

# Does LLM Focus on the Right Words? Diagnosing Language Bias in LLM-based Recommenders

Anonymous Author(s)

## Abstract

Large language models (LLMs), owing to their extensive open-domain knowledge and semantic reasoning capabilities, have been increasingly integrated into recommender systems (RS). However, a substantial gap remains between the pre-training objectives of LLMs and the specific requirements of recommendation tasks. To address this gap, supervised fine-tuning (SFT) is commonly performed on specially curated recommendation datasets to further enhance their predictive ability. Despite its success, SFT exhibits a critical limitation: it induces **Language Bias**, whereby the model over-relies on auxiliary tokens—such as task descriptions and prefix-generated tokens—while underutilizing core user interaction tokens that encode user-specific preferences. This bias not only undermines recommendation accuracy but also raises unfairness concerns.

To address this issue, we propose **Group Distributionally Robust Optimization-based Tuning (GDRT)**, a novel fine-tuning paradigm that enforces consistent model performance across token groups with varying degrees of relevance to auxiliary tokens. By adaptively upweighting underperforming groups, typically those weakly correlated with auxiliary tokens, GDRT shifts the model’s attention from superficial auxiliary cues to informative user interaction tokens, thereby mitigating language bias. Extensive experiments conducted on three public datasets demonstrate that GDRT effectively mitigates language bias, yielding substantial improvements in recommendation accuracy (with an average NDCG@10 gain of 24.29%) and significantly enhancing recommendation fairness. The code is available at <https://anonymous.4open.science/r/GDRT-7604>.

## CCS Concepts

• Information systems → Recommender systems.

## Keywords

Sequential Recommendation; Large Language Model; Bias

## 1 Introduction

With remarkable open-domain knowledge and semantic reasoning capabilities [1, 18], Large Language Models (LLMs) have been extensively explored for integration into recommendation systems (RS) [59]. One prominent approach involves positioning LLMs as the central recommendation backbone [2, 4, 22, 27, 31, 35, 48, 52–54, 67, 69]. These methods express items as textual descriptions (e.g., titles), and construct language prompts based on users’ past interactions, which are then used to instruct LLMs to predict users’ future interactions. Figure 1 illustrates the mechanism of such LLM-based recommendation. LLMs operate at a fine-grained semantic token level, sequentially generating tokens of the predicted items by analyzing their nuanced semantic relations with the user’s previous interactions and other auxiliary information (e.g., task descriptions,

prefix tokens of the prediction). This fine-grained paradigm enables the capture of subtle semantic patterns in user interests and thus represents a promising direction for advancing recommender systems [2].

To better align LLMs with recommendation objectives and capture collaborative filtering signals, Supervised Fine-Tuning (SFT) is commonly employed [2, 4, 35, 69]. In this strategy, each prompt is paired with the target item description, and the LLM is fine-tuned to generate the correct prediction. This procedure enables the model to capture semantic relationships between prompts and targets present in the training data, often resulting in substantial performance gains.

However, we find that SFT introduces a significant **Language Bias**. Specifically, SFT drives the model to over-rely on auxiliary tokens (e.g., task descriptions or prefix tokens) while under-utilizing core interaction tokens that encode user personalized preferences as shown in Figure 1. To verify this, we conduct Feature Ablation Analysis [28, 40], a standard approach to quantify each token’s contribution to model predictions. Figure 2 shows that SFT dramatically amplifies the relative impact of auxiliary tokens while significantly suppressing the influence of interaction tokens, with the influence ratio shifting from about 1:1 before fine-tuning to more than 6:1 afterward as measured on typical Amazon datasets.

This over-reliance reveals shortcut learning: LLMs simply memorize correlations with frequently occurring auxiliary tokens rather than grounding predictions in user-specific preferences. Such bias not only undermines recommendation accuracy but also raises serious fairness concerns [8]. Specifically, this bias skews recommendations toward a narrow subset of items whose tokens exhibit higher semantic relevance to auxiliary tokens. To verify this, we divided items into five groups based on their relevance to auxiliary tokens. As shown in Figure 3, over 80% of recommended items fall into the most relevant group, while this group comprises only 20% of target items in the test set. These findings motivate our core research question: **How can we mitigate language bias in LLM-based recommenders?**

To address this, we propose a novel fine-tuning strategy GDRT, that leverages Group Distributionally Robust Optimization (Group DRO) [46] to mitigate language bias. We first group training instances according to the semantic relevance between the target token and the auxiliary tokens, which can be evaluated by the LLM’s predictive probability when user history is masked. Group DRO then enforces LLMs to perform consistently well across all groups, regardless of their relevance strength with auxiliary tokens. This optimization objective naturally shifts the LLM’s attention away from auxiliary tokens toward user-specific interaction tokens, as simple reliance on shortcut auxiliary tokens results in poor performance on groups with weaker correlations. Importantly, GDRT is easy to implement and computationally efficient, requiring only high-efficiency group construction and dynamic sample weighting

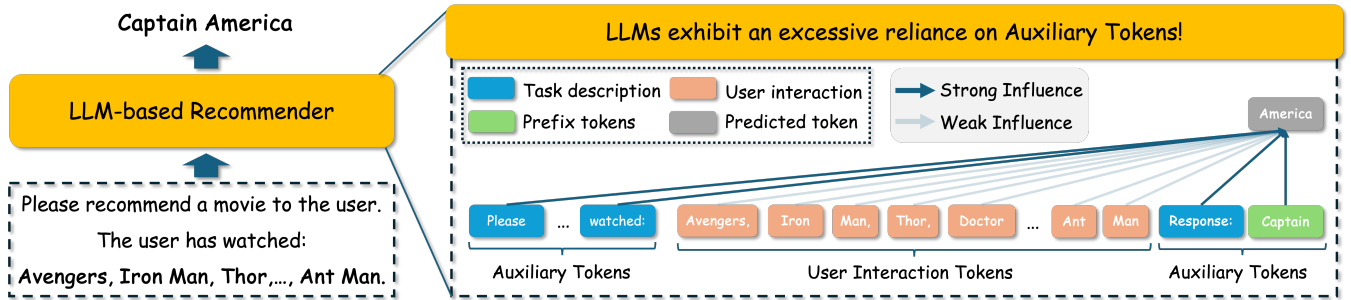


Figure 1: Illustration of LLM-based recommendations and language bias, wherein the model exhibits an over-reliance on auxiliary tokens and insufficient utilization of interaction tokens during generation.




Figure 2: Ratio of attribution values between auxiliary tokens and user interaction tokens before and after SFT. Left: task description vs. user interaction tokens. Right: prefix tokens of predicted item (take the first token) vs. user interaction tokens.

during training. It can be seamlessly integrated into various LLM-based recommenders, yielding improvements in both accuracy and fairness.

Our main contributions are:

- We provide a comprehensive empirical analysis revealing that SFT in LLM-based recommendation induces significant language bias, negatively affecting both accuracy and fairness.
- We propose GDRT, a Group DRO-based fine-tuning strategy, to effectively mitigate language bias in LLM-based recommendation.
- We conduct extensive experiments demonstrating that GDRT achieves state-of-the-art recommendation performance in both accuracy and fairness metrics.

## 2 Preliminary

### 2.1 LLM-based Recommendation

Following previous work [2, 4, 35, 36, 41, 66], this paper also focuses on sequential recommendation [25], a conventional recommendation scenario in practice. Let  $\mathcal{V}$  denote the set of items in the recommendation system. Given a user’s historical interaction sequence  $S = \{s_1, s_2, \dots, s_n\}$ , where each  $s_i \in \mathcal{V}$  represents the  $i$ -th interacted item, the goal of the RS is to predict the user’s next interaction  $s_{n+1}$  that the user is likely to interact with.

The remarkable success of LLMs across diverse domains [43, 50, 57] has spurred growing interest in their application to recommendation systems (RS) [59]. A prominent approach is to directly




Figure 3: Distribution of Top-1 recommended items generated by SFT-trained LLMs across five groups defined according to the semantic relevance of items to the auxiliary tokens (Group 1: highest relevance, Group 5: lowest relevance). We also present the distribution of the test set across the same five groups for comparison.

leverage LLMs as recommenders [31]. As shown in Figure 1, this paradigm constructs a language prompt  $x = [x^{\text{task}}, x^{\text{user}}]$ , where  $x^{\text{task}}$  represents the task description and  $x^{\text{user}}$  denotes the textual form (e.g., titles) of a user’s historical interactions. This prompt then guides the LLM to generate the descriptions of recommended items  $y$ . Notably, LLMs operate at a fine-grained semantic token level, sequentially generating tokens of the predicted items according to the model estimated probability  $P_\theta(y_t | x, y_{<t})$ , where  $y_t$  denotes the  $t$ -th predictive tokens and  $y_{<t}$  denotes the prefix tokens of the prediction. This fine-grained token-level paradigm has the potential to capture subtle semantic patterns in user preferences.

To align LLMs with recommendation objectives, supervised fine-tuning (SFT) is commonly applied, fine-tuning all or part of the model parameters using recommendation data [2, 4, 35, 69]. In this process, the training data is reorganized into a set of prompt–target pairs  $\mathcal{D} = \{(x_i, y_i^*)\}_{i=1}^N$  where each  $x_i$  is the constructed prompt and  $y_i^*$  is the textual description of the target item. The LLM is optimized with the following log-likelihood objective:

$$\mathcal{L}_{SFT} = -\frac{1}{N} \sum_{i=1}^N \sum_{t=1}^{|y_i^*|} \log P_\theta(y_{i,t}^* | x_i, y_{i,<t}^*) \quad (1)$$

where  $|y_i^*|$  denotes the token length of the target item description. SFT increases the generative probability of the target item, encouraging the LLM to capture the inherent token-level semantic




Figure 4: Proportion of Top-1 recommendations belonging to Item Group 1 (most relevant with auxiliary tokens) over the course of SFT.

correlations between each target token  $y_{i,t}^*$  and the user interactions  $x_i^{\text{user}}$ , task descriptions  $x^{\text{task}}$  and the prefix tokens  $y_{i,<t}^*$ . This process often yields substantial performance improvements [4].

## 2.2 Analyses on Language Bias

In this section, we first identify the language bias in fine-tuning LLM for recommendation, followed by discussing its negative effect. We then analyze the underlying causes of this bias and discuss why existing methods can not effectively address this issue.

**2.2.1 Empirical Evidence Demonstrating Language Bias.** We conduct Feature Ablation Attribution (FAA) [28, 40] analysis to quantify contribution of different token types to the model predictions. FAA is an attribution method used to evaluate the importance of individual input components by measuring how the model’s output changes when specific inputs are masked, and it has been widely adopted in the LLMs for interpreting token-level contributions [5, 64, 68]. A higher attribution value indicates a greater influence of input tokens on the model’s output. Figure 2 shows the ratio of attribution values between auxiliary tokens (e.g., task descriptions or prefix tokens) and user-specific interaction tokens. We report this ratio both before and after fine-tuning to enable direct comparison. From these results, we make the following observation:

**Language Bias:** *Supervised fine-tuning (SFT) can bias LLMs to over-rely on auxiliary tokens while under-utilizing core interaction tokens that encode user personalized preferences.*

Before fine-tuning, the attribution value ratio between task description and interaction tokens is approximately 1:1. After fine-tuning, this ratio is markedly amplified across all datasets (e.g., Toy: 4.69:1; Clothing: 6.58:1). A similar phenomenon is observed when comparing prefix tokens with interaction tokens. These results clearly indicate the presence of language bias, whereby the model disproportionately relies on auxiliary tokens rather than more informative interaction tokens.

**2.2.2 Negative Effect of Language Bias.** Language bias can substantially hinder the effectiveness of LLM-based recommenders, leading not only to decreased recommendation accuracy but also to pronounced unfairness issues. On the one hand, critical user-item interaction signals that capture user preferences may be ignored by the model, severely impairing its ability to deliver personalized recommendations. On the other hand, the recommendation output becomes inherently skewed toward a limited subset of items whose textual tokens exhibit strong correlations with auxiliary tokens.




Figure 5: The co-occurrence rate of different types of token pairs in the training set.

To empirically verify this phenomenon, we conduct experiments on the typical Amazon datasets. Specifically, we partition items into five groups based on their semantic relevance to auxiliary tokens, measured by the model’s estimated probability:

$$r(y) = \frac{1}{|y|} \sum_{t=1}^{|y|} \log P_{\theta}(y_t | x^{\text{task}}, y_{<t}) \quad (2)$$

Here we mask the user historical information in the original prompt template with the placeholder 'N/A', and compute the probability of the tokens for each item. This measure serves as an indicator of the relevance between the input and the output, a practice commonly adopted for estimating semantic relevance [26]. Based on this metric, items are sorted and evenly divided into five groups, where Group 1 contains the most relevant items, and Group 5 the least.

Next, we compute the proportion of Top-1 recommended items from each group, with the results shown in Figure 3. The analysis reveals that fine-tuned LLMs display a pronounced inclination for items in Group 1, whose tokens exhibit the highest relevance to auxiliary tokens. Nearly 80% of Top-1 recommendations fall within this group, despite such items comprising only about 20% of the test set. As illustrated in Figure 4, this bias becomes progressively more pronounced as the tuning proceeds. This bias not only undermines recommendation accuracy but also exacerbates exposure unfairness by systematically over-promoting items misaligned with individual user preferences. Such skew can severely degrade user experience and distort the recommendation ecosystem. For example, it would incentivize content providers to adopt clickbait-like titles or other superficial textual strategies to artificially increase token relevance to the auxiliary tokens.

**2.2.3 Origins of Language Bias.** The emergence of language bias can be traced to biases inherent in the training data. As shown in Figure 5, the co-occurrence rate between auxiliary tokens and target item tokens is significantly higher than those between interaction tokens and target item tokens. This is expected, as task prompts appear in every training instance, and prefix tokens are always accompanied for item tokens. Consequently, during fine-tuning, the LLM tends to capture shortcut patterns, simply memorizing frequent correlations with auxiliary tokens, while neglecting the more important user-specific interaction signals. This incurs language bias and motivates the need for improved fine-tuning strategies.

**2.2.4 Limitations of Related Strategies.** Several recent studies have explored bias and fairness issues in LLM-based RS. However, the

characteristics of these biases differ from those of *language bias*. Specifically, language bias arises during fine-tuning and reflects the inherent over-reliance of LLMs on auxiliary tokens. In contrast, earlier work has primarily focused on biases such as *popularity bias* [16, 17, 24, 34, 37], arising from imbalanced item frequency in training data; *position bias* [6, 8, 12, 22, 23, 38, 39, 61], caused by the model’s sensitivity to the ordering of candidate items; *length bias* [3], induced by length normalization, which makes items with longer token sequences more likely to be recommended. These biases stem from distinct sources and mechanisms. Consequently, this newly identified form of language bias warrants explicit and targeted mitigation, yet recent work has fallen short in addressing this issue (cf. Table 2).

Another related work, CFT [63], aims to encourage LLMs to better leverage user interaction information. Notably, this study is parallel to ours, currently available only on arXiv and not formally published. CFT introduces counterfactual learning to forcibly enhance the influence of interaction tokens for each training instance. However, this strategy suffers from multiple aspects of limitations: **(1) Objective misalignment.** The counterfactual objective is not directly aligned with improving recommendation accuracy or fairness, and its contribution to these aspects remains uncertain. In fact, CFT only treats this objective as an auxiliary loss and over-emphasis on this term has been observed to lead to substantial performance degradation. **(2) Difficulty in weight selection.** CFT relies on manually specified weights to determine the degree to which a training instance should rely on interaction tokens. In fact, the optimal weights can vary substantially across instances and are difficult to estimate accurately. Manual specification is therefore challenging, prone to deviation from the ideal value, and ultimately detrimental to model effectiveness. Empirically, even when we directly employ the official source code of CFT and conduct a fine-grained hyperparameter search, CFT yields only limited performance gains on some datasets (cf. Table 2). **(3) High computational overhead.** CFT requires to process counterfactual instances, resulting in more than double the training time compared with SFT (cf. Figure 12). In contrast, our GDRT is explicitly aligned with the target objective, enabling the model to consistently perform well across different instance groups. This naturally encourages greater focus on interaction tokens without requiring manual specification of influence strengths.

### 3 Methodology

The above analyses reveal a significant language bias inherent in fine-tuning LLMs for recommendation. To address this issue, we propose a novel fine-tuning framework, termed GDRT, which leverages Group Distributionally Robust Optimization (Group DRO) [46] to mitigate such bias. This section first outlines the general idea of GDRT, and then describes the group partitioning strategy and customized loss function.

**General Idea.** Rather than directly intervening the learning process to manually enhance the effect of interaction tokens, we pursue an alternative objective: *ensuring that the model achieves consistently strong performance across target tokens regardless of their degree of relevance to auxiliary tokens*. It naturally shifts the LLM’s attention away from auxiliary tokens towards user-specific




Figure 6: Loss of token groups with varying degrees of relevance to auxiliary tokens in the SFT-trained biased model.

interaction tokens, as simple reliance on shortcut auxiliary tokens results in poor performance on the instances with weak correlations as shown in Figure 6. Besides, this objective is directly aligned with the goals of high recommendation accuracy and fairness.

To implement this idea, we adopt Group DRO, which partitions the training data into multiple groups and uses an adversarial training mechanism to encourage great performance across all groups. Group DRO has been widely applied in various domains and shown effectiveness in mitigating group disparities and shortcut correlations [42, 46, 65]. In applying Group DRO to our problem, we address two key questions: (1) how to construct groups that capture varying degrees of relevance to auxiliary tokens; and (2) how to design a loss function that is computationally efficient.

**Token Grouping by Relevance with Auxiliary Tokens.** Our objective is to ensure consistent model performance across groups that differ in their degree of relevance to auxiliary tokens. Accordingly, the constructed groupings should capture variations in this relevance. Towards this end, we compute the predictive probability of each target token conditioned solely on the corresponding auxiliary tokens:

$$r(y_{i,t}^*) = \log P_{\theta} \left( y_{i,t}^* \mid x^{\text{task}}, y_{i,<t}^* \right). \quad (3)$$

where  $y_{i,t}^*$  denotes the  $t$ -th token of the target item in the  $i$ -th training sample, and  $y_{i,<t}^*$  denotes its prefix tokens. Based on these relevance scores, all target tokens are partitioned into  $G$  disjoint groups  $\{\mathcal{G}_1, \dots, \mathcal{G}_G\}$  using the K-means algorithm<sup>1</sup>, where  $G$  is a hyperparameter that determines the granularity of grouping. Each group  $\mathcal{G}_g$  contains the index  $(i, t)$  of target tokens exhibiting similar degrees of relevance to auxiliary tokens, whereas tokens belonging to different groups demonstrate distinct relevance levels.

Unlike Eq. 2, which evaluates *item-wise* relevance between the entire target item and the auxiliary tokens, Eq. 3 computes *token-wise* relevance. This distinction is crucial, as different tokens within the same target item may exhibit varying degrees of relevance to the auxiliary tokens. Such finer-grained grouping can more effectively capture variations in this relevance.

**The Objective of GDRT.** Subsequently, we employ Group DRO to enforce consistent model performance across groups with different strength of relevance. Specifically, the training objective of

<sup>1</sup>We can simply employ the *scikit-learn* package to implement the K-means algorithm. Considering that  $r(y_{i,t}^*)$  is a numerical value, it can be evenly partitioned according to its magnitude, yielding comparable results.

GDRT is defined as:

$$\mathcal{L}_{GDRT} = \max_Q \sum_{g=1}^G Q(g) \mathcal{L}(g) \quad \text{s.t.} \quad D_{KL}(Q, U) \leq \eta \quad (4)$$

$$\mathcal{L}(g) = \frac{1}{|\mathcal{G}_g|} \sum_{(i,t) \in \mathcal{G}_g} -\log P_{\theta} \left( y_{i,t}^* \mid x_i, y_{i,<t}^* \right)$$

where  $\mathcal{L}(g)$  represents the vanilla generative loss for group  $\mathcal{G}_g$ , and  $|\mathcal{G}_g|$  represents the number of target tokens in group  $\mathcal{G}_g$ . The term  $Q$  denotes the weight distribution over groups, acting as a flexible adversarial perturbation to the original empirical distribution, with  $Q(g)$  being the weight assigned to the  $g$ -th group. This weight distribution is regularized via a Kullback–Leibler divergence term  $D_{KL}(Q, U)$  between the weight distribution  $Q$  and the uniform distribution  $U$ , with the parameter  $\eta$  controlling the perturbation magnitude.

Comparing  $\mathcal{L}_{GDRT}$  (Eq. 4) with the original SFT objective  $\mathcal{L}_{SFT}$  (Eq. 1), the main difference lies in the introduction of the additional group weighting term  $Q(g)$ . Intuitively, Group DRO imposes an adversarial shift to the group distribution by perturbing the relevance distribution with respect to auxiliary tokens, thereby compelling the model to perform consistently well across groups under different adversarial re-weightings. This mechanism naturally mitigates the model’s reliance on auxiliary tokens, and encourages the model to perform consistently well across different groups.

**Efficient Implementation.** While Group DRO involves the complex adversarial optimization, it can be simplified to an equivalent closed-form objective. We have the following lemma:

LEMMA 1. Equation 4 can be reformulated as the following objective:

$$\mathcal{L}_{GDRT} = \sum_{g=1}^G Q(g) \mathcal{L}(g), \quad Q(g) = \frac{e^{\mathcal{L}(g)/\tau}}{\sum_{g'=1}^G e^{\mathcal{L}(g')/\tau}} \quad (5)$$

The parameter  $\tau$  is the dual Lagrange coefficient associated with the constraint  $D_{KL}(Q, U) \leq \eta$ .

The lemma yields an explicit closed-form solution for the group weights  $Q(g)$ . From the perspective of these weights, the effect of Group DRO can be well understood: groups with higher losses, which typically correspond to tokens exhibit weaker relevance with auxiliary tokens (see Figure 6), receive larger weights. This adaptive reweighting encourages the model to allocate greater optimization emphasis to underperforming groups, thereby enhancing its learning on samples with low auxiliary-token relevance. As a result, the model attains more balanced performance across groups and effectively mitigates its over-reliance on auxiliary tokens.

This closed-form also facilitates an efficient implementation of GDRT. Compared to SFT, the additional steps just involve: partitioning the token groups based on Eq. 3, and dynamically updating the group weights according to Eq. 5. The time complexity of GDRT is the same as the basic SFT, and our empirical experiments also demonstrate their close running time (*cf.* Figure 12). Besides, this simple reformulation makes the integration of GDRT into existing LLM-based RS straightforward, requiring only minimal code changes.

Table 1: Statistics of the datasets.

Dataset	#Users	#Items	#Interactions	#Density
Toy	19124	11758	165247	0.0735%
Clothing	39230	22948	277534	0.0308%
Book	16559	6344	151928	0.1446%

## 4 Experiments

We aim to answer the following research questions:

- RQ1: How does GDRT perform compared to existing SOTA methods?
- RQ2: Can GDRT be integrated into other LLM-based RS?
- RQ3: Does GDRT mitigate language bias?
- RQ4: How do the hyperparameters affect the performance of GDRT?
- RQ5: How does the efficiency of GDRT compare with baseline methods?

### 4.1 Experimental Settings

**4.1.1 Datasets.** Three widely used real-world datasets—*Amazon Toys and Games*, *Amazon Clothing, Shoes and Jewelry*, and *Amazon Books*—are employed in our experiments. These datasets are commonly adopted in studies on LLM-based recommenders [2, 7, 11, 29, 29, 32, 52]. To ensure fair comparison, we follow the data pre-processing procedures employed in recent literature [2, 3, 11, 52]. Specifically, we first apply the 5-core setting to the raw datasets. For user interaction sequences exceeding 11 interactions, we segment the sequences using a sliding window of length 11. The segmented sequences are then sorted in ascending order by timestamp and partitioned into training, validation, and test sets in an 8:1:1 ratio. Due to the large size of *Amazon Books*, we randomly retain 100,000 items prior to 5-core processing. The statistics of the processed datasets are summarized in Table 1.

**4.1.2 Baselines.** The methods compared fall into several categories:

- **Traditional RS (SASRec [25], SASRec++ [30], DROS [60]):** SASRec employs a self-attention-based model to effectively capture user interests. SASRec++ further improves performance using the softmax loss. DROS integrates DRO to enhance the model’s robustness against distributional shifts.
- **LLM-based RS (SFT [2], CFT [63], MSL [52], LLaRA [35], A-LLM [27]):** SFT introduces instruction-tuning templates to fine-tune LLMs on RS datasets. CFT incorporates a causal loss to strengthen the behavior sequence modeling capabilities of LLMs. MSL improves the loss function specifically tailoring it to optimize recommendation-oriented objectives. LLaRA incorporates embeddings from traditional models to enhance the ability of LLMs to capture collaborative signals. A-LLM further aligns these embeddings with their corresponding textual information.
- **Debiasing Methods for LLM-based RS (Reweight [24], SPRec [16], D3 [3]):** Reweight addresses popularity bias by balancing recommendations using pre-calculated item weights. SPRec develops negative sampling techniques based on Direct Preference Optimization (DPO) [44] to reduce popularity bias. D3 mitigates length bias through enhanced decoding strategies.

**Table 2: The performance comparison on three real-world datasets. The best result is bolded. Lower MGU and DGU indicate better fairness.**

Dataset	Metric	SASRec	DROS	SASRec++	SFT	Reweight	D3	SPRec	CFT	GDRT
Toy	NDCG@5 ↑	0.0057	0.0095	0.0120	0.0118	0.0084	0.0115	0.0148	0.0119	<b>0.0152</b>
	NDCG@10 ↑	0.0073	0.0118	0.0144	0.0158	0.0121	0.0160	0.0175	0.0158	<b>0.0203</b>
	HIT@5 ↑	0.0090	0.0167	0.0200	0.0202	0.0146	0.0186	0.0234	0.0188	<b>0.0246</b>
	HIT@10 ↑	0.0138	0.0240	0.0271	0.0330	0.0261	0.0325	0.0317	0.0311	<b>0.0405</b>
	MGU@5 ↓	/	/	/	0.1870	0.2119	0.1740	0.2255	0.1964	<b>0.1288</b>
	DGU@5 ↓	/	/	/	0.6122	0.6655	0.5722	0.7197	0.6257	<b>0.4031</b>
Clothing	NDCG@5 ↑	0.0016	0.0043	0.0024	0.0038	0.0038	0.0049	0.0045	0.0042	<b>0.0054</b>
	NDCG@10 ↑	0.0024	0.0053	0.0033	0.0063	0.0066	0.0075	0.0072	0.0068	<b>0.0096</b>
	HIT@5 ↑	0.0032	0.0080	0.0042	0.0078	0.0074	0.0100	0.0084	0.0086	<b>0.0118</b>
	HIT@10 ↑	0.0058	0.0110	0.0070	0.0156	0.0160	0.0180	0.0168	0.0168	<b>0.0246</b>
	MGU@5 ↓	/	/	/	0.1553	0.2491	0.1211	0.1728	0.1410	<b>0.0522</b>
	DGU@5 ↓	/	/	/	0.4908	0.7988	0.4069	0.5443	0.4686	<b>0.1996</b>
Book	NDCG@5 ↑	0.0054	0.0060	0.0065	0.0067	0.0051	0.0073	0.0033	0.0080	<b>0.0139</b>
	NDCG@10 ↑	0.0071	0.0084	0.0081	0.0103	0.0070	0.0112	0.0071	0.0105	<b>0.0145</b>
	HIT@5 ↑	0.0089	0.0110	0.0100	0.0116	0.0100	0.0130	0.0057	0.0137	<b>0.0226</b>
	HIT@10 ↑	0.0141	0.0185	0.0153	0.0228	0.0160	0.0240	0.0178	0.0219	<b>0.0244</b>
	MGU@5 ↓	/	/	/	0.1433	0.1382	0.1229	0.1709	0.1657	<b>0.0579</b>
	DGU@5 ↓	/	/	/	0.4742	0.3406	0.4238	0.5477	0.4643	<b>0.2240</b>

**Table 3: Performance comparison under different LLM-based RS. The best result is bolded.**

Method	Toy		Clothing	
	NDCG@5	DGU@5	NDCG@5	DGU@5
MSL	0.0198	0.7162	0.0077	0.5815
MSL+GDRT	<b>0.0253</b>	<b>0.2697</b>	<b>0.0096</b>	<b>0.1469</b>
LLaRA	0.0131	0.6485	0.0041	0.5412
LLaRA+GDRT	<b>0.0168</b>	<b>0.4249</b>	<b>0.0053</b>	<b>0.2320</b>
A-LLM	0.0129	0.6122	0.0045	0.5238
A-LLM+GDRT	<b>0.0160</b>	<b>0.4192</b>	<b>0.0053</b>	<b>0.1817</b>

**4.1.3 Implementation Details.** For all LLM-based methods, we adopt LLaMA3-3B [14] as the backbone, with the number of training epochs set to 5. For inference, we employ Constrained Beam Search (CBS), following prior work [3, 52], to ensure that the recommended items are drawn from the item set. The beam size is fixed at 10. For evaluation, we evaluate the NDCG@5 on the validation set for each epoch’s checkpoint and select the checkpoint with the highest score for testing. The corresponding results on the test set are reported as the final performance. For the hyperparameters in GDRT, the number of groups  $G$  is tuned from  $\{2, 5, 10\}$ , and  $\tau$  is tuned from  $\{0.1, 0.2, 0.3, 0.5, 1.0\}$ . To ensure fair comparisons, we utilize the source code provided by the original authors and tune the hyperparameters of all baseline methods according to the guidelines specified in their respective publications.

**4.1.4 Metrics.** Four widely used evaluation metrics are employed in this study:  $NDCG@K$  and  $Hit\ Ratio@K$  are used to evaluate recommendation accuracy, while  $MGU@K$  and  $DGU@K$  are adopted to evaluate fairness [24] ( $K=5, 10$ ). Specifically, for the fairness metrics,

items are divided into five groups according to Equation 2. We then calculate the discrepancy between each group’s proportion in the Top-K recommendations and its proportion in the user’s interaction history.  $MGU@K$  measures the average of these discrepancies, whereas  $DGU@K$  quantifies the gap between the maximum and minimum discrepancies across groups. Accordingly, smaller values of  $MGU@K$  and  $DGU@K$  indicate that the distribution of recommended items is more aligned with that of the user’s historical interactions, reflecting better fairness.

## 4.2 Performance Comparison (RQ1 & RQ2)

Table 2 presents a comparative analysis of the proposed GDRT method against the baselines. Overall, GDRT demonstrates substantial performance improvements across all datasets. By mitigating language bias, reducing excessive reliance on auxiliary tokens, and enhancing the model’s ability to capture user-specific behavioral patterns, GDRT delivers marked gains in both accuracy and fairness. Specifically, GDRT achieves an average improvement of 24.29% in NDCG@10 compared with SOTA methods. In terms of fairness, GDRT significantly alleviates the model’s tendency to recommend items strongly correlated with auxiliary information, leading to an average reduction of 37.43% in MGU@5 relative to the SOTA baselines. In contrast, other comparison methods yield only marginal improvements, and their lack of explicit mechanisms to address language bias results in limited fairness enhancement.

Furthermore, we evaluate the performance of GDRT when integrated into other advanced LLM-based RS. As shown in Table 3, GDRT consistently achieves significant performance gains and improves fairness across all methods, indicating that it can be seamlessly incorporated into diverse LLM-based recommenders while exhibiting strong generalization capability.




Figure 7: Ratio of attribution values between auxiliary tokens and user interaction tokens using SFT and GDRT. Left: task description vs. user interaction tokens. Right: prefix tokens of predicted item (take the first token) vs. user interaction tokens.




Figure 8: Distribution of Top-1 recommended items from GDRT-trained LLMs across five groups defined by auxiliary-token relevance (Group 1: highest relevance, Group 5: lowest relevance), with test-set distribution for comparison.

### 4.3 In-depth Analysis (RQ3)

In this section, we present an empirical analysis of the effectiveness of GDRT in mitigating language bias. As shown in Figure 7, GDRT considerably reduces the ratio of attribution values assigned to auxiliary tokens relative to interaction tokens, compared with SFT, indicating that GDRT effectively alleviates the problem of over-reliance on auxiliary tokens. Furthermore, GDRT markedly mitigates the unfairness induced by language bias. As shown in Figure 8, the recommendation distributions generated by GDRT more closely align with the item distribution in the test set. In addition, Figure 9 reveals that GDRT alleviates the issue observed in SFT, where the training process continuously amplifies the recommendation of items highly correlated with auxiliary tokens.

### 4.4 Hyper-Parameter Sensitivities (RQ4)

In this section, we perform a sensitivity analysis on the hyperparameters of GDRT: the number of groups  $G$  in K-means (Figure 10) and the coefficient  $\tau$  in DRO (Figure 11). For both parameters, model performance exhibits a trend of initially increasing and then decreasing as the parameter values rise. This phenomenon can be explained as follows. For  $G$ , an excessively small value may group highly heterogeneous data together, thereby diminishing the differences in auxiliary-token relevance between groups. Conversely, an overly large  $G$  results in groups containing too few samples to reliably represent the underlying distribution. As for  $\tau$ , a smaller value places greater emphasis on groups with higher losses; however,




Figure 9: Proportion of Top-1 recommendations belonging to Item Group 1 (most relevant with auxiliary tokens) during training with SFT and GDRT.




Figure 10: Hyperparameter sensitivity analysis on group number  $G$  (dashed: SFT baseline).




Figure 11: Hyperparameter sensitivity analysis on parameter  $\tau$  in DRO (dashed: SFT baseline).

overemphasizing the worst-performing group can cause overfitting and impair overall generalization. In contrast, a larger  $\tau$  balances optimization across all groups, but may reduce group-wise consistency in performance. Besides, the model demonstrates strong performance across a wide range of parameter settings, indicating the robustness of GDRT to hyperparameter selection.

### 4.5 Efficiency Comparison (RQ5)

This section compares the efficiency and performance of GDRT with baseline methods. As illustrated in Figure 12, GDRT demonstrates both optimal performance and high computational efficiency. In contrast, the DPO-based method SPRec requires an additional reference model and incorporates extra negative samples during training, leading to a considerable increase in runtime. CFT incurs higher computational costs by introducing additional counterfactual samples during training. Although Reweight and D3 incur negligible training overhead, their performance improvements remain limited.




Figure 12: Performance comparisons in terms of both recommendation accuracy and efficiency.

## 5 Related Work

### 5.1 Sequential Recommendation

Sequential recommendation focuses on predicting the next item a user will be interested in based on their historical interactions. Compared to collaborative filtering, sequential recommendation incorporates temporal information and places greater emphasis on capturing the evolving patterns of user interests. With the advancement of deep learning, numerous architectures based on deep neural networks have been introduced into the sequential recommendation. For example, GRU4Rec [21] employs RNNs, while Caser [49] utilizes CNNs to effectively capture long-term dependencies and modeling user interest patterns from historical behavior. More advanced models such as SASRec [25] and BERT4Rec [47] leverage self-attention mechanisms [51], enabling the identification of the most relevant parts within the sequence. Due to the dynamic nature of data distributions as time evolves, DROS [60] introduces DRO [45, 58] to further enhance the model’s robustness against distributional shifts caused by temporal changes. The readers may refer to the survey [15, 55] for more details.

### 5.2 Biases in LLM-based Recommendation

Large Language Models (LLMs), with their powerful capabilities in comprehension, reasoning, and extensive knowledge [1, 14, 18], have been widely applied to recommendation systems (RS) [59]. One prominent paradigm is LLM-based RS [31]. In this paradigm, LLMs are directly employed as the recommendation backbone, with prompts used to instruct the models to perform recommendation tasks. However, due to the discrepancy between the pre-training objectives of LLMs and the goals of recommendation systems, subsequent studies have explored fine-tuning LLMs on domain-specific recommendation datasets to enhance their task-specific capabilities [2, 4, 9, 27, 31, 35, 52, 53, 69].

Recent studies have extensively explored bias and fairness issues in LLM-based RS, such as popularity bias [16, 17, 24, 34, 37], position bias [6, 8, 12, 22, 23, 38, 39, 61], length bias [3], and bias stemming from LLMs’ preferences for specific item attributes [13, 24, 33, 62]. Popularity bias refers to the model’s tendency to disproportionately recommend popular items. Mitigation strategies typically reweight samples at the item [24] or token level [17, 37] to prevent over-optimization toward popular items, or increasing the negative sampling probability for such items based on DPO [16, 34]. Position bias

reflects the model’s sensitivity to the ordering of candidate items and is commonly addressed by incorporating position-invariant priors during training [38, 61] or inference [6, 22, 39] stage. Length bias refers to the tendency of LLMs to favor items with longer token sequences during inference. D3 [3] alleviates this issue by removing length normalization from the score computation process. Additionally, LLMs have been shown to exhibit unequal preferences for items with specific attributes (e.g., genre, creator identity, or recency) [13]. To address this issue, some methods employ prompt-engineering techniques to encourage fairer recommendation behavior [33], whereas others reweight items across attribute categories to alleviate attribute-level bias [24, 62].

Nevertheless, existing research has largely overlooked language bias, which arises during fine-tuning and reflects the inherent over-reliance of LLMs on auxiliary tokens. Since existing debiasing methods fail to account for this factor, their effectiveness remains limited. CFT [63] seeks to enhance the modeling of users’ historical interactions. However, it suffers from notable limitations, including objective misalignment, difficulties in weight selection, and high computational overhead, which constrain its effectiveness. These limitations are examined in detail in Section 2.2.4.

### 5.3 Group Distributionally Robust Optimization

Group Distributionally Robust Optimization (Group DRO) [46] is an optimization framework that operates over predefined sample groups, aiming to achieve consistent and reliable performance across them by emphasizing the optimizing of the worst-performing group during training. It has been widely applied in various domains [42, 46, 65] and has demonstrated strong effectiveness in mitigating group disparities [20, 65] and reducing models’ reliance on shortcut correlations [10, 19]. Several studies have applied Group DRO to RS. For example, S-DRO [56] is employed to improve user experience for less-represented user groups that prefer interacting with unpopular items. PDRO [65] boosts the representation learning of both sparse users and items without sacrificing the performance for popular ones.

## 6 Conclusion

In this work, we identify a key limitation of supervised fine-tuning (SFT) in LLM-based recommenders: it often induces **Language Bias**, whereby the model over-relies on auxiliary tokens (e.g., task descriptions and prefix-generated tokens) while underutilizing core user interaction information. This bias undermines recommendation accuracy and raises unfairness concerns. To address this issue, we introduce Group Distributionally Robust Optimization-based Tuning (GDRT), which aims to reduce the model’s over-reliance on auxiliary tokens by applying Group DRO across token groups with varying degrees of relevance to auxiliary tokens. Extensive experiments on multiple public datasets demonstrate that GDRT effectively mitigates language bias, thereby significantly improving recommendation accuracy and enhancing fairness.

This work investigates a novel form of bias introduced by the integration of LLMs into recommenders, which is not present in traditional recommendation models. A promising avenue for future work is to examine whether LLM-based RS exhibit additional, as-yet unidentified biases.

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Keqin Bao, Jizhi Zhang, Wenjie Wang, Yang Zhang, Zhengyi Yang, Yanchen Luo, Chong Chen, Fuli Feng, and Qi Tian. 2025. A bi-step grounding paradigm for large language models in recommendation systems. *ACM Transactions on Recommender Systems* 3, 4 (2025), 1–27.
- [3] Keqin Bao, Jizhi Zhang, Yang Zhang, Xinyue Huo, Chong Chen, and Fuli Feng. 2024. Decoding matters: Addressing amplification bias and homogeneity issue for llm-based recommendation. *arXiv preprint arXiv:2406.14900* (2024).
- [4] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1007–1014.
- [5] Oren Barkan, Yonatan Toib, Yehonatan Elisha, Jonathan Weill, and Noam Koenigstein. 2024. LLM Explainability via Attributive Masking Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. 9522–9537.
- [6] Ethan Bito, Yongli Ren, and Estrid He. 2025. Evaluating Position Bias in Large Language Model Recommendations. *arXiv preprint arXiv:2508.02020* (2025).
- [7] Yuwei Cao, Nikhil Mehta, Xinyang Yi, Raghunandan Keshavan, Lukasz Heldt, Lichan Hong, Ed H Chi, and Maheswaran Sathiamoorthy. 2024. Aligning Large Language Models with Recommendation Knowledge. *arXiv preprint arXiv:2404.00245* (2024).
- [8] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems* 41, 3 (2023), 1–39.
- [9] Yuxin Chen, Junfei Tan, An Zhang, Zhengyi Yang, Leheng Sheng, Enzhi Zhang, Xiang Wang, and Tat-Seng Chua. 2024. On Softmax Direct Preference Optimization for Recommendation. *arXiv preprint arXiv:2406.09215* (2024).
- [10] Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. 2021. Environment inference for invariant learning. In *International Conference on Machine Learning*. PMLR, 2189–2200.
- [11] Yu Cui, Feng Liu, Pengbo Wang, Bohao Wang, Heng Tang, Yi Wan, Jun Wang, and Jiawei Chen. 2024. Distillation Matters: Empowering Sequential Recommenders to Match the Performance of Large Language Models. In *Proceedings of the 18th ACM Conference on Recommender Systems*. 507–517.
- [12] Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Bias and unfairness in information retrieval systems: New challenges in the llm era. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 6437–6447.
- [13] Yashar Deldjoo. 2024. Understanding biases in ChatGPT-based recommender systems: Provider fairness, temporal stability, and recency. *ACM Transactions on Recommender Systems* (2024).
- [14] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783* (2024).
- [15] Hui Fang, Danning Zhang, Yiheng Shu, and Guibing Guo. 2020. Deep learning for sequential recommendation: Algorithms, influential factors, and evaluations. *ACM Transactions on Information Systems (TOIS)* 39, 1 (2020), 1–42.
- [16] Chongming Gao, Ruijun Chen, Shuai Yuan, Kexin Huang, Yuanqing Yu, and Xiangnan He. 2025. Sprec: Self-play to debias llm-based recommendation. In *Proceedings of the ACM on Web Conference 2025*. 5075–5084.
- [17] Chongming Gao, Mengyao Gao, Chenxiao Fan, Shuai Yuan, Wentao Shi, and Xiangnan He. 2025. Process-supervised llm recommenders via flow-guided tuning. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1934–1943.
- [18] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature* 645, 8081 (2025), 633–638.
- [19] Hyeonjeon Han, Sehwan Kim, Hyungjun Joo, Sangwoo Hong, and Jungwoon Lee. 2024. Mitigating spurious correlations via disagreement probability. *Advances in Neural Information Processing Systems* 37 (2024), 74363–74382.
- [20] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. 2018. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*. PMLR, 1929–1938.
- [21] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [22] Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European Conference on Information Retrieval*. Springer, 364–381.
- [23] Chumeng Jiang, Jiayin Wang, Weizhi Ma, Charles LA Clarke, Shuai Wang, Chuhan Wu, and Min Zhang. 2025. Beyond Utility: Evaluating LLM as Recommender. In *Proceedings of the ACM on Web Conference 2025*. 3850–3862.
- [24] Meng Jiang, Keqin Bao, Jizhi Zhang, Wenjie Wang, Zhengyi Yang, Fuli Feng, and Xiangnan He. 2024. Item-side fairness of large language model-based recommendation system. In *Proceedings of the ACM Web Conference 2024*. 4717–4726.
- [25] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [26] Carina Kauf, Emmanuele Chersoni, Alessandro Lenci, Evelina Fedorenko, and Anna Ivanova. 2024. Log Probabilities Are a Reliable Estimate of Semantic Plausibility in Base and Instruction-Tuned Language Models. In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*. 263–277.
- [27] Sein Kim, Hongseok Kang, Seungyeon Choi, Donghyun Kim, Minchul Yang, and Chanyoung Park. 2024. Large language models meet collaborative filtering: An efficient all-round llm-based recommender system. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1395–1406.
- [28] Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, et al. 2020. Captum: A unified and generic model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896* (2020).
- [29] Dong-Ho Lee, Adam Kraft, Long Jin, Nikhil Mehta, Taibai Xu, Lichan Hong, Ed H Chi, and Xinyang Yi. 2024. STAR: A Simple Training-free Approach for Recommendations using Large Language Models. *arXiv preprint arXiv:2410.16458* (2024).
- [30] Simon Lepage, Jeremie Mary, and David Picard. 2025. Closing the Performance Gap in Generative Recommenders with Collaborative Tokenization and Efficient Modeling. *arXiv preprint arXiv:2508.14910* (2025).
- [31] Lei Li, Yongfeng Zhang, Dugang Liu, and Li Chen. 2023. Large language models for generative recommendation: A survey and visionary discussions. *arXiv preprint arXiv:2309.01157* (2023).
- [32] Xinhang Li, Chong Chen, Xiangyu Zhao, Yong Zhang, and Chunxiao Xing. 2023. E4srec: An elegant effective efficient extensible solution of large language models for sequential recommendation. *arXiv preprint arXiv:2312.02443* (2023).
- [33] Xinyi Li, Yongfeng Zhang, and Edward C Malthouse. 2023. A preliminary study of chatgpt on news recommendation: Personalization, provider fairness, fake news. *arXiv preprint arXiv:2306.10702* (2023).
- [34] Jiayi Liao, Xiangnan He, Ruobing Xie, Jiancan Wu, Yancheng Yuan, Xingwu Sun, Zhanhui Kang, and Xiang Wang. 2024. RosePO: Aligning LLM-based Recommenders with Human Values. *arXiv preprint arXiv:2410.12519* (2024).
- [35] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2024. Llara: Large language-recommendation assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1785–1795.
- [36] Jianghao Lin, Rong Shan, Chenxu Zhu, Kounianhua Du, Bo Chen, Shigang Quan, Ruiming Tang, Yong Yu, and Weinan Zhang. 2024. Rella: Retrieval-enhanced large language models for lifelong sequential behavior comprehension in recommendation. In *Proceedings of the ACM on Web Conference 2024*. 3497–3508.
- [37] Sijin Lu, Zhibo Man, Fangyuan Luo, and Jun Wu. 2025. Dual Debiasing in LLM-based Recommendation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2685–2689.
- [38] Sichun Luo, Bowei He, Haohan Zhao, Wei Shao, Yanlin Qi, Yinya Huang, Aojun Zhou, Yuxuan Yao, Zongpeng Li, Yuanzhang Xiao, et al. 2025. Recranker: Instruction tuning large language model as ranker for top-k recommendation. *ACM Transactions on Information Systems* 43, 5 (2025), 1–31.
- [39] Tianhui Ma, Yuan Cheng, Hengshu Zhu, and Hui Xiong. 2023. Large language models are not stable recommender systems. *arXiv preprint arXiv:2312.15746* (2023).
- [40] Vivek Miglani, Aobo Yang, Aram H Markosyan, Diego Garcia-Olano, and Narine Kokhlikyan. 2023. Using captum to explain generative language models. *arXiv preprint arXiv:2312.05491* (2023).
- [41] Hyunsoo Na, Minseok Gang, Youngrok Ko, Jinseok Seol, and Sang-goo Lee. 2024. Enhancing Large Language Model Based Sequential Recommender Systems with Pseudo Labels Reconstruction. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. 7213–7222.
- [42] Yonatan Oren, Shiori Sagawa, Tatsunori B Hashimoto, and Percy Liang. 2019. Distributionally Robust Language Modeling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 4227–4237.
- [43] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [44] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems* 36 (2024).

- [45] Hamed Rahimian and Sanjay Mehrotra. 2019. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659* (2019).
- [46] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. 2019. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731* (2019).
- [47] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.
- [48] Juntao Tan, Shuyuan Xu, Wenyue Hua, Yingqiang Ge, Zelong Li, and Yongfeng Zhang. 2024. Idgenrec: Llm-recsys alignment with textual id learning. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*. 355–364.
- [49] Jiayi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the eleventh ACM international conference on web search and data mining*. 565–573.
- [50] Jiabin Tang, Yuhao Yang, Wei Wei, Lei Shi, Lixin Su, Suqi Cheng, Dawei Yin, and Chao Huang. 2024. Graphgpt: Graph instruction tuning for large language models. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 491–500.
- [51] Ashish Vaswani. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017).
- [52] Bohao Wang, Feng Liu, Jiawei Chen, Xingyu Lou, Changwang Zhang, Jun Wang, Yuegang Sun, Yan Feng, Chun Chen, and Can Wang. 2025. Msl: Not all tokens are what you need for tuning llm as a recommender. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1912–1922.
- [53] Hangyu Wang, Jianghao Lin, Xiangyang Li, Bo Chen, Chenxu Zhu, Ruiming Tang, Weinan Zhang, and Yong Yu. 2024. Flip: Fine-grained alignment between id-based models and pretrained language models for ctr prediction. In *Proceedings of the 18th ACM conference on recommender systems*. 94–104.
- [54] Lei Wang and Ee-Peng Lim. 2023. Zero-shot next-item recommendation using large pretrained language models. *arXiv preprint arXiv:2304.03153* (2023).
- [55] Shoujin Wang, Liang Hu, Yan Wang, Longbing Cao, Quan Z Sheng, and Mehmet Orgun. 2019. Sequential recommender systems: challenges, progress and prospects. *arXiv preprint arXiv:2001.04830* (2019).
- [56] Hongyi Wen, Xinyang Yi, Tiansheng Yao, Jiayi Tang, Lichan Hong, and Ed H Chi. 2022. Distributionally-robust recommendations for improving worst-case user experience. In *Proceedings of the ACM Web Conference 2022*. 3606–3610.
- [57] Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. 2023. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671* (2023).
- [58] Junkang Wu, Jiawei Chen, Jiancan Wu, Wentao Shi, Xiang Wang, and Xiangnan He. 2023. Understanding contrastive learning via distributionally robust optimization. *Advances in Neural Information Processing Systems* 36 (2023), 23297–23320.
- [59] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. 2024. A survey on large language models for recommendation. *World Wide Web* 27, 5 (2024), 60.
- [60] Zhengyi Yang, Xiangnan He, Jizhi Zhang, Jiancan Wu, Xin Xin, Jiawei Chen, and Xiang Wang. 2023. A generic learning framework for sequential recommendation with distribution shifts. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 331–340.
- [61] Junjie Zhang, Yupeng Hou, Ruobing Xie, Wenqi Sun, Julian McAuley, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2024. Agentcf: Collaborative learning with autonomous language agents for recommender systems. In *Proceedings of the ACM Web Conference 2024*. 3679–3689.
- [62] Jiaming Zhang, Yuyuan Li, Yiqun Xu, Li Zhang, Xiaohua Feng, Zhifei Ren, and Chaochao Chen. 2025. BiFair: A Fairness-aware Training Framework for LLM-enhanced Recommender Systems via Bi-level Optimization. *arXiv preprint arXiv:2507.04294* (2025).
- [63] Yang Zhang, Juntao You, Yimeng Bai, Jizhi Zhang, Keqin Bao, Wenjie Wang, and Tat-Seng Chua. 2024. Causality-enhanced behavior sequence modeling in LLMs for personalized recommendation. *arXiv preprint arXiv:2410.22809* (2024).
- [64] Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology* 15, 2 (2024), 1–38.
- [65] Jujia Zhao, Wenjie Wang, Xinyu Lin, Leigang Qu, Jizhi Zhang, and Tat-Seng Chua. 2023. Popularity-aware distributionally robust optimization for recommendation system. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4967–4973.
- [66] Bowen Zheng, Yupeng Hou, Hongyu Lu, Yu Chen, Wayne Xin Zhao, Ming Chen, and Ji-Rong Wen. 2024. Adapting large language models by integrating collaborative semantics for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 1435–1448.
- [67] Zhi Zheng, Wenshuo Chao, Zhaopeng Qiu, Hengshu Zhu, and Hui Xiong. 2024. Harnessing large language models for text-rich sequential recommendation. In *Proceedings of the ACM Web Conference 2024*. 3207–3216.
- [68] Wei Zhou, Heike Adel, Hendrik Schuff, and Ngoc Thang Vu. 2024. Explaining pre-trained language models with attribution scores: An analysis in low-resource settings. *arXiv preprint arXiv:2403.05338* (2024).
- [69] Yaochen Zhu, Liang Wu, Qi Guo, Liangjie Hong, and Jundong Li. 2024. Collaborative large language model for recommender systems. In *Proceedings of the ACM on Web Conference 2024*. 3162–3172.

## A Appendices

### A.1 The proof of Lemma 1

The original formulation of GDRT

$$\mathcal{L}_{GDRT} = \max_Q \sum_{g=1}^G Q(g) \mathcal{L}(g), \quad \text{s.t. } D_{KL}(Q, U) \leq \eta, \quad (6)$$

can be rewritten in the expectation form as

$$\begin{aligned} & \max_Q \mathbb{E}_{g \sim Q} [\mathcal{L}(g)] \\ & \text{s.t. } \mathbb{E}_{g \sim Q} \left[ \log \frac{Q(g)}{U(g)} \right] \leq \eta \end{aligned} \quad (7)$$

In the following, we focus on how to eliminate the inner maximization optimization problem and the KL constraint term. Assume  $W(g) = Q(g)/U(g)$  and define a convex function  $\phi(x) = x \log x - x + 1$ . Then the divergence  $D_{KL}(Q, U)$  can be written as  $\mathbb{E}_U[\phi(W)]$ . The inner layer maximization optimization problem can be reformulated as follow:

$$\begin{aligned} & \max_W \mathbb{E}_U [\mathcal{L}W] \\ & \text{s.t. } \mathbb{E}_U [\phi(W)] \leq \eta, \mathbb{E}_U [W] = 1 \end{aligned} \quad (8)$$

As a convex optimization problem, we use the Lagrangian function to solve it:

$$\begin{aligned} & \min_{\tau \geq 0, \beta} \max_W \mathbb{E}_U [\mathcal{L}W] - \tau (\mathbb{E}_U [\phi(W)] - \eta) + \beta (\mathbb{E}_U [W] - 1) \\ & = \min_{\tau \geq 0, \beta} \left\{ \tau \eta - \beta + \tau \max_W \mathbb{E}_U \left[ \frac{\mathcal{L} + \beta}{\tau} W - \phi(W) \right] \right\} \\ & = \min_{\tau \geq 0, \beta} \left\{ \tau \eta - \beta + \tau \mathbb{E}_U \left[ \max_W \left( \frac{\mathcal{L} + \beta}{\tau} W - \phi(W) \right) \right] \right\} \end{aligned} \quad (9)$$

Notice that  $\max_W \left( \frac{\mathcal{L} + \beta}{\tau} W - \phi(W) \right) = \phi^* \left( \frac{\mathcal{L} + \beta}{\tau} \right)$  is the convex conjugate function of  $\phi(x)$  and we have  $\phi^*(x) = e^x - 1$ .  $W(g) = e^{\frac{\mathcal{L}(g) + \beta}{\tau}}$  when the maximum value is obtained.

$$\begin{aligned} & \min_{\tau \geq 0, \beta} \left\{ \tau \eta - \beta + \tau \mathbb{E}_U \left[ \max_W \left( \frac{\mathcal{L} + \beta}{\tau} W - \phi(W) \right) \right] \right\} \\ & = \min_{\tau \geq 0, \beta} \left\{ \tau \eta - \beta + \tau \mathbb{E}_U \left[ e^{\frac{\mathcal{L} + \beta}{\tau}} - 1 \right] \right\} \\ & = \min_{\tau \geq 0} \left\{ \tau \eta + \tau \log \mathbb{E}_U \left[ e^{\frac{\mathcal{L}}{\tau}} \right] \right\} \end{aligned} \quad (10)$$

where  $\beta = -\tau \log \mathbb{E}_{g \sim U} \left[ e^{\frac{\mathcal{L}(g)}{\tau}} \right]$  and  $W(g) = \frac{e^{\frac{\mathcal{L}(g)}{\tau}}}{\mathbb{E}_{g' \sim U} \left[ e^{\frac{\mathcal{L}(g')}{\tau}} \right]}$  when the minimum value is obtained. We consider the Lagrange multiplier  $\tau$  as a hyperparameter related to the robustness radius  $\eta$ . Then we can get the unconstrained optimization problem as follows,

$$\mathcal{L}_{GDRT} = \tau \eta + \tau \log \mathbb{E}_{g \sim U} \exp \left( \frac{\mathcal{L}(g)}{\tau} \right) \quad (11)$$

where the worst-case distribution

$$Q^*(g) = U(g) \frac{\exp(\mathcal{L}(g)/\tau)}{\mathbb{E}_{i \sim U} [\exp(\mathcal{L}(i)/\tau)]} \quad (12)$$

Since  $U$  is a uniform distribution

$$Q^*(g) = \frac{\exp(\mathcal{L}(g)/\tau)}{\sum_{g'} [\exp(\mathcal{L}(g')/\tau)]} \quad (13)$$

Thus lemma 1 is proven.

1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187  
1188  
1189  
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  
1244  
1245  
1246  
1247  
1248  
1249  
1250  
1251  
1252  
1253  
1254  
1255  
1256  
1257  
1258  
1259  
1260  
1261  
1262  
1263  
1264  
1265  
1266  
1267  
1268  
1269  
1270  
1271  
1272  
1273  
1274  
1275  
1276