting misspecified instances. This correction improves stability by emphasizing samples with well-aligned predicted rewards. For illustration, we display soft utility scores from Eqn 17; however, all experiments use binary utility signals as defined in Eqn 16. We find that $\tau = 0.5$ and $\lambda = 0.7$ achieve the best trade-off between stability and performance. Full details of these ablations are reported in Appendix D.

5.2 BAIDU-ULTR DATASET WITH REAL USER CLICKS

While we previously used the Baidu-ULTR dataset within the PO-Eval framework under IPS-Oracle supervision, here we instead rely directly on the real click signals provided in the data.

Table 2: **Baidu-ULTR with real clicks.** REWARDRANK achieves SOTA performance. \dagger metrics are taken from Hager et al. (2024)

Method	DCG _{rel} @5	DCG _{rel} @10
Point IPS [†] (Hager et al., 2024)	4.79	7.43
List IPS [†] (Hager et al., 2024)	5.20	7.88
LambdaRank [†] (Hager et al., 2024)	5.45	8.23
ListNet (Cao et al., 2007)	5.05	7.64
ListMLE (Xia et al., 2008)	5.13	7.88
PiRank (Swezey et al., 2021)	5.23	8.01
URCC* Xi et al. (2024)	5.01	7.44
PG-rank* Gao et al. (2023)	5.09	7.62
REWARDRANK	5.83	8.42

highlight both the robustness of our approach and its ability to generalize to real human feedback in large-scale search settings.

6 CONCLUSION

We present REWARDRANK, a counterfactual ranking framework that directly optimizes a behaviorally grounded utility instead of relying on proxy click-based surrogates. Notably, our approach accom-

¹We report DCG rather than NDCG for consistency with (Hager et al., 2024)

540 plishes this without imposing any explicit modeling assumptions. Architecturally, REWARDRANK
 541 uses *SoftSort* to produce a differentiable soft permutation matrix, enabling end-to-end learning with
 542 *soft item embeddings* (convex combinations over items) that feed a utility model. To guard against
 543 reward model misspecification, we include a *misspecification regularization* term which is an explicit
 544 λ -weighted correction that penalizes over-reliance on noisy preference signals and stabilizes updates
 545 against spurious gains. Through the proposed *PO-Eval* and *LAU-Eval* protocols, we showed a systematic
 546 mismatch between offline/surrogate metrics (e.g., $\text{NDCG}_{\text{purchase}}$) and true decision outcomes, and
 547 demonstrated that REWARDRANK achieves the highest purchase rates while remaining competitive
 548 on surrogate metrics. Unlike URCC*, REWARDRANK does *not* require a pretrained ranker and
 549 can leverage sessions without purchase labels, extracting useful signal in sparse-feedback regimes.
 550 Ablations further indicate that auxiliary per-item losses (including on purchase-free QGs) provide
 551 consistent, moderate gains. Overall, aligning training and evaluation with counterfactual utility yields
 552 models that better capture decision-relevant user behavior than traditional LTR or locally exploratory
 553 counterfactual baselines.

554 REFERENCES

556 Aman Agarwal, Kenta Takatsu, Ivan Zaitsev, and Thorsten Joachims. A general framework for
 557 counterfactual learning-to-rank. In *Proceedings of the 42nd International ACM SIGIR Conference
 558 on Research and Development in Information Retrieval*, pp. 5–14, 2019.

559 Qingyao Ai, Kai Bi, Yiqun Zhang, Yiqun Liu, and Maosong Ma. Unbiased learning to rank with
 560 unbiased propensity estimation. In *Proceedings of the 41st International ACM SIGIR Conference
 561 on Research and Development in Information Retrieval*, 2018.

562 Qingyao Ai, Tao Yang, Huazheng Wang, and Jiaxin Mao. Unbiased learning to rank: online or
 563 offline? *ACM Transactions on Information Systems (TOIS)*, 39(2):1–29, 2021.

564 Alexander Buchholz, Ben London, Giuseppe Di Benedetto, Jan Malte Lichtenberg, Yannik Stein, and
 565 Thorsten Joachims. Counterfactual ranking evaluation with flexible click models. In *Proceedings
 566 of the 47th International ACM SIGIR Conference on Research and Development in Information
 567 Retrieval*, pp. 1200–1210, 2024.

568 Chris Burges. From ranknet to lambdarank to lambdamart: An overview. In *Microsoft Research
 569 Technical Report*, 2006.

570 Chris Burges, Tal Rago, and Quoc Viet Le. Learning to rank using gradient descent. In *ICML*,
 571 2005a.

572 Chris JC Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg
 573 Hullender. Ranknet: Gradient descent learning for rank function optimization. In *SIGIR*, 2005b.

574 Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise
 575 approach to listwise approach. In *Proceedings of the 24th international conference on Machine
 576 learning*, pp. 129–136, 2007.

577 Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Filippo Belletti, and Ed Chi. Top-k off-
 578 policy correction for a REINFORCE recommender system. In *Proceedings of the Twelfth ACM
 579 International Conference on Web Search and Data Mining*, 2019.

580 Yufei Chen, Wei Li, Jianyang Sun, and Ruiming Wang. Reducing popularity influence by addressing
 581 position bias in recommender systems. *arXiv preprint arXiv:2412.08780*, 2024. URL <https://arxiv.org/abs/2412.08780>.

582 Jack Clark and Dario Amodei. Faulty reward functions in the wild. *Internet: https://blog.openai.
 583 com/faulty-reward-functions*, 2016.

584 Thomas Coste, Usman Anwar, Robert Kirk, and David Krueger. Reward model ensembles help
 585 mitigate overoptimization. *arXiv preprint arXiv:2310.02743*, 2023.

586 Peter W Day. Rearrangement inequalities. *Canadian Journal of Mathematics*, 24(5):930–943, 1972.

594 Yufei Feng, Yu Gong, Fei Sun, Junfeng Ge, and Wenwu Ou. Revisit recommender system in the
 595 permutation prospective. *arXiv preprint arXiv:2102.12057*, 2021.
 596

597 Yoav Freund, Raj Iyer, Robert E Schapire, and Yoram Singer. An efficient boosting algorithm for
 598 combining preferences. In *JMLR*, 2003.

599 Ge Gao, Jonathan D Chang, Claire Cardie, Kianté Brantley, and Thorsten Joachim. Policy-gradient
 600 training of language models for ranking. *arXiv preprint arXiv:2310.04407*, 2023.
 601

602 Aditya Grover, Eric Wang, Aaron Zweig, and Stefano Ermon. Stochastic optimization of sorting
 603 networks via continuous relaxations. *arXiv preprint arXiv:1903.08850*, 2019.
 604

605 Shashank Gupta, Harrie Oosterhuis, and Maarten de Rijke. Practical and robust safety guarantees
 606 for advanced counterfactual learning to rank. In *Proceedings of the 33rd ACM International
 607 Conference on Information and Knowledge Management*, pp. 737–747, 2024a.
 608

609 Shashank Gupta, Harrie Oosterhuis, and Maarten de Rijke. Proximal ranking policy optimization for
 610 practical safety in counterfactual learning to rank. *arXiv preprint arXiv:2409.09881*, 2024b.
 611

612 Philipp Hager, Romain Deffayet, Jean-Michel Renders, Onno Zoeter, and Maarten de Rijke. Unbiased
 613 learning to rank meets reality: Lessons from baidu’s large-scale search dataset. In *Proceedings
 614 of the 47th International ACM SIGIR Conference on Research and Development in Information
 Retrieval*, pp. 1546–1556, 2024.
 615

616 Charlie Hou, Kiran Koshy Thekumparampil, Michael Shavlovsky, Giulia Fanti, Yesh Dattatreya, et al.
 617 Pretrained deep models outperform gbdts in learning-to-rank under label scarcity. *Transactions on
 618 Machine Learning Research*, 2024.

619 Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM
 620 Transactions on Information Systems (TOIS)*, 20(4):422–446, 2002.
 621

622 Olivier Jeunen, Ivan Potapov, and Aleksei Ustimenko. On (normalised) discounted cumulative gain
 623 as an off-policy evaluation metric for top-n recommendation. In *Proceedings of the 30th ACM
 SIGKDD conference on knowledge discovery and data mining*, pp. 1222–1233, 2024.
 624

625 Thorsten Joachims. Optimizing search engines using clickthrough data. In *KDD*, 2002.
 626

627 Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. Unbiased learning-to-rank with biased
 628 feedback. In *Proceedings of the tenth ACM international conference on web search and data
 mining*, pp. 781–789, 2017.
 629

630 Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 reinforce samples, get a baseline for free!
 631 2019.
 632

633 Ying Li, Yu Wang, and Kevin C.C. Chan. We match! building online brand engagement behaviours
 634 through social media influencers. *Journal of Retailing and Consumer Services*, 76:103640,
 635 2025. doi: 10.1016/j.jretconser.2024.103640. URL <https://www.sciencedirect.com/science/article/pii/S0969698924004429>.
 636

637 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint
 638 arXiv:1711.05101*, 2017.
 639

640 Dongyue Luo, Lixin Zou, Qingyao Ai, Zhen Chen, Dawei Yin, and Brian D. Davison. Model-
 641 based unbiased learning-to-rank. In *Proceedings of the 29th ACM International Conference on
 642 Information and Knowledge Management*, 2020.
 643

644 Jinwen Luo, Jiuding Yang, Weidong Guo, Chenglin Li, Di Niu, and Yu Xu. Matrank: Text re-ranking
 645 by latent preference matrix. In *Findings of the Association for Computational Linguistics: EMNLP
 2022*, pp. 2011–2023, 2022.
 646

647 Zechun Niu, Lang Mei, Chong Chen, and Jiaxin Mao. Distributionally robust optimization for
 648 unbiased learning to rank. In *Proceedings of the 48th International ACM SIGIR Conference on
 649 Research and Development in Information Retrieval*, pp. 2266–2275, 2025a.

648 Zechun Niu, Zhilin Zhang, Jiaxin Mao, Qingyao Ai, and Ji-Rong Wen. Investigating the robustness
 649 of counterfactual learning to rank models: A reproducibility study. In *Proceedings of the 48th*
 650 *International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp.
 651 3265–3275, 2025b.

652 Harrie Oosterhuis. Computationally efficient optimization of plackett-luce ranking models for
 653 relevance and fairness. In *Proceedings of the 44th International ACM SIGIR Conference on*
 654 *Research and Development in Information Retrieval*, pp. 1023–1032, 2021.

655 Harrie Oosterhuis. Doubly robust estimation for correcting position bias in click feedback for
 656 unbiased learning to rank. *ACM Transactions on Information Systems*, 41(3):1–33, 2023.

657 Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. Deeprank: A new deep architecture
 658 for relevance ranking in information retrieval. In *CIKM*, 2017.

659 Jialin Pei, Jiafeng Wang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. Unirank: Unifying
 660 listwise and pairwise learning to rank using deep neural networks. In *SIGIR*, 2021.

661 Przemysław Pobrotyn and Radosław Białobrzeski. Neuralndcg: Direct optimisation of a ranking
 662 metric via differentiable relaxation of sorting. *arXiv preprint arXiv:2102.07831*, 2021.

663 Sebastian Prillo and Julian Eisenschlos. Softsort: A continuous relaxation for the argsort operator. In
 664 *International Conference on Machine Learning*, pp. 7793–7802. PMLR, 2020.

665 Tao Qin, Tie-Yan Liu, and Hang Li. A general approximation framework for direct optimization of
 666 information retrieval measures. *Information retrieval*, 13:375–397, 2010.

667 Zhen Qin, Le Yan, Honglei Zhuang, Yi Tay, Rama Kumar Pasumarthi, Xuanhui Wang, Michael
 668 Bendersky, and Marc Najork. Are neural rankers still outperformed by gradient boosted decision
 669 trees? In *International Conference on Learning Representations*, 2021.

670 Chandan K Reddy, Lluís Márquez, Fran Valero, Nikhil Rao, Hugo Zaragoza, Sambaran Bandyopad-
 671 hyay, Arnab Biswas, Anlu Xing, and Karthik Subbian. Shopping queries dataset: A large-scale
 672 esci benchmark for improving product search. *arXiv preprint arXiv:2206.06588*, 2022.

673 Yuxin Ren, Qiya Yang, Yichun Wu, Wei Xu, Yalong Wang, and Zhiqiang Zhang. Non-autoregressive
 674 generative models for reranking recommendation. In *Proceedings of the 30th ACM SIGKDD*
 675 *Conference on Knowledge Discovery and Data Mining*, pp. 5625–5634, 2024.

676 Yuta Saito and Thorsten Joachims. Off-policy evaluation for large action spaces via embeddings.
 677 *arXiv preprint arXiv:2202.06317*, 2022.

678 Xiaowen Shi, Fan Yang, Ze Wang, Xiaoxu Wu, Muzhi Guan, Guogang Liao, Wang Yongkang,
 679 Xingxing Wang, and Dong Wang. Pier: Permutation-level interest-based end-to-end re-ranking
 680 framework in e-commerce. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge*
 681 *Discovery and Data Mining*, pp. 4823–4831, 2023.

682 Robin Sweeney, Aditya Grover, Bruno Charron, and Stefano Ermon. Pirank: Scalable learning to rank
 683 via differentiable sorting. *Advances in Neural Information Processing Systems*, 34:21644–21654,
 684 2021.

685 Yiteng Tu, Zhichao Xu, Tao Yang, Weihang Su, Yujia Zhou, Yiqun Liu, Fen Lin, Qin Liu, and
 686 Qingyao Ai. Reinforcement learning to rank using coarse-grained rewards. *arXiv e-prints*, pp.
 687 arXiv–2208, 2022.

688 Amos Tversky and Itamar Simonson. Extremeness aversion and attribute-balance effects in choice.
 689 *Journal of Consumer Research*, 31(2):249–257, 2004. doi: 10.1086/172209. URL <https://academic.oup.com/jcr/article/31/2/249/1824942>.

690 Ali Vardasbi, Harrie Oosterhuis, and Maarten de Rijke. When inverse propensity scoring does
 691 not work: Affine corrections for unbiased learning to rank. In *Proceedings of the 29th ACM*
 692 *International Conference on Information & Knowledge Management*, pp. 1475–1484, 2020.

702 Shuheng Wang, Weijie Zhang, Yiqun Xu, and Ji-Rong Wen. Decoy effect in search interaction:
 703 Understanding user behavior. In *Proceedings of the ACM Web Conference 2025*. ACM, 2025a.
 704 doi: 10.1145/3708884. URL <https://dl.acm.org/doi/10.1145/3708884>.

705 Shuli Wang, Xue Wei, Senjie Kou, Chi Wang, Wenshuai Chen, Qi Tang, Yinhua Zhu, Xiong
 706 Xiao, and Xingxing Wang. Nlgr: Utilizing neighbor lists for generative rerank in personalized
 707 recommendation systems. *arXiv preprint arXiv:2502.06097*, 2025b.

708 Xiaojie Wang, Ruoyuan Gao, Anoop Jain, Graham Edge, and Sachin Ahuja. How well do offline
 709 metrics predict online performance of product ranking models? In *Proceedings of the 46th
 710 International ACM SIGIR conference on Research and Development in Information Retrieval*, pp.
 711 3415–3420, 2023.

712 Xuanhui Wang, Cheng Li, Nadav Golbandi, Michael Bendersky, and Marc Najork. The lambdaloss
 713 framework for ranking metric optimization. In *Proceedings of the 27th ACM international
 714 conference on information and knowledge management*, pp. 1313–1322, 2018.

715 Qiang Wu, Chris JC Burges, Krysta M Svore, and Jianfeng Gao. Adapting boosting for information
 716 retrieval measures. In *Information Retrieval*, 2010.

717 Yunjia Xi, Weiwen Liu, Xinyi Dai, Ruiming Tang, Qing Liu, Weinan Zhang, and Yong Yu. Utility-
 718 oriented reranking with counterfactual context. *ACM Transactions on Knowledge Discovery from
 719 Data*, 18(8):1–22, 2024.

720 Fei Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. Listwise approach to learning to
 721 rank: theory and algorithm. In *ICML*, 2008.

722 Zhichao Xu, Anh Tran, Tao Yang, and Qingyao Ai. Reinforcement learning to rank with coarse-
 723 grained labels. *arXiv preprint arXiv:2208.07563*, 2022.

724 Hailan Yang, Zhenyu Qi, Shuchang Liu, Xiaoyu Yang, Xiaobei Wang, Xiang Li, Lantao Hu, Han
 725 Li, and Kun Gai. Comprehensive list generation for multi-generator reranking. In *Proceedings
 726 of the 48th International ACM SIGIR Conference on Research and Development in Information
 727 Retrieval*, pp. 2298–2308, 2025.

728 Lixin Zou, Haitao Mao, Xiaokai Chu, Jiliang Tang, Wenwen Ye, Shuaiqiang Wang, and Dawei
 729 Yin. A large scale search dataset for unbiased learning to rank. *Advances in Neural Information
 730 Processing Systems*, 35:1127–1139, 2022.

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A PROOFS AND CONCEPTUAL DETAILS
757758 A.1 CLICK-BASED UTILITY FOR PO-EVAL.
759760 The IPS-Oracle simulates user clicks as a probabilistic function of both position-dependent examination
761 and item-specific relevance. Specifically, the click probability for item $x_{\pi(\ell)}$ at position ℓ under
762 ranking π is modeled as:

763
$$P(C_{q,x_{\pi(\ell)},\ell}) = P(E_\ell) \cdot \sigma(R_{q,x_{\pi(\ell)}}) \quad (14)$$

764

765 where $P(E_\ell)$ denotes the examination probability at position ℓ , and $\sigma(R_{q,x_{\pi(\ell)}})$ is the probability of
766 a click given examination. Given a query group $(q, \{x_\ell\}_{\ell=1}^L, \pi)$, the click indicator for each item is
767 sampled as:

768
$$c_{q,x_{\pi(\ell)},\ell} \sim \text{Bernoulli}(P(C_{q,x_{\pi(\ell)},\ell})) \quad (15)$$

769

770 We define the *group-level utility* under the logged policy as a binary signal indicating whether at least
771 one item in the list was clicked:

772
$$U(q, \{x_\ell\}, \pi) = \begin{cases} 1, & \text{if } \sum_{\ell=1}^L c_{q,x_{\pi(\ell)},\ell} > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

773

774 The corresponding observed utility in the dataset, u , is a realization of $U(q, \{x_\ell\}, \pi_{\log})$ under the
775 logged ranking π_{\log} .776 To obtain a differentiable approximation, we define the expected probability of at least one click as:
777

778
$$U_{\text{IPS}}(q, \{x_\ell\}, \pi) = 1 - \prod_{\ell=1}^L (1 - P(E_\ell) \cdot \sigma(R_{q,x_{\pi(\ell)}})) \quad (17)$$

779

780 This smoothed utility represents the expected engagement for ranking π and serves as a continuous
781 training signal. The reward model is trained to predict the binary group-level utility $u \in \{0, 1\}$ from
782 the logged policy, while the ranker maximizes the expected soft utility U_{IPS} under its own predicted
783 rankings. This formulation bridges synthetic click modeling with realistic counterfactual feedback,
784 enabling effective utility-based optimization even without direct supervision on full permutations.
785786 A.2 IDEAL IPS-ORACLE: REARRANGEMENT INEQUALITY
787788 **Theorem 1** (Ideal Ranking Maximizes Utility via Rearrangement Inequality). *Let $\mathbf{r} = (r_1, \dots, r_n) \in \mathbb{R}_{\geq 0}^n$ be a vector of predicted relevance scores, and let $\mathbf{e} = (e_1, \dots, e_n) \in \mathbb{R}_{\geq 0}^n$ be a non-increasing
789 sequence of examination probabilities: $e_1 \geq e_2 \geq \dots \geq e_n$. Let π^* be the permutation that sorts \mathbf{r}
790 in descending order: $r_{\pi^*(1)} \geq r_{\pi^*(2)} \geq \dots \geq r_{\pi^*(n)}$. Then, for any permutation $\pi \in \mathcal{S}_n$, we have:*

791
$$\sum_{i=1}^n e_i \cdot r_{\pi^*(i)} \geq \sum_{i=1}^n e_i \cdot r_{\pi(i)}$$

792

793 *Proof.* This is a direct consequence of the classical rearrangement inequality (Day, 1972). Among all
794 permutations π of the relevance scores, the weighted sum $\sum_i e_i \cdot r_{\pi(i)}$ is maximized when the $r_{\pi(i)}$
795 are ordered in the same way as the e_i , i.e., both decreasing. Hence, sorting \mathbf{r} in descending order and
796 aligning it with the already sorted \mathbf{e} gives the maximal utility. \square 802 Above analysis shows that ideal ranking order under the IPS Oracle is ordering the items such the
803 sorting of item relevance scores and examination probabilities result in the same permutation.
804805 A.3 DETAILS OF BASELINES
806807 A.3.1 PG-RANK* : PG-RANK WITH LEARNED REWARD MODEL.
808809 We extend the PG-Rank framework (Gao et al., 2023) by replacing the handcrafted reward (e.g.,
810 NDCG) with a learned reward model $g(q, \{i\}_L, \pi)$ that scores entire permutations based on user

utility. The goal is to maximize the expected reward under the Plackett–Luce distribution induced by the ranker’s scores:

$$\mathcal{L}_{\text{PG-reward}}(\theta) = \mathbb{E}_{\pi \sim \mathbb{P}_\theta} [g(q, \{i\}_L, \pi)] \quad (18)$$

where $\mathbb{P}_\theta(\pi)$ is the Plackett–Luce distribution over permutations, parameterized by model scores s_1, \dots, s_L for each item in the query group. To enable backpropagation through the sampled permutations, we adopt the Gumbel-Softmax trick as in the original PG-Rank implementation, which provides a continuous relaxation of the discrete sampling process.

The gradient of this objective is estimated using the REINFORCE trick with a baseline b for variance reduction (adopted from PG-rank (Gao et al., 2023; Kool et al., 2019)):

$$\nabla_\theta \mathcal{L}_{\text{PG-reward}} \approx \frac{1}{K} \sum_{k=1}^K \left[(g(\pi^{(k)}) - b) \cdot \nabla_\theta \log \mathbb{P}_\theta(\pi^{(k)}) \right] \quad (19)$$

where $\pi^{(k)} \sim \mathbb{P}_\theta$ are K Monte Carlo samples drawn from the Plackett–Luce model.

The log-probability of a sampled permutation π under this model is given by:

$$\log \mathbb{P}_\theta(\pi) = \sum_{k=1}^L \left[s_{\pi(k)} - \log \sum_{j=k}^L \exp(s_{\pi(j)}) \right] \quad (20)$$

This formulation allows us to train the ranking model directly on learned, utility-aligned reward signals using fully differentiable, sample-based policy gradients.

A.3.2 URCC* WITH LEARNED REWARD MODEL.

URCC (Xi et al., 2024) proposes a two-stage counterfactual reranking framework that jointly learns a set-aware utility function and a context-aware reranker. The utility model in URCC is itself learned from data and used to guide the optimization of the reranker via a pairwise ranking loss over permutations. Since the official implementation of URCC is not publicly available, we re-implemented the method using our own architecture.

In our version of URCC*, we retain the core two-stage structure but implement the utility model $g(q, \{i\}_L, \pi)$ as a Transformer-based encoder trained to predict user utility over full permutations. Given a query q and a set of items $\{i\}_L$, the reward model assigns a scalar score to a permutation π :

$$S_\pi = g(q, \{i\}_L, \pi) \quad (21)$$

Following URCC, we then train the ranker f_θ to maximize this learned reward by optimizing a context-aware pairwise loss. For a pair of permutations (π^+, π^-) such that $g(q, \{i\}_L, \pi^+) > g(q, \{i\}_L, \pi^-)$, we minimize the following objective:

$$\mathcal{L}_{\text{URCC-reward}}(\theta) = \mathbb{E}_{(\pi^+, \pi^-) \sim \mathcal{P}} [\log (1 + \exp(-(S_{\pi^+} - S_{\pi^-})))] \quad (22)$$

Here, \mathcal{P} denotes the set of sampled permutation pairs with preference orderings induced by the reward model. Our implementation uses neighborhood-based sampling (e.g., pairwise swaps) to construct π^+ and π^- from the base ranking.

Thus, while our training procedure is structurally consistent with the original URCC framework, we employ a more expressive Transformer-based reward model to capture user behavior better and align optimization with utility-oriented objectives.

A.4 COMPARISON OF TIME COMPLEXITY AND COUNTERFACTUAL SPACE EXPLORATION

Table 3 compares the time complexity of three methods: URCC*, PG-rank*, and REWARDRANK. The per-iteration time complexity is analyzed based on the number of calls to the reward model.

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918

Table 3: **Comparison of Time Complexity** for URCC*, PG-rank*, and REWARDRANK in term of number of calls to the reward model per iteration on Baidu-ULTR dataset.

Method	Time Complexity	Wall-Clock Time	Description
PiRank	1	~6 hours	No call to the reward model
URCC*	n^2	~34 hours	Neighborhood search, pessimistic
PG-rank*	k	~16 hours ($k = 10$)	Needs large k for convergence
RewardRank	1	~7 hours	Full counterfactual space exploration

- **URCC*** : n^2 , where n is the number of items in the list. URCC* explores the neighborhood of factual permutations, leading to quadratic complexity due to pairwise comparisons. URCC* only explores the neighborhood of factual permutations, meaning it performs limited counterfactual exploration. This approach is considered pessimistic because it does not explore the entire space of possible rankings, which could miss potentially better arrangements.
- **PG-rank*** : k , where k is the number of Monte Carlo (MC) samples. While k is typically smaller than n , PG-rank* requires large k values and variance reduction baselines to converge. PG-Rank uses Monte Carlo (MC) sampling to explore a broader counterfactual space, but this approach requires large MC samples to converge effectively. To ensure stable and accurate exploration, PG-Rank relies on variance-reduction baselines. However, it still faces challenges in accurately capturing all potential counterfactual configurations without a very large number of samples.
- **REWARDRANK**: 1, as it performs a single call to the reward model. RewardRank explores the entire counterfactual space efficiently and can focus on more certain regions with reward misspecification mitigation.

In Table 3, we also provide the overall wall-clock time to train the model under the above method for the Baidu-ULTR dataset. Each model is trained for 21 epochs.

A.5 EXTENDED RELATED WORK

Unbiased Learning to Rank (ULTR). Unbiased Learning to Rank aims to correct biases in user interaction data such as position bias and examination bias. Classical methods rely on inverse propensity weighting or dual learning schemes to debias logged clicks Joachims et al. (2017); Ai et al. (2018). Model-based ULTR introduces explicit click models to recover unbiased relevance estimates Luo et al. (2020). A recent survey provides a comprehensive comparison of online and offline ULTR frameworks Ai et al. (2021), and more recent work applies distributionally robust optimization to improve stability under click-model misspecification Niu et al. (2025a). A reproducibility study further examines the robustness of counterfactual LTR methods under various click models Niu et al. (2025b). These ULTR methods focus primarily on correcting biased *relevance* signals under specific click models, whereas our work aims to learn and optimize a *utility* function that captures list-level behavior beyond pointwise relevance.

Reinforcement Learning To Rank (RLTR). RL-based ranking methods optimize listwise rewards using stochastic policies and importance weighting. Coarse-grained RLTR methods learn from session-level reward signals Tu et al. (2022), and other approaches apply policy gradients to optimize top- k or click-based objectives Chen et al. (2019). However, RLTR methods require behavior-policy estimation, suffer from high variance, and cannot evaluate permutations outside the support of the logged policy. In contrast, our framework evaluates arbitrary permutations deterministically through SoftSort, without requiring a stochastic policy or behavior-policy estimation.

Utility-Based and List-Level Modeling. Utility-based or list-level modeling has a long history in learning-to-rank, especially in RL-based frameworks that optimize listwise reward signals directly Tu et al. (2022). Beyond classical RLTR, several recent approaches focus explicitly on modeling list-level interactions, such as multi-generator reranking systems that capture global list structure Yang et al. (2025) and utility-oriented reranking models that aim to directly optimize user utility rather

than relevance Xi et al. (2024). Policy-gradient-based ranking approaches such as PG-Rank Gao et al. (2023) and PiRank Swezey et al. (2021) similarly optimize listwise objectives using differentiable sorting or sampled permutations. Other differentiable listwise frameworks—such as NeuralNDCG Pobrotyn & Białobrzeski (2021), stochastic relaxations of sorting networks Grover et al. (2019), and permutation-level re-ranking systems Shi et al. (2023); Feng et al. (2021); Ren et al. (2024)—also emphasize modeling dependencies across the entire slate.

In counterfactual settings, recent work has explored utility-aware optimization under safety constraints Gupta et al. (2024b;a) and introduced affine or doubly-robust corrections to improve stability of listwise estimators Vardasbi et al. (2020); Saito & Joachims (2022). However, these methods typically rely on stochastic exploration, importance weighting, or sampling-based reranking to traverse the space of permutations. In contrast, our approach learns a parametric utility model that can score arbitrary counterfactual permutations and couples this with deterministic SoftSort-based optimization. This enables full-permutation exploration without the variance, behavior-policy dependence, or sampling overhead characteristic of RL-based and counterfactual reranking approaches.

Transformer-Based Ranking Models. Transformer architectures have become standard in large-scale retrieval, re-ranking, and recommendation systems. Modern neural rankers such as Non-Autoregressive Re-ranking models Ren et al. (2024), UniRank Pei et al. (2021), DASALC Qin et al. (2021), DeepRank Pang et al. (2017), and MatRank Luo et al. (2022) leverage self-attention or attention-inspired mechanisms to capture rich query–item and item–item interactions. These models typically optimize relevance-oriented objectives or matching scores at the item level. Our work differs in focus: instead of modeling pointwise or pairwise relevance, we use a transformer to parameterize a list-level *utility* function and combine it with a differentiable ranking operator, allowing direct optimization toward counterfactual utility rather than traditional relevance-based objectives.

Counterfactual Evaluation Protocols. ULTR research often evaluates on BaiduULTR using simulated click models such as PBM, DCM, and cascade models Ai et al. (2021); Niu et al. (2025b). These simulations test bias correction under predefined user-behavior assumptions. In contrast, PO-Eval provides a parametric oracle for evaluating counterfactual utility, and LAU-Eval uses LLM-based list-level assessments to capture behavioral preferences (e.g., redundancy aversion, brand consistency) that cannot be expressed with pointwise human labels. These evaluation methods complement, rather than replace, traditional ULTR simulations by focusing on list-level *utility* rather than click-model fidelity.

While prior model-based ULTR and coarse-grained RL methods also learn reward estimates, they differ from our approach in key ways. Model-based ULTR focuses on bias-correcting click signals and then optimizes standard relevance-based objectives, whereas RewardRank learns a permutation-aware utility function that captures list-level behavioral effects beyond relevance and directly optimizes this utility through a differentiable soft-permutation operator. Coarse-grained RL approaches still rely on stochastic policies and importance weighting, limiting them to permutations explored by the behavior policy; RewardRank removes this dependence entirely by using a deterministic SoftSort-based optimization that can evaluate and optimize any permutation. Finally, prior ULTR methods optimize traditional metrics on BaiduLTR, while RewardRank introduces a misspecification-robust objective and two counterfactual evaluation suites (PO-Eval and LAU-Eval). These components provide capabilities not present in existing work and lead to consistent improvements in our experiments.

B EXPERIMENTATION DETAILS

B.1 DATASETS

Baidu-ULTR Reranking Dataset. The Baidu-ULTR dataset (Hager et al., 2024), a large-scale subset of the Baidu-ULTR corpus (Zou et al., 2022), contains user click interactions over web search queries. It includes 1.8M query groups (11.7M query-document pairs) and 590K validation/test sessions (4.8M pairs). The authors of (Hager et al., 2024) provide BERT-based CLS embeddings for each query-document pair.

We use the large-scale reranking dataset introduced by (Hager et al., 2024): publicly available at: https://huggingface.co/datasets/philippager/baidu-ultr_

972 uva-mlm-ctr, derived from the original Baidu-ULTR corpus (Zou et al., 2022). This dataset is
 973 constructed from real-world user interactions on Baidu’s production search engine and is designed to
 974 support robust evaluation of learning-to-rank models in counterfactual settings.
 975

976 Each session consists of a user query, a candidate list of documents retrieved by an upstream ranker,
 977 the original presented ranking, and user interaction logs (e.g., clicks and dwell time). For each
 978 query-document pair, the dataset provides both sparse lexical features (e.g., BM25, TF-IDF, query
 979 likelihood) and dense semantic representations.
 980

981 To generate the dense features, the authors pretrain a BERT-style model, referred to as MonoBERT,
 982 from scratch using masked language modeling (MLM) on the full Baidu corpus. This model is trained
 983 in a mono-encoder configuration and outputs a [CLS] token embedding for each query-document pair.
 984 These CLS embeddings are included in the dataset and serve as fixed, high-quality dense features for
 985 downstream reranking. The pretrained MonoBERT model and inference code are publicly available
 986 at: <https://github.com/philippager/baidu-bert-model>.
 987

988 **Amazon KDD-cup.** The KDD-Cup dataset (Reddy et al., 2022) contains 130K queries and 2.6M
 989 annotated query-product pairs in English, Japanese, and Spanish. Each query is linked to up to 40
 990 products with rich textual metadata (titles, descriptions, bullet points), making it well-suited for LLM-
 991 based evaluation, unlike Baidu-ULTR. Although the presentation order is not recorded, the dataset
 992 primarily consists of relevant query-product pairs that were shown to users. For training, validation,
 993 and testing, we sample five random permutations of length 8 per query, resulting in 400,000 training
 994 and 50,000 validation/test groups. We use the English subset of the product search dataset released as
 995 part of the KDD Cup 2022 challenge (Reddy et al., 2022), which contains real-world queries and
 996 associated candidate products from Amazon. Each query-product pair is annotated using the ESCI
 997 labeling scheme: Exact match, Substitute, Complement, or Irrelevant.
 998

999 Each query group is identified by a unique `query_id` and paired with 10–40 product candidates.
 1000 For each product, the dataset provides structured metadata including:
 1001

- 1002 • `product_title`,
- 1003 • `product_brand`,
- 1004 • `product_color`
- 1005 • `product_description`,
- 1006 • `product_bullet_point` (optional fields)
- 1007 • `product_id`,
- 1008 • `product_locale`, and
- 1009 • ESCI relevance label

1010 To construct our training and evaluation sets, we sample 5 random permutations of length 8 from
 1011 each query group. Note that we do not use the human-annotated ESCI labels provided in the
 1012 dataset. Instead, we leverage the LLM’s capability for contextual understanding to generate relevance
 1013 labels automatically. Ideally, the relevance judgments produced by the LLM should align closely
 1014 with those of human annotators. This yields approximately 392K query groups for training and
 1015 20K for validation, and 20K for testing. For a given query group, we encode each query-item
 1016 pair into sentence embeddings using the all-MiniLM-L6-v2: <https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2> model from Sentence Transformers. The
 1017 input format for the sentence transformer is constructed as:
 1018

1019 1020 **Table 4: Statistics of the Baidu-ULTR reranking dataset** (Hager et al., 2024).
 1021

1022 <code>Split</code>	1023 <code>#Query Groups</code>	1024 <code>#Query-Document Pairs</code>
1025 Training	1,857,248	11,738,489
Validation/Test	590,612	4,797,378
Total	2,447,860	16,535,867

1026 {query} [SEP] {product_title} Brand:{brand}
 1027 Color:{color}
 1028

1029 **Table 5: Statistics of the Amazon KDD Cup (ESCI) dataset (English subset).**
 1030

1031 Split	1032	1033 #Query groups	1034 #Query-Product Pairs
1033 Training	1034	78,447	627,576
1034 Validation	1035	4,000	32,000
1035 Test		4,000	32,000
1036 Total		86,447	691,576
1037 Total (including 5 random permutations)		392,235	3,137,880

1040 **B.2 IMPLEMENTATION DETAILS**

1041 We use a transformer architecture for both the reward model and the ranker across all methods to
 1042 ensure a consistent architectural backbone. The model contains 12 transformer layers, 768 hidden
 1043 dimensions, 12 attention heads, and approximately 110M parameters. All models are trained with a
 1044 learning rate of 2×10^{-5} using the AdamW optimizer (Loshchilov & Hutter, 2017) with a weight
 1045 decay of 10^{-2} . We use a batch size of 512 and train for 21 epochs, applying a learning rate decay at
 1046 epoch 12 via a step-based learning rate scheduler. All experiments are conducted using 2 NVIDIA
 1047 A100 GPUs (40GB each).

1048 For our method, REWARDRANK, we use a soft permutation temperature $\tau = 0.5$ and reward
 1049 correction term $\lambda = 0.7$. In the PG-rank* baseline, which replaces the handcrafted NDCG utility with
 1050 our learned reward model, we apply Gumbel-Softmax sampling with temperature 0.1 to approximate
 1051 permutation sampling from the Plackett–Luce distribution. We report PG-rank* results for different
 1052 Monte Carlo samples (MC = 1, 5, 10) to evaluate variance in reward estimation.

1053 In our URCC* implementation, we follow the original two-stage design: a set-aware utility model
 1054 and a pairwise ranker. The utility model is trained with a binary cross-entropy loss computed over
 1055 per-item logits derived from the transformer encoder outputs. Specifically, for each item in the
 1056 permutation, we pool its embedding from the encoder, apply dropout, and project it through a shared
 1057 per-item classifier. The per-item predictions are matched to click labels, and their aggregated loss
 1058 forms the utility supervision.

1059 As an additional baseline, we include a Naive-ranker trained with a relaxed NDCG objective following
 1060 the PiRank formulation (Swezey et al., 2021), allowing listwise supervision using soft permutation
 1061 matrices. All baselines are trained using the same reward data and input embeddings to isolate the
 1062 impact of the learning objective.

1063 Representative code for our implementations of REWARDRANK, PG-rank*, URCC* baselines, and
 1064 evaluation procedures is included in the supplementary material.

1066 **C COUNTERFACTUAL EVALUATION PROTOCOLS**1069 **C.1 PO-EVAL DETAILS**

1070 PO-Eval provides a click-based framework for counterfactual evaluation of ranking models. Using the
 1071 pre-trained Inverse Propensity Scoring model (IPS-Oracle) (Hager et al., 2024)² on the Baidu-ULTR
 1072 dataset, it generates soft click probabilities for items in a ranked query group. These probabilities
 1073 serve as counterfactual labels, enabling the evaluation of how effectively a ranker can model user
 1074 engagement patterns reflected in clicks.

1075 As the Baidu-ULTR dataset is derived from user interaction logs, click activity is heavily concentrated
 1076 in the top-ranked positions, reflecting strong position bias (see Figure 4b). In contrast, the distilled
 1077 soft utility (\hat{y}) generated by the IPS-Oracle exhibits a more uniform distribution across positions

1078
 1079 ²<https://github.com/philippbayer/baidu-bert-model>

(Figure 4a), indicating that the oracle has successfully learned to correct for position bias. Under the PO-Eval protocol, ranking methods aim to implicitly learn position debiasing from the IPS-Oracle’s soft utility, as indicated by high $U_{\text{IPS-O}}(q, \{i\}_L, \hat{\pi})$.

Training and evaluating ranking schemes. Using the learned reward model, any ranker f can be optimized via the reward maximization objective defined in Eqn 12. To evaluate its performance under the IPS-Oracle, we define the following metric: Given a query group $(q, \{i\}_L)$ and predicted relevance scores $s = [s_1, \dots, s_L]$, the induced permutation is $\hat{\pi} = \text{argsort}(s)$. For each position $\hat{\pi}_\ell$, the examination probability is $P(E_{\hat{\pi}_\ell})$, and the associated relevance score R_{q, i_ℓ} is provided by the IPS-Oracle. The overall utility is computed as the probability of at least one click: $U_{\text{IPS}}(q, \{i\}_L, \hat{\pi})$, which serves as the primary evaluation metric. It reflects how well f aligns with the user behavior modeled by the IPS-Oracle; higher values indicating better alignment. Additionally, we report $\text{NDCG}_{\text{rel}}@10$, which measures how much the predicted ranking respects the relevance scores $R_{q, i}$.

We incorporate the examination probabilities from (Hager et al., 2024), which are defined as:

$$P(E) = \{1 : 1.0000, 2 : 0.6738, 3 : 0.4145, 4 : 0.2932, 5 : 0.2079, 6 : 0.1714, 7 : 0.1363, 8 : 0.1166\}$$

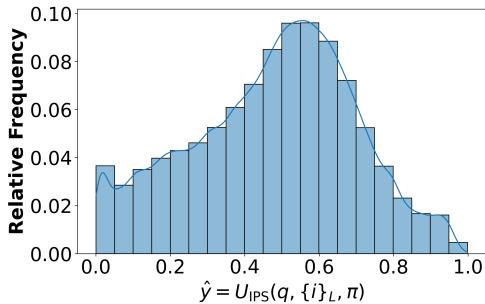
1097 C.2 LAU-EVAL DETAILS

We use Claude 3.5 Sonnet v2 with a temperature of 0.5 and a context window of 5,000 tokens. The LLM is prompted using a consistent instruction template, as illustrated in Figure 6. To evaluate a ranker with LAU-Eval, its predicted scores are converted into item positions, which are then used to reorder the input list. This reordered list is passed to the LLM alongside the original query, and the LLM outputs a binary decision regarding purchase. We include representative query groups and the corresponding LLM responses to demonstrate this pipeline.

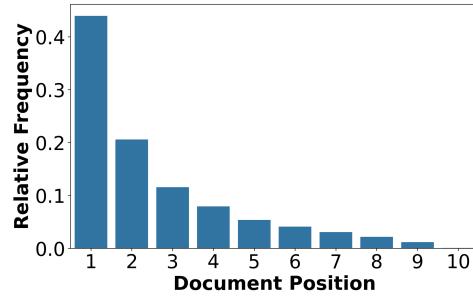
To assess the robustness of LAU-Eval under different sampling conditions, we examine how varying the temperature of the LLM decoding process affects its outputs. Figure 5 shows the distributions of LLM-simulated purchase decisions and selected item positions at temperatures 0.1, 0.5, and 0.75. While purchase rates exhibit slight variation, the LLM consistently favors top-ranked items—reflecting realistic user behavior in shopping scenarios.

Instruction prompt for LLM. We design the LLM-Eval instruction to incorporate behavioral biases such as position bias, brand preference, irrelevance filtering, similarity aversion, and color bias, guiding the LLM to consider both relevance and context-dependent preferences. Given a query and an ordered product list, the LLM estimates (i) the probability of purchasing at least one item and (ii) the selected item, without explicit relevance constraints. We illustrate the instruction prompt using an example from the Amazon KDD-Cup dataset (Reddy et al., 2022), as shown in Figure 6.

Ranking Evaluations. We present the LLM’s response to the initial list in Figure 7, including the full reasoning behind the response. It is noteworthy how the LLM is able to reason about the biases



(a) Soft-utility distribution on Baidu-ULTR generated by IPS-Oracle computed using Eqn 17.



(b) Position distribution of clicks in Baidu-ULTR.

Figure 4: Distributions extracted from IPS-Oracle analysis on Baidu-ULTR.

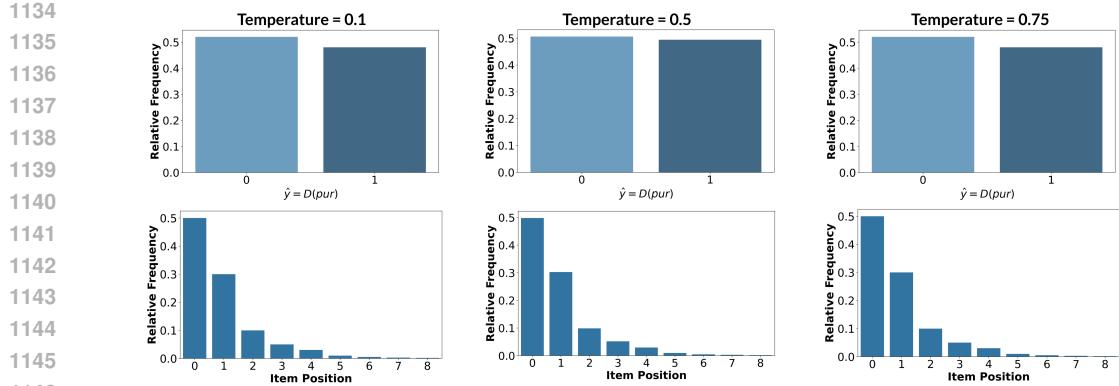


Figure 5: **Effect of Sampling Temperature on LLM-Simulated Behavior in LAU-Eval.** We visualize the distribution of binary purchase decisions (top) and item positions (bottom) generated by Claude Sonnet 3.5 v2 under three sampling temperatures: 0.1, 0.5, and 0.75. Each sample corresponds to a ranked list generated during LAU-Eval. As temperature increases, the purchase signal slightly diversifies, while positional biases remain consistent across settings. These results suggest that LAU-Eval is robust to moderate sampling variability, with LLMs producing stable user-like behavior under soft prompting.

Table 6: **Counterfactual vs. surrogate evaluation of ranking methods (LAU-Eval only).** We report performance on the true counterfactual utility (Eq. 1) and offline/surrogate metrics (Eq. 2). While most methods score highly on surrogate metrics, these gains often fail to align with true user utility.

Method	LAU-Eval		
	Counterfactual (✓)	Offline (✗)	
	Pr(#Purchases ≥ 1)	NDCG _{purchase}	NDCG _{ESCI}
Policy in data	0.497 \pm 0.009	0.496 \pm 0.009	0.995 \pm 0.009
ListNet (Cao et al., 2007)	0.521 \pm 0.009	0.405 \pm 0.009	0.8611 \pm 0.009
ListMLE (Xia et al., 2008)	0.522 \pm 0.008	0.402 \pm 0.008	0.8610 \pm 0.003
LambdaRank (Wang et al., 2018)	0.523 \pm 0.009	0.406 \pm 0.009	0.8610 \pm 0.009
PiRank (Swezey et al., 2021)	0.528 \pm 0.007	0.408 \pm 0.009	0.8623 \pm 0.005
URCC* (Xi et al., 2024)	0.471 \pm 0.008	0.401 \pm 0.007	0.8621 \pm 0.009
PG-rank* (Gao et al., 2023)	0.489 \pm 0.007	0.402 \pm 0.008	0.8630 \pm 0.009
REWARDRANK	0.561 \pm 0.008	0.401 \pm 0.007	0.8628 \pm 0.009

present in the query groups effectively. For each initial list, we also show the LLM’s response to the rearranged list generated by Claude, depicted in Figure 8. As seen, the initial arrangements in Figure 7 lead to a no purchase decision, whereas REWARDRANK generates arrangements that increase the likelihood of purchase according to the LLM. Furthermore, the LLM’s response enhances the interpretability of LLM-Eval, demonstrating how REWARDRANK’s ranking capabilities align with the LLM’s reasoning process.

Initially, we experimented with smaller language models such as Llama-3.1-8B: meta-llama/Llama-3.1-8B-Instruct and DeepSeek-R1-Distill: deepseek-ai/DeepSeek-R1-Distill-Llama-8B. However, these models were unable to generate appropriate responses to the instructions. Our experiments revealed that larger models were better at understanding the context.

It is important to note that LAU-Eval is used to simulate user behavior dynamics that may influence user decisions. Our selection of biases and instruction prompt serves as a proof-of-concept demonstrating that an LLM can be used as a proxy user to study counterfactual ranking strategies. We

1188 acknowledge that there are likely many variants of instruction prompts that could be designed to
1189 simulate user behavior. This area of exploration could be a direction for future work.
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

1242
1243**LLM Prompt (Probability Estimation Task)**

1244 You are shopping for a given query. Your task is to estimate the likelihood of purchasing any item
 1245 in a provided list. Please answer yes or no, indicating whether you wish to purchase any item from
 1246 the given list. Consider the relative relevance of items in the list when making your decisions.
 1247 Be frugal, as a typical human user would be—most users buy when the list is highly relevant,
 1248 and often make no purchase when following behavioral criteria are not met. You enter a 'query'
 1249 into the shopping system, and it returns some items mentioned in the 'products'. The items are
 1250 presented in the given order, with 1st item shown at the top of the list and the last item shown at
 1251 the bottom.

1251 Your query-products shopping list:

1252 **Query:** 11 iphone pro screen protector

1253 **Products:**

```
1254 { "B09CGJ8RW1": "title": "JETech Screen Protector and Camera Lens",  

  1255   "brand": "JETech", "color": "Transparent",  

  1256   "B07515P7PT": "title": "JETech Screen Protector for iPhone 11 Pro",  

  1257     "brand": "JETech", "color": "Clear",  

  1258   "B075S8V728": "title": "Ailun for Apple iPhone 11 Pro/iPhone",  

  1259     "brand": "Ailun", "color": "NA",  

  1260   "B07STC633H": "title": "UNBREAKable Screen Protector for iPhone 11",  

  1261     "brand": "UNBREAKable", "color": "NA",  

  1262   "B07D6XR7FM": "title": "TETHYS Glass Screen Protector for iPhone 11",  

  1263     "brand": "TETHYS", "color": "Transparent",  

  1264   "B073DLZWX7": "title": "Maxboost Screen Protector for Apple iPhone",  

  1265     "brand": "Maxboost", "color": "Clear",  

  1266   "B07FP41MC5": "title": "Trianium (3 Packs) Screen Protector",  

  1267     brand: "Trianium", color: "Clear",  

  1268   "B09BQRWG15": "title": "YRMJK Screen Protector Compatible iPhone",  

  1269     brand: "YRMJK", color: "NA"}
```

1268 **Relevance Score:** The relevance score shows how relevant the item is given the query. For every
 1269 query-item pair, it is a numerical value between 0 and 1. You should consider the following
 1270 criteria:

1. **Position bias:** where the items appearing near the top are more likely to be clicked. The position score decreases based on the following examination probabilities: $\text{position_scores} = \{ 1: 1.0000, 2: 0.6738, 3: 0.4145, 4: 0.2932, 5: 0.2079, 6: 0.1714, 7: 0.1363, 8: 0.1166 \}$ If the relevant item is not near the top, it will reduce the probability of purchase irrespective of its relevance.
2. **Brand bias:** If items from the same brand are placed adjacent to each other, the user is less likely to make a purchase. High brand bias means adjacent items are from the same brand.
3. **Irrelevance bias:** Multiple irrelevant items near the top reduce the chance of purchasing any item. This measures contextual dissimilarity at top positions.
4. **Color bias:** A lack of visual diversity (e.g., similar-colored items grouped together) discourages purchases.

1280 Note that high brand-bias, irrelevance-bias, or color-bias harms the user experience and should
 1281 lower the chance of purchasing. Act like a frugal user: only purchase if all criteria are satisfied;
 1282 otherwise, do not purchase.

1283 **Task:** Given a query and its candidate product list, decide whether any item should be purchased
 1284 and report the decision (no exact totals needed). Use a strict policy: purchase only if all criteria
 1285 are satisfied; otherwise, output no purchase. **Decision rule:** Determine the Final decision (yes/no)
 1286 by jointly considering relevance score, position bias, brand bias, irrelevance bias, and color bias.
 1287 Proceed to purchase only if all criteria are satisfactorily met. If not, do not purchase. Be especially
 1288 cautious about brand and color biases—purchase only when the biases are very low. If the decision
 1289 is yes, select exactly one item to purchase.

1290 **Output:** $D(\text{purchase}) = \langle \text{answer} \rangle$

1291 If decision = yes, then Item to be purchased = $\langle \text{answer} \rangle$

1292

1293

1294 **Figure 6: LLM instructions for a sample taken from the Amazon KDD-Cup dataset (Reddy
 1295 et al., 2022).** A total of 8 products are displayed for the query "11 iphone pro screen protector". For
 1296 visual clarity, we have shortened the titles and product descriptions.

1296 **LLM's response to the initial item list**
 1297
 1298 Let me analyze this shopping scenario step by step:
 1299 1. **Relevance Analysis**:- Most products are screen protectors for iPhone 11 Pro/X/XS (5.8-inch),
 1300 which is highly relevant to the query - Top items (JETech, Ailun) are particularly relevant as they
 1301 specifically mention iPhone 11 Pro.
 1302 **Position Bias**:- First few positions have higher probability of being selected- Top 3 positions
 1303 (JETech, JETech, Ailun) have the highest position scores.
 1304 **Brand Bias**:- First two positions are JETech brand (high brand bias at top)- This creates a negative
 1305 shopping experience- Having same brand adjacent to each other reduces purchase likelihood.
 1306 **Irrelevance Bias**:- Most items are relevant to the query- However, some products like Trianium
 1307 (4th position) is for iPhone 11 Pro Max (6.5"), which is incorrect size- Overall irrelevance bias is
 1308 low as most products match the query.
 1309 **Color Bias**:- Many items are listed as "Clear" or "Transparent"- High color similarity among
 1310 products- This creates a moderate color bias.
 1311 **Final Analysis**: Despite high relevance of products, there are several concerns: 1. High brand
 1312 bias at top positions (JETech repeated)2. High color bias (many transparent/clear options)3. The
 1313 position bias favors top items, but those have brand repetition 4. Some size mismatches in the list.
 1314 Given these factors, especially the high brand bias at top positions and color similarity, a frugal
 1315 user would likely be hesitant to make a purchase from this list.
 1316 **D(purchase)** = no; The brand repetition at top positions and lack of diversity in colors significantly
 1317 harm the shopping experience, making it less likely for a careful shopper to make a purchase, even
 1318 though the products themselves are relevant to the query.

Figure 7: LLM response to the initial list of items shown in Figure 6.

1319
 1320 **LLM's response to the ranked list provided by REWARDRANK**
 1321
 1322 Let me analyze this shopping scenario carefully:
 1323 1. **Relevance Analysis**:- All products are screen protectors for iPhone 11 Pro/X/XS, which
 1324 matches the query perfectly- Most products offer multiple packs (2-3 pieces)- All are tempered
 1325 glass protectors- High relevance across items.
 1326 2. **Position Bias**:- Top positions have higher visibility- First three items (UNBREAKable,
 1327 TETHYS, Ailun) have strongest position advantage.
 1328 3. **Brand Bias**:- JETech appears twice (positions 5 and 6) - creates brand bias- Other brands are
 1329 well distributed- Moderate brand bias due to JETech repetition.
 1330 4. **Irrelevance Bias**:- Low irrelevance bias as all products are highly relevant to query- All
 1331 products are specifically designed for the requested phone model- Core features (tempered glass,
 1332 screen protection) are consistent.
 1333 5. **Color Bias**:- Most items are clear/transparent- Very similar color schemes across products-
 1334 High color bias due to lack of variety.
 1335 **Final Decision Analysis**:- While relevance is high and position bias favors top items- The high
 1336 color bias (all transparent) and moderate brand bias (JETech repetition) are concerning- However,
 1337 given this is a screen protector where transparency is expected/desired, color bias should be
 1338 weighted less heavily- The first three positions offer different brands and good specifications-
 1339 TETHYS (2nd position) offers 3-pack with edge-to-edge coverage and installation frame.
 1340
 1341 **D(purchase)** = yes; Item to be purchased = TETHYS Glass Screen Protector
 1342 (B07D6XR7FM)Reasoning: Selected TETHYS due to good position (2nd), differentiated brand,
 1343 comprehensive features, and 3-pack value proposition. The high relevance and acceptable biases
 1344 make this a reasonable purchase despite some concerns.

Figure 8: LLM response to the ranked list generated by REWARDRANK.

1350 Table 7: **Ablation studies for counterfactual evaluation of LTR methods.** This table presents ablation results
 1351 for RewardRank under different configurations, including variations in SoftSort temperature τ , misspecification
 1352 correction regularization λ , and the addition of the auxiliary reward loss term from Eqn 6. We also report
 1353 results for PG-rank* using different numbers of Monte Carlo samples. The counterfactual evaluation metrics are:
 1354 $\text{Pr}(\#\text{Clicks} \geq 1)$ for *PO-Eval* and $\text{Pr}(\#\text{Purchase} \geq 1)$ for *LAU-Eval*.

1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403	1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403	1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403	
	Method	PO-Eval $\text{Pr}(\#\text{Clicks} \geq 1)$	LAU-Eval $\text{Pr}(\#\text{Purchase} \geq 1)$
	Upper-Bound	0.553	-
	ListNet Cao et al. (2007)	0.523 ± 0.0007	0.521 ± 0.009
	ListMLE Xia et al. (2008)	0.522 ± 0.0007	0.522 ± 0.008
	LambdaRank Wang et al. (2018)	0.524 ± 0.0007	0.523 ± 0.009
	PiRank Swezey et al. (2021)	0.525 ± 0.0007	0.528 ± 0.007
	URCC*	0.462 ± 0.0005	0.471 ± 0.008
	PG-rank* (mc=1)	0.481 ± 0.0006	0.441 ± 0.006
	PG-rank* (mc=5)	0.495 ± 0.0005	0.465 ± 0.007
	PG-rank* (mc=10)	0.501 ± 0.0005	0.489 ± 0.007
	NAR4Rec Ren et al. (2024)	0.527 ± 0.0007	-
	GRPO Xu et al. (2022)	0.518 ± 0.0005	-
	SoftSort Temperature τ		
	RewardRank ($\tau = 0.1, \lambda = 0.0$)	0.531 ± 0.0005	0.548 ± 0.008
	RewardRank ($\tau = 0.2, \lambda = 0.0$)	0.532 ± 0.0005	0.550 ± 0.008
	RewardRank ($\tau = 0.5, \lambda = 0.0$)	0.533 ± 0.0005	0.551 ± 0.007
	RewardRank ($\tau = 0.7, \lambda = 0.0$)	0.531 ± 0.0005	0.550 ± 0.008
	RewardRank ($\tau = 1.0, \lambda = 0.0$)	0.530 ± 0.0005	0.549 ± 0.009
	Misspecification Correction λ		
	RewardRank ($\tau = 0.5, \lambda = 0.1$)	0.532 ± 0.0005	0.549 ± 0.007
	RewardRank ($\tau = 0.5, \lambda = 0.3$)	0.534 ± 0.0007	0.554 ± 0.007
	RewardRank ($\tau = 0.5, \lambda = 0.7$)	0.536 ± 0.0007	0.561 ± 0.008
	RewardRank ($\tau = 0.5, \lambda = 1.0$)	0.533 ± 0.0007	0.553 ± 0.006
	Auxiliary Per-Item Regularizer Eqn 6		
	RewardRank (reward loss = Eqn 3)	0.528 ± 0.0005	0.553 ± 0.008
	RewardRank (reward loss = Eqn 3 + Eqn 6)	0.536 ± 0.0005	0.561 ± 0.008
	Using pretrained ranker: PiRank		
	URCC*	0.521 ± 0.0005	-
	PG-rank*	0.503 ± 0.0006	-
	RewardRank	0.538 ± 0.0005	-

D FURTHER ABLATION STUDIES

We use the Baidu-ULTR dataset to study how the performance of REWARDRANK varies with two key hyperparameters: the temperature τ of the *SoftSort* operator, which controls permutation sharpness, and the regularization strength λ for reward misspecification correction introduced in Eqn 13. Varying $\tau \in \{0.1, 0.2, 0.7, 1.0\}$ shows that moderate temperature ($\tau = 0.2 - 0.5$) achieves the best utility and relevance alignment. Too low a temperature leads to unstable gradients due to near-hard permutations, while higher values oversmooth rankings, diluting learning signals. Fixing $\tau = 0.5$, we ablate the correction term with $\lambda \in \{0.0, 0.1, 0.3, 0.7, 1.0\}$. As shown in Table 7 and visualized in Figure 3, moderate correction ($\lambda = 0.5 - 0.7$) yields the best trade-off, by down-weighting unreliable samples without discarding informative ones. This results in higher IPS utility, confirming the benefit of explicitly mitigating reward misspecification.

We explore the impact of incorporating an auxiliary item-level reward loss (Eqn 6) into the training objective of the reward model. As shown in Table 7, adding this auxiliary loss to the list-level cross-entropy objective (Eqn 3) improves expected utility from 0.528 to 0.536. This indicates that learning

1404 to predict the per-item feedback as an auxiliary task enhances the reward model’s generalization and
 1405 improves the downstream utility-optimized ranking.

1406
 1407 Table 7 presents the results for the pretrained ranker, which is the ranker trained with PiRank Swezey
 1408 et al. (2021) LTR loss. URCC*, being dependent on the pretrained ranker, demonstrates larger
 1409 performance improvements. However, the gains from the pre-trained ranker are not as significant,
 1410 suggesting that URCC*’s performance is more sensitive to the quality of the pretrained model. On
 1411 the other hand, REWARDRANK and PG-rank* show limited improvements when using the pretrained
 1412 ranker, as their performance is not heavily reliant on the presence of a strong pretrained model. These
 1413 methods are more robust in their ranking capabilities and do not exhibit substantial gains from a
 1414 pretrained ranker.

1415 E INFERENCE COST AND LIMITATIONS

1416
 1417 **Inference Cost.** The main inference cost in our work arises from using large language models
 1418 (LLMs) for ranking and purchase probability estimation. These models require significant computa-
 1419 tional resources, especially for large datasets and permutations of items. Optimizations like batch
 1420 processing and multi-GPU use help manage costs, but scalability remains a challenge. Caching
 1421 frequently accessed queries can further reduce repeated computation costs.

1422
 1423 **Limitations.** While both PO-Eval and LAU-Eval provide valuable insights into ranking quality
 1424 and user preferences, there are inherent limitations in each approach. These limitations arise from
 1425 their reliance on specific biases and the quality of input data, which may affect their performance in
 1426 diverse real-world scenarios. Below, we outline the key limitations of each method:

- 1427 • **PO-Eval Limitations:** While PO-Eval provides a robust baseline for position-debiasing,
 1428 it is limited in behavioral scope. It primarily focuses on mitigating position bias without
 1429 considering other nuanced user preferences, such as brand bias or contextual relevance,
 1430 which can lead to suboptimal performance in more complex scenarios.
- 1431 • **LAU-Eval Limitations:** LAU-Eval captures richer heuristics and offers more context-aware
 1432 ranking, but it depends heavily on the quality and stability of the LLM outputs. Inconsistent
 1433 or noisy outputs from the LLM can negatively affect the reliability of the evaluation, as the
 1434 method assumes that the LLM accurately reflects user preferences in all scenarios.

1435
 1436 These limitations highlight areas for future improvement, such as incorporating additional user
 1437 behavior modeling and enhancing the robustness of the LLM outputs.