

# SPRec: Self-Play to Debias LLM-based Recommendation

Chongming Gao\*  
chongminggao@ustc.edu.cn  
University of Science and Technology  
of China  
Hefei, China

Ruijun Chen\*  
rjchen20@mail.ustc.edu.cn  
University of Science and Technology  
of China  
Hefei, China

Shuai Yuan  
syuanaf@connect.ust.hk  
Hong Kong University of Science and  
Technology  
Hong Kong, China

Kexin Huang  
huangkx@mail.ustc.edu.cn  
University of Science and Technology  
of China  
Hefei, China

Yuanqing Yu  
yyq23@mails.tsinghua.edu.cn  
Tsinghua University  
Beijing, China

Xiangnan He<sup>†</sup>  
xiangnanhe@gmail.com  
MoE Key Lab of BIPC, University of  
Science and Technology of China  
Hefei, China

## Abstract

Large language models (LLMs) have attracted significant attention in recommendation systems. Current work primarily applies supervised fine-tuning (SFT) to adapt the model for recommendation tasks. However, SFT on positive examples only limits the model’s ability to align with user preference. To address this, researchers recently introduced Direct Preference Optimization (DPO), which explicitly aligns LLMs with user preferences using offline preference ranking data. However, we found that DPO inherently biases the model towards a few items, exacerbating the filter bubble issue and ultimately degrading user experience.

In this paper, we propose SPRec, a novel self-play framework designed to mitigate over-recommendation and improve fairness without requiring additional data or manual intervention. In each self-play iteration, the model undergoes an SFT step followed by a DPO step, treating offline interaction data as positive samples and the predicted outputs from the previous iteration as negative samples. This effectively re-weights the DPO loss function using the model’s logits, adaptively suppressing biased items. Extensive experiments on multiple real-world datasets demonstrate SPRec’s effectiveness in enhancing recommendation accuracy and fairness. The code is available via <https://github.com/RegionCh/SPRec>.

## CCS Concepts

• Information systems → Recommender systems.

## Keywords

Large Language Model-based Recommendation, Homogeneity Issue, Fairness, Self-play, Direct Preference Optimization

\*Both authors contributed equally to this research.

<sup>†</sup>Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

WWW '25, April 28-May 2, 2025, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1274-6/25/04


<https://doi.org/10.1145/3696410.3714524>

### • Token-level biases in LLM-based recommendations



☹️ Favoring movies titles with frequent tokens, e.g., “the”.

### • Item-level biases in fine-tuned LLMs



☹️ Focusing exclusively on the Batman film series.

**Figure 1: Homogeneity issues in LLM-based recommendation results caused by token-level and item-level biases.**

## ACM Reference Format:

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 Web Conference 2025 (WWW '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3696410.3714524>

## 1 Introduction

Recently, large language models (LLMs) have demonstrated significant success across numerous domains, showcasing advanced capabilities in learning, reasoning, and generalizing to downstream tasks [12, 38]. In the field of recommender systems, there has been growing interest in leveraging the potential of LLMs [41]. One prominent approach involves positioning LLMs as the central recommendation backbone, utilizing users’ past interactions and current needs to generate personalized recommendations [3, 40]. Compared to traditional methods, LLM-based recommendation systems (LRSs) offer distinct advantages, including a deeper contextual understanding and the flexibility to adapt to users’ evolving preferences.

To enable LLMs to learn collaborative filtering signals and effectively perform item recommendations, a prevalent strategy is to fine-tune pre-trained LLMs via Supervised Fine-Tuning (SFT) [3]. This approach allows LLMs to efficiently internalize user preferences from offline data by adjusting their parameters to align with

the recommendation task. Building on SFT, recent research has adopted Direct Preference Optimization (DPO) [31] to further refine user preferences [2, 10, 25]. While SFT relies solely on desirable answers, DPO incorporates both chosen and rejected response pairs, allowing the LLM to learn user ranking preferences and gain a more nuanced understanding of fine-grained, personalized information. This approach mirrors the common practice in recommendation models, which utilize both positive and negative samples for effective training [7, 32].

Despite these advancements, we find that employing DPO to align user preferences in recommender systems inherently introduces significant biases due to its underlying mechanisms. These biases can lead to serious homogeneity issues, where LLMs recommend items with similar names or content. Fig. 1 illustrates how token-level and item-level biases manifest in the Top-K movie recommendation. Token-level biases arise due to the fact that LLMs generate item names in a tokenized fashion. Since LLMs are usually tuned to maximize the likelihood of target tokens, items with more common tokens (e.g., movies with the word “the” in their titles) may be overrepresented, regardless of user relevance. At the item level, biases can emerge from multiple factors, particularly after fine-tuning, where LLMs may disproportionately recommend popular items, such as the Batman film series. This can lead to filter bubbles [16, 17], where users are repeatedly exposed to a narrow range of popular content, limiting the diversity of recommendations and degrading the user experience.

Some research has been conducted on bias and unfairness issues in LRSs [12, 15, 45]. Dai et al. [12] provided a comprehensive overview of the various types of biases that emerge across different stages of these models and outlined strategies to mitigate them. For example, Jiang et al. [22] proposed re-weighting the fine-tuning loss for each item and re-ranking the generated results to ensure equitable treatment across genre groups. Similarly, Bao et al. [4] adjusted the LLM decoding process by removing the length normalization term and incorporating predictions from a text-free model, helping to reduce amplification bias and address homogeneity issues. However, these methods often rely on carefully crafted rules or external knowledge, limiting their broader applicability in general recommendation systems.

To this end, we propose a self-play recommendation tuning framework, SPRec, to adaptively suppress biases and improve fairness in LRSs without the need for additional data or expert knowledge. The core idea of SPRec is straightforward: each tuning iteration begins with an SFT round using positive samples from offline data, followed by a DPO. In the DPO step, the SFT data is treated as positive samples, while the predicted outputs from the previous iteration are treated as negative samples. The philosophy is to let the model “play” with its own output by re-weighting the DPO loss function based on its predictions. As a result, items that rank higher in the model’s predictions are penalized, while the SFT process reinforces the ranking of positive items. Over time, this self-play learning process adaptively suppresses undesirable items (biases) while maintaining alignment with positive samples. Extensive experiments on public datasets demonstrate that SPRec effectively improves both accuracy and fairness, showcasing its potential as a practical and efficient solution for LRSs.

The main contributions of this paper are as follows:

- We analyze how current LRSs tuned through DPO inevitably exhibit biases due to their underlying learning mechanisms, leading to the homogeneity issue.
- We propose SPRec, a self-play recommendation tuning framework that addresses these biases and improves fairness without the need for external knowledge.
- Experiments validate that SPRec improves accuracy, diversity, and fairness, with ablation studies indicating that the self-play negative samples contribute significantly to the improvements.

## 2 Related work

We provide a brief overview of LLM-based recommender systems and their associated bias issues, followed by an introduction to the self-play mechanism employed in our method.

### 2.1 LLMs for Recommendation

LLMs have shown exceptional generative, generalization, and reasoning capabilities in NLP, driving research into their applications for personalized recommendations. Their integration into recommendation tasks follows three main paradigms: (1) acting as decision makers [4, 27], (2) assisting by providing contextual information [19, 26], and (3) serving as user simulators [5, 46]. Early studies explored prompt engineering to leverage LLMs for recommendation tasks [18, 20].

Later, fine-tuning methods emerged, demonstrating that adapting LLM parameters on recommendation data significantly enhances performance. These approaches primarily rely on SFT [4, 9]. To further align LLMs with user preferences, DPO has been employed for post-training [2, 10, 25]. However, prior work has overlooked DPO’s inherent tendency to introduce severe biases, favoring only frequently exposed items and degrading user experience. In this work, we are the first to identify this issue and propose a mitigation strategy.

### 2.2 Biases in Recommender Systems

Bias and fairness issues are pervasive in recommender systems and have been extensively studied. Chen et al. [8] provide a comprehensive survey on biases such as popularity bias, selection bias, and position bias. These biases can significantly impact user satisfaction, promoting clickbait content or reinforcing filter bubbles that reduce engagement [17]. Additionally, algorithmic decisions may favor certain items, raising fairness concerns [24, 39], disproportionately affecting user groups and discouraging content creators [21, 44].

These challenges persist in LLM-based recommender systems [36, 41, 43]. Research shows that LLMs can inherit social biases, leading to unfair recommendations related to sensitive attributes like gender and race [45]. Dai et al. [12] provide a unified distribution mismatch perspective on bias and fairness in information retrieval. Existing bias mitigation methods in LRSs typically rely on predefined target distributions or external guidance for LLM alignment. In contrast, we introduce the first self-play framework for mitigating bias in LRSs, requiring neither prior knowledge nor additional models. By simply modifying the tuning process, our approach reduces long-tail effects and improves fairness.

## 2.3 Self-Play Mechanism

Machine learning models are often data-driven, relying heavily on the availability of offline data. However, offline data is inherently limited, raising an important question: can algorithms improve themselves iteratively without the need for additional data? This is precisely the challenge that the self-play mechanism aims to address. The concept of self-play originated from board games such as Go and chess, exemplified by groundbreaking systems like AlphaGo Zero [34] and AlphaZero [33].

In the era of LLMs, early preference alignment algorithms like RLHF and DPO operate as single optimization procedures. Building on this, Chen et al. [11] proposed that LLMs can refine their capabilities through self-play by interacting with instances of themselves. More specifically, the LLM generates its own training data from prior iterations and subsequently learns a new policy that outperforms the old one [42]. This iterative fine-tuning framework has been shown to be equivalent to finding the Nash equilibrium of a two-player game, gaining significant recognition due to its solid theoretical foundation and simplicity [6]. In recommender systems, we leverage the self-play mechanism to adaptively reduce biased items in a simple and effective manner.

## 3 Preliminary

In this section, we provide a brief overview of the technologies for aligning LLMs with the recommendation task. We then introduce the idea of evaluating the biases and unfairness in LRSs.

### 3.1 Supervised Fine-tuning (SFT)

To enable an open-source LLM to learn recommendation tasks effectively, a practical approach is fine-tuning all or part of its parameters using demonstration data from offline recommendation logs. The objective is to align the model’s behavior with the recommendation task by maximizing the log-likelihood over the training dataset  $\mathcal{D}$ :

$$\pi_{SFT} = \arg \max_{\pi_{\theta}} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}} \log \pi_{\theta}(y_i | x_i), \quad (1)$$

where  $(x_i, y_i)$  are input-output pairs from  $\mathcal{D}$ , with  $x_i$  representing user context and interaction history, and  $y_i$  the target item. Defining  $p_{\mathcal{D}}(y|x)$  as the empirical probability (i.e., item popularity), SFT aligns model predictions by minimizing the forward KL-divergence:

$$\begin{aligned} \pi_{SFT} &= \arg \min_{\pi_{\theta}} \mathbb{D}_{KL}(p_{\mathcal{D}}(y|x), \pi_{\theta}(y|x)) \\ &= \arg \min_{\pi_{\theta}} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}} -\log \pi_{\theta}(y_i | x_i) + H(p_{\mathcal{D}}), \end{aligned} \quad (2)$$

where  $H(p_{\mathcal{D}})$  is the constant entropy of  $p_{\mathcal{D}}$ .

### 3.2 Direct Preference Optimization (DPO)

To ensure that model outputs align with intricate user preferences, researchers have proposed Direct Preference Optimization (DPO) [31], which optimizes the following objective function:

$$\min_{\pi_{\theta}} -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \log \sigma \left[ \beta \log \left( \frac{\pi_{\theta}(y_w | x)}{\pi_{\text{ref}}(y_w | x)} \right) - \beta \log \left( \frac{\pi_{\theta}(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right) \right], \quad (3)$$

where  $(x, y_w, y_l)$  denotes a prompt  $x$  with a chosen (preferred) answer  $y_w$  and a rejected answer  $y_l$ . The parameter  $\beta$  acts as a



**Figure 2: Distribution of cold-start recommendation results. Group 0: least popular, group 4: most popular.**

regularization factor, controlling the extent to which the learned policy  $\pi_{\theta}$  deviates from the reference policy  $\pi_{\text{ref}}$ .

In the context of recommendation tasks,  $x$  represents the user context, typically comprising user features and historical interaction sequences, while  $y_w$  and  $y_l$  correspond to positive and negative samples, respectively. The goal is to encourage the model to assign higher probabilities to preferred items ( $y_w$ ) over less desirable ones ( $y_l$ ), effectively capturing user preferences.

DPO offers an efficient and stable solution for preference alignment, eliminating the need for complex reward models often required in reinforcement learning-based approaches. Its natural ability to incorporate both positive and negative samples makes it particularly well-suited for recommender systems, where learning from contrasting user interactions is crucial [2, 10, 25].

### 3.3 Evaluating Bias via Distribution Alignment

When aligning user preferences, LLMs may inadvertently learn biased or unfair outcomes. To assess the bias and fairness issues in LRSs, a mainstream perspective is to formulate the problems as a mismatch distribution problem [12]. Specifically, let  $R$  denote the ground-truth user preference (e.g., an item list), following the distribution  $P(R)$ , and let  $\hat{R}$  represent the model-predicted preferences, drawn from  $P(\hat{R})$ . Bias or unfairness is then quantified by the mismatch between these two distributions:  $P(R) \neq P(\hat{R})$ .

To apply this framework, we follow Jiang et al. [22], approximating  $P(R)$  using the category distribution from offline training data. We further employ their MGU metric to systematically measure the degree of mismatch in our experiments.


## 4 Problem of DPO: Amplify Popularity Bias

We present an empirical analysis to demonstrate how DPO exacerbates popularity bias in LRSs, followed by a theoretical examination of the underlying mechanisms driving this phenomenon.

### 4.1 Empirical Analysis

To systematically evaluate how DPO amplifies recommendation bias, we design a cold-start recommendation task using two widely-used benchmark datasets: MovieLens and Goodreads<sup>1</sup>. In this task, the LLM generates recommendations without access to user interaction history. We randomly sample 100 items and partition them into

<sup>1</sup>Dataset details are provided in Section 6.1.1



**Figure 3: Illustration of SFT, DPO, and SPRec in LLM-based recommendations. (a) SFT generates mass-covering results but retains inherent biases. (b) DPO amplifies these biases by over-representing certain items. (c) SPRec mitigates bias through self-play, leveraging model outputs as negative samples to achieve balanced recommendations.**

five groups based on interaction probabilities, ensuring balanced popularity levels. Positive samples are drawn accordingly, while negative samples for DPO training are randomly selected. The training and validation sets contain 4096 and 512 samples, respectively. After training, we analyze the distribution of recommendations across the five groups.

Fig. 2 shows the proportion of recommendations allocated to each group before and after training, where group 0 represents the least popular items and group 4 the most popular. The results reveal three key insights: (1) SFT introduces a slight bias toward group 4; (2) DPO significantly amplifies this bias, causing recommendations to concentrate almost entirely on the most popular items; and (3) our proposed method, SPRec (introduced in detail later), effectively mitigates the bias amplification induced by DPO.

## 4.2 Theoretical Analysis

In recommendation tasks, the input and positive samples  $(x, y_w)$  are derived from logged interactions in offline data, while negative samples  $y_l$  are drawn from non-interacted items. Given an input  $x$ , the probability  $p_{\mathcal{D}}(y|x)$  represents the conditional popularity of item  $y$  in the dataset, and  $q_{\mathcal{D}}(y|x)$  denotes the probability of the same item  $y$  being selected as a negative sample. Using these definitions, the DPO loss in Eq. (3) can be rewritten as:

$$\begin{aligned} \mathcal{L}_{\text{DPO}}(\pi_{\theta}; \pi_{\text{ref}}) &= -\mathbb{E}_{(x, y_w) \sim \mathcal{D}, y_l \sim q_{\mathcal{D}}(\cdot|x)} \ell(\pi_{\theta}, \pi_{\text{ref}}, x, y_w, y_l), \\ \text{with } \ell(\cdot) &= \log \sigma \left[ \beta \log \left( \frac{\pi_{\theta}(y_w|x)}{\pi_{\text{ref}}(y_w|x)} \right) - \beta \log \left( \frac{\pi_{\theta}(y_l|x)}{\pi_{\text{ref}}(y_l|x)} \right) \right]. \end{aligned} \quad (4)$$

In this setting, the optimal policy has a closed-form solution, as stated in the following theorem:

**THEOREM 1.** *The optimal policy  $\pi_{\theta}^*(\cdot|x)$  for the DPO loss defined in Eq. (4) is given by:*

$$\pi_{\theta}^*(y|x) \propto \pi_{\text{ref}}(y|x) \cdot \left( \frac{p_{\mathcal{D}}(y|x)}{q_{\mathcal{D}}(y|x)} \right)^{1/\beta}.$$

The proof is deferred to Appendix A. This result highlights that the optimal policy is proportional to the reference policy  $\pi_{\text{ref}}(y|x)$ , adjusted by the relative likelihood ratio  $\left( \frac{p_{\mathcal{D}}(y|x)}{q_{\mathcal{D}}(y|x)} \right)^{1/\beta}$ .

In most recommendation settings, negative samples are uniformly distributed [2, 10], i.e.,  $q_{\mathcal{D}}(y|x) = \mathcal{U} = \frac{1}{|\mathcal{I}|}$ , where  $\mathcal{I}$  is the set of all candidate items. Additionally, in typical DPO-based preference alignment scenarios,  $\beta$  is constrained to  $0 < \beta < 1$ . Under these conditions, *the DPO loss inherently biases the model toward popular items with higher  $p_{\mathcal{D}}(y|x)$ , exacerbating popularity bias*. In the extreme case where  $\beta \rightarrow 0$ , the optimal policy collapses to recommending only the most popular items, effectively disregarding less frequent but potentially valuable recommendations.

**Remark:** This result is a byproduct of DPO’s loss function. Unlike the forward KL-divergence  $\mathbb{D}_{\text{KL}}(p_{\mathcal{D}}(y|x), \pi_{\theta}(y|x))$  used in the SFT loss in Eq. (2), DPO optimizes the reverse KL-divergence  $\mathbb{D}_{\text{KL}}(\pi_{\theta}(y|x), \pi_{\text{ref}}(y|x))$ . Forward KL-divergence is known for its mass-covering property, which encourages learning an average behavior and is less sensitive to subtle differences in the preference distribution (as illustrated in Fig. 3(a)) [35]. In contrast, the reverse KL-divergence used in DPO promotes mode-seeking behavior [28, 37], guiding the model to focus on the “peaks” of the distribution (Fig. 3(b)).

The issue of DPO has also been highlighted by other researchers [29]. Specifically, Feng et al. [14] derive that DPO suppresses negative samples more aggressively than it elevates positive samples during optimization. Moreover, Azar et al. [1] demonstrate that the empirical optimal policy often drives  $\pi_{\theta}(y_l|x) \rightarrow 0$  for all  $\beta$ , stemming from an underfitting of the potential reward.

In the context of recommendation, this behavior can be detrimental, as it exacerbates the filter bubble issue and undermines user interests by limiting exposure to diverse items [16].

## 5 Method

We present how to address the popularity bias in DPO by utilizing the self-play philosophy. Then we detail the proposed SPRec architecture.

### 5.1 Solution: Suppress Biases through Self-Play

Since the DPO loss inherently causes the policy  $\pi_{\theta}$  to learn sharp “peaks”, leading to bias, an intuitive solution is to directly suppress

these learned peaks. To address this, we utilize a self-play framework, dubbed SPDPO, which iteratively alternates between policy learning and bias suppression. Specifically, in the  $(t+1)$ -th iteration, negative samples are drawn from the model’s predictive distribution  $\pi_{\theta_t}(\cdot|x)$  at iteration  $t$ , resulting in the following learning paradigm:

$$\pi_{\theta_{t+1}} \leftarrow \arg \max_{\pi_{\theta}} \mathbb{E}_{(x, y_w) \sim D, y_l \sim \pi_{\theta_t}(\cdot|x)} l(\pi_{\theta}; \pi_{\theta_t}; x, y_w, y_l). \quad (5)$$

By comparing Eq. (5) with Eq. (4), we obtain that the objective function  $\mathcal{L}_{SPDPO}$  in the  $(t+1)$ -th iteration can be viewed as  $\mathcal{L}_{DPO}$  weighted by  $\frac{\pi_{\theta_t}(y_l|x)}{q_{\mathcal{D}}(y_l|x)}$ , which can be expressed as follows:

$$\mathcal{L}_{SPDPO} = -\mathbb{E}_{(x, y_w) \sim D, y_l \sim q_{\mathcal{D}}(\cdot|x)} \frac{\pi_{\theta_t}(y_l|x)}{q_{\mathcal{D}}(y_l|x)} l(\pi_{\theta}; \pi_{\theta_t}; x, y_w, y_l). \quad (6)$$

Again, if DPO uses negative samples from a discrete uniform distribution  $q_{\mathcal{D}}(y|x) = \mathcal{U} = \frac{1}{|I|}$ , then the objective function  $\mathcal{L}_{SPDPO}$  can be viewed as  $\mathcal{L}_{DPO}$  weighted by  $\pi_{\theta_t}(y_l|x)$ . This highlights that the objective adaptively pays more attention to biased items by increasing their learning rates if they have higher probabilities in the model’s output distribution.

**Remark:** Unlike traditional recommendation methods that predefine negative samples or allocate weights in advance, our approach dynamically selects negative samples during the learning process. This provides a significant advantage, enabling the model to adaptively adjust its learning paradigm for effective bias suppression. As a result, this approach mitigates the filter bubble issue and enhances the diversity of recommendations.

## 5.2 Architecture of SPRec

Utilizing the loss function in Eq. (5), we propose a self-play recommendation tuning framework, SPRec, which generally includes multiple iterations of both an SFT step and a DPO step. The workflow is illustrated in Fig. 3(c), in which three key steps are conducted sequentially in each iteration:

- (1) **Dataset Construction:** For each positive sample  $\{(x^i, y_w^i)\}$  in the offline dataset, sample a negative sample  $y_l^i$  by running the current model  $\pi_{\theta_t}$  and using its predicted recommendation as  $y_l^i$ . Thus we obtain pairwise preference data for each sample as  $\{(x^i, y_w^i, y_l^i)\}$ .
- (2) **SFT Step:** Use only the positive sample  $\{(x^i, y_w^i)\}$  to refine the model  $\pi_{\theta_t}$  through SFT techniques such as instruction learning.
- (3) **DPO Step:** Align the model  $\pi_{\theta_t}$  by perform DPO step using pairwise dataset  $\{(x^i, y_w^i, y_l^i)\}$ , and obtain  $\pi_{\theta_{t+1}}$ .

This process repeats for  $T$  iterations per epoch. The self-play mechanism is adaptable to any LLM-based recommender system. To ensure comparability with existing DPO-based recommenders [2, 10, 25], we can extend from a single to multiple negative samples, with results analyzed in experiments.

**Remark:** Although the loss function in Eq. (5) is inherently capable of aligning the model with positive samples, our experiments reveal that incorporating an SFT step in each self-play iteration can further enhance performance.

In fact, combining SFT and DPO has been shown to be an effective practice in recent research and open-sourced LLM models.

For example, each iteration of post-training for Llama 3 includes an SFT stage followed by a DPO stage [13]. Similarly, Pang et al. [30] demonstrate the effectiveness of an iterative preference optimization algorithm using a modified DPO loss with an additional negative log-likelihood (NLL) term, which mirrors the SFT loss defined in Eq. (1).

## 6 Experiments

In this section, we conduct experiments to address the following research questions:

- **RQ1:** How does the SPRec training framework compare to baseline methods in terms of accuracy, diversity, and fairness?
- **RQ2:** What are the contributions of different components within the SPRec framework?
- **RQ3:** How do the random sampling ratio and the number of negative samples impact the performance?

### 6.1 Experimental Setup

**6.1.1 Datasets.** We conducted extensive experiments on four real-world datasets: *MovieLens*<sup>2</sup>, *Steam*<sup>3</sup>, *Goodreads*<sup>4</sup>, and the *CDs* and *Vinyl* category of the Amazon Review Dataset<sup>5</sup>. Additional details about the datasets are provided in the Appendix B. Following the data processing approach in [4, 10], interaction sequences with fewer than 10 entries were excluded. The datasets were then split chronologically into training, validation, and test sets in an 8:1:1 ratio, ensuring mutual exclusivity and preventing data leakage. To ensure comparability across different LLM-based methods, we further sampled 4,096 interactions from each dataset’s training set as the training samples for all methods, 512 interactions from the validation set, and 1,000 interactions from the test set.

To process category information, we extracted category metadata from each dataset and identified the most 10 popular categories within the training sets. To ensure category independence, we removed categories with clear hierarchical relationships, such as “FPS” and “Shooting” in the Steam dataset, and “Rock” and “Classical Rock” in the CDs and Vinyl dataset.

**6.1.2 Evaluation Setting.** To leverage the strengths of LLMs in generative recommendation tasks, we prompt the LLM to generate a predicted item based on the input history sequence. Then, following the procedures in BIGRec [3], we calculate scores and rankings for the entire item space and ground our predicted item to an exact item in the dataset.

**6.1.3 Metrics.** We evaluate the model on 1,000 randomly sampled test cases per iteration using four key metrics. Accuracy is measured by NDCG@5 and HR@5, averaged across results. Diversity is assessed via DivRatio, representing the proportion of unique recommendations. Over-recommendation is quantified by ORRatio, indicating the proportion of results dominated by the three most frequently recommended items. Fairness is evaluated using MGU [22], capturing category-level discrepancies between recommendations and user history.

<sup>2</sup><https://grouplens.org/datasets/movielens/>

<sup>3</sup>[https://cseweb.ucsd.edu/~jmcauley/datasets.html#amazon\\_reviews](https://cseweb.ucsd.edu/~jmcauley/datasets.html#amazon_reviews)

<sup>4</sup><https://mengtingwan.github.io/data/goodreads>

<sup>5</sup>[http://jmcauley.ucsd.edu/data/amazon/index\\_2014.html](http://jmcauley.ucsd.edu/data/amazon/index_2014.html)






Figure 4: Comparison of models across genres on Group Unfairness (GU) in top-1 recommendation.

Table 1: Overall performance comparison of SPRec (green), SFT-based (brown), and DPO-based (blue) methods. Best results are bold, sub-optimal ones underlined.  $\uparrow$  indicates higher is better, while  $\downarrow$  indicates lower is better.

Dataset	Model	DivRatio $\uparrow$	ORRatio $\downarrow$	MGU $\downarrow$	HR $\uparrow$	NDCG $\uparrow$
MovieLens	SASRec	0.0031	1.0000	0.1209	0.0225	0.0136
	BIGRec	0.1939	0.2561	0.0620	0.0347	0.0281
	RW	0.1918	0.2551	0.0577	0.0327	0.0276
	D <sup>3</sup>	0.1246	0.3238	0.0664	0.0266	0.0197
	DMPO	0.1827	0.2561	0.0529	0.0310	0.0264
	SDPO	0.1816	0.2449	0.0462	0.0310	0.0258
	RosePO	0.1857	0.2378	0.0538	0.0290	0.0244
	SPRec	<b>0.2806</b>	<b>0.1510</b>	<b>0.0432</b>	<b>0.0388</b>	<b>0.0319</b>
	SASRec	0.0030	1.0000	0.0458	0.0202	0.0139
	BIGRec	0.1420	0.3170	0.0175	0.0310	0.0236
RW	0.1050	0.3730	0.0238	0.0380	0.0281	
D <sup>3</sup>	0.1199	0.2581	0.0260	0.0413	<b>0.0324</b>	
DMPO	0.1560	0.3310	0.0164	0.0410	0.0314	
SDPO	0.1580	0.3270	0.0161	<b>0.0420</b>	0.0315	
RosePO	0.1860	0.3230	0.0181	0.0300	0.0215	
SPRec	<b>0.2090</b>	<b>0.2170</b>	<b>0.0099</b>	0.0330	0.0250	
CDs_and_Vinyl	SASRec	0.0021	1.0000	0.1205	0.0226	0.0170
	BIGRec	0.3198	0.2597	0.0461	0.0132	0.0130
	RW	0.2821	0.2648	0.0432	0.0112	0.0103
	D <sup>3</sup>	0.3135	0.2672	0.0320	0.0103	0.0088
	DMPO	0.3116	0.1740	<b>0.0195</b>	0.0090	0.0090
	SDPO	0.3218	0.2118	0.0380	0.0110	0.0106
	RosePO	0.3625	0.1853	0.0426	0.0090	0.0090
	SPRec	<b>0.3859</b>	<b>0.1670</b>	0.0242	<b>0.0143</b>	<b>0.0140</b>
	SASRec	0.0010	1.0000	0.1094	0.0660	0.0379
	BIGRec	0.1940	0.3910	0.0650	0.0780	0.0766
RW	0.2890	0.2710	<b>0.0313</b>	0.0760	0.0735	
D <sup>3</sup>	0.1580	0.4560	0.0511	0.0730	0.0718	
DMPO	0.2270	0.2990	0.0443	0.0850	0.0834	
SDPO	0.2080	0.3510	0.0475	0.0820	0.0810	
RosePO	0.2310	0.3160	0.0499	0.0820	0.0805	
SPRec	<b>0.2930</b>	<b>0.2560</b>	0.0367	<b>0.0910</b>	<b>0.0893</b>	

6.1.4 *Baseline.* For traditional recommendation models, we select SASRec [23], a widely used baseline employing a sequential method with a self-attention mechanism. For LLM-based models, we consider several baselines. (1) For SFT-based methods, BIGRec [3]

serves as an instruction-tuning LLM framework for sequential recommendations and forms the foundation for SPRec. **Re-weighting (RW)** [22] improves fairness in BIGRec by balancing recommendations across categories through dataset-based training weights. **Debiasing-Diversifying Decoding (D<sup>3</sup>)** [4] enhances diversity in BIGRec using a decoding strategy guided by SASRec. (2) For DPO-based models, **DMPO** [2] introduces DPO into LRSs by sampling multiple negative items as rejected responses, while **Softmax-DPO (SDPO)** [10] follows a similar approach but incorporates a softmax loss over multiple negative samples. Finally, **RosePO** [25] is a preference optimization framework that combines negative sampling strategies and personalized uncertainty to achieve fairness, unbiasedness, and robustness. The implementation details are listed in Appendix C.

## 6.2 Overall Performance Comparison (RQ1)

The experimental results are presented in Table 1, leading to the following observations. The non-LLM baseline, SASRec, performs poorly with the given training size, which is expected as SASRec requires large datasets to achieve effective fitting. In this study, we primarily focus on LLM-based methods, and SASRec’s results are included only for reference and as the assistant model for D<sup>3</sup> during the decoding stage.

6.2.1 *Limitations of SFT-based Methods.* Fine-tuning LLMs with instruction-based methods results in recommendations heavily favoring popular items, leading to a lack of diversity. For example, in the Goodreads dataset, the DivRatio of BIGRec is only 0.142, meaning the model provides just 14 distinct recommendations per 100 tasks. Similarly, in the Steam dataset, BIGRec’s ORRatio reaches 0.391, with over 39% of recommendations concentrated on the 3 most popular items. These findings highlight that relying solely on SFT introduces severe biases, significantly overexposing certain popular items.

6.2.2 *Limitations of DPO-based Methods.* For DPO methods using random sampling, such as SDPO and DMPO, while multiple negative samples improve recommendation accuracy, they perform poorly on diversity and fairness metrics. On the Goodreads and MovieLens datasets, SDPO and DMPO have minimal impact on DivRatio and ORRatio and may even degrade model performance. On the CDs and Steam datasets, although ORRatio decreases, diversity metrics remain largely unchanged, suggesting that the model favors moderately popular items but fails to effectively recommend new ones. In contrast, RosePO performs well on the CD dataset due to

**Table 2: Ablation results. “with RN” for random negative samples, “w/o” for without specific components.**

Dataset	Model	DivRatio $\uparrow$	ORRatio $\downarrow$	MGU $\downarrow$	HR $\uparrow$	NDCG $\uparrow$
MovieLens	w/o SFT	<b>0.3020</b>	<b>0.0837</b>	<b>0.0198</b>	0.0184	0.0149
	w/o DPO	0.1959	0.2714	0.0637	0.0316	0.0260
	with RN	0.2194	0.2224	0.0544	0.0286	0.0230
	SPRec	0.2806	0.1510	0.0432	<b>0.0388</b>	<b>0.0319</b>
Goodreads	w/o SFT	0.2010	0.2390	<b>0.0044</b>	0.0270	0.0206
	w/o DPO	0.1350	0.2970	0.0142	0.0350	0.0274
	with RN	0.1380	0.3380	0.0188	<b>0.0420</b>	<b>0.0310</b>
	SPRec	<b>0.2090</b>	<b>0.2170</b>	0.0099	0.0330	0.0250
CDs_and_Vinyl	w/o SFT	0.3381	0.2373	<b>0.0216</b>	0.0132	0.0126
	w/o DPO	0.3136	0.2363	0.0333	0.0143	0.0136
	with RN	0.3625	0.2536	0.0359	<b>0.0163</b>	<b>0.0150</b>
	SPRec	<b>0.3859</b>	<b>0.1670</b>	0.0242	0.0143	0.0140
Steam	w/o SFT	0.2900	<b>0.2260</b>	<b>0.0173</b>	0.0880	0.0868
	w/o DPO	0.2220	0.3910	0.0620	0.0790	0.0776
	with RN	0.2860	0.2530	0.0351	0.0860	0.0837
	SPRec	<b>0.2930</b>	0.2560	0.0367	<b>0.0910</b>	<b>0.0893</b>

its negative sampling strategy based on semantic information. However, this approach heavily relies on the semantic characteristics of the dataset’s structure, resulting in relatively poor performance on other datasets and limiting its generalizability for debiasing.

In summary, existing DPO-based methods fail to address fairness issues in LRS.


**6.2.3 Superiority of SPRec.** As shown in Table 1, SPRec significantly improves both DivRatio and ORRatio metrics across all datasets compared to BIGRec, demonstrating its effectiveness in mitigating the over-recommendation of popular items and enhancing diversity. Additionally, SPRec outperforms BIGRec on most fairness metrics, reducing the discrepancies between the model’s recommendations and users’ historical sequences, thereby providing fairer recommendations.

SPRec also surpasses all baseline models on DivRatio and ORRatio, showcasing its superior ability to balance recommendation distributions. For fairness, SPRec achieved the highest MGU scores on the MovieLens and Goodreads datasets, and the second-highest on the Steam and CD datasets. Moreover, as shown in Fig. 4, SPRec alleviates category-level unfairness on the MovieLens dataset, achieving the best results in 7 out of 8 categories, further underscoring its effectiveness in improving fairness.

While RosePO performs well on the CDs and Vinyl dataset, leveraging semantic-based negative sampling to address fairness in music recommendations, and Re-weighting shows strong performance on the Steam dataset by employing category-based re-weighting for gaming recommendations, these methods are tailored to specific datasets and lack generalizability. In contrast, SPRec’s self-play framework provides a universal solution, overcoming dataset-specific challenges and delivering fairer recommendations across diverse scenarios.

### 6.3 Ablation Study (RQ2)

We conducted a series of ablation experiments to explore the impact of each component of the SPRec training framework.

**Figure 5: Performance on the MovieLens dataset across different ablation experiments.**

**6.3.1 SPRec without SFT.** As shown in Table 2, SPRec w/o SFT achieves the lowest recommendation accuracy across all datasets except Steam. This indicates that, during the self-play process, the model’s excessive focus on fairness compromises its accuracy. In the MovieLens dataset (Fig. 5), the absence of the SFT stage leads to a steady decline in recommendation accuracy (NDCG) throughout training. These findings highlight the critical role of SFT in maintaining SPRec’s recommendation performance.

**6.3.2 SPRec without DPO.** Removing DPO reduces SPRec to further SFT training, ensuring that performance gains are not due to incorporating additional data. As shown in Table 2, additional SFT fails to improve diversity or fairness metrics, and the recommendations remain biased. Furthermore, Fig. 5 reveals minimal fluctuations during training, indicating that the prior SFT training has already converged. This experiment underscores the limitations of SFT-based methods in addressing recommendation fairness and diversity.

**6.3.3 Randomly sampling negative items.** As observed in Table 2, when the negative sampling strategy is replaced with random sampling, SPRec-RN fails to achieve further improvements in DivRatio and ORRatio metrics on the MovieLens and Goodreads datasets. Additionally, SPRec-RN’s fairness metrics perform worse compared to SPRec. Although SPRec-RN shows a significant improvement in DivRatio on the CDs and Vinyl dataset, its ORRatio still performs poorly. This suggests that random sampling of negative samples during training is ineffective at suppressing popular items, and the recommendation results continue to exhibit a significant long-tail effect. This ablation experiment demonstrates that our Self-play negative sampling strategy effectively balances the distribution of the model’s output, leading to debiasing in recommendations. Replacing the negative sampling strategy with random sampling (SPRec-RN) fails to improve DivRatio and ORRatio metrics on the MovieLens and Goodreads datasets (Table 2). Additionally, SPRec-RN exhibits worse fairness metrics compared to SPRec. While it achieves a significant boost in DivRatio on the CDs dataset, its




Figure 6: Effect of random sampling ratio.




Figure 7: Effect of negative sample size.

ORRatio remains poor. These results suggest that random negative sampling is ineffective in suppressing popular items, leaving a pronounced long-tail effect in the recommendations. This experiment demonstrates that our self-play negative sampling strategy effectively balances the model’s output distribution, resulting in debiased recommendations.

#### 6.4 Impact of Negative Samples (RQ3)

We investigate the role of negative samples in SPRec by introducing a proportion of random negative samples to contaminate SPRec’s original self-play samples. Additionally, we examine the impact of increasing the number of negative samples in SPRec’s loss function (Eq. (6)). To achieve this, we adopt the SDPO loss function to expand to multiple negative samples. For efficiency, we limit the training sample size to 1,024, keeping other experimental settings consistent with Section 6.2. To ensure result stability, the model’s performance is averaged over the last three training iterations. We report the results on the MovieLens dataset.

**6.4.1 Effect of Random Sampling Ratio.** In each training iteration, we randomly replace a proportion of negative samples with randomly selected items while leaving the remaining negative samples unchanged as the model’s recommendation outputs. As shown in Fig. 6, increasing the proportion of random negative samples leads to a steady decline in diversity and accuracy. Fairness also deteriorates, with recommendations becoming more skewed toward

a small subset of popular items. These results highlight the superiority of our proposed self-play negative sampling strategy over random sampling.

**6.4.2 Effect of Negative Sample Size.** To generate  $N$  negative samples, we use beam search decoding to sample  $2N$  items from the model’s output. After deduplication, the top  $N$  items with the highest predicted probabilities are selected as negative samples. As shown in Fig. 7, increasing the number of negative samples results in stable recommendation accuracy but significantly improves diversity and fairness, reducing the focus on popular items. This demonstrates the versatility of the self-play negative sampling strategy, which can be effectively combined with multi-negative sampling approaches to further debias LRS.

## 7 Conclusion & Discussion

Our work establishes a critical bridge between preference alignment techniques and fairness-aware recommendation in the era of LLMs. Through both theoretical analysis and empirical validation, we demonstrate that conventional DPO-based tuning fundamentally conflicts with the principles of equitable recommendation, creating self-reinforcing popularity biases that traditional debiasing approaches fail to address. The proposed SPRec framework represents a paradigm shift in recommendation alignment - rather than treating bias mitigation as a post-hoc correction, we redesign the core learning mechanism to enable autonomous bias suppression through self-regulated competition between model generations. This approach not only achieves state-of-the-art performance across accuracy and fairness metrics but more importantly, provides a blueprint for developing self-correcting AI systems that maintain alignment with both user preferences and ethical constraints.

Despite its effectiveness, our work primarily addresses bias in DPO-based tuning, while overlooking the popularity bias already present in SFT due to its cross-entropy loss. Future research should focus on mitigating bias at the SFT stage to ensure fairness from the start of fine-tuning. Additionally, optimizing preferences in recommendation is a long-term challenge, requiring alignment across sequential recommendations rather than individual predictions. However, LLMs generate outputs token by token, making it difficult to optimize preferences from token-level to item-level and ultimately list-level recommendations. Tackling this issue will require new datasets, benchmarks, and models capable of long-term alignment. A promising direction is reinforcement learning with process-level rewards, shifting optimization from short-term token likelihood to long-horizon user engagement.

## Acknowledgements

This work is supported by the National Key Research and Development Program of China (2021ZD0111802), the National Natural Science Foundation of China (62402470, 62272437, U24B20180, 62121002), the Fundamental Research Funds for the Central Universities of China (WK210000053, PA2024GDSK0107), Anhui Provincial Natural Science Foundation (2408085QF189), and the Postdoctoral Fellowship Program of CPSF (GZC20241643). This research is supported by the advanced computing resources provided by the Supercomputing Center of the USTC.



## References

- [1] Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics (AISTATS '24)*. PMLR, 4447–4455.
- [2] Zhuoxi Bai, Ning Wu, Fengyu Cai, Xinyi Zhu, and Yun Xiong. 2024. Aligning Large Language Model with Direct Multi-Preference Optimization for Recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*. 76–86.
- [3] 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 (TORS)* (2025).
- [4] 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. *EMNLP* (2024).
- [5] Shihao Cai, Jizhi Zhang, Keqin Bao, Chongming Gao, and Fuli Feng. 2024. FLOW: A Feedback LooP FrameWork for Simultaneously Enhancing Recommendation and User Agents. *arXiv preprint arXiv:2410.20027* (2024).
- [6] Daniele Calandriello, Zhaohan Daniel Guo, Remi Munos, Mark Rowland, Yunhao Tang, Bernardo Avila Pires, Pierre Harvey Richemond, Charline Le Lan, Michal Valko, Tianqi Liu, Rishabh Joshi, Zeyu Zheng, and Bilal Piot. 2024. Human alignment of large language models through online preference optimisation. In *Proceedings of the 41st International Conference on Machine Learning (ICML '24)*. Article 211, 27 pages.
- [7] Chong Chen, Weizhi Ma, Min Zhang, Chenyang Wang, Yiqun Liu, and Shaoping Ma. 2023. Revisiting Negative Sampling vs. Non-sampling in Implicit Recommendation. *ACM Transactions on Information Systems (TOIS)* 41, 1, Article 12 (Feb. 2023), 25 pages.
- [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 Trans. Inf. Syst.* 41, 3, Article 67 (Feb. 2023), 39 pages.
- [9] Jiayu Chen, Chongming Gao, Shuai Yuan, Shuchang Liu, Qingpeng Cai, and Peng Jiang. 2025. DLCCRec: A Novel Approach for Managing Diversity in LLM-Based Recommender Systems. *The 18th ACM International Conference on Web Search and Data Mining (WSDM '25)* (2025).
- [10] 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. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS '24)*.
- [11] Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-Play Fine-Tuning Converts Weak Language Models to Strong Language Models. In *Forty-first International Conference on Machine Learning (ICML '24)*.
- [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 (KDD '24)*. 6437–6447.
- [13] 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).
- [14] Duanyu Feng, Bowen Qin, Chen Huang, Zheng Zhang, and Wenqiang Lei. 2024. Towards Analyzing and Understanding the Limitations of DPO: A Theoretical Perspective. *arXiv preprint arXiv:2404.04626* (2024).
- [15] Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. Bias and Fairness in Large Language Models: A Survey. *Computational Linguistics* 50, 3 (Sept. 2024), 1097–1179.
- [16] Chongming Gao, Kexin Huang, Jiawei Chen, Yuan Zhang, Biao Li, Peng Jiang, Shiqi Wang, Zhong Zhang, and Xiangnan He. 2023. Alleviating Matthew Effect of Offline Reinforcement Learning in Interactive Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*. 11 pages.
- [17] Chongming Gao, Shiqi Wang, Shijun Li, Jiawei Chen, Xiangnan He, Wenqiang Lei, Biao Li, Yuan Zhang, and Peng Jiang. 2023. CIRS: Bursting Filter Bubbles by Counterfactual Interactive Recommender System. *ACM Transactions on Information Systems (TOIS)* 42, 1, Article 14 (aug 2023), 27 pages.
- [18] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. 2023. Chat-rec: Towards interactive and explainable llms-augmented recommender system. *arXiv preprint arXiv:2303.14524* (2023).
- [19] Binzong Geng, Zhaoxin Huan, Xiaolu Zhang, Yong He, Liang Zhang, Fajie Yuan, Jun Zhou, and Linjian Mo. 2024. Breaking the Length Barrier: LLM-Enhanced CTR Prediction in Long Textual User Behaviors. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. 2311–2315.
- [20] 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 *Advances in Information Retrieval: 46th European Conference on Information Retrieval, ECIR 2024, Glasgow, UK, March 24–28, 2024, Proceedings, Part II*. 364–381.
- [21] Meena Jagadeesan, Nikhil Garg, and Jacob Steinhardt. 2024. Supply-side equilibria in recommender systems. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NeurIPS '23)*. Article 642, 12 pages.
- [22] 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 on Web Conference 2024 (WWW '24)*. 4717–4726.
- [23] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [24] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. 2023. Fairness in Recommendation: Foundations, Methods, and Applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* 14, 5, Article 95 (Oct. 2023), 48 pages.
- [25] 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).
- [26] Qidong Liu, Xian Wu, Xiangyu Zhao, Yejing Wang, Zijian Zhang, Feng Tian, and Yefeng Zheng. 2024. Large Language Models Enhanced Sequential Recommendation for Long-tail User and Item. *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [27] Wenyu Mao, Jiancan Wu, Weijian Chen, Chongming Gao, Xiang Wang, and Xiangnan He. 2024. Reinforced Prompt Personalization for Recommendation with Large Language Models. *ACM Transactions on Information Systems (TOIS)* (2024).
- [28] Motoki Omura, Yasuhiro Fujita, and Toshiki Kataoka. 2024. Entropy Controllable Direct Preference Optimization. *arXiv preprint arXiv:2411.07595* (2024).
- [29] Arka Pal, Deep Karkhanis, Samuel Dooley, Manley Roberts, Siddhartha Naidu, and Colin White. 2024. Smaug: Fixing failure modes of preference optimisation with dpo-positive. *arXiv preprint arXiv:2402.13228* (2024).
- [30] Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and Jason Weston. 2024. Iterative Reasoning Preference Optimization. *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [31] 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).
- [32] Wentao Shi, Jiawei Chen, Fuli Feng, Jizhi Zhang, Junkang Wu, Chongming Gao, and Xiangnan He. 2023. On the Theories Behind Hard Negative Sampling for Recommendation. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*. 812–822.
- [33] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmarajan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419 (2018), 1140–1144.
- [34] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *nature* 550, 7676 (2017), 354–359.
- [35] Hao Sun and Mihaela van der Schaar. 2024. Inverse-RLignment: Inverse Reinforcement Learning from Demonstrations for LLM Alignment. *arXiv preprint arXiv:2405.15624* (2024).
- [36] Antonela Tommasel. 2024. Fairness Matters: A look at LLM-generated group recommendations. In *Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24)*. 993–998.
- [37] Chaoqi Wang, Yibo Jiang, Chenghao Yang, Han Liu, and Yuxin Chen. 2024. Beyond Reverse KL: Generalizing Direct Preference Optimization with Diverse Divergence Constraints. In *The Twelfth International Conference on Learning Representations (ICLR '24)*.
- [38] Qi Wang, Jindong Li, Shiqi Wang, Qianli Xing, Runliang Niu, He Kong, Rui Li, Guodong Long, Yi Chang, and Chengqi Zhang. 2024. Towards Next-Generation LLM-based Recommender Systems: A Survey and Beyond. *arXiv preprint arXiv:2410.19744* (2024).
- [39] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Transactions on Information Systems (TOIS)* 41, 3, Article 52 (Feb. 2023), 43 pages.
- [40] Wei Wei, Xubin Ren, Jiabin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. LLMRec: Large Language Models with Graph Augmentation for Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*. 806–815.
- [41] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, Hui Xiong, and Enhong Chen. 2024. A Survey on Large Language Models for Recommendation. *World Wide Web* 27, 5 (Aug. 2024), 31 pages.
- [42] Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. 2025. Self-Play Preference Optimization for Language Model Alignment. In

- The Thirteenth International Conference on Learning Representations (ICLR '2025).*
- [43] Chen Xu, Wenjie Wang, Yuxin Li, Liang Pang, Jun Xu, and Tat-Seng Chua. 2024. A Study of Implicit Ranking Unfairness in Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. 7957–7970.
- [44] Fan Yao, Yiming Liao, Jingzhou Liu, Shaoliang Nie, Qifan Wang, Haifeng Xu, and Hongning Wang. 2024. Unveiling User Satisfaction and Creator Productivity Trade-Offs in Recommendation Platforms. *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [45] Jizhi Zhang, Keqin Bao, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Is ChatGPT Fair for Recommendation? Evaluating Fairness in Large Language Model Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems (RecSys '23)*. 993–999.
- [46] 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 (WWW '24)*. 3679–3689.

## A Mathematical Derivations

PROOF OF THEOREM 1. The DPO loss is derived from the objective of Reinforcement Learning with Human Feedback (RLHF):

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\theta}(\cdot|x)} [r(x, y)] - \beta \text{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}), \quad (7)$$

where the reward model is defined via the BT model:

$$\max_r \mathbb{E}_{(x, y_w) \sim \mathcal{D}, y_l \sim Y_u} \log \sigma(r(x, y_w) - r(x, y_l)).$$

In the original paper of DPO [31], the authors proved that the optimal policy  $\pi_{\theta}^*$  for DPO loss in Eq. (4) and the solution to the optimization problem in Eq. (7) are the same. Thus, we can analyze the solution to Eq. (7), equivalent to examining the DPO loss.

Consider a fixed context  $x$  and define  $\hat{r}(y|x)$  as:

$$\hat{r}(y|x) = \exp(r(x, y)), \quad (8)$$

then our goal is to optimize  $\hat{r}(\cdot|x) \in \mathbb{R}_+^{\mathcal{I}}$ , which is a  $|\mathcal{I}|$ -dim vector representing the latent rewards for all items in the recommendation dataset  $\mathcal{I}$ . we can rewrite the reward model's optimization as:

$$\max_{\hat{r}(\cdot|x) \in \mathbb{R}_+^{\mathcal{I}}} \sum_{y_w \in \mathcal{I}} \sum_{y_l \in \mathcal{I}} p_{\mathcal{D}}(y_w|x) q_{\mathcal{D}}(y_l|x) \log \left( \frac{\hat{r}(y_w|x)}{\hat{r}(y_w|x) + \hat{r}(y_l|x)} \right).$$

Then we calculate the gradients:

$$\begin{aligned} & \partial p_{\mathcal{D}}(y_w|x) q_{\mathcal{D}}(y_l|x) \log \left( \frac{\hat{r}(y_w|x)}{\hat{r}(y_w|x) + \hat{r}(y_l|x)} \right) / \partial \hat{r}(y|x) \\ &= \begin{cases} p_{\mathcal{D}}(y|x) q_{\mathcal{D}}(y_l|x) \left( \frac{1}{\hat{r}(y|x)} - \frac{1}{\hat{r}(y|x) + \hat{r}(y_l|x)} \right) & y_w = y, y_l \neq y, \\ -p_{\mathcal{D}}(y_w|x) q_{\mathcal{D}}(y|x) \frac{1}{\hat{r}(y_w|x) + \hat{r}(y|x)} & y_w \neq y, y_l = y, \\ 0 & \text{else.} \end{cases} \end{aligned}$$

Hence, the objective's gradient w.r.t.  $\hat{r}(y|x)$  can be written as:

$$\begin{aligned} & \partial \sum_{y_w \in \mathcal{I}} \sum_{y_l \in \mathcal{I}} p_{\mathcal{D}}(y_w|x) q_{\mathcal{D}}(y_l|x) \log \left( \frac{\hat{r}(y_w|x)}{\hat{r}(y_w|x) + \hat{r}(y_l|x)} \right) / \partial \hat{r}(y|x) \\ &= 0 + \sum_{y_l \neq y} \left[ p_{\mathcal{D}}(y|x) q_{\mathcal{D}}(y_l|x) \left( \frac{1}{\hat{r}(y|x)} - \frac{1}{\hat{r}(y|x) + \hat{r}(y_l|x)} \right) \right] \\ & \quad - \sum_{y_w \neq y} p_{\mathcal{D}}(y_w|x) q_{\mathcal{D}}(y|x) \frac{1}{\hat{r}(y_w|x) + \hat{r}(y|x)} \\ &= \sum_{y_i \in \mathcal{I}} \left[ p_{\mathcal{D}}(y|x) q_{\mathcal{D}}(y_i|x) \left( \frac{1}{\hat{r}(y|x)} - \frac{1}{\hat{r}(y|x) + \hat{r}(y_i|x)} \right) \right] \\ & \quad - \sum_{y_i \in \mathcal{I}} p_{\mathcal{D}}(y_i|x) q_{\mathcal{D}}(y|x) \frac{1}{\hat{r}(y_i|x) + \hat{r}(y|x)} \quad (\text{add } 0) \\ &= \sum_{y_i \in \mathcal{I}} \left[ \frac{p_{\mathcal{D}}(y|x) q_{\mathcal{D}}(y_i|x)}{\hat{r}(y|x)} - \frac{p_{\mathcal{D}}(y_i|x) q_{\mathcal{D}}(y|x) + p_{\mathcal{D}}(y|x) q_{\mathcal{D}}(y_i|x)}{\hat{r}(y_i|x) + \hat{r}(y|x)} \right]. \end{aligned}$$

By setting the gradients to be 0, we obtain that for  $\forall y \in \mathcal{I}$ , the optimal reward  $r^*(y|x)$  is:

$$\hat{r}^*(y|x) \propto \frac{p_{\mathcal{D}}(y|x)}{q_{\mathcal{D}}(y|x)}.$$

By plugging it into Eq. (8), we have:

$$r^*(x, y) = \log p_{\mathcal{D}}(y|x) - \log q_{\mathcal{D}}(y|x) + \text{Constant}.$$

Back to the RLHF objective, we have:

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\theta}(\cdot|x)} [\log p_{\mathcal{D}}(y|x) - \log q_{\mathcal{D}}(y|x)] - \beta \text{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}),$$

which has a well-known closed form solution [31]:

$$\pi_{\theta}^*(y|x) \propto \pi_{\text{ref}}(y|x) \cdot \left( \frac{p_{\mathcal{D}}(y|x)}{q_{\mathcal{D}}(y|x)} \right)^{1/\beta}.$$

□

## B Dataset Statistics

Our datasets span diverse domains, including movies, books, music, and games, offering varied sizes and user interaction patterns to provide a comprehensive basis for evaluating LRSs. Note that while we report the full dataset statistics, only a subset of interaction sequences is sampled for LLM fine-tuning, as detailed in Section 6.1.1.

Table 3: Statistics of Datasets.

Datasets	#Items	#Interactions	#Sequences
MovieLens	10,682	10,000,054	9,301,274
Goodreads	4,058	160,398	6,031
CDs_and_Vinyl	13,078	185,855	21,347
Steam	32,094	178,961	29,876

## C Implementation Details

For LLM-based methods, we adopted Llama-3.2-1B-Instruct as the backbone LLM. Considering the ability of LLMs to quickly adapt to downstream tasks with limited data, we followed BIGRec [3] and used relatively smaller datasets. To ensure fairness in comparison, all baseline methods and SPRec utilize the same dataset as used in the SFT training phase. For SPRec, the total number of iterations was set to 5, with each SFT and DPO phase trained for 1 epoch. To ensure that the training data used in each iteration is not identical, we further randomly sample half of the training data (i.e., 2048 interactions) for training in each iteration. All experiments were carried out on four RTX 3090 GPUs, each with 24GB of VRAM.

For the traditional model SASRec, we use the same training and validation datasets as other LLM-based methods, with dataset sizes of 4,096 and 512, respectively. The embedding size was fixed at 64, and the dropout ratio was set to 0.1. Negative samples were randomly sampled in training, with Adam as the optimizer and a learning rate of  $4e-3$ . More details of the implementation are available via <https://github.com/RegionCh/SPRec>.