

Quantifying and Mitigating Popularity Bias in Conversational Recommender Systems

Allen Lin

Texas A&M University
College Station, Texas, USA
al001@tamu.edu

Ziwei Zhu

George Mason University
Fairfax, Virginia, USA
zzhu20@gmu.edu

Jianling Wang

Texas A&M University
College Station, Texas, USA
jlwang@tamu.edu

James Caverlee

Texas A&M University
College Station, Texas, USA
caverlee@tamu.edu

ABSTRACT

Conversational recommender systems (CRS) have shown great success in accurately capturing a user's current and detailed preference through the multi-round interaction cycle while effectively guiding users to a more personalized recommendation. Perhaps surprisingly, conversational recommender systems can be plagued by popularity bias, much like traditional recommender systems. In this paper, we systematically study the problem of popularity bias in CRSs. We demonstrate the existence of popularity bias in existing state-of-the-art CRSs from an exposure rate, a success rate, and a conversational utility perspective, and propose a suite of popularity bias metrics designed specifically for the CRS setting. We then introduce a debiasing framework with three unique features: (i) *Popularity-Aware Focused Learning* to reduce the popularity-distorting impact on preference prediction; (ii) *Cold-Start Item Embedding Reconstruction* via Attribute Mapping, to improve the modeling of cold-start items; and (iii) *Dual-Policy Learning*, to better guide the CRS when dealing with either popular or unpopular items. Through extensive experiments on two frequently used CRS datasets, we find the proposed model-agnostic debiasing framework not only mitigates the popularity bias in state-of-the-art CRSs but also improves the overall recommendation performance.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Conversational Recommender System, Popularity Bias, Debiasing

ACM Reference Format:

Allen Lin, Jianling Wang, Ziwei Zhu, and James Caverlee. 2022. Quantifying and Mitigating Popularity Bias in Conversational Recommender Systems. In *Proceedings of the 31st ACM International Conference on Information and*

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 ACM 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.

CIKM '22, October 17–21, 2022, Atlanta, GA, USA.

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9236-5/22/10...\$15.00

<https://doi.org/10.1145/3511808.3557423>

Knowledge Management (CIKM '22), October 17–21, 2022, Atlanta, GA, USA.
ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3511808.3557423>

1 INTRODUCTION

Recommender systems have become an indispensable tool in everyone's daily life. Ranging from E-commerce, to multimedia services, and to E-education platforms, recommender systems alleviate information overload by connecting users to their items of interest. While they have been one of the success stories of AI in practice, one long-lasting challenge faced by recommender systems is *popularity bias*, which refers to the phenomenon of the minority popular items being overly exposed to users while the majority unpopular items do not receive their deserved attention [1, 5, 6, 8, 10, 11, 19, 30, 38]. This popularity bias greatly limits the opportunities for users to discover these less exposed items [11, 27, 30] and the system's potential to learn an unbiased view of the user's true preferences [24, 38]. To solve such a critical issue, extensive efforts have been made to investigate and mitigate popularity bias in recommender systems [2, 19, 27, 38]. Recently, a new form of interactive recommendation system – the conversational recommender system (CRS) – has shown great success in enhancing personalization through multi-turn interactions (i.e., dialogue). In this way, the system can elicit users' current and detailed preferences, which can lead to justifiable and highly personalized recommendations. Similar to traditional recommender systems, a CRS also takes into consideration the past user-item interactions when making recommendations. Therefore, it would be reasonable to assume that a CRS is also affected by popularity bias similar to a traditional recommender. However, despite the promising functionality and wide range of potential applications of CRSs, *how to quantify and mitigate the issue of popularity bias in CRSs* remains understudied.

Hence, this paper first illustrates the existence of popularity bias in a multi-round CRS from three perspectives: (i) Exposure Rate, where we find popular items are exposed to users at a significantly higher rate than others, even when these other items match all preferences specified by the user; (ii) Success Rate, where we find items with high popularity have much higher chances of being successfully recommended; and (iii) Conversational Utility, where we find popular items are recommended in earlier rounds than other items, violating the assumption of an unbiased recommender. Based on these observations, we propose a suite of popularity bias metrics designed specifically for the CRS setting.




Figure 1: Evidence of different perspectives of popularity bias in CRSs

We then diagnose the root causes of the observed popularity bias. First, we find that item popularity can significantly impact the magnitude of the item embedding in the recommender component which drives the predicted preference score. Since popular items occur more in the training set (through user-item interactions), the embeddings of popular items are updated more frequently, leading to greater magnitudes, and hence higher predicted scores. Second, we find that cold-start items are typically ill-modeled in existing CRS, meaning that they are rarely successfully recommended to users. Finally, we find that the policy agent which guides the actions of a CRS at each step may further amplify popularity bias by prioritizing specific attributes associated with popular items. Based on this diagnosis, we propose a three-stage debiasing framework that generalizes across specific CRS implementations. To combat the item magnitude problem, we introduce Popularity-Aware Focused Learning (PAL) which is designed to help the recommender component learn a more fine-grained set of embeddings for the unpopular items without sacrificing the effectiveness of the learned embeddings for popular items. To better model cold-start items, we propose a Cold-Start Item Embedding Reconstruction via Attribute Mapping (CSM) to avoid the problem of randomly initialized embeddings for cold-start items by transferring representations from warm-start items to the cold-start ones. To improve the policy agent, we design Dual-Policy Learning (DPL) to train one policy network on popular items and one policy network on unpopular items. Through extensive experiments on both Yelp and Lastfm, we demonstrate that the proposed three-stage debiasing framework can significantly mitigate popularity bias across existing baseline CRS methods by an average of 42.5% while improving their overall recommendation success rate by 12.3%, all while keeping the average conversational turns the same.

2 PRELIMINARIES

The goal of a CRS is to improve the preference elicitation process by directly asking the user’s current preference on a set of domain specific attributes that characterize every item in the itemset. In this work, we follow the System Ask - User Responds (SAUR) setting [35] to investigate popularity bias in multi-round CRS. Formally let U denote the user set, V denote the itemset, and $A = a_1, a_2, \dots, a_m$ denote a set of m domain-specific attributes used to systematically characterize all items in V .

At each turn, a CRS calls upon its conversation component, which is typically a policy agent, to decide whether to ask the user’s preference on a specific attribute or make a recommendation. If the policy agent decides not enough preference evidence has been collected, it will pick one attribute a from the set of unasked attributes to prompt the user. When prompted with the question, the user is assumed to provide her preference p on the asked attribute a . Upon receiving the user’s preference, the policy agent updates its current belief state s by adding the (attribute, preference) pair to it. If the policy agent decides enough information has been collected after t turns of interaction, the CRS then calls upon its recommender component to make a list of recommendations for the user. Unlike a traditional static recommender which ranks all items in the entire itemset, the recommender component of a CRS only ranks items in the candidate itemset V_{cand} – the itemset that contains only items with attributes perfectly matching all (attribute, preference) pairs in the current belief state $s_t = [(a_1, p_1), \dots, (a_t, p_t)]$. After ranking all items in V_{cand} , the CRS recommends the top K items to display to the user. If the user accepts the recommendation, then the system quits. If the user rejects all the recommended items, the system repeats the above-introduced cycle by calling upon its policy agent to decide the next action to take. This process continues until the user quits due to impatience or a predefined maximum number of allowed conversational turns has been reached.

3 EVIDENCE OF POPULARITY BIAS IN CRS

In this section, we conduct data-driven analysis on the Lastfm dataset for music artist recommendation with three state-of-the-art CRSs: SCPR [21], EAR [20], and CRM [28], to illustrate the existence of popularity bias in a multi-round CRS from three perspectives:

Exposure Rate Perspective. A CRS starts by collecting a set of user preferences, from which the system forms a candidate item set containing all items matching the collected user preferences so far. Once the system determines enough information has been collected, it selects the top- N items from the candidate set to generate the recommendation set. As shown in Figure 1(a), the average popularity of items that get recommended by CRM is almost 8 times higher than the average popularity of items in the candidate set. A similar pattern is observed in EAR and SCPR, indicating popular items are much more likely to get recommended (exposed) than unpopular items. However, in a completely unbiased CRS, since

both the candidate set and the recommendation set contain items perfectly matching all preferences specified by the user (i.e., equally qualified items), the overall average popularity for both sets should not deviate excessively. The observed phenomenon suggests the existence of a strong exposure based bias in all three CRS models.

Success Rate Perspective. Next, we investigate whether there exists a correlation between an item's popularity and recommendation success rate (i.e., how likely the item will be recommended to the user within the allowed turns of interactions). In a completely unbiased system, an item's popularity should be completely irrelevant to its recommendation success rate. Any correlation, either positive or negative, is an indication of popularity bias in the CRS. As we can observe from Figure 1(b), all three models – SCPR, EAR, and CRM – exhibit strong positive correlations between an item's popularity and its recommendation success rate. For example, in EAR, items with popularity greater than 0.01 – the 75th percentile of the sorted popularity of all items – have over 75% recommendation success rate, while items with popularity lesser than 0.01 having significantly lower. That is, items with high popularity have much higher chances of being successfully recommended, indicating the existence of strong success rate bias in SCPR, EAR, and CRM.

Conversational Utility Perspective. Unlike traditional static recommenders, a CRS operates under a multi-round interaction setting, in which a user might leave due to impatience when the session is taking too long. Intuitively, items requiring fewer turns of user interactions (i.e., dialogue) to get successfully recommended should be considered more advantageous than those requiring more. Thus in an unbiased CRS, an item's required turns of user interactions should be completely independent of its popularity. However, as shown in Figure 1(c), all three models exhibit strong inverse correlations between an item's popularity and its required turns of user interaction to get successfully recommended. This observation shows that an item's conversational utility is heavily dependent upon its popularity, indicating the existence of strong conversational utility bias in SCPR, EAR, and CRM.

4 QUANTIFYING POPULARITY BIAS IN CRS

After demonstrating the existence of popularity bias in multi-round CRSs, we formally define three bias metrics for quantifying the degree of popularity bias from each of the three perspectives introduced in Section 3.

Popularity Correlation with Exposure Rate (PER). Given the deviation of average item popularity between the recommended items and the candidate items as illustrated in Figure 1(a), we seek to measure popularity bias from a more comprehensive perspective that also accounts for the rank of popular items on the recommendation list (ranking utility). To start, we consider the correlation between popularity and exposure rate first formally quantified in [11]. Let A be a set of m attributes associated with an itemset V , [11] defines the exposure-based bias at the n^{th} turn to be:

$$E_{A_n}[\hat{V}_{u,n}] = \sum_{i=1}^m \mathbb{1}[a_i \in A_n] \rho(a_i|\hat{V}_{u,n}) P(a_i|\hat{V}_{u,n}), \quad (1)$$

where $\hat{V}_{u,n}$ denotes the set of potential recommendations made by the CRS at the n^{th} interaction (turn). A_n denotes the set of




Figure 2: Success Rate Bias Measured via Gini Coefficient

attributes associated with all invoked questions till the n^{th} turn, and $\mathbb{1}[\cdot]$ denotes whether an attribute a exists in A_n . $\rho(a_i|\hat{V}_{u,n})$ measures the ranking utility of all popular items, defined based on a certain attribute a_i , on the recommendation list. The higher the popular items are ranked, the greater $\rho(a_i|\hat{V}_{u,n})$ will be. $P(a_i|\hat{V}_{u,n})$ measures the number of popular items on the recommendation list. The more frequently popular items are recommended, the greater $P(a_i|\hat{V}_{u,n})$ will be. Together, $E_{A_n}[\hat{V}_{u,n}]$ considers both the ranking utilities of popular items and the frequency of them being recommended.

However, Equation 1 cannot be easily adapted to the conventional interaction-frequency based definition of popularity. Referring back to Equation 1, the notion of popularity is based upon a particular attribute a_i . For instance, in a *Movie* dataset, the attribute *production* has over 15 values, but *Columbia Pictures, Disney, and Lions Gate Films* alone account for more than 50% of the movies in the training data. In this case, items (movies) that are produced by *Columbia Pictures, Disney, or Lions Gate Films* would be deemed as popular items; however, such a definition has a slight limitation. Referring back to the same *Movie* dataset, the attribute *Duration* has 4 different values, but $>=90\text{min}$ alone accounts for more than 60% of the movies in the training data. In this case, items (movies) with durations greater or equal to 90 minutes would be classified as popular items; however, there could be an item (movie) that has a duration of less than 90 minutes but is produced by *Columbia Pictures*. In this case, this movie would be classified as popular by the attribute *production* but unpopular by the attribute *duration*. Therefore, to adapt to a more generalizable definition of popularity, we modify the existing metric to be:

$$E_n[\hat{V}_{u,n}] = \rho(\hat{V}_{u,n}) P(\hat{V}_{u,n}), \quad (2)$$

with $\rho(a_i|\hat{V}_{u,n})$ updated to be:

$$\rho(\hat{V}_{u,n}) = \sum_{v \in \hat{V}_u} \frac{\mathbb{1}[v \in \hat{V}_u^{pop}(t^{pop})]}{\log(\text{rank}(v) + 1)},$$

and $P(a_i|\hat{V}_{u,n})$ updated to be:

$$P(\hat{V}_{u,n}) = \frac{|\hat{V}_{u,n}^{pop}(t^{pop})|}{|\hat{V}_u|},$$

where t^{pop} denotes the popularity threshold and $\hat{V}_{u,t}^{pop}(t^{pop})$ denotes recommended items that have popularity greater or equal to

the threshold. We adopt the conventional definition of $\frac{|U_{interacted}|}{|U|}$ for measuring the popularity of an item and categorize items with popularity greater than the 80th percentile as popular (head) items [3–5]. Similar to Equation 1, this metric (Equation 2) also measures how frequently popular items are recommended (exposed) to the user along with their ranking utilities; however, this metric is attribute independent, making it more generalizable and consistent across different datasets. Formally we define the *popularity correlation with exposure rate (PER)* to be:

$$PER = \frac{1}{|I|} \sum_{(u,i) \in I} \frac{1}{N} \sum_{n=1}^N E_n [\hat{V}_{u,n}], \quad (3)$$

where I denotes the user-item interactions and N denotes the number of total interactions (turns).

Popularity Correlation with Success Rate (PSR). While Figure 1(b) provides the illustration, to the best of our knowledge, there has not been a well-defined metric for quantifying the degree of popularity bias in a multi-round CRS from a recommendation success rate perspective. However, a few previous works have demonstrated the effectiveness of using the Gini Coefficient for measuring the correlation between recommendation accuracy and item popularity [9, 26, 37]. Inspired by these works, we formally define the *popularity correlation with success rate (PSR)* to be:

$$PSR = 1 - 2 \int_0^1 SR(V_{sort}) dV_{sort}, \quad (4)$$

where V_{sort} denotes a set of items sorted in ascending popularity, and $SR(V_{sort})$ denotes the Lorenz curve constructed with the x-axis representing items sorted in ascending popularity and the y-axis representing the average success rate accumulated so far. As illustrated in Figure 2, in a perfectly unbiased CRS, the average success rate should be accumulated at the same rate for all items regardless of their popularity, making $SR(V_{sort})$ a straight line with a constant slope of 1. In this case, the AUC would be 0.5, resulting in a Gini Coefficient of 0. However, when a positive correlation exists between popularity and success rate, $SR(V_{sort})$ would instead be a concave hyperbolic curve with increasing slope, making the AUC less than 0.5. In this case, the Gini Coefficient would be a positive number ranging from (0,1), the closer to 1 the stronger the bias.

Popularity Correlation with Conversational Utility (PCU). To the best of our knowledge, this is the first work that investigates popularity bias in a multi-round CRS from a conversational utility perspective. Therefore, there does not exist any well-defined metric that can be applied. Building upon the notion that an item's required turns of user interactions should be completely independent of its popularity, we calculate the averaged turns of user interactions for both the popular (head) items and the unpopular (tail) items and use the difference (gap) between the two as our metric. Formally we define the *popularity correlation with conversational utility (PCU)* to be:

$$PCU = \frac{1}{(T)} \left[\frac{1}{|V^{unpop}|} \sum_{v \in V^{unpop}} U_{turn}(v_i) - \frac{1}{|V^{pop}|} \sum_{v \in V^{pop}} U_{turn}(v_i) \right], \quad (5)$$

where V^{pop} and V^{unpop} respectively denote the popular (head) and unpopular (tail) itemset, $U_{turn}(v)$ denotes the needed turns of user

interactions to get successfully recommended item v . T denotes the total turns of allowed interactions which serves as the normalizing constant. Following the convention [3–5], we define items with popularity greater than the 80th percentile to be popular and items with popularity less than the 20th percentile to be unpopular.

5 DEBIASING FRAMEWORK

In this section, we identify the three primary causes of popularity bias in CRSs and propose a three-stage debiasing framework that can be easily generalized. The framework starts with the *popularity-aware focused learning* stage where the recommender component learns to fairly model all items by taking into account their popularity. After learning an effective set of embeddings for all items, the framework builds a *attribute to item embedding mapping* to reconstruct embeddings for items with zero user interactions (the cold-start items). Lastly, to help reduce the number of *popularity-based false assumptions* that the policy agent makes, the framework separately trains two policy agents, one for the popular items and one for the unpopular items.

5.1 Popularity-Aware Focused Learning (PAL)

Cause. A CRS relies upon its recommender component to rank all the candidate items when making recommendations to a user. This component is typically first trained separately as a static recommendation system then fine-tuned via reinforcement learning in an offline simulation [20, 28, 32]. Similar to a typical collaborative recommendation system, the recommender component learns a set of user, item, and attribute embeddings that will be used to predict a user's interest in a particular item. As shown in [20], one way to calculate the predicted rating is:

$$\hat{y}(u, v, A_u) = \mathbf{u}^\top \mathbf{v} + \sum_{a_i \in A_u} \mathbf{v}^\top \mathbf{a}_i \quad (6)$$

where \mathbf{u} and \mathbf{v} denote the embedding for user u and item v , and \mathbf{a}_i denotes the embedding for a specific attribute in the user's preferred attribute set A_u . From the above equation, we can see that the item embedding \mathbf{v} plays a vital role in the rating prediction, as it is used both to predict the general interest of the user on the target item and the affinity between the target item and the user's preferred attributes. However, most existing CRS neglect the potential issue that popularity bias could bring to item embeddings. As shown in Figure 3, both EAR and CRM show a strong positive correlation between the magnitude of an item's learned embedding – calculated as the squared sum of all features – and its popularity. Especially for EAR, the magnitude of the learned embeddings for popular items is significantly greater than the learned embeddings of the unpopular items. Since an item's popularity is calculated by the frequency that such an item appears in the training set (via past user-item interaction records). Higher popularity is equivalent to more occurrences (or samples) in the training set. Consequently, during the training process of the recommender component, the embeddings for popular items are more frequently updated (e.g., through back-propagation or other learning schemes), resulting in a greater magnitude. Such a phenomenon is problematic since this component ranks all candidate items based upon their predicted ratings. As a result, popular items will have higher predicted ratings




Figure 3: item popularity vs item embedding magnitude

due to the greater magnitude of their embeddings, making them easier to get recommended by the system.

Treatment. To address this issue, we propose a popularity-aware focused learning formulation for the recommender component. This formulation could be easily generalized to any objective function used to optimize the recommender used in a multi-round CRS. For clarity, we show the formulation of our popularity-aware focused learning on the pairwise Bayesian Personalized Ranking (BPR) [25] objective function due to its wide usage in recommendation tasks. Formally the loss function of BPR is defined as:

$$L_{bpr} = \sum_{(u, v, v')} -\ln \sigma(\hat{y}(u, v, A_u)) - \sigma(\hat{y}(u, v', A_u)) + \lambda_\theta \|\theta\|^2, \quad (7)$$

where v and v' respectively denote an item that has been interacted by the user and an item that has not. σ is the sigmoid function and λ_θ is the regularization parameter. $\hat{y}(\cdot)$ could be any function that calculates the predicted rating based on the user embedding \mathbf{u} , item embedding \mathbf{v} , and all attribute embeddings in the user preferred attribute set A_u .

Similar to any loss function, the standard BPR finds a global optimal solution which, as shown in [7], is typically not locally optimal. In the case of CRS, since the popular items dominate the training set, the global optimal solution that BPR finds will be heavily tailored towards them, ignoring the less relevant unpopular items. Consequently, popular items will have a much higher recommendation success rate compared to the unpopular items, introducing severe popularity bias into the system. The goal of the proposed popularity aware focused learning is to help the recommender component learn a more fine-grained set of embeddings for the unpopular items without sacrificing the effectiveness of the learned embeddings for popular items. Formally, our formulation could be generalized to BPR as the following:

$$L_{bpr} = \frac{1}{e^{n_1 p_i}} [-\ln \sigma(\hat{y}(u, v, A_u)) - \sigma(\hat{y}(u, v', A_u))] + e^{n_2 p_i} \|\mathbf{v}\|^2 + \lambda_\theta (\|\theta\|^2) \quad (8)$$

where p_i is the popularity of the item v . $\frac{1}{e^{n_1 p_i}}$ controls how much the sample is weighted in the loss function, and $e^{n_2 p_i}$ controls the scale of the regularization on the learned item embedding. n_1 and n_2 are hyperparameters that control the impact of popularity on an item's weight and regularization scale. Greater n_1 denotes lower weights on the popular items and greater weights on unpopular

items; while n_2 denotes a more significant regularization penalty on the popular items and a more minor penalty on unpopular items. Note that it is sufficient to only apply the popularity-correlated regularization parameter $e^{n_2 p_i}$ to the norm of the item embedding. The recommender component is only responsible for ranking items within a candidate itemset, so the magnitude of the user and attribute embeddings becomes irrelevant since the same user and attribute embeddings are used to calculate the predicted rating. Applying the popularity-correlated regularization parameter also to the norm of the user and the attribute embeddings would bring no additional benefits and might cause the loss function to fail to converge. Adjusting the importance and scale of regularization of an item based on its popularity helps the recommender component to focus also on the unpopular items which would have been poorly-modeled otherwise.

5.2 Cold-Start Item Embedding Reconstruction via Attribute Mapping (CSM)

Cause. Cold-start items refer to items with no past user interactions, which are sometimes excluded from the start in the traditional static recommendation system. However, one of the intentions for a CRS is to directly elicit preferences from users so that even cold-start items could be recommended accurately. Thus it is reasonable to keep cold-start items in the testing set for a more holistic evaluation. Referring back to Equation 6, we know the item embedding \mathbf{v} plays a significant role in rating prediction; however, cold-start items were never included in the training of the recommender component. Their embeddings are often randomly initialized or simply zeros. Therefore, the predicted rating for a cold-start item is rarely an accurate reflection of the user's true interests in the item. Consequently, these cold-start items are rarely successfully recommended or exposed to users, defeating the purpose of a CRS.

Treatment. While popularity-aware focused learning alleviates the popularity bias introduced by the recommender component, the embeddings of cold-start items remain randomly initialized or simply zeros since they are never included in the training of the recommender. However, utilizing a unique property of a CRS, we propose a simple mapping mechanism to reconstruct the embeddings for cold-start items. In the setting of a CRS, every item, regardless of its popularity, is associated with a set of predefined attributes or features that characterizes the item. Based on this property, for the warm start items, we train a simple feed-forward neural network that maps items from their one-hot attribute embeddings to their item embeddings learned in the previous stage. Using this trained mapping, we reconstruct the embeddings for cold-start items by feeding their one-hot attribute embeddings to the trained mapping neural network. Formally, Let V^+ denote the set of warm-start items. Let $emb_{i^+}^{item}$ and $emb_{i^-}^{item}$ denote the item embedding for a warm-start and a cold-start item respectively. Let $emb_{i^+}^{onehot}$ and $emb_{i^-}^{onehot}$ denote the one-hot attribute embedding for a warm-start and a cold-start item respectively. We train a mapping function f_θ with the following loss function:

$$\operatorname{argmin}_\theta \sum_{i \in V^+} (f_\theta(emb_i^{onehot}) - emb_i^{item})^2 + \lambda \|\theta\|^2 \quad (9)$$

where λ is the regularization parameter to prevent overfitting. The item embedding for the cold-start item can then be constructed as:

$$\text{emb}_{i^-}^{\text{item}} = f_{\theta}(\text{emb}_{i^-}^{\text{onehot}}) \quad (10)$$

In this way, instead of having randomly initialized embeddings, cold-start items now have embeddings that effectively encode their unique properties, making their predicted ratings a more realistic reflection of a user's true interest.

5.3 Dual-Policy Learning (DPL)

Cause. A CRS relies upon a policy agent to decide which action to take at each interaction with the user. The action space typically includes whether to elicit the user's preference on any of the unasked attributes or to make a recommendation if the agent feels enough information has been collected. However, without careful modification, the policy agent brings a new form of popularity bias into the system. In CRS, policy agents are typically first pretrained with a max entropy attribute selecting strategy, with the goal of minimizing the number of needed interactions to reach a successful recommendation. Under the max entropy strategy, the policy agent is trained to pick the attribute that reduces the candidate set the most, skipping attributes with a high disparity in their value distribution. For instance, if most of the movies in a *movie* dataset have value *Action* for the attribute *genre*, then the entropy of the attribute *genre* would be quite low since knowing the value of it does not help the system gain as much information as asking about an attribute with lower disparity in its value distribution. By skipping the attribute *genre*, the policy agent makes the false assumption that users desire only the popular *Action* movies when in fact the user is actually seeking a *Sci-Fi* movie.

Besides making false assumptions on specific user preferences, the training of the policy agent is also affected by the popularity bias introduced by the recommender component. Since the recommender component learns relatively uninformative embeddings for unpopular items, the policy agent learns to elicit more user preferences before making a recommendation. In addition, the lower magnitude of their learned embeddings makes them rank lower in the candidate set, and thus they are less likely to be recommended to the user. Thus, for unpopular items, the agent either exceeds the number of allowed interactions before making a recommendation or could never successfully recommend the target items.

Treatment. As discussed above, affected by the max entropy attribute selection strategy used in pretraining, the policy network tends to make false assumptions on specific user preferences by skipping attributes with a low disparity in their value distributions. To tackle this issue, we propose a dual policy learning scheme. First, we split the entire dataset into two smaller datasets – D_1 and D_2 such that D_1 contains only the items with popularity greater than the 80th percentile, and D_2 contains only the items with popularity lesser than the 20th percentile. Then, we train two policy networks, PN_1 and PN_2 , on D_1 and D_2 respectively. When interacting with the user, we select PN_1 if the item's popularity is greater than the 25th percentile and PN_2 if otherwise. Both policy networks are first

Table 1: Dataset statistics

Dataset	users	items	interactions	attributes
Yelp	27,675	70,311	1,368,606	590
Lastfm	1,801	7,432	76,693	33

pretrained as classifiers to avoid optimization failures, then fine-tuned with the standard policy gradient method as the following:

$$\theta \leftarrow \theta - \alpha \nabla \log \pi_{\theta}(a^n | s^n) R_n$$

where θ denotes the parameter of the policy network, α denotes the learning rate of the policy network and R_n represents the total accumulated reward from the n^{th} turn to the final turn. Inspired by [11], we define R_n to be:

$$R_n = \sum_{n=1}^N \gamma^n (w_{rec} R_n^{rec} + w_{conv} R_n^{conv} + w_{bias} R_n^{bias}),$$

where γ denotes the discounted factor, R_n^{rec} , R_n^{conv} , and R_n^{bias} respectively denotes the recommendation success state, the user experience of the conversation, and the degree of popularity bias defined in Equation 2 at the n -th of the interaction. And w_{rec} , w_{conv} , and w_{bias} denotes the weight for R_n^{rec} , R_n^{conv} , and R_n^{bias} respectively.

The entire framework functions as an entity. The embeddings, learned in the PAL stage, are used in the CSM stage to reconstruct embeddings for all the cold-start items. Only with embeddings of all items properly learned, can the framework train a dual-policy agent that decides the next most appropriate action to take. To achieve the best performance, the framework must be executed in the order presented above since each stage has a significant impact on the following.

6 EXPERIMENTS

To validate the proposed debiasing framework, we showcase experiments over real-world datasets designed to answer three key research questions: **RQ1**. Does the proposed framework effectively mitigate popularity bias in existing conversational recommendation methods? **RQ2**. How does the proposed framework impact the overall recommendation performance of existing conversational recommendation methods? **RQ3**. How does each stage contribute to the mitigation of popularity bias and impact the overall recommendation performance?

6.1 Experiments Setup

Dataset. We conduct experiments on two datasets widely adopted to evaluate CRSs [20, 21, 28, 33] – **Yelp**¹ business recommendation dataset and **Lastfm**² music artist recommendation dataset. Following previous works [15, 25], we only keep users with at least 10 reviews to alleviate data sparsity. The user-item interactions are split in the ratio of 7:2:1 for training, validation, and testing. For the Yelp dataset, we perform the same hierarchical attribute prepossessing as in previous works [20, 21, 28]. The two datasets are summarized in Table 1.

¹<https://www.yelp.com/dataset/>

²<https://grouplens.org/datasets/hetrec-2011/>

Implementation Details. For PAL, we perform grid search to find the best values for the hyperparameters n_1 and n_2 . Specifically, we set n_1 to 7 and n_2 to 8. For CSM, we use a two layer neural network with parameter sizes of $|\text{emb}^{\text{onehot}}| \times 128$ and $128 \times |\text{emb}^{\text{item}}|$. For DAL, each policy network is modeled as a two-layer neural network with parameter sizes of $|s_t| \times 64$ and $64 \times |A|$, where s_t denotes the state vector (see Section 2) and A denotes the action space. we use the REINFORCE algorithm to train [31] the two policy networks. For calculating the rewards, we set R^{rec} to 1 if the user accepts the recommendation or to -1 if the user quits due to impatience or the allowed conversational turns have been reached. R^{conv} is set to 0.1 if the user replies to the prompted attribute and to -0.1 if otherwise. R^{bias} is calculated via Equation 2, and the discount factor γ is set to 0.7. We report the best results of 15 conversational rounds on the Lastfm dataset and 5 conversational rounds on the Yelp dataset.

Baselines. To examine the proposed framework, we investigate its effectiveness on the following state-of-the-art baseline conversational recommendation approaches.

- **Max Entropy (MaxEnt):** MaxEnt is a rule-based attribute selection strategy. At each turn, the policy agent computes the entropy for each unknown attribute and selects the attribute with the highest entropy to be the next one to ask. Recommendation is made either when the candidate space is small enough, or the policy agent runs out of attributes to ask.
- **CRM [28]:** CRM is a CRS that uses a belief tracker to record a user’s preference conveyed during the conversation, and trains a policy network via reinforcement learning to decide how to interact with the user. The policy network takes the output of the belief tracker and decides the most appropriate subsequent action. We follow [20, 21] to adapt it to the multi-round conversational setting to make fair comparisons.
- **EAR [20]:** Similar to CRM, EAR also learns a predictive model to estimate a user’s preference and trains a policy network to determine whether to ask more attributes or make recommendations; however, different from CRM, EAR also considers the feedback from the user to further fine-tune the learned predictive model, achieving better recommendation performance.
- **SCPR [21]:** SCPR is the state-of-the-art multi-round CRS. Extending from EAR, SCPR leverages the concept of adjacent attributes to reduce the search space of attributes and builds a knowledge graph to learn a more efficient policy network.

Besides applying our proposed debiasing framework on the aforementioned CRSs, we also compare it with **Popcorn** [11], which focuses on attribute-based popularity bias in CRSs. We adapt this approach to the conventional interaction frequency-based popularity definition for a fair comparison. Note that **Popcorn** is *not a model-agnostic framework* and we report its performance following the workflow in the original paper.

Metrics. To measure the *degree of popularity bias* in the CRS, we use *PER*, *PSR*, and *PCU* introduced in Section 4. In addition, since this work intends to design a framework that mitigates the popularity bias while preserving the overall recommendation performance, we also include two widely adopted [20, 21, 28, 33] *recommendation performance* metrics: success rate (*SR@t*) and average turns (*AT*).

Table 2: Overall debiasing effectiveness on two benchmark datasets. Popularity correlation with exposure rate (PER↓), Popularity correlation with success rate (PSR↓), and Popularity correlation with conversational utility (PCU↓) are used as evaluation metrics. The best performance is in bold-face.

	Lastfm			Yelp		
	PER	PSR	PCU	PER	PSR	PCU
Popcorn	3.19	.448	.146	4.26	.195	.255
MaxEnt	3.99	.503	.172	2.83	.110	.132
w/ Debias	.192	.417	.077	2.29	.058	.029
CRM	4.16	.474	.148	4.81	.253	.347
w/ Debias	1.24	.331	.072	3.84	.103	.163
EAR	2.71	.552	.224	4.99	.269	.402
w/ Debias	1.11	.318	.119	3.71	.141	.170
SCPR	5.79	.440	.281	5.97	.281	.440
w/ Debias	3.61	.322	.142	4.39	.179	.257

Success rate measures if the target item can be successfully recommended within the allowed turns of interactions while average turns measures the number of needed interactions to successfully recommend an item. Additionally, to measure the ratio of successful recommendation for popular and unpopular items respectively, we also include popular (head) item success rate (*HSR*) and unpopular (tail) item success rate (*TSR*).

6.2 Mitigation of Popularity Bias (RQ1)

In this section, we evaluate the debiasing effectiveness of the proposed framework on each baseline CRS model with different popularity bias metrics introduced in Section 4. As shown in Table 2, by applying the debiasing framework on different baseline CRS models, it notably reduces the degree of popularity bias quantified by all metrics. Meanwhile, the proposed framework also achieves lower values for PSR and PCU compared to Popcorn on all baseline models, which further demonstrates its effectiveness in mitigating the popularity bias in CRSs. On Lastfm, the framework exhibits strong debiasing performance on all baseline models. For MaxEnt, we observe a lower decrement in PSR. This is because MaxEnt relies on a rule-based policy agent to select attributes to ask, which vastly limits the effectiveness of the PAL stage. However, since MaxEnt is still affected by the issues of biased recommender and cold-start negligence, applying the CSM stage still dramatically decreases both PER and PCU. Compared to MaxEnt, the proposed framework demonstrates an even stronger bias mitigation performance on the RL-based CRS approaches – CRM, EAR, and SCPR. For CRM, we observe a particularly sharp decrease in PER, which indicates unpopular items now have higher chances of being recommended (exposed) to users. For EAR, our framework decreases the PCU by over 68% which significantly lowers the frequency of unpopular items getting unsuccessfully recommended due to lengthy conversations. And for SCPR, we observe a 49% decrease in PCU and a 38% decrease in PER. In addition, our framework decreases the PSR across EAR, CRM, and SCPR by 30%, 42% and 27% respectively, significantly weakening the correlation between an item’s popularity

Table 3: Overall recommendation performance on two benchmark datasets. Success rate (SR↑), Popular (Head) item success rate (HSR↑), Unpopular (Tail) item success rate (TSR↑), and Average Turns (AT↓) are used as evaluation metrics.

	Lastfm				Yelp			
	SR	HSR	TSR	AT	SR	HSR	TSR	AT
Popcorn	.396	.710	.104	12.9	.637	.711	.358	3.56
MaxEnt	.286	.541	.068	13.6	.507	.530	.424	3.97
w/ Debias	.311	.532	.102	13.5	.509	.522	.459	3.99
CRM	.324	.598	.089	13.3	.604	.703	.228	3.59
w/ Debias	.438	.645	.162	12.3	.633	.694	.402	3.62
EAR	.421	.742	.061	12.5	.629	.719	.295	3.61
w/ Debias	.479	.711	.187	12.1	.647	.708	.422	3.64
SCPR	.457	.793	.109	12.5	.631	.715	.301	3.67
w/ Debias	.546	.839	.216	11.5	.667	.717	.501	3.75

and its recommendation success rate. It is important to note that all baseline models produce a lesser degree of popularity bias on the Yelp dataset. One key reason is that the Yelp dataset adopts an enumerated question setting in which a user can provide values for multiple attributes at each conversational turn. Compared to the Lastfm dataset, which adopts a binary question setting, the enumerated question setting facilitates user preference elicitation. As a result, all items, in general, have higher chances for being successfully recommended, weakening the correlation between an item’s popularity and its recommendation success rate, exposure rate, and conversational utility. Even so, our framework still exhibits strong bias mitigation performance on the Yelp dataset. In particular, the PSR value for CRM decreases by 59% and the PCU value for MaxEnt decreases by more than 28%.

6.3 Recommendation Performance (RQ2)

Since this work intends to mitigate the undesirable effects of popularity bias while preserving the recommendation performance of the CRS, we also report the overall recommendation performance on all three baseline CRSs after applying our framework. In addition to the Recommendation Success Rate (SR) and Average Turns (AT), we also report the success rate for popular (head) items and the unpopular (tail) items respectively (HSR and TSR). Following Table 2, we include the overall recommendation performance of Popcorn [11] as our baseline and use boldface to denote statistical significance of $p < 0.01$. As shown in Table 3, our framework significantly increases the overall recommendation performance of all baseline models by greatly preserving the recommendation success rate for the popular (head) items while boosting the recommendation success rate for the unpopular (tail) items. In general, we observe a more minor improvement on the Yelp dataset. This is because all baseline models already achieve high recommendation performance due to the enumerated user preference elicitation process. On the contrary, we observe a higher performance boost across all the baseline models on the Lastfm dataset. In particular, both the HSR and TSR of CRM increase after applying the framework. In addition to the DPL stage, the increase in overall recommendation performance is primarily due to the more effective item embeddings




Figure 4: The performance of Item Prediction Before and After Applying the Debiasing Framework.

learned in the PAL stage (for warm-start items) and re-constructed in the CSM stage (for cold-start items). To compare the effectiveness of the item embeddings learned with and without the proposed debasing framework, we calculate their AUC scores on item prediction as in [25]. As shown in Figure 4, the AUC scores of unpopular items significantly increase across three baseline models on both the the Lastfm and the Yelp dataset. This is because many of the unpopular items are either cold-start or have extremely low occurrences in the training set. Re-constructing their embeddings in the CSM stage increases their effectiveness which in turn increases their recommendation success rate. In addition, the AUC scores of popular items also increase on all baseline models, especially for CRM. This is because even amongst the popular items, there exists a large variation in an item’s popularity. By strategically re-weighting an item’s relevance based on its popularity, the PAL stage refrains the recommender from focusing on optimizing the embeddings for the few extremely popular items, thus learning an overall more effective set of item embeddings. Note SPCR shares the same recommender component as EAR, making their learned item embeddings the same; thus we report its AUC score with EAR.

6.4 Ablation Studies (RQ3)

In this section, we investigate how each stage in the proposed framed work contributes to the mitigation of popularity bias and the improvement of recommendation performance via the ablation studies conducted on both EAR and SPCR with the Lastfm dataset. Note we also perform ablation study on EAR to serve as a comparison between knowledge-graph based CRS (SPCR) and non-knowledge-graph based CRS (EAR). For EAR, we find that while the CSM stage contributes the most to lowering both the PSR and PCU, it also produces a strong trade-off between the recommendation performance and the debiasing effectiveness. As shown in Table 5, skipping the CSM stage gives a significant recommendation performance boost to the system but also limits the overall bias mitigation effectiveness. Such trade-off happens because since the CSM reconstructs the embeddings for all cold-start items, the cold-start items will also predict ratings that reflect a user’s true interests on them, making them rank significantly higher within the candidate itemset. Therefore, when the CRS is trying to recommend a popular target item, all the cold-start items with attributes matching the collected user preferences will now be deemed as equally qualified as the target item. Consequently, the number of conversational

Table 4: Debiasing effectiveness and Recommendation Performance of skipping (-) one stage of the proposed debiasing framework on EAR.

	PER	PSR	PCU	SR	HSR	TSR	AT
EAR w/ Debias	1.11	.318	.119	.479	.711	.187	12.1
- PAL	1.31	.325	.166	.445	.689	.171	12.4
- CSM	1.21	.439	.204	.516	.814	.110	11.9
- DPL	1.23	.338	.145	.459	.709	.166	12.3
EAR	2.71	.552	.224	.421	.742	.061	12.5

Table 5: Debiasing effectiveness and Recommendation Performance of skipping (-) one stage of the proposed debiasing framework on SCPR.

	PER	PSR	PCU	SR	HSR	TSR	AT
SCPR w/ Debias	3.61	.322	.142	.546	.839	.216	11.5
- PAL	5.05	.373	.205	.458	.761	.160	12.3
- CSM	4.43	.469	.275	.495	.831	.097	11.9
- DPL	4.19	.422	.257	.484	.805	.156	11.7
SCPR	5.79	.440	.281	.457	.793	.109	12.5

turns needed to recommend the target item increases, and some popular items might even fail to get successfully recommended. Interestingly, skipping the DPL stage also results in a notable increase in PCU which backs up our intuition of the policy agent prioritizing popular items during the attribute selection process. Compared to EAR, skipping the DPL stage led to greater recommendation performance drop in SCPR. This is because SCPR deploys a knowledge graph that significantly reduces the decision space of the policy network, which makes our DPL stage more effective. Interestingly for SCPR, skipping the CSM stage did not lead to better recommendation performance. This is because the DPL stage functions better on SCPR such that re-constructing the embeddings for the cold-start items have less impact on the recommendation performance of the popular items.

7 RELATED WORK

Conversational Recommender Systems. Traditional static recommender systems that are trained using offline historical user-item interactions face two inherent challenges: (1) the inability to capture users' precise preferences; and (2) failure to provide a human-interpretable justification for their recommendations [13]. Although many existing works have attempted to solve these problems, most rely on a large amount of auxiliary data to better interpret user intentions. The emergence of CRSs provides an intrinsic solution to these problems. By dynamically interacting with users through real-time interactions (e.g., conversations, form fields, buttons and even gestures [17]), CRSs are able to elicit current and detailed preferences directly from users, learning precisely what the users' interests are and thus making highly personalized recommendations that are justifiable. While early works on CRSs primarily resort to choice-based questions for collecting user preferences [14, 18],

recent works have explored the possibilities of more effective preference elicitation methods and conversational strategies with the help of reinforcement learning. For example, [28] first proposes a unified deep RL framework that builds a personalized preference elicitation process by optimizing a session-based utility function. Inspired by [28], research in [20] further strengthens the interaction between the recommender component and the conversation component by proposing a three-stage framework that inherently adapts to a multi-round setting. Extending upon [20], work proposed in [21] first integrates a knowledge graph to improve the reasoning ability of the system and reduce the learning difficulty of the policy agent. In addition, other research directions in CRSs include dialogue understanding and generation [35, 39], response generation [23, 36], and the exploration-exploitation trade-offs [22, 34].

Popularity Bias in Recommender Systems. Many existing works have studied the impact of popularity bias in traditional static recommender systems. [27] first examines the trade-off between item popularity and recommendation accuracy. [24] introduces the importance of long tail items in promoting user satisfaction and preventing monopoly by big brands. More comprehensively, [16] empirically demonstrated that different recommendation algorithms have different vulnerabilities to popularity bias. To mitigate the harmful effects of popularity bias, a variety of debiasing approaches have been proposed. [30] studies the popularity bias in recommender systems from a cause-effect perspective and proposes a model-agnostic counterfactual debiasing method that amends the learning process of recommendation. And [29] proposes to augment learning for casual users. However, the above introduced works focus on discovering and mitigating undesirable effects of item popularity in traditional recommender systems and cannot be directly applied to CRSs. Recently, [12] introduces the notion of human-in-the-loop (HitL) reasoning in CRS and [11] first investigates the issue of HitL bias in CRSs. Specifically, [11] defines a metric that quantifies the degree of HitL popularity bias, and modifies the reward function to dynamically mitigate it during the preference elicitation process. However, it adopts an attribute-based definition of popularity which is less generalizable and consistent than the conventional interaction-frequency-based definition. To this end, our work presents a more comprehensive study on popularity bias in CRSs and proposes a generalizable framework that interactively mitigates the harmful effects of popularity bias in the entire system.

8 CONCLUSION AND FUTURE WORK

In this work, we present the first systematic study on popularity bias in conversational recommender systems. We illustrate the existence of popularity bias in the setting of CRSs from three perspectives. Building upon this analysis, we propose three metrics for quantifying the degree of popularity bias in CRSs and a three-staged debiasing framework that can be easily applied to any CRS. The experimental results on two frequently adopted conversational recommendation datasets show that this framework not only mitigates the undesirable effects brought by popularity bias in CRSs but also improves the overall recommendation performance. In the future, we will explore more effective mapping schemes for cold-start items.

9 ACKNOWLEDGEMENTS

This work is supported in part by NSF grant IIS-1939716 and an Amazon Research Award.

REFERENCES

- [1] Himan Abdollahpouri. 2019. Popularity bias in ranking and recommendation. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 529–530.
- [2] Himan Abdollahpouri and Robin Burke. 2019. Reducing Popularity Bias in Recommendation Over Time. *CoRR* abs/1906.11711 (2019). arXiv:1906.11711 <http://arxiv.org/abs/1906.11711>
- [3] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the eleventh ACM conference on recommender systems*. 42–46.
- [4] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2019. Managing popularity bias in recommender systems with personalized re-ranking. In *The thirty-second international flairs conference*.
- [5] Himan Abdollahpouri, Masoud Mansouri, Robin Burke, and Bamshad Mobasher. 2019. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286* (2019).
- [6] Himan Abdollahpouri, Masoud Mansouri, Robin Burke, Bamshad Mobasher, and Edward Malthouse. 2021. User-Centered Evaluation of Popularity Bias in Recommender Systems. In *UMAP*.
- [7] Alex Beutel, Ed H. Chi, Zhiyuan Cheng, Hubert Pham, and John Anderson. 2017. Beyond Globally Optimal: Focused Learning for Improved Recommendations. In *The WebConf*.
- [8] Ludovico Boratto, Gianni Fenu, and Mirko Marras. 2021. Connecting user and item perspectives in popularity debiasing for collaborative recommendation. *Information Processing & Management* 58, 1 (2021), 102387.
- [9] M. Craig Brown. 1994. Using Gini-style indices to evaluate the spatial patterns of health practitioners: theoretical considerations and an application based on Alberta data. *Social science & medicine* (1994).
- [10] Erik Brynjolfsson, Yu Jeffrey Hu, and Michael D Smith. 2006. From niches to riches: Anatomy of the long tail. *Sloan management review* 47, 4 (2006), 67–71.
- [11] Zuohui Fu, Yikun Xian, Shijie Geng, Gerard de Melo, and Yongfeng Zhang. 2021. Popcorn: Human-in-the-Loop Popularity Debiasing in Conversational Recommender Systems. In *CIKM*.
- [12] Zuohui Fu, Yikun Xian, Yaxin Zhu, Shuyuan Xu, Zelong Li, Gerard De Melo, and Yongfeng Zhang. 2021. HOOPS: Human-in-the-Loop Graph Reasoning for Conversational Recommendation. In *SIGIR*.
- [13] Chongming Gao, Wenqiang Lei, Xiangnan He, M. de Rijke, and Tat-Seng Chua. 2021. Advances and Challenges in Conversational Recommender Systems: A Survey. *ArXiv* abs/2101.09459 (2021).
- [14] Mark P. Graus and Martijn C. Willemse. 2015. Improving the User Experience during Cold Start through Choice-Based Preference Elicitation. In *RecSys*.
- [15] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [16] Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. 2015. What Recommenders Recommend: An Analysis of Recommendation Biases and Possible Countermeasures. *User Modeling and User-Adapted Interaction* (2015).
- [17] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A survey on conversational recommender systems.
- [18] Hai Jiang, Xin Qi, and He Sun. 2014. Choice-based recommender systems: a unified approach to achieving relevancy and diversity. In *Operations Research*.
- [19] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2014. Correcting Popularity Bias by Enhancing Recommendation Neutrality. In *RecSys Posters*.
- [20] Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-action-reflection: Towards deep interaction between conversational and recommender systems. In *WSDM*.
- [21] Wenqiang Lei, Gangyi Zhang, Xiangnan He, Yisong Miao, Xiang Wang, Liang Chen, and Tat-Seng Chua. 2020. Interactive path reasoning on graph for conversational recommendation. In *KDD*.
- [22] Shijun Li, Wenqiang Lei, Qingyun Wu, Xiangnan He, Peng Jiang, and Tat-Seng Chua. 2021. Seamlessly unifying attributes and items: Conversational recommendation for cold-start users. In *TOIS*.
- [23] Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. In *ACL*.
- [24] Yoon-Joo Park and Alexander Tuzhilin. 2008. The long tail of recommender systems and how to leverage it. In *RecSys*.
- [25] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [26] Guy Shani and Asela Gunawardana. 2011. Evaluating Recommendation Systems. *Recommender Systems Handbook*.
- [27] Harald Steck. 2011. Item popularity and recommendation accuracy. In *Proceedings of the fifth ACM conference on Recommender systems*. 125–132.
- [28] Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *SIGIR*.
- [29] Jianling Wang, Ya Le, Bo Chang, Yuyan Wang, Ed H Chi, and Minmin Chen. 2022. Learning to Augment for Casual User Recommendation. In *TheWebConf*.
- [30] Tianxin Wei, Fuli Feng, Jiawei Chen, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2021. Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system. In *KDD*.
- [31] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3 (1992), 229–256.
- [32] Kerui Xu, Jingxuan Yang, Jun Xu, Sheng Gao, Jun Guo, and Ji-Rong Wen. 2021. *Adapting User Preference to Online Feedback in Multi-Round Conversational Recommendation*.
- [33] Kerui Xu, Jingxuan Yang, Jun Xu, Sheng Gao, Jun Guo, and Ji-Rong Wen. 2021. Adapting user preference to online feedback in multi-round conversational recommendation. In *WSDM*.
- [34] Xiaoying Zhang, Hong Xie, Hang Li, and John C.S. Lui. 2020. *Conversational Contextual Bandit: Algorithm and Application*.
- [35] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. Towards Conversational Search and Recommendation: System Ask, User Respond. In *CIKM*.
- [36] Kun Zhou, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen, and Jingsong Yu. 2020. Improving Conversational Recommender Systems via Knowledge Graph Based Semantic Fusion. In *KDD*.
- [37] Ziwei Zhu, Yun He, Xing Zhao, and James Caverlee. 2021. Popularity Bias in Dynamic Recommendation. In *KDD*.
- [38] Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang, and James Caverlee. 2021. Popularity-Opportunity Bias in Collaborative Filtering. In *WSDM*.
- [39] Jie Zou, Yifan Chen, and Evangelos Kanoulas. 2020. Towards Question-Based Recommender Systems. In *SIGIR*.