Provider Fairness Across Continents in Collaborative Recommender Systems

Elizabeth Gómez[a], Ludovico Boratto[b], Maria Salamó[a]

[a] *Facultat de Matemàtiques i Informàtica, Universitat de Barcelona, Spain*
[b] *Department of Mathematics and Computer Science, University of Cagliari, Italy*

Abstract

When a recommender system suggests items to the end-users, it gives a certain exposure to the providers behind the recommended items. Indeed, the system offers a possibility to the items of those providers of being reached and consumed by the end-users. Hence, according to how recommendation lists are shaped, the experience of under-recommended providers in online platforms can be affected. To study this phenomenon, we focus on movie and book recommendation and enrich two datasets with the continent of production of an item. We use this data to characterize imbalances in the distribution of the user-item observations and regarding where items are produced (*geographic imbalance*). To assess if recommender systems generate a disparate impact and (dis)advantage a group, we divide items into groups, based on their continent of production, and characterize how represented is each group in the data. Then, we run state-of-the-art recommender systems and measure the visibility and exposure given to each group. We observe disparities that favor the most represented groups. We overcome these phenomena by introducing equity with a re-ranking approach that regulates the share of recommendations given to the items produced in a continent (*visibility*) and the positions in which items are ranked in the recommendation list (*exposure*), with a negligible loss in effectiveness, thus controlling fairness of providers coming from different continents. A comparison with the state of the art shows that our approach can provide more equity for providers, both in terms of visibility and of exposure.

Email addresses: egomezye13@alumnes.ub.edu (Elizabeth Gómez), ludovico.boratto@acm.org (Ludovico Boratto), maria.salamo@ub.edu (Maria Salamó)

Keywords:
Recommender systems, Bias, Provider fairness, Geographic groups, Data Imbalance, Disparate impact

1. Introduction

Recommender systems support users by suggesting items that might be of interest to them [1]. This is usually done by learning behavioral patterns from historical data, usually in the form of user-item interactions. However, imbalances in the input data can lead to biases in the results these algorithms produce [2]. The main example of this type of phenomenon is popularity bias, where popular items get over-recommended, to the detriment of long-tail ones [3]. If bias is associated with sensitive attributes of the users (such as gender or race), biased results might lead to unethical consequences, such as discrimination (*unfairness*) [2, 4]. Discrimination might affect both the end-users (often referred to as consumers), when those belonging to legally protected groups or certain individuals receive systematically worse recommendations (*consumer fairness*), and content producers, in case the items of those belonging to legally protected groups or individuals are under-recommended by an algorithm (*provider fairness*) [2, 5]. However, there are scenarios in which a recommender system works with imbalanced data not only because of a biased data collection, but because of the way an industry is composed. A clear example of this is the modern film industry, where the United States Cinema (Hollywood) takes the largest share of the market, both in terms of produced movies and of revenues[1]. Moreover, as observed by Bauer and Schedl, users belonging to different geographic areas have different item consumption patterns [6].

Given these considerations, it becomes natural to ask ourselves *if data imbalances associated with an industry can lead to unfairness for providers, according to the way recommendations are produced.* Specifically, we consider providers belonging to different geographic areas and assess if recommender systems exacerbate the natural imbalances existing in the input data, thus affecting the producers of smaller industries from a geographic point of view. In our recent work [7], we considered a binary setting, in which item pro-

[1]https://www.boxofficepro.com/mpa-2019-global-box-office-and-home-entertainment-surpasses-100-billion/

ducers were divided into two groups, a *majority* containing the items coming from the main country of production of the items, and a *minority* containing items produced in the rest of the world. We assessed how state-of-the-art collaborative filtering algorithms distributed the recommendations, and observed that the majority items are over-represented in the recommendation lists, both in terms of the number of recommendations (*visibility*) and in their position in the rankings (*exposure*). We presented an approach that redistributes the recommendations, so that the majority group has a representation in the recommendations that corresponds to that in the input data.

However, when dealing with provider fairness, it is important to *understand how recommendations are distributed across different provider groups.* Indeed, even if we ensure that the providers in the majority group are not over-recommended (as we did in [7]), we still do not have guarantees that the different provider groups belonging to the minority are recommended in equitable ways. In the context of geographic groups, this means that *the problem of how recommendations of items produced by providers in small regions are distributed, and of how to mitigate possible disparities, remains open.*

Unfair recommendations for providers, based on their geographic provenience, is an issue that goes beyond a biased functioning of an automated decision-support system and has consequences at multiple levels, by denying the opportunity to providers to offer their items (*ethical perspective*), thus limiting their possibility to work (*business perspective*); on top of this, unfair outputs are also forbidden by current regulations, such as GDPR (*legal perspective*). Hence, ensuring that providers coming from different parts of the world are not affected by the fact that they belong to a region that has a low share in the market, is a problem of central importance.

To address this problem, in this paper, we move from a binary to a multi-group setting, to study *unfairness for providers belonging to different continents.* We consider two of the main recommendation domains, namely movies and books, and assess how state-of-the-art collaborative filtering algorithms distribute the recommendations. We observe that both the original models and the mitigation we introduced for binary groups create disparities in both the visibility and exposure given to content providers in different continents and that the less represented is a group in the data, the worse is this disparity created by a recommendation model. To overcome these phenomena, we propose an approach that introduces fairness for providers belonging to different geographic areas, by re-distributing the recommen-

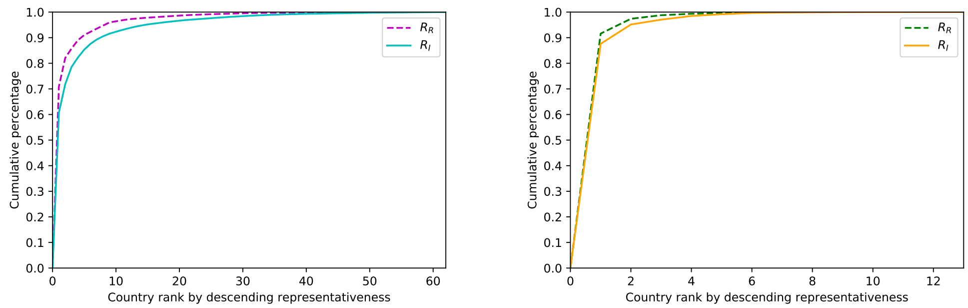

(a) Country representation in MovieLens-1M.

(b) Country representation in Book-Crossing.

Figure 1: **Country representation in the input data**. Representation of each country in the MovieLens-1M (a) and Book-Crossing (b) datasets. Representation is computed either considering the amount of items produced by a country ($\mathcal{R}_I$) or the amount of ratings it attracted ($\mathcal{R}_R$).

dations across continents following a notion of *equity* [8]. Concretely, our mitigation strategy gives a provider group a visibility and an exposure equal to the representation of the group in the input data.

Our choice to introduce equity considering a continent as our granularity level is motivated by Figure 1, where we show the representation of each country of production in the input data. The first thing that emerges is that the dataset we considered in the movie domain (MovieLens-1M) contains items produced in over 60 countries. Hence, introducing equity at the country level would require adjusting the recommendation lists to ensure that each of these countries received a representation equal to its representation in the input data. While this is challenging, due to the large number of countries and the limited size of a recommendation list, a second issue emerges when we observe how represented is each country in the data, both considering the number of items it produced ($\mathcal{R}_I$) and the number of ratings it received ($\mathcal{R}_R$)[2]. Again, from the MovieLens-1M dataset, we can see that the imbalance in the representation is severe, with the main country (the United States) who produced over 60% of the movies and attracted around 70% of the ratings. Given the very small representation of almost all the other countries, a regulation at the country level would lead to a severe drop in the

[2]Our two notions of group representation, $\mathcal{R}_I$ and $\mathcal{R}_R$, are formally presented in Section 3.1.

recommendation effectiveness for the users. Readers can also see that, while our second dataset, Book-Crossing, has much fewer countries, the imbalance in the data is even more severe. Hence, introducing provider fairness at the continent level allows us to work with a stable setting and to contrast among them the results obtained with the two datasets.

Our contributions can be summarized as follows:

- We assess unfairness for groups of providers belonging to different geographic continents, considering state-of-the-art recommendation models;
- We propose a re-ranking algorithm to introduce provider fairness for multiple groups, following a notion of equity that distributes the recommendations according to the representation of the groups in the input data;
- We evaluate our algorithm in two recommendation domains and study its effectiveness at producing fair but effective recommendations.

Concretely, we extend the study presented in [7] in the following ways: (*i*) we extend our related work, to improve the coverage of the existing literature; (*ii*) we analyze how our previous proposal deals with more than two groups; (*iii*) we introduce a new problem setting, which has a multi-group fairness goal; (*iv*) we introduce a new algorithm to introduce equity for more than two provider groups; (*v*) we add a comparison with the state of the art, to show why a mitigation tacking both visibility and exposure is needed to ensure provider fairness.

The rest of the paper is structured as follows: in Section 2, we present related work and in Section 3 we provide the foundation to our work. We continue by presenting, in Section 4, a summary of our work in [7], to have a reference on provider fairness for two provider groups in this context. Given this setting, in Section 5, we assess the capability of the state-of-the-art models and of our binary mitigation to provide fairness to groups shaped at continent level. In Section 6, we propose a mitigation algorithm to overcome unfairness scenarios in presence of multiple groups and we evaluate it in Section 7. Finally, we provide concluding remarks in Section 8.

2. Related Work

This section covers related studies. First of all, Section 2.1 starts from the concepts of visibility and exposure in ranking. Next, in Section 2.2, we continue with the impact of recommendation for providers. Finally, Section 2.3 concludes by contextualizing our work with the existing studies.

2.1. Visibility and Exposure in Rankings

Given a ranking, visibility measures the amount of times an item is presented in the rankings [9, 10] whereas exposure assesses *where* an item is ranked [11, 12]. Visibility and exposure were first introduced in the context of non-personalized rankings, where the objects being ranked are individual users, such as job candidates. These metrics can operate at the *individual* or *group* levels.

At the *individual* level these metrics are devoted to guaranteeing that similar individuals are treated similarly [11, 13]. For instance, Biega et al. [11] defined measures to capture unfairness at the level of individual subjects. Conversely, at the *group* level these metrics make sure that users belonging to different groups are given adequate visibility or exposure [13, 10, 12]. One example is the ranked group fairness definition presented in [10], which extends group fairness using the standard notion of protected groups. Zehlike and Castillo [12] describe an approach that measures discrimination and unequal opportunity in rankings at training time in terms of discrepancies in the average group exposure.

Under the group setting, the visibility/exposure of a group is proportional to its representation in the data [14, 15, 16, 17]. Since our study considers group settings, we embrace this class of metrics, assessing both the visibility and exposure in the recommendation lists.

2.2. Impact of Recommendations for Providers

The concepts of visibility and exposure have a direct impact on the providers behind the recommended items. *Provider fairness (P-fairness)* is the impact of the generated recommendations on the item providers. It guarantees that the providers of the recommended objects that belong to different groups are similar at the individual level, will get recommended according to their representation in the data. Provider fairness was mostly tackled through post-processing approaches.

Defining when a user or a group of users gets discriminated by an Artificial Intelligence (AI) system highly depends on the context that is being studied [18, 19, 20, 21].

In the context of books recommendation, Ekstrand et al. [22] assessed that collaborative filtering algorithms recommend author's books of a given gender with a distribution that differs from that of the original user profiles. Liu and Burke [23] consider P-fairness in the Kiva.org platform, which grants loans to low-income entrepreneurs. It is achieved through a re-ranking function (based on xQuad), which balances recommendation accuracy and fairness, by dynamically adding a bonus to the items of the uncovered providers. In the same domain, Sonboli and Burke [24] define the concept of local fairness, to identify protected groups through consideration of local conditions. This is done to avoid discriminating between types of loans and to equalize access to capital across all types of businesses. Abdollahpouri et al. [25] analyze the unfairness of popularity bias in movies recommendation, while Kowald et al. [26] analyze the same problem in the music domain.

Mehrotra et al. [27] assess unfairness based on the popularity of the providers. More specifically, they focus on a two-sided marketplace, with the consumers being users who listen to music, and the artists being the providers. If only highly popular artists are recommended to users, this creates a disadvantage for the less popular ones. For this reason, artists are divided into ten bins based on their popularity, and a fairness metric that rewards recommendation lists that are diverse in terms of popularity bins is defined. Several policies are defined to study the trade-offs between user relevance and fairness, with the ones that balance the two aspects being those who achieve the best trade-off.

Several policies are defined to study the trade-offs between user relevance and fairness. Kamishima et al. [28] introduce the concept of recommendation independence. Given a sensitive feature (which can be associated with the consumers, the providers, or the items), they present a framework to generate fair recommendations, in the sense that the outcome is statistically independent of a specified sensitive feature. Specifically, an objective function with three components (a loss function, an independence term, and a regularization term) is introduced, so that the prediction function returns an expected value of the loss function as small as possible and an independent term as large as possible.

2.3. Contextualizing our Work

To the best of our knowledge, this is the first time that unfairness phenomena for content providers belonging to different continents are tackled. Considering the UNESCO[3], our two study domains (i.e, cinema and literature) are powerful vehicles for culture, education, leisure, and propaganda. This report also highlights the importance of smaller cinematographic industries at the global level. In our movie dataset, India represents 0.004% of the total amount of items.

Moreover, both domains have an impact on the economy of a country, with (sometimes public) investments for the production of movies/books that are expected to generate a return. Hence, considering how recommender systems can push the consumption of items of a country is a related but different problem w.r.t. provider fairness.

In conclusion, studying the disparities emerging from the geographic imbalances in the composition of an industry is a problem that goes beyond the impact for content providers. Denying visibility and exposure to the items of a continent has a negative impact (i) on the cultural impact that a country can have and (ii) at an economic level.

3. Preliminaries

In this section we present the preliminaries, to provide foundations to our work. First, Section 3.1 details the recommendation scenario. Next, the metrics are described in Section 3.2. In Section 3.3, we present the recommendation algorithms. Finally, we describe the datasets used in this study in Section 3.4.

3.1. Recommendation Scenario

We consider a set of users, $U = \{u_1, u_2, ..., u_n\}$, a set of items, $I = \{i_1, i_2, ..., i_j\}$, and let V be a totally ordered set of values that can be used to express a preference together with a special symbol $\perp$. The set of ratings result from a map $r : U \times I \rightarrow V$, where V is the ratings' domain. If $r(u, i) = \perp$ then we say that user, u, did not rate item, i. To simplify notation, we denote $r(u, i)$ by r_{ui}. We define the set of ratings as $R =$

[3] https://publications.parliament.uk/pa/cm200203/cmselect/cmcumeds/667/667.pdf

$\{(u, i, r_{ui}) : u \in U, i \in I, r_{ui} \neq \perp\}$. These ratings can directly feed an algorithm in the form of triplets (point-wise approaches) or shape learner-course observations (pair-wise approaches).

To assess the real impact of the recommendations, we consider a temporal split of the data, where a fixed percentage of the ratings of the learners (ordered by timestamp) goes to the training and the rest goes to the test set [29].

The recommendation goal is to learn a function f that estimates the relevance ($\hat{r}_{ui}$) of the user-item pairs that do not appear in the training data (i.e., $r_{ui} = \perp$). We denote as $\hat{R}$ the set of recommendations, and as $\hat{R}_G$ those involving items of a group G, i.e., $\hat{R}_G = \{\hat{r}_{ui} : u \in U, i \in G \subseteq I\}$.

Let $A = \{a_1, a_2, ..., a_g\}$ denote the set of g geographic areas in which items are organized. Specifically, we consider a geographic area as the continent of provenience of each item provider. We denote as A_i the set of geographic areas associated with an item i. Note that, since the providers of an item could be from different geographical areas, several geographic areas may appear in an item, and thus, $|A_i| \geq 1$. In case two providers belong to the same geographic area, it appears only once. We use the geographic areas to shape k demographic groups, where the t^{th} demographic group is defined as $G_t = \{i \in I : a_t \in A_i\}$, for $t = 1, \ldots, g$. Finally, Table 1 summarizes the terminology used in this article.

3.2. Metrics

This section describes the metrics used in our analysis and experiments, i.e., the representation of a group, disparate visibility, and disparate exposure.

Representation. The representation of a group is the amount of times that group appears in the data. We consider two forms of representation, based on (i) the amount of items offered by a group and (ii) the amount of ratings collected for that group. We define with $\mathcal{R}$ the *representation* of a group G ($\mathcal{R}_I$ denotes an item-based representation, while $\mathcal{R}_R$ a rating-based representation):

$$\mathcal{R}_I(G) = |G|/|I| \tag{1}$$

$$\mathcal{R}_R(G) = |\{r_{ui} : u \in U, i \in G \subseteq I\}| / |R| \tag{2}$$

Table 1: Summary of the terminology used in the article. First column details the concept, while the second presents the notation for this concept.

Concept	**Term**
Set of users	U
Set of items	I
Set of preferences	V
Set of ratings	R
Rating of user u over item i	r_{ui}
Predicted relevance of item i for user u	$\hat{r}_{ui}$
Set of recommendations	$\hat{R}$
Set of recommendations involving items of a group G	$\hat{R}_G$
Set of geographic areas	A
Set of geographic areas associated with a item i	A_i
Demographic group	G_t
Item-based representation of a group	$\mathcal{R}_I(G)$
Rating-based representation of a group	$\mathcal{R}_R(G)$
Disparate visibility of a group	$\Delta\mathcal{V}(G)$
Disparate exposure of a group	$\Delta\mathcal{E}(G)$

Eq. (1) accounts for the proportion of items of a group, while Eq. (2) for the proportion of ratings associated with a group. Both metrics are between 0 and 1. We compute the representation of a group only considering the training set. Trivially, given a perspective (either item- or rating-based), the sum of the representations of all groups is equal to 1, $\sum_{i=1}^{k} \mathcal{R}_*(G_i) = 1$ (where '*' refers to I or R).

Disparate Impact. We assess unfairness with two notions of *disparate impact* generated by a recommender system. Specifically, we assess disparate impact with two metrics.

Definition 1 (Disparate visibility). *Given a group G, the* disparate visibility *returned by a recommender system for that group is measured as the difference between the share of recommendations for items of that group and the representation of that group in the input data:*

$$\Delta\mathcal{V}(G) = \left(\frac{1}{|U|} \sum_{u \in U} \frac{|\{\hat{r}_{ui} : \hat{r}_{ui} \in \hat{R}_G,\, i \in G \subseteq I\}|}{|\hat{R}|} \right) - \mathcal{R}_*(G) \tag{3}$$

where '*' refers to I (i.e., item-based representation) or R (i.e., rating-based representation). The range of values for this score is $[-\mathcal{R}_*(G), 1-\mathcal{R}_*(G)]$; specifically, it is 0 when the recommender system has no disparate visibility, while negative/positive values indicate that the group received a share of recommendations that is lower/higher than its representation. This metric is based on that considered by Fabbri et al. in [9].

Definition 2 (Disparate exposure). *Given a group G, the* disparate exposure *returned by a recommender system for that group is measured as the difference between the exposure given to that group in the recommendation lists [30] and its representation:*

$$\Delta\mathcal{E}(G) = \left(\frac{1}{|U|} \sum_{u \in U} \frac{\sum_{pos=1}^{k} \frac{1}{log_2\left(\hat{r}_G^u(pos)+1\right)}}{\sum_{pos=1}^{k} \frac{1}{log_2\left(\hat{r}_I^u(pos)+1\right)}} \right) - \mathcal{R}_*(G), \tag{4}$$

where $\hat{r}_G^u(pos)$ denotes the rating $\hat{r}_{ui}$ that takes position pos in the list $\hat{R}_G^u = \{\hat{r}_{vi} : v = u, i \in G \subseteq I\}$, $u \in U$, sorted by decreasing order.

This metric also ranges in $[-\mathcal{R}_*(G), 1-\mathcal{R}_*(G)]$ range; concretely, a value equal to 0 indicates that the recommender system has no disparate exposure, while negative/positive values indicate that the exposure given to the group is lower/higher than its representation in the data.

Remark. *Since the goal of our paper is to* allow the items of a group to be recommended enough times (visibility) and with enough exposure*, a unique "disparate impact" metric would not allow us to balance both perspectives, by combining everything in a single number. For this reason, we keep both disparate visibility as disparate exposure as goals to enable provider fairness in the context of geographic imbalance.*

3.3. Recommendation Algorithms

In this study, we focus on well-known state-of-the-art Collaborative Filtering approaches. In particular, we focus on both classes of point-wise and pair-wise approaches, and considered memory-based and model-based algorithms.

As *memory-based approaches*, we consider two approaches: **UserKNN** [31] and **ItemKNN** [32] algorithms. UserKNN [31] selects the K neighbors closest to the target user, and recommends the elements like by other users more similar to the target one. Similarly, ItemKNN [32] recommends to the target user those items that are more similar to other items that they liked before.

For the class of *matrix factorization based approaches*, we consider the **BPR** [33], **BiasedMF** [34], and **SVD++** [35] algorithms. Matrix factorization algorithms divide the data into matrices, representing them in latent factors to determine the degree of affinity that users and items have with those factors. In particular, Bayesian Personalized Ranking (in short, BPR) [33] is an algorithm optimized to generate recommendation lists, based on a Bayesian probability function. The preference function is based on the ratings of pairs of items. BiasedMF [34] performs basic factorization of the matrix that includes global mean, user bias, and item bias whereas SVD++ [35] takes into account the implicit interactions, as well as the user's latent factors and the item's latent factors.

Our baselines are two *non-personalized algorithms* (**MostPopular** and **RandomGuess**), which allow us to contextualize our results. MostPopular recommends items based on how their popularity in the dataset, by counting the number of times an item was rated. In this way, the algorithm considers only the item perspective, without associating the ratings to the individual users and their preferences. On the other hand, RandomGuess establishes the maximum and minimum rating in the data and returns a random rating for each user-item pair to predict.

3.4. Datasets

We analyze data from two different contexts, movies and books, exploring the role of the geographic provenience of providers in the recommendation process. In what follows, we describe the characteristics of each dataset:

- **MovieLens-1M (Movies).** The dataset provides 1M ratings (in the range [1-5]), given by 6,040 users to 3,600 movies. Each user rated at least 20 ratings. For each user, the dataset provides demographic information (namely, their gender, age, occupation, and zip code), which have not been considered for this study to focus on provider fairness; these attributes will be considered in future work, to study the interplay between provider fairness and the characteristics of the users. For each movie, the dataset provides its IMDb ID, which allowed us to associate it to its continent of production, thanks to the OMDB APIs[4]

[4] http://www.omdbapi.com/

(note that *each movie may have more than one continent of production*). The dataset also offers the title and genre of each movie, which are not relevant in the context of this study and, thus, have not been considered.

- **Book-Crossing (Books).** The dataset contains 356k ratings (in the range [1-10]), given by 10,409 users, to 14,137 books. Also this dataset provides demographic information about its users, by offering age and location attributes, if the user has provided them. Also in this case, these attributes are not relevant for our study, so any additional information of the users offered by this study will be considered in future work. For each book, the dataset provides its ISBN code; this code allowed us to retrieve information about the production continent, by exploiting the APIs offered by the Global Register of Publishers[5]. Additional book information such as its title, author, publisher, and cover image, is not relevant for our study and, thus, has not been considered.

We shape demographic groups considering the following continents: Africa, Asia, Europe, North America, Oceania, and South America for Movies, and Europe, North America, Oceania, and South America for Books. We remark that no items from Africa or Asia were available in the Books dataset and there are no items from the seventh continent (Antarctica) in both datasets.

4. Provider Fairness with a Binary Perspective

This section frames our previous study, which deals with provider fairness for two groups in presence of geographic imbalance. It summarizes the main observations obtained from our experiments to enhance the need for a cross-continent provider perspective.

In our previous work [7], we observed that imbalances in the data distribution can affect the visibility and exposure given to providers. Our study was focused on analyzing the items into two different groups based on the country of production, in a majority-versus-rest setting, and assessed if recommender systems generate a disparate impact and (dis)advantage a group. Recall that, in this work, our setting is focused on cross-continent provider fairness. Another observation extracted from our study is that the produced

[5] `https://grp.isbn-international.org/search/piid_cineca_solr`

recommendations by the recommender systems can amplify these imbalances and create biases. To study this phenomenon, we enriched two datasets and characterize data imbalance w.r.t. the country of production of an item (geographic imbalance). We conducted the experiments in two domains, movies (MovieLens-1M) and books (Book-Crossing), where both datasets are imbalanced towards the United States.

4.1. Group Representation

We assessed disparate impact by comparing the visibility and exposure given to a group of providers with the representation of the group in the data. As we are doing in this study, we studied two forms of representation, based on (*i*) the number of items a group offers, or (*ii*) the number of ratings given to the items of a group.

Let C_i be the set of production countries of an item i. We use it to shape two groups, a majority $G_M = \{i \in I : 1 \in C_i\}$, and a minority $G_m = \{i \in I : 1 \notin C_i\}$. Note that 1 identifies the country associated with the majority group.

In the binary perspective, the representation of the minority group in the Movies dataset is $\mathcal{R}_I(m) = 0.3$ and $\mathcal{R}_R(m) = 0.23$, considering item and rating, respectively. In the Books dataset, instead, the representation of the minority group is $\mathcal{R}_I(m) = 0.12$ and $\mathcal{R}_R(m) = 0.08$. As it can be observed, both datasets show a strong geographic imbalance, with the majority group covering 70% of the items in the first dataset and 88% in the second. This imbalance is worsened when we consider the ratings, since in the movie context the ratings associated with the majority are 77%, while in the book data the rating representation for the majority is 92%. However, the minority items are not considered as of lower quality for the users, since the average rating for both groups is nearly the same in both datasets. In the Movies dataset, the average rating for the majority group is 3.56, while that of the minority group is 3.61. In the Books dataset, we observed an average rating of 4.38 for the majority, and of 4.43 for the minority. This shows that the preference of the users for the two groups does not differ.

4.2. Metrics and Algorithms

We characterized both the visibility and exposure given to the providers of a group by a recommendation algorithm. To evaluate recommendation effectiveness, we measured the ranking quality of the lists by measuring the

Normalized Discounted Cumulative Gain (NDCG). We ran the state-of-the-art recommender systems described in Section 3.3, using the LibRec library. The test set was composed of the most recent 20% of the ratings of each user.

4.3. Assessment

In our initial analysis of both datasets (i.e, Movies and Books), the phenomenon that emerges is that both groups can be affected by disparate impact and that, when one group receives more visibility, it also receives more exposure; hence, when a group is favored in the number of recommendations, it is also ranked higher.

Concretely, the results showed the presence of a disparate impact that mostly favors the majority, since we feed algorithms with much more instances than their counterpart. However, factorization approaches are still capable of capturing the preferences for the minority items with latent factors, thus creating a positive impact for the group. But, if the imbalance is too severe, the minority is always affected by the disparate impact.

4.4. Approach

To mitigate disparities, we proposed a binary re-ranking approach that optimizes both the visibility and exposure given to providers in a binary (i.e., majority-vs-rest) setting, based on their representation in the data. Specifically, our approach introduces, in the recommendations, items that increase the visibility and exposure for the disadvantaged group, causing the minimum possible loss in user relevance. For each user, we generated 150 recommendations (denoted as the top-n) so that we can mitigate the disparate impact through a re-ranking algorithm. The final recommendation list for each user is composed of 20 items (denoted as top-k) and measured the visibility and exposure given to each group.

In what follows, we provide a summary of the steps followed by our re-ranking approach:

1. We identify the user causing the minimal loss in terms of items' predicted relevance;
2. We select two items in the list of the user, namely the last item of the advantaged group in the top-k and the first item of a disadvantaged group out of the top-k;
3. We swap the items and move to step 1 until the target visibility is reached.

After the target visibility is reached, we consider the top-k to regulate the exposure of the disadvantaged group. We swap inside the list the pair of items belonging to different groups that cause the minimum loss of predicted relevance, until the desired exposure for the disadvantaged group is reached.

4.5. Impact of Mitigation

Briefly, the impact of our proposed re-ranking algorithm for mitigating disparities in the binary setting is three-fold. First, our approach leads to the goal visibility and exposure. Given a target representation and a dataset, all algorithms achieve the same disparate visibility/exposure. Second, thanks to our mitigation based on the minimum-loss principle, the loss in NDCG was negligible. Finally, the most effective approach before mitigation is confirmed as such also after mitigation.

5. Disparate Impact Assessment

The first goal of this study is to evaluate the presence of unfairness in the state-of-the-art collaborative recommendation models, so as to understand if and where a problem exists. Concretely, our task is to analyze the recommendations these models generate and assess the presence of disparities in the way recommendations are distributed across different provider groups.

To accomplish this goal, in this section, we run the algorithms presented in Section 3.3 and measure their effectiveness and the disparate impact they generate for providers belonging to different continents.

At the end of this section, we will be able to understand which models create disparities and under which conditions.

5.1. Experimental Setting

For both datasets presented in Section 3.4, the test set was composed of the most recent 20% of the ratings of each user. To run the recommendation algorithms presented in Section 3.3, we considered the LibRec library (version 2). For each user, we generate 150 recommendations (denoted in the paper as the top-n) to then mitigate disparities through a re-ranking algorithm. The final recommendation list for each user is composed of 20 items (denoted as the top-k).

Each algorithm was run with the following hyper-parameters:

- **UserKNN.** We used Pearson similarity and 50 neighbors. The similarity shrinkage was set up to 10;

- **ItemKNN.** We used the Cosine similarity for the Movies dataset and Pearson similarity for the Books one. The number of neighbors was 200 for Movies and 50 for Books dataset. The similarity shrinkage was set up to 10;

- **BPR.** We configured the iterator learnrate to 0.1, the iterator learnrate maximum to 0.01, the iterator maximum to 150, the user regularization to 0.01; the item regularization to 0.01; the factor number to 10; the learnrate bolddriver to false, and the learnrate decay to 1.0;

- **BiasedMF.** We adjusted the iterator learnrate to 0.01, the iterator learnrate maximum to 0.01, the iterator maximum to 20 for the Movies dataset and 1 for the Books one, the user regularization to 0.01, the item regularization to 0.01; the bias regularization to 0.01, the number of factors to 10, the learnrate bolddriver to false, and the learnrate decay to 1.0;

- **SVD++.** We set up the iterator learnrate to 0.01, the iterator learnrate maximum to 0.01, the iterator maximum to 10 for the Movies dataset and 1 for the Books one, the user regularization to 0.01, the item regularization to 0.01, the impItem regularization to 0.001, the number of factors to 10, the learnrate bolddriver to false, and the learnrate decay to 1.0.

Recommendation effectiveness is assessed by measuring the ranking quality of the list, using the *Normalized Discounted Cumulative Gain* (NDCG) [36].

$$DCG@k = \sum_{u \in U} \hat{r}_{ui}^{pos} + \sum_{pos=2}^{k} \frac{\hat{r}_{ui}^{pos}}{log_2(pos)} \qquad NDCG@k = \frac{DCG@k}{IDCG@k} \tag{5}$$

where $\hat{r}_{ui}^{pos}$ is relevance of item i recommended to user u at position pos. The ideal DCG ($IDCG$) is calculated by sorting items based on decreasing true relevance (true relevance is 1 if the user interacted with the item in the test set, 0 otherwise). The higher the better.

5.2. Characterizing Representation

The first step towards the assessment of disparate impact is to characterize the representation of the different groups in the data, which we present

Table 2: **Group representation.** item-based ($\mathcal{R}_I$) and rating-based ($\mathcal{R}_R$) representations of each group (AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Groups appear in alphabetical order by the name of the continent.

	Movies		**Books**	
	$\mathcal{R}_I$	$\mathcal{R}_R$	$\mathcal{R}_I$	$\mathcal{R}_R$
AF	0.0038	0.0028	-	-
AS	0.0392	0.0234	-	-
EU	0.2469	0.1946	0.1043	0.0698
NA	0.6937	0.7659	0.8951	0.9299
OC	0.0139	0.0128	0.0005	0.0002
SA	0.0025	0.0003	0.0001	0.0001

in Table 2. Note that the Books dataset does not contain books from Africa and Asia.

We can observe that North America represents the most represented continent in both datasets, covering 69% of the produced items ($\mathcal{R}_I$) in the Movies data and almost 90% of the items in Books. This existing imbalance is increased if we consider the rating-based representation ($\mathcal{R}_R$), where North America has a share of 76.6% and 93% of the ratings, respectively in the Movies and Books data. This leads us to our first observation.

Observation 1. *Both datasets have a strong geographic imbalance towards North America, which is the most represented group from both item- and rating-based perspectives. The imbalance is strengthened when considered the rating-based representation, meaning that the largest groups attract a share of ratings that is even higher than the amount of items it offers. This clearly has a price for the smaller groups, which are able to attract a percentage of ratings that is lower than the amount of items they offer. Hence,* user-item interactions favor the largest group and exacerbate imbalances that already existed in the item offer, even before we run a recommendation algorithm.

5.3. Assessing Effectiveness and Disparate Impact

In this section, we assess the effectiveness (in terms of NDCG) and the disparate impact (both in terms of visibility and exposure) returned by the state-of-the-art algorithms. Moreover, we assess if the binary mitigation (for two groups) proposed in [7] is capable of enabling fairness for multiple groups shaped at the continent level.

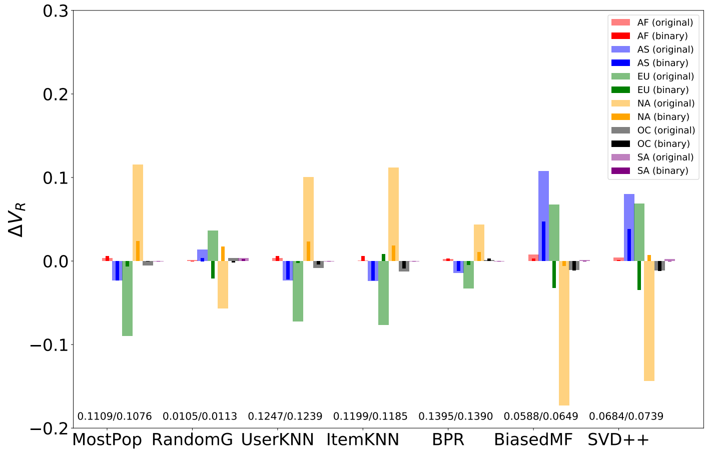

(a) Disparate visibility (ratings).

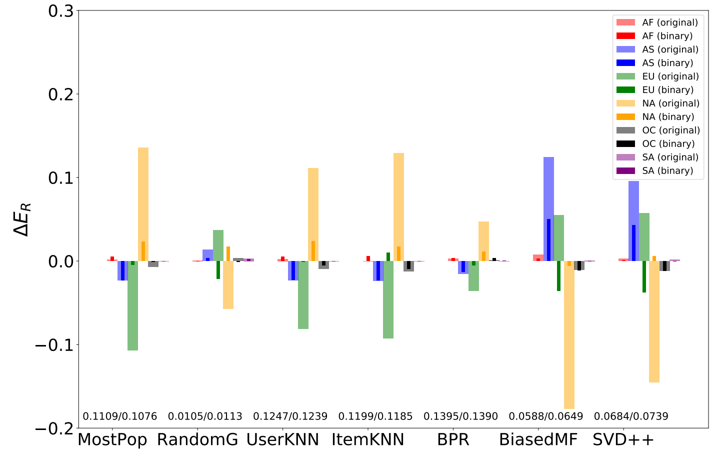

(b) Disparate exposure (ratings).

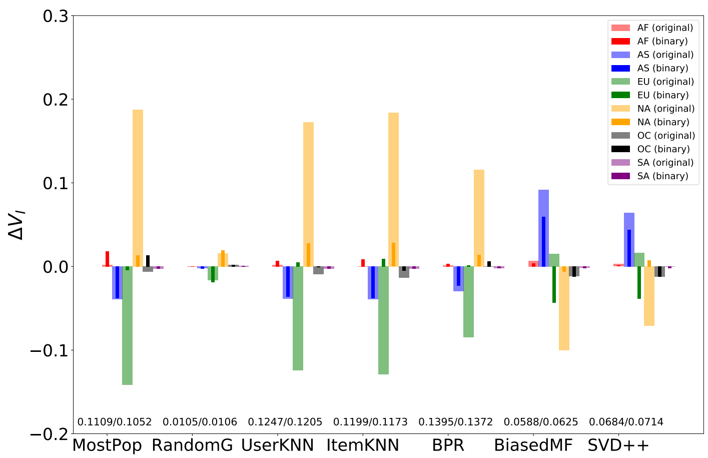

(c) Disparate visibility (items).

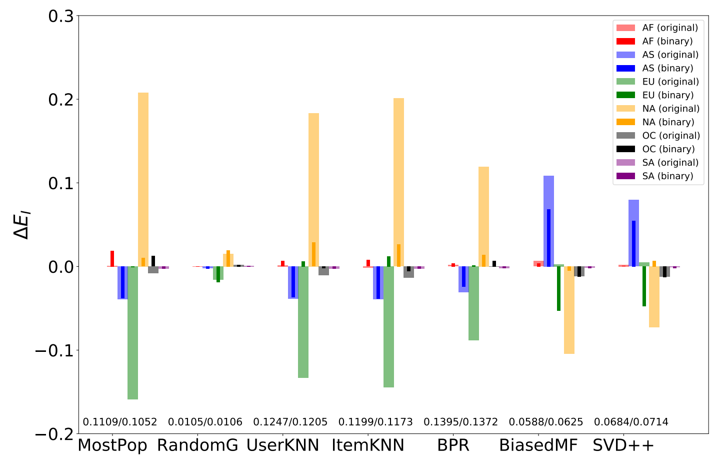

(d) Disparate exposure (items).

Figure 2: **Disparate impact in the Movies dataset**. Disparate impact returned by the state-of-the-art models (thick bars) and by the binary mitigation proposed in [7] (thin bars). Each figure contains one section for each algorithm and one color for each continent. The text at the bottom of each figure contains the NDCG returned by the original model and after the binary mitigation, separated by a "/". In (a) and (b), we report the disparate visibility and disparate exposure obtained when considering a rating-based representation, while in (c) and (d), the disparate visibility and disparate exposure obtained when considering an item-based representation.

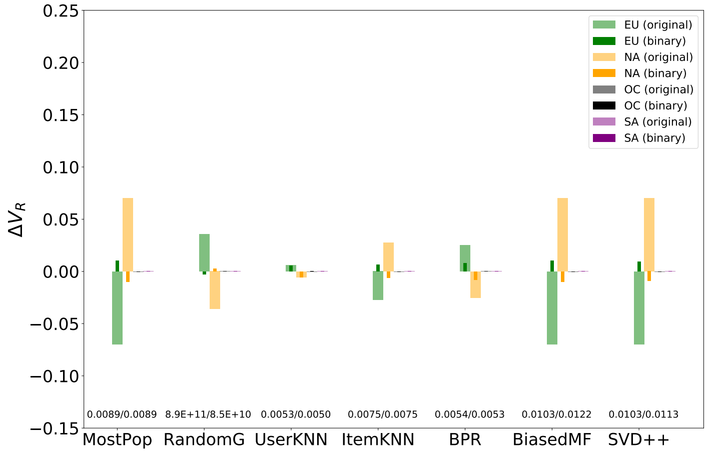

(a) Disparate visibility (ratings).

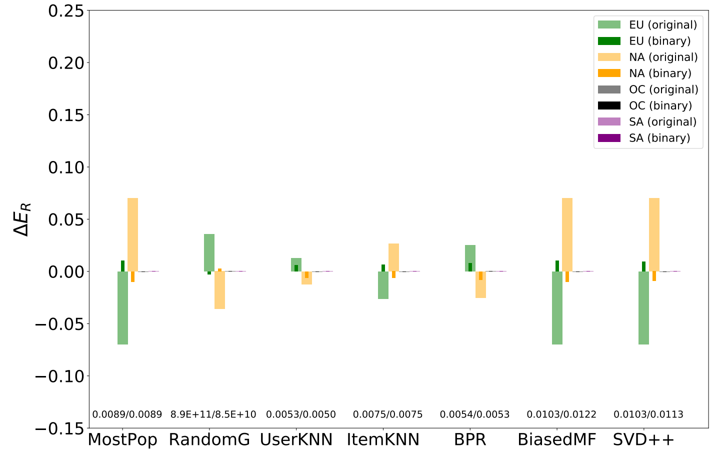

(b) Disparate exposure (ratings).

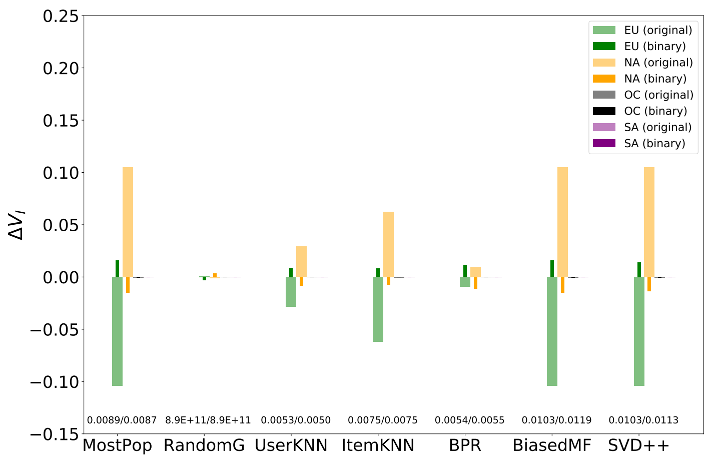

(c) Disparate visibility (items).

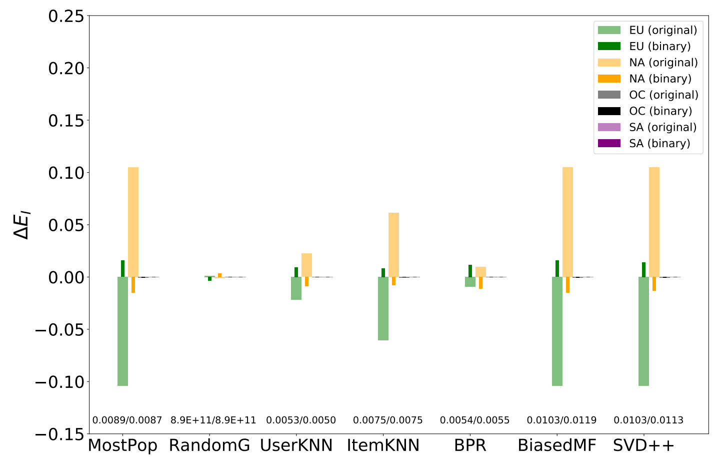

(d) Disparate exposure (items).

Figure 3: **Disparate impact in the Books dataset**. Disparate impact returned by the state-of-the-art models (thick bars) and by the binary mitigation proposed in [7] (thin bars). Each figure contains one section for each algorithm and a color for each continent. The text at the bottom of each figure contains the NDCG returned by the original model and after the binary mitigation, separated by a "/". In (a) and (b), we report the disparate visibility and disparate exposure obtained when considering a rating-based representation, while in (c) and (d), the disparate visibility and disparate exposure obtained when considering an item-based representation.

Results are visually reported in Figure 2 for the Movies dataset and in Figure 3 for the Books one. To present our results in a reproducible way, Tables 7 and 8 (placed in the Appendix) report the values that shape our figures, for Movies and Books respectively. If we look at how the original models (thick bars) behave according to the different types of representation, we can observe that they adjust better to the rating-based representation of the groups. Indeed, in both figures we can observe that in (c) and (d) the disparities returned when considering an item-based representation are more prominent than their rating-based counterparts, in (a) and (b). In other words, recommendation models adapt better to the interactions between users and items than to the amount of items a group has to offer. The only exception to this is RandomGuess which, by picking items at random, better adapts to the distribution of the items. Indeed, as subfigures (c) and (d) show, it is the approach that returns the most equitable results, in both Figures 2 and 3. Nevertheless, as it is shown by the NDCG values at the bottom, it is also the least effective approach. One can notice that the thin bars, reporting the results after the binary mitigation, are generally closer to 0 than the original models, indicating that, even though we are not providing fairness at the continent level, still the approach in [7] distributes the recommendation in a more equitable way than the original models.

Going more in-depth to the behavior of the system in each domain, Figure 2, shows that for Movies the algorithm that better adjusts to the different continents, when considering the rating-based representation is BPR; indeed, the algorithm is the most effective one, returning the highest NDCG, and the one that adjusts better to the desired equity in terms of disparate visibility and exposure. One interesting phenomenon we can observe is that, when a model over- or under-recommends one of the two most represented groups (i.e., North America and Europe), the other is directly affected. We can see this clear pattern for MostPop, RandomGuess, UserKNN, and ItemKNN, who create disparate visibility and exposure at the advantage of North America, while affecting the most the second most represented country (Europe). On the contrary, when Europe is over-recommended, North America is the most affected country, as shown by BiasedMF and SVD++. This last phenomenon we observed for the point-wise recommendation models (BiasedMF and SVD+) connects to the observations of Cremonesi et al. [37], who showed the capability of factorization approaches to recommend long-tail items. Disparate impact clearly affects less the smallest groups, since they have a representation almost equal to 0, which is reflected in the visibility and exposure they are given; nevertheless, when a country is under- or over-recommended, with disparate visibility and exposure values lower/higher than 0, (e.g., Oceania and Africa), all models follow the same pattern.

Moving to the behavior of the binary mitigation proposed in [7], clearly, since our original mitigation was based on providing the United States with an equitable number of recommendations, North America is the continent that benefits the most from our original mitigation, especially with BiasedMF and SVD++.

Moving to our Books data, in Figure 3 we can observe that BiasedMF is the most effective approach. One interesting aspect we can notice here is that not even the point-wise matrix factorization based models are able to contrast the imbalance, favoring North America in terms of visibility and exposure. This leads us to our second observation.

> **Observation 2**. *Recommendation models better adjust to the rating distribution than to the item offer associated with a group. Factorization-based approaches are able to account for the needs of smaller groups, unless the imbalance in the input data is too severe. Even though we are working in a multi-group setting, recommender systems mostly operate as if two big groups existed; when one group is favored, the other is affected, both in terms of visibility and exposure. A mitigation for binary groups helps reducing disparities, but is not enough to introduce fairness for groups shaped at the continent level. Indeed, integrating more recommendations of the items from the minority group does not ensure that these recommendations are equally distributed among the different continents, so disparities still emerge.*

6. Mitigating Disparate Impact

In the previous section, we assessed the presence of disparities, having a negative impact mainly for the less represented provider groups. To overcome this limitation and introduce provider fairness for the different geographic groups, in this section, we detail the motivation of our approach and describe the re-ranking algorithm proposed to mitigate disparities at the continent level.

6.1. Motivation Behind our Approach

From the previous section, considering the representation of each group in the data, we noticed that some groups receive disproportional visibility and exposure. As a result of this observation, we propose to mitigate disparities with a re-ranking algorithm. Specifically, the goal of the proposed algorithm is to reach a visibility and exposure for each group proportional to their representation by moving items of the disadvantaged groups in the recommendation list.

A re-ranking algorithm is the unique option when optimizing metrics such as visibility and exposure. We cannot perform an in-processing regularization (e.g., [28, 38]) because at the prediction stage it is not possible to know if and where an item is ranked in a recommendation list. For this reason, it is not possible to do a comparison with this class of approaches. The reason why this comparison is not possible is not due to the algorithms we chose in our study, as this consideration also applies to list-wise approaches. Note that re-ranking algorithms have been introduced in the context of recommendation [27, 39, 40] as well as in non-personalized rankings [10, 30, 11, 41, 12, 14], but all of them are optimizing just one metric (i.e, visibility or exposure). Hence, in this section, we present an approach that provides fairness guarantees to all the provider groups in the data, considering both visibility and exposure metrics.

6.2. Algorithm

Our mitigation algorithm is based on the idea of *promoting in the recommendation list the item, that considering all the users, minimizes the loss in prediction*. Algorithm 1 describes pipeline followed by our mitigation method and Algorithm 2 presents our regulation of visibility and exposure in the recommendation lists. Finally, Algorithm 3 presents the support methods that are called by our mitigation method. Algorithm 1 takes as inputs (i) the recommendation list, *recList*, for all the users (consisting of the top-n items) and (ii) how recommendations should be distributed across continents after the mitigation ($targetProportions$). The output is the new list of re-ranked items, $reRankedList$.

Algorithm 1 consists of one main method, called $optimizeContinentsVisibilityExposure$ (lines 1-6). It makes two interventions, one based on visibility and the second one based on exposure. After each method is called, it returns the recommendation list, optimized for visibility (line 3) and exposure (line 4).

Algorithm 2 contains the method that performs our mitigation process, called $mitigationContinent$ (lines 1-34). Concretely, the method regulates, in a recommendation list, the visibility or exposure, so that it reaches the representation of each continent.

```
Input: recList: ranked list (records contain user, item, prediction, exposure, continent,
       position)
       targetProportions: list with the target proportions of each continent
Output: reRankedList: ranked list adjusted by visibility and exposure
define optimizeContinentsVisibilityExposure (recList, targetProportions)
begin
    // mitigation to target the desired visibility
    reRankedList ← mitigationContinent(recList, "visibility", targetProportions);
    // mitigation to regulate the exposure
    reRankedList ← mitigationContinent(reRankedList, "exposure", targetProportions);
    return reRankedList ; // re-ranked list adjusted by visibility and exposure
end
```

Algorithm 1: Muticlass mitigation algorithm based on Visibility and Exposure

After setting some supporting data structures (line 3) and assessing the current disparity we observe for each continent (lines 4 and 5), in lines 6-20, we create two lists of candidate items, respectively to be removed from and added in the recommendation list, named *itemsOut* and *itemsIn*. Concretely, the first list contains items currently recommended to the user that belong to an advantaged group, while the second contains items of a disadvantaged group currently not recommended to the user. In lines 14-19, we create a list, named *possibleSwaps*, containing pairs of candidate items that cause the minimum possible loss in terms of predicted relevance for the users. This list is sorted by loss in line 21. Finally, in lines 23-32, we swap the items and update the proportions, until we reach the desired visibility or exposure in the recommendation list. The re-ranked list is returned in line 33.

Finally, Algorithm 3, contains the methods we call in Algorithm 2. Concretely, the *checkPosition* method (lines 1-5) is responsible for checking the position of an item in the list, taking into account if we perform a visibility- or exposure-based mitigation. The *checkDisadvantagedGroup* method (lines 6-10) verifies whether the item belongs to a disadvantaged continent or not. Note that the method contains a for loop, since multiple continents may occur in an item. In that case, we compute the total sum of disparities to define a global disparity of the item. The method returns true when the disparity is positive, false otherwise. The *initialProportions* method (lines 11-24), returns the visibility and exposure of each continent before running our mitigation. The last method, *updatePositions* (lines 25-33) is responsible for updating the visibility and exposure given to a group after an item is added to the recommendation list.

7. Impact of Mitigation

The final goal of our study is to assess if the approach we presented in Section 6 is capable of providing fairness, by distributing the recommendations in equitable ways between the different provider groups. Moreover, we want to analyze if our approach can accomplish this goal in a better way than the existing approaches at the state of the art. Concretely, our task is to analyze the recommendations generated after running our mitigation strategy and a well-known state-of-the-art algorithm, to assess if disparities are still present and where and how recommendation effectiveness is impacted by our mitigation.

To accomplish this goal, in this section we analyze the impact of our mitigation approach, by assessing recommendation effectiveness and the presence of disparate impact for providers belonging to different continents. Section 7.1 shows the results of our mitigation algorithm and the advantages of employing an approach that can account for the presence of multiple groups, rather than the binary-group perspective proposed in [7]. Next, in Section 7.2, we compare our proposal against a well-known re-ranking approach, proposed in [23].

7.1. Impact of Mitigating for Multiple Groups

In this section, we analyze the impact of our mitigation algorithm for multiple groups, analyzing both the recommendation effectiveness and the visibility and exposure given to the different groups.

We report our results in Figure 4 for the Movies dataset, and in Figure 5 for the Books one. To present our results in a reproducible way, Tables 9 and 10 (see Appendix) report the values that shape our figures, for Movies and Books respectively.

One aspect that can be appreciated is that, given a reference representation and a dataset, all algorithms disparities are almost equal to 0, indicating we can provide a fair distribution of the recommendations, based on the distribution of the continents in the input data. This can be noticed by observing the thin bars in each subfigure.

Let us consider the trade-off between disparate visibility/exposure and effectiveness. Considering the Movies data (Figure 4), in both the rating- (subfigures a and b) and item-based (subfigures c and d) representations of the groups, BPR is the algorithm with the best trade-off between effectiveness and equity of visibility and exposure. It was already the most accurate

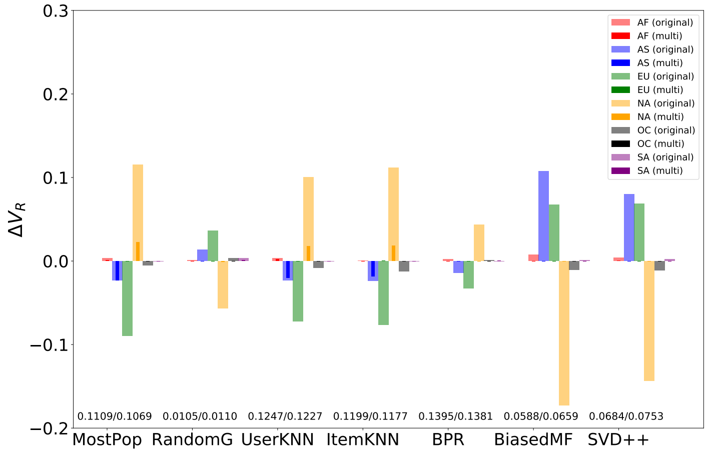

(a) Disparate visibility (ratings).

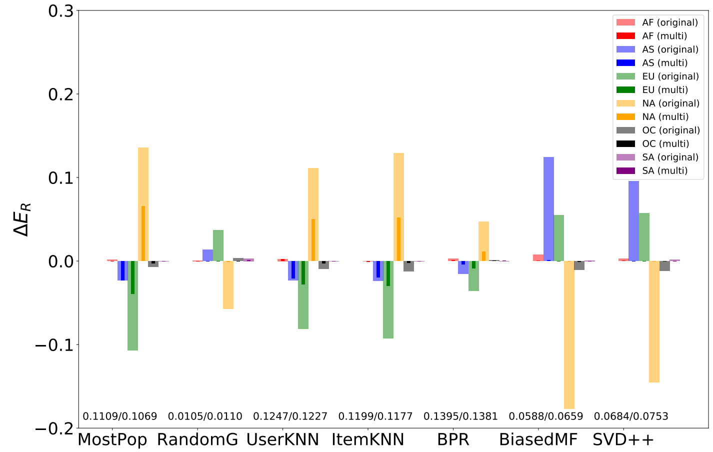

(b) Disparate exposure (ratings).

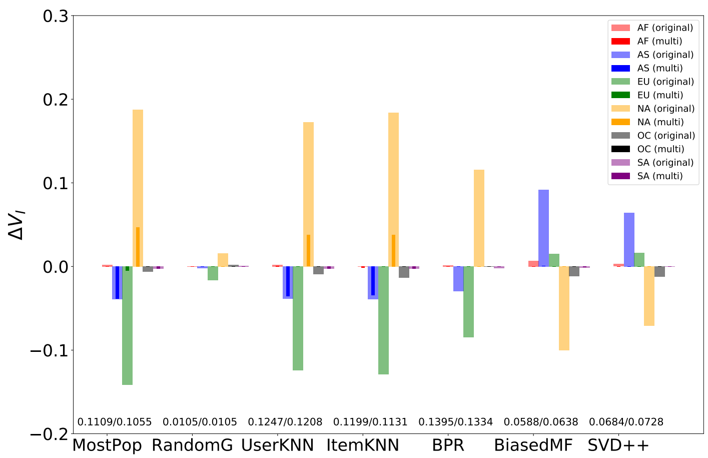

(c) Disparate visibility (items).

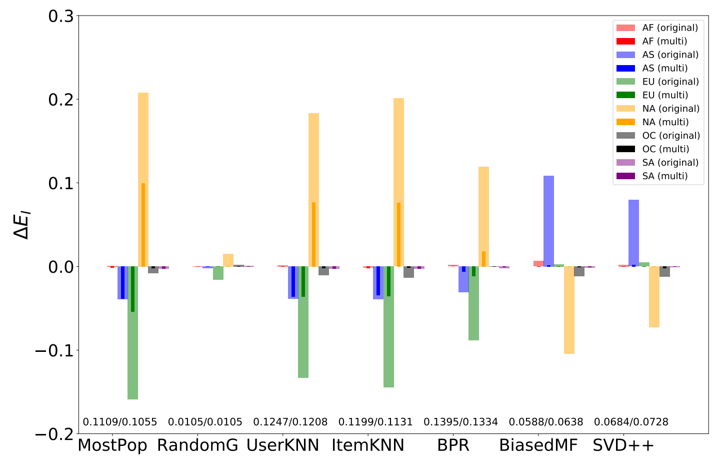

(d) Disparate exposure (items).

Figure 4: **Disparate impact in the Movies dataset after mitigation**. Disparate impact returned by the state-of-the-art models (thick bars) and by our multi-group mitigation algorithm (thin bars). Each figure contains one section for each algorithm and a color for each continent. The text at the bottom of each figure contains the NDCG returned by the original model and after the binary mitigation, separated by a "/". In (a) and (b), we report the disparate visibility and disparate exposure obtained when considering a rating-based representation, while in (c) and (d), the disparate visibility and disparate exposure obtained when considering an item-based representation.

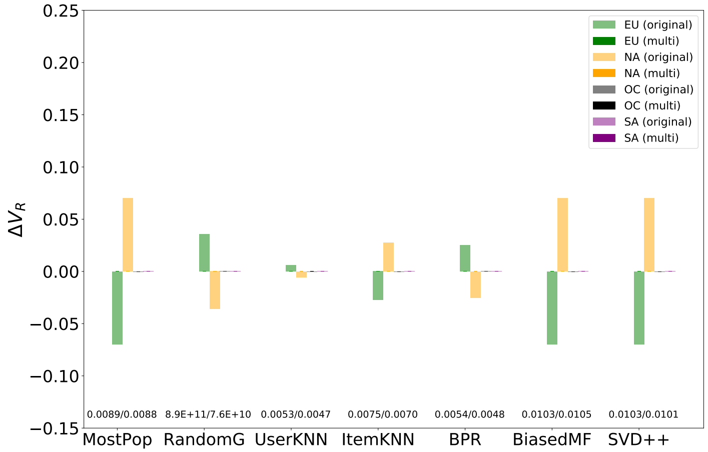

(a) Disparate visibility (ratings).

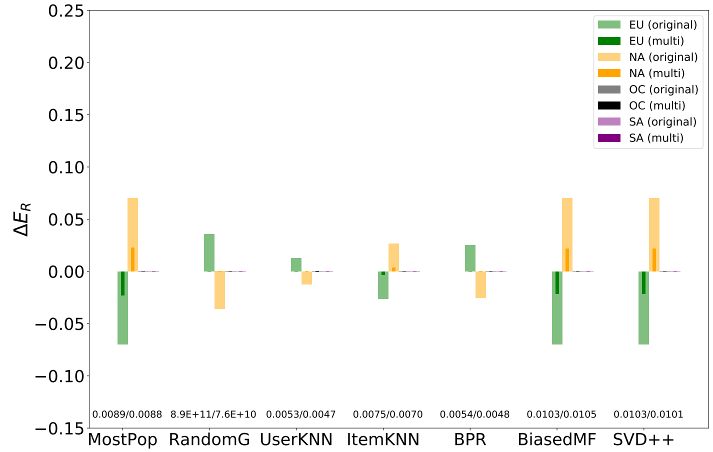

(b) Disparate exposure (ratings).

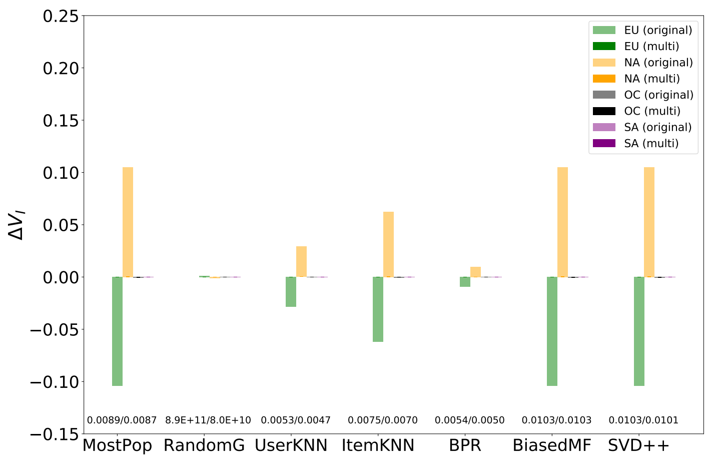

(c) Disparate visibility (items).

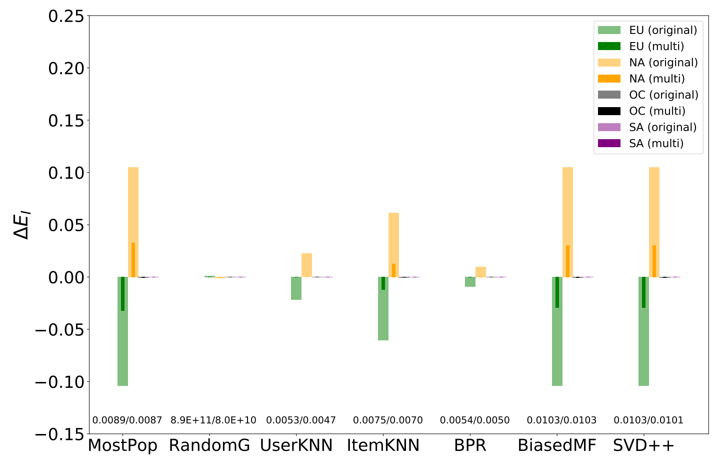

(d) Disparate exposure (items).

Figure 5: **Disparate impact in the Books dataset after mitigation**. Disparate impact returned by the state-of-the-art models (thick bars) and by our multi-group mitigation algorithm (thin bars). Each figure contains one section for each algorithm and a color for each continent. The text at the bottom of each figure contains the NDCG returned by the original model and after the binary mitigation, separated by a "/". In (a) and (b), we report the disparate visibility and disparate exposure obtained when considering a rating-based representation, while in (c) and (d), the disparate visibility and disparate exposure obtained when considering an item-based representation.

algorithm, and thanks to our mitigation based on the minimum-loss principle, the loss in NDCG was negligible. Moving to the Books dataset (Figure 5), BiasedMF confirms to be the best approach, in both effectiveness and equity of visibility and exposure. It is interesting to observe that, in both scenarios, MostPop is the second most effective algorithm and now provides the same visibility and exposure as the other algorithms; we conjecture that this might be due to popularity bias phenomena [3], and their analysis is left as future work.

> **Observation 3**. *When providing a re-ranking based on minimal predicted loss, the effectiveness remains stable, but disparate visibility and disparate exposure are mitigated. The most effective approach remains the best one after the mitigation.*

In Section 5.3, we have shown that a mitigation considering only two groups is not enough to introduce equity in the presence of multiple groups. To assess the difference the benefits of considering multiple groups in a mitigation strategy, in Tables 3 and 4, we compare the results we can obtain with our multi-group mitigation and the one considering only a binary perspective, in the Movies and Books datasets, respectively. The results clearly show that a mitigation at a more fine-grained granularity can provide fairness to providers in different groups.

7.2. Comparison with the State of the Art

We compare the results of our mitigation with that proposed in [23]. This approach aims at introducing provider fairness via a re-ranking approach, as our approach. Differently from us, in the mitigation proposed in [23] the predicted relevance is increased if a provider has not appeared yet in the top-k of a user. Since we are dealing with a provider fairness setting, we increase the predicted rating if a geographic area has not appeared yet in the ranking of a user. We remind readers to [23] for the technical details of the re-ranking approach we compare with. Hyperparameter λ of the original algorithm proposed in [23] was set to 2.

Tables 5 and 6 report the obtained results, where *multi* refers to our re-ranking multi-group mitigation algorithm and *baseline* is the compared algorithm.

We observe that our approach in most cases is capable of introducing equity by mitigating both disparate visibility and exposure in all algorithms. In general, our algorithm achieves better disparities than the baseline (indeed, in our results disparities are almost always close to 0). The baseline algorithm is able to minimize the disparities for those groups that are more represented but not for the less represented ones. Our proposal reduces slightly disparities with respect to the baseline in Most Popular, UserKnn, and ItemKnn only in South America (SA), which is the group with the smallest representation. In the remaining continents and algorithms, our proposal is highly effective in mitigating the disparities. We consider that the baseline is not mitigating both visibility and exposure to a greater extent because it favors the introduction, in the top-k, of items produced in more than one geographic group. This means that, while a disadvantaged group might gain visibility and/or exposure, the accompanying group also receives the same treatment, even though it might be advantaged.

Observation 4. *Introducing provider fairness requires interventions at the recommendation-list level. Mitigating by boosting predicted relevance for the disadvantaged groups does not provide guarantees of equity of visibility and exposure are fully mitigated. Disparities are only partially mitigated.*

8. Conclusions and Future Work

Recommender systems usually emphasize biases that emerge because of the way data has been collected. In this work, we focused on a scenario in which imbalances are associated with the way an industry is composed, with certain geographic areas that produce more items of certain types. This is the case for movies and books, which have been the main use-cases in our work.

Concretely, we assessed how recommender systems dealt with data imbalances, studying their capability to recommend items of providers coming from different continents and possible unfairness phenomena emerging from the way recommendations are distributed. We considered state-of-the-art collaborative filtering models and assessed that all of them create disparities in the way recommendations are produced, both in terms of visibility and of exposure given to providers.

To overcome these phenomena, we analyzed a binary re-ranking approach [7], which improves geographic imbalance in a binary (i.e., majority-vs-rest) setting by maintaining recommendation effectiveness. However, we have observed that in a group setting the approach does not reach equity for all groups. Accordingly to this observation, we proposed a multi-group re-ranking approach that re-distributes the recommendation across provider groups (i.e., geographic continents) based on a notion of equity, that assigns to each group a share of recommendation proportional to its representation in the input data. Experimental results show that our approach can introduce provider fairness without affecting recommendation effectiveness.

Considering that in this study we observed that the mitigation of data imbalances needs intervention at a fine-grained level, in future work we will assess the interplay between the representation of individual providers and the geographic area they belong to. Concretely, we will consider settings with more fine-grained groups (e.g., at country level), to assess if, with more groups, each with a lower representation in the data, our approach can still enable fairness for provider groups, and possibly refine our approach. Moreover, we will consider additional domains such as education, to explore disparities for teachers [21, 42, 20, 43]. Finally, we will also consider other issues emerging from imbalanced groups, such as bribing [44, 45].

References

[1] F. Ricci, L. Rokach, B. Shapira, Recommender systems: Introduction and challenges, in: Recommender Systems Handbook, Springer, 2015, pp. 1–34. doi:10.1007/978-1-4899-7637-6_1.

[2] L. Boratto, M. Marras, Advances in bias-aware recommendation on the web, in: Proceedings of the 14th ACM International Conference on Web Search and Data Mining, WSDM '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 1147–1149. doi:10.1145/3437963.3441665.
URL https://doi.org/10.1145/3437963.3441665

[3] L. Boratto, G. Fenu, M. Marras, Connecting user and item perspectives in popularity debiasing for collaborative recommendation, Inf. Process. Manag. 58 (1) (2021) 102387. doi:10.1016/j.ipm.2020.102387.
URL https://doi.org/10.1016/j.ipm.2020.102387

[4] S. Hajian, F. Bonchi, C. Castillo, Algorithmic bias: From discrimination discovery to fairness-aware data mining, in: Proceedings of the

22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2016, pp. 2125–2126. doi:10.1145/2939672.2945386.

[5] C. Bauer, E. Zangerle, Leveraging multi-method evaluation for multi-stakeholder settings, CoRR abs/2001.04348 (2020). arXiv:2001.04348.

[6] C. Bauer, M. Schedl, Global and country-specific mainstreaminess measures: Definitions, analysis, and usage for improving personalized music recommendation systems, PLOS ONE 14 (6) (2019) 1–36. doi:10.1371/journal.pone.0217389.

[7] E. Gómez, L. Boratto, M. Salamó, Disparate impact in item recommendation: A case of geographic imbalance, in: D. Hiemstra, M. Moens, J. Mothe, R. Perego, M. Potthast, F. Sebastiani (Eds.), Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 - April 1, 2021, Proceedings, Part I, Vol. 12656 of Lecture Notes in Computer Science, Springer, 2021, pp. 190–206. doi:10.1007/978-3-030-72113-8_13.
URL https://doi.org/10.1007/978-3-030-72113-8_13

[8] E. Walster, E. Berscheid, G. W. Walster, New directions in equity research., Journal of personality and social psychology 25 (2) (1973) 151.

[9] F. Fabbri, F. Bonchi, L. Boratto, C. Castillo, The effect of homophily on disparate visibility of minorities in people recommender systems, in: Proceedings of the Fourteenth International AAAI Conference on Web and Social Media, ICWSM 2020, AAAI Press, 2020, pp. 165–175.

[10] M. Zehlike, F. Bonchi, C. Castillo, S. Hajian, M. Megahed, R. Baeza-Yates, Fa*ir: A fair top-k ranking algorithm, in: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, ACM, 2017, pp. 1569–1578. doi:10.1145/3132847.3132938.

[11] A. J. Biega, K. P. Gummadi, G. Weikum, Equity of attention: Amortizing individual fairness in rankings, in: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, ACM, 2018, pp. 405–414. doi:10.1145/3209978.3210063.

[12] M. Zehlike, C. Castillo, Reducing disparate exposure in ranking: A learning to rank approach, in: WWW '20: The Web Conference 2020, ACM / IW3C2, 2020, pp. 2849–2855. doi:10.1145/3366424.3380048.

[13] F. Diaz, B. Mitra, M. D. Ekstrand, A. J. Biega, B. Carterette, Evaluating stochastic rankings with expected exposure, CoRR abs/2004.13157 (2020). arXiv:2004.13157.

[14] G. K. Patro, A. Biswas, N. Ganguly, K. P. Gummadi, A. Chakraborty, Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms, in: WWW '20: The Web Conference 2020, ACM / IW3C2, 2020, pp. 1194–1204. doi:10.1145/3366423.3380196.

[15] P. Sapiezynski, W. Zeng, R. E. Robertson, A. Mislove, C. Wilson, Quantifying the impact of user attentionon fair group representation in ranked lists, in: Companion of The 2019 World Wide Web Conference, WWW 2019, ACM, 2019, pp. 553–562. doi:10.1145/3308560.3317595.

[16] K. Yang, J. Stoyanovich, Measuring fairness in ranked outputs, in: Proceedings of the 29th International Conference on Scientific and Statistical Database Management, ACM, 2017, pp. 22:1–22:6. doi:10.1145/3085504.3085526.

[17] G. Ramos, C. Caleiro, A novel similarity measure for group recommender systems with optimal time complexity, in: L. Boratto, S. Faralli, M. Marras, G. Stilo (Eds.), Bias and Social Aspects in Search and Recommendation - First International Workshop, BIAS 2020, Lisbon, Portugal, April 14, 2020, Proceedings, Vol. 1245 of Communications in Computer and Information Science, Springer, 2020, pp. 95–109. doi:10.1007/978-3-030-52485-2_10.

[18] K. Holstein, S. Doroudi, Fairness and equity in learning analytics systems (fairlak), in: Companion Proceedings of the Ninth International Learning Analytics & Knowledge Conference (LAK 2019), 2019.

[19] B. Green, L. Hu, The myth in the methodology: Towards a recontextualization of fairness in machine learning, in: Proceedings of the machine learning: the debates workshop, 2018.

[20] D. Dessì, G. Fenu, M. Marras, D. Reforgiato Recupero, Leveraging cognitive computing for multi-class classification of e-learning videos, in:

The Semantic Web: ESWC 2017 Satellite Events, Revised Selected Papers, Vol. 10577 of Lecture Notes in Computer Science, Springer, 2017, pp. 21–25. doi:10.1007/978-3-319-70407-4_5.

[21] S. Barra, M. Marras, G. Fenu, Continuous authentication on smartphone by means of periocular and virtual keystroke, in: M. H. Au, S. Yiu, J. Li, X. Luo, C. Wang, A. Castiglione, K. Kluczniak (Eds.), Network and System Security - 12th International Conference, NSS 2018, Hong Kong, China, August 27-29, 2018, Proceedings, Vol. 11058 of Lecture Notes in Computer Science, Springer, 2018, pp. 212–220. doi:10.1007/978-3-030-02744-5_16.

[22] M. D. Ekstrand, M. Tian, M. R. I. Kazi, H. Mehrpouyan, D. Kluver, Exploring author gender in book rating and recommendation, in: Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, ACM, 2018, pp. 242–250. doi:10.1145/3240323.3240373.

[23] W. Liu, R. Burke, Personalizing fairness-aware re-ranking, CoRR abs/1809.02921 (2018). arXiv:1809.02921.

[24] N. Sonboli, R. Burke, Localized fairness in recommender systems, in: Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization, UMAP 2019, ACM, 2019, pp. 295–300. doi:10.1145/3314183.3323845.

[25] H. Abdollahpouri, M. Mansoury, R. Burke, B. Mobasher, The unfairness of popularity bias in recommendation, in: Workshop on Recommendation in Multi-stakeholder Environments (RMSE2019), in conjunction with the 13th ACM Conference on Recommender Systems, RecSys 2019, 2019.

[26] D. Kowald, M. Schedl, E. Lex, The unfairness of popularity bias in music recommendation: A reproducibility study, in: J. M. Jose, E. Yilmaz, J. Magalhães, P. Castells, N. Ferro, M. J. Silva, F. Martins (Eds.), Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part II, Vol. 12036 of Lecture Notes in Computer Science, Springer, 2020, pp. 35–42. doi:10.1007/978-3-030-45442-5_5.
URL https://doi.org/10.1007/978-3-030-45442-5_5

[27] R. Mehrotra, J. McInerney, H. Bouchard, M. Lalmas, F. Diaz, Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, ACM, 2018, pp. 2243–2251. doi:10.1145/3269206.3272027.

[28] T. Kamishima, S. Akaho, H. Asoh, J. Sakuma, Recommendation independence, in: Conference on Fairness, Accountability and Transparency, FAT 2018, Vol. 81 of Proceedings of Machine Learning Research, PMLR, 2018, pp. 187–201.

[29] A. Bellogín, P. Castells, I. Cantador, Statistical biases in information retrieval metrics for recommender systems, Inf. Retr. Journal 20 (6) (2017) 606–634.

[30] A. Singh, T. Joachims, Fairness of exposure in rankings, in: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, ACM, 2018, pp. 2219–2228. doi:10.1145/3219819.3220088.

[31] J. L. Herlocker, J. A. Konstan, J. Riedl, An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms, Inf. Retr. 5 (4) (2002) 287–310. doi:10.1023/A:1020443909834.

[32] B. M. Sarwar, G. Karypis, J. A. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in: Proceedings of the Tenth International World Wide Web Conference, WWW 10, ACM, 2001, pp. 285–295. doi:10.1145/371920.372071.

[33] S. Rendle, C. Freudenthaler, Z. Gantner, L. Schmidt-Thieme, BPR: bayesian personalized ranking from implicit feedback, in: UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, AUAI Press, 2009, pp. 452–461.

[34] Y. Koren, R. M. Bell, C. Volinsky, Matrix factorization techniques for recommender systems, IEEE Computer 42 (8) (2009) 30–37. doi:10.1109/MC.2009.263.

[35] Y. Koren, Factorization meets the neighborhood: a multifaceted collaborative filtering model, in: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2008, pp. 426–434. doi:10.1145/1401890.1401944.

[36] K. Järvelin, J. Kekäläinen, Cumulated gain-based evaluation of IR techniques, ACM Trans. Inf. Syst. 20 (4) (2002) 422–446. doi:10.1145/582415.582418.

[37] P. Cremonesi, Y. Koren, R. Turrin, Performance of recommender algorithms on top-n recommendation tasks, in: Proceedings of the 2010 ACM Conference on Recommender Systems, RecSys 2010, ACM, 2010, pp. 39–46. doi:10.1145/1864708.1864721.

[38] A. Beutel, J. Chen, T. Doshi, H. Qian, L. Wei, Y. Wu, L. Heldt, Z. Zhao, L. Hong, E. H. Chi, C. Goodrow, Fairness in recommendation ranking through pairwise comparisons, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019., ACM, 2019, pp. 2212–2220. doi:10.1145/3292500.3330745.

[39] R. Burke, N. Sonboli, A. Ordonez-Gauger, Balanced neighborhoods for multi-sided fairness in recommendation, in: Conference on Fairness, Accountability and Transparency, FAT 2018, Vol. 81 of Proceedings of Machine Learning Research, PMLR, 2018, pp. 202–214.

[40] J. G. Carbonell, J. Goldstein, The use of mmr, diversity-based reranking for reordering documents and producing summaries, in: SIGIR '98: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 1998, pp. 335–336. doi:10.1145/290941.291025.

[41] L. E. Celis, D. Straszak, N. K. Vishnoi, Ranking with fairness constraints, in: 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, Vol. 107 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 28:1–28:15. doi:10.4230/LIPIcs.ICALP.2018.28.

[42] D. Dessì, M. Dragoni, G. Fenu, M. Marras, D. Reforgiato Recupero, Evaluating neural word embeddings created from online course reviews

for sentiment analysis, in: Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC 2019, ACM, 2019, pp. 2124–2127. doi:10.1145/3297280.3297620.

[43] D. Dessì, G. Fenu, M. Marras, D. Reforgiato Recupero, COCO: semantic-enriched collection of online courses at scale with experimental use cases, in: Trends and Advances in Information Systems and Technologies - Volume 2 [WorldCIST'18], Vol. 746 of Advances in Intelligent Systems and Computing, Springer, 2018, pp. 1386–1396. doi:10.1007/978-3-319-77712-2_133.

[44] J. Saúde, G. Ramos, C. Caleiro, S. Kar, Reputation-based ranking systems and their resistance to bribery, in: V. Raghavan, S. Aluru, G. Karypis, L. Miele, X. Wu (Eds.), 2017 IEEE International Conference on Data Mining, ICDM 2017, New Orleans, LA, USA, November 18-21, 2017, IEEE Computer Society, 2017, pp. 1063–1068. doi:10.1109/ICDM.2017.139.

[45] G. Ramos, L. Boratto, C. Caleiro, On the negative impact of social influence in recommender systems: A study of bribery in collaborative hybrid algorithms, Inf. Process. Manag. 57 (2) (2020) 102058. doi:10.1016/j.ipm.2019.102058.

Appendix

This appendix contains Tables that will help the reader to reproduce the results obtained in our experiments.

```
define mitigationContinent (list, reRankingType, targetProportions)
begin
    // initializes four empty lists to store candidate items to add, candidate
    items to remove, all possible swaps of items, and the disparities per
    continent, respectively
    itemsIn, itemsOut, possibleSwaps, continentList ← list(), list(), list(), list() ;
    proportions ← initialProportions(list, reRankingType); // compute continents'
      proportions in the ranked list
    continentList ← proportions − targetProportions ; // updates disparity of each
      continent
    foreach user ∈ list do // for each user
        foreach list.item ∈ top-n do // we loop over all items that belong to this
          user
            if checkPosition(list.item, itemsOut, reRankingType)==True and
              checkDisadvantagedGroup(list.continent,continentList)==False then
                itemsOut.add(list.item) ; // adds the item as possible candidate to
                  move out if it belongs to an advantaged group and belongs to the
                  top-k
            else if checkPosition(list.item, itemsOut, reRankingType)==False and
              checkDisadvantagedGroup(list.continent,continentList)==True then
                itemsIn.add(list.item) ; // adds the item as possible candidate to
                  move in if it belongs to a disadvantaged group and it is not in
                  the top-k
            end
        end
        while !itemsIn.empty() and !itemsOut.empty() do
            itemIn ← itemsIn.pop(first); // item ranked higher in the top-n, outside
              the top-k
            itemOut ← itemsOut.pop(last); // item ranked lower in the top-k
            loss ← itemOut.prediction − itemIn.prediction ; // computes the loss
            possibleSwaps.add(id, user, itemOut, itemIn, loss); // adds the possible swap
        end
    end
    sortByLoss(possibleSwaps); // sort candidate swaps by loss, from minor to major
    i ← 0;
    // do swaps until the target proportions are reached or no more swaps
    while proportions < targetProportions and i < len(possibleSwaps) do
        elem ← possibleSwaps.get(i) ; // gets candidate swap with the minor loss
        if checkPosition(elem.id, elem.itemOut, reRankingType)==True and
          checkDisadvantagedGroup(elem.itemIn.continent,continentList)==False then
            list ← swap(list, elem.itemOut, elem.itemIn); // makes the swap of items
            // computes exposure difference
            exp ← itemOut.exposure − itemIn.exposure ;
            // reduces continents' proportions for the itemOut
            proportions ← updateProportions(elem.itemOut, reRankingType, exp, −1);
            // adds continents' proportions for the itemIn
            proportions ← updateProportions(elem.itemIn, reRankingType, exp, 1);
            // updates continent's disparities
            continentList ← proportions − targetProportions ;
        i ← i + 1 ; // advances to the next possible swap with minor loss
    end
    return list ; // re-ranked list
end
```

Algorithm 2: Support method to the multiclass mitigation algorithm

```
define checkPosition(item, itemsOut, reRankingType) // check the position of an item
begin
    if reRankingType == "visibility" then return item.position < top-k ;
    else if reRankingType == "exposure" then return
      item.position < itemsOut.last.position ;
end
define checkDisadvantagedGroup (continent, continentList) // check disadvantaged
  continent
begin
    for cont ∈ continent do sumDeltas += continentList.get(cont) ; // adds the disparity
      of the continent
    return (sumDeltas > 0);
end
define initialProportions(list, reRankingType) // check initial continents'
  proportions
begin
    proportions ← 0; // set up each continent' proportion to 0
    foreach user ∈ list do // for each user
        foreach list.item ∈ top-k do // we loop over the top-k items that belong to
          this user
            if reRankingType == "visibility" then
                for cont ∈ list.continent do proportions[cont] += 1 ;
            else if reRankingType == "exposure" then
                for cont ∈ list.continent do proportions[cont] += list.exposure ;
            end
        end
    end
    return proportions
end
define updateProportions(item, reRankingType, exp, value) // update proportions after
  a swap
begin
    if reRankingType == "visibility" then
        for cont ∈ item.continent do proportions[cont] += (1 × value) ;
    else if reRankingType == "exposure" then
        for cont ∈ item.continent do proportions[cont] += ( exp × value) ;
    end
    return proportions
end
```

Algorithm 3: Support methods for the mitigationContinent method

Table 3: **Disparate impact with different mitigation strategies in the Movies dataset.** Disparate impact metrics returned by the different models for each continent (AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America) considering the Movies data. For each algorithm, we report the results obtained by the binary and by our multi-group mitigation, in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines). Under each metric, we report the gain or loss we obtained when moving from the binary to our multi-group mitigation.

		Movies											
		AF		**AS**		**EU**		**NA**		**OC**		**SA**	
		binary	**multi**	**binary**	**multi**	**binary**	**multi**	**binary**	**multi**	**binary**	**multi**	**binary**	**multi**
MostPop	$\Delta\mathcal{V}_R$	0.0060	0.0007	-0.0228	-0.0230	-0.0062	0.0001	0.0237	0.0226	-0.0003	0.0000	-0.0003	-0.0003
	(gain/loss)	0.0028	-0.0025	0.0005	0.0003	0.0831	0.0894	-0.0914	-0.0925	0.0050	0.0053	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0051	-0.0005	-0.0229	-0.0231	-0.0042	-0.0392	0.0233	0.0655	-0.0009	-0.0024	-0.0003	-0.0003
	(gain/loss)	0.0033	-0.0023	0.0005	0.0002	0.1026	0.0676	-0.1124	-0.0702	0.0060	0.0046	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0182	-0.0005	-0.0378	-0.0389	-0.0047	-0.0053	0.0133	0.0468	0.0136	0.0003	-0.0025	-0.0025
	(gain/loss)	0.0161	-0.0026	0.0012	0.0002	0.1369	0.1363	-0.1741	-0.1406	0.0200	0.0067	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0188	-0.0015	-0.0380	-0.0389	-0.0007	-0.0543	0.0100	0.0995	0.0124	-0.0022	-0.0025	-0.0025
	(gain/loss)	0.0180	-0.0023	0.0011	0.0002	0.1584	0.1047	-0.1979	-0.1085	0.0205	0.0058	0.0000	0.0000
RandomG	$\Delta\mathcal{V}_R$	-0.0002	0.0000	0.0036	0.0000	-0.0209	0.0000	0.0168	-0.0011	-0.0012	0.0000	0.0020	0.0011
	(gain/loss)	-0.0009	-0.0007	-0.0100	-0.0136	-0.0569	-0.0360	0.0734	0.0555	-0.0045	-0.0033	-0.0011	-0.0019
	$\Delta\mathcal{E}_R$	-0.0003	0.0000	0.0036	0.0000	-0.0211	-0.0001	0.0170	-0.0009	-0.0011	0.0000	0.0020	0.0010
	(gain/loss)	-0.0009	-0.0006	-0.0099	-0.0135	-0.0577	-0.0367	0.0740	0.0561	-0.0044	-0.0033	-0.0011	-0.0020
	$\Delta\mathcal{V}_I$	-0.0003	0.0000	-0.0027	0.0000	-0.0191	0.0000	0.0193	0.0000	0.0020	0.0000	0.0009	0.0000
	(gain/loss)	0.0000	0.0003	-0.0005	0.0022	-0.0029	0.0163	0.0036	-0.0157	-0.0002	-0.0022	0.0000	-0.0009
	$\Delta\mathcal{E}_I$	-0.0004	0.0000	-0.0029	0.0000	-0.0191	0.0000	0.0196	0.0000	0.0020	0.0000	0.0008	0.0000
	(gain/loss)	0.0000	0.0004	-0.0006	0.0023	-0.0034	0.0157	0.0042	-0.0153	-0.0002	-0.0022	-0.0001	-0.0008
UserKNN	$\Delta\mathcal{V}_R$	0.0056	0.0023	-0.0220	-0.0199	-0.0022	0.0000	0.0231	0.0179	-0.0042	0.0000	-0.0003	-0.0003
	(gain/loss)	0.0025	-0.0007	0.0008	0.0029	0.0697	0.0719	-0.0771	-0.0824	0.0041	0.0083	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0052	0.0019	-0.0222	-0.0208	-0.0009	-0.0277	0.0235	0.0499	-0.0053	-0.0029	-0.0003	-0.0003
	(gain/loss)	0.0028	-0.0006	0.0008	0.0022	0.0802	0.0534	-0.0878	-0.0614	0.0040	0.0064	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0069	0.0009	-0.0365	-0.0359	0.0049	-0.0002	0.0278	0.0377	-0.0006	0.0000	-0.0025	-0.0025
	(gain/loss)	0.0048	-0.0012	0.0021	0.0027	0.1290	0.1239	-0.1447	-0.1349	0.0087	0.0094	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0065	0.0004	-0.0368	-0.0364	0.0064	-0.0360	0.0286	0.0768	-0.0021	-0.0023	-0.0025	-0.0025
	(gain/loss)	0.0050	-0.0011	0.0020	0.0024	0.1397	0.0974	-0.1550	-0.1068	0.0083	0.0081	0.0000	0.0000
ItemKNN	$\Delta\mathcal{V}_R$	0.0059	0.0000	-0.0230	-0.0184	0.0083	0.0002	0.0180	0.0185	-0.0088	0.0000	-0.0003	-0.0003
	(gain/loss)	0.0053	-0.0006	0.0004	0.0050	0.0847	0.0767	-0.0937	-0.0932	0.0033	0.0121	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0055	-0.0008	-0.0231	-0.0192	0.0097	-0.0294	0.0173	0.0520	-0.0091	-0.0023	-0.0003	-0.0003
	(gain/loss)	0.0058	-0.0005	0.0003	0.0042	0.1021	0.0630	-0.1115	-0.0768	0.0032	0.0101	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0085	-0.0015	-0.0382	-0.0341	0.0090	0.0000	0.0281	0.0381	-0.0049	0.0000	-0.0025	-0.0025
	(gain/loss)	0.0089	-0.0011	0.0010	0.0050	0.1378	0.1288	-0.1559	-0.1459	0.0084	0.0132	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0078	-0.0022	-0.0383	-0.0345	0.0123	-0.0354	0.0264	0.0764	-0.0057	-0.0018	-0.0025	-0.0025
	(gain/loss)	0.0091	-0.0009	0.0009	0.0047	0.1570	0.1092	-0.1747	-0.1247	0.0077	0.0117	0.0000	0.0000
BPR	$\Delta\mathcal{V}_R$	0.0028	0.0000	-0.0117	-0.0001	-0.0047	0.0000	0.0106	0.0000	0.0029	0.0000	0.0000	0.0001
	(gain/loss)	0.0006	-0.0022	0.0024	0.0140	0.0280	0.0326	-0.0330	-0.0436	0.0020	-0.0008	0.0000	0.0001
	$\Delta\mathcal{E}_R$	0.0033	0.0009	-0.0129	-0.0041	-0.0050	-0.0086	0.0109	0.0113	0.0035	0.0003	0.0001	0.0002
	(gain/loss)	0.0007	-0.0017	0.0023	0.0110	0.0308	0.0271	-0.0363	-0.0359	0.0024	-0.0007	0.0000	0.0001
	$\Delta\mathcal{V}_I$	0.0033	0.0001	-0.0229	0.0000	0.0016	0.0000	0.0137	0.0000	0.0063	0.0002	-0.0021	-0.0003
	(gain/loss)	0.0021	-0.0011	0.0069	0.0298	0.0865	0.0849	-0.1022	-0.1159	0.0066	0.0005	0.0001	0.0018
	$\Delta\mathcal{E}_I$	0.0038	0.0005	-0.0243	-0.0065	0.0015	-0.0117	0.0142	0.0183	0.0068	0.0001	-0.0021	-0.0007
	(gain/loss)	0.0022	-0.0012	0.0067	0.0244	0.0895	0.0763	-0.1053	-0.1012	0.0069	0.0001	0.0001	0.0015
BiasedMF	$\Delta\mathcal{V}_R$	0.0026	0.0000	0.0468	0.0000	-0.0322	0.0000	-0.0060	0.0000	-0.0113	0.0000	0.0000	0.0000
	(gain/loss)	-0.0051	-0.0077	-0.0608	-0.1076	-0.0996	-0.0674	0.1669	0.1729	-0.0005	0.0108	-0.0009	-0.0010
	$\Delta\mathcal{E}_R$	0.0026	0.0001	0.0501	0.0009	-0.0358	0.0000	-0.0056	0.0000	-0.0112	-0.0009	0.0000	-0.0001
	(gain/loss)	-0.0048	-0.0073	-0.0744	-0.1236	-0.0907	-0.0549	0.1711	0.1767	-0.0005	0.0098	-0.0007	-0.0007
	$\Delta\mathcal{V}_I$	0.0040	0.0003	0.0594	0.0005	-0.0432	0.0000	-0.0060	0.0000	-0.0122	0.0000	-0.0020	-0.0007
	(gain/loss)	-0.0027	-0.0065	-0.0324	-0.0913	-0.0583	-0.0151	0.0945	0.1005	-0.0003	0.0119	-0.0007	0.0005
	$\Delta\mathcal{E}_I$	0.0038	0.0003	0.0685	0.0015	-0.0532	0.0000	-0.0049	0.0000	-0.0122	-0.0011	-0.0021	-0.0007
	(gain/loss)	-0.0026	-0.0061	-0.0402	-0.1072	-0.0558	-0.0026	0.0994	0.1044	-0.0003	0.0107	-0.0005	0.0009
SVD++	$\Delta\mathcal{V}_R$	0.0012	0.0004	0.0381	0.0000	-0.0343	0.0000	0.0070	0.0000	-0.0116	-0.0004	-0.0003	0.0000
	(gain/loss)	-0.0027	-0.0034	-0.0419	-0.0800	-0.1029	-0.0686	0.1503	0.1433	-0.0004	0.0109	-0.0024	-0.0021
	$\Delta\mathcal{E}_R$	0.0009	0.0005	0.0427	0.0000	-0.0374	0.0000	0.0058	0.0000	-0.0117	-0.0004	-0.0003	-0.0001
	(gain/loss)	-0.0020	-0.0024	-0.0526	-0.0954	-0.0943	-0.0569	0.1510	0.1452	-0.0003	0.0111	-0.0017	-0.0015
	$\Delta\mathcal{V}_I$	0.0011	0.0000	0.0440	0.0000	-0.0384	0.0000	0.0076	0.0000	-0.0125	0.0000	-0.0018	0.0000
	(gain/loss)	-0.0018	-0.0029	-0.0202	-0.0642	-0.0547	-0.0164	0.0786	0.0710	-0.0001	0.0124	-0.0017	0.0001
	$\Delta\mathcal{E}_I$	0.0005	0.0000	0.0547	0.0019	-0.0474	0.0001	0.0069	0.0000	-0.0127	-0.0019	-0.0021	-0.0001
	(gain/loss)	-0.0014	-0.0019	-0.0249	-0.0777	-0.0521	-0.0046	0.0798	0.0729	-0.0001	0.0106	-0.0013	0.0007

Table 4: **Disparate impact with different mitigation strategies in the Books dataset.** Disparate impact metrics returned by the different models for each continent (EU: Europe, NA: North America, OC: Oceania, SA: South America) considering the Books data. For each algorithm, we report the results obtained by the binary and by our multi-group mitigation, in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines). Under each metric, we report the gain or loss we obtained when moving from the binary to our multi-group mitigation.

		Books							
		EU		**NA**		**OC**		**SA**	
		binary	**multi**	**binary**	**multi**	**binary**	**multi**	**binary**	**multi**
MostPop	$\Delta\mathcal{V}_R$	0.0102	0.0000	-0.0099	0.0003	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0800	0.0697	-0.0800	-0.0697	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0102	-0.0227	-0.0099	0.0230	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0800	0.0471	-0.0800	-0.0471	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0157	0.0000	-0.0151	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.1199	0.1042	-0.1199	-0.1042	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0157	-0.0322	-0.0151	0.0328	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.1200	0.0720	-0.1200	-0.0720	0.0000	0.0000	0.0000	0.0000
RandomG	$\Delta\mathcal{V}_R$	-0.0026	0.0000	0.0025	0.0000	0.0001	0.0000	0.0000	0.0000
	(gain/loss)	-0.0383	-0.0357	0.0385	0.0360	-0.0002	-0.0003	0.0000	0.0000
	$\Delta\mathcal{E}_R$	-0.0028	0.0000	0.0027	0.0000	0.0001	0.0000	0.0000	0.0000
	(gain/loss)	-0.0384	-0.0356	0.0386	0.0359	-0.0002	-0.0003	0.0000	-0.0001
	$\Delta\mathcal{V}_I$	-0.0033	0.0000	0.0033	0.0000	0.0000	0.0000	0.0000	0.0000
	(gain/loss)	-0.0045	-0.0012	0.0045	0.0012	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_I$	-0.0035	0.0000	0.0035	0.0000	-0.0001	0.0000	0.0000	0.0000
	(gain/loss)	-0.0046	-0.0011	0.0046	0.0011	0.0000	0.0000	0.0000	0.0000
UserKNN	$\Delta\mathcal{V}_R$	0.0057	0.0001	-0.0055	0.0000	-0.0001	-0.0001	-0.0001	-0.0001
	(gain/loss)	-0.0002	-0.0058	0.0002	0.0057	0.0000	0.0001	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0062	0.0001	-0.0060	0.0000	-0.0001	-0.0001	-0.0001	-0.0001
	(gain/loss)	-0.0064	-0.0125	0.0064	0.0124	0.0000	0.0001	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0088	0.0000	-0.0082	0.0004	-0.0004	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0373	0.0286	-0.0374	-0.0288	0.0000	0.0002	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0095	0.0000	-0.0089	0.0004	-0.0004	-0.0003	-0.0001	-0.0001
	(gain/loss)	0.0313	0.0219	-0.0314	-0.0220	0.0000	0.0002	0.0000	0.0000
ItemKNN	$\Delta\mathcal{V}_R$	0.0062	0.0000	-0.0060	0.0002	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0335	0.0273	-0.0336	-0.0273	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0065	-0.0032	-0.0063	0.0034	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0328	0.0231	-0.0328	-0.0231	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0081	0.0000	-0.0075	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.0699	0.0618	-0.0699	-0.0618	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0085	-0.0121	-0.0079	0.0127	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.0692	0.0486	-0.0693	-0.0486	0.0000	0.0000	0.0000	0.0000
BPR	$\Delta\mathcal{V}_R$	0.0080	0.0000	-0.0079	0.0000	-0.0001	0.0000	0.0000	0.0000
	(gain/loss)	-0.0172	-0.0252	0.0172	0.0251	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0081	0.0000	-0.0080	0.0000	-0.0001	0.0000	0.0000	0.0000
	(gain/loss)	-0.0171	-0.0252	0.0171	0.0251	0.0000	0.0001	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0114	0.0000	-0.0110	0.0000	-0.0003	0.0000	-0.0001	0.0000
	(gain/loss)	0.0207	0.0093	-0.0208	-0.0098	0.0000	0.0003	0.0000	0.0001
	$\Delta\mathcal{E}_I$	0.0116	0.0000	-0.0112	0.0001	-0.0003	-0.0001	-0.0001	0.0000
	(gain/loss)	0.0209	0.0093	-0.0209	-0.0096	0.0000	0.0003	0.0000	0.0001
BiasedMF	$\Delta\mathcal{V}_R$	0.0102	0.0000	-0.0099	0.0003	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0800	0.0698	-0.0800	-0.0698	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0102	-0.0215	-0.0099	0.0218	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0800	0.0483	-0.0800	-0.0483	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0157	0.0000	-0.0151	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.1200	0.1043	-0.1200	-0.1043	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0157	-0.0296	-0.0151	0.0302	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.1200	0.0747	-0.1200	-0.0747	0.0000	0.0000	0.0000	0.0000
SVD++	$\Delta\mathcal{V}_R$	0.0094	0.0000	-0.0091	0.0003	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0792	0.0698	-0.0792	-0.0698	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0094	-0.0213	-0.0091	0.0216	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)	0.0792	0.0485	-0.0792	-0.0485	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0141	0.0000	-0.0134	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.1183	0.1043	-0.1183	-0.1043	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_I$	0.0140	-0.0296	-0.0134	0.0302	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)	0.1183	0.0747	-0.1183	-0.0747	0.0000	0.0000	0.0000	0.0000

Table 5: **Disparate impact with different mitigation strategies.** Disparate impact metrics returned by the different models for each continent (AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). For each algorithm, we report the results obtained by the baseline and by our multi-group mitigation, in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines).

		Movies											
		AF		**AS**		**EU**		**NA**		**OC**		**SA**	
		multi	**baseline**	**multi**	**baseline**	**multi**	**baseline**	**multi**	**baseline**	**multi**	**baseline**	**multi**	**baseline**
MostPop	$\Delta\mathcal{V}_R$	0.0007	0.0096	-0.0230	-0.0218	0.0001	-0.0853	0.0226	0.1017	0.0000	-0.0039	-0.0003	-0.0003
	$\Delta\mathcal{E}_R$	-0.0005	0.0058	-0.0231	-0.0225	-0.0392	-0.1043	0.0655	0.1274	-0.0024	-0.0060	-0.0003	-0.0003
	$\Delta\mathcal{V}_I$	-0.0005	0.0086	-0.0389	-0.0375	-0.0053	-0.1376	0.0468	0.1740	0.0003	-0.0050	-0.0025	-0.0025
	$\Delta\mathcal{E}_I$	-0.0015	0.0048	-0.0389	-0.0383	-0.0543	-0.1566	0.0995	0.1997	-0.0022	-0.0071	-0.0025	-0.0025
RandomG	$\Delta\mathcal{V}_R$	0.0000	0.0136	0.0000	0.0359	0.0000	0.0402	-0.0011	-0.1336	0.0000	0.0280	0.0011	0.0159
	$\Delta\mathcal{E}_R$	0.0000	0.0072	0.0000	0.0255	-0.0001	0.0387	-0.0009	-0.0971	0.0000	0.0163	0.0010	0.0095
	$\Delta\mathcal{V}_I$	0.0000	0.0127	0.0000	0.0202	0.0000	-0.0121	0.0000	-0.0613	0.0000	0.0269	0.0000	0.0137
	$\Delta\mathcal{E}_I$	0.0000	0.0062	0.0000	0.0097	0.0000	-0.0136	0.0000	-0.0248	0.0000	0.0152	0.0000	0.0073
UserKNN	$\Delta\mathcal{V}_R$	0.0023	0.0075	-0.0199	-0.0211	0.0000	-0.0681	0.0179	0.0866	0.0000	-0.0045	-0.0003	-0.0003
	$\Delta\mathcal{E}_R$	0.0019	0.0050	-0.0208	-0.0220	-0.0277	-0.0788	0.0499	0.1031	-0.0029	-0.0070	-0.0003	-0.0003
	$\Delta\mathcal{V}_I$	0.0009	0.0065	-0.0359	-0.0369	-0.0002	-0.1204	0.0377	0.1589	0.0000	-0.0056	-0.0025	-0.0025
	$\Delta\mathcal{E}_I$	0.0004	0.0041	-0.0364	-0.0378	-0.0360	-0.1310	0.0768	0.1754	-0.0023	-0.0081	-0.0025	-0.0025
ItemKNN	$\Delta\mathcal{V}_R$	0.0000	0.0060	-0.0184	-0.0226	0.0002	-0.0722	0.0185	0.0996	0.0000	-0.0105	-0.0003	-0.0003
	$\Delta\mathcal{E}_R$	-0.0008	0.0028	-0.0192	-0.0230	-0.0294	-0.0898	0.0520	0.1217	-0.0023	-0.0114	-0.0003	-0.0003
	$\Delta\mathcal{V}_I$	-0.0015	0.0050	-0.0341	-0.0384	0.0000	-0.1245	0.0381	0.1719	0.0000	-0.0115	-0.0025	-0.0025
	$\Delta\mathcal{E}_I$	-0.0022	0.0018	-0.0345	-0.0388	-0.0354	-0.1421	0.0764	0.1940	-0.0018	-0.0124	-0.0025	-0.0025
BPR	$\Delta\mathcal{V}_R$	0.0000	0.0046	-0.0001	-0.0091	0.0000	-0.0293	0.0000	0.0292	0.0000	0.0043	0.0001	0.0002
	$\Delta\mathcal{E}_R$	0.0009	0.0041	-0.0041	-0.0122	-0.0086	-0.0337	0.0113	0.0386	0.0003	0.0031	0.0002	0.0002
	$\Delta\mathcal{V}_I$	0.0001	0.0037	0.0000	-0.0249	0.0000	-0.0815	0.0000	0.1015	0.0002	0.0032	-0.0003	-0.0019
	$\Delta\mathcal{E}_I$	0.0005	0.0031	-0.0065	-0.0280	-0.0117	-0.0860	0.0183	0.1109	0.0001	0.0020	-0.0007	-0.0020
BiasedMF	$\Delta\mathcal{V}_R$	0.0000	0.0177	0.0000	0.1061	0.0000	0.0707	0.0000	-0.1994	0.0000	-0.0065	0.0000	0.0114
	$\Delta\mathcal{E}_R$	0.0001	0.0130	0.0009	0.1238	0.0000	0.0566	0.0000	-0.1909	-0.0009	-0.0085	-0.0001	0.0060
	$\Delta\mathcal{V}_I$	0.0003	0.0168	0.0005	0.0903	0.0000	0.0185	0.0000	-0.1271	0.0000	-0.0076	-0.0007	0.0092
	$\Delta\mathcal{E}_I$	0.0003	0.0120	0.0015	0.1080	0.0000	0.0043	0.0000	-0.1186	-0.0011	-0.0096	-0.0007	0.0038
SVD++	$\Delta\mathcal{V}_R$	0.0004	0.0163	0.0000	0.0769	0.0000	0.0765	0.0000	-0.1838	-0.0004	-0.0066	0.0000	0.0207
	$\Delta\mathcal{E}_R$	0.0005	0.0098	0.0000	0.0938	0.0000	0.0610	0.0000	-0.1669	-0.0004	-0.0090	-0.0001	0.0112
	$\Delta\mathcal{V}_I$	0.0000	0.0153	0.0000	0.0611	0.0000	0.0243	0.0000	-0.1116	0.0000	-0.0076	0.0000	0.0185
	$\Delta\mathcal{E}_I$	0.0000	0.0089	0.0019	0.0780	0.0001	0.0088	0.0000	-0.0946	-0.0019	-0.0101	-0.0001	0.0091

Table 6: **Disparate impact with different mitigation strategies.** Disparate impact metrics returned by the different models for each continent (AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). For each algorithm, we report the results obtained by the baseline and by our multi-group mitigation, in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines).

		Books							
		EU		**NA**		**OC**		**SA**	
		multi	**baseline**	**multi**	**baseline**	**multi**	**baseline**	**multi**	**baseline**
MostPop	$\Delta\mathcal{V}_R$	0.0000	-0.0592	0.0003	0.0595	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{E}_R$	-0.0227	-0.0643	0.0230	0.0646	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{V}_I$	0.0000	-0.0937	0.0006	0.0943	-0.0005	-0.0005	-0.0001	-0.0001
	$\Delta\mathcal{E}_I$	-0.0322	-0.0988	0.0328	0.0995	-0.0005	-0.0005	-0.0001	-0.0001
RandomG	$\Delta\mathcal{V}_R$	0.0000	0.0378	0.0000	-0.0395	0.0000	0.0013	0.0000	0.0004
	$\Delta\mathcal{E}_R$	0.0000	0.0369	0.0000	-0.0379	0.0000	0.0007	0.0000	0.0002
	$\Delta\mathcal{V}_I$	0.0000	0.0034	0.0000	-0.0047	0.0000	0.0010	0.0000	0.0003
	$\Delta\mathcal{E}_I$	0.0000	0.0024	0.0000	-0.0030	0.0000	0.0004	0.0000	0.0002
UserKNN	$\Delta\mathcal{V}_R$	0.0001	0.0094	0.0000	-0.0093	-0.0001	-0.0001	-0.0001	-0.0001
	$\Delta\mathcal{E}_R$	0.0001	0.0157	0.0000	-0.0156	-0.0001	-0.0001	-0.0001	-0.0001
	$\Delta\mathcal{V}_I$	0.0000	-0.0251	0.0004	0.0256	-0.0002	-0.0004	-0.0001	-0.0001
	$\Delta\mathcal{E}_I$	0.0000	-0.0187	0.0004	0.0193	-0.0003	-0.0004	-0.0001	-0.0001
ItemKNN	$\Delta\mathcal{V}_R$	0.0000	-0.0248	0.0002	0.0251	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{E}_R$	-0.0032	-0.0248	0.0034	0.0251	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{V}_I$	0.0000	-0.0593	0.0006	0.0599	-0.0005	-0.0005	-0.0001	-0.0001
	$\Delta\mathcal{E}_I$	-0.0121	-0.0593	0.0127	0.0599	-0.0005	-0.0005	-0.0001	-0.0001
BPR	$\Delta\mathcal{V}_R$	0.0000	0.0262	0.0000	-0.0262	0.0000	0.0000	0.0000	0.0000
	$\Delta\mathcal{E}_R$	0.0000	0.0259	0.0000	-0.0258	0.0000	-0.0001	0.0000	0.0000
	$\Delta\mathcal{V}_I$	0.0000	-0.0083	0.0000	0.0087	0.0000	-0.0003	0.0000	-0.0001
	$\Delta\mathcal{E}_I$	0.0000	-0.0086	0.0001	0.0091	-0.0001	-0.0004	0.0000	-0.0001
BiasedMF	$\Delta\mathcal{V}_R$	0.0000	-0.0698	0.0003	0.0701	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{E}_R$	-0.0215	-0.0698	0.0218	0.0701	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{V}_I$	0.0000	-0.1043	0.0006	0.1049	-0.0005	-0.0005	-0.0001	-0.0001
	$\Delta\mathcal{E}_I$	-0.0296	-0.1043	0.0302	0.1049	-0.0005	-0.0005	-0.0001	-0.0001
SVD++	$\Delta\mathcal{V}_R$	0.0000	-0.0698	0.0003	0.0700	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{E}_R$	-0.0213	-0.0698	0.0216	0.0700	-0.0002	-0.0002	-0.0001	-0.0001
	$\Delta\mathcal{V}_I$	0.0000	-0.1042	0.0006	0.1049	-0.0005	-0.0005	-0.0001	-0.0001
	$\Delta\mathcal{E}_I$	-0.0296	-0.1042	0.0302	0.1049	-0.0005	-0.0005	-0.0001	-0.0001

Table 7: **Disparate impact in the Movies dataset.** Disparate impact metrics returned by the different models for each continent (AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America) considering the Movies data. For each algorithm, we report the results obtained by the original state-of-the-art algorithm and the binary mitigation proposed in [7], in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines). Under each metric, we report the gain or loss we obtained when moving from the original model to the binary mitigation.

		Movies											
		AF		**AS**		**EU**		**NA**		**OC**		**SA**	
		original	**binary**	**original**	**binary**	**original**	**binary**	**original**	**binary**	**original**	**binary**	**original**	**binary**
MostPop	$\Delta\mathcal{V}_R$	0.0031	0.0060	-0.0233	-0.0228	-0.0893	-0.0062	0.1151	0.0237	-0.0053	-0.0003	-0.0003	-0.0003
	(gain/loss)		0.0028		0.0005		0.0831		-0.0914		0.0050		0.0000
	$\Delta\mathcal{E}_R$	0.0018	0.0051	-0.0233	-0.0229	-0.1068	-0.0042	0.1357	0.0233	-0.0070	-0.0009	-0.0003	-0.0003
	(gain/loss)		0.0033		0.0005		0.1026		-0.1124		0.0060		0.0000
	$\Delta\mathcal{V}_I$	0.0022	0.0182	-0.0391	-0.0378	-0.1416	-0.0047	0.1874	0.0133	-0.0064	0.0136	-0.0025	-0.0025
	(gain/loss)		0.0161		0.0012		0.1369		-0.1741		0.0200		0.0000
	$\Delta\mathcal{E}_I$	0.0008	0.0188	-0.0391	-0.0380	-0.1591	-0.0007	0.2080	0.0100	-0.0080	0.0124	-0.0025	-0.0025
	(gain/loss)		0.0180		0.0011		0.1584		-0.1979		0.0205		0.0000
RandomG	$\Delta\mathcal{V}_R$	0.0007	-0.0002	0.0135	0.0036	0.0360	-0.0209	-0.0566	0.0168	0.0033	-0.0012	0.0031	0.0020
	(gain/loss)		-0.0009		-0.0100		-0.0569		0.0734		-0.0045		-0.0011
	$\Delta\mathcal{E}_R$	0.0006	-0.0003	0.0135	0.0036	0.0366	-0.0211	-0.0570	0.0170	0.0033	-0.0011	0.0030	0.0020
	(gain/loss)		-0.0009		-0.0099		-0.0577		0.0740		-0.0044		-0.0011
	$\Delta\mathcal{V}_I$	-0.0003	-0.0003	-0.0022	-0.0027	-0.0163	-0.0191	0.0157	0.0193	0.0022	0.0020	0.0009	0.0009
	(gain/loss)		0.0000		-0.0005		-0.0029		0.0036		-0.0002		0.0000
	$\Delta\mathcal{E}_I$	-0.0004	-0.0004	-0.0023	-0.0029	-0.0157	-0.0191	0.0153	0.0196	0.0022	0.0020	0.0008	0.0008
	(gain/loss)		0.0000		-0.0006		-0.0034		0.0042		-0.0002		-0.0001
UserKNN	$\Delta\mathcal{V}_R$	0.0031	0.0056	-0.0228	-0.0220	-0.0719	-0.0022	0.1003	0.0231	-0.0083	-0.0042	-0.0003	-0.0003
	(gain/loss)		0.0025		0.0008		0.0697		-0.0771		0.0041		0.0000
	$\Delta\mathcal{E}_R$	0.0024	0.0052	-0.0230	-0.0222	-0.0811	-0.0009	0.1113	0.0235	-0.0093	-0.0053	-0.0003	-0.0003
	(gain/loss)		0.0028		0.0008		0.0802		-0.0878		0.0040		0.0000
	$\Delta\mathcal{V}_I$	0.0021	0.0069	-0.0386	-0.0365	-0.1241	0.0049	0.1726	0.0278	-0.0094	-0.0006	-0.0025	-0.0025
	(gain/loss)		0.0048		0.0021		0.1290		-0.1447		0.0087		0.0000
	$\Delta\mathcal{E}_I$	0.0015	0.0065	-0.0388	-0.0368	-0.1333	0.0064	0.1836	0.0286	-0.0104	-0.0021	-0.0025	-0.0025
	(gain/loss)		0.0050		0.0020		0.1397		-0.1550		0.0083		0.0000
ItemKNN	$\Delta\mathcal{V}_R$	0.0006	0.0059	-0.0234	-0.0230	-0.0765	0.0083	0.1117	0.0180	-0.0121	-0.0088	-0.0003	-0.0003
	(gain/loss)		0.0053		0.0004		0.0847		-0.0937		0.0033		0.0000
	$\Delta\mathcal{E}_R$	-0.0003	0.0055	-0.0234	-0.0231	-0.0924	0.0097	0.1288	0.0173	-0.0123	-0.0091	-0.0003	-0.0003
	(gain/loss)		0.0058		0.0003		0.1021		-0.1115		0.0032		0.0000
	$\Delta\mathcal{V}_I$	-0.0004	0.0085	-0.0391	-0.0382	-0.1288	0.0090	0.1840	0.0281	-0.0132	-0.0049	-0.0025	-0.0025
	(gain/loss)		0.0089		0.0010		0.1378		-0.1559		0.0084		0.0000
	$\Delta\mathcal{E}_I$	-0.0013	0.0078	-0.0392	-0.0383	-0.1447	0.0123	0.2011	0.0264	-0.0134	-0.0057	-0.0025	-0.0025
	(gain/loss)		0.0091		0.0009		0.1570		-0.1747		0.0077		0.0000
BPR	$\Delta\mathcal{V}_R$	0.0022	0.0028	-0.0140	-0.0117	-0.0326	-0.0047	0.0436	0.0106	0.0008	0.0029	0.0000	0.0000
	(gain/loss)		0.0006		0.0024		0.0280		-0.0330		0.0020		0.0000
	$\Delta\mathcal{E}_R$	0.0026	0.0033	-0.0152	-0.0129	-0.0357	-0.0050	0.0472	0.0109	0.0011	0.0035	0.0000	0.0001
	(gain/loss)		0.0007		0.0023		0.0308		-0.0363		0.0024		0.0000
	$\Delta\mathcal{V}_I$	0.0012	0.0033	-0.0298	-0.0229	-0.0849	0.0016	0.1159	0.0137	-0.0003	0.0063	-0.0022	-0.0021
	(gain/loss)		0.0021		0.0069		0.0865		-0.1022		0.0066		0.0001
	$\Delta\mathcal{E}_I$	0.0016	0.0038	-0.0310	-0.0243	-0.0880	0.0015	0.1195	0.0142	0.0000	0.0068	-0.0021	-0.0021
	(gain/loss)		0.0022		0.0067		0.0895		-0.1053		0.0069		0.0001
BiasedMF	$\Delta\mathcal{V}_R$	0.0077	0.0026	0.1076	0.0468	0.0674	-0.0322	-0.1728	-0.0060	-0.0108	-0.0113	0.0009	0.0000
	(gain/loss)		-0.0051		-0.0608		-0.0996		0.1669		-0.0005		-0.0009
	$\Delta\mathcal{E}_R$	0.0074	0.0026	0.1245	0.0501	0.0549	-0.0358	-0.1767	-0.0056	-0.0108	-0.0112	0.0006	0.0000
	(gain/loss)		-0.0048		-0.0744		-0.0907		0.1711		-0.0005		-0.0007
	$\Delta\mathcal{V}_I$	0.0067	0.0040	0.0918	0.0594	0.0151	-0.0432	-0.1005	-0.0060	-0.0119	-0.0122	-0.0012	-0.0020
	(gain/loss)		-0.0027		-0.0324		-0.0583		0.0945		-0.0003		-0.0007
	$\Delta\mathcal{E}_I$	0.0064	0.0038	0.1087	0.0685	0.0026	-0.0532	-0.1044	-0.0049	-0.0118	-0.0122	-0.0016	-0.0021
	(gain/loss)		-0.0026		-0.0402		-0.0558		0.0994		-0.0003		-0.0005
SVD++	$\Delta\mathcal{V}_R$	0.0039	0.0012	0.0800	0.0381	0.0686	-0.0343	-0.1433	0.0070	-0.0113	-0.0116	0.0021	-0.0003
	(gain/loss)		-0.0027		-0.0419		-0.1029		0.1503		-0.0004		-0.0024
	$\Delta\mathcal{E}_R$	0.0029	0.0009	0.0954	0.0427	0.0569	-0.0374	-0.1452	0.0058	-0.0114	-0.0117	0.0014	-0.0003
	(gain/loss)		-0.0020		-0.0526		-0.0943		0.1510		-0.0003		-0.0017
	$\Delta\mathcal{V}_I$	0.0029	0.0011	0.0642	0.0440	0.0164	-0.0384	-0.0710	0.0076	-0.0124	-0.0125	-0.0001	-0.0018
	(gain/loss)		-0.0018		-0.0202		-0.0547		0.0786		-0.0001		-0.0017
	$\Delta\mathcal{E}_I$	0.0019	0.0005	0.0796	0.0547	0.0047	-0.0474	-0.0729	0.0069	-0.0125	-0.0127	-0.0008	-0.0021
	(gain/loss)		-0.0014		-0.0249		-0.0521		0.0798		-0.0001		-0.0013

Table 8: **Disparate impact in the Books dataset.** Disparate impact metrics returned by the different models for each continent (EU: Europe, NA: North America, OC: Oceania, SA: South America) considering the Books data. For each algorithm, we report the results obtained by the original state-of-the-art algorithm and the binary mitigation proposed in [7], in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines). Under each metric, we report the gain or loss we obtained when moving from the original model to the binary mitigation.

		Books							
		Europe		**NA**		**OCE**		**SA**	
		original	**binary**	**original**	**binary**	**original**	**binary**	**original**	**binary**
MostPop	$\Delta\mathcal{V}_R$	-0.0697	0.0102	0.0700	-0.0099	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0800		-0.0800		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0697	0.0102	0.0700	-0.0099	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0800		-0.0800		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.1042	0.0157	0.1049	-0.0151	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1199		-0.1199		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.1042	0.0157	0.1049	-0.0151	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1200		-0.1200		0.0000		0.0000
RandomG	$\Delta\mathcal{V}_R$	0.0357	-0.0026	-0.0360	0.0025	0.0003	0.0001	0.0001	0.0000
	(gain/loss)		-0.0383		0.0385		-0.0002		0.0000
	$\Delta\mathcal{E}_R$	0.0356	-0.0028	-0.0359	0.0027	0.0003	0.0001	0.0001	0.0000
	(gain/loss)		-0.0384		0.0386		-0.0002		0.0000
	$\Delta\mathcal{V}_I$	0.0012	-0.0033	-0.0012	0.0033	0.0000	0.0000	0.0000	0.0000
	(gain/loss)		-0.0045		0.0045		0.0000		0.0000
	$\Delta\mathcal{E}_I$	0.0011	-0.0035	-0.0011	0.0035	0.0000	-0.0001	0.0000	0.0000
	(gain/loss)		-0.0046		0.0046		0.0000		0.0000
UserKNN	$\Delta\mathcal{V}_R$	0.0059	0.0057	-0.0057	-0.0055	-0.0002	-0.0001	-0.0001	-0.0001
	(gain/loss)		-0.0002		0.0002		0.0000		0.0000
	$\Delta\mathcal{E}_R$	0.0126	0.0062	-0.0124	-0.0060	-0.0002	-0.0001	-0.0001	-0.0001
	(gain/loss)		-0.0064		0.0064		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.0286	0.0088	0.0292	-0.0082	-0.0005	-0.0004	-0.0001	-0.0001
	(gain/loss)		0.0373		-0.0374		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.0219	0.0095	0.0225	-0.0089	-0.0005	-0.0004	-0.0001	-0.0001
	(gain/loss)		0.0313		-0.0314		0.0000		0.0000
ItemKNN	$\Delta\mathcal{V}_R$	-0.0273	0.0062	0.0276	-0.0060	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0335		-0.0336		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0263	0.0065	0.0265	-0.0063	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0328		-0.0328		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.0618	0.0081	0.0624	-0.0075	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0699		-0.0699		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.0607	0.0085	0.0614	-0.0079	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0692		-0.0693		0.0000		0.0000
BPR	$\Delta\mathcal{V}_R$	0.0252	0.0080	-0.0251	-0.0079	0.0000	-0.0001	0.0000	0.0000
	(gain/loss)		-0.0172		0.0172		0.0000		0.0000
	$\Delta\mathcal{E}_R$	0.0252	0.0081	-0.0251	-0.0080	-0.0001	-0.0001	0.0000	0.0000
	(gain/loss)		-0.0171		0.0171		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.0093	0.0114	0.0097	-0.0110	-0.0003	-0.0003	-0.0001	-0.0001
	(gain/loss)		0.0207		-0.0208		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.0093	0.0116	0.0097	-0.0112	-0.0004	-0.0003	-0.0001	-0.0001
	(gain/loss)		0.0209		-0.0209		0.0000		0.0000
BiasedMF	$\Delta\mathcal{V}_R$	-0.0698	0.0102	0.0701	-0.0099	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0800		-0.0800		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0698	0.0102	0.0701	-0.0099	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0800		-0.0800		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.1043	0.0157	0.1049	-0.0151	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1200		-0.1200		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.1043	0.0157	0.1049	-0.0151	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1200		-0.1200		0.0000		0.0000
SVD++	$\Delta\mathcal{V}_R$	-0.0698	0.0094	0.0701	-0.0091	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0792		-0.0792		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0698	0.0094	0.0701	-0.0091	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0792		-0.0792		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.1043	0.0141	0.1049	-0.0134	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1183		-0.1183		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.1043	0.0140	0.1049	-0.0134	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1183		-0.1183		0.0000		0.0000

Table 9: **Disparate impact after mitigation in the Movies dataset.** Disparate impact metrics returned by the different models for each continent (AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America) considering the Movies data. For each algorithm, we report the results obtained by the original algorithm and by our multi-group mitigation, in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines). Under each metric, we report the gain or loss we obtained when moving from the original model to our multi-group mitigation.

		Movies											
		AF		**AS**		**EU**		**NA**		**OC**		**SA**	
		original	**multi**	**original**	**multi**	**original**	**multi**	**original**	**multi**	**original**	**multi**	**original**	**multi**
MostPop	$\Delta\mathcal{V}_R$	0.0031	0.0007	-0.0233	-0.0230	-0.0893	0.0001	0.1151	0.0226	-0.0053	0.0000	-0.0003	-0.0003
	(gain/loss)		-0.0025		0.0003		0.0894		-0.0925		0.0053		0.0000
	$\Delta\mathcal{E}_R$	0.0018	-0.0005	-0.0233	-0.0231	-0.1068	-0.0392	0.1357	0.0655	-0.0070	-0.0024	-0.0003	-0.0003
	(gain/loss)		-0.0023		0.0002		0.0676		-0.0702		0.0046		0.0000
	$\Delta\mathcal{V}_I$	0.0022	-0.0005	-0.0391	-0.0389	-0.1416	-0.0053	0.1874	0.0468	-0.0064	0.0003	-0.0025	-0.0025
	(gain/loss)		-0.0026		0.0002		0.1363		-0.1406		0.0067		0.0000
	$\Delta\mathcal{E}_I$	0.0008	-0.0015	-0.0391	-0.0389	-0.1591	-0.0543	0.2080	0.0995	-0.0080	-0.0022	-0.0025	-0.0025
	(gain/loss)		-0.0023		0.0002		0.1047		-0.1085		0.0058		0.0000
RandomG	$\Delta\mathcal{V}_R$	0.0007	0.0000	0.0135	0.0000	0.0360	0.0000	-0.0566	-0.0011	0.0033	0.0000	0.0031	0.0011
	(gain/loss)		-0.0007		-0.0136		-0.0360		0.0555		-0.0033		-0.0019
	$\Delta\mathcal{E}_R$	0.0006	0.0000	0.0135	0.0000	0.0366	-0.0001	-0.0570	-0.0009	0.0033	0.0000	0.0030	0.0010
	(gain/loss)		-0.0006		-0.0135		-0.0367		0.0561		-0.0033		-0.0020
	$\Delta\mathcal{V}_I$	-0.0003	0.0000	-0.0022	0.0000	-0.0163	0.0000	0.0157	0.0000	0.0022	0.0000	0.0009	0.0000
	(gain/loss)		0.0003		0.0022		0.0163		-0.0157		-0.0022		-0.0009
	$\Delta\mathcal{E}_I$	-0.0004	0.0000	-0.0023	0.0000	-0.0157	0.0000	0.0153	0.0000	0.0022	0.0000	0.0008	0.0000
	(gain/loss)		0.0004		0.0023		0.0157		-0.0153		-0.0022		-0.0008
UserKNN	$\Delta\mathcal{V}_R$	0.0031	0.0023	-0.0228	-0.0199	-0.0719	0.0000	0.1003	0.0179	-0.0083	0.0000	-0.0003	-0.0003
	(gain/loss)		-0.0007		0.0029		0.0719		-0.0824		0.0083		0.0000
	$\Delta\mathcal{E}_R$	0.0024	0.0019	-0.0230	-0.0208	-0.0811	-0.0277	0.1113	0.0499	-0.0093	-0.0029	-0.0003	-0.0003
	(gain/loss)		-0.0006		0.0022		0.0534		-0.0614		0.0064		0.0000
	$\Delta\mathcal{V}_I$	0.0021	0.0009	-0.0386	-0.0359	-0.1241	-0.0002	0.1726	0.0377	-0.0094	0.0000	-0.0025	-0.0025
	(gain/loss)		-0.0012		0.0027		0.1239		-0.1349		0.0094		0.0000
	$\Delta\mathcal{E}_I$	0.0015	0.0004	-0.0388	-0.0364	-0.1333	-0.0360	0.1836	0.0768	-0.0104	-0.0023	-0.0025	-0.0025
	(gain/loss)		-0.0011		0.0024		0.0974		-0.1068		0.0081		0.0000
ItemKNN	$\Delta\mathcal{V}_R$	0.0006	0.0000	-0.0234	-0.0184	-0.0765	0.0002	0.1117	0.0185	-0.0121	0.0000	-0.0003	-0.0003
	(gain/loss)		-0.0006		0.0050		0.0767		-0.0932		0.0121		0.0000
	$\Delta\mathcal{E}_R$	-0.0003	-0.0008	-0.0234	-0.0192	-0.0924	-0.0294	0.1288	0.0520	-0.0123	-0.0023	-0.0003	-0.0003
	(gain/loss)		-0.0005		0.0042		0.0630		-0.0768		0.0101		0.0000
	$\Delta\mathcal{V}_I$	-0.0004	-0.0015	-0.0391	-0.0341	-0.1288	0.0000	0.1840	0.0381	-0.0132	0.0000	-0.0025	-0.0025
	(gain/loss)		-0.0011		0.0050		0.1288		-0.1459		0.0132		0.0000
	$\Delta\mathcal{E}_I$	-0.0013	-0.0022	-0.0392	-0.0345	-0.1447	-0.0354	0.2011	0.0764	-0.0134	-0.0018	-0.0025	-0.0025
	(gain/loss)		-0.0009		0.0047		0.1092		-0.1247		0.0117		0.0000
BPR	$\Delta\mathcal{V}_R$	0.0022	0.0000	-0.0140	-0.0001	-0.0326	0.0000	0.0436	0.0000	0.0008	0.0000	0.0000	0.0001
	(gain/loss)		-0.0022		0.0140		0.0326		-0.0436		-0.0008		0.0001
	$\Delta\mathcal{E}_R$	0.0026	0.0009	-0.0152	-0.0041	-0.0357	-0.0086	0.0472	0.0113	0.0011	0.0003	0.0000	0.0002
	(gain/loss)		-0.0017		0.0110		0.0271		-0.0359		-0.0007		0.0001
	$\Delta\mathcal{V}_I$	0.0012	0.0001	-0.0298	0.0000	-0.0849	0.0000	0.1159	0.0000	-0.0003	0.0002	-0.0022	-0.0003
	(gain/loss)		-0.0011		0.0298		0.0849		-0.1159		0.0005		0.0018
	$\Delta\mathcal{E}_I$	0.0016	0.0005	-0.0310	-0.0065	-0.0880	-0.0117	0.1195	0.0183	0.0000	0.0001	-0.0021	-0.0007
	(gain/loss)		-0.0012		0.0244		0.0763		-0.1012		0.0001		0.0015
BiasedMF	$\Delta\mathcal{V}_R$	0.0077	0.0000	0.1076	0.0000	0.0674	0.0000	-0.1728	0.0000	-0.0108	0.0000	0.0009	0.0000
	(gain/loss)		-0.0077		-0.1076		-0.0674		0.1729		0.0108		-0.0010
	$\Delta\mathcal{E}_R$	0.0074	0.0001	0.1245	0.0009	0.0549	0.0000	-0.1767	0.0000	-0.0108	-0.0009	0.0006	-0.0001
	(gain/loss)		-0.0073		-0.1236		-0.0549		0.1767		0.0098		-0.0007
	$\Delta\mathcal{V}_I$	0.0067	0.0003	0.0918	0.0005	0.0151	0.0000	-0.1005	0.0000	-0.0119	0.0000	-0.0012	-0.0007
	(gain/loss)		-0.0065		-0.0913		-0.0151		0.1005		0.0119		0.0005
	$\Delta\mathcal{E}_I$	0.0064	0.0003	0.1087	0.0015	0.0026	0.0000	-0.1044	0.0000	-0.0118	-0.0011	-0.0016	-0.0007
	(gain/loss)		-0.0061		-0.1072		-0.0026		0.1044		0.0107		0.0009
SVD++	$\Delta\mathcal{V}_R$	0.0039	0.0004	0.0800	0.0000	0.0686	0.0000	-0.1433	0.0000	-0.0113	-0.0004	0.0021	0.0000
	(gain/loss)		-0.0034		-0.0800		-0.0686		0.1433		0.0109		-0.0021
	$\Delta\mathcal{E}_R$	0.0029	0.0005	0.0954	0.0000	0.0569	0.0000	-0.1452	0.0000	-0.0114	-0.0004	0.0014	-0.0001
	(gain/loss)		-0.0024		-0.0954		-0.0569		0.1452		0.0111		-0.0015
	$\Delta\mathcal{V}_I$	0.0029	0.0000	0.0642	0.0000	0.0164	0.0000	-0.0710	0.0000	-0.0124	0.0000	-0.0001	0.0000
	(gain/loss)		-0.0029		-0.0642		-0.0164		0.0710		0.0124		0.0001
	$\Delta\mathcal{E}_I$	0.0019	0.0000	0.0796	0.0019	0.0047	0.0001	-0.0729	0.0000	-0.0125	-0.0019	-0.0008	-0.0001
	(gain/loss)		-0.0019		-0.0777		-0.0046		0.0729		0.0106		0.0007

Table 10: **Disparate impact after mitigation in the Books dataset.** Disparate impact metrics returned by the different models for each continent (EU: Europe, NA: North America, OC: Oceania, SA: South America) considering the Books data. For each algorithm, we report the results obtained by the original algorithm and by our multi-group mitigation, in terms of disparate visibility and exposure when considering the rating-based representation as a reference ($\Delta\mathcal{V}_R$ and $\Delta\mathcal{E}_R$ lines) and with the item-based representation ($\Delta\mathcal{V}_I$ and $\Delta\mathcal{E}_I$ lines). Under each metric, we report the gain or loss we obtained when moving from the original model to our multi-group mitigation.

		Books							
		EU		**NA**		**OC**		**SA**	
		original	**multi**	**original**	**multi**	**original**	**multi**	**original**	**multi**
MostPop	$\Delta\mathcal{V}_R$	-0.0697	0.0000	0.0700	0.0003	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0697		-0.0697		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0697	-0.0227	0.0700	0.0230	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0471		-0.0471		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.1042	0.0000	0.1049	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1042		-0.1042		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.1042	-0.0322	0.1049	0.0328	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0720		-0.0720		0.0000		0.0000
RandomG	$\Delta\mathcal{V}_R$	0.0357	0.0000	-0.0360	0.0000	0.0003	0.0000	0.0001	0.0000
	(gain/loss)		-0.0357		0.0360		-0.0003		0.0000
	$\Delta\mathcal{E}_R$	0.0356	0.0000	-0.0359	0.0000	0.0003	0.0000	0.0001	0.0000
	(gain/loss)		-0.0356		0.0359		-0.0003		-0.0001
	$\Delta\mathcal{V}_I$	0.0012	0.0000	-0.0012	0.0000	0.0000	0.0000	0.0000	0.0000
	(gain/loss)		-0.0012		0.0012		0.0000		0.0000
	$\Delta\mathcal{E}_I$	0.0011	0.0000	-0.0011	0.0000	0.0000	0.0000	0.0000	0.0000
	(gain/loss)		-0.0011		0.0011		0.0000		0.0000
UserKNN	$\Delta\mathcal{V}_R$	0.0059	0.0001	-0.0057	0.0000	-0.0002	-0.0001	-0.0001	-0.0001
	(gain/loss)		-0.0058		0.0057		0.0001		0.0000
	$\Delta\mathcal{E}_R$	0.0126	0.0001	-0.0124	0.0000	-0.0002	-0.0001	-0.0001	-0.0001
	(gain/loss)		-0.0125		0.0124		0.0001		0.0000
	$\Delta\mathcal{V}_I$	-0.0286	0.0000	0.0292	0.0004	-0.0005	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0286		-0.0288		0.0002		0.0000
	$\Delta\mathcal{E}_I$	-0.0219	0.0000	0.0225	0.0004	-0.0005	-0.0003	-0.0001	-0.0001
	(gain/loss)		0.0219		-0.0220		0.0002		0.0000
ItemKNN	$\Delta\mathcal{V}_R$	-0.0273	0.0000	0.0276	0.0002	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0273		-0.0273		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0263	-0.0032	0.0265	0.0034	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0231		-0.0231		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.0618	0.0000	0.0624	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0618		-0.0618		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.0607	-0.0121	0.0614	0.0127	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0486		-0.0486		0.0000		0.0000
BPR	$\Delta\mathcal{V}_R$	0.0252	0.0000	-0.0251	0.0000	0.0000	0.0000	0.0000	0.0000
	(gain/loss)		-0.0252		0.0251		0.0000		0.0000
	$\Delta\mathcal{E}_R$	0.0252	0.0000	-0.0251	0.0000	-0.0001	0.0000	0.0000	0.0000
	(gain/loss)		-0.0252		0.0251		0.0001		0.0000
	$\Delta\mathcal{V}_I$	-0.0093	0.0000	0.0097	0.0000	-0.0003	0.0000	-0.0001	0.0000
	(gain/loss)		0.0093		-0.0098		0.0003		0.0001
	$\Delta\mathcal{E}_I$	-0.0093	0.0000	0.0097	0.0001	-0.0004	-0.0001	-0.0001	0.0000
	(gain/loss)		0.0093		-0.0096		0.0003		0.0001
BiasedMF	$\Delta\mathcal{V}_R$	-0.0698	0.0000	0.0701	0.0003	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0698		-0.0698		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0698	-0.0215	0.0701	0.0218	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0483		-0.0483		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.1043	0.0000	0.1049	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1043		-0.1043		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.1043	-0.0296	0.1049	0.0302	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0747		-0.0747		0.0000		0.0000
SVD++	$\Delta\mathcal{V}_R$	-0.0698	0.0000	0.0701	0.0003	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0698		-0.0698		0.0000		0.0000
	$\Delta\mathcal{E}_R$	-0.0698	-0.0213	0.0701	0.0216	-0.0002	-0.0002	-0.0001	-0.0001
	(gain/loss)		0.0485		-0.0485		0.0000		0.0000
	$\Delta\mathcal{V}_I$	-0.1043	0.0000	0.1049	0.0006	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.1043		-0.1043		0.0000		0.0000
	$\Delta\mathcal{E}_I$	-0.1043	-0.0296	0.1049	0.0302	-0.0005	-0.0005	-0.0001	-0.0001
	(gain/loss)		0.0747		-0.0747		0.0000		0.0000