MOReGIn: Multi-Objective Recommendation at the Global and Individual Levels

Elizabeth Gómez[1[0000−0003−2698−3984]], David Contreras[2[0000−0002−6906−485X]], Ludovico Boratto[3[0000−0002−6053−3015]], and Maria Salamó[1[0000−0003−1939−8963]]

[1] Facultat de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain
egomezye13@alumnes.ub.edu,maria.salamo@ub.edu
[2] Facultad de Ingeniería y arquitectura, Universidad Arturo Prat, Iquique, Chile
david.contreras@unap.cl
[3] Department of Mathematics and Computer Science, University of Cagliari, Cagliari, Italy ludovico.boratto@acm.org

Abstract. Multi-Objective Recommender Systems (MORSs) emerged as a paradigm to guarantee multiple (often conflicting) goals. Besides accuracy, a MORS can operate at the *global* level, where additional beyond-accuracy goals are met for the system as a whole, or at the *individual* level, meaning that the recommendations are tailored to the needs of each user. The state-of-the-art MORSs either operate at the global or individual level, without assuming the co-existence of the two perspectives. In this study, we show that when global and individual objectives co-exist, MORSs are not able to meet both types of goals. To overcome this issue, we present an approach that regulates the recommendation lists so as to guarantee both global and individual perspectives, while preserving its effectiveness. Specifically, as individual perspective, we tackle genre calibration and, as global perspective, provider fairness. We validate our approach on two real-world datasets, publicly released with this paper[4].

Keywords: Multi-Objective Recommendation · Calibration · Provider Fairness.

1 Introduction

Motivation. Since the goal of recommender systems is to provide relevant suggestions for the users, the main focus has been the effectiveness of the results [37]. Nevertheless, users might be interested in properties of the items besides their effectiveness, and there are other stakeholders who can benefit from how recommendations are produced (e.g., content providers). Hence, beyond-accuracy perspectives are central to the generation and evaluation of recommendations.

Multi-Objective Recommender Systems (MORSs) support the provision of perspectives that go beyond item relevance, such as, e.g., diversity, calibration, and fairness [49]. The optimization for these objectives can happen at the *global* (*aggregate*) level, thus ensuring that the system as a whole can guarantee certain properties (e.g., all providers receive a certain exposure in the recommendation lists). In alternative, a MORS can operate at the *individual* (*local*) level, and

shape results that are consider the prominence of individual users towards the different goals (e.g., each user can receive a different level of diversity or the recommended genres can be calibrated to the preferences in the training set) [21].

When analyzing the current literature, a MORS either operates at the global level [15,26,28,32,44] or at the local level [34,8,7].

Open issues. There might be scenarios in which both global and individual objectives co-exist. Indeed, a platform might decide that, as a whole, the recommendations should offer certain properties (e.g., be fair to providers of different demographic groups, or enable a certain level of novelty). Moreover, specific goals might be set for the individual users (e.g., the calibration of the genres or the diversity of the recommended items might need to follow what is observed in the training set of each user). As we show in Section 6, when a MORS tackles *only* global or individual perspectives, the other perspective trivially remains under-considered and cannot be guaranteed by the system.

Our contributions. To overcome the aforementioned challenges, in this paper, we present a MORS that produces recommendations with both global and individual objectives. As a use case, we consider, as a global objective, *provider fairness* and, as an individual one, *calibrated recommendations*. This aligns our study with the rest of the MORS literature, where two beyond-accuracy objectives are considered. For the sake of clarity, we will talk about *provider-fair and calibrated recommendations* but, as we discuss in Section 3.2, **our approach can be generalized to any global or individual objective**.

Besides accounting for beyond-accuracy perspectives involving both global and individual objectives, the problem of providing provider-fair and calibrated recommendations becomes interesting also from a practical point of view. As we will show in Section 4, users tend to rate items of certain genres and that are produced in certain geographic areas, suggesting that we can account for both perspectives at the same time when generating the recommendations. Hence, at the technical level, we would need a unique solution that (i) produces effective results for the users, (ii) can provide fairness for providers belonging to different groups at the *global* level, i.e., by distributing, over the entire user base, the recommendation of items belonging to different provider groups in equitable ways, and (iii) can calibrate the recommendation lists of each *individual* user.

Our approach involves a post-processing strategy. To enable a form of provider fairness that can consider demographic groups that are not necessarily characterized by a binary group (e.g., males and females), we consider, as a sensitive attribute, the geographic provenance of the providers and have the different continents as the granularity with which we split the groups; this is aligned with recent literature on provider fairness [17,18]. As in classic calibrated fashion, we distribute the recommendations according to the item genre. Based on this characterization of the data, we present an approach that makes use of buckets to associate the continents in which the items are produced and the genre of the items. We use these buckets to post-process the recommendation lists (we will later discuss that this is the best way to regulate both aggregate- and individual-level properties) and regulate how the recommendations are distributed across

the users. Thanks to the fact that each bucket contains (i) the continent in which the item is produced, to regulate provider fairness, and (ii) the genre of the item, to regulate calibration, both global and individual perspectives are captured at once by our approach. To validate our proposal, we apply it to the recommendations produced by five algorithms, and study the effectiveness of our approach on two datasets (including a novel one, released with this study), and against state-of-the-art approaches for calibrated recommendation and provider fairness.

Concretely, our contributions can be summarized as follows:

- After the identification of the research gaps (Section 2) and characterization of our setting (Section 3), we provide the foundations to our use case, by showing that calibration and provider fairness are related problems, since the genres of the items and their country of production are connected (Section 4);
- We present an approach to post-process the recommendation lists to meet both global and individual goals. We calibrate the results for the individual users in terms of genres, and are fair towards providers (Section 5);
- We face the limitation of evaluating this problem, due to the scarcity of data offering both the category of the items and the sensitive attributes of the providers, so we i) extend the MoviLens-1M dataset, to integrate the continent of production of each item, and ii) we collect and present (in Section 3) a novel dataset. Both resources are publicly available here[4];
- We perform experiments (Section 6) to validate our proposal when applied to the recommendation produced by five algorithms, covering both memory- and model-based approaches, and point-wise and pair-wise approaches. To evaluate its effectiveness in different domains, we consider movie and song recommendation as application scenarios. Based on our outcomes, we highlight possible research paths that might emerge from it (Section 7).

2   Related Work

MORSs. Recent literature has studied how to account for multi-objective goals from different angles. The user perspective was tackled by Li *et al* [26], which balance recommendation accuracy for users with different levels of activities. From an item perspective, Ge *et al.* [15] proposed an approach to balance item relevance and exposure. Considering both the user and item perspectives, Naghiaei *et al.* [32] propose a re-ranking approach to account for consumer and provider fairness. Other studies blend the multiple objectives into a single function, in order to obtain a Pareto-optimal solution [28,44]. Recent advances have also proposed MORSs in sequential settings, by optimizing the results for accuracy, diversity, and novelty [41]. MORS that operate at the individual level have optimized the recommendation process mainly via online interactions, such as conversational approaches [25] or via critiquing [43,12], but approaches aiming at learning individual propensities from past interactions also exist, e.g., [22,34,8,7].

---

[4] https://tinyurl.com/yc6nnx5v

Calibrated recommendation. *Calibration* is a well-studied technique commonly used to solve the problem of unfair output [33,46,42] in recommender systems. Seymen *et al.* [40] address the problem of providing calibration in the recommendations from a constrained optimization perspective. Abdollahpouri *et al.* [1] study the connection between popularity bias, calibration, and consumer fairness in recommendation. Recently, Rojas *et al.* [38] analyze how the calibration method in [42] deals with the bias in different recommendation models. Other studies focus on analyzing user profiles to mitigate miscalibrated recommendations [27] or to mitigate popularity bias from the user's perspective [6]. Existing metrics have some limitations when applying a user-centered approach to evaluate popularity bias and calibrated recommendations. To address these limitations, Abdollahpouri *et al.* [2] present a new metric.

Provider fairness. *Provider fairness* [9] has been studied in many common scenarios, e.g., [29,18,11,10,16,14,30]. It is usually assessed by considering metrics such as the visibility and the exposure that respectively assess the amount of times an item is present in the rankings [13,47] and *where* an item is ranked [5,48], for users to whom each provider's items are recommended. Other approaches, such as that by Karakolis *et al.* [23], consider diversity and coverage for users. Raj *et al.* [35] present a comparative analysis among several fairness metrics recently introduced to measure fair ranking. Wu *et al.* [45] formalize a family of exposure fairness metrics that model the problem of fairness jointly from the perspective of both types of stakeholders.

Contextualizing our work. No MORS can address both calibrated recommendation lists for the users and provider fairness. Our algorithm's aims are to provide i) each user with calibrated recommendations, ii) fair recommendations for the providers, iii) aiming at a minimum loss in effectiveness.

3   Preliminaries

3.1   Recommendation Scenario

Let $U = \{u_1, u_2, ..., u_n\}$ be a set of users, $I = \{i_1, i_2, ..., i_j\}$ 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 $r : U \times I \to V$, 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 $R = \{(u, i, r_{ui}) : u \in U, i \in I, r_{ui} \neq \perp\}$ and they can directly feed an algorithm in the form of triplets (point-wise approaches) or shape user-item observations (pair-wise approaches). We denote with R_u the ratings associated with a user $u \in U$. We consider a temporal split of the data, where a fixed percentage of the ratings of the users (ordered by timestamp) goes to the training and the rest goes to the test set [4]. The 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.

Let C denote the set of geographic continents in which items are organized. We consider a geographic continent as the provenance of an item provider. We denote as C_i the set of geographic continents associated with an item i. Note that, since an item could be produced by more than one provider, it might be associated with several geographic continents, and thus, $|C_i| \geq 1$ and $C_i \subseteq C$. In case two providers belong to the same geographic continent, that continent appears only once; indeed, we are dealing with group fairness so, when a group of providers is associated with an item (once or multiple times), we account for its presence. We use the geographic continents to shape demographic groups, which can be defined to group the ratings of the items produced in a continent (we denote the items in I produced in a continent $c \in C$ as I_c, where $I_c \subseteq I$).

Let G denote the set of genres in which items are organized. We denote as G_i the set of genres associated with an item i. Note that, an item can be of one or more genres, and thus, $|G_i| \geq 1$ and $G_i \subseteq G$. We denote the items in I that have a genre $g \in G$ as I_g, where $I_g \subseteq I$.

3.2   Metrics

Provider-group Representation. In order to enable provider fairness, we should understand the attention received by a provider group in the training data. For this reason, we compute the representation of a demographic group in the data as the number of ratings for items associated with that group in the data. We define with $\mathcal{R}$ the *representation* of a group $c \in C$ as follows:

$$\mathcal{R}_c = \frac{|\{r_{ui} \,:\, u \in U,\, i \in I_c\}|}{|R|} \tag{1}$$

Eq. (1) accounts for the proportion of ratings given to the items of a demographic group associated with a continent. This metric ranges between 0 and 1. We compute the representation of a group only considering the training set. Trivially, the sum of the representations of all groups is equal to 1.

User-based genre propensity. In order to calibrate the results for the users, we need to understand how the preferences for the different item genres are distributed. For this reason, we define with $\mathcal{P}$ the *propensity* of a user of $u \in U$ to rate items of a genre $g \in G$, as follows:

$$\mathcal{P}_{ug} = \frac{|\{r_{ui} \,:\, g \in G_i\}|}{|R_u|} \tag{2}$$

Eq. (2) accounts for the proportion of ratings associated with a genre for a given user. This metric ranges between 0 and 1. Trivially, the sum of the propensities of all genres for a user is equal to 1. This metric is equivalent to the distribution $p(g|u)$ [42].

Disparate Impact. We assess unfairness with the notion of *disparate impact* generated by a recommender system. Specifically, we assess disparate visibility.

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

$$\Delta\mathcal{V}_c = \left(\frac{1}{|U|} \sum_{u \in U} \frac{|\{\hat{r}_{ui} : i \in I_c\}|}{|\hat{R}|} \right) - \mathcal{R}_c \tag{3}$$

The range of values for this score is $[-\mathcal{R}_c, 1 - \mathcal{R}_c]$; 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 defined by Fabbri *et al.* [13].

Miscalibration. We assess the tendency of a system to recommend a user items whose genres are distributed differently from those they prefer via *miscalibration.*

Definition 2 (Miscalibration). *Given a user $u \in U$ and a genre $g \in G$, the* miscalibration *returned by a recommender system for that user is measured as the difference between the share of recommendations for items of that genre and the propensity of the user for that genre in the training data:*

$$\Delta\mathcal{M}_{ug} = \frac{|\{\hat{r}_{ui} : i \in I_g\}|}{|\hat{R}_u|} - \mathcal{P}_{ug} \tag{4}$$

Generalizability. The rest of our paper will consider disparate visibility ($\Delta\mathcal{V}_c$) as the *global* perspective and miscalibration ($\Delta\mathcal{M}_{ug}$) as the *individual* perspective our MORS considers. Nevertheless, our approach can be generalized to *any* metric that assesses the difference between (i) the distribution of the recommendations and (ii) what can be observed in the training set or an objective set by the platform via a policy (e.g., a given amount of content novelty or diversity).

4   Matching Item Providers and Genre Propensity

4.1   Real-world Datasets

First, we extended the MovieLens-1M dataset, so as to integrate the continent of production of each movie. Second, a domain that fits our problem is song recommendation. However, existing music datasets, such as LastFM-2B [31], do not contain song genres and sensitive attributes of the artists, so they do not fit our problem. Thus, we collected a dataset from an online music platform.

In particular, the **MovieLens-1M (Movies)** dataset comprises 1M ratings (range 1-5), from 6,040 users for 3,600 movies across 18 genres. The dataset provides its IMDB ID, which allowed us to associate it to its continent of production, thanks to the OMDB APIs (`http://www.omdbapi.com/`). Keep in mind that *a movie may be produced on more than one continent.* On the other hand, **BeyondSongs (Songs)** contains 1,777,981 ratings (range 1-5), provided by 30,759 users, to 16,380 songs. For each song, we collected the continent of provenance of the artist, and 14 music genres. Both resources are available online[4].

(a) *Movies $\mathcal{R}_c$*

(b) *Songs $\mathcal{R}_c$*

(c) *Movies $\mathcal{P}_{ug}$*

(d) *Songs $\mathcal{P}_{ug}$*
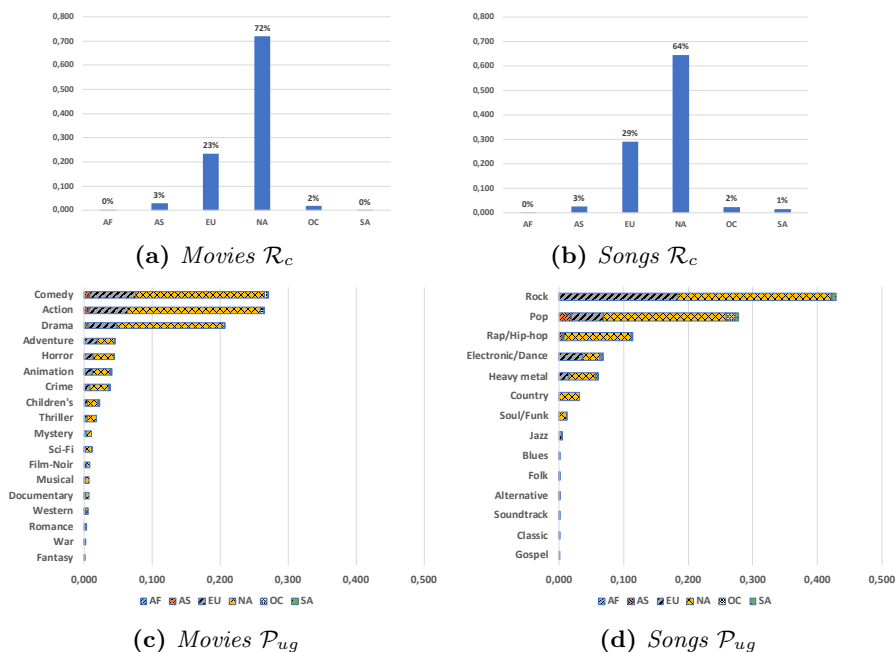
Fig. 1: Group representation (a and b) and genre propensity (c and d) in the Movies and Songs data. *Acronyms stand for AF: Africa, AS: Asia, EU: Europe, NA: North America, OC: Oceania, SA: South America.*

4.2   Characterizing Group Representation and Genre Propensity

We consider the temporal split of the data, where 80% of the ratings are considered for the training set and have been used to measure $\mathcal{R}_c$ and $\mathcal{P}_{ug}$. Note that, while the representation of a demographic group covers the entire training set, the propensity is measured at the user level. Hence, to characterize the link between the two phenomena we aggregate the propensity of all the users for a given genre by summing their values.

Figures 1a and 1b show the $\mathcal{R}_c$ for Movies and Songs, respectively. Both datasets depict a similar representation by continents, where the highest representation is of items from NA providers (72% in movies and 64% in songs) and the second place is for EU providers (23% and 29%). In the rest of the continents, for both datasets, it is less than 10%. Figures 1c and 1d show the $\mathcal{P}_{ug}$ in both datasets; three genres attract most of the ratings by users.

We can also observe that ratings seem to be clustered between certain genre-continent pairs. In other words, different genres are distributed differently across continents. In the Movies data (Fig. 1c), Comedy movies are largely preferred when produced by EU producers, just as Action attracted the majority of ratings for movies by NA producers. In the Songs data (Fig. 1d), the Electronic/Dance genre was consumed much more heavily when produced by EU artists than by

those in the rest of the world, and Heavy metal songs are mostly consumed when they come from NA. In both datasets, users' preferences for the minority provider groups (AF, AS, OC, and SA) are also concentrated on a few selected genres, confirming this rating aggregation in certain genre-continent pairs.

Observation 1. *Users have the propensity to rate items of certain genres and that are produced by certain geographic groups (i.e., in certain continents). Calibration and provider fairness are related problems so, when producing recommendation lists, both perspectives should be accounted at the same time, in MORS fashion.*

5   Individually Calibrated and P-Fair Recommendation

5.1   Algorithm

MOREGIN adjusts the recommendations according to the continent of the providers and the representation of each demographic group and seeks to make a calibration at the individual level, following the propensity of each user to rate items of a given genre. Formally, MOREGIN (see Algorithm 1) works following four main steps. **Steps 1 and 2** are devoted to compute $\mathcal{R}_c$ and $\mathcal{P}_{ug}$, considering the ratings in the training set. **Step 3** computes the items that were predicted as relevant for a user by the recommender system and creates a bucket list, $joinBucket$, considering each continent-genre pair, which will store the predicted items. Each bucket comes with two attributes: $\mathcal{R}_c$ and $\mathcal{P}_{ug}$. 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, $\hat{r}_{ui}$. That item will be stored in the buckets associated with each genre $g \in G_i$ and each continent $c \in C_i$ (even though an item may appear in more than one bucket, it can only be recommended only once). Each element in the bucket is a record that contains the item ID and the $\hat{r}_{ui}$. We sort each bucket considering three values. We sort out $\mathcal{R}_c$ and $\mathcal{P}_{ug}$, in ascending order to ensure the inclusion in the recommendation lists of items from genres and continents that are less represented in the dataset, and we sort in descending order by rating to enhance those products that are relevant to the user. Finally, **Step 4** performs a three-phase re-ranking based on the generated bucket lists. Phase 1 is where we begin, and subsequent phases occur until the top-k is complete. In detail, **Phase 1** selects items starting from the least represented continents to the most represented ones in their corresponding buckets. The algorithm selects items with these conditions: (1) the percentage of items in the recommendation list for a continent is lower or equal to the representation of the continent ($\mathcal{R}_c$); (2) the percentage of items of a given genre in the top-k is lower or equal than $\mathcal{P}_{ug} \cdot k$; and (3) the number of recommended items so far is lower than k. **Phase 2** relaxes the restrictions of phase 1 and here condition 2 is not applied. **Phase 3** selects the items that have the greater relevance for the user, until we complete the top-k. That is, conditions 1 and 2 are not considered.

Input: *recList*: ranked list (records contain *user*, *item*, *rating*, *position*,
    genre , *continent*), 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, *genre*, *continent*), which is sorted by user and rating.
    topk: top k recommendations, we set up $k = 10$.
    topn: top n recommendations, we set up $n= 1000$.
Output: *reRankedList*: ranked list with Individually Calibrated and P-Fair
    Recommendation.

1 define **MOReGIn** (*recList*, *trainList*, *topk*, *topn*)

2 **begin**

  // Step 1. Compute R_c

3 $recBucketRep \leftarrow$ **computeRepresentation**(*topk*, *recList*, *trainList*);

  // Step 2. Compute $\mathcal{P}_{ug}$

4 $recBucketUserProp \leftarrow$ **computePropensity**(*topk*, *recList*, *trainList*);

  // Step 3. Create a bucket list

5 $joinBucket \leftarrow recList + recBucketRep + recBucketUserProp$;

6 $joinBucket \leftarrow$ sort(joinBucket);

  // Step 4. Perform selection of items with three phases

7 $userCounts$, $userGenCounts$, $contCounts \leftarrow \emptyset$;

8 $joinBucket \leftarrow$ **selectWithHardConstraints**(*joinBucket*,
    recBucketRep, *recBucketUserProp*, *userCounts*, *userGenCounts*,
    contCounts) ; // Phase 1

9 $joinBucket \leftarrow$ **selectWithSoftConstraints**(*joinBucket*, *recBucketRep*,
    recBucketUserProp, 2, *userCounts*, *userGenCounts*, *contCounts*) ;
    // Phase 2

10 $joinBucket \leftarrow$ **selectWithSoftConstraints**(*joinBucket*, *recBucketRep*,
    recBucketUserProp, 3, *userCounts*, *userGenCounts*, *contCounts*) ;
    // Phase 3

11 $reRankedList \leftarrow$ chooseSelectedItems(joinBucket);

12 $reRankedList \leftarrow$ sort(reRankedList) ; // sort by user and rating

13   **return** *reRankedList*;

14 **end**

<div align="center">Algorithm 1: Pseudocode of MOReGIn algorithm</div>

6 Experimental Evaluation

6.1 Experimental Methodology

In this work, we focus on well-known state-of-the-art Collaborative Filtering algorithms: **ItemKNN** [39], **UserKNN** [20], **BPRMF** [36], **SVDpp** [24], and **NeuMF** [19]). We will report the results of the original recommendation algorithm (denoted as **OR**). We also consider two comparison baselines: (i) a greedy calibration algorithm [42] (denoted as **CL**) with a λ value of 0.99 (setup defined in [42]), which post-processes the recommendation lists generated by traditional recommender systems; and (ii) a provider fairness algorithm [17] (denoted as **PF**) that considers the providers' continent provenance as a sensitive attribute, with a re-ranking approach that regulates the share of recommendations given

1 define **selectWithHardConstraints(**$joinBucket$, $recBucketRep$,
$recBucketUserProp$, $userCounts$, $uGenCounts$, $contCounts$**)** **begin**

2 $expectedRecordsCont \leftarrow getExpectedRecordsCont(recBucketRep)$;

3 $expectedRecUserGen \leftarrow getRecordsUserGen(recBucketUserProp)$;

4   **foreach** $rec \in joinBucket$ **do** // `for each record`

5 $userGen \leftarrow rec.user + " - " + rec.genre$;

6   **if** $userGen \in expectedRecUserGen$ and
$rec.cont \in expectedRecordsCont$ **then**

7 $userCounts[rec.user] \leftarrow userCounts[rec.user] + 1$;

8 $uGenCounts[rec.userGen] \leftarrow uGenCounts[rec.userGen] + 1$;

9 $contCounts[rec.cont] \leftarrow contCounts[rec.cont] + 1$;

10   **if** $expectedRecUserGen[rec.userGen] \geq userGenCounts$ and
$expectedRecordsCont[rec.cont] \geq contCounts$ and $topk \geq$
$userCounts[rec.user]$ **then**

11 $rec.phase \leftarrow 1$; // `selects element in phase 1`

12 $joinBucket.update(rec)$; // `updates the element`

13   **end**

14   **end**

15   **end**

16   **return** joinBucket

17 **end**

18 define **selectWithSoftConstraints(**$joinBucket$, $recBucketRep$,
$recBucketUserProp$, $phaseMOReGIn$, $userCounts$, $userGenCounts$,
$contCounts$**)** **begin**

19 $expectedRecordsCont \leftarrow getExpectedRecordsCont(recBucketRep)$;

20   **foreach** $rec \in joinBucket$ **do** // `for each record`

21   **if** $rec.cont \in expectedRecordsCont$ **then**

22   **if** $phaseMOReGIn == 2$ **then**

23   **if** $expectedRecordsCont[rec.cont] \geq contCounts$ and $topk \geq$
$userCounts[rec.user]$ **then**

24 $userCounts[rec.user] \leftarrow userCounts[rec.user] + 1$;

25 $contCounts[rec.cont] \leftarrow contCounts[rec.cont] + 1$;

26 $rec.phase \leftarrow 2$; // `selects element in phase 2`

27 $joinBucket.update(rec)$; // `updates the element`

28   **end**

29   **end**

30   **if** $phaseMOReGIn == 3$ **then**

31   **if** $topk \geq userCounts[rec.user]$ **then**

32 $contCounts[rec.cont] \leftarrow contCounts[rec.cont] + 1$;

33 $rec.phase \leftarrow 3$; // `selects element in phase 3`

34 $joinBucket.update(rec)$; // `updates the element`

35   **end**

36   **end**

37   **end**

38   **end**

39   **return** joinBucket

40 **end**

Algorithm 2: Selection methods for the MOReGIn algorithm

to the items produced in a continent (visibility) and the positions in which items are ranked in the recommendation list (exposure).

To run the recommendation models, we used the *Elliot* framework [3], which generated the recommendations for each user that fed the input of MOReGIn. As noted in Section 4.2, the dataset was divided into two sets, one for training (80%) and the other for testing with the most recent ratings of each user (20%).

For each user, we generated the top-1000 recommendations (denoted in the paper as the top-n; we remind the reader that these $n = 1000$ results are not shown to the users, they are only used internally by our algorithm) to then re-rank the top-k (set up to 10) through the proposed MOReGIn algorithm. We performed a grid search of the hyper-parameters for each model in the two datasets. For ItemKNN and UserKNN, in both datasets, we use 50 *neighbors*, a cosine *similarity*, and the classical *implementation*. For BPRMF, SVDpp, and NeuMf we defined 10 *epochs* and 10 *factors* on each dataset, except NeuMF in Movies that uses 12 *factors*. The *batch size* is 512 for SVDpp and NeuMF and is 1 for BPRMF on both datasets. Moreover, for BPRMF in Movies∼Songs, *learning rate*=0.1∼1.346, *bias regularization*=0∼1.236, *user regularization*=0.01∼1.575, *positive item regularization*=0.01∼1.376, and *negative item regularization*=0.01∼1.624; for SVDpp in Movies∼Songs, *learning rate*=0.01∼0.001, *factors regularization*=0.1∼0.001, and *bias regularization*=0.001 in both datasets; NeuMF in Movies∼Songs, the *multi-layer perceptron*=10 in both, *learning rate*=0.0025 in both, and *factors regularization*=0.1∼0.001.

6.2 Assessment of Disparities and Mitigation

Table 1 compares MOReGIn with the baselines in terms of the overall disparate visibility, $\Delta Total$, for each continent. It is computed as $\forall c \in C$, $\Delta Total = \sum \Delta \mathcal{V}_c$. MOReGIn almost entirely reduces the disparities in both movies and song datasets, where most results are $\Delta Total = 0.0000$. Although there is a little difference in the $\Delta Total$ between some approaches, these differences are more explicit, considering the provider provenance. For example, in the movie domain with the BPRMF algorithm, the $\Delta Total$ value in the OR approach is similar to that of PF. However, in a more detailed analysis of more representative continents such as NA and EU, there are notorious differences between the two approaches (i.e., 0.0075 for OR in contrast to -0.0066 for PF in the NA continent, see the example shown in Figure 2a). It is important to highlight that in both domains, our proposal mitigates the disparity regardless of the provenance of the provider, in contrast to the other algorithms that show a clear dependence on the data (i.e., the continent attribute).

Regarding the item genres, Table 2 compares MOReGIn with the baselines in terms of the overall miscalibration, $\Delta Genre$, for each continent. It is computed as $\forall g \in G$ and each user $u \in U$, $\Delta Genre = \sum \Delta \mathcal{M}_{ug}$. For both datasets, MOReGIn obtained the best $\Delta Genre$ (i.e., lowest miscalibration) in all the recommendation models. An analysis of how the algorithms behave with the different genres is shown in Fig. 2b. Although miscalibration never reaches values of $\Delta Genre$ equal to zero, our proposal always calibrates better than the baselines.

(a) *Movies* $\Delta\mathcal{V}_c$

(b) *Movies* $\Delta\mathcal{M}_{ug}$
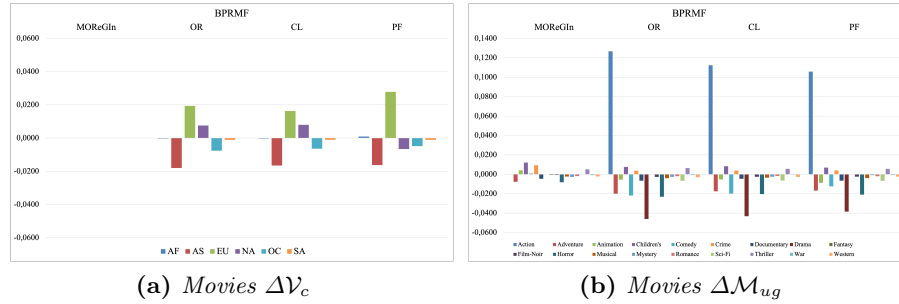
Fig. 2: Disparity mitigation per continent (a) and miscalibration per genre (b) in BPRMF.

Table 1: **Results of disparity mitigation of continents in the Movies and Songs datasets.** Each value represents the sum of disparities, $\Delta Total$.

	MOVIES				SONGS			
	OR	CL	PF	MOReGIn	OR	CL	PF	MOReGIn
BPRMF	0.0539	<u>0.0485</u>	0.0576	**0.0000**	0.2637	<u>0.0840</u>	0.2628	**0.0000**
SVDpp	0.1154	0.1085	<u>0.1059</u>	**0.0000**	0.2678	<u>0.1063</u>	0.2445	**0.0000**
NeuMF	<u>0.0395</u>	0.0421	0.0638	**0.0000**	0.4434	0.4516	<u>0.3990</u>	**0.0000**
UserKNN	0.0345	<u>0.0327</u>	0.0328	**0.0000**	<u>0.0361</u>	0.0575	0.0370	**0.0000**
ItemKNN	0.0431	0.0418	<u>0.0412</u>	**0.0000**	<u>0.0392</u>	0.0583	0.0420	**0.0000**

Observation 2. MOREGIN, *by taking action on the distribution of the items per genre at the user level and provider provenance at the same time, can both calibrate and be fair to the providers. This joint effort allows us to improve the capability to calibrate the results and to be fair to providers with respect to baselines devoted solely to these purposes.*

6.3   Impact on the Quality of Recommendations

We evaluate the accuracy for the different approaches via the NDCG metric.

Table 3 shows its values for MOREGIN and the rest of the baselines, in all the recommendation algorithms, for Movies and Songs. MOREGIN obtained a better NDCG than the PF model, except for BPRMF and ItemKNN in the Movies dataset, and UserKNN and ItemKNN in the Songs dataset. Similar results are obtained with the CL method. Comparing MOREGIN to a non-fair approach, MOReGIn outperforms OR models, with the exception of BPRMF and ItemKNN in the Movie domain. Except for UserKNN and ItemKNN, MORE-GIN also outperforms the OR model in the Songs domain.

All recommendation quality results show that the need for fairer and calibrated recommendations impacts the recommendation quality. However, beyond-accuracy perspectives, such as those offered by MOREGIN allows for compensating for the minimal loss in quality with more unbiased recommendations.

Table 2: **Results of miscalibration of genres in the Movies and Songs datasets.** Each value represents the sum of miscalibrations, $\Delta Genre$.

	MOVIES				SONGS			
	OR	CL	PF	**MOReGIn**	OR	CL	PF	**MOReGIn**
BPRMF	0.2892	0.2606	<u>0.2454</u>	**0.0634**	5.5107	<u>0.0772</u>	0.4610	**0.0289**
SVDpp	0.5792	<u>0.5026</u>	0.5694	**0.1184**	0.5773	0.1031	<u>0.5029</u>	**0.0256**
NeuMF	0.4596	0.3962	<u>0.3735</u>	**0.2901**	1.2886	<u>0.7494</u>	1.2202	**0.0787**
UserKNN	0.0743	0.0862	<u>0.0580</u>	**0.0392**	0.0298	0.0989	<u>0.0291</u>	**0.0208**
ItemKNN	0.2102	0.1966	<u>0.1954</u>	**0.0559**	0.0890	<u>0.0601</u>	0.0879	**0.0205**

Table 3: **NDCG for each approach and recommendation algorithm**.

	MOVIES				SONGS			
	OR	CL	PF	**MOReGIn**	OR	CL	PF	**MOReGIn**
BPRMF	**0.3204**	0.3144	<u>0.3195</u>	0.3057	0.0034	**0.0067**	0.0031	<u>0.0055</u>
SVDpp	0.0830	<u>0.0888</u>	0.0812	**0.1024**	0.0050	<u>0.0103</u>	0.0051	**0.0138**
NeuMF	<u>0.1963</u>	0.1931	0.1956	**0.2050**	<u>0.0183</u>	0.0098	0.0179	**0.0314**
UserKNN	<u>0.3051</u>	0.2954	0.3030	**0.3053**	**0.3760**	0.1925	<u>0.3759</u>	0.2648
ItemKNN	**0.3229**	0.3145	<u>0.3211</u>	0.3131	**0.3860**	0.1668	<u>0.3857</u>	0.2864

7   Conclusions and Future Work

Global and individual objectives in MORs have never been studied jointly. To study this problem, we provided data, by i) extending the MovieLens-1M dataset and ii) collecting a new dataset for song recommendation. The analysis of this data showed that when users rate items of a given genre, the geographic provenance of that item matters. Based on these insights, we proposed a new post-processing approach, named MOReGIn, that aggregates the recommended items into buckets, pairing item genres and their continent of production. Results show that MOReGIn outperforms the existing approaches at producing effective, calibrated, and provider-fair recommendations. Future work will explore different strategies to generate recommendation lists given the generated buckets. Moreover, we will consider consumer fairness as a global perspective.

Acknowledgments

References

1. Abdollahpouri, H., Mansoury, M., Burke, R., Mobasher, B.: The connection between popularity bias, calibration, and fairness in recommendation. In: Fourteenth ACM conference on recommender systems. pp. 726–731 (2020)

2. Abdollahpouri, H., Mansoury, M., Burke, R., Mobasher, B., Malthouse, E.: User-centered evaluation of popularity bias in recommender systems. In: Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization. pp. 119–129 (2021)

3. Anelli, V.W., Bellogin, A., Ferrara, A., Malitesta, D., Merra, F.A., Pomo, C., Donini, F.M., Di Noia, T.: Elliot: A Comprehensive and Rigorous Framework for Reproducible Recommender Systems Evaluation, p. 2405–2414. Association for Computing Machinery, New York, NY, USA (2021)

4. Bellogín, A., Castells, P., Cantador, I.: Statistical biases in information retrieval metrics for recommender systems. Inf. Retr. Journal **20**(6), 606–634 (2017). `https://doi.org/https://doi.org/10.1007/s10791-017-9312-z`

5. Biega, A.J., Gummadi, K.P., Weikum, G.: Equity of attention: Amortizing individual fairness in rankings. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018. pp. 405–414. ACM, New York, NY, USA (2018). `https://doi.org/10.1145/3209978.3210063`

6. Chen, J., Wu, W., Shi, L., Zheng, W., He, L.: Long-tail session-based recommendation from calibration. Appl. Intell. **53**(4), 4685–4702 (2023). `https://doi.org/10.1007/s10489-022-03718-7`, `https://doi.org/10.1007/s10489-022-03718-7`

7. Dokoupil, P., Peska, L., Boratto, L.: Looks can be deceiving: Linking user-item interactions and user's propensity towards multi-objective recommendations. CoRR **abs/2307.00654** (2023). `https://doi.org/10.48550/arXiv.2307.00654`, `https://doi.org/10.48550/arXiv.2307.00654`

8. Dokoupil, P., Peska, L., Boratto, L.: Rows or columns? minimizing presentation bias when comparing multiple recommender systems. In: Chen, H., Duh, W.E., Huang, H., Kato, M.P., Mothe, J., Poblete, B. (eds.) Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023. pp. 2354–2358. ACM (2023). `https://doi.org/10.1145/3539618.3592056`, `https://doi.org/10.1145/3539618.3592056`

9. Ekstrand, M.D., Das, A., Burke, R., Diaz, F.: Fairness in recommender systems. In: Ricci, F., Rokach, L., Shapira, B. (eds.) Recommender Systems Handbook, pp. 679–707. Springer US (2022). `https://doi.org/10.1007/978-1-0716-2197-4_18`, `https://doi.org/10.1007/978-1-0716-2197-4_18`

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

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

12. Elahi, M., Ge, M., Ricci, F., Massimo, D., Berkovsky, S.: Interactive food recommendation for groups. In: 8th ACM Conference on Recommender Systems, RecSys 2014. CEUR-WS (2014)

13. Fabbri, F., Bonchi, F., Boratto, L., Castillo, C.: The effect of homophily on disparate visibility of minorities in people recommender systems. In: Proceedings of the Fourteenth International AAAI Conference on Web and Social Media, ICWSM 2020. pp. 165–175. AAAI Press, USA (2020), `https://ojs.aaai.org/index.php/ICWSM/issue/view/262`

14. Ferraro, A., Serra, X., Bauer, C.: What is fair? exploring the artists' perspective on the fairness of music streaming platforms. In: Human-Computer Interaction - INTERACT 2021 - 18th IFIP TC 13 International Conference, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12933, pp. 562–584. Springer (2021). `https://doi.org/10.1007/978-3-030-85616-8_33`, `https://doi.org/10.1007/978-3-030-85616-8_33`

15. Ge, Y., Zhao, X., Yu, L., Paul, S., Hu, D., Hsieh, C., Zhang, Y.: Toward pareto efficient fairness-utility trade-off in recommendation through reinforcement learning. In: Candan, K.S., Liu, H., Akoglu, L., Dong, X.L., Tang, J. (eds.) WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022. pp. 316–324. ACM (2022). `https://doi.org/10.1145/3488560.3498487`, `https://doi.org/10.1145/3488560.3498487`

16. Gharahighehi, A., Vens, C., Pliakos, K.: Fair multi-stakeholder news recommender system with hypergraph ranking. Information Processing & Management **58**(5), 102663 (2021). `https://doi.org/https://doi.org/10.1016/j.ipm.2021.102663`, `https://www.sciencedirect.com/science/article/pii/S0306457321001515`

17. Gómez, E., Boratto, L., Salamó, M.: Provider fairness across continents in collaborative recommender systems. Inf. Process. Manag. **59**(1), 102719 (2022). `https://doi.org/10.1016/j.ipm.2021.102719`

18. Gómez, E., Zhang, C.S., Boratto, L., Salamó, M., Ramos, G.: Enabling cross-continent provider fairness in educational recommender systems. Future Gener. Comput. Syst. **127**, 435–447 (2022). `https://doi.org/10.1016/j.future.2021.08.025`

19. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., Chua, T.S.: Neural collaborative filtering. In: Proceedings of the 26th international conference on world wide web. pp. 173–182 (2017)

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

21. Jannach, D.: Multi-objective recommendation: Overview and challenges. In: Abdollahpouri, H., Sahebi, S., Elahi, M., Mansoury, M., Loni, B., Nazari, Z., Dimakopoulou, M. (eds.) Proceedings of the 2nd Workshop on Multi-Objective Recommender Systems co-located with 16th ACM Conference on Recommender Systems (RecSys 2022), Seattle, WA, USA, 18th-23rd September 2022. CEUR Workshop Proceedings, vol. 3268. CEUR-WS.org (2022), `https://ceur-ws.org/Vol-3268/paper1.pdf`

22. Jugovac, M., Jannach, D., Lerche, L.: Efficient optimization of multiple recommendation quality factors according to individual user tendencies. Expert Systems with Applications **81**, 321–331 (2017). `https://doi.org/https://doi.org/10.1016/j.eswa.2017.03.055`, `https://www.sciencedirect.com/science/article/pii/S0957417417302075`

23. Karakolis, E., Kokkinakos, P., Askounis, D.: Provider fairness for diversity and coverage in multi-stakeholder recommender systems. Applied Sciences

12(10) (2022). `https://doi.org/10.3390/app12104984`, `https://www.mdpi.com/2076-3417/12/10/4984`

24. Koren, Y.: 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. ACM (2008). `https://doi.org/10.1145/1401890.1401944`

25. Li, R., Kahou, S., Schulz, H., Michalski, V., Charlin, L., Pal, C.: Towards deep conversational recommendations. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. p. 9748–9758. NIPS'18, Curran Associates Inc., Red Hook, NY, USA (2018)

26. Li, Y., Chen, H., Fu, Z., Ge, Y., Zhang, Y.: User-oriented fairness in recommendation. In: Leskovec, J., Grobelnik, M., Najork, M., Tang, J., Zia, L. (eds.) WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021. pp. 624–632. ACM / IW3C2 (2021). `https://doi.org/10.1145/3442381.3449866`, `https://doi.org/10.1145/3442381.3449866`

27. Lin, K., Sonboli, N., Mobasher, B., Burke, R.: Calibration in collaborative filtering recommender systems: a user-centered analysis. In: Proceedings of the 31st ACM Conference on Hypertext and Social Media. pp. 197–206 (2020)

28. Lin, X., Chen, H., Pei, C., Sun, F., Xiao, X., Sun, H., Zhang, Y., Ou, W., Jiang, P.: A pareto-efficient algorithm for multiple objective optimization in e-commerce recommendation. In: Bogers, T., Said, A., Brusilovsky, P., Tikk, D. (eds.) Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019. pp. 20–28. ACM (2019). `https://doi.org/10.1145/3298689.3346998`, `https://doi.org/10.1145/3298689.3346998`

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

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

31. Melchiorre, A.B., Rekabsaz, N., Parada-Cabaleiro, E., Brandl, S., Lesota, O., Schedl, M.: Investigating gender fairness of recommendation algorithms in the music domain. Inf. Process. Manage. **58**(5) (sep 2021)

32. Naghiaei, M., Rahmani, H.A., Deldjoo, Y.: Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. In: SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 770–779. ACM (2022). `https://doi.org/10.1145/3477495.3531959`, `https://doi.org/10.1145/3477495.3531959`

33. Nixon, J., Dusenberry, M.W., Zhang, L., Jerfel, G., Tran, D.: Measuring calibration in deep learning. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 38–41. Computer Vision Foundation / IEEE (2019), `http://openaccess.thecvf.com/content_CVPRW_2019/html/Uncertainty_and_Robustness_in_Deep_Visual_Learning/Nixon_Measuring_Calibration_in_Deep_Learning_CVPRW_2019_paper.html`

34. Peska, L., Dokoupil, P.: Towards results-level proportionality for multi-objective recommender systems. In: Amigó, E., Castells, P., Gonzalo, J., Carterette, B., Culpepper, J.S., Kazai, G. (eds.) SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022. pp. 1963–1968. ACM (2022). https://doi.org/10.1145/3477495.3531787, https://doi.org/10.1145/3477495.3531787

35. Raj, A., Ekstrand, M.D.: Measuring fairness in ranked results: An analytical and empirical comparison. In: SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 726–736. ACM, New York, NY, USA (2022). https://doi.org/10.1145/3477495.3532018, https://doi.org/10.1145/3477495.3532018

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

37. Ricci, F., Rokach, L., Shapira, B.: Recommender systems: Techniques, applications, and challenges. In: Ricci, F., Rokach, L., Shapira, B. (eds.) Recommender Systems Handbook, pp. 1–35. Springer US (2022). https://doi.org/10.1007/978-1-0716-2197-4_1, https://doi.org/10.1007/978-1-0716-2197-4_1

38. Rojas, C., Contreras, D., Salamó, M.: Analysis of biases in calibrated recommendations. In: Advances in Bias and Fairness in Information Retrieval. pp. 91–103. Springer International Publishing, Cham (2022)

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

40. Seymen, S., Abdollahpouri, H., Malthouse, E.C.: A constrained optimization approach for calibrated recommendations. In: Fifteenth ACM Conference on Recommender Systems. pp. 607–612 (2021)

41. Stamenkovic, D., Karatzoglou, A., Arapakis, I., Xin, X., Katevas, K.: Choosing the best of both worlds: Diverse and novel recommendations through multi-objective reinforcement learning. In: Candan, K.S., Liu, H., Akoglu, L., Dong, X.L., Tang, J. (eds.) WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022. pp. 957–965. ACM (2022). https://doi.org/10.1145/3488560.3498471, https://doi.org/10.1145/3488560.3498471

42. Steck, H.: Calibrated recommendations. In: Proceedings of the 12th ACM Conference on Recommender Systems. p. 154–162. RecSys '18, Association for Computing Machinery, New York, NY, USA (2018). https://doi.org/10.1145/3240323.3240372

43. Wang, Z., Meng, C., Ji, S., Li, T., Zheng, Y.: Food package suggestion system based on multi-objective optimization: A case study on a real-world restaurant. Applied Soft Computing 93, 106369 (2020). https://doi.org/https://doi.org/10.1016/j.asoc.2020.106369, https://www.sciencedirect.com/science/article/pii/S1568494620303094

44. Wu, H., Ma, C., Mitra, B., Diaz, F., Liu, X.: Multi-fr: A multi-objective optimization framework for multi-stakeholder fairness-aware recommendation. In: Transactions on Information Systems (TOIS). ACM (2022)

45. Wu, H., Mitra, B., Ma, C., Diaz, F., Liu, X.: Joint multisided exposure fairness for recommendation. In: Proceedings of the 45th International ACM SIGIR Conference

on Research and Development in Information Retrieval. p. 703–714. SIGIR '22, Association for Computing Machinery, New York, NY, USA (2022). `https://doi.org/10.1145/3477495.3532007`, `https://doi.org/10.1145/3477495.3532007`

46. Zadrozny, B., Elkan, C.: Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In: Icml. vol. 1, pp. 609–616. Citeseer (2001)

47. Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M., Baeza-Yates, R.: Fa*ir: A fair top-k ranking algorithm. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017. pp. 1569–1578. ACM, New York, NY, USA (2017). `https://doi.org/10.1145/3132847.3132938`

48. Zehlike, M., Castillo, C.: Reducing disparate exposure in ranking: A learning to rank approach. In: WWW '20: The Web Conference 2020. pp. 2849–2855. ACM / IW3C2, New York, NY, USA (2020). `https://doi.org/10.1145/3366424.3380048`

49. Zheng, Y., Wang, D.X.: A survey of recommender systems with multi-objective optimization. Neurocomputing **474**, 141–153 (2022). `https://doi.org/10.1016/j.neucom.2021.11.041`, `https://doi.org/10.1016/j.neucom.2021.11.041`