

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

12-2022

Mitigating popularity bias in recommendation with unbalanced interactions: A gradient perspective

Weijieying REN

University of Central Florida

Lei WANG

Singapore Management University, lei.wang.2019@phdcs.smu.edu.sg

Kunpeng LIU

Portland State University

Ruocheng GUO

ByteDance AI Lab

Ee-peng LIM

Singapore Management University, eplim@smu.edu.sg

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#), and the [Data Storage Systems Commons](#)

Citation

REN, Weijieying; WANG, Lei; LIU, Kunpeng; GUO, Ruocheng; LIM, Ee-peng; and FU, Yanjie. Mitigating popularity bias in recommendation with unbalanced interactions: A gradient perspective. (2022). *2022 IEEE International Conference on Data Mining (ICDM): Orlando, FL, November 28 - December 1: Proceedings*. 438-447.

Available at: https://ink.library.smu.edu.sg/sis_research/7510

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Author

Weijieying REN, Lei WANG, Kunpeng LIU, Ruocheng GUO, Ee-peng LIM, and Yanjie FU

Mitigating Popularity Bias in Recommendation with Unbalanced Interactions: A Gradient Perspective

Weijieying Ren^{*‡}, Lei Wang^{*^b}, Kunpeng Liu[#], Ruocheng Guo[§], LIM Ee Peng^{b^l} and Yanjie Fu^{#^l}

[‡]Department of Computer Science, University of Central Florida, USA

^bDepartment of Computer Science, Singapore Management University, Singapore

[#]Department of Computer Science, Portland State University, USA

[§]ByteDance AI lab, UK

wjyren@knights.ucf.edu, lei.wang.2019@phdcs.smu.edu.sg, kunpeng@pdx.edu

ruocheng.guo@bytedance.com, eplim@smu.edu.sg, yanjie.fu@ucf.edu

Abstract—Recommender systems learn from historical user-item interactions to identify preferred items for target users. These observed interactions are usually unbalanced following a long-tailed distribution. Such long-tailed data lead to popularity bias to recommend popular but not personalized items to users. We present a gradient perspective to understand two negative impacts of popularity bias in recommendation model optimization: (i) the gradient direction of popular item embeddings is closer to that of positive interactions, and (ii) the magnitude of positive gradient for popular items are much greater than that of unpopular items. To address these issues, we propose a simple yet efficient framework to mitigate popularity bias from a gradient perspective. Specifically, we first normalize each user embedding and record accumulated gradients of users and items via popularity bias measures in model training. To address the popularity bias issues, we develop a gradient-based embedding adjustment approach used in model testing. This strategy is generic, model-agnostic, and can be seamlessly integrated into most existing recommender systems. Our extensive experiments on two classic recommendation models and four real-world datasets demonstrate the effectiveness of our method over state-of-the-art debiasing baselines.

Index Terms—Recommendation system, Popularity bias.

I. INTRODUCTION

Recommender systems (RS) provide users with newly suggested items by mining user preferences from historical user-item interactions [1]–[4]. Each observed user-item interaction represents an implicit feedback (e.g., purchases, clicks, browses) indicating the user’s preference on the item. We call this a *positive* interaction. To train a RS model, we can uniformly select a few items having no interaction with the user as *negative* items. A RS model based on such implicit feedback data is designed to predict positive items with higher prediction scores than negative items. Generally, each item will receive a *positive gradient* if it has observed interactions with a user and obtain a *negative gradient* if it is selected as a negative sample.

Nevertheless, recent studies [1], [4], [5] found that in most real-world application scenarios, user-item interactions




Figure 1. (a) Directional similarity between accumulated positive gradient (negative gradient) and combined gradient of item embeddings shown as orange dots (green dots) for the Movielens-10M dataset. We define directional similarity by cosine similarity function. Note that items are sorted in inverse popularity order, i.e., more popular items are assigned smaller index on the x -axis. (b) Magnitude (L2 norm) of accumulated positive gradient (negative gradient) of item embeddings shown as orange dots (green dots) for items sorted in inverse popularity order.

in recommender systems are highly imbalanced following a long-tailed distribution [6], [7]. Only a small number of items are popular as they frequently interact with users and thus are more likely favored by users. When such long-tailed data are used for training a RS model, *popularity bias* arises in the trained model, i.e, popular items are overly recommended to all users, while some less popular items may be more relevant to some specific users. Popularity bias can result in low visibility of unpopular items, reduced diversity of markets and damage consumer experience. Mitigating popularity bias is therefore highly demanded.

To deal with popularity bias, one prominent direction is to design a re-weighting loss, e.g., inverse propensity score (IPS), to balance overall item distributions [8], [9]. Another alternative is to conduct knowledge transfer [6] with the assumption that knowledge represented in popular items should be beneficial to the user preference mining in less popular items. Recently, more sophisticated models are designed to preserve user specific tastes via the usage of side information [10], autoencoder model [11] and to disentangle user preference following pre-defined prior heuristics [12]. However, previous works either require complex model design [11], degenerate embedding representation or rely on strong assumption [12],

* denotes equal contribution.

^l denotes corresponding author.




Figure 2. For a popular item i , let \overrightarrow{AC} be its accumulated positive gradient and \overrightarrow{AB} denotes the accumulated negative gradient. The overall gradient (learned vector) is thus $\overrightarrow{AD} = \overrightarrow{AB} + \overrightarrow{AC}$ which is still very similar to \overrightarrow{AC} . Our gradient based embedding adjustment approach is to reduce \overrightarrow{AC} to $\overrightarrow{AC'}$ leading to a new overall gradient $\overrightarrow{AD'}$ which effectively incorporates the negative gradient.

which is lack of generalization ability. Unlike the aforementioned strategies, we approach the popularity bias by analyzing gradient distortion in long-tailed data, attributed to the combination of gradients generated by positive and negative items. We discover that (i) the overall gradient direction of popular items’ embeddings tend to be closer to the gradient of positive interactions, and (ii) the magnitude of positive gradient for popular items is far larger than that of unpopular items.

To illustrate the above challenges, we visualize the related variables in the training process of Movielens-10M dataset in Fig.1. Specifically, the cosine similarity between the accumulated positive gradient (given by $\sum_{t=1}^T \nabla_i \mathcal{L}^t(u_+, i_+)$, where t is the iteration step) and overall gradient of each item is visualized in Fig. 1(a) for items sorted by inverse popularity order. We also show the cosine similarity between the negative gradient and overall gradient of items. The high similarity (i.e., close to 1) between the overall gradient and the accumulated positive gradient of popular items indicates that the overall gradient of each popular item is shifted towards the direction of its positive interactions. The opposite can be observed for unpopular items. Fig. 1(b) shows the norm of accumulated positive and negative gradients for items sorted in the descending popularity order. In this figure, popular items exhibit large accumulated positive gradient norm values than their accumulated negative gradient norm values. In other words, the accumulated gradients of popular items increase the norm of their corresponding item embedding vectors and results in tendency to rank high in recommendation results.

To further analyze the underlying reasons of the above empirical findings, we review the optimization process of the classical pairwise ranking loss used in implicit RS, i.e., BPR loss [13]. It is clear that item (or user) embedding vectors are mainly determined by their accumulated gradients. Considering a popular item i illustrated in Fig. 2, vectors \overrightarrow{AC} and \overrightarrow{AB} represent its accumulated positive and negative gradient respectively. The overall gradients \overrightarrow{AD} is shifted towards \overrightarrow{AC} due to the latter’s large magnitude, making the learned item embedding only facilitate the learning of positive interaction (vice versa for unpopular items). The accumulated gradients are derived by summing up gradient values across all iterations when optimizing the objective function $\mathcal{L}^t(u, i)$. Consequently, We identify three factors which may account for the reason of the existence of popularity bias in implicit RS: 1) item popularity, 2) user activity, and 3) user conformity.

Our Approach. Motivated by mitigating popularity bias, we found that the combined analysis for gradients generated by positive and negative interactions are key solutions. Hence, we propose the concept of user/item embedding adjustment. In the training stage, we only normalize user embeddings to mitigate the influence of high activity users and record the accumulated gradients of users and items. In the testing stage, we propose to mitigate bias by adjusting the learned user and item embeddings to mitigate the influence by item popularity and user conformity. Our main contributions are as follows:

- **A Novel Gradient Perspective to Understand Popularity Bias Mechanism.** Classical methods focus on understanding popularity bias from balancing the overall data distribution, e.g., popular and unpopular groups. In this paper, we provide new findings that 1) the overall gradient direction of popular (unpopular) items is shifted to be closer to the gradient on positive (negative) interactions; 2) positive gradient magnitude of popular items is usually much larger than that of unpopular items. We analyze such phenomenon from the gradient perspective, and provide a new insight on balancing the positive and negative gradients for the same item.
- **Gradient Based Embedding Adjustment Approach.** Motivated by the new findings, we identify three essential factors, namely item popularity, user conformity, and user activeness and analyze how user and item embeddings are related to them during the optimization steps. We develop a simple yet novel embedding post-doc adjustment method. This framework is explainable and interpretable.
- **Extensive Experiments with Real-world Datasets.** We conduct extensive experiments to evaluate the effectiveness of our proposed methods in terms of overall accuracy, generalisation, and efficiency against state-of-the-art baselines. We also conducted ablation studies to evaluate the components of our proposed methods.

II. PRELIMINARIES

Problem Statement. Suppose we have a user set \mathcal{U} , an item set \mathcal{I} and an user-item interaction matrix $Y \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{I}|}$, where Y_{ui} indicates whether $u \in \mathcal{U}$ has made an interaction (e.g., adoption or purchase) with item $i \in \mathcal{I}$ or not. The goal of a recommender is to learn a scoring function $r(u, i)$ to recommend a list of k items to a user u :

$$\hat{y}_{ui} = r(u, i) = \mathbf{P}_u^T \mathbf{Q}_i \quad (1)$$

where \hat{y}_{ui} is the predicted score, \mathbf{P}_u is the user representation vector for the user u , and \mathbf{Q}_i is the item representation vector for the item i . \mathbf{P}_u (\mathbf{Q}_i) is derived from the user encoder $f(u|\Theta_U)$ (the item encoder $f(i|\Theta_I)$). Note that the user or item encoder is agnostic to model, which can be the embedding matrix [14], graph neural network [15], or other models. For simplicity, we refer to \mathbf{P}_u (\mathbf{Q}_i) as user embedding (item embedding) in the rest of the paper.

Why does the embedding magnitude matter in RS problems? In the testing stage, the recommender system returns the personalized ranking of items R_u for a user u :

$$\begin{aligned} R_u &= \{(i, \text{rank}(\hat{y}_{ui}))\}_{i \in I} \\ &= \{(i, \text{rank}(\mathbf{P}_u^T \mathbf{Q}_i))\}_{i \in I} \\ &= \{(i, \text{rank}(\|\mathbf{Q}_i\| \cos(\mathbf{P}_u, \mathbf{Q}_i)))\}_{i \in I} \end{aligned} \quad (2)$$

where rank is a ranking function based on a set of scoring function values between u and all items $i \in I$. In Eq. 2, we omit $\|\mathbf{P}_u\|$ as R_u involves the same user u . In other words, the personalized ranking result during the testing stage only depends on the magnitude of item embedding \mathbf{Q}_i and its cosine similarity with the user embedding \mathbf{P}_u .

How are the embedding magnitudes generated in RS? In the training stage, we often adopt a pairwise ranking loss, e.g., Bayesian Personalized Ranking (BPR) loss [13] and Binary cross entropy (BCE) [16] loss. Without loss of generality, we revisit BPR loss functions for recommender system with implicit feedback. In realistic scenarios, only positive implicit feedback (e.g., clicks and purchases) are observed. In implicit RS, we assume that user u 's preference over an item i with an observed (u, i) interaction is ranked higher than the unobserved (u, j) interaction as a negative sample, where the item j is uniformly selected from item set I that does not have observed interactions with u . Without loss of generality, we represent BPR loss here to maximize the likelihood of observing such pairwise ranking relations:

$$L_{BPR} = \sum_{(u, i, j) \in D_s} \log(\sigma(\hat{y}_{ui} - \hat{y}_{uj})) - \lambda_\Theta \|\Theta\|^2 \quad (3)$$

where D_s is a set of (u, i, j) triplets with (u, i) and (u, j) as positive and negative samples respectively. Θ is the model parameters and $\sigma(x) = \frac{1}{1 + \exp(-x)}$ is the sigmoid function. The gradient of BPR loss with respect to the model parameters Θ (e.g., user/item embeddings) can be written as:

$$\begin{aligned} \frac{\partial L_{BPR}}{\partial \Theta} &= \sum_{(u, i, j) \in D_s} \frac{\partial \log \sigma(\hat{y}_{ui} - \hat{y}_{uj})}{\partial \Theta} - \lambda_\Theta \frac{\partial \|\Theta\|^2}{\partial \Theta} \\ &\propto \frac{-e^{-(\hat{y}_{ui} - \hat{y}_{uj})}}{1 + e^{-(\hat{y}_{ui} - \hat{y}_{uj})}} \cdot \frac{\partial (\hat{y}_{ui} - \hat{y}_{uj})}{\partial \Theta} - \lambda_\Theta \Theta \end{aligned} \quad (4)$$

More concretely, the gradients for updating a matrix factorization model are:

$$\frac{\partial (\hat{y}_{ui} - \hat{y}_{uj})}{\partial \Theta} = \begin{cases} (\mathbf{Q}_i - \mathbf{Q}_j) & \text{if } \Theta = \mathbf{P}_u \\ \mathbf{P}_u & \text{if } \Theta = \mathbf{Q}_i \\ -\mathbf{P}_u & \text{if } \Theta = \mathbf{Q}_j \\ \mathbf{0} & \text{otherwise} \end{cases} \quad (5)$$

Failure case analysis. For BPR loss optimization, we notice that for each user-item pair (u, i) with positive interaction, the model will randomly sample a negative pair (u, j) . For a popular item i which matches mainstream preference, it is frequently updated by a positive gradient \mathbf{P}_u . It is only updated by a negative gradient $-\mathbf{P}_u$ when it serves as a negative sample. In other words, item i will be updated with positive gradients more frequently than negative gradients during the optimization.

Table I

STATISTICS OF THE TRAINING DATA IN OUR MOVIELENS-10M DATASET. $\text{pop}(\text{unp})_i4u$ AND $\text{act}(\text{ina})_u4i$ INDICATE THE AVERAGE NUMBER OF POPULAR (UNPOPULAR) ITEMS FOR A USER AND ACTIVE (INACTIVE) USERS FOR AN ITEM RESPECTIVELY.

User Group	#Users	pop_i4u	unp_i4u
All	37,962	17.3	4.3
Active	20,801	25.1	6.6
Inactive	17,161	7.9	1.7
Item Group	#Items	act_u4i	ina_u4i
All	4,819	136.6	34.2
Popular	328	1,591.6	414.0
Unpopular	4,491	30.3	6.4

III. MOTIVATION: POPULARITY BIAS FROM A GRADIENT PERSPECTIVE

In this section, we discuss popularity bias from a gradient update perspective, i.e., how the item popularity, user activity and user conformity influence the gradient magnitude and shift the overall gradient to the positive interaction direction for popular items.

Based on Eq.5, we can derive that user (item) embeddings affect the magnitude and direction of the item (user) embeddings. For simplicity, we define the update of the latent vector of positive items \mathbf{Q}_i^t , negative items \mathbf{Q}_j^t , and users \mathbf{P}_u^t in the t^{th} iteration step for a data example (u, i, j) with the learning rate lr as follows:

$$\begin{aligned} \mathbf{Q}_i^t &= \mathbf{Q}_i^{t-1} - lr \cdot \mathbf{P}_u^{t-1}, \\ \mathbf{Q}_j^t &= \mathbf{Q}_j^{t-1} - lr \cdot (-\mathbf{P}_u^{t-1}), \\ \mathbf{P}_u^t &= \mathbf{P}_u^{t-1} - lr \cdot (\mathbf{Q}_i^{t-1} - \mathbf{Q}_j^{t-1}) \end{aligned} \quad (6)$$

For supporting the analysis of popularity bias from a gradient perspective. We take the Movielens-10M as an example, and statistics is shown in Table I. We sort items by their number of interactions in descending order and select sorted items until the number of interactions of selected items up to 80% of total interactions in the training set (in fact, less than 20% items account for more than 80% interactions in many real scenarios). We assume this selected items as popular items and the remaining items as unpopular items. In the same way, we can get active users and inactive users.

Influence of User Activity on Item Side. Let the latent vector of item i in the 0-th iteration be \mathbf{Q}_i^0 . Then we can obtain the equivalent form of \mathbf{Q}_i^t as:

$$\mathbf{Q}_i^t = \mathbf{Q}_i^0 - lr \cdot \left(\sum_{(u, i) \in Y_i^{+, t-1}} \mathbf{P}_u - \sum_{(u', i) \in Y_i^{-, t-1}} \mathbf{P}_{u'} \right) \quad (7)$$

where $Y_i^{+, t-1}$ and $Y_i^{-, t-1}$ denote the set of use-item pairs in which the item i works as the positive and negative items in training set before the iteration step t , respectively. Without loss of generalization, we simplify all equations later by ignoring iteration steps t and $t-1$. We let Δ_i represent the increment of \mathbf{Q}_i^t :

$$\Delta_i = \sum_{(u, i) \in Y_i^+} \mathbf{P}_u - \sum_{(u', i) \in Y_i^-} \mathbf{P}_{u'} \quad (8)$$

Since users in both Y_i^+ and Y_i^- consist of active users and inactive users, equation 8 can be extended as:

$$\begin{aligned} \Delta_i = & \sum_{(u,i) \in Y_{i,act}^+} \mathbf{P}_u + \sum_{(u,i) \in Y_{i,inact}^+} \mathbf{P}_u \\ & - \sum_{(u',i) \in Y_{i,act}^-} \mathbf{P}_{u'} - \sum_{(u',i) \in Y_{i,inact}^-} \mathbf{P}_{u'} \end{aligned} \quad (9)$$

where $Y_{i,act}^+$ ($Y_{i,inact}^+$) is the set of use-item pairs for item i in which users are active (inactive) users.

As shown in Table I, for items in any group (all, popular, and unpopular), the number of positive active users is much larger than inactive users. Besides, the number of negative active users is much smaller than inactive users via negative sampling over uniform distribution. Thus, the items embeddings are easily dominated by active users Inferential Failure, according to the cumulative gradients calculated by Eq. 8. We can derive that the higher the popularity of the item is, the larger the magnitude of its corresponding embedding will be. Moreover, the distance between popular items might be close in latent space due to the user conformity.

Influence of Item Popularity on User Side. Similar to analyze items from the gradient perspective, we can obtain accumulated gradients Δ_u of user u by:

$$\begin{aligned} \Delta_u = & \sum_{(u,i) \in Y_{u,pop}^+} \mathbf{Q}_i + \sum_{(u,i) \in Y_{u,unp}^+} \mathbf{Q}_i \\ & - \sum_{(u,j) \in Y_{u,pop}^-} \mathbf{Q}_j - \sum_{(u,j) \in Y_{u,unp}^-} \mathbf{Q}_j \end{aligned} \quad (10)$$

For users in any group (all, active, inactive), the number of positive popular items is much larger than unpopular items. The number of negative popular items is however much smaller than unpopular items in randomly negative sampling. Therefore, popular items significantly dominate user embeddings due to the accumulated gradients defined in Eq. 10.

IV. PROPOSED SOLUTION

In this section, we present our method. Specifically, in the training stage, our method normalizes only user embeddings after learning the user and item embeddings using LGN [15] or MF [14]. We then apply post-hoc debiasing on user and item embeddings in the testing stage to mitigate item popularity bias. We first give our definition of an unbiased item (user) embedding in implicit RS. We next give a detailed description about how to mitigate user activity in the training stage, and how to mitigate item popularity and user conformity influence.

- **Disentangled embedding learning.** We first define four factors influencing user and item embedding learning.
- **A Post-hoc debiasing method.** Following the gradient optimization process, we analyze how the four factors propagate popularity bias into embedding updating phase and design a specific debiasing term for each of the four factors accordingly.

A. Unbiased Embedding Definition

As described in Section 3, there are three factors exacerbating the imbalance issue between accumulated positive and negative gradients of the each item: (1) user's high activity interacting with many items; (2) item's high popularity involved in more positive interactions. (3) user's conformity towards other mainstream users. For popular items, its overall gradient is shifted towards the positive gradient direction, making the popular items heavily exposed. The unpopular items in contrast are under-exposed. The large magnitude of popular item vectors results in higher ranking scores for positive interactions. Since updating of user and item latent vectors is intertwined in the implicit RS optimization process, we propose to adjust both user and item latent vectors to mitigate popular bias.

Given the item embedding \mathbf{Q}_i and user embedding \mathbf{P}_u trained using vanilla MF [14] (or LGN [15]) model, we represent the item embedding as:

$$\mathbf{Q}_i = \mathbf{Q}_i^{pop} + \mathbf{Q}_i^{int} \quad (11)$$

where \mathbf{Q}_i^{pop} captures the item's popularity and \mathbf{Q}_i^{int} captures the item's genuine features. Similarly, we represent a user embedding \mathbf{P}_u as:

$$\mathbf{P}_u = \|\mathbf{P}_u^{act}\|_2 (\mathbf{P}_u^{conf} + \mathbf{P}_u^{int}) \quad (12)$$

where \mathbf{P}_u^{act} captures the users' activeness, \mathbf{P}_u^{conf} signifies the user's conformity to other mainstream users, and \mathbf{P}_u^{int} denotes the user's genuine interest. Notably, taking the product of \mathbf{P}_u^{act} and $\mathbf{P}_u^{conf} + \mathbf{P}_u^{int}$ is justified by the fact that an active user is inclined to have more observed interactions [17]. Formally, the unbiased matching score for a user and an item is:

$$\hat{y}_{ui} = \mathbf{P}_u^{int} \cdot \mathbf{Q}_i^{int} \quad (13)$$

In the following sections, we will describe how we mitigate \mathbf{P}_u^{act} in the training stage and eliminate \mathbf{Q}_i^{pop} and \mathbf{P}_u^{conf} in the testing stage.

B. Normalizing User Embeddings in Training

As an active user will have more observed interactions with items, their accumulated gradients based on these items will likely be larger than those of other less active users. For example, if users u_1 and u_2 share similar item preferences with u_1 being more active, the embedding vectors of u_1 and u_2 will have similar updating direction but the vector norm of user u_1 will be much larger. Such large embedding contributes to the gradient updating of its corresponding items, aggravating the imbalance issue between accumulated positive and negative gradients of items.

To reduce the bias propagation of highly active users, we normalize each user embedding $\hat{\mathbf{P}}_u$ into a unit vector during the training stage:

$$\hat{\mathbf{P}}_u = \mathbf{P}_u / \|\mathbf{P}_u^{act}\|_2 \quad (14)$$

where $\hat{\mathbf{P}}_u$ is the learned user embedding after normalization. To keep the notation simple, we use $\|\mathbf{P}_u\|$ to denote $\|\mathbf{P}_u^{act}\|_2$. The complete training procedure is shown in Algorithm 1.

Algorithm 1 Training

Require: Backbone user encoder $f(\cdot|\Theta_U)$, backbone item encoder $f(\cdot|\Theta_I)$, user set \mathcal{U} , item set \mathcal{I} , and user-item interactions Y .

- 1: Initialize the user and item encoder parameters Θ_U and Θ_I respectively;
- 2: $\bar{P} = 0$ and $\bar{Q} = 0$
- 3: **for** $t = 1$ to T **do**
- 4: $Y_t^+ \leftarrow \text{SampleMiniBatch}(Y, m)$
- 5: $Y_t^- \leftarrow \text{NegativeSampler}(Y_t^+)$
- 6: **for** $k = 1$ to m **do**
- 7: $u, i, j = Y_{t,k}^+, Y_{t,k}^-$
- 8: $P_u, Q_i, Q_j = f(u|\Theta_U), f(i|\Theta_I), f(j|\Theta_I)$
- 9: $\hat{P}_u = P_u / \|P_u\|_2$
- 10: $\bar{P}_u = \bar{P}_u - lr \cdot (Q_i - Q_j)$
- 11: $\bar{Q}_i = \bar{Q}_i - lr \cdot (\hat{P}_u)$
- 12: $\bar{Q}_j = \bar{Q}_j - lr \cdot (-\hat{P}_u)$
- 13: $\hat{y}_{ui} = \hat{P}_u \cdot Q_i$
- 14: $\hat{y}_{uj} = \hat{P}_u \cdot Q_j$
- 15: $\mathcal{L}_k = \log(\sigma(\hat{y}_{ui} - \hat{y}_{uj})) - \lambda_\Theta \|\Theta\|^2$
- 16: **end for**
- 17: $\mathcal{L}_t = \frac{1}{m} \sum_{k=1}^m \mathcal{L}_k$
- 18: $\Theta \leftarrow \Theta - lr \cdot \nabla_\Theta \mathcal{L}(f_\Theta)$
- 19: **end for**
- 20: $\bar{P} = \frac{1}{U} \sum_{u=1}^U \bar{P}_u$
- 21: $\bar{Q} = \frac{1}{I} \sum_{i=1}^I \bar{Q}_i$

C. Adjusting Popularity Bias in Inference

1) *Mitigating item popularity influence.*: For a popular item, its positive interactions with users are frequently observed among the training samples, shifting the learned embedding to the positive interaction direction. As depicted in Fig.3, the motivation here is to modify the overall learned item embedding \overrightarrow{AD} to an unbiased item vector $\overrightarrow{AD'}$. There are a number of ways to do so. In this paper, we propose to modify the biased accumulated gradient \overrightarrow{AD} only through its accumulated positive gradients \overrightarrow{AC} and keep \overrightarrow{AB} unchanged. Nevertheless, the iterative optimization of user/item embed-




Figure 3. \overrightarrow{AG} denotes the popular direction of item embeddings. Our approach adjusts \overrightarrow{AC} to $\overrightarrow{AC'}$ through \overrightarrow{AG} , leading to a new modified item embedding $\overrightarrow{AD'}$.

ding makes it unclear how much other items and relevant user embeddings contribute to the accumulated positive gradients for each item, worsening the imbalance between \overrightarrow{AB} and \overrightarrow{AC} . Hence, the important question here is how to approximate the biased embedding P_u^{pop} in Eq.12. Our analysis in Section 3 suggests that 1) for popular items, their overall accumulated gradient (learned) embeddings (\overrightarrow{AD}) are shifted towards its positive gradient direction (\overrightarrow{AC}); 2) Iterative optimization of user and item embedding makes item embeddings are intertwined with other relevant item information. Besides, user

Algorithm 2 Adjusting Popularity Bias in Inference

Input: Backbone user encoder $f(\cdot|\Theta_U)$, backbone item encoder $f(\cdot|\Theta_I)$, user u , item i , recorded user conformity \bar{P} , and recorded item popularity \bar{Q} .

Output: \hat{y}_{ui}

- 1: $P_u, Q_i, Q_j = f(u|\Theta_U), f(i|\Theta_I), f(j|\Theta_I)$
- 2: $Q_i^{int} = Q_i - \alpha_1 \cdot \cos(Q_i, \bar{Q}) \cdot \|Q_i\| \cdot \bar{Q}$
- 3: $P_u^{int} = P_u - \alpha_2 \cdot \cos(P_u, \bar{P}) \cdot \|P_u\| \cdot \bar{P}$
- 4: $\hat{y}_{ui} = P_u^{int} \cdot Q_i^{int}$

conformity and activity makes their item embeddings reflects similar user preference. Thus, popular items share similar positive gradient directions. Based on the two findings, we use the mean embedding of popular items' embedding to represent such biased embedding Q^{pop} . The direction of such biased vector is named as popular direction \bar{Q} .

Thus, we decompose the an item embedding Q_i into a *genuine features embedding* Q_i^{int} and a projection on the popular direction \bar{Q} . Projection vector here refers to item popularity embedding Q_i^{pop} in Eq.11. In the test phase, we adjust the learned item embedding Q_i and derive its genuine representation Q_i^{int} :

$$Q_i^{int} = Q_i - Q_i^{pop} = Q_i - \alpha_1 \cdot \cos(Q_i, \bar{Q}) \cdot \|Q_i\| \cdot \bar{Q} \quad (15)$$

where α_1 is a scaling parameter which controls the magnitude of the modified positive vector AC' . In this paper, we choose the value of α_1 based on the validation set.

2) *Mitigating user conformity.*: Different from item embeddings which are updated by both positive and negative gradients, a user embedding is only updated by the gradient $Q_i - Q_j$, in which items i and j appear in observed and non-observed interactions with the user, respectively. A user embedding is inevitably influenced by the preference of other users during the gradient updates with BPR loss, and thus account for user conformity.

For example, consider a scenario where three users interact with an item i_1 , while the first user u also interacts with item i_2 and i_3 . Since i_1 has multiple observed interactions, it receives the accumulated gradients based on three users. Whenever we update user embedding $\nabla_u \mathcal{L}(u, i) = Q_{i_1}$. The embedding of Q_{i_1} with large magnitude will increase the magnitude of latent vector of u .

The situation here is similar to item popularity bias, we therefore adopt a similar way to extract user's true interest by subtracting the user conformity projection:

$$P_u = P_u - \alpha_2 \cdot \cos(P_u, \bar{P}) \cdot \|P_u\| \cdot \bar{P} \quad (16)$$

where \bar{P} denotes the mean or normalized embedding of mainstream users' embedding. The term $\alpha_2 \cdot \cos(P_u, \bar{P}) \cdot \|P_u\| \cdot \bar{P}$ captures the user conformity embedding P_u^{conf} in Eq. 12. α_2 is a scaling parameter. We show inference procedure in Algorithm 2.

V. EXPERIMENTS

In this section, we conduct experiments to show the effectiveness and efficiency of the proposed framework. Specifically, we aim to answer the following research questions:

Table II
STATISTICS OF DATASETS.

Dataset	User	Item	Interaction	Sparsity
Movielens-10M	37,962	4819	1,371,473	0.007496
Netflix	32,450	8432	2,212,690	0.008086
Adresa	13,485	744	116,321	0.011594
Gowalla	29,858	40,981	1,027,370	0.000840

- **RQ1:** (a) Does our proposed method outperform state-of-the-art recommendation methods under the non-IID setting? (b) Furthermore, does it generalize well on both BPR and BCE loss? (c) Is it more efficient than state-of-the-art methods?
- **RQ2:** How do different components (e.g., normalization, user conformity subtrahend, item popularity subtrahend, and these subtrahends' hyper-parameters¹) affect the result of our method?
- **RQ3:** (a) How well does our method mitigate the popularity bias? (b) How does the performance of our method change with the intervention proportion of test data?

A. Experiment Settings

Datasets and Data Preprocessing. Our experiments utilize four publicly available datasets: 1) Movielens-10M² is a dataset with 10M movie ratings. 2) Netflix [18] is from an open competition and it covers film ratings. 3) Adresa [19] is a popular dataset for news recommendation. 4) Gowalla [20] is a dataset covering location-based user check-in behaviors. We summarize statistics of the datasets in Table II.

To investigate the extent to which the debiased models reduce popularity bias, we replace the conventional evaluation strategy, where the test set consists of IID samples from the observed interactions, with a sampling strategy to construct a debiased test set [21]. Specifically, we randomly sample interactions with equal probability given to all items to construct the test and validation data. We then use the remaining data that come with popularity bias as the training data. We follow DICE [12] to process Movielens-10M and Netflix datasets and MACR [11] to process Adresa and Gowalla for a fair comparison. More specifically, we sample 30% unbiased interactions regarding items as the test set, 10% as the validation set, and the remaining 60% as training data for Movielens-10M and Netflix. We prepare 10% unbiased data as the test set, another 10% unbiased data as the validation set, and the rest as the training set for Adresa and Gowalla.

Baselines. We implement our method over the classical MF and the state-of-the-art LGN to investigate the effectiveness and generalization ability of our proposed approach across different backbone recommendation models. We compare our methods with the following baselines. Matrix factorization (MF) [14] is a representative collaborative filtering method. LGN [15] is a state-of-the-art collaborative filtering based on light graph convolution network. IPS [4] is a widely used re-weighting method based on causal inference. The

weighting score, i.e., inverse propensity score is set as the inverse of corresponding item popularity value. It eliminates popularity bias by imposing large weight to unpopular items. **CausE** [21] is a domain adaptation algorithm that learns a causal embedding from a small uniform dataset. **DICE** [12] is the state-of-the-art disentangled method for learning causal embedding to cope with popularity bias problem. It designs a framework with causal-specific data to disentangle interest and popularity with pre-defined guidelines. DICE leverage popularity bias in the inference. We use the code provided by authors. **MACR** [11] addresses popularity bias from a causal-effect view. The counterfactual inference is performed to estimate the direct effect from item properties to the ranking score, which is removed to eliminate the popularity bias.

Implementation Details. We evaluate the effectiveness performance by three widely-used metrics, i.e., *top-k Recall* (Recall@k), *top-k Hit Ratio* (HR@k), and *top-k Normalized Discounted Cumulative Gain* (NDCG@k), which are used in DICE [12] and MACR [11]. For MACR on Movielens-10M and Netflix, we use the source code provided by their authors and set the batch size and embedding size as 128 for a fair comparison. For the reported results of the remaining baselines, we refer to DICE and MACR directly. Note that our method is based on two backbones (i.e., MF and LGN). The hyper-parameters of backbones are set following the suggestions from DICE for Movielens-10M and Netflix and MACR for Adresa and Gowalla to guarantee a fair comparison with DICE and MACR. α_1 and α_2 are determined on validation set based on grid search in the search list formed by taking values from 0 to 2 with a 0.2 increment.

B. Overall Comparison (RQ1)

The recommendation performance (i.e., Recall@20, HR@20, and NDCG@20) of different models are presented in Table III. The methods with best performance are marked with boldface fonts and the second best ones are underlined. Based on the results, we make the following observations:

- As shown in Table III, our proposed method consistently outperforms all the baselines on all the datasets w.r.t. all the evaluation metrics. The state-of-the-art baselines DICE and MACR incorporate extra modules to reduce popularity bias, which demands additional prior knowledge and efforts to exploit good architecture by trial and error during the training stage. In contrast, our method only needs to store the aggregated influence of the dominating items and users that is caused by gradient based updates. It efficiently handles popularity bias by performing a post-hoc subtraction calculation in the inference step, which eventually leads to better performance than baselines.
- Our proposed method is both general and effective. It can easily be incorporated into various backbone RS models such as MF, LGN, and any other general recommendation models. At the same time, it achieves promising performance on debiased test set. Compared to classical MF, our proposed method that uses MF as the backbone model achieves average 37.33%, 33.87%, and 51.79% improve-

¹We show the effect of hyper-parameters α_1 and α_2 in detail in Appendix.

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

Table III
OVERALL PERFORMANCE COMPARISON (RECALL@20, HR@20, AND NDCG@20) OF DIFFERENT METHODS ON FOUR DATASETS.

Dataset		Movielens-10M			Netflix			Adressa			Gowalla		
Model	Method	Recall	HR	NDCG	Recall	HR	NDCG	Recall	HR	NDCG	Recall	HR	NDCG
MF	Original	0.128	0.442	0.084	0.112	0.519	0.094	0.085	0.111	0.034	0.046	0.174	0.032
	IPS	0.133	0.443	0.085	0.105	0.488	0.086	0.096	0.128	0.039	0.048	0.174	0.033
	CausE	0.115	0.406	0.074	0.093	0.464	0.078	0.084	0.112	0.037	0.045	0.166	0.032
	DICE	0.163	0.519	0.108	0.125	0.554	0.107	0.098	0.133	0.041	0.052	0.177	0.033
	MACR	0.138	0.469	0.089	0.119	0.533	0.101	0.109	0.140	0.050	0.077	0.252	0.050
	Ours	0.173	0.536	0.116	0.138	0.575	0.119	0.116	0.147	0.056	0.081	0.259	0.057
LGN	Original	0.137	0.462	0.089	0.102	0.490	0.087	0.098	0.123	0.040	0.045	0.172	0.032
	IPS	0.139	0.464	0.091	0.110	0.509	0.095	0.107	0.139	0.047	0.045	0.174	0.032
	CausE	0.102	0.372	0.063	0.083	0.428	0.067	0.082	0.115	0.037	0.046	0.173	0.033
	DICE	0.181	0.556	0.122	0.142	0.591	0.121	0.111	0.141	0.046	0.054	0.185	0.036
	MACR	0.157	0.508	0.101	0.119	0.531	0.094	0.127	0.158	0.052	0.077	0.254	0.051
	Ours	0.186	0.564	0.127	0.145	0.595	0.127	0.134	0.166	0.067	0.081	0.263	0.053

ments in Recall@20, HR@20, and NDCG@20, respectively. Similarly, our proposed method boosts the performance of original LGN about 51.33%, 32.83%, 56.34% w.r.t. Recall@20, HR@20, and NDCG@20, respectively. These significant results demonstrate that the method effectively mitigates the popularity bias in the training and test stages.

- Our proposed method has stronger generalization ability, i.e., it is more robust, as it consistently outperforms others on all four datasets with diverse context and data distribution. While the state-of-the-art methods DICE and MACR achieve comparable performance, their performance may not always be consistent across the different datasets. A possible reason is that DICE and MACR require elaborate configuration of incorporated modules and take much effort to tune the optimal hyper-parameters. Another possible reason to explain the inconsistency of DICE is that DICE depends on two rule-based strategies to obtain negative examples and to employ negative or positive BPR loss to update the model. The strategies are unfortunately sensitive to data distribution.
- As to the remaining baselines for mitigating popularity bias, CausE has limited improvement over the basic models and sometimes performs even worse. The reason should be that CausE aligns user and item embeddings with a tiny uniform dataset which hardly allows the model to learn enough user preference information. IPS performs poorly in most cases as it simply up-weighs unpopular items with a static score, making the model overly-confident to recommend unpopular items.

C. Further Analysis

In this section, we continue to evaluate our proposed method by answering the remaining research questions:

1) *Does it generalize well on both BPR and BCE loss? (RQ1(b))*: Figure 4 shows how our method using MF as the backbone performs when applied with BPR or BCE loss on Movielens-10M and Netflix. Overall, subtracting user conformity bias and item popularity bias in the testing stage provides substantially better performance than the models using vanilla user and item embeddings. This observation validates our view that the optimization of BPR and BCE incorporates the user conformity bias and the item popularity bias into user and item embeddings, respectively, making the recommendation




Figure 4. Performance (Recall@20, HR@20, and NDCG@20) of MF and Our Method with BPR and BCE loss on Movielens-10m and Netflix datasets.

Table IV
COMPARISON OF COMPUTING TIME FOR OUR METHOD VERSUS DICE AND MACR. ALL METHODS ARE RUN ON 1 TESLA V100 GPU CHIP.

Model	Method	Seconds per epoch	
		Movielens-10M	Netflix
MF	Original	47	77
	DICE	124 (2.63x)	207 (2.68x)
	MACR	57 (1.21x)	87(1.29x)
	Ours	51 (1.08x)	78 (1.01x)
LGN	Original	126	201
	DICE	307 (2.43x)	454 (2.25x)
	MACR	304 (2.41x)	447 (2.22x)
	Ours	129 (1.02x)	203 (1.01x)

results deviating from users' true interest. The proposed post-hoc design is an elegant and effective way to remove bias.

2) *Is our method more efficient than state-of-the-art baselines? (RQ1(c))*: We investigate the efficiency of our proposed method by comparing the running time on the same device against the two state-of-the-art methods, i.e., DICE and MACR, with MF and LGN backbones. As shown in Table IV, we observe that our method achieves comparable running time with classical MF and LGN methods, and it is significantly more efficient than DICE and MACR. With special causal learning design and direct supervision on disentanglement, DICE takes twice more time than our proposed method due to the additional cost of parameter updating. MACR includes additional modules to reconstruct user and item frequencies, and is therefore less efficient. Another interesting side observation is that DICE always suffers from excessive computation time compared with MACR. We conjecture that the additional two rule-based strategies of DICE are partially responsible for the

Table V

ABLATION STUDY (RECALL@20, HR@20, AND NDCG@20) OF OUR METHOD ON MOVIELENS-10M AND ADRESSA DATASETS.

Datasets		Movielens-10M			Adressa		
Model	Method	Recall	HR	NDCG	Recall	HR	NDCG
MF	Ours	0.173	0.536	0.116	0.116	0.147	0.056
	w/o norm	0.164	0.517	0.109	0.098	0.125	0.044
	$\alpha_1 = 0$	0.153	0.497	0.101	0.088	0.113	0.043
	$\alpha_2 = 0$	0.160	0.520	0.111	0.089	0.133	0.049
LGN	Ours	0.186	0.564	0.127	0.134	0.166	0.067
	w/o norm	0.171	0.532	0.114	0.122	0.148	0.052
	$\alpha_1 = 0$	0.154	0.498	0.101	0.119	0.151	0.061
	$\alpha_2 = 0$	0.175	0.540	0.117	0.113	0.144	0.051

excess time consumed.

3) *Is it necessary to normalize user embeddings and is it desired to subtract user conformity and item popularity term? (RQ2):* We conduct an ablation study with our method using MF and LGN backbones on Movielens-10M and Adressa datasets to analyze the contribution of different components. Specifically, we compare our method with its three special variants: our method without (w/o) user norm (i.e., MF w/o norm and LGM w/o norm), where normalization on user embedding is removed; our method with $\alpha_1 = 0$, where we just simply drop the user conformity embedding term in the inference; our method with $\alpha_2 = 0$ that drops the item popularity embedding term.

The results in Table V show that removing any component would lead to worse performance. It indicates that all three components contribute to improving the model performance under non-IID circumstances. Omitting user embedding normalization would reduce accuracy in all metrics on both Movielens-10M and Adressa datasets. This validates the usefulness of normalizing user embedding magnitudes in the training stage, i.e., promoting items with fewer interactions benefits recommendation performance. It is clear that compared with the variant without item popularity bias, the model performs much worse when removing the user conformity bias in most cases. One possible reason is that user has more impact than item for these two datasets.

4) *How does our method mitigate the popularity bias? (RQ3(a)):* In this study, we investigate whether our method can remove the effect of popularity bias by the methods, i.e., MF (LGN), Our_MF (Our_LGN) w/o norm, and Our_MF (Our_LGN) on Movielens-10M. For fine-grained debias analysis, we first sort items in inverse popularity order. We then divide items into five groups, each of the first four groups accounts for 5% of items, and we put the rest into the last group. Group 1 represents the most popular item group, and Group 5 represents the least popular item group. For each group, we calculate the average recommended frequency and recall@20. We show the recommended frequency of each group by methods in Figure 5 and recall in Figure 6.

Conventional recommendation methods have the most significant recommended frequency and perform best in group 1 but yield a significant performance drop in group 2 and further drops in other groups. It indicates that conventional recommender systems tend to recommend popular items to




Figure 5. Recommended frequency of different item groups by MF (LGN), Our_MF (Our_LGN) w/o norm, and Our_MF (Our_LGN) on Movielens-10M.




Figure 6. Performance (Recall@20) of MF (LGN), Our_MF (Our_LGN) w/o norm, and Our_MF (Our_LGN) in different item groups on Movielens-10M.




Figure 7. Performance of MF and Ours_MF with different proportion of intervened test data.

unrelated users due to user conformity and popularity bias. In contrast, conventional recommender systems with our proposed post-hoc method have less recommended frequency and performance in the most popular group but achieve a significantly higher recommended frequency and better recommendation performance from group 2 to group 4. It shows that our proposed method effectively reduces user conformity and item popularity bias.

In group 5, all methods achieve similar recommended frequency and performance. The possible reason is that these least popular items rarely appear in the training set so that the representation of these items cannot be learned effectively. In some cases, our method achieves less recommended frequency but higher recall than our method w/o norm. It means that our method can still more accurately recommend items to users.

5) *How does our method's performance change along with intervention proportion of test data? (RQ3(b)):* In previous experiments, all methods are tested on the 100% intervened test data. This study investigates how methods perform as the proportion of intervened test data changes. We evaluate MF and our method (MF) with five different proportions of intervened test data, i.e., 0%, 50%, 75%, 90%, 100%. 0% means training and test data have the same distribution. 100% means all test data are intervened data. As shown in

Figure 7, our method starts to outperform MF on 75% of intervened test data and outperforms most significantly under 100% intervened test data, which verifies the effectiveness of our method under non-IID circumstances.

6) *Effect of the hyper-parameters of subtrahends*: Our post-hoc method has two hyper-parameters α_1 and α_2 for the user conformity subtrahend and the item popularity subtrahend respectively. We experiment with our method using MF as the backbone on MovieLens-10M to investigate their effect. Figure 8 shows how our method’s performance (recall@20) changes as α_1 and α_2 change from 0 to 2 over the interval.

As we can see, the model performs well when α_1 is between 0.4 to 0.6 and performs increasingly better while increasing α_2 in most cases. It suggests that the appropriate amount of mitigating user conformity bias and item popularity bias benefits the recommendation system under non-IID circumstances. Large α_1 yields a drop in performance while large α_2 improves the recommendation performance, which indicates that excessive user conformity debiasing hurts the performance. One possible reason is that user normalization overly constrains the magnitude of user embeddings.




Figure 8. Effect of α_1 and α_2 in our method (MF) for mitigating the user conformity bias and the item popularity bias respectively.

VI. RELATED WORK

In recent years, recommendation systems have achieved great success due to the merit of deep learning methods, most of which aim at the development of novel machine learning models to fit user behavior data. When it comes to real-world scenarios, the distribution of user and item interaction is inevitably long-tailed, making popularity bias a long-standing challenge [22]–[24].

The first core idea in mitigating popularity bias is to balance the data distribution such that popular and unpopular items are equally important in the training process. Inverse propensity weighting (IPW) [8], [25] is widely used to assign weights inversely proportional to item popularity in the empirical risk of a RS model [26]. While such reweighting methods achieve significant improvement for under-represented items, they come with a high potential risk of overfitting and suffer from high variance [27], [28]. Recent work [29]–[32] studies this problem from a fairness perspective [33], [34], and enforce

the algorithm to produce a fair allocation of exposure based on merit. Other strategies following the first pipeline is to transfer the knowledge learned from popular groups to unpopular groups [35]. Such methods often introduce extra noise to unpopular items and modify the initial dataset distribution. Recent work [36] investigates the bias from the perspective of causal inference and proposes a counterfactual reasoning method using Do-calculus [37]. They propose a different causal graph based on prior knowledge, however, the essence is to re-balance the distribution for popular and unpopular groups with a causal language.

Another line of research is to preserve specific user and item information intentionally, with the goal to make the unpopular items more representative [10]. With auxiliary text information, [10] an autoencoder layer [38], [39] is added while learning user and item representations with text-based Convolutional Neural Networks. Latest work [21] regards a personalized recommendation system as a treatment policy utility maximization problem and proposes a domain adaptation method named CausE. [40] propose a meta-learning strategy and transfer unbiased data representation from a uniform logged data. However, the availability and high-quality of auxiliary dataset is the bottleneck for such approaches.

The third categories try to solve this problem through ranking adjustment [41]–[43]. These approaches result in a trade-off between the recommendation accuracy and the coverage of unpopular items. They typically suffer from accuracy drop due to give consideration to the long-tail in a brute manner.

Different from previous method, in this paper we do not focus on balancing the distribution of popular and unpopular items groups or making unpopular items more representative. We provide a new insight that unbalancing between the positive and negative gradients of the same item incurs the popularity bias. Some related works [44] discuss the gradient vanish of unpopular items from a gradient perspective in implicit recommendation system. However, they still focus on the unbalanced distribution between popular and unpopular item groups, which actually fall into the former three categories.

VII. CONCLUSION

In this paper, we propose a simple, efficient, and general framework to mitigate popularity bias from a gradient (first-order information) perspective. In our approach, we record accumulated gradients of users and items as popularity bias during the training stage. Then we address popularity bias by gradient-level calibration during the testing stage as a post-hoc debiasing strategy. We implement our proposed method over MF and LGN on four real-world datasets. Extensive experiments demonstrate that our approach outperforms competitive baselines, e.g., DICE and MACR. As part of the future work, we will consider applying our approach to more complex recommendation tasks, such as sequential recommendation and conversational recommendation.

VIII. ACKNOWLEDGEMENT

This research was partially supported by the National Science Foundation (NSF) via the grant numbers: 2040950, 2006889, 2045567.

REFERENCES

- [1] L. Yang, Y. Cui, Y. Xuan, C. Wang, S. Belongie, and D. Estrin, "Unbiased offline recommender evaluation for missing-not-at-random implicit feedback," in *Proceedings of the 12th ACM Conference on Recommender Systems*, 2018, pp. 279–287.
- [2] L. Xu, W. Ouyang, X. Ren, Y. Wang, and L. Jiang, "Enhancing semantic representations of bilingual word embeddings with syntactic dependencies," in *IJCAI*, 2018, pp. 4517–4524.
- [3] W. Ren, L. Zhang, B. Jiang, Z. Wang, G. Guo, and G. Liu, "Robust mapping learning for multi-view multi-label classification with missing labels," in *International Conference on Knowledge Science, Engineering and Management*. Springer, 2017, pp. 543–551.
- [4] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, and X. He, "Bias and debias in recommender system: A survey and future directions," *arXiv preprint arXiv:2010.03240*, 2020.
- [5] L. Wang, E.-P. Lim, Z. Liu, and T. Zhao, "Explanation guided contrastive learning for sequential recommendation," *arXiv preprint arXiv:2209.01347*, 2022.
- [6] J. Yin, C. Liu, W. Wang, J. Sun, and S. C. Hoi, "Learning transferable parameters for long-tailed sequential user behavior modeling," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 359–367.
- [7] K. Lee and K. Lee, "My head is your tail: applying link analysis on long-tailed music listening behavior for music recommendation," in *Proceedings of the fifth ACM conference on Recommender systems*, 2011, pp. 213–220.
- [8] P. C. Austin and E. A. Stuart, "Moving towards best practice when using inverse probability of treatment weighting (iptw) using the propensity score to estimate causal treatment effects in observational studies," *Statistics in medicine*, vol. 34, no. 28, pp. 3661–3679, 2015.
- [9] Y. Ge, S. Liu, R. Gao, Y. Xian, Y. Li, X. Zhao, C. Pei, F. Sun, J. Ge, W. Ou *et al.*, "Towards long-term fairness in recommendation," in *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 2021, pp. 445–453.
- [10] R. Z. Li, J. Urbano, and A. Hanjalic, "Leave no user behind: Towards improving the utility of recommender systems for non-mainstream users," in *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 2021, pp. 103–111.
- [11] T. Wei, F. Feng, J. Chen, C. Shi, Z. Wu, J. Yi, and X. He, "Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system," *arXiv preprint arXiv:2010.15363*, 2020.
- [12] Y. Zheng, C. Gao, X. Li, X. He, Y. Li, and D. Jin, "Disentangling user interest and conformity for recommendation with causal embedding," in *Proceedings of the Web Conference 2021*, 2021, pp. 2980–2991.
- [13] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "Bpr: Bayesian personalized ranking from implicit feedback," *arXiv preprint arXiv:1205.2618*, 2012.
- [14] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [15] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "Lightgcn: Simplifying and powering graph convolution network for recommendation," in *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 2020.
- [16] H.-J. Xue, X. Dai, J. Zhang, S. Huang, and J. Chen, "Deep matrix factorization models for recommender systems," in *IJCAI*, vol. 17. Melbourne, Australia, 2017, pp. 3203–3209.
- [17] J. Liu, P. Dolan, and E. R. Pedersen, "Personalized news recommendation based on click behavior," in *Proceedings of the 15th international conference on Intelligent user interfaces*, 2010, pp. 31–40.
- [18] J. Bennett, S. Lanning *et al.*, "The netflix prize," in *Proceedings of KDD cup and workshop*, vol. 2007. Citeseer, 2007, p. 35.
- [19] J. A. Gulla, L. Zhang, P. Liu, Ö. Özgöbek, and X. Su, "The adressa dataset for news recommendation," in *Proceedings of the international conference on web intelligence*, 2017, pp. 1042–1048.
- [20] E. Cho, S. A. Myers, and J. Leskovec, "Friendship and mobility: user movement in location-based social networks," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2011, pp. 1082–1090.
- [21] S. Bonner and F. Vasile, "Causal embeddings for recommendation," in *Proceedings of the 12th ACM conference on recommender systems*, 2018, pp. 104–112.
- [22] W. Ren, K. Liu, T. Zhao, and Y. Fu, "Fair and effective policing for neighborhood safety: understanding and overcoming selection biases," *Frontiers in big data*, vol. 4, 2021.
- [23] T. Joachims, A. Swaminathan, and T. Schnabel, "Unbiased learning-to-rank with biased feedback," in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 2017, pp. 781–789.
- [24] W. Ren, P. Wang, X. Li, C. E. Hughes, and Y. Fu, "Semi-supervised drifted stream learning with short lookback," *arXiv preprint arXiv:2205.13066*, 2022.
- [25] X. Ren, J. Jiang, L. M. S. Khoo, and H. L. Chieu, "Cross-topic rumor detection using topic-mixtures," in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 2021, pp. 1534–1538.
- [26] T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, and T. Joachims, "Recommendations as treatments: Debiasing learning and evaluation," in *international conference on machine learning*. PMLR, 2016, pp. 1670–1679.
- [27] R. Guo, L. Cheng, J. Li, P. R. Hahn, and H. Liu, "A survey of learning causality with data: Problems and methods," *ACM Computing Surveys (CSUR)*, vol. 53, no. 4, pp. 1–37, 2020.
- [28] L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, and A. Zhang, "A survey on causal inference," *arXiv preprint arXiv:2002.02770*, 2020.
- [29] D. Sacharidis, "Building user trust in recommendations via fairness and explanations," in *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, 2020, pp. 313–314.
- [30] T. Zhao, G. Liu, C. Ma, E. Chen *et al.*, "Zero-shot learning: An energy based approach," in *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2018, pp. 797–806.
- [31] L. Xiao, Z. Min, Z. Yongfeng, G. Zhaoquan, L. Yiqun, and M. Shaoping, "Fairness-aware group recommendation with pareto-efficiency," in *Proceedings of the eleventh ACM conference on recommender systems*, 2017, pp. 107–115.
- [32] T. Zhao, X. Zhang, and S. Wang, "Exploring edge disentanglement for node classification," in *Proceedings of the ACM Web Conference 2022*, 2022, pp. 1028–1036.
- [33] N. Francez, *Fairness*. Springer Science & Business Media, 2012.
- [34] T. Zhao, L. Liu, G. Huang, H. Li, Y. Liu, L. GuiQuan, and S. Shi, "Balancing quality and human involvement: An effective approach to interactive neural machine translation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 9660–9667.
- [35] D. Liu, P. Cheng, Z. Dong, X. He, W. Pan, and Z. Ming, "A general knowledge distillation framework for counterfactual recommendation via uniform data," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020, pp. 831–840.
- [36] Y. Zhang, F. Feng, X. He, T. Wei, C. Song, G. Ling, and Y. Zhang, "Causal intervention for leveraging popularity bias in recommendation," *arXiv preprint arXiv:2105.06067*, 2021.
- [37] J. Pearl, *Causality*. Cambridge university press, 2009.
- [38] T. Zhao, X. Zhang, and S. Wang, "Graphsmote: Imbalanced node classification on graphs with graph neural networks," in *Proceedings of the 14th ACM international conference on web search and data mining*, 2021, pp. 833–841.
- [39] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [40] J. Chen, H. Dong, Y. Qiu, X. He, X. Xin, L. Chen, G. Lin, and K. Yang, "Autodebias: Learning to debias for recommendation," *arXiv preprint arXiv:2105.04170*, 2021.
- [41] H. Abdollahpouri, R. Burke, and B. Mobasher, "Controlling popularity bias in learning-to-rank recommendation," in *Proceedings of the eleventh ACM conference on recommender systems*, 2017, pp. 42–46.
- [42] T. Zhao, X. Tang, X. Zhang, and S. Wang, "Semi-supervised graph-to-graph translation," in *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2020, pp. 1863–1872.

- [43] X. Ren, L. Xu, T. Zhao, C. Zhu, J. Guo, and E. Chen, "Tracking and forecasting dynamics in crowdfunding: A basis-synthesis approach," in *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2018, pp. 1212–1217.
- [44] S. Rendle and C. Freudenthaler, "Improving pairwise learning for item recommendation from implicit feedback," in *Proceedings of the 7th ACM international conference on Web search and data mining*, 2014, pp. 273–282.