International Journal of Information Management Data Insights 5 (2025) 100311

Contents lists available at ScienceDirect

International Journal of Information Management Data Insights

journal homepage: www.elsevier.com/locate/jjimei

Enhancing recommender systems with provider fairness through preference distribution awareness

Elizabeth Gómez [a], David Contreras [b], Ludovico Boratto [c,*], Maria Salamó [a]

[a] Facultat de Matemàtiques i Informàtica, Universitat de Barcelona, Spain
[b] Facultad de Ingeniería y Arquitectura, Universidad Arturo Prat, Iquique, Chile
[c] Department of Mathematics and Computer Science, University of Cagliari, Italy

ARTICLE INFO

Keywords:
Recommender systems
Re-ranking algorithm
Algorithmic fairness
Geographic demographic groups
Data imbalance
Disparate impact

ABSTRACT

Going beyond recommendations' effectiveness, by ensuring properties such as unbiased and fair results, is an aspect that is receiving more and more attention in the literature. This means not only providing accurate recommendations but also ensuring that the visibility of providers aligns with user preferences and demographic representation, which has been identified as a key aspect of fairness in recommender systems. In particular, *provider fairness* enables the generation of results which are equitable for different (groups of) providers. In this paper, we raise the problem of *how recommendations are distributed when enabling provider fairness*. Indeed, on the one hand, users have clear preferences with respect to which providers they choose (*e.g.*, Italian users mostly buy Italian food), so recommendations should reflect these preferences. On the other hand, content providers should be able to reach a diverse audience, and be visible across the different user groups that expressed a preference for them. Specifically, we consider demographic groups based on their continent of origin for both users and providers, and assess how the preferences of the user groups are distributed across the provider groups. We first show that the state-of-the-art models and the existing approaches that enable provider fairness do not reflect the original distribution of the user preferences. To enable this property, we propose a re-ranking approach that, thanks to the use of buckets associating users and items, favors what we call *preference distribution-aware provider fairness*. Results on two real-world datasets (*i.e.*, the Book-Crossing and COCO) show that our approach can enable provider fairness and tailor the recommendations to the original distribution of the user preferences, with negligible losses in effectiveness. In particular, in the Books dataset, our approach obtains an overall disparity that is around 6%. On the other hand, in the case of the COCO dataset, the disparities are reduced to 2%.

1. Introduction

Motivation. The impact of modern information access systems on their stakeholders is a critical concern (Ekstrand, Das, Burke, & Diaz, 2022b). This impact is increasingly investigated, particularly in efforts to minimize bias (Baeza-Yates, 2018) and ensure fairness according to decision criteria (Bauer, Hinz, van der Aalst, & Weinhardt, 2021; Deldjoo, Jannach, Bellogín, Difonzo, & Zanzonelli, 2022). **However, achieving unbiased and fair outcomes remains a challenge**, as fairness is subjective and context-dependent. For instance, in human resource management, as Hatif Abd Almajed, Nassreddine, and Younis (2024) illustrate, machine learning models are introduced to improve accuracy in candidate selection, but these models often risk perpetuating biases embedded in training data.

The importance of these issues is highlighted by legal regulations such as Article 5(1)(a) of Europe's General Data Protection Regulation (GDPR), which stipulates that personal data must be processed "lawfully, fairly and transparently". Given that recommender systems are central to modern information access, it becomes crucial to align them with these legal and ethical expectations. This has directed significant attention toward fairness, particularly algorithmic fairness, which seeks to prevent negative societal impacts on both item consumers and providers.

These challenges align closely with the Design Science paradigm in Information Systems, as outlined by Hevner, March, Park, and Ram (2004). Design Science seeks to create and evaluate innovative artifacts (Ken Peffers, Tuunanen, & Chatterjee, 2007) (such as algorithms

* Corresponding author.
E-mail addresses: egomezye13@alumnes.ub.edu (E. Gómez), david.contreras@unap.cl (D. Contreras), ludovico.boratto@acm.org (L. Boratto), maria.salamo@ub.edu (M. Salamó).

https://doi.org/10.1016/j.jjimei.2024.100311

Available online 28 December 2024

or AI systems) that address both organizational and societal challenges. In the context of AI fairness, the problem is not just technical – ensuring accurate predictions – but also ethical, balancing multiple stakeholder needs such as fairness, transparency, and accountability.

While many approaches have focused on fairness for item consumers, **the challenge of ensuring fairness for item providers has received less emphasis**. Ensuring provider fairness means that items from different providers, especially from underrepresented groups, are recommended equitably. In this work, we focus on **provider fairness**, specifically from the perspective of group fairness. This means ensuring that providers from different demographic groups receive a proportionate number of recommendations, in line with their representation in the training data. Although progress has been made in this area (Ashokan & Haas, 2021; Boratto, Fenu, Marras, & Medda, 2022; Geyik, Ambler, & Kenthapadi, 2019; Gharahighehi, Vens, & Pliakos, 2021; Gómez, Boratto and Salamó, 2021, 2022; Gómez, Shui Zhang, Boratto, Salamó and Marras, 2021; Gómez, Zhang, Boratto, Salamó and Ramos, 2022; Liu & Burke, 2018; Marras, Boratto, Ramos, & Fenu, 2021; Patro, Biswas, Ganguly, Gummadi, & Chakraborty, 2020; Sonboli & Burke, 2019), current work primarily focuses on ensuring enough visibility or exposure to a provider (group), by considering how many times we recommend the items of a provider group and in which positions in the ranking, without adequately considering **to whom** these recommendations are made.

Open issue. Despite these advances, a critical gap remains: **existing methods fail to account for the audience to whom items from different provider groups are recommended.** Understanding **who** receives the recommendations of a given provider group is essential, as user preferences often follow specific patterns influenced by cultural, linguistic, or geographic factors. Indeed, as we will later show in Section 4, users have clear patterns when expressing their preferences, regarding which providers they choose for the items they rate. This can range from a pure attachment to the local culture (e.g., in cinema and literature), to linguistic reasons (e.g., learners of online courses favoring items in languages they are familiar with). As a result, ignoring these patterns risks recommending items to user groups that may have no interest or cultural connection.

Moreover, ensuring that a provider's items reach a diverse audience, aligned with user preferences, can help providers expand their reach across different user segments, ultimately benefiting their visibility and business prospects. For instance, enabling Swiss chocolate producers to be recommended to all countries with an interest in chocolate products increases their market exposure. This not only supports business growth but also aligns with broader ethical imperatives, as developing AI-driven recommender systems that consider fairness, inclusion, and equality can enhance consumer trust and brand credibility in marketing (Dwivedi et al., 2021). In an era where fairness and transparency are paramount, ensuring that recommendations are both equitable and aligned with user preferences helps providers maintain visibility in a way that respects these values (Tsao, 2013). Therefore, the dual perspective of ensuring that recommendations are made to users who exhibit preferences for specific providers, while maintaining equity in provider visibility, is a critical and unresolved challenge. We introduce this dual concern as **preference distribution-aware provider fairness**.

Our contributions. To address this challenge, we propose a novel approach to provider fairness that considers both **the demographic origin of providers** and **the distribution of user preferences** across provider groups. We focus on the continent of origin to create demographic groups of providers and users, but our approach is adaptable to other attributes. This alignment of recommendations with user preferences allows us to achieve preference distribution-aware provider fairness, ensuring that the exposure of provider groups matches the preferences expressed in the training data.

The contextualization of our work with the state of the art, presented in Section 4, shows that neither the original recommendation models, nor the recent advances in provider fairness, can enable *distribution-aware provider fairness*. Indeed, the classic recommendation models exacerbate original imbalances in the data (thus leading to provider unfairness), while existing provider fairness approaches do not respect the distribution of the preferences for the different user groups. This means that *(i)* users will receive redundant recommendations with respect to what providers they are recommended, not being exposed to smaller but interesting provider segments and that *(ii)* providers will have limited possibilities to reach users segments potentially interested in their items, thus limiting their business.

To mitigate these phenomena, in this paper, we propose an approach that builds buckets, containing the recommendations associated with a user continent and an item continent pair (e.g., the Europe-North America bucket will contain all the recommendations to European users of items produced in North America). The buckets are used to build the recommendation lists, respecting the distributions of preferences of each user group, and following the representation in the data of each provider group.

To validate our approach, we perform an extensive analysis on two real-world datasets. Our results show that our approach achieves provider fairness, targeting the recommendation of a provider group to the user groups that are interested in those providers, thus enabling preference distribution-aware provider fairness. This perspective can be guaranteed with negligible losses in terms of effectiveness.

Our contributions can be summarized as follows:

- We formulate the problem of **preference distribution-aware provider fairness** and assess that the state-of-the-art models cannot enable this property in their results;
- We devise a re-ranking algorithm that, thanks to the use of buckets, builds new rankings that, keeping items' relevance into account, distributes recommendations following the preferences of the user groups and the representation of the provider groups in the data;
- We validate our approach on two real-world datasets, in the book and education domains, showing the capability of our approach to enable *preference distribution-aware provider fairness* while preserving recommendation effectiveness.

The rest of the paper is structured as follows: in Section 2, we cover related work and, in Section 3, we provide the foundations to our study. We assess the limits at the state of the art in Section 4 and mitigate disparities in Section 5. We validate our proposal in Section 6, and provide concluding remarks in Section 8.

2. Related work

There exists a considerable body of literature related to the concern of *fairness* in recommender systems. A nice and detailed introduction to the concepts related to fairness is presented in Ekstrand, Das, Burke, and Diaz (2022a) and Ekstrand et al. (2022b). This section aims at providing references related to our work. We have focused on ranking systems although there is also literature on rating systems, such as Ashokan and Haas (2021).

Provider Fairness. A critical overview of fair ranking systems can be found in the recent work by Patro et al. (2022). Ranking systems are present in many of today's online content applications, such as LinkedIn Geyik et al. (2019). Note that *provider fairness* is concerned with how different providers, either individually or as members of protected groups, have their items appear (or not) in the rankings produced by a recommender system (Ekstrand et al., 2022b).

Since there is no universal definition of what constitutes a fair system due to its dependence on the studied scenario (Barra, Marras, & Fenu, 2018; Dessì, Fenu, Marras, & Reforgiato Recupero, 2017;

Green & Hu, 2018), many proposals have emerged in the literature. Some of them discover and measure discrimination, while others, in addition, deal with fairness by mitigating or reducing discrimination. Just to mention some, common scenarios in recommender systems are course recommender systems (Gómez, Zhang et al., 2022; Marras et al., 2021), book recommendation (Ekstrand & Kluver, 2021; Ekstrand, Tian, Kazi, Mehrpouyan, & Kluver, 2018), grant loans to low-income entrepreneurs (Liu & Burke, 2018; Sonboli & Burke, 2019), and news recommendation (Gharahighehi et al., 2021), among others.

There exists approaches, such as Karakolis, Kokkinakos, and Askounis (2022), whose goal is to provide fair recommendations across item providers in terms of diversity and coverage[1] for users to whom each provider's items are recommended. The authors addressed the solution as a mathematical optimization problem under constraints for coverage and diversity and proposed a new heuristic algorithm. Previous studies have also emphasized the use of visibility and exposure[2] to assess fairness. Given a ranking, visibility and exposure metrics respectively assess the amount of times an item is present in the rankings (Fabbri, Bonchi, Boratto, & Castillo, 2020; Zehlike et al., 2017) and *where* an item is ranked (Biega, Gummadi, & Weikum, 2018; Zehlike & Castillo, 2020). Over time, a large number of existing studies in the broader literature have examined the use of these metrics on guaranteeing *individual* or *group* fairness. For instance, Biega et al. (2018) defined measures to capture unfairness at the level of individual subjects, and subsumed group fairness as a special case. On the other hand, at the *group* level, these metrics make sure that users belonging to different groups are given adequate visibility or exposure (Diaz, Mitra, Ekstrand, Biega, & Carterette, 2020; Zehlike et al., 2017; Zehlike & Castillo, 2020). One example is the ranked group fairness definition presented in Zehlike et al. (2017), which extends group fairness using the standard notion of protected groups. Zehlike and Castillo (2020) describe an approach that measures discrimination and unequal opportunity in rankings at training time in terms of discrepancies in the average group exposure. On the other hand, under the group setting, the visibility/exposure of a group is proportional to its representation in the data (Patro et al., 2020; Ramos & Caleiro, 2020; Sapiezynski, Zeng, Robertson, Mislove, & Wilson, 2019; Yang & Stoyanovich, 2017).

Fairness metrics. Several metrics have been proposed to assess fairness, some focusing on the aforementioned visibility and exposure, others on defining a different kind of metrics. Mehrotra, McInerney, Bouchard, Lalmas, and Diaz (2018) assess unfairness based on the popularity of the providers, they define a fairness metric that rewards recommendation lists that are diverse in terms of popularity bias. Raj and Ekstrand present a comparative analysis, considering both the theoretical and practical perspectives, among several fairness metrics recently introduced to measure fair ranking (Raj & Ekstrand, 2022). Wu, Mitra, Ma, Diaz and Liu (2022), based on the *expected exposure* metric, formalize a family of exposure fairness metrics that model the problem of fairness jointly from the perspective of both types of stakeholders. In particular, they consider group attributes on both the user side and item side to identify and mitigate fairness concerns.

User relevance and fairness. 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. Kamishima, Akaho, Asoh, and Sakuma (2018) introduced the concept of recommendation independence and define an objective function with three components (a loss function, an independence term, and a regularization term) that is in charge of returning a value with the smallest loss and the largest independent term. Zehlike et al. (2022) presented the extension of FA*IR to multiple groups that guarantees ranked group fairness, without introducing a large utility loss. The proposed algorithm can be used when more than one protected group is present, by leveraging a multinomial distribution. In the same line of research, the analysis of disadvantaged groups in the ranking list, Tahery, Aftabi, and Farzi (2021) extended the analysis to those items that belong to more than two protected groups. The authors proposed an algorithm named FARGO, which is a genetic algorithm enhanced by the simulated annealing that can handle any number of protected groups, and also define a new evaluation metric – named Expected Gain Ratio (EGR) – to assess the output of a fair ranking algorithm. More recently, Naghiaei, Rahmani, and Deldjoo (2022) proposed a *CP-fairness* method, which integrates fairness constraints from both the consumer and the provider-side in an optimization-based re-ranking approach. Burke, Mattei, Grozin, Voida, and Sonboli (2022) introduce a framework, whose architecture is named SCRUF-D, in which the interest of providers and other stakeholders are presented as agents that participate in the production of recommendations through a two-stage social choice mechanism. The framework allows for many different types of fairness concerns. Finally, Wu, Ma, Mitra, Diaz and Liu (2022) propose a multi-objective optimization framework, named *Multi-FR*, for fairness-aware recommendation that adaptively balances accuracy and fairness in multi-sided marketplaces, where the final solution is guaranteed to be Pareto optimal.

Contextualizing our work. A closer look to the literature on *fairness* reveals that in provider fairness the goal of the proposals is to ensure enough visibility to a provider (group) without paying attention to whom the recommendation of a given provider are proposed. This paper addresses *preference distribution-aware provider fairness*, so far lacking in the scientific literature.

Concretely, this is the first time that unfairness phenomena for content providers and users belonging to different continents are tackled. In this work, we propose an algorithm that ensures minority providers have a diverse audience, while aligning the original preferences of the users. Our approach lets providers reach their business to a greater extent while maintaining the users' preferences.

3. Preliminaries

In this section, we provide the foundations to our work. It is organized as follows: Section 3.1 details the recommendation scenario and Section 3.2 describes the metrics used in our study. Next, the recommendation algorithms are presented in Section 3.3. We conclude this section with a description of the dataset used in this study, in Section 3.4.

3.1. Recommendation scenario

Let U be a set of users, I be a set of items, and 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 results from a map $r : U \times I \rightarrow V$, where V is the rating domain. If $r(u,i) = \perp$, then we say that u did not rate i. To easy notation, we denote $r(u,i)$ by r_{ui}. We define the set of ratings as R and they can directly feed an algorithm in the form of triplets (point-wise approaches) or shape user-item observations (pair-wise approaches).

We consider a temporal split of the data, where a fixed percentage of the ratings of users (ordered by timestamp) goes to the training and the rest goes to the test set (Bellogín, Castells, & Cantador, 2017). This is done to assess the real impact of the recommendations.

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. $\hat{R}$ is a

[1] Diversity is often used in reference to the system's ability to recommend different items to different users, while coverage reflects the degree to which the generated recommendations cover the catalog of available items (Ge, Delgado-Battenfeld, & Jannach, 2010; Kaminskas & Bridge, 2016).

[2] Exposure measures the value of recommendation opportunities given to a particular item or group of items.

Table 1
Summary of the terminology used in the article. First column details the notation of the term, second column presents the formal definition, while the last column describes the concept.

Term	Definition	Concept
U	$U = \{u_1, u_2, \dots, u_n\}$	Set of users
I	$I = \{i_1, i_2, \dots, i_j\}$	Set of items
V	V is the rating domain	Set of preferences
R	$R = \{(u, i, r_{ui}) : u \in U, i \in I, r_{ui} \neq \perp\}$	Set of ratings
r_{ui}	$r(u, i)$	Rating of user u over item i
$\hat{r}_{ui}$	$\hat{r}(u, i)$ when $r_{ui} = \perp$	Predicted relevance of item i for user u
$\hat{R}$	$\hat{R} = \{(u, i, \hat{r}_{ui}) : u \in U, i \in I\}$	Set of recommendations
P	$P = \{p_1, p_2, \dots, p_g\}$	Provenance (set of geographic continents)
P_i	$P_i \subseteq P$	Provenance (set of geographic continents associated with item i)
p_u	$p_u \in P$	Provenance (geographic continent associated with user u)
G_j^U	$G_j^U = \{r_{ui} : p_u = p_j\}$	User-based demographic group
G_j^I	$G_j^I = \{r_{ui} : p_j \in P_i\}$	Item-based demographic group

set of items ordered from most relevant to least relevant for each user. More details of definitions are shown in Table 1.

Let $P = \{p_1, p_2, \dots, p_g\}$ denote the provenance, either of an item or of a user. As mentioned in Section 1, this paper focuses on the continent of origin. Thus, P is the set of g geographic continents associated with the items and the users. According to the entity we are modeling, we consider a geographic continent as the continent of origin of either a user that has rated an item, or of an item provider. We denote as P_i the set of geographic continents associated with an item i. Note that, since an item could be produced by more than one provider, several geographic continents may appear in a item, and thus, $|P_i| \geq 1$. In case two providers belong to the same geographic continent, that continent appears only once; this choice was made since we are dealing with group fairness so, when a group of providers is associated with an item (either once or multiple times), we account for the presence of that group. Moreover, we denote as $p_u \in P$ the geographic continent associated with a user u (clearly, each user is associated only with a continent). We use the geographic continents to shape g demographic groups, which can be defined either from the user side, to group the ratings of the users in a continent (we denote these groups as G_j^U, where j indicates a continent $p_j \in P$), or from the item side to group the ratings of the items produced in a continent (we denote these groups as G_j^I, where j indicates a continent $p_j \in P$).

Table 1 summarizes the formal definitions of the terms described above.

3.2. Metrics

This section describes the metrics used in our analysis and experiments, i.e., the representation of the items of a group in relation to the ratings received by the users belonging to a group, and disparate visibility.

Representation. Based on the demographic groups characterizing the user preferences and the item production, we can define the *representation of a group of providers belonging to a continent q (G_q^I) when considering the ratings given by a user demographic group belonging to a continent j (G_j^U)*:

$$\mathcal{R}(G_j^U, G_q^I) = \frac{|G_j^U \cap G_q^I|}{|R|} \tag{1}$$

The metric produces values between 0 and 1.

We compute the representation only considering the training set.

It should be clear that Eq. (1) defines our goal to achieve a preference distribution-aware provider fairness. Indeed, it measures to what extent the users in a continent p_j have rated the items produced in a continent p_q. If the recommendations follow the same distribution of the representation, we are able to provide equity to providers and respect the preferences of the users for the providers. We assess unfairness with the notion of *disparate impact* generated by a recommender system. Specifically, we assess disparate impact with the Disparate visibility metric.

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 a group G_q^I to the users of a group G_j^U and the representation in the input data:

$$\Delta\mathcal{V}(G_j^U, G_q^I) = \frac{|\{\hat{r}_{ui} : u \in G_j^U, i \in G_q^I\}|}{|\hat{R}|} - \mathcal{R}(G_j^U, G_q^I) \tag{2}$$

The range of values for this score is $[-\mathcal{R}(G_j^U, G_q^I), 1 - \mathcal{R}(G_j^U, G_q^I)]$; 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 Fabbri et al. (2020).

3.3. Recommendation algorithms

In this study, we focus on well-known state-of-the-art Collaborative Filtering algorithms. Specifically, our experiments are focused on the comparison of algorithms based on *memory-based approaches* (i.e., item-based k-nearest neighbors, **ItemKNN** (Sarwar, Karypis, Konstan, & Riedl, 2001), and user-based k-nearest neighbors, **UserKNN** (Herlocker, Konstan, & Riedl, 2002)) and *matrix factorization based approaches* (i.e., Bayesian Personalized Ranking with Matrix Factorization, **BPRMF** (Rendle, Freudenthaler, Gantner, & Schmidt-Thieme, 2012), and Singular Value Decomposition, **SVDpp** (Koren, 2008)).

Our baseline is a proposal (Gómez, Boratto et al., 2022) that introduces equity for providers, 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), thus controlling fairness of providers coming from different continents.

3.4. Datasets

Our analysis focuses on two different domains, *books* and *courses*, exploring the role of the geographic provenance of providers and users in the recommendation process. It is important to note that both datasets are real-world data and they have been widely used in previous studies to evaluate recommender systems and fairness recommendation (Ekstrand & Kluver, 2021; Ekstrand et al., 2018; Gómez, Zhang et al., 2022; Marras et al., 2021).

Books. The Book-Crossing (Ziegler, McNee, Konstan, & Lausen, 2005) dataset contains 394,163 ratings (in the range [1–10]), given by 46,899 users, to 12,314 books. In addition, this dataset provides demographic information of its users, by offering age and location attributes, if the user has provided them. We only use the location of the user and the remaining attributes are not relevant to our study, so any additional information of the users offered by this dataset will be considered in future work. For each book, the dataset provides its ISBN code, which has been used to retrieve information about the production continent, by exploiting the APIs offered by the Global Register of Publishers.[3]

Courses. The COCO (Dessì, Fenu, Marras, & Reforgiato Recupero, 2018) dataset comes from the education context. We use it to explore

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

Table 2
Group representation, $\mathcal{R}$, in the Books and Courses data. Percentage Representations of each pair of Users-based and Item-based Demographic Groups (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Groups appear in alphabetical order by the name of the continent.

Items											
		Books				Courses					
		EU	NA	OC	SA	AF	AS	EU	NA	OC	SA
Users	AF	0.003	0.022	0.000	0.000	0.720	0.688	1.524	8.089	0.542	0.051
	AN	0.001	0.006	0.000	0.000	–	–	–	–	–	–
	AS	0.103	1.117	0.001	0.000	1.143	1.740	4.125	11.591	1.577	0.044
	EU	2.036	5.063	0.004	0.001	0.790	1.213	4.289	9.372	1.085	0.192
	NA	3.807	86.106	0.009	0.000	1.638	3.137	6.871	25.808	2.294	0.270
	OC	0.219	1.359	0.008	0.000	0.122	0.287	0.445	1.617	0.397	0.007
	SA	0.009	0.127	0.000	0.000	0.310	0.236	0.894	3.061	0.229	3.601
		6.177	**93.801**	**0.022**	**0.001**	**4.723**	**7.301**	**18.147**	**59.539**	**6.124**	**4.166**

the role of the geographic provenance of teachers and learners in the recommendation process. As far as we know, COCO is the only educational dataset that contains the geographic provenance of the users. The dataset was collected from an online course platform, and includes 11,454 courses and 75,896 learners who gave 544,833 ratings (in the range [0.5–5]). It is important to remark that each course is associated with one or more teachers. The dataset contains the time zone of the teachers and learners and, with this information, we have inferred the continent they belong to.

4. Problem assessment

We first characterize, in Section 4.1, the representation per continent of origin of both users and items. Next, in Section 4.2, we present the outcome of different recommendation models considering both the geographic provenance of both users and items.

4.1. Group representation

Table 2 depict the rating-based representation, considering each possible pair of user- and item-based demographic groups in the Books and Courses data, respectively. Both sides of the table depict the groups in alphabetical order by the name of the continent. Note that, in Table 2, some continents in the Books dataset are not shown. This happens because, in the training set, there are no items produced by Africa (AF), Antarctica (AN), and Asia (AS). In addition, Oceania (OC) has an overall representation of 0.022% and for South America (SA) the percentage is negligible, with a value of 0.001% in Table 2. It is also worth mentioning that 93% of rated items are from North America (NA); hence, in our scenario, unfairness is linked to data imbalances related to selection and representation bias (Gómez, Boratto et al., 2021). From this percentage, 86% of the items are mainly rated by users that also are in NA, thus confirming the cultural attachment of the users of a continent to the items produced in the same continent. In fact, in this domain, most of the items and users are from NA and Europe (EU).

Conversely, in Table 2, we can observe the group representation in the Courses data. Similarly to the Books domain, in this dataset, there are neither items nor users from Antarctica. Nevertheless, as depicted in Table 2, the representation has a broader geographical distribution between items and users. We notice that there is a more distributed representation among the items of the different continents coming from users from different continents. Again, the majority of items (i.e., 59%) are produced in NA but only 25% have been rated by users from NA (this 25%, however, still represents the largest share of ratings received by NA items from the users of a single continent).

Finally, we observe that, in both datasets, North American users are the ones who provided the majority of ratings. A similar effect occurs in the providers of items, being North Americans providers the most rated (with the exception of SA in the Courses dataset, where the majority of ratings were given by users from the same continent, i.e., SA).

4.2. Disparate visibility in the state-of-the-art models

This section examines the disparities in continent distribution for each state-of-the-art model in comparison to the training set across both the books and courses datasets. Tables 3 and 4 measure, for each recommendation model described in Section 3.3, the disparate visibility metric $\Delta\mathcal{V}$, to assess the difference between the distribution of the recommendations and that of the preferences in the training data.

The algorithm that presents the greatest disparities is BPRMF in both datasets. It can be observed that BPRMF gives more recommendations to NA items than expected, impairing the visibility of providers from the other continents. This also occurs in the Books dataset with the SVDpp algorithm. In contrast, ItemKNN shows an opposite behavior to that of BPRMF and SVDpp in the Books dataset. Since ItemKNN gives less visibility to NA, items from the other continents are over-represented in the recommendation lists. Note that UserKNN also under-represents NA items. Regarding the Courses dataset, we can see that NA is under-represented by SVDpp, UserKNN, and ItemKNN.

Observation 1. *The state-of-the-art recommendation models we have considered are unable to guarantee provider fairness and to respect the distribution of the original preferences of the user, by usually exacerbating the majority group. In other words, the items produced by the majority groups are over-recommended and the distribution of the recommendations to the different user groups does not respect that of the original preferences. Nevertheless, different models generate different disparities, highlighting the need for algorithmic interventions that mitigate these phenomena.*

5. Preference distribution-aware provider fairness

In this section, we detail the motivation behind our proposal, see Section 5.1, and we describe the *preference Distribution-Aware Provider fairness* (DAP-fair) algorithm to mitigate continent distribution imbalances in the recommendation lists, see Section 5.2.

5.1. Motivation behind our approach

In Section 4.2, we have assessed the presence of disparities in the way recommendations are distributed across different continents, for both the users and the providers. We identified imbalances in the recommendation lists of well-known recommender systems, focusing on users' preferences and item provenance based on geographic distribution. One common way to solve the imbalance in the recommendations is to use a provider fairness algorithm.

In this section, to motivate the need of our approach, we analyze the behavior of a provider fairness algorithm that does not consider how the preferences of users from different countries are distributed. We have addressed this issue with the analysis of the results of an algorithm (Gómez, Boratto et al., 2022) whose basis is founded in enabling provider fairness in the context of geographic imbalance in

Table 3
Disparate visibility, $\Delta\mathcal{V}$, in the continent distribution of original models to the training set in the Books dataset. Values in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Underlined numbers denote a large disparity, either positive or negative.

Books									
Items									
		BPRMF				SVDpp			
		EU	NA	OC	SA	EU	NA	OC	SA
	AF	−0.003	0.039	0.000	0.000	0.000	0.035	0.000	0.000
U	AN	−0.001	−0.002	0.000	0.000	−0.001	−0.002	0.000	0.000
s	AS	−0.103	−0.047	−0.001	0.000	−0.070	−0.079	−0.001	0.000
e	EU	−2.036	4.984	−0.004	−0.001	−1.474	4.420	−0.002	−0.001
r	NA	−3.807	0.840	−0.009	0.000	−1.199	−1.772	−0.005	0.000
s	OC	−0.219	0.311	−0.008	0.000	−0.147	0.239	−0.008	0.000
	SA	−0.009	0.076	0.000	0.000	−0.003	0.069	0.000	0.000
		UserKNN				ItemKNN			
	AF	0.004	0.032	0.000	0.000	0.005	0.030	0.000	0.000
U	AN	−0.001	−0.002	0.000	0.000	−0.001	−0.002	0.000	0.000
s	AS	−0.018	−0.132	−0.001	0.000	−0.008	−0.142	−0.001	0.000
e	EU	1.687	1.257	−0.001	0.001	2.022	0.909	0.000	0.012
r	NA	0.117	−3.091	−0.001	0.000	1.685	−4.691	0.018	0.013
s	OC	−0.017	0.109	−0.008	0.000	0.018	0.071	−0.005	0.000
	SA	0.007	0.058	0.001	0.000	0.008	0.057	0.000	0.000

Table 4
Disparate visibility, $\Delta\mathcal{V}$, in the continent distribution of original models to the training set in the Courses dataset. Values in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Underlined numbers denote a large disparity, either positive or negative.

Courses													
Items													
		BPRMF						SVDpp					
		AF	AS	EU	NA	OC	SA	AF	AS	EU	NA	OC	SA
U	AF	−0.720	−0.688	−1.477	3.964	−0.359	−0.051	−0.679	0.109	0.256	0.308	0.434	0.010
s	AS	−1.143	−1.740	−4.109	8.712	−1.467	−0.044	−1.026	0.187	0.443	0.082	0.579	0.000
e	EU	−0.790	−1.213	−4.273	7.067	−0.992	−0.192	−0.718	0.568	−0.858	0.151	0.421	−0.024
r	NA	−1.638	−3.137	−6.834	13.213	−2.082	−0.270	−1.465	1.215	0.880	−2.740	1.523	−0.151
s	OC	−0.122	−0.287	−0.443	1.119	−0.384	−0.007	−0.109	−0.013	0.074	0.050	−0.127	−0.003
	SA	−0.310	−0.236	−0.888	5.610	−0.187	−3.601	−0.114	0.148	−0.030	−0.071	0.232	0.458
		UserKNN						ItemKNN					
U	AF	−0.030	−0.076	0.048	0.588	0.006	0.003	−0.036	0.120	−0.020	0.424	0.073	0.003
s	AS	0.057	−0.170	0.304	0.001	0.007	0.005	−0.021	0.095	0.062	−0.156	0.143	0.014
e	EU	−0.085	−0.130	0.039	−0.138	−0.074	0.021	−0.131	0.070	0.010	−0.353	−0.007	0.026
r	NA	0.043	−0.178	0.517	−1.229	−0.073	0.037	−0.159	0.256	0.137	−1.241	0.001	0.062
s	OC	−0.010	0.037	−0.012	−0.075	−0.056	0.001	−0.030	0.006	−0.030	−0.116	0.044	0.003
	SA	0.021	−0.032	−0.004	−0.016	−0.014	0.666	0.007	−0.010	0.009	0.249	−0.007	0.504

Table 5
Disparate visibility, $\Delta\mathcal{V}$, in the continent distribution of the baseline (Gómez, Boratto et al., 2022) models in the Books Dataset. Values are in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Underlined numbers denote a large disparity, either positive or negative.

Books									
Items									
		BPRMF				SVDpp			
		EU	NA	OC	SA	EU	NA	OC	SA
	AF	−0.003	0.039	0.000	0.000	0.000	0.036	0.000	0.000
U	AN	−0.001	−0.002	0.000	0.000	−0.001	−0.002	0.000	0.000
s	AS	−0.103	−0.047	−0.001	0.000	−0.072	−0.077	−0.001	0.000
e	EU	−2.036	4.984	−0.004	−0.001	−1.516	4.462	−0.001	−0.001
r	NA	−3.807	0.840	−0.009	0.000	−1.257	−1.714	−0.005	0.000
s	OC	−0.219	0.311	−0.008	0.000	−0.150	0.242	−0.008	0.000
	SA	−0.009	0.076	0.000	0.000	−0.004	0.070	0.000	0.000
		UserKNN				ItemKNN			
	AF	0.005	0.031	0.000	0.000	0.006	0.030	0.000	0.000
U	AN	−0.001	−0.002	0.000	0.000	−0.001	−0.002	0.000	0.000
s	AS	−0.020	−0.130	−0.001	0.000	0.001	−0.152	0.000	0.000
e	EU	1.717	1.227	−0.002	0.001	2.036	0.892	0.001	0.014
r	NA	0.346	−3.323	0.001	0.000	1.964	−4.980	0.026	0.015
s	OC	−0.006	0.098	−0.008	0.000	0.028	0.061	−0.006	0.000
	SA	0.008	0.057	0.001	0.000	0.011	0.054	0.000	0.000

Table 6
Disparate visibility, $\Delta\mathcal{V}$, in the continent distribution of the baseline (Gómez, Boratto et al., 2022) models in the Courses Dataset. Values are in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Underlined numbers denote a large disparity, either positive or negative.

Courses													
Items													
		BPRMF						SVDpp					
		AF	AS	EU	NA	OC	SA	AF	AS	EU	NA	OC	SA
U	AF	−0.720	−0.637	−0.962	3.222	−0.182	−0.051	−0.651	0.088	0.530	−0.113	0.512	0.010
s	AS	−1.143	−1.652	−3.259	7.372	−1.063	−0.044	−0.986	0.166	0.761	−0.637	0.897	−0.001
e	EU	−0.790	−1.141	−3.585	5.990	−0.675	−0.192	−0.631	0.534	−0.551	−0.490	0.690	−0.025
r	NA	−1.638	−2.974	−5.209	10.642	−1.307	−0.270	−1.231	1.150	1.778	−4.408	2.117	−0.152
s	OC	−0.122	−0.275	−0.326	0.934	−0.327	−0.007	−0.091	−0.017	0.130	−0.076	−0.070	−0.001
	SA	−0.310	−0.197	−0.531	5.041	−0.010	−3.601	−0.100	0.144	0.013	0.053	0.263	0.395
		UserKNN						ItemKNN					
U	AF	−0.040	−0.071	0.005	0.639	0.004	0.004	−0.052	0.120	−0.027	0.449	0.078	0.000
s	AS	0.037	−0.132	0.167	0.130	−0.009	0.004	−0.057	0.114	0.004	−0.103	0.148	0.012
e	EU	−0.093	−0.096	−0.089	−0.032	−0.078	0.019	−0.133	0.090	−0.062	−0.304	0.003	0.023
r	NA	0.025	−0.099	0.269	−1.018	−0.086	0.036	−0.181	0.317	0.013	−1.159	0.028	0.054
s	OC	−0.011	0.042	−0.024	−0.063	−0.058	0.002	−0.030	0.009	−0.041	−0.105	0.041	0.003
	SA	0.032	−0.023	−0.007	−0.017	−0.015	0.647	0.018	−0.002	0.009	0.247	−0.008	0.480

a multi-group setting. Tables 5 and 6 show the differences in the continent distribution of the baseline algorithm to the input data. It can be observed that the baseline algorithm is able to introduce fairness for providers belonging to different geographic areas, by re-distributing the recommendations across continents following a notion of *equity*, and it lowers the geographic imbalances. However, the results show that there is still imbalances in the users' preferences. We conjecture that this is largely due to the fact that this algorithm does not contemplate this aspect.

Observation 2. *When geographic imbalances are present in both users' preferences and providers, it is not enough just to consider provider fairness to mitigate these imbalances. For this reason, we need to keep both perspectives (i.e, the users' preferences and the providers' provenance) as goals to enable provider fairness in the context of geographic imbalance.*

5.2. Algorithm

Our analysis in Section 5.1 reveals that provider fairness algorithms fail to adequately meet users' preferences, indicating the need for an alternative approach. In order to address both perspectives (i.e., the user preferences and the providers' provenance), we propose an algorithm – named DAP-fair – that is aware of the distribution of all of them in the data.

Our DAP-fair algorithm is based on the idea of including in the recommendation lists first the items that are (i) produced in an under-represented continent and (ii) rated by an under-represented user group, by selecting the items that cause the minimum loss in prediction for the users. This process continues until their target percentage in the training set is reached. Our approach to introducing provider fairness via a re-ranking is the only possible one to provide equity guarantees since we keep changing the recommendation list until the desired distribution is reached.

Fig. 1 depicts the high-level description of the process for providing recommendation lists with distribution-aware provider fairness. Formally, the pseudocode of DAP-Fair is described in Algorithm 1. DAP-Fair works following two main steps:

- **Step 1** – *Create a bucket list*: Given the items that were predicted as relevant for a user by the recommender system, we create a bucket list considering each user-continent and item-continent pair, which will store the predicted items. We remind the reader that buckets contain the recommendations associated with a user continent and an item continent pair (e.g., the Europe-North America bucket will contain all the recommendations to European users of items produced in North America). Each bucket is characterized by an attribute representing the group in the data, $\mathcal{R}(G^U_j, G^I_q)$ (computed in Eq. (1)). Moreover, we also compute the loss of changing any item from the $top-n$ to the $top-k$. In other words, we measure the impact that the incorporation of any element, not currently in the recommendation list, would have if we added it to the list. In detail, we calculate the loss of each recommendation in the $top-n$ as $loss = \frac{itemRating - baseRating}{baseRating}$, where $baseRating$ corresponds to the rating of the last item in $top-k$ and $itemRating$ is the rating of the item considered to be moved to the $top-k$. Note that all items belonging to the $top-k$ have a loss of 0.
Specifically, the recommender system returns a list of top-n recommendations (where n is much larger than the cut-off value k, so as to be able to perform a re-ranking). Our starting point to fill a bucket is the relevance predicted for a user u and an item i, $\hat{r}_{ui}$. Each element in the bucket is a record that contains the user ID, the item ID, and the relevance predicted by the recommender system for that user, $\hat{r}_{ui}$. We sort each bucket by item relevance.
- **Step 2** – *Perform a selection of items*: we perform the re-ranking on the basis of the created bucket lists. The goal is to guarantee fairness for providers and to correctly distribute the recommendations among the different user groups. This step includes three phases, the first one is the most constrained and the conditions for selection are relaxed in the second and third phases, so to complete the recommendation list of the user.
  - **Phase 1**: we select items starting from the least represented continents to the most represented ones in their corresponding buckets. The algorithm selects items that satisfy the following **conditions**: **(1)** the loss is lower than the tolerance threshold; **(2)** the number of recommended items so far is lower than k; **(3)** the percentage of items in the recommendation list for a continent is lower or equal to the representation of the continent ($\mathcal{R}(G^U_j, G^I_q)$; **(4)** the continent of the item is the same as that of the bucket; and **(5)** the continent of the user is the same as that of the bucket. According to the way the algorithm works, it never chooses an element that has been previously selected.
  - **Phase 2**: we start this phase to include more items in the recommendation list when phase 1 finishes but the top-k is not complete. In this phase, the restrictions are relaxed. Specifically, the selection is made in the same way as done in phase 1, but this time we do not mind if the continent of the user is different. That is, condition 5 presented in phase 1 is not applied. With this, we enlarge the set of buckets where an item can be selected. Again, the recommendation

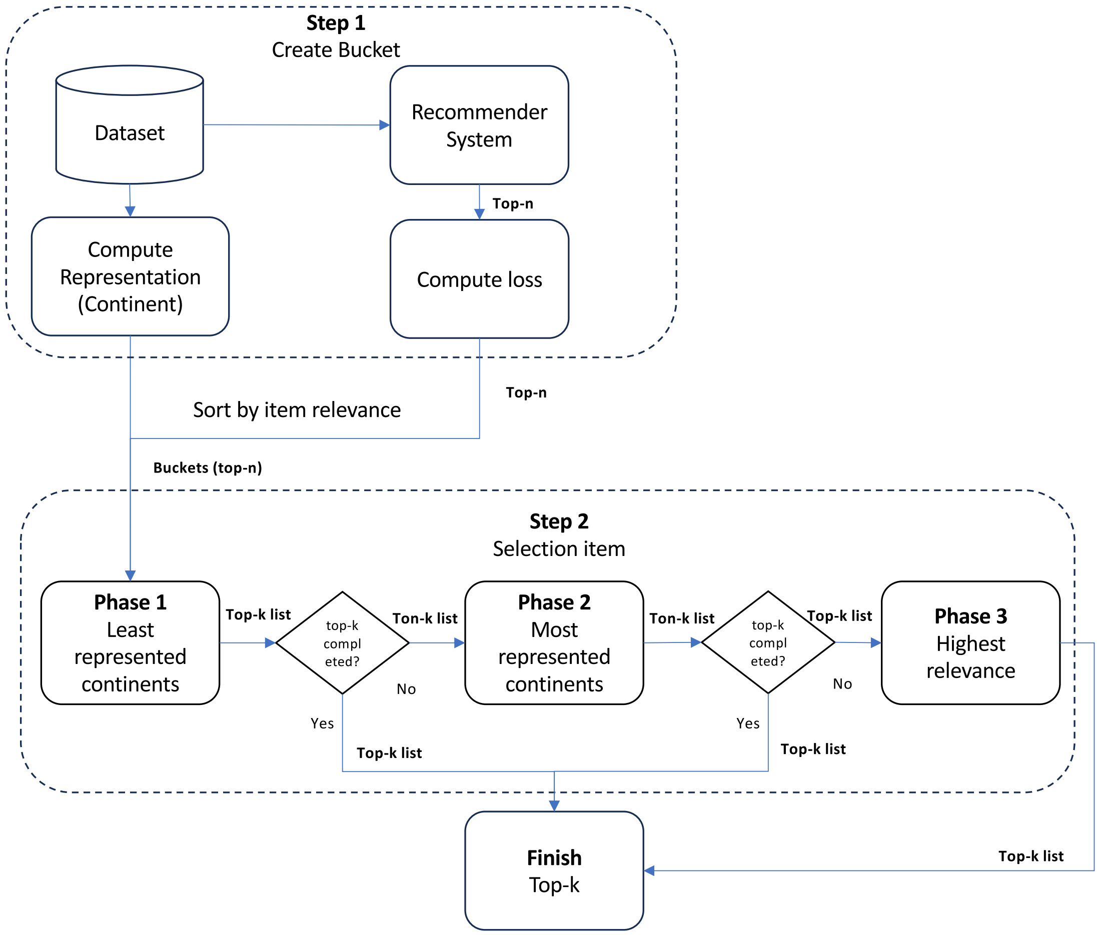

Fig. 1. General overview of the DAP-fair process for providing distribution-aware provider fairness.

list, top-k, may be completed or not. If completed, the process finishes, otherwise, it is necessary to move to phase 3.

- **Phase 3**: we complete the recommendation list if the top-k recommendations could not be reached due to the constraints in the previous phases. In this phase, we select the items from any bucket that have the greater relevance for the user, until we complete the top-k. That is, conditions 4 and 5 explained in phase 1 are not considered.

As mentioned above, the formal definition of the pseudo-code is depicted in Algorithm 1, and its support methods are described in Algorithms 2 and 3.

6. Assessment of the impact of DAP-fair

The goal of our study is to assess whether DAP-fair proposal is able to provide fairness and, at the same time, keep the distribution of the user preferences of the different continents into consideration. In this section, we analyze the impact of our mitigation approach by assessing the distribution in both perspectives in the data.

6.1. Experimental setting

We exploit the implementation of the state-of-the-art models provided by the *Elliot* framework (Anelli et al., 2021) to generate the recommendations for each user. This framework was chosen because it simplifies the process of implementing the algorithms (by having all the methods we considered already implemented) and it generates consistently formatted lists that can be fed to our DAP-fair algorithm.

For each user, we generated 1000 recommendations (denoted in the algorithm as the $top - n$) to then re-rank the $top - k$ (set to 10 in our experiments) through the proposed DAP-fair algorithm. We have performed a grid search of the hyper-parameters for each recommendation model in the two datasets.

For what concerns the ItemKNN and UserKNN models, the optimal hyper-parameters are nearly the same in both datasets. Both algorithms use a cosine *similarity*, and the classical *implementation*. ItemKNN uses 50 *neighbors* whereas UserKNN use 100.

The hyper-parameters for BPRMF and SVDpp, are defined with 10 *epochs* and 10 *factors* on each dataset. The *batch size* is 512 for SVDpp and is 1 for BPRMF on both datasets. Additional parameters are: for BPRMF *learning rate* = 0.1 in both datasets, *bias regularization* = 0.001 in Books and 0.1 in Courses, *user regularization* = 0.001 in Books and 0.1 in Courses, *positive item regularization* = 0.1 in both datasets, and *negative item regularization* = 0.001 in Books and 0.01 in Courses; for SVDpp *learning rate* = 0.001 in Books and 0.001 in Courses, *factors regularization* = 0.1 in both datasets, and *bias regularization* = 0.01 in both datasets.

To evaluate recommendation effectiveness, we measure the ranking quality of the lists by measuring the *Normalized Discounted Cumulative Gain* (NDCG) (Järvelin & Kekäläinen, 2002).

$$DCG@k = \sum_{u \in U} \hat{r}_{ui}^{pos} + \sum_{pos=2}^{k} \frac{\hat{r}_{ui}^{pos}}{log_2(pos)} \qquad NDCG@k = \frac{DCG@k}{IDCG@k} \tag{3}$$

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.

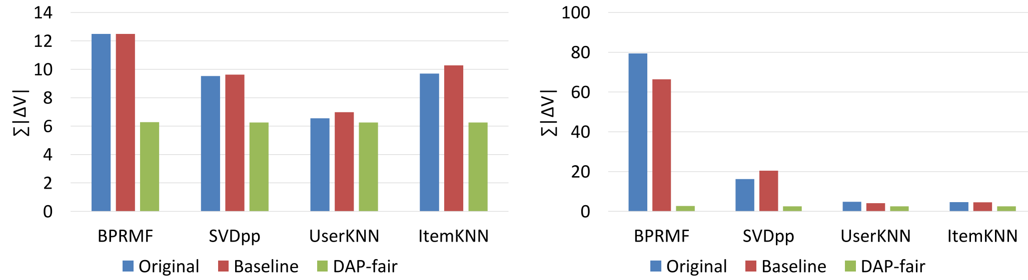

(a) Books (b) Courses

Fig. 2. Overall percentage of disparity between the models with different mitigation techniques (original corresponds to no mitigation, the baseline Gómez, Boratto et al., 2022, and the DAP-fair algorithm) in the books and courses datasets.

Table 7
Disparate visibility, $\Delta\mathcal{V}$, in the continent distribution of DAP-fair models with 100% loss tolerance in the Books Dataset. Values are in percentage of each group (AF: Africa, AN: Antartica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Bold numbers denote where the disparity has been mitigated.

Books									
Items									
		BPRMF				SVDpp			
		EU	NA	OC	SA	EU	NA	OC	SA
	AF	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**
U	AN	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**
s	AS	**0.000**	−0.150	**−0.001**	**0.000**	**0.000**	−0.151	**0.000**	**0.000**
e	EU	**0.000**	2.948	**−0.004**	**−0.001**	**0.000**	2.943	**0.000**	**0.000**
r	NA	**0.000**	−2.967	**−0.009**	**0.000**	**0.000**	−2.975	**0.000**	**0.000**
s	OC	**0.000**	0.092	**−0.008**	**0.000**	**0.000**	0.084	**0.000**	**0.000**
	SA	**0.000**	0.066	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**
		UserKNN				ItemKNN			
	AF	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**
U	AN	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**
s	AS	**0.000**	−0.151	**0.000**	**0.000**	**0.000**	−0.151	**0.000**	**0.000**
e	EU	**0.000**	2.943	**0.000**	**0.000**	**0.000**	2.943	**0.000**	**0.000**
r	NA	**0.000**	−2.975	**0.000**	**0.000**	**0.000**	−2.975	**0.000**	**0.000**
s	OC	**0.000**	0.084	**0.000**	**0.000**	**0.000**	0.084	**0.000**	**0.000**
	SA	**0.000**	0.066	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**

6.2. Disparate visibility of DAP-fair algorithm

Here, we analyze the disparate visibility, $\Delta\mathcal{V}$, after running the DAP-fair recommendations. Concretely, Tables 7 and 8 show, for each recommendation model, the difference in the continent distribution of the DAP-fair algorithm to the representation in the input data in the Books and Courses, respectively. As the tables show, although at the group level (i.e., continent of the provider), the providers almost reach in most cases the desired ratio (i.e., close to zero). At the bucket level (i.e., continent of the user and continent of the provider), we observe in both datasets that there is still some disparity (mainly in the buckets that have in common the NA provider items); the reason is that there are not enough elements in the recommendation lists to guarantee that the ratio is reached perfectly. We assume that with a recommendation list greater than 1000, the desired ratios can be achieved, but note that this goal will also reduce the quality of the recommendation lists.

6.2.1. Detailed analysis on the books dataset

In particular, in the Books dataset, we have observed that BPRMF has modified the recommendation lists with items from EU and NA continents, as the shortage of OC and SA items remained despite continuing to go down in the recommendation lists. Note that, in EU, the disparities were almost completely mitigated, considering both the user and item perspectives. In SVDpp, UserKNN, and ItemKNN the disparities at the group level (provider continent) were almost completely mitigated, however, at the bucket level in NA they are not; this is especially true at the item level. Again, we notice that when the recommendation model does not provide enough items to complete the groups, it is not possible to completely mitigate disparities, but we get close to the ideal value. It is important to remark that, in the Books dataset, the disparity in the original models (see Table 3) range in BPRMF from 4.984% to −3.807%, in SVDpp from 4.420% to −1.772%, in UserKNN from 1.687% to −3.091% and ItemKNN from 2.022% to −4.691%. When applying our mitigation algorithm (see Table 7) they range from 2.948% to −2.967% in BPRMF, from 2.943% to −2.975% in SVDpp, UserKNN and ItemKNN. Hence, the preference distribution-aware provider fairness algorithm is able to mitigate disparities.

6.2.2. Detailed analysis on the courses dataset

On the other hand, in the Courses dataset (Table 8), we have observed that BPRMF made changes in the recommendation lists in all the provider groups. Although disparities are not completely reduced in all the recommendation algorithms, they are almost close to 0, denoting that they nearly reach the input data distribution at the provider level, with the greatest differences being in the NA items. As shown in Table 4, in the case of the Course dataset, the disparity in the original models range in BPRMF from 13.213% to −6.834%, in SVDpp from 1.523% to −2.740%, in UserKNN from 0.666% to −1.229% and ItemKNN from 0.504% to −1.241%. Once applied our DAP-fair algorithm (see Table 8), the disparities range from 0.581% to −0.769% in BPRMF, from 0.559% to −0.830% in SVDpp, and from 0.560% to −0.832% UserKNN and ItemKNNN.

6.2.3. DAP-fair comparison to state-of-the-art models

Going more in depth to the behavior of the DAP-fair algorithm, Fig. 2 show the overall percentage of disparity returned by the different recommendation models when applying different mitigation

Table 8
Disparate visibility, $\Delta\mathcal{V}$, in the continent distribution of DAP-fair models with 100% loss tolerance in the Courses Dataset. Values are in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Bold numbers denote where the disparity has been mitigated.

Courses													
Items													
		BPRMF						SVDpp					
		AF	AS	EU	NA	OC	SA	AF	AS	EU	NA	OC	SA
U	AF	−0.014	−0.013	**0.001**	0.581	**0.003**	**0.000**	**0.000**	**0.000**	**0.000**	0.559	**0.000**	**0.000**
s	AS	−0.019	−0.013	**0.005**	0.203	**−0.005**	**−0.001**	**0.000**	**−0.001**	**0.001**	0.167	**0.001**	**0.000**
e	EU	−0.014	−0.012	**0.003**	−0.287	**−0.001**	**−0.002**	**0.000**	**−0.001**	**0.000**	−0.314	**0.000**	**0.000**
r	NA	−0.031	−0.023	**0.006**	−0.769	**−0.004**	**−0.001**	**0.000**	**−0.001**	**0.000**	−0.830	**0.000**	**−0.001**
s	OC	**−0.002**	**−0.002**	**0.000**	−0.106	**−0.001**	**0.000**	**0.000**	**−0.001**	**0.000**	−0.109	**−0.001**	**0.000**
	SA	**−0.004**	**−0.003**	−0.017	0.557	**−0.001**	−0.014	**0.000**	**0.000**	**0.000**	0.530	**0.000**	**0.000**
		UserKNN						ItemKNN					
U	AF	**0.000**	**0.000**	**0.000**	0.560	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	0.560	**0.000**	**0.000**
s	AS	**0.000**	**0.000**	**0.000**	0.167	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	0.167	**0.000**	**0.000**
e	EU	**0.000**	**0.000**	**0.000**	−0.315	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	−0.315	**0.000**	**0.000**
r	NA	**0.000**	**0.000**	**0.000**	−0.832	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	−0.832	**0.000**	**0.000**
s	OC	**0.000**	**0.000**	**0.000**	−0.111	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	−0.111	**0.000**	**0.000**
	SA	**0.000**	**0.000**	**0.000**	0.530	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	0.530	**0.000**	**0.000**

Table 9
Percentage of disparity between the models with different mitigation techniques (original corresponds to no mitigation, the baseline Gómez, Boratto et al., 2022, and the DAP-fair algorithm) in the Books dataset, taking into account only the provider.

Books					
		EU	NA	OC	SA
BPRMF	Original	6.197	−6.174	−0.022	−0.001
	Baseline	**−0.002**	0.025	−0.022	−0.001
	DAP-fair	**−0.002**	**0.003**	**−0.001**	**0.000**
SVDpp	Original	2.907	−2.890	−0.016	−0.001
	Baseline	**−0.002**	**0.002**	**0.000**	**0.000**
	DAP-fair	−0.003	0.003	**0.000**	**0.000**
UserKNN	Original	−1.772	1.781	−0.010	0.001
	Baseline	**−0.002**	**0.002**	**0.000**	**0.000**
	DAP-fair	−0.003	0.003	**0.000**	**0.000**
ItemKNN	Original	−3.771	3.732	0.012	0.026
	Baseline	**−0.002**	**0.002**	**0.000**	**0.000**
	DAP-fair	−0.003	0.003	**0.000**	**0.000**

strategies (i.e., no mitigation in the original version, the baseline approach Gómez, Boratto et al., 2022, and our DAP-fair algorithm), in the Books and Courses datasets, respectively. Values within the figure correspond to the sum of the absolute values of the disparities between the models and the input data set. For example, Fig. 2 shows that the overall percentage of disparity is 12.497% in the BPRMF for the original model, and this percentage of disparity has been reduced to 6.284% with the DAP-fair algorithm. In the Books dataset, DAP-fair obtains an overall disparity that is around 6%. On the other hand, Fig. 2 shows that, in the case of the Courses dataset, the disparities have fallen to 2% in all recommenders. In both datasets, we have noticed that a provider fairness mitigation algorithm reduces disparities; however, the DAP-fair is more effective since it considers a broader set of perspectives (i.e., the distribution of the user preferences and the geographic provenance of providers) for distributing items. We have shown that mitigating by considering only the provider fairness perspective is insufficient to reduce most disparities.

Considering the results obtained in Fig. 2, we clearly observe that the DAP-Fair algorithm reduces quite a lot of the disparities taking into account both users' preferences and geographic location of providers; however, this does not seem to be the case with the algorithm used as Baseline. In particular, in the Books dataset and considering the SVDpp, UserKNN, and ItemKNN models, the results are higher, and the same happens with SVDpp with the dataset of Courses.

Observation 3. *By providing a broader set of perspectives, i.e., not only the provider fairness but also the user preferences, disparities are considerably reduced. Our results clearly show that DAP-fair is necessary and effective when trying to generate geographically-aware recommendation lists according to the preferences of users from different continents.*

We have also included the analysis of all algorithms considering only the providers without considering the provenance of the users. As it can be seen in Tables 9 and 10, the Baseline model works just as well as DAP-Fair, and this happens since the Baseline was designed only to consider the provenance of the providers; hence, there is a need to implement our DAP-Fair model to handle a broader set of perspectives.

Observation 4. *Recommendation models that only adjust provider provenance cannot capture all disparities since they do not consider users' preferences. As a result, disparities are decreased but still emerge.*

6.3. Impact of loss tolerance

Our algorithm includes a hyper-parameter to manage the maximum loss allowed by the algorithm when mitigating disparities. In our experiments, we have tested different loss tolerance values, namely 25%, 50%, 75%, and 100%. Note that, when we increase the loss tolerance, we achieve recommendation lists with a geographic distribution more similar to the input dataset.

Fig. 3 depicts the differences in the continent distribution of the original model and DAP-fair algorithm for different values of loss tolerance (from 25% to 100%) in the Books and Courses datasets. It is important to highlight that Fig. 3 provides a summary of our results. Detailed outcomes based on user preferences and provider distributions are presented in Tables 12 and 13 in Appendix. The results in Fig. 3 reveal that, in both datasets, BPRMF with a 25% value of loss tolerance reduces, but is not able to entirely adjust, the disparities. This happens because the recommendation lists do not contain a sufficient variety. Note that after 50% tolerance, no changes are perceived and the results are in line with the remaining recommendation models. In the case

Table 10
Percentage of disparity between the models with different mitigation techniques (original corresponds to no mitigation, the baseline Gómez, Boratto et al., 2022, and the DAP-fair algorithm) in the courses dataset, taking into account only the provider.

Courses							
		AF	AS	EU	NA	OC	SA
BPRMF	Original	2.635	14.795	6.456	−13.638	−6.124	−4.125
	Baseline	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	DAP-fair	**0.000**	−0.003	0.003	0.003	−0.002	−0.002
SVDpp	Original	−4.110	2.214	0.765	−2.220	3.061	0.290
	Baseline	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	DAP-fair	**0.000**	−0.004	0.002	0.003	−0.001	−0.001
UserKNN	Original	−0.003	−0.550	0.892	−0.869	−0.204	0.734
	Baseline	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	DAP-fair	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
ItemKNN	Original	−0.370	0.537	0.167	−1.193	0.248	0.612
	Baseline	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
	DAP-fair	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**

Table 11
NDCG for every recommendation model before and after mitigation in the books and courses datasets. Original corresponds to the original model, Baseline (Gómez, Boratto et al., 2022), and DAP-fair mitigation algorithm with different loss tolerance values.

NDCG						
Books						
	Original	Baseline	DAP-fair (loss tolerance)			
			25%	50%	75%	100%
BPRMF	0.0231	0.0232	**0.0233**	**0.0233**	**0.0233**	**0.0233**
SVDpp	0.0142	**0.0153**	0.0140	0.0140	0.0140	0.0140
UserKNN	**0.0323**	0.0320	0.0320	0.0319	0.0314	0.0313
ItemKNN	**0.0332**	0.0328	0.0331	0.0325	0.0323	0.0319
Courses						
BPRMF	0.0350	0.0337	0.0362	**0.0373**	**0.0373**	**0.0373**
SVDpp	**0.0848**	0.0838	0.0722	0.0708	0.0706	0.0706
UserKNN	**0.1924**	0.1910	0.1870	0.1760	0.1680	0.1667
ItemKNN	**0.2046**	0.2036	0.1979	0.1910	0.1894	0.1890

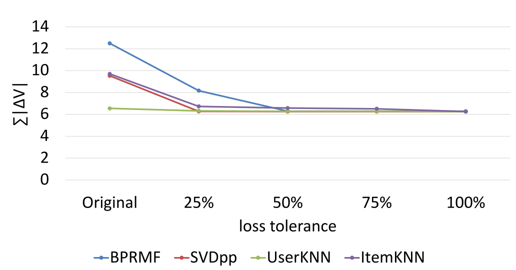

(a) Books

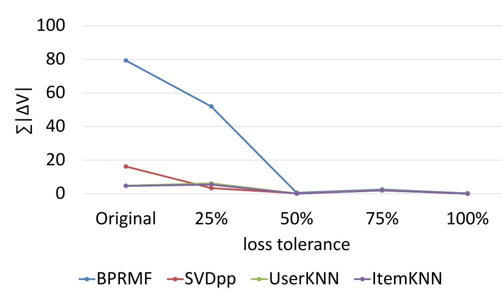

(b) Courses

Fig. 3. Overall percentage of disparity between the original model and DAP-fair algorithm with different loss tolerance values in the Books and Courses datasets.

of SVDpp, after a 50% loss tolerance, there are no changes either. In UserKNN and ItemKNN, in the Courses dataset, we achieve almost the desired proportions at the bucket and provider level by tolerating 100% losses.

Observation 5. *DAP-fair effectively reduces disparities while permitting a small margin of loss tolerance and achieves favorable results across both small and large percentages of loss tolerance.*

6.4. Impact on the recommendation quality

In the preceding sections, we examined how DAP-fair effectively mitigates disparities. However, re-ranking the recommendation lists through this intervention may influence the overall quality of the recommendations. In this section, we analyze the impact of our DAP-fair algorithm on the recommendation quality. We report our results in Table 11 for the Books and Courses datasets. We compare the NDCG achieved by the recommendation models without a mitigation (i.e., original in the table) and that with two mitigation approaches (i.e., the baseline and our proposal, DAP-fair, with several loss tolerance ranges).

In Section 4.2 we have shown that BPRMF and SVDpp are the two recommendation algorithms with the largest geographical disparities in both datasets. In fact, these algorithms achieve the worst NDCG in Table 11 for the original setting. On the other hand, the recommendation models that perform the best at the level of geographical distribution have been UserKNN and ItemKNN and, specifically, these algorithms are the ones that obtained the best NDCG.

We compare the results of our distribution-aware provider fairness (DAP-fair) algorithm with that proposed in Gómez, Boratto et al. (2022) (denoted as Baseline in Table 11). As described in Section 6.3, the baseline is able to capture provider disparities and to mitigate them, but not the users' preferences whereas DAP-fair focuses on both aspects minimizing in this way the overall disparities. It is interesting to observe that, in both scenarios, the loss in NDCG is negligible for the baseline algorithm and the loss in NDCG is not significant for the DAP-fair algorithm. Note that this reduction of NDCG in DAP-fair is expected because when the constraints are relaxed (loss tolerance is increased), the quality of the recommendation lists is reduced. This is largely due to the fact that the more changes are made to the lists, the more the quality is affected. Interestingly, however, these reductions are not significant in NDCG (the highest is only −0.0013 in Books and −0.0257 in Courses – see Table 11–), but DAP-fair achieved a significant better geographical distribution, so that, it is the most effective in reducing the disparities between the recommendation lists and the input data, as denoted in Section 6.3.

These results clearly show that a mitigation considering both the provider and the user perspective can achieve greater heterogeneity in the geographic distribution, at least to the extent that it reflects past users' preferences, without significantly affecting the quality and satisfaction of the recommendations. In addition, our DAP-fair proposal focuses on setting constraints that allow us to control how much loss of relevance we wish to allow in exchange for tighter, more geographically aware lists of users and providers.

Observation 6. *When providing a re-ranking based on satisfying users' preferences while adjusting provider provenance, the disparity is mitigated, and the effectiveness remains stable.*

7. Discussion

In this section, we outline the contributions of our proposal and discuss its practical implications.

7.1. Contributions to literature

This paper makes significant contributions by addressing a critical gap in the literature on provider fairness in recommender systems. While existing approaches focus on ensuring adequate visibility for providers, they often neglect to consider *to whom* the recommendations are being made. We introduce the concept of **preference distribution-aware provider fairness**, which ensures that recommendations align with users' original preferences while promoting visibility for minority providers. To the best of our knowledge, this is the first work to address unfairness in content providers and users across different continents.

Our proposed algorithm tackles this challenge by ensuring that minority providers reach a diverse audience, allowing them to expand their business opportunities while maintaining user preference alignment. This dual consideration of provider visibility and user preferences enables more equitable recommendation outcomes.

Furthermore, our study reveals key observations about the limitations of current recommendation models, as shown in Section 4.2. These models often exacerbate imbalances by over-recommending items from majority groups, failing to respect the original distribution of user preferences. Our results emphasize the need for algorithmic interventions to mitigate these disparities. In Section 5.1, we also demonstrate that simply focusing on provider fairness is insufficient in the presence of geographic imbalances; both user preferences and provider provenance must be incorporated to achieve true fairness.

Finally, in Section 6, we show that by re-ranking recommendations based on both user preferences and provider provenance, our approach significantly reduces disparities while maintaining the effectiveness of the recommendations. The trade-off between fairness and relevance of recommendations can be maintained within acceptable levels of accuracy loss, as shown in Fig. 3. Even with 100% loss tolerances, as Table 11 shows the impact on overall recommendation quality was minimal, suggesting that it is possible to ensure fairness without compromising system utility for end users. The introduction of DAP-fair proves essential for generating geographically aware recommendation lists, offering a novel solution to the issue of unfairness in recommender systems.

7.2. Implications for practice

The purpose of our research is to develop a fairness algorithm, DAP-Fair, which considers both provider fairness and to whom the recommendations are suggested. Although we have evaluated the algorithm considering provider provenance, it can be easily adapted to other sensitive attributes. We have demonstrated its effectiveness using well-known real-world datasets, but it can be successfully applied across various domains. Additionally, since the proposal is a post-processing re-ranking algorithm, it is compatible with any type of recommender system. In our experiments, we used well-known state-of-the-art Collaborative Filtering algorithms (see Section 3.3), ranging from *memory-based* to *matrix factorization* approaches. The DAP-fair algorithm simultaneously considers both the provider and user perspectives, which may increase processing time in comparison to algorithms that uniquely focus on provider fairness perspective. However, the issue of decreasing the processing time can be mitigated in two ways: by precomputing the representation of each group in the training set, and by parallelizing the computation of the buckets to calculate the expected loss.

8. Conclusions and future work

While provider fairness is receiving important attention in the recommender systems literature, the problem of how fair recommendations are distributed across consumers remains open.

To address this gap, we have first assessed that the problem occurs in the real world, by considering two datasets. We have also analyzed whether a provider fairness approach can consider the users' preferences and the results have shown that a dual perspective should be considered in an algorithm to address the issue. In this paper, we proposed DAP-fair, a re-ranking approach capable of performing the dual perspective of distributing the recommendations in a fair way for the providers, and of considering the preferences of the users. Specifically, we considered user and provider groups based on their continent of origin and presented a re-ranking approach that shapes recommendation lists that, thanks to the use of buckets (i.e., containing the recommendations associated with a user continent and an item continent pair), is able to (i) respect the distribution of the users' preferences in terms of providers, (ii) recommend each provider group according to their representation in the data, and (iii) take items' relevance into account. Our results on two real-world datasets show that our approach can achieve our goal, thus producing *preference distribution-aware provider fairness*.

As future work, we are planning to consider preference-distribution awareness at the individual user level, thus creating calibrated and provider-fair recommendations. Moreover, we will evaluate the DAP-fair algorithm on different types of recommender systems. Finally, we will consider additional recommendation domains, to assess these issues in different application areas and explore additional demographic factors.

CRediT authorship contribution statement

David Contreras: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Data curation, Conceptualization. **Ludovico Boratto:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix

See Algorithms 1 to 3 and Tables 12 and 13.

```
Input: recList: ranked list (records contain user, item, rating, position, usercont, itemcont), which arrives sorted by user and rating and contains topn recommendations to the user.
       trainList: list with the training set (records contain user, item, rating, usercont, itemcont).
       topn: top n recommendations, we set up n = 1000.
       topk: top k recommendations, we set up k = 10.
Output: reRankedList: ranked list with distribution-aware recommendations.
define DAP-fair (recList, trainList, topn, topk)
begin
    // Step 1. Create a bucket list considering each user-continent and item-continent pair, which will store the predicted items (recList)
    recBucketRep ← computeRepresentationPerContinent(recList, trainList, topk); // Compute R considering the ratings in trainList, see Eq. (1)
    lossBucketc ← calculateLoss(recList, topn, topk); // Calculate the loss for each user
    joinBucket ← recList + recBucketRep + lossBucket; // Create a bucket list considering each user-continent and item-continent pair
    joinBucket ← sort(joinBucket) ; // sort by percentages and relevance
    // Step 2. Perform selection of items
    tolerance ← 0.75 ; // hyper-parameter with the maximum tolerance, e.g., 75%
    userCounts ← ∅ ; // list with #recommendations each user received
    bucketCounts ← ∅ ; // list with #items recommended at each bucket
    joinBucket ← phasesDAP(joinBucket, recBucketRep, topk, tolerance, 1, userCounts, bucketCounts) ; // Phase 1
    joinBucket ← phasesDAP(joinBucket, recBucketRep, topk, tolerance, 2, userCounts, bucketCounts) ; // Phase 2
    joinBucket ← phasesDAP(joinBucket, recBucketRep, topk, tolerance, 3, userCounts, bucketCounts) ; // Phase 3
    reRankedList ← chooseSelectedItems(joinBucket) ; // add items selected to the re-ranked list, we only need user, item and rating
    reRankedList ← sort(reRankedList) ; // sort by user and rating
    return reRankedList ; // list adjusted by distribution-aware preferences
end
```

Algorithm 1: Pseudocode of DAP-fair algorithm

```
define calculateLoss (recList, topn, topk)
begin
    lossBucket ← recList; // define initial set of buckets
    recList ← minmax_scale(recList.rating); // normalize ratings with a min max scaler
    i ← 0; // first position
    while rec ∈ recList do // for each record in recList
        if i ≤ topk then
            lossBucket[rec].loss ← 0; // items belonging to topk has loss 0
        else
            if i == topk then
                baseLoss ← recList[rec].rating; // base rating is last item in topk
            else
                lossBucket[rec].loss ← (recList[rec].rating − baseLoss)/baseLoss;
            end
        end
        lossBucket[rec].bucketId ← recList[rec].usercont + "−" + recList[rec].itemcont;
        if i == topn then
            i ← 0;
        else
            i ← i + 1;
        end
    end
    return lossBucket
end
define computeRepresentationPerContinent (recList, trainList, topk)
begin
    recBucketRep ← ∅; // define an empty set of buckets
    continents ← extractContinents(trainList); // obtain list of continents
    while cn1 ∈ continents do // for each user continent
        while cn2 ∈ continents do // for each item continent
            countTrain ← quantifyElements(trainList, cn1, cn2) ; // number of records in trainList with usercont = cn1 and itemcont = cn2
              percentageTrain ← countTrain/trainList.size(); // compute percentage countRecs ← quantifyElements(recList, cn1, cn2); // number of records in recList
              with usercont = cn1 and itemcont = cn2 percentageRecs ← countRecs/recList.size(); // % interval [0..1] res.bucketId ← cn1 + "−" + cn2; // save the
              bucket, e.g., Europe-Asia res.percentageTrain ← percentage; // save the percentage res.percentageRecs ← percentage; // save the percentage
              res.expectedRecords ← round(recList.getNumUsers() × topk × percentageTrain); // number of expected records in train recBucketRep.add(res); // store
              the bucket
        end
    end
    return recBucketRep
end
```

Algorithm 2: Pre-processing support methods to create a bucket

```
define phasesDAP(joinBucket, recBucketRep, topk, tolerance, phaseDAP, userCounts, bucketCounts)
begin
    expectedRecords ← getExpectedRecords(recBucketRep) ; // returns a list with the number of items expected on each bucket
    while ele ∈ recBucket do // for each element
        updateRecord ← 0;
        while rec ∈ joinBucket do // for each record
            // conditions 1, 2 and 3
            if |rec.loss| ≤ tolerance and expectedRecords[rec.bucket] ≥ bucketCounts[rec.bucket] and topk ≥ userCounts[rec.user] then
                if phaseDAP == 1andrec.ItemCont == ele.itemContandrec.ItemCont == rec.user.Cont then
                    rec.phase ← 1 ; // selects element in phase 1
                    updateRecord ← 1;
                else
                    if phaseDAP == 2andrec.ItemCont == ele.itemCont then
                        rec.phase ← 2 ; // selects element in phase 2
                        updateRecord ← 1;
                    else
                        if phaseDAP == 3 then
                            rec.phase ← 3 ; // selects element in phase 3
                            updateRecord ← 1;
                        end
                    end
                end
            end
            if updateRecord == 1 then
                userCounts[rec.user] ← userCounts[rec.user] + 1 bucketCounts[rec.bucket] ← bucketCounts[rec.bucket] + 1 joinBucket.update(rec) ; // updates the element
            end
            if bucketCount[ele.bucketId] == ele.expectedRecords then
                break;
            end
        end
    end
    return joinBucket
end
```

Algorithm 3: Support method for the DAP-fair algorithm

Table 12
Distribution DAP-fair Models with different loss tolerance in the Books Dataset. Representations in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Groups appear in alphabetical order by the name of the continent.

Books													
Items													
BPRMF													
		25%				50%				75%			
		EU	NA	OC	SA	EU	NA	OC	SA	EU	NA	OC	SA
	AF	**0.004**	0.032	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**
U	AN	**0.000**	**−0.002**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**
s	AS	**0.000**	−0.150	**−0.001**	**0.000**	**0.000**	−0.150	**−0.001**	**0.000**	**0.000**	−0.150	**−0.001**	**0.000**
e	EU	−0.909	3.857	**−0.004**	**−0.001**	**0.000**	2.948	**−0.004**	**−0.001**	**0.000**	2.948	**−0.004**	**−0.001**
r	NA	**0.000**	−2.967	**−0.009**	**0.000**	**0.000**	−2.967	**−0.009**	**0.000**	**0.000**	−2.967	**−0.009**	**0.000**
s	OC	−0.034	0.126	**−0.008**	**0.000**	**0.000**	0.092	**−0.008**	**0.000**	**0.000**	0.092	**−0.008**	**0.000**
	SA	0.013	0.053	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**
SVDpp													
	AF	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**
U	AN	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**
s	AS	**0.000**	−0.151	**0.000**	**0.000**	**0.000**	−0.151	**0.000**	**0.000**	**0.000**	−0.151	**0.000**	**0.000**
e	EU	**0.001**	2.938	**0.005**	−0.001	**0.000**	2.943	**0.000**	**0.000**	**0.000**	2.943	**0.000**	**0.000**
r	NA	**0.000**	−2.975	**0.000**	**0.000**	**0.000**	−2.975	**0.000**	**0.000**	**0.000**	−2.975	**0.000**	**0.000**
s	OC	**0.000**	0.090	**−0.006**	**0.000**	**0.000**	0.084	**0.000**	**0.000**	**0.000**	0.084	**0.000**	**0.000**
	SA	**0.000**	0.066	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**
UserKNN													
	AF	**0.000**	0.035	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**
U	AN	**−0.001**	**−0.002**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**
s	AS	**0.002**	−0.153	**0.000**	**0.000**	**0.001**	−0.152	**0.000**	**0.000**	**0.000**	−0.151	**0.000**	**0.000**
e	EU	0.702	2.239	**0.001**	**0.000**	0.358	2.584	**0.001**	**0.000**	0.177	2.766	**0.000**	**0.000**
r	NA	0.025	−3.000	**0.000**	**0.000**	**0.007**	−2.982	**0.000**	**0.000**	**0.004**	−2.980	**0.000**	**0.000**
s	OC	**0.004**	0.086	**−0.006**	**0.000**	**0.003**	0.083	**−0.002**	**0.000**	**0.001**	0.083	**0.000**	**0.000**
	SA	**0.003**	0.063	**0.001**	**0.000**	**0.000**	0.065	**0.001**	**0.000**	**0.000**	0.066	**0.000**	**0.000**
ItemKNN													
	AF	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**	**0.000**	0.036	**0.000**	**0.000**
U	AN	**−0.001**	**−0.002**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**	**0.000**	**−0.003**	**0.000**	**0.000**
s	AS	**0.003**	−0.154	**0.000**	**0.000**	**0.001**	−0.152	**0.000**	**0.000**	**0.001**	−0.152	**0.000**	**0.000**
e	EU	1.380	1.559	**0.000**	**0.004**	1.112	1.830	**0.000**	**0.001**	0.972	1.971	**0.000**	**0.001**
r	NA	0.229	−3.208	**0.003**	**0.000**	0.153	−3.132	**0.003**	**0.000**	0.121	−3.099	**0.003**	**0.000**
s	OC	0.019	0.065	**0.000**	**0.000**	0.012	0.072	**0.000**	**0.000**	**0.009**	0.075	**0.000**	**0.000**
	SA	**0.001**	0.065	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**	**0.000**	0.066	**0.000**	**0.000**

Table 13
Distribution DAP-fair Models with different loss tolerance in the Courses Dataset. Representations in percentage of each group (AF: Africa, AN: Antarctica, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America). Groups appear in alphabetical order by the name of the continent.

Courses																			
Items																			
BPRMF																			
		25%						50%						75%					
		AF	AS	EU	NA	OC	SA	AF	AS	EU	NA	OC	SA	AF	AS	EU	NA	OC	SA
U	AF	−0.719	−0.542	**−0.005**	1.899	0.012	−0.050	−0.038	−0.012	**−0.004**	0.597	**−0.004**	**−0.001**	−0.014	−0.013	**−0.002**	0.585	**−0.003**	**−0.001**
s	AS	−1.142	−1.519	**−0.003**	2.927	**−0.009**	−0.043	−0.042	−0.010	0.012	0.223	**−0.003**	**−0.001**	−0.019	−0.013	**−0.006**	0.207	**−0.005**	**−0.001**
e	EU	−0.789	−1.026	0.011	1.710	**−0.009**	−0.192	−0.042	−0.010	0.010	−0.270	**−0.001**	**−0.003**	−0.015	−0.011	**−0.004**	−0.283	**−0.001**	**−0.003**
r	NA	−1.635	−2.699	**−0.009**	3.899	0.018	−0.269	−0.078	−0.018	0.018	−0.724	**−0.000**	**−0.002**	−0.031	−0.022	**−0.009**	−0.759	**−0.003**	**−0.002**
s	OC	−0.122	−0.259	**−0.001**	0.273	**−0.002**	**−0.007**	**−0.007**	**−0.007**	**−0.001**	−0.103	**−0.000**	**−0.000**	**−0.001**	**−0.002**	**−0.001**	−0.105	**−0.001**	**−0.000**
	SA	−0.310	−0.161	−0.016	4.348	**−0.002**	−3.601	−0.042	**−0.003**	−0.016	0.597	**−0.000**	−0.041	−0.037	**−0.003**	−0.017	0.592	**−0.001**	−0.046
SVDpp																			
U	AF	−0.038	**−0.000**	0.018	0.565	0.012	**−0.002**	−0.013	−0.012	**−0.002**	0.574	**−0.004**	**−0.001**	**−0.000**	**−0.000**	**−0.000**	0.560	**−0.000**	**−0.000**
s	AS	−0.042	**−0.007**	0.076	0.127	0.012	**−0.004**	−0.013	−0.010	0.015	0.169	**−0.008**	**−0.001**	**−0.000**	**−0.001**	**−0.001**	0.167	**−0.001**	**−0.001**
e	EU	−0.042	−0.010	**−0.009**	−0.282	**−0.001**	0.010	−0.012	−0.012	**−0.002**	−0.294	**−0.001**	**−0.001**	**−0.000**	**−0.000**	**−0.000**	−0.314	**−0.000**	**−0.000**
r	NA	−0.078	−0.019	0.017	−0.732	**−0.000**	**−0.000**	−0.026	−0.025	**−0.004**	−0.786	**−0.001**	**−0.001**	**−0.001**	**−0.001**	**−0.000**	−0.829	**−0.000**	**−0.001**
s	OC	**−0.007**	**−0.002**	**−0.001**	−0.103	**−0.001**	**−0.000**	**−0.001**	**−0.002**	**−0.000**	−0.107	**−0.001**	**−0.000**	**−0.000**	**−0.001**	**−0.000**	−0.109	**−0.001**	**−0.000**
	SA	−0.042	**−0.001**	−0.011	0.585	**−0.005**	−0.042	**−0.002**	**−0.002**	**−0.003**	0.522	**−0.003**	0.012	**−0.000**	**−0.000**	**−0.001**	0.530	**−0.000**	**−0.000**
UserKNN																			
U	AF	0.021	0.025	0.046	0.451	**−0.001**	**−0.003**	**−0.005**	**−0.004**	**−0.005**	0.511	**−0.002**	0.033	**−0.001**	**−0.001**	**−0.002**	0.490	**−0.001**	0.064
s	AS	0.049	0.044	0.151	−0.119	0.047	**−0.003**	**−0.004**	**−0.008**	0.027	0.123	**−0.002**	**−0.002**	**−0.001**	**−0.004**	0.023	0.108	**−0.001**	0.030
e	EU	0.015	0.033	0.161	−0.557	0.024	0.011	**−0.002**	**−0.006**	0.045	−0.373	**−0.003**	**−0.003**	**−0.001**	**−0.002**	0.016	−0.335	**−0.001**	**−0.000**
r	NA	0.049	0.080	0.230	−1.260	0.047	0.021	**−0.009**	0.011	0.019	−0.882	**−0.004**	**−0.006**	**−0.004**	**−0.006**	**−0.009**	−0.855	**−0.002**	**−0.004**
s	OC	**−0.003**	0.016	0.010	−0.151	0.011	**−0.001**	**−0.000**	**−0.006**	**−0.001**	−0.128	0.011	**−0.000**	**−0.000**	**−0.001**	**−0.001**	−0.117	**−0.004**	**−0.000**
	SA	0.012	**−0.005**	0.029	0.449	**−0.003**	0.038	**−0.001**	**−0.001**	**−0.002**	0.570	**−0.000**	−0.043	**−0.000**	**−0.001**	**−0.000**	0.627	**−0.000**	−0.098
ItemKNN																			
U	AF	0.015	0.032	0.031	0.440	0.038	**−0.003**	**−0.001**	**−0.004**	**−0.001**	0.525	**−0.002**	0.025	**−0.001**	**−0.003**	**−0.000**	0.553	**−0.001**	**−0.000**
s	AS	0.032	0.051	0.111	−0.119	0.096	**−0.003**	**−0.000**	**−0.001**	0.019	0.138	**−0.005**	**−0.006**	**−0.000**	**−0.000**	0.012	0.157	**−0.000**	**−0.000**
e	EU	0.012	0.037	0.105	−0.524	0.037	0.022	**−0.000**	**−0.001**	0.014	−0.337	**−0.001**	**−0.008**	**−0.000**	**−0.000**	**−0.009**	−0.322	**−0.000**	**−0.000**
r	NA	0.030	0.094	0.148	−1.215	0.080	0.027	**−0.001**	**−0.000**	**−0.001**	−0.845	**−0.004**	**−0.005**	**−0.000**	**−0.000**	**−0.001**	−0.834	**−0.000**	**−0.000**
s	OC	**−0.000**	0.011	**−0.007**	−0.147	0.017	**−0.001**	**−0.000**	**−0.001**	**−0.000**	−0.111	**−0.000**	**−0.000**	**−0.000**	**−0.000**	**−0.000**	−0.111	**−0.000**	**−0.000**
	SA	**−0.005**	**−0.003**	0.011	0.505	**−0.007**	**−0.005**	**−0.001**	**−0.000**	**−0.000**	0.574	**−0.000**	−0.044	**−0.000**	**−0.000**	**−0.000**	0.530	**−0.000**	**−0.000**

References

Anelli, V. W., Bellogin, A., Ferrara, A., Malitesta, D., Merra, F. A., Pomo, C., et al. (2021). *Elliot: a comprehensive and rigorous framework for reproducible recommender systems evaluation* (pp. 2405–2414). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3404835.3463245.

Ashokan, A., & Haas, C. (2021). Fairness metrics and bias mitigation strategies for rating predictions. *Information Processing & Management*, *58*(5), Article 102646. http://dx.doi.org/10.1016/j.ipm.2021.102646.

Baeza-Yates, R. (2018). Bias on the web. *Communications of the ACM*, *61*(6), 54–61. http://dx.doi.org/10.1145/3209581.

Barra, S., Marras, M., & Fenu, G. (2018). 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.), *Lecture notes in computer science*: *vol. 11058*, *Network and system security - 12th international conference, NSS 2018, Hong kong, China, August 27-29, 2018, proceedings* (pp. 212–220). USA: Springer, http://dx.doi.org/10.1007/978-3-030-02744-5_16.

Bauer, K., Hinz, O., van der Aalst, W. M. P., & Weinhardt, C. (2021). Expl(ai)n it to me - explainable AI and information systems research. *Business & Information Systems Engineering*, *63*(2), 79–82. http://dx.doi.org/10.1007/s12599-021-00683-2.

Bellogín, A., Castells, P., & Cantador, I. (2017). Statistical biases in information retrieval metrics for recommender systems. *Information Retrieval Research Journal*, *20*(6), 606–634. http://dx.doi.org/10.1007/s10791-017-9312-z.

Biega, A. J., Gummadi, K. P., & Weikum, G. (2018). Equity of attention: Amortizing individual fairness in rankings. In *The 41st international ACM SIGIR conference on research & development in information retrieval* (pp. 405–414). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3209978.3210063.

Boratto, L., Fenu, G., Marras, M., & Medda, G. (2022). Consumer fairness in recommender systems: Contextualizing definitions and mitigations. In M. Hagen, S. Verberne, C. Macdonald, C. Seifert, K. Balog, K. Nørvåg, & V. Setty (Eds.), *Advances in information retrieval* (pp. 552–566). Cham: Springer International Publishing, http://dx.doi.org/10.1007/978-3-030-99736-6_37.

Burke, R., Mattei, N., Grozin, V., Voida, A., & Sonboli, N. (2022). Multi-agent social choice for dynamic fairness-aware recommendation. In *UMAP '22: 30th ACM conference on user modeling, adaptation and personalization, Barcelona, Spain, July 4-7, 2022, adjunct proceedings* (pp. 234–244). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3511047.3538032.

Deldjoo, Y., Jannach, D., Bellogín, A., Difonzo, A., & Zanzonelli, D. (2022). A survey of research on fair recommender systems. http://dx.doi.org/10.48550/arXiv.2205.11127, CoRR abs/2205.11127. arXiv:2205.11127.

Dessì, D., Fenu, G., Marras, M., & Reforgiato Recupero, D. (2017). Leveraging cognitive computing for multi-class classification of e-learning videos. In *Lecture notes in computer science*: *vol. 10577*, *The semantic web: ESWC 2017 satellite events, revised selected papers* (pp. 21–25). Springer, http://dx.doi.org/10.1007/978-3-319-70407-4_5.

Dessì, D., Fenu, G., Marras, M., & Reforgiato Recupero, D. (2018). COCO: semantic-enriched collection of online courses at scale with experimental use cases. In *Advances in intelligent systems and computing*: *vol. 746*, *Trends and advances in information systems and technologies - volume 2* (pp. 1386–1396). Cham: Springer, http://dx.doi.org/10.1007/978-3-319-77712-2_133.

Diaz, F., Mitra, B., Ekstrand, M. D., Biega, A. J., & Carterette, B. (2020). Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM international conference on information & knowledge management* (pp. 275–284). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3340531.3411962.

Dwivedi, Y. K., Ismagilova, E., Hughes, D. L., Carlson, J., Filieri, R., Jacobson, J., et al. (2021). Setting the future of digital and social media marketing research: Perspectives and research propositions. *International Journal of Information Management*, *59*, Article 102168. http://dx.doi.org/10.1016/j.ijinfomgt.2020.102168.

Ekstrand, M. D., Das, A., Burke, R., & Diaz, F. (2022a). *Fairness in recommender systems* (pp. 679–707). New York, NY: Springer US, http://dx.doi.org/10.1007/978-1-0716-2197-4_18.

Ekstrand, M. D., Das, A., Burke, R., & Diaz, F. (2022b). Fairness and discrimination in information access systems. CoRR abs/2105.05779. arXiv:2105.05779.

Ekstrand, M. D., & Kluver, D. (2021). Exploring author gender in book rating and recommendation. *User Modeling and User-Adapted Interaction*, *31*(3), 377–420. http://dx.doi.org/10.1007/s11257-020-09284-2.

Ekstrand, M. D., Tian, M., Kazi, M. R. I., Mehrpouyan, H., & Kluver, D. (2018). Exploring author gender in book rating and recommendation. In *Proceedings of the 12th ACM conference on recommender systems* (pp. 242–250). ACM, http://dx.doi.org/10.1145/3240323.3240373.

Fabbri, F., Bonchi, F., Boratto, L., & Castillo, C. (2020). 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* (pp. 165–175). USA: AAAI Press, http://dx.doi.org/10.1609/icwsm.v14i1.7288.

Ge, M., Delgado-Battenfeld, C., & Jannach, D. (2010). Beyond accuracy: Evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on recommender systems* (pp. 257–260). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/1864708.1864761.

Geyik, S. C., Ambler, S., & Kenthapadi, K. (2019). Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 2221–2231). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3292500.3330691.

Gharahighehi, A., Vens, C., & Pliakos, K. (2021). Fair multi-stakeholder news recommender system with hypergraph ranking. *Information Processing & Management*, *58*(5), Article 102663. http://dx.doi.org/10.1016/j.ipm.2021.102663.

Gómez, E., Boratto, L., & Salamó, M. (2021). Disparate impact in item recommendation: A case of geographic imbalance. In D. Hiemstra, M. Moens, J. Mothe, R. Perego, M. Potthast, & F. Sebastiani (Eds.), *Lecture notes in computer science*: *vol. 12656*, *Advances in information retrieval - 43rd European conference on IR research, ECIR 2021, virtual event, March 28 - April 1 2021, proceedings, part i* (pp. 190–206). Cham: Springer, http://dx.doi.org/10.1007/978-3-030-72113-8_13.

Gómez, E., Boratto, L., & Salamó, M. (2022). Provider fairness across continents in collaborative recommender systems. *Information Processing & Management*, *59*(1), Article 102719. http://dx.doi.org/10.1016/j.ipm.2021.102719.

Gómez, E., Shui Zhang, C., Boratto, L., Salamó, M., & Marras, M. (2021). The winner takes it all: Geographic imbalance and provider (un)fairness in educational recommender systems. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval* (pp. 1808–1812). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3404835.3463235.

Gómez, E., Zhang, C. S., Boratto, L., Salamó, M., & Ramos, G. (2022). Enabling cross-continent provider fairness in educational recommender systems. *Future Generation Computer Systems*, *127*, 435–447. http://dx.doi.org/10.1016/j.future.2021.08.025.

Green, B., & Hu, L. (2018). The myth in the methodology: Towards a recontextualization of fairness in machine learning. In *International conference on machine learning*. URL https://api.semanticscholar.org/CorpusID:49544563.

Hatif Abd Almajed, I., Nassreddine, G., & Younis, J. (2024). New recruitment approach based on logistic regression model. *Mesopotamian Journal of Computer Science*, *2024*, 15–27. http://dx.doi.org/10.58496/MJCSC/2024/002, URL https://mesopotamian.press/journals/index.php/cs/article/view/127.

Herlocker, J. L., Konstan, J. A., & Riedl, J. (2002). An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms. *Information Retrieval*, *5*(4), 287–310. http://dx.doi.org/10.1023/A:1020443909834.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1), 75–105. http://dx.doi.org/10.2307/25148625, URL http://www.jstor.org/stable/25148625.

Järvelin, K., & Kekäläinen, J. (2002). Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, *20*(4), 422–446. http://dx.doi.org/10.1145/582415.582418.

Kaminskas, M., & Bridge, D. (2016). Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems*, *7*(1), http://dx.doi.org/10.1145/2926720.

Kamishima, T., Akaho, S., Asoh, H., & Sakuma, J. (2018). Recommendation independence. In C. Wilson S. A. Friedler (Ed.), *Proceedings of machine learning research*: *vol. 81*, *Proceedings of the 1st conference on fairness, accountability and transparency* (pp. 187–201). PMLR, URL https://proceedings.mlr.press/v81/kamishima18a.html.

Karakolis, E., Kokkinakos, P., & Askounis, D. (2022). Provider fairness for diversity and coverage in multi-stakeholder recommender systems. *Applied Sciences*, *12*(10), http://dx.doi.org/10.3390/app12104984.

Ken Peffers, M. A. R., Tuunanen, Tuure, & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77. http://dx.doi.org/10.2753/MIS0742-1222240302, arXiv:https://doi.org/10.2753/MIS0742-1222240302.

Koren, Y. (2008). Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 426–434). New York, NY, USA: ACM, http://dx.doi.org/10.1145/1401890.1401944.

Liu, W., & Burke, R. (2018). Personalizing fairness-aware re-ranking. CoRR abs/1809.02921. arXiv:1809.02921. URL http://arxiv.org/abs/1809.02921.

Marras, M., Boratto, L., Ramos, G., & Fenu, G. (2021). Equality of learning opportunity via individual fairness in personalized recommendations. *International Journal of Artificial Intelligence in Education*, http://dx.doi.org/10.1007/s40593-021-00271-1.

Mehrotra, R., McInerney, J., Bouchard, H., Lalmas, M., & Diaz, F. (2018). 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* (pp. 2243–2251). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3269206.3272027.

Naghiaei, M., Rahmani, H. A., & Deldjoo, Y. (2022). Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. In E. Amigó, P. Castells, J. Gonzalo, B. Carterette, J. S. Culpepper, & G. Kazai (Eds.), *SIGIR '22: the 45th international ACM SIGIR conference on research and development in information retrieval, Madrid, Spain, July 11-15, 2022* (pp. 770–779). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3477495.3531959.

Patro, G. K., Biswas, A., Ganguly, N., Gummadi, K. P., & Chakraborty, A. (2020). Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the web conference 2020* (pp. 1194–1204). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3366423.3380196.

Patro, G. K., Porcaro, L., Mitchell, L., Zhang, Q., Zehlike, M., & Garg, N. (2022). Fair ranking: A critical review, challenges, and future directions. In *2022 ACM conference on fairness, accountability, and transparency* (pp. 1929–1942). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3531146.3533238.

Raj, A., & Ekstrand, M. D. (2022). Measuring fairness in ranked results: An analytical and empirical comparison. In E. Amigó, P. Castells, J. Gonzalo, B. Carterette ad J.S. Culpepper, & G. Kazai (Eds.), *SIGIR '22: the 45th international ACM SIGIR conference on research and development in information retrieval, Madrid, Spain, July 11-15, 2022* (pp. 726–736). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3477495.3532018.

Ramos, G., & Caleiro, C. (2020). A novel similarity measure for group recommender systems with optimal time complexity. In L. Boratto, S. Faralli, M. Marras, & G. Stilo (Eds.), *Communications in computer and information science*: *vol. 1245*, *Bias and social aspects in search and recommendation - first international workshop, BIAS 2020, Lisbon, Portugal, April 14 2020, proceedings* (pp. 95–109). USA: Springer, http://dx.doi.org/10.1007/978-3-030-52485-2_10.

Rendle, S., Freudenthaler, C., Gantner, Z., & Schmidt-Thieme, L. (2012). BPR: Bayesian personalized ranking from implicit feedback. abs/1205.2618. arXiv:1205.2618. URL http://arxiv.org/abs/1205.2618.

Sapiezynski, P., Zeng, W., Robertson, R. E., Mislove, A., & Wilson, C. (2019). Quantifying the impact of user attentionon fair group representation in ranked lists. In *Companion of the 2019 world wide web conference* (pp. 553–562). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3308560.3317595.

Sarwar, B. M., Karypis, G., Konstan, J. A., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the tenth international world wide web conference* (pp. 285–295). New York, NY, USA: ACM, http://dx.doi.org/10.1145/371920.372071.

Sonboli, N., & Burke, R. (2019). Localized fairness in recommender systems. In *Adjunct publication of the 27th conference on user modeling, adaptation and personalization* (pp. 295–300). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3314183.3323845.

Tahery, S., Aftabi, S. Z., & Farzi, S. (2021). A ga-based algorithm meets the fair ranking problem. *Information Processing & Management*, *58*(6), Article 102711. http://dx.doi.org/10.1016/j.ipm.2021.102711.

Tsao, W.-Y. (2013). The fitness of product information: Evidence from online recommendations. *International Journal of Information Management*, *33*(1), 1–9. http://dx.doi.org/10.1016/j.ijinfomgt.2012.04.003.

Wu, H., Ma, C., Mitra, B., Diaz, F., & Liu, X. (2022). A multi-objective optimization framework for multi-stakeholder fairness-aware recommendation. *ACM Transactions on Information Systems*, *41*(2), http://dx.doi.org/10.1145/3564285.

Wu, H., Mitra, B., Ma, C., Diaz, F., & Liu, X. (2022). Joint multisided exposure fairness for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval* (pp. 703–714). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3477495.3532007.

Yang, K., & Stoyanovich, J. (2017). Measuring fairness in ranked outputs. In *Proceedings of the 29th international conference on scientific and statistical database management* (pp. 22:1–22:6). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3085504.3085526.

Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M., & Baeza-Yates, R. (2017). Fa*ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on conference on information and knowledge management* (pp. 1569–1578). New York, NY, USA: ACM, http://dx.doi.org/10.1145/3132847.3132938.

Zehlike, M., & Castillo, C. (2020). Reducing disparate exposure in ranking: A learning to rank approach. In *WWW '20: the web conference 2020* (pp. 2849–2855). New York, NY, USA: ACM / IW3C2, http://dx.doi.org/10.1145/3366424.3380048.

Zehlike, M., Sühr, T., Baeza-Yates, R., Bonchi, F., Castillo, C., & Hajian, S. (2022). Fair top-k ranking with multiple protected groups. *Information Processing & Management*, *59*(1), Article 102707. http://dx.doi.org/10.1016/j.ipm.2021.102707.

Ziegler, C.-N., McNee, S. M., Konstan, J. A., & Lausen, G. (2005). Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on world wide web* (pp. 22–32). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/1060745.1060754.