

Jan 17th, 12:00 AM

## **Beyond the Rating Matrix: Debiasing Implicit Feedback Loops in Collaborative Filtering**

Thorsten Krause

*Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Germany, thorsten.krause@dfki.de*

Daniel Stattkus

*Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Germany, daniel.stattkus@dfki.de*

Alina Deriyeva

*Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Germany, alina.deriyeva@dfki.de*

Jan Heinrich Beinke

*Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Germany, jan.beinke@dfki.de*

Oliver Thomas

*Deutsches Forschungszentrum für Künstliche Intelligenz GmbH, Germany, oliver.thomas@dfki.de*

Follow this and additional works at: <https://aisel.aisnet.org/wi2022>

---

### **Recommended Citation**

Krause, Thorsten; Stattkus, Daniel; Deriyeva, Alina; Beinke, Jan Heinrich; and Thomas, Oliver, "Beyond the Rating Matrix: Debiasing Implicit Feedback Loops in Collaborative Filtering" (2022). *Wirtschaftsinformatik 2022 Proceedings*. 12.

<https://aisel.aisnet.org/wi2022/ai/ai/12>

This material is brought to you by the Wirtschaftsinformatik at AIS Electronic Library (AISeL). It has been accepted for inclusion in Wirtschaftsinformatik 2022 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).

# Beyond the Rating Matrix: Debiasing Implicit Feedback Loops in Collaborative Filtering

Thorsten Krause<sup>1</sup>, Daniel Stattkus<sup>1</sup>, Alina Deriyeva<sup>1</sup>, Jan Heinrich Beinke<sup>1</sup> and Oliver Thomas<sup>1,2</sup>

<sup>1</sup> German Research Center for Artificial Intelligence (DFKI), Osnabrück, Germany  
{thorsten.krause, daniel.stattkus, alina.deriyeva, jan.beinke, oliver.thomas}@dfki.de

<sup>2</sup> University of Osnabrück, Osnabrück, Germany  
oliver.thomas@uni-osnabrueck.de

**Abstract.** Implicit feedback collaborative filtering recommender systems suffer from exposure bias that corrupts performance and creates filter bubbles and echo chambers. Our study aims to provide a practical method that does not inherit any exposure bias from the data given the information about the user, the choice, and the choice set associated with each observation. We validated the model’s functionality and capability to reduce bias and compared it to baseline mitigation strategies by simulation. Our model inherited little to no bias, while the other approaches failed to mitigate all bias. To the best of our knowledge, we are first to identify a feasible approach to tackle exposure bias in recommender systems that does not require arbitrary parameter choices or large model extensions. With our findings, we encourage the recommender systems community to move away from rating-matrix-based towards discrete-choice-based models.

**Keywords:** recommender systems, implicit feedback, exposure bias, feedback loop, discrete choice

## 1 Introduction

Recommender systems support users’ decisions from what movies to watch over what items to buy to what to learn next. The industry leaders rely dramatically on recommender systems: Gomez-Uribe et al. (2015) [1] revealed that recommendations contribute to almost 80% of watch time on Netflix. However, what items the system exposes the users to can significantly affect system performance. Collaborative filtering recommender systems operating on implicit feedback (as navigation patterns) rather than explicit feedback (as user-specified ratings) are known to suffer from exposure bias. Exposure bias happens as users are only exposed to specific items so that unobserved interactions do not always represent negative preference [2]. As users are more likely to engage with a few top recommended items [3], the system’s recommendations can bias their behavior. Once the system infers biased knowledge from these observations, it propagates this bias into its recommendations. For example, if a user is overexposed to a particular item not among their top preferences, they will likely still engage with it. Based on the high engagement rate, the system will

overestimate the item’s popularity and keep frequently recommending the item, even if more relevant items enter the domain. A feedback loop of exposure, choice, and inference emerges. Such feedback loops further affect user behavior and amplify pre-existing system biases [4]. The causes of exposure bias include an unfair policy of the previous recommender system [5], new items entering the domain, or even the user’s socio-demographic background [6]. In addition to the definition by Chen et al. (2020) [2], we define exposure bias as any systematic shift in the observed choice behavior due to altering exposure.

Past studies have proven the occurrence of exposure bias and corresponding feedback loops in recommender systems and analyzed their effects. As the need to reduce bias in information systems has moved into the spotlight in recent years [7], researchers have proposed various mitigation strategies [8–13]. However, these approaches are either heuristic, demand large model extensions, or require arbitrary choices of parameters and estimators that can affect performance if chosen incorrectly [2]. Other approaches have only considered a single-user setting [14]. To our knowledge, the recommender systems community lacks models that are naturally resistant to exposure bias. Also, the proposed methods assume no available information about what items a user has been exposed to but not engaged with. Most methods estimate exposure while this information could be retrieved by the system and provided to the model directly. The benefit of providing such information is still unknown.

This study aims to provide a feasible method that does not inherit any exposure bias from the data given the information about the user, the choice, and the choice set associated with each observation. We hypothesize that a variant of collaborative competitive filtering [15] is resistant against exposure bias without any loss of accuracy compared to baseline collaborative filtering algorithms on implicit feedback data. This paper is structured as follows: First, we present the method, elaborate on our design choices, and discuss why the training algorithm should consider the entire choice set. Then, we evaluate our prototype and compare it to baseline models and existing bias mitigation strategies. Our approach differs from past studies because it does not rely on arbitrary parameter choices, large model extensions, or heuristics and takes advantage of the full information about the decision process in a multi-user setting. On systems with sufficient implicit choice data, our approach prevents items from falsely claiming dominance through over-exposure. It is easier to implement than other bias mitigation strategies and utilizes more information than current techniques to improve accuracy.

## **2 Research Approach**

In our interdisciplinary research project KUPPEL, we explore and design new learning environments. Specifically, our research project aims to provide students with learning resources and assign learning peers in an online learning environment. The first analysis of our partners' existing learning environments revealed that they are currently only incorporating implicit feedback. As implicit feedback is prone to exposure bias [2], we conducted a literature review according to Webster and Watson (2002) [16] to obtain an overview of current bias mitigation approaches. For this, we searched AISEL,

EBSCOHost, ScienceDirect, SpringerLink, and Google Scholar databases using the following search string: “recommender system” AND “implicit feedback” AND “exposure bias” The main results of our literature review are summarized in Section 3. It is important to note that Section 3 is not the central focus of this paper and does not aim at systematically covering the entire scientific discourse but rather serves as an argumentative starting point for developing a new prototype. All bias mitigation strategies we encountered rely on large model extensions or arbitrary choices of heuristics and parameters that – if chosen incorrectly – can greatly affect performance [2]. Introducing explicit feedback to the system could result in selection bias and would therefore not have solved the issue. This observation inspired us to develop a prototype that we hypothesized was resistant against exposure bias, and that does not rely on any arbitrary settings. This development process was carried out in an iterative approach and was improved by frequent evaluation cycles. Finally, we compared our prototype to baseline bias mitigation strategies. In this paper, we report on the prototype and the results from that evaluation cycle as we deem the findings relevant for research on exposure bias in recommender systems.

### **3 Related Work**

Past research has shown how bias can corrupt the performance of recommender systems[4]. It can create filter bubbles [16], echo chambers [17] and lead to false consensus effect [18]. In particular, the exposure bias defined in the previous section can lead to the miscalibration of recommendations for groups of users who are less interested in the popular items [19]. Moreover, it affects other parties involved in the recommender systems, for example, content providers. An amplification of popularity of some items and a severe underexposure to alternatives can decrease providers' reputation [20], consequently disincentivizing new providers from entering the market.

Chen et al. (2020) [2] summarized the biases related to recommender systems: selection bias, conformity bias, exposure bias, position bias, inductive bias, popularity bias, and unfairness. For each form of bias, they provide an overview of past research on its occurrence, effects, and mitigation strategies. Research on bias in recommender systems has also focused on the relationship between bias and the user-model feedback loop. Mansoury et al. (2020) [4] show how bias amplifies through the user-model feedback loop. They conclude that bias can shift the users' tastes over time, homogenize recommendations, and therefore neglect the preferences of minority groups. This shift might cause concept drift, making the model unsuitable for the evolved data [21]. Moreover, with a lack of diversity of item exposure and unreliable non-positive data, exposure bias can affect the overall quality of the recommendations long-term. For instance, the resulting imbalance of the inherent audience size can exacerbate popularity bias [22].

Common practices for exposure bias mitigation are heuristic weights [8], [12], [13], sophisticated sampling [9], [10] or exposure-based models [11]. An approach for bias mitigation, similar to the one presented in this paper, has been considered by Çapan et al. (2019) [14], who apply a multinomial logit model to estimate the choice probability

of a single user among a set of items to reduce exposure bias. However, to our knowledge, only the authors’ follow-up paper [23] has been peer-reviewed. In comparison, our model proposes a significant expansion to the method and its evaluation. The authors of [14] and [23] claim that their model can extend to a multi-user setting but do not provide further details. Most importantly, their approach implies exposing the user to every available item at least once, which is impossible in many applied settings. It remains ambiguous whether an extension to a multi-user approach where not every user interacts with every item also prevents exposure bias. Lastly, the authors do not compare their approach to existing bias mitigation strategies. Instead, our approach is naturally suited for a multi-user setting, and we compare proneness against exposure bias to alternative approaches. We are, to the best of our knowledge, the first to propose this approach in the context of exposure bias. To differentiate our results from previous work, Table 1 details the weaknesses of the research discussed above.

**Table 1.** Summary of selected work

| Publication                 | Heuristic or arbitrary design | Large model extension | Single-user Setting |
|-----------------------------|-------------------------------|-----------------------|---------------------|
| Saito et al. (2020) [8]     | ×                             |                       |                     |
| Carraro & Bridge (2020) [9] | ×                             |                       |                     |
| Ding et al. (2019) [10]     |                               | ×                     |                     |
| Liang et al. (2015) [11]    | ×                             | ×                     |                     |
| Pan & Scholz (2009) [12]    | ×                             |                       |                     |
| Pan et al. (2008) [13]      | ×                             | ×                     |                     |
| Çapan et al. (2019) [14]    |                               |                       | ×                   |
| Our work                    |                               |                       |                     |

## 4 Prototype

We assume a system in that a user iteratively receives a choice set out of which they can pick exactly one item for a total of  $n_{\text{users}}$  and  $n_{\text{items}}$ . Let  $O$  be a set of  $m$  observations of past interactions, where each observation is a 3-tuple  $(i, j, C) \in O$  of a user  $i$  and their choice  $j$  out of a choice set  $C$ . For example, if user 5 was presented items 1, 2, and 4, and chose item 2, then  $O$  would contain the observation  $(5, 2, \{1, 2, 4\})$ . Traditional methods would only record that user 5 picked item 2. We constructed a model to estimate which item should be recommended to which user. We designed our prototype to combine the advantages of classic matrix factorization and the approach of Çapan et al. (2019) [14]. Our prototype can also independently be seen as a version of collaborative competitive filtering [15]. Therefore, our work may be interpreted as building on collaborative competitive filtering to show that it provides a practical, exposure bias-free recommendation method, which is the aim of our research.

For our prototype, we defined four major requirements. It should not inherit any bias induced by past recommendations (req. 1), be accurate (req 2.), be scalable as the

number of users, items, and observations increases (req. 3), and incorporate observed user-, item- or context-specific attributes (req. 4). First, we describe our method, and then we discuss how it meets each requirement. For the recommendation phase, we employed a matrix-factorization model that stores latent parameters  $u_i \in \mathbb{R}^k$  for each user  $i$  and  $v_j \in \mathbb{R}^k$  for each item  $j$  for some  $k \in \mathbb{N}$  to estimate the perceived utility  $\mu_{ij}$  for user  $i$  and item  $j$  as

$$\mu_{ij} = u_i^\top v_j.$$

The system recommends the highest utility items to each user. To fit the matrix factorization model, we embedded it into a multinomial logit model that estimates the probability of user  $i$  choosing item  $j$  from the choice set  $C$  as

$$P_{ijc} = \frac{e^{\mu_{ij}}}{\sum_{j' \in C} e^{\mu_{ij'}}}, \quad (1)$$

providing a likelihood function over the observed choices given the matrix factorization model's parameters. The distribution of  $P_{ijc}$  is common in the discrete-choice literature, which assumes a deterministic utility component  $v_{ij}$  and an overall perceived utility

$$\pi_{ij} = v_{ij} + \epsilon_{ij}$$

for *iid*. Gumbel-distributed variables  $\epsilon_{ij}$  [24]. Then, the probability that  $j$  has the highest perceived utility for user  $i$  among all items in  $C$  equals

$$P(\pi_{ij} \geq \pi_{ij'} \quad \forall j' \in C) = \frac{e^{v_{ij}}}{\sum_{j' \in C} e^{v_{ij'}}},$$

which equation (1) is derived from [24]. We define a batch gradient descent rule to adjust the parameters and minimize the negative-logarithmic-likelihood function of the multinomial logit model, which serves as a cost function. Let  $U$  and  $V$  be matrices with row vectors  $u_i$  and  $v_j$  respectively. By writing the minimization problem as

$$\begin{aligned} \min_{U, V} -\ln \left( \prod_{(i,j,c) \in O} P_{ijc} \right) &= \min_{U, V} - \sum_{(i,j,c) \in O} \ln (P_{ijc}) \\ &= \min_{U, V} \sum_{(i,j,c) \in O} \ln \left( \sum_{j' \in C} e^{u_i^\top v_{j'}} \right) - u_i^\top v_j, \end{aligned}$$

one obtains the following update-rules for each for the  $l$ -th coefficients of the user vectors  $u_i$  or item vectors  $v_j$  with  $l \leq k$ :

$$u_{il} \rightarrow u_{il} + \alpha \sum_{\{(i^*, j, c) \in O \mid i^* = i\}} v_{jl} - \frac{\sum_{j' \in C} v_{j'l} e^{\mu_{ij'}}}{\sum_{j' \in C} e^{\mu_{ij'}}$$

$$v_{jl} \rightarrow v_{jk} + \alpha \sum_{\{(i,j^*,C) \in O \mid j \in C\}} u_{il} \delta_{j=j^*} - \frac{u_{ik} e^{\mu_{ij}}}{\sum_{j' \in C} e^{\mu_{ij'}}},$$

where  $\alpha$  denotes a learning-rate factor and  $\delta_{(\cdot)}$  the indicator function. To prevent overfitting, we added an L2-regularization term with a regularization factor  $\lambda$  and adjusted the update rule to

$$\begin{aligned} u_{il} &\rightarrow \frac{(1-\lambda)}{n_{\text{users}}} u_{il} + \alpha \sum_{\{(i^*,j,C) \in O \mid i=i^*\}} v_{jl} - \frac{\sum_{j' \in C} v_{j'l} e^{\mu_{ij'}}}{\sum_{j' \in C} e^{\mu_{ij'}}} \\ v_{jl} &\rightarrow \frac{(1-\lambda)}{n_{\text{items}}} v_{jl} + \alpha \sum_{\{(i,j^*,C) \in O \mid j \in C\}} u_{il} \delta_{j=j^*} - \frac{u_{il} e^{\mu_{ij}}}{\sum_{j' \in C} e^{\mu_{ij'}}}. \end{aligned}$$

We validated our implementation of the gradient descent algorithm by comparing it to the numerically approximated gradient. To break the symmetry between the latent dimensions, the initial parameters follow a uniform distribution over the  $\epsilon$ -sphere for a positive hyperparameter  $\epsilon$ . This fixes the variance of  $\mu_{ij}$  independently of  $k$ . If we intuitively drew the entries of  $U$  and  $V$  as, for example, independent Gaussians, then  $\text{Var}(\mu_{ij})$  would increase in  $k$ . We also implemented optimization through mini-batch gradient descent. Optionally, a bias term can be added by increasing  $k$  by one and fixing all  $u_{i,k}$  to 1 during training. The choice of  $k$  is determined by cross-validation.

The multinomial logit cost function supposedly prevents the model from inheriting bias (req. 1), as it considers which items are included in the choice set and which are not. It acknowledges which alternative was chosen over another and which were observed but not picked. Because items that occur over-proportionally often are also more often decided against by a user, we expect this model no longer favors frequently recommended items. We purposely decided against using a simpler logistic matrix factorization model to maximize

$$\prod_{(i,j,C) \in O} \prod_{j' \in C} \frac{(e^{\mu_{ij}})^{\delta_{j=j'}}}{1 + e^{\mu_{ij'}}}.$$

Such an approach would falsely ignore the dependence of the decisions to pick one item and not to pick another corresponding to the same observation. For example, assume a user  $i$  and five items 1, ..., 5 with utility

$$\mu_{i1} > \mu_{i2} > \mu_{i3} > \mu_{i4} > \mu_{i5}.$$

If we successively present the user with the choice sets  $\{1, 2\}$ ,  $\{2, 3\}$ ,  $\{3, 4\}$  and  $\{3, 5\}$ , they will pick item 2 for 50 % and item 3 for 66% of the opportunity to do so. Due to the biased exposure, a logistic matrix factorization model would ignore the clear decision hierarchy and consider item 3 more relevant than item 2. On the other hand, our model would decrease its cost function if it assigned a higher relevance to item 2 than to item 3. More advanced discrete choice models may as well be considered [24]. The matrix factorization component allows incorporating observed features (req. 2) as

a user’s age or an item’s price by replacing specific components of the latent vectors with observed variables. As matrix factorization techniques are well established in the recommender systems domain, we expected the model to produce accurate recommendations (req. 3). It only accurately estimates the true values of  $U$  and  $V$  –if there are any– up to linear transformations under that the scalar product is invariant. We report accuracy in Subsection 5.2. Scalability in memory space (req. 4) is ensured by the embedded matrix factorization model as its parameter space scales linearly in each component, i.e., with  $\mathcal{O}(k(n_{users} + n_{items}))$  and the training data size scales with  $\mathcal{O}(m\bar{C})$ , where  $\bar{C}$  denotes the mean choice set size. Training the model through batch gradient descent requires a duration of order  $\mathcal{O}(mk\bar{C})$ . However, mini-batch gradient descent decreases the computational cost significantly. More advanced gradient-based algorithms may accelerate convergence further.

## 5 Simulation and Evaluation

### 5.1 Simulation

We validated the model’s functionality and its capability to reduce bias by simulation. For comparison, we employed a k-nearest-neighbors (KNN), logistic matrix factorization (MF), exposure-based matrix factorization (EXMF) [11] and weighted matrix factorization (WMF) [13] model. EXMF and WMF are both designed to mitigate exposure bias. The EXMF approach explicitly models exposure, while the WMF approach reweighs the impact of unobserved user-item interactions. The simulation was repeated for each model and can be structured into four steps:

1. *Setup*: Initially, we chose  $k = 3$  and generated latent vectors  $v_j \in \mathbb{R}^k$  for 25 items and  $u_i \in \mathbb{R}^k$  for 5000 users. The choice of  $k$  was large enough to incorporate a bias variable, and the number of items ensured that no user compared all items. Simultaneously, both parameters were picked small to save computational resources. We chose the high number of users so that the model’s parameters would converge as many users imply many observations. The first two components of each vector were uniformly distributed over the 2-dimensional unit circle. The third component of each user vector  $u_{i,3}$  was set to 1, and the corresponding component of each item vector  $v_{j,3}$  was drawn uniformly from (0,1) to obtain an item-wise bias term. We then divided the items into two disjoint subsets  $I_A$  and  $I_B$  of size 10 and 15 respectively and separated the users into ten disjoint subsets  $N_1, \dots, N_{10}$  of 500 users each. We declared one randomly picked item  $i_{bias}$  from  $I_A$  the “promoted item” that users were later disproportionately often exposed to. Also, we uniformly drew a probability  $p_{bias} \in (0.2, 0.8)$  that determined the degree of unfair exposure.
2. *All users made unbiased choices on  $I_B$* : Each user was then provided with a random choice set of three items from  $I_B$  and picked one item according to the distribution defined in equation (1). We repeated this process four times for a total of five observations per user. For each user, picked items were excluded from the following



choice sets. Based on these unbiased observations, users later received recommendations from  $I_A$ . Overall, we observed five choices per user on  $I_B$ .

3. *Users from  $N_1$  made biased choices on  $I_A$* : As in step 2, each user from  $N_1$  iteratively made choices on three random choice sets of three items according to equation (1). However, the probability that a choice set contained the item  $i_{\text{bias}}$  was set to  $p_{\text{bias}}$  if they had not yet picked  $i_{\text{bias}}$ . As in step 2, any item the user picked was excluded from the following subsets. Last, we trained the model on all data from steps 2 and 3. The overexposure to the promoted item inserted exposure bias into the feedback loop.
4. *Users from  $N_2, \dots, N_{10}$  received recommendations on  $I_A$* : For nine periods, users received recommendations by the model and made choices as in step 3. Step 2 was considered period 1. Period 2 to 10 contained three stages:
  - (a) *Benchmarking on  $I_A$  and  $N_i$* : At the beginning of Period  $i$ , we measured the model’s accuracy as the normalized discounted cumulative gain (nDCG) [25] of the estimated user preferences and the true user preferences for the items from  $I_A$  on  $N_i$  with a discount factor based on the logarithm of base 2. The nDCG serves as a measure for how well an estimated ranking matches the corresponding true ranking, returning positive values up to one. It emphasizes accurately estimating the ranks of the highest preference items, making it a popular choice for evaluating recommender systems. We also measured the difference between the estimated preference rank of the item  $i_{\text{bias}}$  and its true preference rank on  $I_A$  for each user from  $N_i$  to obtain a measure for the model’s bias towards the promoted item. These measures served as an estimate of the model’s performance after period  $i - 1$ .
  - (b) *Users from  $N_i$  received recommendations on  $I_A$* : The model recommended three items to each user from  $N_i$  and each user picked an item according to equation (1). This process was repeated twice following the procedure from step 3 so that each user made three choices in total.
  - (c) *Retraining the model*: The model was retrained on all available data.

The observations in step 4 showed how the exposure bias inserted in step 3 affects the system’s recommendations. We repeated the simulation for 100 iterations and averaged the results to control for outliers. Every iteration was initialized on a different random seed. The seeds were stored to ensure reproducibility.

All four models chosen for comparison received a rating matrix  $R \in \mathbb{R}^{n_{\text{users}} \times n_{\text{items}}}$  with its entries set to 1 if the corresponding user-item choice had been observed and 0 otherwise. Except for the KNN model, all models only considered entries for which observations were possible; in period  $i$ , all observations for users from  $N_{i+1}, \dots, N_{10}$  on  $I_A$  were ignored as these had to be negative. This method ensured that the matrix factorization component did not falsely underestimate the relevance of items from  $I_A$ . The KNN model relied on cosine-similarity. Our implementation of the EXMF model contained a logistic component as the target variables were binary and not normally distributed. We deemed this adaptation necessary as violating the EXMF model’s assumptions might give it an unfair disadvantage. Instead of the m-step from Liang et al. (2015) [11], we optimized the model by mini-batch gradient descent. Exposure was

encoded via item popularity, as the authors suggest. The WMF model downweighed the negative observations proportionally to propensity scores  $\theta_{ij}$  in period  $l$  with

$$\theta_{ij} = \frac{2}{500l} \sum_{m=1}^{n_{\text{users}}} R_{mj}.$$

This way, the zero-entries of the rating matrix weighed less during training if there were few positive entries. As stated earlier, the choice for the propensity factors remains arbitrary. All models used a latent dimension  $k$  of 3, where the last user coefficient was set to 1 to obtain a bias term.


Hyperparameters were selected to maximize the mean nDCG over all periods of five iterations of the simulation. Different seeds to those that the simulation was later run on were used to prevent overfitting. Table 2 lists the selected hyperparameters. The matrix-factorization-based models used a different learning rate and several learning epochs for period 1 than for periods 2-10, as only fine-tuning was required during the later periods. Mini-batch size was set to 32 for all models. For the EXMF model, we initialized  $p_{ij}$  at 0.8,  $\alpha_1$  as 2 and  $\alpha_2$  as 3.

**Table 2.** Hyperparameters

| Model     | $n_{\text{neigh}}$ | Period 1            |          | Periods 2-10        |          |           |
|-----------|--------------------|---------------------|----------|---------------------|----------|-----------|
|           |                    | $n_{\text{epochs}}$ | $\alpha$ | $n_{\text{epochs}}$ | $\alpha$ | $\lambda$ |
| KNN       | 100                | -                   | -        | -                   | -        | -         |
| MF        | -                  | 10                  | 3        | 2                   | 1        | 0         |
| EXMF      | -                  | 10                  | 3        | 2                   | 1        | 3         |
| WMF       | -                  | 15                  | 3        | 5                   | 1        | 1         |
| Our model | -                  | 10                  | 3        | 2                   | 1        | 3         |


## 5.2 Evaluation

Fig. 1 shows that our model inherited little to no bias towards the promoted item while the other models overestimated its relevance. Each line represents the mean difference of the estimated rank and the true rank of the unfairly promoted item by the period of the corresponding model over all iterations. On the last sub-iteration, our model overestimated the promoted item by an average of 0.01 in rank, compared to 1.57 to 2.81 for the other models. Bias in the other models was the highest on the second sub-iteration, from where it decreased until it stagnated.



**Figure 1.** Bias by model

Our model also consistently yielded a higher nDCG, as fig. 2 shows. Each line represents the mean nDCG by the period of the corresponding model over all iterations. On the last sub-iteration, our model achieved 92% nDCG on average compared to 89% to 90%, corresponding to a 20% to 25% smaller error. For comparison, a model recommending random items would have achieved an expected nDCG of about 80%. Among the competing methods, the KNN and MF achieved the highest accuracy. However, WMF overestimated the promoted item by 0.9 ranks less than



**Figure 2.** Accuracy by model

logistic matrix factorization. Our implementation of EXMF performed worse both in accuracy and bias than MF.

## 6 Discussion

We found that our approach did not inherit any exposure bias from the data while all comparative models overestimated the promoted item’s true relevance. Despite its resistance against bias, our model performed most accurately.

We attribute our model’s high bias resistance to its ability to consider an alternative preferred over the promoted item. The other models were incapable of considering this information as they relied on a rating matrix where unobserved interactions could not be distinguished from observed negative ones. Also, as in the example given in Section 4, the other models cannot consider any decision hierarchy. For these reasons, we assume the WMF model could not eliminate all bias from the system. The EXMF model may have also had issues adapting to the varying exposure probability from period to period. The fact that our model includes more information may also have contributed to its higher accuracy. On top of their weaknesses discussed above, the EXMF and WMF model’s low accuracy may result from their additional probabilistic components. Inaccurate estimates of exposure may affect their output and the resulting nDCG. As our model does not require any additional components to model exposure as the EXMF and WMF models do, we found our model easier to implement.

Our findings are subject to key limitations. We only validated and compared our model through a simulation that fulfilled the model’s assumptions. Therefore, we cannot derive any prospects on its performance on real-world data. We expect our model to work well as the underlying principles have been a common choice in the recommender systems and discrete choice domain, but further experiments on real-world data are required. Also, in real-world applications, users often choose between more than three alternatives. They might then select the first item of interest before comparing the entire set available. Future research needs to investigate the effect of larger choice sets on model performance. The multinomial choice model we employed has been criticized for its so-called assumption of independence of irrelevant alternatives [24]. More advanced discrete choice models that do not make this assumption may be more appropriate. In this study, we chose the simpler multinomial logit model as it sufficed to prove the concept. As the two bias mitigation approaches tested for comparison rely on arbitrary initialization parameters or estimators, other choices for these components might have been more successful. However, the reliance on arbitrary parameter choices has been the main point of criticism about these methods. The poor performance of the EXMF model should be investigated. However, as our model showed almost no bias, we do not expect a revised EXMF model to perform better. The EXMF model also cannot be expected to perform more accurately in the simulation as it does not model the underlying assumptions as precisely.

Several extensions to our approach are possible. For example, by setting a specific component of the user vectors to an observed attribute as the users’ age, the effect of this attribute on preference can be directly observed via the corresponding components

of the item vectors. Such procedure is common in the discrete choice literature and would render the model's output interpretable as is desired in modern AI-enabled systems [7]. Also, other models as neural networks can replace the scalar product to allow for more complex interactions of the user and item vectors. Finally, researchers need to evaluate our model's performance on real-world data. For implementation appropriate choice sets will need to be defined. In real-world applications, a choice set may correspond to the items shown on the user's screen. Items that the user cannot observe because they e.g., require browsing to the next page of search results, should be excluded.

Our findings imply that the recommender systems community should move away from models that rely on rating matrices as these cannot capture the full information about the users' decision process. Accordingly, researchers should focus less on mitigating exposure bias in these models. Instead, we need to investigate how discrete choices should be modeled for a recommendation. Comparing the various discrete choice models [24] in a recommendation setting is mandatory. We also require efficient and accurate mechanisms to identify the choice set a user encountered prior to their choice. Our approach allows practitioners to avoid exposure bias in implicit feedback collaborative filtering systems like e-commerce, digital entertainment [26], or online learning environments [27]. For implementation, systems must begin tracking and storing all information about users' decision processes. Because companies already possess large treasures of user behavior that lack choice set data, implementing our model requires measures to reconstruct the choice sets associated with past choices.

## **7 Conclusion**

In this study, we identified a method that does not inherit any exposure bias given full information about the users, the choices, and the choice sets associated with previous observations. We showed by simulation that this method is resistant against exposure bias without any disadvantage in accuracy compared to baseline collaborative filtering algorithms and other bias mitigation strategies. To the best of our knowledge, our approach is the first to tackle exposure bias in a multi-user setting that does not require any arbitrary parameter choices or large model extensions. Moreover, we have illustrated that the choice set presented to a user can play a crucial role in providing accurate, unbiased recommendations. We conclude that current systems need to start monitoring which items a user encountered in their decision process. With our findings, we encourage the recommender systems community to move away from rating-matrix-based towards discrete-choice-based models.

## **8 Acknowledgments**

This contribution originates from the research project KUPPEL (KI-unterstützte plattformübergreifende Professionalisierung erwachsenen-pädagogischer Lehrkräfte), funded by the Federal Ministry of Education and Research (BMBF), ref. no. 21INVI0803.

## References

1. Gomez-Uribe, C.A., Hunt, N.: The netflix recommender system: Algorithms, business value, and innovation. *ACM Trans. Manag. Inf. Syst.* 6, 1–19 (2015).
2. Chen, J., Dong, H., Wang, X., Feng, F., Wang, M.-C., He, X.: Bias and Debias in Recommender System: A Survey and Future Directions. *ArXiv. abs/2010.03240*, (2020).
3. Agarwal, A., Hosanagar, K., Smith, M.: Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets. *J. Mark. Res.* 48, (2008).
4. Mansoury, M., Abdollahpouri, H., Pechenizkiy, M., Mobasher, B., Burke, R.: Feedback loop and bias amplification in recommender systems. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. pp. 2145–2148 (2020).
5. Liu, D., Cheng, P., Dong, Z., He, X., Pan, W., Ming, Z.: 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*. pp. 831–840. Association for Computing Machinery, New York, NY, USA (2020).
6. Chen, J., Wang, C., Zhou, S., Shi, Q., Feng, Y., Chen, C.: Samwalker: Social recommendation with informative sampling strategy. In: *The World Wide Web Conference*. pp. 228–239 (2019).
7. Maedche, A., Legner, C., Benlian, A., Berger, B., Gimpel, H., Hess, T., Hinz, O., Morana, S., Söllner, M.: AI-Based Digital Assistants: Opportunities, Threats, and Research Perspectives. *Bus. Inf. Syst. Eng.* 61, 535–544 (2019).
8. Saito, Y., Yaginuma, S., Nishino, Y., Sakata, H., Nakata, K.: Unbiased Recommender Learning from Missing-Not-At-Random Implicit Feedback. *Proc. 13th Int. Conf. Web Search Data Min.* (2020).
9. Carraro, D., Bridge, D.: Debaised Offline Evaluation of Recommender Systems: A Weighted-Sampling Approach. In: *Proceedings of the 35th Annual ACM Symposium on Applied Computing*. pp. 1435–1442. Association for Computing Machinery, New York, NY, USA (2020).
10. Ding, J., Quan, Y., He, X., Li, Y., Jin, D.: Reinforced Negative Sampling for Recommendation with Exposure Data. In: *IJCAI* (2019).
11. Liang, D., Charlin, L., McInerney, J., Blei, D.M.: Modeling user exposure in recommendation. *25th Int. World Wide Web Conf. WWW 2016*. 951–961 (2016).
12. Pan, R., Scholz, M.: Mind the gaps: Weighting the unknown in large-scale one-class collaborative filtering. In: *Proc. 15th AMC Co Nf. Knowledge Discovery and Data Mining*. pp. 667–676 (2009).
13. Pan, R., Zhou, Y., Cao, B., Liu, N., Lukose, R., Scholz, M., Yang, Q.: One-Class Collaborative Filtering. Presented at the August (2008).
14. Çapan, G., Gündoğdu, İ., Turkmen, A.C., Sofuoğlu, Ç., Cemgil, A.: A Bayesian Choice Model for Eliminating Feedback Loops, (2019).
15. Yang, S.-H., Long, B., Smola, A.J., Zha, H., Zheng, Z.: Collaborative competitive filtering: learning recommender using context of user choice. In: *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. pp. 295–304 (2011).
16. Chaney, A.J., Stewart, B.M., Engelhardt, B.: How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. *Proc. 12th ACM Conf. Recomm. Syst.* (2018).

17. Jiang, R., Chiappa, S., Lattimore, T., György, A., Kohli, P.: Degenerate Feedback Loops in Recommender Systems. Proc. 2019 AAAI/ACM Conf. AI, Ethics, Soc. (2019).
18. Luzsa, R., Mayr, S.: False consensus in the echo chamber: Exposure to favorably biased social media news feeds leads to increased perception of public support for own opinions. *J. Psychosoc. Res.* 15, (2021).
19. Abdollahpouri, H., Mansoury, M., Burke, R., Mobasher, B.: The Connection between Popularity Bias, Calibration, and Fairness in Recommendation, (2020).
20. Abdollahpouri, H., Mansoury, M.: Multi-sided Exposure Bias in Recommendation. In: International Workshop on Industrial Recommendation Systems in conjunction with ACM KDD 2020 (2020).
21. Baier, L., Hofmann, M., Kühl, N., Mohr, M., Satzger, G.: Handling concept drifts in regression problems - The error intersection approach. Proc. 15th Int. Conf. Bus. Inf. Syst. 2020 "Developments, Oppor. Challenges Digit. WIRTSCHAFTSINFORMATIK 2020. (2020).
22. Zhu, Z., He, Y., Zhao, X., Caverlee, J.: Popularity Bias in Dynamic Recommendation. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 2439–2449 (2021).
23. Çapan, G., Bozal, Ö., Gündoğdu, İ., Cemgil, A.T.: Towards Fair Personalization by Avoiding Feedback Loops, (2020).
24. Train, K.E.: Discrete choice methods with simulation. Cambridge university press (2003).
25. Järvelin, K., Kekäläinen, J.: Cumulated Gain-Based Evaluation of IR Techniques. *ACM Trans. Inf. Syst.* 20, 422–446 (2002).
26. Jannach, D., Lerche, L., Zanker, M.: Recommending Based on Implicit Feedback. In: Brusilovsky, P. and He, D. (eds.) *Social Information Access: Systems and Technologies*. pp. 510–569. Springer International Publishing, Cham (2018).
27. Gössling, H., Dreesbach, T., to Brinke, J.S., Thomas, O.: Smart Cooperative Distance Learning: Ein multiagentenbasierter Ansatz zur Unterstützung des kooperativen Lernens im Distanzunterricht. *HMD Prax. der Wirtschaftsinformatik.* 58, 870–883 (2021).