

# Unmasking Gender Bias in Recommendation Systems and Enhancing Category-Aware Fairness

Tahsin Alamgir Kheya  
t.kheya@deakin.edu.au  
Deakin University  
Geelong, VIC, Australia

Mohamed Reda Bouadjenek  
reda.bouadjenek@deakin.edu.au  
Deakin University  
Geelong, VIC, Australia

Sunil Aryal  
sunil.aryal@deakin.edu.au  
Deakin University  
Geelong, VIC, Australia

## Abstract

Recommendation systems are now an integral part of our daily lives. We rely on them for tasks such as discovering new movies, finding friends on social media, and connecting job seekers with relevant opportunities. Given their vital role, we must ensure these recommendations are free from societal stereotypes. Therefore, evaluating and addressing such biases in recommendation systems is crucial. Previous work evaluating the fairness of recommended items fails to capture certain nuances as they mainly focus on comparing performance metrics for different sensitive groups. In this paper, we introduce a set of comprehensive metrics for quantifying gender bias in recommendations. Specifically, we show the importance of evaluating fairness on a more granular level, which can be achieved using our metrics to capture gender bias using categories of recommended items like genres for movies. Furthermore, we show that employing a category-aware fairness metric as a regularization term along with the main recommendation loss during training can help effectively minimize bias in the models' output. We experiment on three real-world datasets, using five baseline models alongside two popular fairness-aware models, to show the effectiveness of our metrics in evaluating gender bias. Our metrics help provide an enhanced insight into bias in recommended items compared to previous metrics. Additionally, our results demonstrate how incorporating our regularization term significantly improves the fairness in recommendations for different categories without substantial degradation in overall recommendation performance.

## CCS Concepts

• **Computing methodologies** → **Artificial intelligence**; • **Information systems** → **Personalization**; **Recommender systems**.

## Keywords

Fairness Metrics, Bias in Recommendations, Societal Stereotypes, Recommender System

## ACM Reference Format:

Tahsin Alamgir Kheya, Mohamed Reda Bouadjenek, and Sunil Aryal. 2025. Unmasking Gender Bias in Recommendation Systems and Enhancing Category-Aware Fairness. In *Proceedings of the ACM Web Conference 2025 (WWW '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3696410.3714528>



This work is licensed under a Creative Commons Attribution-NoDerivatives 4.0 International License.

WWW '25, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s).

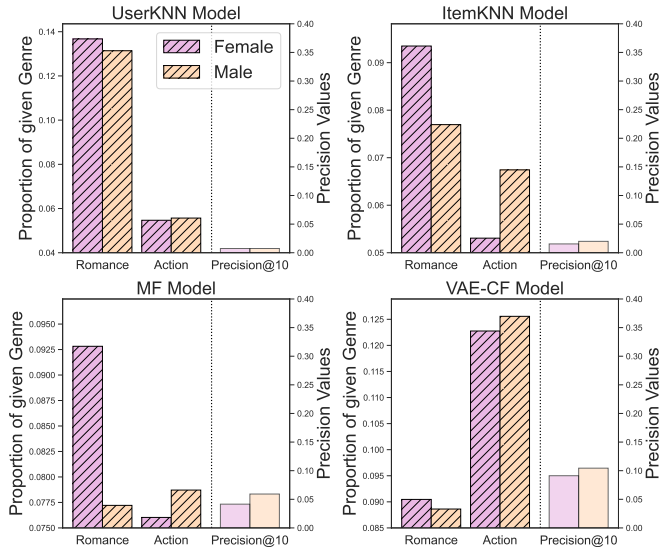
ACM ISBN 979-8-4007-1274-6/25/04

<https://doi.org/10.1145/3696410.3714528>

## 1 Introduction

Recommender Systems (RS) personalize item selections for users, providing suggestions based on individual preferences and behaviors. These systems have an important impact on our decision-making, as they are widely employed across diverse platforms like e-commerce, social media, streaming services, and news outlets, shaping the content and products we encounter. For several online platforms, recommendation systems help create an engaging experience for users by diversifying and personalizing the content and interactions. This would help users avoid information overload and help them focus on options that reflect their past behaviors. Amid the promise held by RS, however, there are concerns about potential bias in these systems. For example, research has shown that, on certain recommender systems, simply changing the gender in a job search—while keeping qualifications constant—can significantly influence access to high-paying positions [11, 33].

RS algorithms are traditionally evaluated using measures like RMSE (Root Mean Squared Error), NDCG (Normalized Discounted Cumulative Gain), precision, and diversity [61]. If only these metrics are considered, then the recommended items are deemed good when they align with user preferences. These metrics can help evaluate the performance of the RS, but in recent years there is a growing emphasis on evaluating and ensuring fairness as well [1, 2, 16, 23, 43, 53]. For instance, the authors in [13] introduce a fairness evaluation metric that can be sensitive to different fairness notions like user-centric or item-centric. Although substantial work has been done in this field in recent years, there is a notable gap in research specifically focused on robust ways to quantify consumer bias accurately. Most of the metrics used are deficient in the following ways: (i) they over-simplify the concept of fairness, (ii) they fail to consider the ranking of recommended items, and (iii) they rely exclusively on a single type of fairness metric. A popular approach to evaluating consumer-side fairness in recommendation systems involves an adaptation of *equal opportunity* (refer to Section: 2.2), which aims to balance performance metrics (such as recall) or utility scores across different groups, such as male vs. female users [43, 51, 52, 54, 63, 75]. However, while this approach is quite straightforward, it may overlook disparities in recommendations across different item categories. To assess these disparities in recommendations, we refer to Figure 1, where we present, for various recommendation algorithms, an analysis of the proportion of action and romance movies among the top 10 recommendations for male and female user groups, along with the corresponding Precision@10 values for each group. There are three notable observations here: (i) across all models, there is a noticeable disparity in the proportion of romance and action movies recommended, as romance movies tend to be recommended more



**Figure 1: Comparison of action and romance movie recommendations among male and female users across four recommendation algorithms, along with corresponding Precision@10 values. The graphs highlight disparities in genre recommendations by gender, with romance movies being more frequently suggested to female users and action movies to male users, despite similar Precision@10 metrics.**

frequently to female users, while action movies are more often recommended to male users; (ii) precision@10 values for both male and female users are similar across all models, suggesting that the models appear to perform equally well for both genders based on this metric; and (iii) different models exhibit varying levels of bias. These observations suggest that simply comparing performance metrics across sensitive attributes is insufficient to assess fairness in recommendation systems. More granular metrics are necessary to accurately quantify bias and ensure fair recommendations across different user groups. To further emphasize the significance of granular evaluation, we provide a brief overview here with a more detailed example provided in Section 3.1. Bias can arise in recommender systems due to stereotypical interactions in the dataset used to train them, which leads to biased recommendations based on sensitive attributes of the users. This bias can trap users in filter bubbles, with a focus on stereotypical categories or narrowly defined preferences. Providing more of a certain type of content can have unintended consequences; for instance, if a young male user is predominantly recommended action movies, where violence is glorified and shown in a consequence-free way, then it could contribute to adverse effects, including the desensitization of violence. This is just one example, but the stakes are much higher for more sensitive domains like news, job recommendations, etc., where there can be catastrophic consequences if similar biases manifest in them.

In this paper, our focus is to overcome the limitations of current fairness assessment metrics by designing a set of evaluation measures to help us quantify gender bias in recommendation models.

Specifically, our aim is to ensure that they are diverse so as to capture nuances of the bias that might not be immediately obvious but are significant. We propose to do this by incorporating item categories and rank (when relevant). Next, we utilize one of these metrics as part of the loss function to optimize, which results in an effective increase of fairness for recommendations made to users for each category. This not only shows that our metrics are effective in addressing fairness concerns but also proves their utility for quantifying bias. Additionally, we use our metrics to evaluate a variety of recommendation models, including fairness-aware ones, using three real-world datasets. The contribution of this paper can be summarized as follows:

- We introduce a set of metrics that considers both categories and rankings of the recommendations made, in turn providing a more nuanced evaluation.
- We evaluate different recommendation algorithms to check for gender bias using the proposed set of metrics.
- We employ an in-processing technique to ensure the recommended items are fair, not only in a general sense but in a category-aware sense.

## 2 Related work

We categorize and discuss the related work existing in this field in the subsections below.

### 2.1 Gender Fairness in RS

Gender bias in recommendation systems can exist as systematic discrepancies in the algorithms when recommending items to users of different genders. The challenges of gender bias can have a wide-ranging set of implications, including imbalanced representations of items for different genders, stereotypical recommendations (male-dominated occupations recommended to males more than females), and limitations in user personalization (being recommended items that are not related to users' preferences just because they are male or female). Research on evaluating and addressing gender bias in recommendation systems is an ongoing process. Melchiorre et al. [52] investigate the impact of a common de-biasing strategy called resampling on RS algorithms. This strategy marginally decreases gender bias, with a slight decrease in performance. The authors in [27, 74, 79] introduce ways to re-rank items to offer a balanced solution that caters to group fairness (for gender and other sensitive attributes) and user preferences. Historical data that can contain stereotypical movie preferences can intensify certain biases further when used to train traditional recommendation models. For instance, male users display a bias towards *action* movies, which is amplified by recommendation algorithms like UserKNN [64]. The authors in [3] introduce a framework that fairly predicts the quality of TedTalk speeches by using causal models and counterfactuals to mitigate gender and racial bias. Adversarial fairness where the recommender system is trained to not only make accurate predictions but also make it difficult for the adversary to guess the sensitive attribute, has also been employed to make such systems more fair [48, 56]. Recent advancements in this field, have led to the development of fairness-aware recommendation models, with special emphasis on gender bias [5, 68, 73, 76, 80, 81].

## 2.2 Evaluating Gender Bias

For evaluating gender bias, the most common fairness definitions employed are the concepts of *Demographic Parity* and *Equal Opportunity*, which are both related to group fairness. Fairness in this context, pertains to equitable treatment across different groups (which can be measured in classification and recommendation tasks). For group fairness, the idea is to ensure that the predicted outcomes  $\hat{Y}$  of a model should not be dependent on sensitive attributes like gender  $S$ . For demographic parity, the proportions of each sensitive group (like male and female) receiving positive predictions should be equal. For binary classification, demographic parity can be formalized as:

$$P(\hat{Y} = 1|S = 1) = P(\hat{Y} = 1|S = 0)$$

Essentially what this implies is that the positive outcome should be the same for both genders, where 0 may represent male and 1 may represent female or vice versa. Equal Opportunity, on the other hand, holds when the model has equal true positive rates across different demographic groups [30]. This concept can be formalized as:

$$P(\hat{Y} = 1|S = 1, Y = 1) = P(\hat{Y} = 1|S = 0, Y = 1)$$

where  $Y$  represents the true outcome.

Besides these two methods to quantify fairness, some additional concepts (as discussed in [40]) used include Equalized Odds [30], Balance for Negative Class [65], Balance for Positive Class [65], Intersectional Fairness [26], Equal Calibration [10] and Causal-based notions [3, 42].

## 2.3 Evaluating Consumer-Side Fairness

In recommendation systems, fairness can be seen as a multi-sided concept and categorized into three groups: consumers (C-fairness) [13, 15, 25, 27, 29, 52, 66, 71] which is related to the impact of recommendations of the system on protected classes of user, provider (P-fairness) which focuses on ensuring fairness for providers/sellers on a platform and both (CP-fairness) [8]. Our work focuses on the consumer-side fairness concept because we want to ensure that recommendations made by models are not biased against a certain gender. Prior research has shown how recommendations can differ in an unfair way based on sensitive attributes of users like gender, age, race, etc. [28, 38, 67]. Evaluating bias in these systems, before deploying is thus essential to stop the reinforcement of stereotypes and limiting diverse content. When quantifying consumer side bias in recommendation systems, the most common approach is to adopt the concept of equality of opportunity and focus on the differences in metrics such as recall, precision and/or NDCG [21, 43, 52]. Other papers also employ causal-based fairness notions [46, 68] and demographic parity [5, 18, 24]. Unlike these metrics, which assume fairness implies equality, [13] suggests how, for instance, paid users should be provided better recommendations when compared to free users. They design a set of metrics that can take this disparity into account and then measure fairness accordingly. Another interesting way to measure unfairness is the concept of envy-free fairness, which is achieved when no one user prefers another user's recommendations way more significantly than their own [14]. While these metrics can identify consumer-side bias to an extent, they come with some limitations.

Limitations of current metrics used to quantify consumer-side bias in recommendation systems include: (i) over-simplifying the meaning of fairness in RS, which employs various techniques like collaborative and content-based filtering and hybrid methods. Simple notions can fail to capture disparities that exist across different types of items, for example, genres, when considering movie recommendations. Some metrics that can fall prey to this oversimplification issue include [13, 14, 19–21, 34, 43, 52, 69, 70, 77]; (ii) not utilizing ranks when evaluating recommendation quality can yield considerable issues. This is due to the fact that recommendations are displayed one after another, so the items on higher ranks must be more relevant to keep the user satisfied. Thus, capturing the quality of recommendations using the ranks reflects a more complete way of evaluating models. Some metrics that can fall prey to this issue include [13, 14, 69, 70, 77]. It is important to state however, that even considering rank can give rise to positional bias, which refers to the tendency of users to favor the items that appear on top of a ranked list. Evaluating till a certain position like Recall@k can lead to a skewed sense of assessment of how the recommendation model performs. Additionally, metrics like MAP (Mean Average Precision), which normally treats relevance as a binary value (0: not relevant and 1: relevant) can also fail to capture the nuanced relevance that can come from items having multiple categories. So, using more than one metric (both with and without using ranks) to quantify bias is essential; (iii) relying only on one type of fairness metric can obscure underlying biases and give a false impression of fairness in recommendation systems.

Hence, using multiple metrics can help uncover hidden biases. Additionally, a model that is fair according to one metric can fail to hold other fairness metrics and risk overlooking subtle unfairness issues.

We want to highlight some of the works that have taken into account different classes when evaluating recommendations [24, 36, 37, 47, 58, 60, 62, 72]. For instance, [47] groups users on certain attributes and items by category, then measures preference ratio, which is the fraction of liked items by a group across categories. Next, they measure the bias disparity by taking the preference and recommended ratios' relative differences. This is close to our work but still doesn't account for the ranks of items and we evaluate the direct comparison of the recommendations for males and females. Additionally, [62] introduce calibrated recommendations, ensuring the recommended items align with user preferences without overemphasizing particular categories. The work by [24] uses a measure *Skew@k* to evaluate proportions of candidates based on sensitive attributes, and [36] uses a fairness metric called Attention Weighted Ranked Fairness (AWRF) [58] to ensure there is balance in exposure in different groups of providers. While both these works ensure group fairness, our work is more concentrated on evaluating the distribution of content categories for different groups. Unlike the work by [36] that focuses on provider-side fairness, we focus on consumer-side fairness.

## 3 Proposed Evaluation Metrics

### 3.1 Motivating Fairness Concern

Our example is a typical offline setting recommendation system, which is trained using historical user and movie interactions. For

our scenario, let us assume we have  $u_1$ , a male user who has watched numerous action and sci-fi movies, with some romance movies. We also have  $u_2$ , a female user who has watched a lot of drama movies but also a few action movies. Let us say we decide to calculate overall precision for each user group (male and female) and then compare them to ensure the model's fairness.

**3.1.1 Potential Issue.** Machine learning systems tend to learn and amplify bias from the training data [4, 12, 17, 40, 49]. Recommender systems are no different, as they can pick up on stereotypical user-item interactions and make biased recommendations based on sensitive attributes of users [52, 64]. Biased recommendations can have a broad impact on users if the models recommend content just because it aligns with certain stereotypes, for instance, males like action movies, and females like romance movies.

This imbalance in recommendation can lead to a situation in which users are only exposed to items that align with part of their preferences and stereotypical norms. This would potentially filter out diverse content and prevent users from discovering new movies. As time passes,  $u_1$  might stop getting romance movies recommended to them, even though they enjoy them. For  $u_2$ , a similar case could arise for action movies. Essentially, this imbalance can trap the users in a bubble of recommendations with only their established preferences and gender-stereotypical genres. Additionally, if a user's established preferences are already aligned with gender stereotypes, then the bias in recommendation will intensify further giving rise to a filter bubble, with very redundant movie recommendations.

**3.1.2 Falling short when quantifying gender bias.** Moreover, as mentioned earlier, using some performance metrics for both genders and comparing them to evaluate fairness is not a great idea. The two groups can have similar scores, even if the model is making biased recommendations by choosing to neglect certain categories for certain users. Our proposed metrics address this issue by breaking down the recommendations by category to get a more nuanced sense of fairness.

The key takeaway here is the importance of learning fair user and item representations, as well as evaluating fairness in the output. Even if a model is trained only on user-item-rating interactions with no explicit mention of sensitive attributes, the model can still infer this private information due to the correlation between their behavior and their sensitive attributes [6, 9, 22]. So, we have to take precautionary measures when training the model itself, so it is unable to learn these correlations. We also want to discuss how it is important to ensure personalization, including gender-specific preferences to enhance user satisfaction, but such preferences should be balanced against any risk of reinforcing stereotypes.

Please note in our study we focus on binary genders, acknowledging there are many other gender identities not represented here.

## 3.2 Notation

We present all metrics for fairness assessment using the following mathematical notation:

- $u_i$ : A single user, where  $i$  indexes the users.
- $v_j$ : A single item, where  $j$  indexes the items.
- $\mathcal{U}$  and  $\mathcal{V}$ : The set of users and items, respectively.

- $\mathcal{U}_m$  and  $\mathcal{U}_f$ : The set of male and female users, respectively.
- $c$ : An item category, such as Action, Sci-Fi, Romance, etc.
- $C$ : A category matrix where  $C_{j,c} = 1$  if category  $c$  is associated with item  $v_j$ , and  $C_{j,c} = 0$  otherwise.
- $C_{v_j}$ : The list of categories associated with item  $v_j$ .
- $TopK_{u_i}$ : The set of top  $K$  recommended items for user  $u_i$ .
- $\mathcal{C}$ : Represents the set of categories for items.

To assess fairness, we adapt and extend Information Retrieval metrics, introducing both non-ranking and ranking-based metrics, which are detailed in the following subsections.

## 3.3 Non-ranking-based metrics

The first set of metrics we propose evaluates the fairness of a recommender system without considering the ranking of movies.

**3.3.1 Category Coverage (CC).** This metric estimates the proportion of recommended items associated with category  $c$  relative to all categories for all users  $u_i$ . This essentially captures the category-specific performance of the recommended items. It is defined as follows:

$$CC(c, \mathcal{U}) = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{1}{|TopK_{u_i}|} \sum_{v_j \in TopK_{u_i}} \frac{C_{j,c}}{|C_{v_j}|} \quad (1)$$

**3.3.2 Relative Category Representation (RCR).** This metric estimates the proportion of category  $c$  in recommended items relative to the proportions of all categories available from the whole dataset. This helps provide insights on category-specific items and is defined as follows:

$$RCR(c, \mathcal{U}) = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{\sum_{v_j \in TopK_{u_i}} \frac{C_{j,c}}{|C_{v_j}|}}{\sum_{v_k \in \mathcal{V}} \frac{C_{k,c}}{|C_{v_k}|}} \quad (2)$$

These two metrics help us quantify how relevant a category is for items recommended to a given set of users. For instance,  $CC$  would help us quantify the diversity of items recommended by reflecting the overall distribution of recommended content when computed for different  $c$ . Whereas,  $RCR$  (calculated for different  $c$ ) would measure if the recommendations suppress or amplify certain categories by comparing them to the actual distribution of available content. Both  $CC$  and  $RCR$  provide a more nuanced understanding of how the recommendation system performs but don't take into account the position of the recommended items. For recommendation systems, where in most cases items surface one after another, it is vital that the items on the top of the menu are the most relevant. Hence, to ensure that ranking order is also considered for evaluating recommended items, in the next section we introduce rank-based metrics.

## 3.4 Rank-based Metrics

We now introduce our metrics that utilize ranking to provide a comprehensive assessment of fairness.

**3.4.1 Category Mean Average Precision (CMAP).** For a given set of users, this metric estimates the proportion of recommended items associated with category  $c$  by incorporating the rank of the items. A category-specific average precision score is computed for each user,

followed by averaging these scores across all users. This metric is defined as follows:

$$CMAP(c, \mathcal{U}) = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{\sum_{j=1}^{|TopK_{u_i}|} P(j) \cdot \frac{C_{j,c}}{|C_{v_j}|}}{\sum_{v_j \in \mathcal{V}} \frac{C_{j,c}}{|C_{v_j}|}} \quad (3)$$

where  $P(j) = \frac{1}{j} \sum_{k=1}^j \frac{C_{k,c}}{|C_{v_k}|}$ .

**3.4.2 Category Discounted Cumulative Gain (CDCG).** The score for this metric is discounted based on the position of the item in the recommended list. While similar to the previous metric, CDCG uses a logarithmic discount factor, which is higher for items that appear lower down in the list. It is defined as follows:

$$CDCG(c, \mathcal{U}) = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{1}{|TopK_{u_i}|} \sum_{j=1}^{|TopK_{u_i}|} \frac{\frac{C_{j,c}}{|C_{v_j}|}}{\log(j+1)} \quad (4)$$

**3.4.3 Category Mean Reciprocal Rank (CMRR).** This metric is based on MRR, which is essentially the mean reciprocal of the rank of the first relevant item. We adapt it for assessing fairness and we define it as follows:

$$CMRR(c, \mathcal{U}) = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{1}{|TopK_{u_i}|} \sum_{j=1}^{|TopK_{u_i}|} \frac{\frac{C_{j,c}}{|C_{v_j}|}}{j} \quad (5)$$

**3.4.4 Category RPrecision (CRP).** This metric uses the category-specific precision but only for the top  $\lfloor R_c \rfloor$  results, where  $R_c$  represents the proportion of category  $c$  when considering all items in  $\mathcal{V}$  (i.e.,  $R_c = \sum_{j=1}^{|\mathcal{V}|} \frac{C_{j,c}}{|C_{v_j}|}$ , note that item  $v_j$  can belong to multiple categories like movie genres). This metric is defined as follows:

$$CRP(c, \mathcal{U}) = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{\sum_{j=1}^{\lfloor R_c \rfloor} \frac{C_{j,c}}{|C_{v_j}|}}{\lfloor R_c \rfloor} \quad (6)$$

All four of the ranking-based metrics, take into consideration the proportions of categories and use rank as a discounting factor. The discounting factor is different for each of them and has distinct nuances. For instance, *CRP* would only evaluate recommendations till a certain rank  $\lfloor R_c \rfloor$  to measure if the model is providing a proportional amount of recommendations for  $c$  relative to its presence in the dataset. Another metric like *CDCG* for example would capture something different, which is the relevance of  $c$  by giving more weight to items (falling under category  $c$ ) appearing higher in the list.

## 3.5 Group Fairness (Gender)

Gender Balance Score (GBS) for a given category  $c$ , is the absolute difference of category distribution values for  $\mathcal{U}_m$  and  $\mathcal{U}_f$  based on any of the category-aware metrics discussed above. This metric is an adaptation of *demographic parity* as mentioned in Section 2.2. We want to make sure that the category distributions of the recommended items for the groups, of males and females are similar. For each metric  $\mathcal{M}$ , defined above, we want to ensure that the difference of values calculated for the groups male and female

summed up for all categories are close to 0. This can be formalized as:

$$\Delta \mathcal{M}_c = |\mathcal{M}(c, \mathcal{U}_m) - \mathcal{M}(c, \mathcal{U}_f)| \quad (7)$$

$$GBS(\mathcal{M}) = \sum_{c \in \mathcal{C}} \Delta \mathcal{M}_c \approx 0 \quad (8)$$

## 4 Genre Aware Regularization For Gender Fairness

Now that we have introduced our set of metrics for capturing a nuanced sense of fairness for recommended items, we aim to learn fair representations that are category-aware for users of different genders. To promote fairness in recommendation models we employ an in-processing regularization technique inspired by the approach first proposed by [78]. We propose to incorporate Category Coverage (Section 3.3.1) into the loss function alongside the primary recommendation loss, encouraging the model to optimize for category-aware fairness. We want to emphasize that while we select *CC* as an example here, optimizing on any of the other metrics (as long as the implementation is differentiable) is possible. We incorporate this fairness regularizer into the loss function of each baseline model—MF, VAE-CF, and NeuMF.

Specifically, to discourage the model from learning any biased representations, we incorporate the following regularizer:

$$\mathcal{L}_{FairGenreGender} = GBS(CC) = \sum_{c \in \mathcal{C}} |CC(c, \mathcal{U}_m) - CC(c, \mathcal{U}_f)| \quad (9)$$

The goal is to minimize the difference in category distribution of the items being recommended to the users of each sensitive group. The new loss function for MF, VAE-CF and NeuCF models can be written as:

$$\mathcal{L} = \alpha \mathcal{L}_{FairGenreGender} + (1 - \alpha) \mathcal{L}_{Recommendation} \quad (10)$$

where  $\alpha$  is used to calibrate the trade-off between the model loss and the fairness term.  $\mathcal{L}_{Recommendation}$  represents the respective recommendation loss term of each model. By tuning this hyperparameter, we ensure that fairness does not overly compromise the recommender's ability to respect users' personal preferences (even if there are natural differences across liked categories for users of different genders). In order to modulate the gender loss value and to ensure it is able to make a significant contribution to the loss function, we pass it through a sigmoid function, centering at 0.5 and scaling it by 0.1. Additionally, the maximum batch loss value of the actual recommendation loss is used to scale the gender loss up to ensure both losses are on the same scale. In the next section we show how using this fairness regularization helps the models minimize bias in recommendations provided to users of different genders for different categories.

## 5 Experiments

In the sub-sections below, we outline our experimental setup.

### 5.1 Datasets

Our experiments are performed on three recommendation datasets summarized in Table 1. The data is split into training, validation, and test sets with a 70:10:20 ratio following user-based split scheme.

**Table 1: Description of datasets.**

Dataset	Users	Item	Interactions	Categories for Items
ML 100K [31]	943	1,349	99,287	18
ML 1M [31]	6,040	3,416	999,611	18
Yelp [50]	1,316	1,272	97,991	21

It’s important to highlight that items can fall under multiple categories. For clarity and ease of understanding, we focus on four representative categories—*Action*, *Romance*, *Sci-Fi*, and *Drama* from the MovieLens datasets and *Coffee*, *Tea & Desserts*, *Arts & Entertainment*, *Travel & Transportation* and *Asian* from the yelp dataset – when visualizing our metric values; however all categories are included in our experiments.

## 5.2 Evaluating Performance

For measuring the performance of recommendations made, we use HitRatio@k calculated for each user and then averaged. Additionally, we use a ranking-based metric NDCG@k (Normalized Discounted Cumulative Gain) which gives higher relevance to items appearing higher in the ranked list. This too, is calculated for each user using their top  $k$  recommendations and then averaged. We calculate these values for  $k = 50$ .

## 5.3 Base Recommendation Models

For analyzing which models manifest gender bias we evaluate several recommendation approaches. When selecting these models, we include a variety of algorithms, ranging from traditional ones to more modern ones including Matrix Factorization, UserKNN, ItemKNN, NeuMF and VAE-CF (more detailed information on these and the fairness-aware models can be found in Section A.2).

Additionally, we include a regularization-based fairness-aware model BeyondParity [77] and a counterfactually fair recommendation model SM-GBiasedMF [44]. To weigh the fairness-aware loss for SM-GBiasedMF, we use  $\lambda = 20$  for all experiments since this value seems to work well as mentioned in [44]. To make a fair comparison, we use BeyondParity’s  $U_{val}$  (refer to Equation 11) as a regularization term in the same setting as shown in Equation 10, with the same value of alpha we use for our MF model. We use an early stopping strategy for the three models being compared, monitoring NDCG@20 with a delta of 0.0005 and patience of 10 epochs with a maximum number of iterations of 50 and 100 for the 100K and 1M datasets, respectively.

We utilize the Cornac framework [57], with customizations to meet our requirements, including the implementations of differentiable ways to obtain top  $k$  category distribution for calculating our gender loss term for different models. When calculating our regularization term, we generate a top  $k$  recommendation list of size  $k = 50$ . Additionally, all our results for all metrics for GBS (except CRP) use the top 50 recommended items. For GBS(CRP) we use top  $[R_g]$  as explained for Equation 6.

## 6 Analysis of the proposed metrics

In this section, we discuss the results we obtained from our experiments through a comprehensive analysis.

Model	Hit Ratio $\uparrow$	NDCG $\uparrow$	CC $\downarrow$	RCR $\downarrow$	CMAF $\downarrow$	CDCG $\downarrow$	CMRR $\downarrow$	CRP $\downarrow$	
<b>ML 100k Dataset</b>									
UserKNN	0.4931	0.0321	0.0283	0.0340	0.0029	0.0063	0.0017	0.0178	
ItemKNN	0.4698	0.0347	0.0675	0.0749	0.0095	0.0181	0.0070	0.0567	
MF	Original	<b>0.9003</b>	<b>0.1979</b>	0.0288	0.0321	0.0043	0.0089	<b>0.0055</b>	0.0276
	Fair	0.8950	0.1794	<b>0.0228</b>	<b>0.0312</b>	<b>0.0037</b>	<b>0.0085</b>	0.0060	<b>0.0234</b>
VAE-CF	Original	0.8749	<b>0.1643</b>	0.0187	0.0210	0.0030	0.0052	<b>0.0026</b>	0.0137
	Fair	<b>0.8759</b>	0.1575	<b>0.0140</b>	<b>0.0186</b>	<b>0.0026</b>	<b>0.0047</b>	0.0037	<b>0.0128</b>
NEU-MF	Original	<b>0.9470</b>	<b>0.2402</b>	0.1723	0.1452	0.0268	0.0502	0.0220	0.1411
	Fair	0.9194	0.2305	<b>0.0333</b>	<b>0.0405</b>	<b>0.0056</b>	<b>0.0093</b>	<b>0.0059</b>	<b>0.0342</b>
<b>ML 1M Dataset</b>									
UserKNN	0.3997	0.0179	0.0204	0.0110	0.0008	0.0043	0.0008	0.0174	
ItemKNN	0.3565	0.0383	0.0881	0.0351	0.0051	0.0224	0.0082	0.0650	
MF	Original	<b>0.8485</b>	<b>0.1522</b>	0.1775	0.0486	0.0097	0.0495	0.0204	0.1440
	Fair	0.8055	0.1158	<b>0.1160</b>	<b>0.0314</b>	<b>0.0067</b>	<b>0.0344</b>	<b>0.0162</b>	<b>0.0760</b>
VAE-CF	Original	<b>0.8315</b>	<b>0.1473</b>	0.2551	0.0708	0.0154	0.0706	0.0282	0.1900
	Fair	0.8280	0.1423	<b>0.2229</b>	<b>0.0637</b>	<b>0.0132</b>	<b>0.0631</b>	<b>0.0263</b>	<b>0.1502</b>
NEU-MF	Original	<b>0.9108</b>	0.1336	0.3096	0.0998	0.0246	0.0835	0.0316	0.2180
	Fair	0.8333	<b>0.1359</b>	<b>0.1503</b>	<b>0.0415</b>	<b>0.0086</b>	<b>0.0436</b>	<b>0.0191</b>	<b>0.0959</b>
<b>Yelp Dataset</b>									
UserKNN	0.4886	0.0312	0.0068	0.0252	0.0068	0.0026	0.0022	0.0129	
ItemKNN	0.5030	0.0344	0.0242	0.0513	0.0060	0.0065	0.0035	0.0219	
MF	Original	<b>0.8587</b>	<b>0.1209</b>	0.0358	0.0342	0.0055	0.0088	0.0036	0.0296
	Fair	0.7500	0.0773	<b>0.0159</b>	<b>0.0312</b>	<b>0.0019</b>	<b>0.0042</b>	<b>0.0019</b>	<b>0.0130</b>
VAE-CF	Original	<b>0.8131</b>	<b>0.0985</b>	0.0045	0.0041	0.0003	0.0007	0.0003	0.0018
	Fair	0.8032	0.0961	<b>0.0021</b>	<b>0.0019</b>	<b>0.0002</b>	<b>0.0004</b>	<b>0.0001</b>	<b>0.0011</b>
NEU-MF	Original	<b>0.9347</b>	<b>0.1794</b>	0.0592	0.0850	0.0134	0.0152	0.0057	0.0651
	Fair	0.8533	0.1124	<b>0.0309</b>	<b>0.0401</b>	<b>0.0054</b>	<b>0.0079</b>	<b>0.0028</b>	<b>0.0394</b>

**Table 2: Performance and Bias Evaluation values (GBSs) for 5 baseline models, along with the fair models.**

## 6.1 Baselines Comparison

**6.1.1 Bias Evaluation.** We report our results in Figures 2, 3, and Table 2. For the KNN-based models, ItemKNN seems more bias-prone. This can be due to the nature of the model of calculating similarity between items based on the interactions by the users. For instance, in the ML-100k dataset, the average rating for romance movies is higher for female users than for males. This kind of imbalance can reinforce category-related gender bias when items are recommended. Among the three other baselines, the MF-based models can be considered more biased. For the MF model, explicit embeddings are used to store user and item representations, which are likely to capture correlations of user behavior (e.g., liking certain genres) and their sensitive attributes. For VAE-CF, variational autoencoders are used to learn probabilistic latent representations of users, which are still sensitive to capturing biased representations but not as much as NeuMF. NeuMF captures both linear and non-linear relationships between users and items through Generalized Matrix Factorization (GMF) and Multi-Layer Perceptrons (MLP). So, NeuMF can capture biases in the data, especially the intricate patterns about user preferences, which can reflect social stereotypes. The MF model itself is not as complex as NeuMF, so our results of NeuMF being more biased than MF is reasonable. For the two smaller datasets, VAE-CF doesn’t manifest high values for bias, which can be due to the model not being able to fully capture representations of gendered preferences because of the lower number of interactions available. Comparatively, for the 1M dataset, the model can learn the stereotypical patterns available in the dataset more effectively and reinforce these biases in the latent representations. In general, all models manifest more gender bias for the larger dataset than the smaller ones. For all baseline models, the Gender Balance Scores are  $\neq 0$ , which highlights the fact that they are producing recommendations in a discriminative way where they choose certain genres to be relevant for certain groups of users. From Table 2, GBSs are higher for CC or RCR (ref



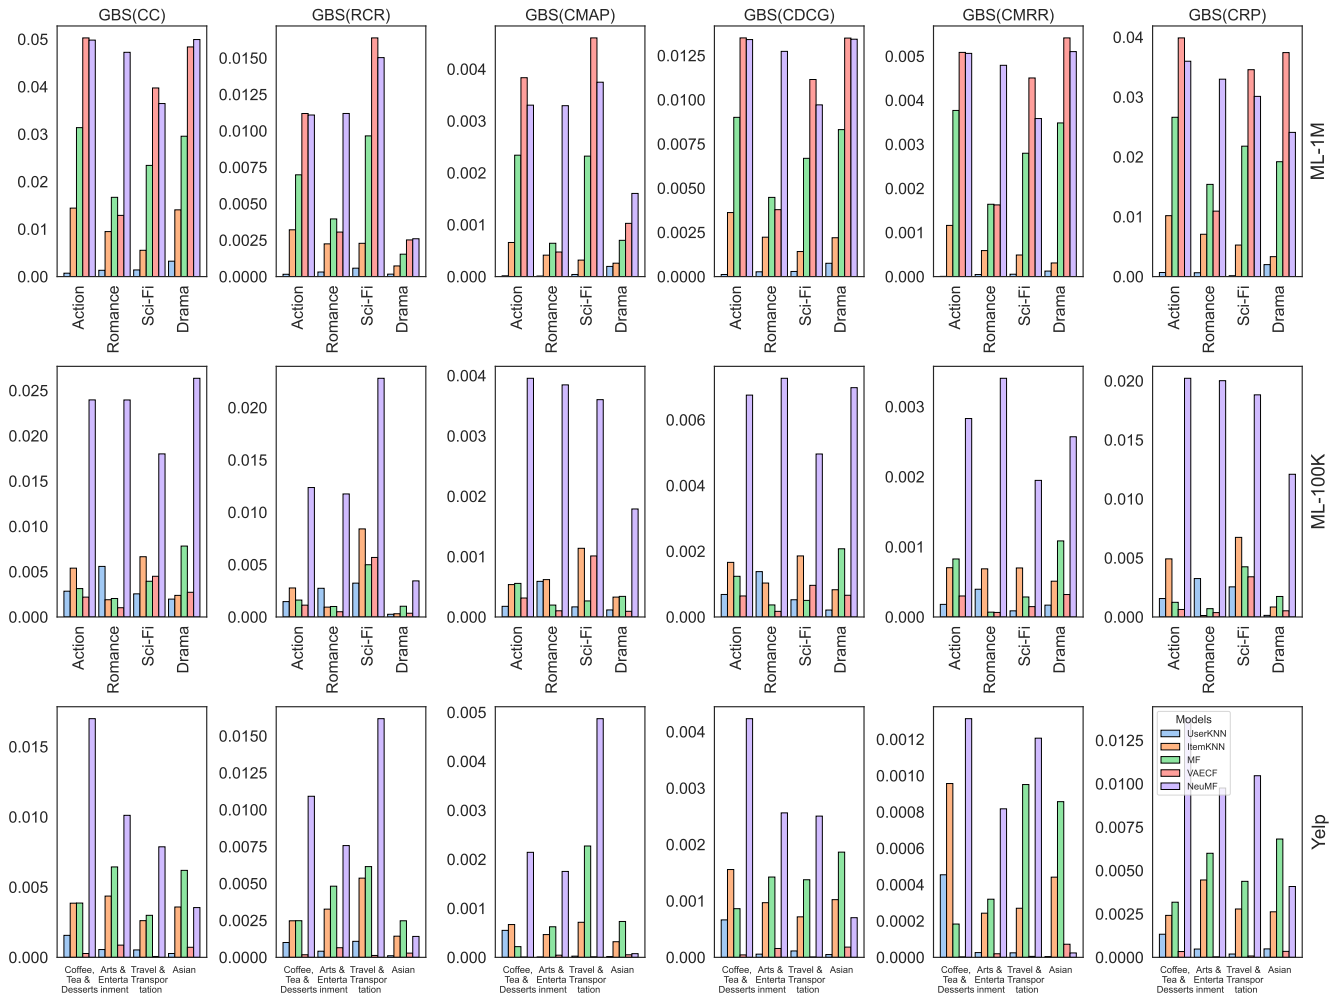


Figure 2: Comparison of Bias values for six of our metrics for all three dataset.

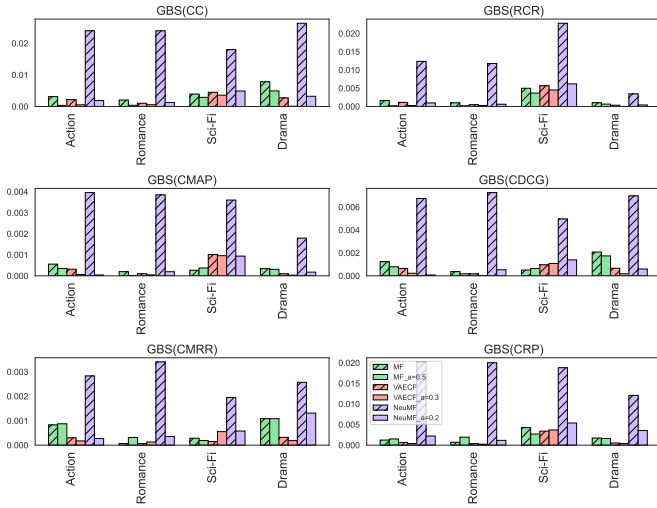
to Equations: 1 and 2). These two are classification-based metrics and do not consider ranks, so there is no discounting done in the value of any item that appears lower in the list but is still relevant (category-wise). GBS(CMRR) has the lowest values for almost all models since it has a strong discounting factor for items that appear lower in rank.

The bias reduction after using the fairness aware loss term is displayed in Figure 3. We can see a notable improvement in the bias scores, especially for the NeuMF model. The regularizer term is most effective for reducing bias in the NeuMF model, with an average decrease across all six metrics of 77%, 53%, and 50% for ML100K, ML1M, and Yelp datasets, respectively. The impact of the fairness regularizer is less pronounced for the VAE-CF model. We believe this is due to the model’s nature of learning probabilistic latent factors for users and items that don’t amplify gender bias like the MF models which use embeddings to learn user and item interactions. As a result, the regularizer has less impact on the model’s learning process.

6.1.2 Performance Evaluation. The performance drop for VAE-CF and NeuMF is less than that of the MF model, however, it all depends on the selection of  $\alpha$  (more on this in Section A.1). For all models, when choosing  $\alpha$  we ensure there is substantial bias reduction without too significant of a drop in performance. In some cases the fair models outperform the original models, this although seems counter-intuitive, but is presumably due to the fairness term acting as a regularization term which can help reduce over-fitting in the model to some extent (as noted in [35, 39]).

## 6.2 Fairness-Aware Baselines

We can observe the results for our MF fair model when compared with two other fairness-aware models in Table 3 and Figure 4. We chose to use our MF fair model since the other two fairness-aware models use MF as their foundational models. Our model has outperformed the other models in terms of bias scores across the majority of the comparisons. While our models don’t excel in terms of performance, they still offer a better balance between performance and

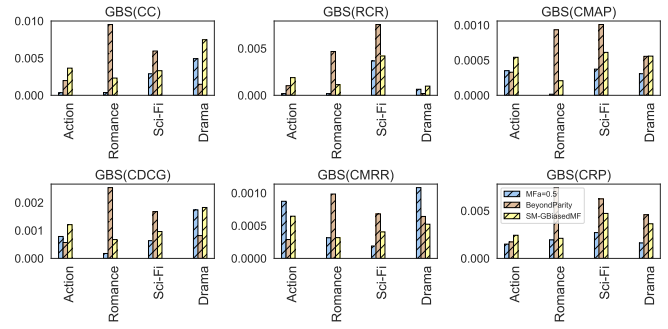


**Figure 3: Reduction in bias scores after using our fairness-aware regularizer. Since all datasets provide similar outcomes, we present the results for only the ML 100K dataset.**

Model	NDCG↑	HitRatio↑	CC↓	RCR↓	CMAP↓	CDCG↓	CMRR↓	CRP↓
<b>ML 100K Dataset</b>								
MF <sub>a=0.5</sub>	<b>0.1794</b>	<b>0.8950</b>	<b>0.0228</b>	<b>0.0312</b>	<b>0.0037</b>	<b>0.0085</b>	0.0060	<b>0.0234</b>
SM-GBiasedMF	0.1230	0.8537	0.03785	0.0378	0.0059	0.0114	<b>0.0053</b>	0.0286
BeyondParity	0.0790	0.6384	0.0528	0.0752	0.0093	0.0147	0.0068	0.0480
<b>ML 1M Dataset</b>								
MF <sub>a=0.6</sub>	<b>0.1158</b>	0.8055	<b>0.1160</b>	<b>0.0314</b>	<b>0.0067</b>	<b>0.0344</b>	0.0162	<b>0.0760</b>
SM-GBiasedMF	0.0846	<b>0.8075</b>	0.2590	0.0726	0.0152	0.0656	0.0211	0.2058
BeyondParity	0.0829	0.6495	0.1307	0.0539	0.0079	0.0369	<b>0.0154</b>	0.0838
<b>Yelp Dataset</b>								
MF <sub>a=0.3</sub>	0.0773	0.7500	0.0159	<b>0.0312</b>	<b>0.0019</b>	0.0042	<b>0.0019</b>	<b>0.0130</b>
SM-GBiasedMF	<b>0.9157</b>	<b>0.1343</b>	0.0515	0.0789	0.0132	0.0140	0.0055	0.0587
BeyondParity	0.0358	0.4901	<b>0.0112</b>	0.0342	0.0066	<b>0.0033</b>	0.0027	0.0241

**Table 3: Performance and Bias Evaluation Metric values for fairness-aware models.**

bias scores when compared with the other fairness-aware baselines. It is important to highlight that BeyondParity uses plain MSE loss, while the other two use BPR loss with negative sampling, which explains the exceptionally low-performance values. While we emphasize the significance of comparing the aggregated differences for each metric  $\mathcal{M}$  for all the categories, we also want to highlight the importance of evaluating bias values separately. For instance, although GBS(CMRR) of SM-GBiasedMF is lower for the ML100K dataset when compared to ours, there is still significant bias for movies recommended that belong to the genres like *Romance* and *Sci-Fi* as observed in Figure 4. This strengthens our idea of having a set of metrics to quantify bias in a nuanced way since the model seems "fair" when considering overall scores for all categories, but closer inspection reveals how it is biased against certain categories. It is worth mentioning that SM-GBiasedMF is more biased than the plain MF model for the Yelp dataset. Our theory is that since this model needs more epochs to satisfy the early stopping criterion (as mentioned before), it likely over-fits the dataset, in turn amplifying the biases present in it.



**Figure 4: Bias score for the fairness-aware models over four stereotypical genres for the ML 100K dataset.**

## 7 Discussion and Conclusion

In this paper, we identify the underlying issues of current metrics for evaluating consumer-side bias in recommendation systems. To better quantify bias, specifically gender bias, in such models, we propose a set of metrics. These metrics help capture a nuanced sense of fairness on recommended items by considering categories of the items. Next, we demonstrate how introducing one of our metrics as a fairness loss term along with the recommendation loss helped minimize the unfairness manifested in models in terms of different categories of items recommended with minimal performance loss. Experiments on three real-world datasets using a variety of recommendation models, including fairness-aware models, show the effectiveness of our metrics in capturing bias. Additionally, after incorporating our loss term, the bias in the models was significantly reduced, with a favorable balance between fairness and accuracy. Our loss function is very flexible which allows for context-sensitive fairness. In domains such as job recommendation and housing, where fairness is vital, the functions can prioritize fairness. Conversely, for areas like movie recommendations it can balance personalization and fairness (ensuring user preferences are respected). Our work aims to spread awareness about how simple metrics that are currently utilized to evaluate bias might be giving researchers a false sense of fairness, and a more refined approach like ours is required to address this issue. Our metrics are very versatile and can be easily adapted to measure provider-side fairness by utilizing proportions of item brands, for example, and can help quantify popularity bias in recommendations. Extending on this, the metrics can be used for measuring CP-fairness since we can measure bias from both the consumer side and provider side, by taking differences in values for different item providers for different demographic groups. Plus we plan to extend our metrics to consider sensitive attributes which are multi-valued, by using a pairwise difference scheme.

## Acknowledgments

This material is based upon work supported by the Air Force Office of Scientific Research under award number FA2386-23-1-4003.

## References

- [1] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The unfairness of popularity bias in recommendation. *arXiv preprint*



- arXiv:1907.13286 (2019).
- [2] Himan Abdollahpour, Masoud Mansoury, Robin Burke, Bamshad Mobasher, and Edward Malthouse. 2021. User-centered Evaluation of Popularity Bias in Recommender Systems. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization* (Utrecht, Netherlands) (UMAP '21). Association for Computing Machinery, New York, NY, USA, 119–129. <https://doi.org/10.1145/3450613.3456821>
  - [3] Rupam Acharyya, Shouman Das, Ankani Chattoraj, and Md. Iftekhar Tanveer. 2020. FairTED: A Fair Rating Predictor for TED Talk Data. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (Apr. 2020), 338–345. <https://doi.org/10.1609/aaai.v34i01.5368>
  - [4] Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Proceedings of the 30th International Conference on Neural Information Processing Systems* (Barcelona, Spain) (NIPS'16). Curran Associates Inc., Red Hook, NY, USA, 4356–4364.
  - [5] Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. 2022. Consumer fairness in recommender systems: Contextualizing definitions and mitigations. In *European Conference on Information Retrieval*. Springer, 552–566.
  - [6] Ludwig Bothmann, Kristina Peters, and Bernd Bischl. 2024. What Is Fairness? On the Role of Protected Attributes and Fictitious Worlds. arXiv:2205.09622 [cs.LG] <https://arxiv.org/abs/2205.09622>
  - [7] John S. Breese, David Heckerman, and Carl Kadie. 1998. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (Madison, Wisconsin) (UAI'98). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 43–52.
  - [8] Robin Burke. 2017. Multisided Fairness for Recommendation. <https://doi.org/10.48550/arXiv.1707.00093> arXiv:1707.00093 [cs].
  - [9] Joyantlya Chakraborty, Suvodeep Majumder, Zhe Yu, and Tim Menzies. 2020. Fairway: a way to build fair ML software. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Virtual Event, USA) (ESEC/FSE 2020). Association for Computing Machinery, New York, NY, USA, 654–665. <https://doi.org/10.1145/3368089.3409697>
  - [10] Alexandra Chouldechova. 2016. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* 5 2 (2016), 153–163. <https://api.semanticscholar.org/CorpusID:1443041>
  - [11] Amit Datta, Michael Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings. *Proceedings on Privacy Enhancing Technologies* 1 (04 2015). <https://doi.org/10.1515/popets-2015-0007>
  - [12] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (FAT\* '19). Association for Computing Machinery, 120–128. <https://doi.org/10.1145/3287560.3287572>
  - [13] Yashar Deldjoo, Vito Walter Anelli, Hamed Zamani, Alejandro Bellogin, and Tommaso Di Noia. 2021. A flexible framework for evaluating user and item fairness in recommender systems. *User Modeling and User-Adapted Interaction* 31, 3 (July 2021), 457–511. <https://doi.org/10.1007/s11257-020-09285-1>
  - [14] Virginie Do, Sam Corbett-Davies, Jamal Atif, and Nicolas Usunier. 2022. Online certification of preference-based fairness for personalized recommender systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 6532–6540.
  - [15] Bora Edizel, Francesco Bonchi, Sara Hajian, Andre Panisson, and Tamir Tassa. 2020. FairRecSys: mitigating algorithmic bias in recommender systems. *International Journal of Data Science and Analytics* 9 (03 2020), 197–213. <https://doi.org/10.1007/s41060-019-00181-5>
  - [16] Marwa El Halabi, Slobodan Mitrović, Ashkan Norouzi-Fard, Jakab Tardos, and Jakub M Tarnawski. 2020. Fairness in Streaming Submodular Maximization: Algorithms and Hardness. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 13609–13622. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/9d752cb08ef466fc480fba981cfa44a1-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/9d752cb08ef466fc480fba981cfa44a1-Paper.pdf)
  - [17] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway Feedback Loops in Predictive Policing. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (Proceedings of Machine Learning Research, Vol. 81), Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 160–171. <https://proceedings.mlr.press/v81/ensign18a.html>
  - [18] Golnoosh Farnadi, Pigi Kouki, Spencer K. Thompson, Sriram Srinivasan, and Lise Getoor. 2018. A Fairness-aware Hybrid Recommender System. *CoRR* abs/1809.09030 (2018). arXiv:1809.09030 <http://arxiv.org/abs/1809.09030>
  - [19] James Foulds, Rashidul Islam, Kamrun Keya, and Shimei Pan. 2019. Bayesian Modeling of Intersectional Fairness: The Variance of Bias. arXiv:1811.07255 [cs.LG] <https://arxiv.org/abs/1811.07255>
  - [20] James R. Foulds, Rashidul Islam, Kamrun Naher Keya, and Shimei Pan. 2020. An Intersectional Definition of Fairness. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. 1918–1921. <https://doi.org/10.1109/ICDE48307.2020.00203>
  - [21] Zuohui Fu, Yikun Xian, Ruoyuan Gao, Jieyu Zhao, Qiaoying Huang, Yingqiang Ge, Shuyuan Xu, Shijie Geng, Chirag Shah, Yongfeng Zhang, and Gerard de Melo. 2020. Fairness-Aware Explainable Recommendation over Knowledge Graphs. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 69–78. <https://doi.org/10.1145/3397271.3401051>
  - [22] Sainyam Galhotra, Yuriy Brun, and Alexandra Meliou. 2017. Fairness testing: testing software for discrimination. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering* (Paderborn, Germany) (ESEC/FSE 2017). Association for Computing Machinery, New York, NY, USA, 498–510. <https://doi.org/10.1145/3106237.3106277>
  - [23] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, and Yongfeng Zhang. 2021. Towards Long-term Fairness in Recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (WSDM '21). Association for Computing Machinery, New York, NY, USA, 445–453. <https://doi.org/10.1145/3437963.3441824>
  - [24] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Anchorage, AK, USA) (KDD '19). Association for Computing Machinery, New York, NY, USA, 2221–2231. <https://doi.org/10.1145/3292500.3330691>
  - [25] Avijit Ghosh, Ritam Dutt, and Christo Wilson. 2021. When Fair Ranking Meets Uncertain Inference. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 1033–1043. <https://doi.org/10.1145/3404835.3462850>
  - [26] Usman Gohar and Lu Cheng. 2023. A Survey on Intersectional Fairness in Machine Learning: Notions, Mitigation, and Challenges. In *International Joint Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:258615326>
  - [27] Sruthi Gorantla, Amit Deshpande, and Anand Louis. 2021. On the Problem of Underranking in Group-Fair Ranking. In *Proceedings of the 38th International Conference on Machine Learning* (Proceedings of Machine Learning Research, Vol. 139), Marina Meila and Tong Zhang (Eds.). PMLR, 3777–3787. <https://proceedings.mlr.press/v139/gorantla21a.html>
  - [28] Manjul Gupta, Carlos M. Parra, and Denis Denhehy. 2022. Questioning Racial and Gender Bias in AI-based Recommendations: Do Espoused National Cultural Values Matter? *Information Systems Frontiers* 24, 5 (Oct. 2022), 1465–1481. <https://doi.org/10.1007/s10796-021-10156-2>
  - [29] Qianxiu Hao, Qianqian Xu, Zhiyong Yang, and Qingming Huang. 2021. Pareto Optimality for Fairness-constrained Collaborative Filtering. In *Proceedings of the 29th ACM International Conference on Multimedia* (Virtual Event, China) (MM '21). Association for Computing Machinery, New York, NY, USA, 5619–5627. <https://doi.org/10.1145/3474085.3475706>
  - [30] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of opportunity in supervised learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems* (Barcelona, Spain) (NIPS'16). Curran Associates Inc., Red Hook, NY, USA, 3323–3331.
  - [31] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. <https://doi.org/10.1145/2827872>
  - [32] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) (WWW '17). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. <https://doi.org/10.1145/3038912.3052569>
  - [33] Basileal Imana, Aleksandra Korolova, and John Heidemann. 2021. Auditing for Discrimination in Algorithms Delivering Job Ads. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 3767–3778. <https://doi.org/10.1145/3442381.3450077>
  - [34] Rashidul Islam, Kamrun Naher Keya, Ziqian Zeng, Shimei Pan, and James Foulds. 2021. Debiasing Career Recommendations with Neural Fair Collaborative Filtering. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 3779–3790. <https://doi.org/10.1145/3442381.3449904>
  - [35] Rashidul Islam, Kamrun Naher Keya, Ziqian Zeng, Shimei Pan, and James Foulds. 2021. Debiasing Career Recommendations with Neural Fair Collaborative Filtering. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 3779–3790. <https://doi.org/10.1145/3442381.3449904>
  - [36] Thomas Jaenic, Graham McDonald, and Iadh Ounis. 2024. Fairness-Aware Exposure Allocation via Adaptive Reranking. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 1504–1513. <https://doi.org/10.1145/3626772.3657794>

- [37] Meng Jiang, Keqin Bao, Jizhi Zhang, Wenjie Wang, Zhengyi Yang, Fuli Feng, and Xiangnan He. 2024. Item-side Fairness of Large Language Model-based Recommendation System. In *Proceedings of the ACM Web Conference 2024* (Singapore, Singapore) (WWW '24). Association for Computing Machinery, New York, NY, USA, 4717–4726. <https://doi.org/10.1145/3589334.3648158>
- [38] Di Jin, Luzhi Wang, He Zhang, Yizhen Zheng, Weiping Ding, Feng Xia, and Shirui Pan. 2023. A survey on fairness-aware recommender systems. *Information Fusion* 100 (2023), 101906. <https://doi.org/10.1016/j.inffus.2023.101906>
- [39] Kamrun Naher Keya, Rashidul Islam, Shimei Pan, Ian Stockwell, and James R. Foulds. 2020. Equitable Allocation of Healthcare Resources with Fair Cox Models. arXiv:2010.06820 [cs.LG] <https://arxiv.org/abs/2010.06820>
- [40] Tahsin Alamgir Khaya, Mohamed Reda Bouadjeneq, and Sunil Aryal. 2024. The Pursuit of Fairness in Artificial Intelligence Models: A Survey. arXiv:2403.17333 [cs.AI] <https://arxiv.org/abs/2403.17333>
- [41] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009), 30–37. <https://doi.org/10.1109/MC.2009.263>
- [42] Matt Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 4069–4079.
- [43] Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2021. User-oriented Fairness in Recommendation. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 624–632. <https://doi.org/10.1145/3442381.3449866>
- [44] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2021. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1054–1063.
- [45] Dawen Liang, Rahul G. Krishnan, Matthew D. Hoffman, and Tony Jebara. 2018. Variational Autoencoders for Collaborative Filtering. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) (WWW '18). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 689–698. <https://doi.org/10.1145/3178876.3186150>
- [46] Chen Lin, Xinyi Liu, Guipeng Xu, and Hui Li. 2021. Mitigating Sentiment Bias for Recommender Systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 31–40. <https://doi.org/10.1145/3404835.3462943>
- [47] Kun Lin, Nasim Sonboli, Bamshad Mobasher, and Robin Burke. 2019. Crank up the volume: preference bias amplification in collaborative recommendation. arXiv:1909.06362 [cs.IR] <https://arxiv.org/abs/1909.06362>
- [48] Haifeng Liu, Nan Zhao, Xiaokun Zhang, Hongfei Lin, Liang Yang, Bo Xu, Yuan Lin, and Wenqi Fan. 2022. Dual constraints and adversarial learning for fair recommenders. *Knowledge-Based Systems* 239 (2022), 108058. <https://doi.org/10.1016/j.knsys.2021.108058>
- [49] Kristian Lum and William Isaac. 2016. To Predict and Serve? *Significance* 13 (10 2016), 14–19. <https://doi.org/10.1111/j.1740-9713.2016.00960.x>
- [50] Masoud Mansoury, Bamshad Mobasher, Robin Burke, and Mykola Pechenizkiy. 2019. Bias Disparity in Collaborative Recommendation: Algorithmic Evaluation and Comparison. arXiv:1908.00831 [cs.IR] <https://arxiv.org/abs/1908.00831>
- [51] Giacomo Medda, Francesco Fabbri, Mirko Marras, Ludovico Boratto, and Gianni Fenu. 2024. GNNUsers: Fairness Explanation in GNNs for Recommendation via Counterfactual Reasoning. *ACM Trans. Intell. Syst. Technol.* (apr 2024). <https://doi.org/10.1145/3655631> Just Accepted.
- [52] Alessandro B. Melchiorre, Navid Rekasaz, Emilia Parada-Cabaleiro, Stefan Brandl, Oleg Lesota, and Markus Schedl. 2021. Investigating gender fairness of recommendation algorithms in the music domain. *Inf. Process. Manage.* 58, 5 (sep 2021), 27 pages. <https://doi.org/10.1016/j.ipm.2021.102666>
- [53] Mohammadmehdi Naghiaei, Hossein A. Rahmani, and Yashar Deldjoo. 2022. CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Madrid, Spain) (SIGIR '22). Association for Computing Machinery, New York, NY, USA, 770–779. <https://doi.org/10.1145/3477495.3531959>
- [54] Hossein A. Rahmani, Mohammadmehdi Naghiaei, and Yashar Deldjoo. 2024. A Personalized Framework for Consumer and Producer Group Fairness Optimization in Recommender Systems. *ACM Trans. Recomm. Syst.* 2, 3, Article 19 (jun 2024), 24 pages. <https://doi.org/10.1145/3651167>
- [55] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) (UAI '09). AUAI Press, Arlington, Virginia, USA, 452–461.
- [56] Clara Rus, Jeffrey Luppess, Harrie Oosterhuis, and Gido H. Schoenmacker. 2022. Closing the Gender Wage Gap: Adversarial Fairness in Job Recommendation. arXiv:2209.09592 [cs.LG] <https://arxiv.org/abs/2209.09592>
- [57] Aghiles Salah, Quoc-Tuan Truong, and Hady W. Lauw. 2020. Cornac: A Comparative Framework for Multimodal Recommender Systems. *Journal of Machine Learning Research* 21, 95 (2020), 1–5. <http://jmlr.org/papers/v21/19-805.html>
- [58] Piotr Sapiezynski, Wesley Zeng, Ronald E Robertson, Alan Mislove, and Christo Wilson. 2019. Quantifying the Impact of User Attention Fair Group Representation in Ranked Lists. In *Companion Proceedings of The 2019 World Wide Web Conference* (San Francisco, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 553–562. <https://doi.org/10.1145/3308560.3317595>
- [59] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2000. Analysis of recommendation algorithms for e-commerce. In *Proceedings of the 2nd ACM Conference on Electronic Commerce* (Minneapolis, Minnesota, USA) (EC '00). Association for Computing Machinery, New York, NY, USA, 158–167. <https://doi.org/10.1145/352871.352887>
- [60] Dimitris Serbos, Shuyao Qi, Nikos Mamoulis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2017. Fairness in Package-to-Group Recommendations. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) (WWW '17). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 371–379. <https://doi.org/10.1145/3038912.3052612>
- [61] Guy Shani and Asela Gunawardana. 2011. *Evaluating Recommendation Systems*. Vol. 12. 257–297. [https://doi.org/10.1007/978-0-387-85820-3\\_8](https://doi.org/10.1007/978-0-387-85820-3_8)
- [62] Harald Steck. 2018. Calibrated recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) (RecSys '18). Association for Computing Machinery, New York, NY, USA, 154–162. <https://doi.org/10.1145/3240323.3240372>
- [63] Jiakai Tang, Shiqi Shen, Zhipeng Wang, Zhi Gong, Jingsen Zhang, and Xu Chen. 2023. When Fairness meets Bias: a Debaised Framework for Fairness aware Top-N Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems* (RecSys '23). Association for Computing Machinery, New York, NY, USA, 200–210. <https://doi.org/10.1145/3604915.3608770>
- [64] Virginia Tsintzou, Evaggelia Pitoura, and Panayiotis Tsaparas. 2019. Bias Disparity in Recommendation Systems. In *Proceedings of the Workshop on Recommendation in Multi-stakeholder Environments co-located with the 13th ACM Conference on Recommender Systems (RecSys 2019), Copenhagen, Denmark, September 20, 2019 (CEUR Workshop Proceedings, Vol. 2440)*, Robin Burke, Himan Abdollahpouri, Edward C. Malthouse, K. P. Thai, and Yongfeng Zhang (Eds.). CEUR-WS.org. <https://ceur-ws.org/Vol-2440/short4.pdf>
- [65] Sahil Verma and Julia Rubin. 2018. Fairness definitions explained. In *Proceedings of the International Workshop on Software Fairness* (Gothenburg, Sweden) (FairWare '18). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3194770.3194776>
- [66] Mengting Wan, Jianmo Ni, Rishabh Misra, and Julian McAuley. 2020. Addressing Marketing Bias in Product Recommendations. In *Proceedings of the 13th International Conference on Web Search and Data Mining* (Houston, TX, USA) (WSDM '20). Association for Computing Machinery, New York, NY, USA, 618–626. <https://doi.org/10.1145/3336191.3371855>
- [67] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Trans. Inf. Syst.* 41, 3, Article 52 (feb 2023), 43 pages. <https://doi.org/10.1145/3547333>
- [68] Tianxin Wei and Jingrui He. 2022. Comprehensive Fair Meta-learned Recommender System. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, New York, NY, USA, 1989–1999.
- [69] Tianxin Wei and Jingrui He. 2022. Comprehensive Fair Meta-learned Recommender System. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (KDD '22). Association for Computing Machinery, New York, NY, USA, 1989–1999. <https://doi.org/10.1145/3534678.3539269>
- [70] Leonard Weydemann, Dimitris Sacharidis, and Hannes Werthner. 2019. Defining and measuring fairness in location recommendations. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Recommendations, Geosocial Networks and GeoAdvertising* (Chicago, Illinois) (LocalRec '19). Association for Computing Machinery, New York, NY, USA, Article 6, 8 pages. <https://doi.org/10.1145/3356994.3365497>
- [71] Chuhan Wu, Fangzhao Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2020. Fairness-aware News Recommendation with Decomposed Adversarial Learning.
- [72] Kun Wu, Jacob Erickson, Wendy Hui Wang, and Yue Ning. 2022. Equipping Recommender Systems with Individual Fairness via Second-order Proximity Embedding. In *2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. 171–175. <https://doi.org/10.1109/ASONAM55673.2022.10068703>
- [73] Le Wu, Lei Chen, Pengyang Shao, Richang Hong, Xiting Wang, and Meng Wang. 2021. Learning Fair Representations for Recommendation: A Graph-based Perspective. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 2198–2208. <https://doi.org/10.1145/3442381.3450015>
- [74] Bin Xia, Junjie Yin, Jian Xu, and Yun Li. 2019. WE-Rec: A fairness-aware reciprocal recommendation based on Walrasian equilibrium. *Knowledge-Based Systems* 182 (2019), 104857. <https://doi.org/10.1016/j.knsys.2019.07.028>
- [75] Hao Yang, Zhining Liu, Zeyu Zhang, Chenyi Zhuang, and Xu Chen. 2023. Towards Robust Fairness-aware Recommendation. In *Proceedings of the 17th ACM*

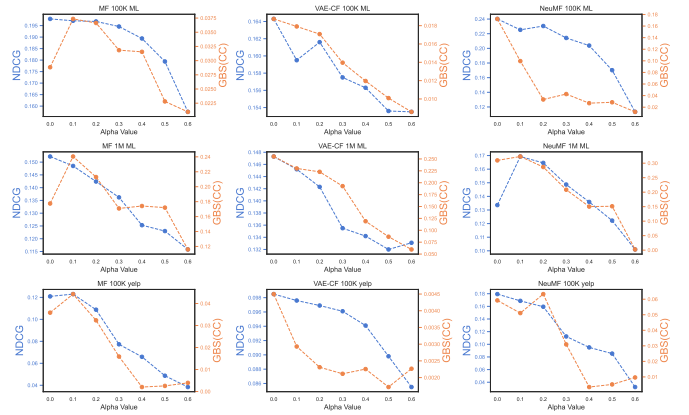
Conference on Recommender Systems (Singapore, Singapore) (RecSys '23). Association for Computing Machinery, New York, NY, USA, 211–222. <https://doi.org/10.1145/3604915.3608784>

- [76] Ke Yang, Joshua R. Loftus, and Julia Stoyanovich. 2020. Causal intersectionality for fair ranking. arXiv:2006.08688 [cs.LG] <https://arxiv.org/abs/2006.08688>
- [77] Sirui Yao and Bert Huang. 2017. Beyond Parity: Fairness Objectives for Collaborative Filtering. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc., Red Hook, NY, USA. [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/e6384711491713d29bc63fc5eeb5ba4f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/e6384711491713d29bc63fc5eeb5ba4f-Paper.pdf)
- [78] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P. Gummadi. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Jerry Zhu (Eds.). PMLR, 962–970. <https://proceedings.mlr.press/v54/zafar17a.html>
- [79] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. FA\*IR: A Fair Top-k Ranking Algorithm. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management (Singapore) (CIKM '17)*. Association for Computing Machinery, New York, NY, USA, 1569–1578. <https://doi.org/10.1145/3132847.3132938>
- [80] Ziwei Zhu, Xia Hu, and James Caverlee. 2018. Fairness-Aware Tensor-Based Recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. Association for Computing Machinery, New York, NY, USA, 1153–1162. <https://doi.org/10.1145/3269206.3271795>
- [81] Ziwei Zhu, Jianling Wang, and James Caverlee. 2020. Measuring and Mitigating Item Under-Recommendation Bias in Personalized Ranking Systems. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, China) (SIGIR '20)*. Association for Computing Machinery, New York, NY, USA, 449–458. <https://doi.org/10.1145/3397271.3401177>

## A Appendix

### A.1 Abalation Study

In this section, we discuss the influence of the hyper-parameter  $\alpha$  on recommendation performance and fairness. As mentioned in Section 4,  $\alpha$  can be used to control the strength of the fairness aware loss. Theoretically as  $\alpha$  increases we would expect a decrease in performance since the recommendation loss would have a decreased weight. We use  $\alpha$  values from 0 to 0.6 with increments of 0.1 for all three models. We did not include values above 0.6, because it does not make sense to overpower the recommendation loss since that is our main task. To verify the expected behavior of alpha on recommendation loss and bias we plot the metric we optimize on: GBS(CC) or  $\Delta$ CategoryCoverage and the NDCG values. As shown in Figure 5, there is a general trend of the performance metric and the bias score dropping as  $\alpha$  increases. There are some fluctuations, where the bias increases for the first few values of alpha. Our intuition for this is that the fairness loss might not be enough to decrease the bias when  $\alpha$  is small. It can essentially end up disturbing the loss function itself, without explicitly decreasing bias. But as the value increases the bias scores drop lower which is expected. There is a slight increase in NDCG value for certain models, as mentioned in Section 6.1.2, this can be the result of the fairness loss term acting as a regularization term that essentially prevents over-fitting in the model. We choose  $\alpha$  values by ensuring there is a decrease in bias, without significant loss in performance. For the ML 100K dataset, we choose 0.4, 0.3, and 0.2 respectively for MF, VAE-CF, and NeuMF models. For the ML 1M dataset, we choose 0.6, 0.2, and 0.4 for MF, VAE-CF, and NeuMF models respectively. Lastly, for the Yelp dataset, we choose 0.3 for all models.



**Figure 5: Impact of  $\alpha$  on recommendation performance wrt NDCG@50 and the bias measure which is the difference of Category Coverage values for male and female. The experiments are performed for all three datasets**

### A.2 Details about Baseline Models

This subsection delves into details about the models we worked on for our experiments.

- MF [41]: A classical Matrix Factorization algorithm where users and items are represented as latent vectors with global bias. In our implementation, we use BPR [55] loss with negative sampling to enhance the performance of the model.
- UserKNN [7]: This is a neighborhood-based method based on users, and items are recommended by discovering similar users based on the cosine similarity of their historical interactions.
- ItemKNN [59]: Another neighborhood-based method that computes the similarity between items instead.
- NeuMF [32]: NeuMF is a deep-learning based extension to MF, which combines the linearity of traditional MF models and the non-linearity of DNNs (Deep Neural Networks).
- VAE-CF [45]: This non-linear probabilistic model is based on the auto-encoder architecture and learns compressed information about data. The encoder helps map user interactions as a low-dimensional latent space, and the decoder decodes this information back to the high-dimensional vector which is used to make predictions.

The  $U_{val}$  regularization term for BeyondParity [77] can be written as:

$$U_{val} = \frac{1}{n} \sum_{j=1}^n |(E_g[y]_j - E_g[r]_j) - (E_{-g}[y]_j - E_{-g}[r]_j)| \quad (11)$$

where  $E_g[y]_j$  is the average predicted score for the  $j$ -th item from disadvantaged users,  $E_{-g}[y]_j$  is the average predicted score for the  $j$ -th item from advantaged users,  $E_g[r]_j$  is the average rating for the  $j$ -th item from disadvantaged users and  $E_{-g}[r]_j$  is the average rating for the  $j$ -th item from advantaged users. This fairness objective is optimized alongside the actual learning objective for a collaborative-based recommendation system. SM-GBiasedMF, is a fair recommendation model which achieves counterfactual fairness

by utilizing adversarial learning. They combine recommendation loss and an adversarial loss and, the trade-off is controlled by  $\lambda$ .

### A.3 Dataset Pre-processing

For all of our datasets, we filter out inactive users. A user is considered inactive if they have less than 5 interactions. Additionally, we remove any user whose gender is not known. The Yelp dataset has over 300 categories. We reduce this set of categories to a condensed group of 21 broader categories, for better interpretability. These categories include *Active Life & Fitness, Arts & Entertainment, Automotive, Bars & Nightlife, Coffee, Tea & Desserts, Drinks & Spirits, Education & Learning, Event Services, Family & Kids, Food & Restaurants, Health & Beauty, Home & Garden, Miscellaneous,*

*Outdoor Activities, Public Services & Community, Shopping & Fashion, Specialty Food & Groceries, Sports & Recreation, Technology & Electronics, Travel & Transportation, and Asian.*

### A.4 Process of generating Figure 1

This plot is generated using the ML-100K dataset. Once the models are done being trained we identify the users by their gender in the test set. To ensure a fair comparison we take the number of users in the smaller group, which was female in our case. We randomly chose the same number of male users from the test set. Once we have an equal number of male and female users, precision@10 values are calculated for each user and averaged by gender. To find the proportions of movies recommended, we generate the top 10 movies for each user. The proportion of each genre is computed using Equation 1.