



Oh, Behave! Country Representation Dynamics Created by Feedback Loops in Music Recommender Systems

Oleg Lesota
 oleg.lesota@jku.at
 Johannes Kepler University Linz and
 Linz Institute of Technology
 Linz, Austria

Jonas Geiger
 jonasgeiger@outlook.de
 Johannes Kepler University Linz and
 Linz Institute of Technology
 Linz, Austria

Max Walder
 max@gstat.eu
 Johannes Kepler University Linz and
 Linz Institute of Technology
 Linz, Austria

Dominik Kowald
 dkowald@know-center.at
 Know-Center GmbH and
 Graz University of Technology
 Graz, Austria

Markus Schedl
 markus.schedl@jku.at
 Johannes Kepler University Linz and
 Linz Institute of Technology
 Linz, Austria

ABSTRACT

Recent work suggests that music recommender systems are prone to disproportionately frequent recommendations of music from countries more prominently represented in the training data, notably the US. However, it remains unclear to what extent feedback loops in music recommendation influence the dynamics of such imbalance. In this work, we investigate the dynamics of representation of local (i.e., country-specific) and US-produced music in user profiles and recommendations. To this end, we conduct a feedback loop simulation study using the LFM-2b dataset. The results suggest that most of the investigated recommendation models decrease the proportion of music from local artists in their recommendations. Furthermore, we find that models preserving average proportions of US and local music do not necessarily provide country-calibrated recommendations. We also look into popularity calibration and, surprisingly, find that the most popularity-calibrated model in our study (ItemKNN) provides the least country-calibrated recommendations. In addition, users from less represented countries (e.g., Finland) are, in the long term, most affected by the under-representation of their local music in recommendations.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; • **Social and professional topics** → *Cultural characteristics*.

KEYWORDS

Music Recommendation, Feedback Loop, Demographic Bias

ACM Reference Format:

Oleg Lesota, Jonas Geiger, Max Walder, Dominik Kowald, and Markus Schedl. 2024. Oh, Behave! Country Representation Dynamics Created by Feedback Loops in Music Recommender Systems. In *18th ACM Conference*



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

RecSys '24, October 14–18, 2024, Bari, Italy
 © 2024 Copyright held by the owner/author(s).
 ACM ISBN 979-8-4007-0505-2/24/10
<https://doi.org/10.1145/3640457.3688187>

on *Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3640457.3688187>

1 INTRODUCTION

Recent research has demonstrated the effectiveness of long-term investigations in studying bias and representation dynamics induced by feedback loops in recommender systems and decision support systems [2, 3, 9, 13, 18]. For instance, Mansoury et al. investigated the amplification of popularity bias [7, 8, 11] by examining the average item popularity and aggregating diversity in movie recommendation lists across multiple iterations [13]. In Lacic et al. [9], the authors studied long-term dynamics in news recommender systems through an online study, measuring the impact of introducing a personalized recommendation algorithm on popularity bias compared to a purely popularity-based one. The results revealed that significant changes in users' news consumption behavior could be observed after the end of the study. Apart from popularity bias, previous research has investigated related phenomena, namely traces of globalization in music consumption patterns amplified by music recommender systems: Lesota et al. identified a strong position of US-generated music (i.e., the majority group) in the consumption behavior of users from other countries [12]. However, it is still unclear to what extent feedback loops in music recommendation influence country representation dynamics (of users and artists) in the long run. To close this research gap, we study the following questions:

- **RQ1:** How do different recommender systems affect representation of local and US-produced music in recommendations and user profiles in the long term?
- **RQ2:** How do effects of feedback loops vary across different countries? Do different recommender systems treat individual countries differently?

To the best of our knowledge this work is the first to (1) investigate the impact recommender systems can have on the representation of music from different countries in the *long term*, and (2) describe the impact in terms of the degree of mismatch between country label distributions over the tracks in user consumption histories and their corresponding recommendations (*miscalibration of country distributions*).

2 METHODOLOGY

We address the research questions by conducting an offline simulation study, following the general procedure introduced by Mansoury et al. [13]. We simulate user interaction with a recommendation model over a long period of time via a feedback loop. At each iteration of the loop, the model produces recommendations. One of the recommended items per user¹ is added to the respective user’s interaction history (simulating item consumption). After that, the model is retrained on the data enriched with the new interactions of all users and then the process is repeated multiple times. The recommendations and simulated user profiles at each iteration are evaluated regarding the representation of local and US music.

Simulation procedure. At each iteration, the following steps are performed for the recommendation algorithm under investigation:

- (1) We randomly split the input dataset into a training set (75%), validation set (20%) and test set (5%).
- (2) We train the recommendation model under investigation on the training set and use NDCG@k as selection criterion on the validation set. After the training process has terminated, the final model is used to produce k recommendations for each user. Items already seen by the user are excluded.
- (3) From the recommended items $\{t\}$ generated by the model, one is added to the user profile based on an acceptance probability concept adapted from [13]. For each user and their list of k recommended items (R), the probability is expressed as $prob(t|R) = \frac{e^{\alpha \cdot rank_t}}{\sum_{j=1}^k e^{\alpha \cdot j}}$ where $rank_t$ is the rank of the item t in the sorted list R and α is a hyperparameter < 0 .
- (4) After an item has been selected for all users, a new dataset is created combining the input dataset with the new (simulated) interactions. The obtained dataset is then used for the next simulation iteration as the new input dataset.

Evaluation. The following indicators and metrics are used to evaluate country representation in user recommendations and profiles. *US and local proportions* show how many items in the profile or recommendations of user u originate from the US or the country of the user. The local proportion p_u^{local} for a given set of interactions or recommendations I_u of a user is defined as $p_u^{local} = \frac{|I_u^{local}|}{|I_u|}$, where I_u^{local} is the set of all tracks where user country and track country coincide. The US proportion p_u^{US} is defined analogously. The proportions are defined on the user level and averaged over users. We analyze the difference in these proportions between the original, unmodified dataset and the state of the recommendations or user profiles at iteration i . *Popularity and country miscalibration* are used to show how much the user profiles deviate from their original states after n iterations of the simulation. Following the assumption that the users prefer calibrated recommendations [19], we compute the miscalibration of their profiles with regard to the music track attributes of country and popularity as Jensen–Shannon divergence (JSD), see Equation 1.

¹We simulate consumption of only one item per user per iteration to avoid the situation of “running out of item supply”, which may occur in our setting when items already consumed by a user can no longer be recommended to them.

Table 1: Basic statistics of our LFM-2b dataset sample.

	Tracks	Track Interactions		Users	User Interactions
		Total	Average		
US	39,614	1,040,360	26.26	1,582	323,072
UK	15,522	422,225	27.20	823	171,469
DE	6,793	107,832	15.87	805	158,642
SE	4,519	107,491	23.79	320	60,993
CA	3,754	95,343	25.40	217	47,490
FR	2,800	56,241	20.09	254	52,850
AU	2,346	53,701	22.89	193	40,767
FI	2,260	45,709	20.23	420	78,819
NO	1,765	36,769	20.83	208	40,545
BR	2,236	35,964	16.08	1,064	205,093
NL	1,738	32,035	18.43	375	89,546
PL	1,709	27,116	15.87	1,040	195,296
RU	1,888	24,086	12.76	1,162	187,876
JP	1,796	21,818	12.15	101	14,411
IT	1,506	21,273	14.13	222	37,421
other	9,651	159,769	16.55	2,990	583,442
Total	99,897	2,287,732	22.90	11,776	2,287,732

$$JSD(H_u, H_u^*) = \frac{1}{2} \sum_c H_u(c) \log_2 \frac{2H_u(c)}{H_u(c) + H_u^*(c)} + \frac{1}{2} \sum_c H_u^*(c) \log_2 \frac{2H_u^*(c)}{H_u(c) + H_u^*(c)} \quad (1)$$

H_u is the distribution of a track attribute over unmodified interaction history of user u and H_u^* is the same for the interaction history after n iterations. $H(c)$ is the (probability) value in the distribution H , corresponding to the track attribute value c . In case of popularity, $c \in [HighPop, MidPop, LowPop]$ defined as in [1, 10]. For country $c \in [local, US, other]$. The per-user JSD scores are averaged to represent miscalibration affecting a user group.

3 EXPERIMENTAL SETUP

To foster reproducibility, we make the source code, data sample and other materials publicly available on GitHub².

Dataset. We conduct our experiments on the LFM-2b dataset [17]. It contains listening events from the online music aggregator platform Last.fm.³ First, we apply the same pre-processing and augmentation operations as Lesota et al. [12]. Specifically, we extend the dataset with information about the country of origin of artists from Musicbrainz⁴ to attribute a country to each music track. We then filter (1) all tracks listened to only once, and (2) all interactions for which no country information about the user or track is available. We consider interactions from the years 2018-2019, the two most recent full years in the dataset. Additionally, due to computational limitations, we only work with a sub-sample of this filtered dataset. We randomly sample 100K tracks, ensuring that the resulting data sample is 5-core filtered. This results in a dataset of ~2290K interactions between ~100K tracks and ~12K users. Detailed statistics are reported in Table 1.

²<https://github.com/hcai-mms/FeedbackLoops4RecSys>

³<https://www.last.fm>

⁴<https://musicbrainz.org>

Models. We perform a feedback loop simulation with $n = 100$ iterations using the method described above for each of the following six models: BPR [16], ItemKNN [4], LightGCN [5], MultiVAE [20], NeuMF [6], and Pop. ItemKNN is a traditional approach that assigns recommendation scores to items based on how similar an item is to those already consumed by a user, depending on the existing interactions of other users. Bayesian Personalized Ranking (BPR) is a matrix factorization [15] approach using a personalized loss, which enforces explicit ranking between pairs of items. LightGCN is a graph-based approach employing a Graph Convolution Network, where the reconstruction of the interaction matrix also considers the embeddings of neighboring users and items on the user-item interaction graph. Neural Matrix Factorization (NeuMF) employs a multi-layer perceptron, allowing it to learn a custom matching function. MultiVAE is a variational autoencoder that learns a latent user representation from users' interaction vectors, and predicts the distribution of the relevance score over all items. The Pop model recommends the most popular unseen items to each user, where popularity is computed as frequency of interactions over all users. It serves as a baseline model against which the more complex models are compared.

Training and evaluation procedure. We run our experiments using the Recbole library [21] and its default settings for each model. We train for a maximum of 200 epochs, with early stopping applied after not observing any improvement in terms of NDCG@10 for 5 subsequent epochs. For the simulated acceptance of new items, we consider the items with top $k = 10$ scores for each user. The acceptance probability is calculated with $\alpha = -0.1$. Following [12], when analyzing the results for individual countries we concentrate on ones that are represented by at least 100 users and 1000 tracks.⁵

4 RESULTS

We summarize the main results concerning **RQ1** in Table 2 and Figure 1. Table 2 presents the simulated impact the six investigated recommenders made on user recommendations and user profiles. The rows Rec_{local} and Rec_{US} show how different the proportions of local and US-produced music recommended at iteration 100 are to the respective proportions in the original user profiles in percent. Similarly, the rows $Prof_{local}$ and $Prof_{US}$ show the difference between user profiles before and after the simulation. For example, the value of $+17.0^*$ for Rec_{US} and NeuMF means that at the last iteration of the simulation, the average proportion of US music recommended by this model is 17% higher than the proportion of the US music consumed by the users before the simulation. Statistically significant⁶ inconsistencies are marked with *. JSD_{Prof} shows the miscalibration in terms of country distribution between user profiles before and after the simulation. $nDCG_1$ shows model accuracy evaluation at iteration 1, verifying that at the initial state, every model has shown adequate performance for the dataset (similar results were achieved on other sub-samples of LFM-2b [10, 14]). In Figure 1 every point corresponds to the average proportion of local

⁵We abbreviate the countries as follows. US: United States, UK: United Kingdom, DE: Germany, SE: Sweden, CA: Canada, FR: France, AU: Australia, FI: Finland, BR: Brazil, RU: Russia, JP: Japan, NO: Norway, PL: Poland, NL: Netherlands, IT: Italy

⁶As determined by a two-sided T-Test with the Bonferroni-corrected threshold of 0.05/12 as each initial proportion is compared to ones produced by each of 6 algorithms both in the recommendations and in user profiles.

Table 2: Proportion of local and US-produced music in recommendations (Rec_*) and simulated user profiles ($Prof_*$) after iteration 100 compared to the respective proportions in the original user profiles (difference in %). Significant changes are marked with *.

	Pop	ItemKNN	BPR	NeuMF	MultiVAE	LightGCN
Rec_{local}	-47.5*	+0.4	-6.6*	-39.7*	-21.2*	-3.5
Rec_{US}	+15.6*	+4.8*	+3.5*	+17.0*	+7.5*	+0.5
$Prof_{local}$	-22.5*	+2.1	-7.0*	-19.5*	-10.5*	-3.6
$Prof_{US}$	+19.5*	+2.2*	+3.7*	+7.6*	+2.6*	+0.8
JSD_{Prof}	0.13	0.12	0.08	0.10	0.09	0.08
$nDCG_1$	0.03	0.26	0.13	0.08	0.12	0.14




Figure 1: Proportion of local items recommended by different algorithms. The dashed line shows the average consumption of local music before the simulation.

music in recommendations across all users per algorithm per iteration. The dashed line (0.18) shows the proportion of consumed local music in the initial dataset. We approach **RQ2** analyzing Table 3 and Table 4. The former shows miscalibration between country distributions over tracks in user profiles before and after simulation for each model and each of the considered countries, expressed as JSD . The countries in the table are arranged in the order of decreasing number of respective tracks. The highest and lowest values of miscalibration among non-baseline algorithms for each country are put in bold and underlined, respectively. Table 4 shows the mismatch in average proportions of local and US-produced music in user profiles before and after the simulation in percent. We make additional observations regarding popularity and country calibration using Figure 2 showing the progression of miscalibration between the simulated user profiles at each iteration and their original state.

4.1 RQ1: Representation of Local and US-produced Music in the Long Term

We observe that the recommendations produced by all models, except for LightGCN, at iteration 100 significantly differ from the initial user profiles in terms of average proportions of US and local music (Table 2); with an overall trend of increasing average proportions for US music and decreasing for local music (except for




Figure 2: Miscalibration between the interaction history at iteration i and the initial interaction history in terms of country (left) and popularity (right). Measured as JSD .

ItemKNN) in recommendations. Figure 1 shows that at iteration 1 all models, except for LightGCN, on average suggest a proportion of local tracks inconsistent with the initial dataset. The proportions of local tracks deviate further away in recommendations by MultiVAE and NeuMF and potentially converge to the initial level for ItemKNN and BPR. NeuMF shows the highest deviation, comparable to Pop. The user profiles’ shift from their original state corresponds to the discussed shift in the recommendations (Table 2). Additionally, LightGCN causes the least country-miscalibrated (JSD_{Prof}) user profiles on average, while the highest miscalibration in terms of country is caused by ItemKNN.

Main findings. We conclude that recommender systems tend to over-represent US and under-represent local music. However, some models (ItemKNN, BPR, LightGCN) may converge to the original levels of local representation in the long term. The overall algorithmic impact varies from the most calibrated LightGCN that also preserves the initial proportions, to the most miscalibrated ItemKNN still preserving the initial proportions, and notably miscalibrated NeuMF significantly affecting both country proportions.

4.2 RQ2: Effects of Feedback Loops Across Different Countries

Table 3 demonstrates that the majority of algorithms treat users from FI, PL, RU, JP, and IT with the highest country miscalibration. The latter four countries are among the ones with the fewest tracks and the least average interactions per track. On the contrary, FI shows above median number of tracks and average track popularity, more than BR, which however is less affected by the miscalibration. Brazilian users are more frequently represented in the training data, which could be the reason for more calibrated recommendations. Among the countries experiencing lower country miscalibration are US, NL, CA, AU, and NO. The latter four are among the least represented. Notably, UK features more tracks and higher average track popularity than the four, while suffering higher miscalibration. Previous work [12] suggests that CA, NL, AU, and NO consume above average US-produced and below average of domestic music, unlike UK or DE. *The results show that feedback loops affect countries differently depending on multiple factors from country representation in the training data to consumption habits of respective users. More research is needed to determine clearer dependencies.*

Table 4 shows that ItemKNN significantly affected average proportions of local or US music only for UK and PL. NeuMF, however, significantly affected at least one of the proportions for all considered countries. Except for NeuMF, all non-baseline models over-represent US-music for a limited number of countries, while the local music is under-represented for more combinations of country and model. This leads us to believe that the local music is gradually replaced by tracks from countries beyond the US and thus the progressive under-representation of local music is not necessarily caused by over-representation of music from one particular country (e.g. US), but rather by its own initial low representation.

Main findings. The list of significantly affected countries varies between algorithms. Overall on per-country level most changes are in local rather than US proportions.

4.3 Notes on Calibration

Comparing the progression of miscalibration of user profiles from their original state in terms of country and popularity distributions (Figure 2), we observe that in both cases at iteration 100 the algorithms can be ranked according to the degree of miscalibration. Notably the rankings for popularity and country miscalibration do not match, for instance while ItemKNN shows the lowest popularity miscalibration it results in the most country-miscalibrated user profiles.

Main findings. This suggests that while country miscalibration is likely connected to popularity bias, the connection is not direct, as the least popularity-biased model provides the most country-miscalibrated results.

5 LIMITATIONS AND DISCUSSION

In this section, we describe the limitations of our work and discuss our main findings mentioned in the previous section.

Country and culture. In this work, we investigate the dynamics of country representation in user profiles and recommendations. It is worth pointing out that in our study, the country labels are assigned to tracks based on “... the area with which an artist is primarily identified with. It is often, but not always, its birth/formation country...”⁷. This means that the country label does not necessarily reflect the language of the track, or to which extent the track is representative of the culture of the corresponding country. This limits the scope of our findings to the representation of artists from certain countries, which, while valuable, can be only a weak indicator of cultural representation. At the same time, formalizing the cultural signature of a music piece is a non-trivial task. The same could be said about assigning a language label to a track when it comes to classical, instrumental, electronic music or music performed in dead languages. With this paper we strive to ignite the discussion about representation in recommender systems in terms of culture, country, language and other attributes for all stakeholders and suggest one possible approach to studying it.

Choice model. When investigating feedback loops, we have to select a user choice model as a rule for simulating user interaction with recommendation lists. In our case, at each iteration the model randomly “consumes” one item per user with slight preference

⁷<https://musicbrainz.org/doc/Artist#Area>

Table 3: Miscalibration between three-bin country distributions over the original user profiles and the profiles after 100 iterations (measured as $JS D$). Higher values indicate higher inconsistency. Darker background indicates higher miscalibration, normalized per algorithm.

	US	UK	DE	SE	CA	FR	AU	FI	NO	BR	NL	PL	RU	JP	IT	all
Pop	0.073	0.126	0.146	0.133	0.108	0.144	0.101	0.186	0.108	0.133	0.096	0.16	0.163	0.196	0.153	0.128
ItemKNN	0.084	0.123	0.139	0.137	0.106	0.128	0.109	0.152	0.127	0.134	0.094	0.155	0.149	0.166	0.174	0.124
BPR	0.054	0.093	0.089	0.09	0.085	0.095	0.088	0.103	0.084	0.083	0.07	0.103	0.098	0.146	0.114	0.083
NeuMF	0.058	0.099	0.114	0.104	0.09	0.111	0.093	0.126	0.087	0.099	0.074	0.121	0.115	0.179	0.135	0.096
MultiVAE	0.062	0.098	0.101	0.099	0.092	0.1	0.094	0.114	0.091	0.1	0.074	0.111	0.108	0.132	0.136	0.092
LightGCN	0.053	0.089	0.079	0.081	0.08	0.086	0.091	0.088	0.084	0.069	0.068	0.093	0.086	0.099	0.13	0.077

Table 4: Deviations in the average proportions of local and US music in user profiles at iteration 100 from the respective average proportions in the original user profiles before simulation (in percent). Statistically significant deviations are marked with *. The highest value for each column for each indicator is depicted in bold, the lowest is underlined.

	US	UK	DE	SE	CA	FR	AU	FI	NO	BR	NL	PL	RU	JP	IT	all
local proportion in user profiles																
Pop	-1.2	-4.2	-41.6*	-40.8*	-31.0*	-42.2*	-31.2*	-53.4*	-47.3*	-48.6*	-44.1*	-54.4*	-57.7*	-67.3*	-58.3*	-22.5*
ItemKNN	-1.5	<u>-15.0*</u>	+5.6	+2.4	-18.3	-20.1	-25.6	+16.0	-1.7	+7.2	-19.6	+28.5*	+10.7	+18.7	+49.6	+2.1
BPR	-2.3	<u>-7.6*</u>	-8.8	-17.9*	<u>-23.7*</u>	-31.3*	-31.3*	+3.4	-39.2*	+3.8	-35.1*	-1.2	-19.1*	-40.3	-5.6	-7.0*
NeuMF	-3.3*	-6.1	<u>-33.6*</u>	<u>-31.4*</u>	-23.1*	<u>-40.2*</u>	-33.8*	<u>-35.0*</u>	-39.9*	<u>-29.5*</u>	<u>-37.3*</u>	<u>-46.7*</u>	<u>-49.0*</u>	<u>-66.5*</u>	<u>-56.6*</u>	<u>-19.5*</u>
MultiVAE	<u>-3.5*</u>	-7.5*	-7.0	-19.0*	-23.4*	-34.8*	-29.8*	-19.0	-38.7*	-0.4	-36.3*	-23.6*	-12.7	-18.5	-53.1*	-10.5*
LightGCN	-3.1	-10.3*	-3.4	-1.7	-21.3*	-3.7	-29.5*	-0.8	-35.4*	-3.5	-31.5*	+12.3	-12.3	+10.0	+48.6*	-3.6
US proportion in user profiles																
Pop	-1.2	+14.9*	+28.8*	+23.8*	+2.3	+30.9*	+8.4*	+44.5*	+17.5*	+17.4*	+16.5*	+34.8*	+36.9*	+28.8*	+24.1*	+19.5*
ItemKNN	-1.5	+8.4*	+2.9	<u>-0.7</u>	-0.4	+7.7	+4.5	<u>-6.3</u>	+1.5	<u>-0.4</u>	+5.7	+3.8	+6.1	<u>-10.8</u>	-1.9	+2.2*
BPR	-2.3	+7.3*	+4.3	+6.0	+0.3	+9.8	+3.2	0.0	+7.0	+2.5	+4.8	+6.9*	+7.0*	-0.4	+3.1	+3.7*
NeuMF	-3.3*	+8.3*	+13.6*	+10.6*	<u>-1.2</u>	+15.4*	+2.9	+15.6*	+6.5	+9.7*	+6.7	+14.2*	+11.8*	+17.3*	+13.0*	+7.6*
MultiVAE	<u>-3.5*</u>	+5.2	+2.4	+3.4	+0.1	+6.0	+2.6	+4.4	+5.4	+2.8	+4.4	+6.4*	+3.0	+0.7	+11.0*	+2.6*
LightGCN	-3.1	+5.2	<u>+0.7</u>	+0.2	-0.9	<u>+3.0</u>	+2.7	-2.6	<u>+0.4</u>	+1.9	<u>+3.4</u>	<u>+3.2</u>	<u>+2.9</u>	-5.6	<u>-8.2</u>	<u>+0.8</u>

towards higher ranking items. In the real world scenario, users may consume more or fewer items per recommendation and are also able to consume items that were not recommended to them. In addition, every user is likely to have their own unique choice model(s). These discrepancies may affect the way the effects we observe in this paper would manifest for actual users. We see the study at hand as a reference point for further investigations considering alternative user choice models.

Confounding factors. We conduct our experiments on a sample of the LFM-2b dataset collected from the Last.fm platform. Naturally, depending on how popular the platform is in different countries, they are represented in the data sample to a different extent in terms of both users and items. While, to an extent, representative of the real world, this, in combination with other factors, does not allow us to identify key factors affecting country representation in the recommendations. One approach to address this would be to conduct an “ablation study”, a series of experiments on sub-samples balanced in terms of various attributes.

6 CONCLUSION AND FUTURE WORK

The work at hand shows that in the context of a feedback loop, recommender systems tend to over-represent US and under-represent local music in their recommendations, potentially causing a shift in user preferences in the long run. We conclude that users whose countries are less represented in the data (e.g., Finland) are more likely to receive recommendations inconsistent with their original

preferences in terms of artist country representation. However, higher representation does not always guarantee the opposite (e.g., Germany). We also find that NeuMF affects the proportions of US and local music the most. At the same time, ItemKNN provides least calibrated recommendations in terms of country representation while showing the highest calibration in terms of popularity. This suggests that country calibration does not necessarily follow popularity calibration.

Future work. We plan to investigate the connection between country-calibrated recommendations and popularity-calibrated recommendations. Other potential directions for future work include investigating alternative choice models and additional recommendation scenarios (repeated item consumption, sequential recommendation), defining causes of inconsistencies for every particular country, exploring various mitigation strategies as well as investigating the language (of a track and spoken by the user) and other cultural aspects of country representation.

ACKNOWLEDGMENTS

This research was funded in whole or in part by the Austrian Science Fund (FWF): <https://doi.org/10.55776/P33526>, <https://doi.org/10.55776/DFH23>, <https://doi.org/10.55776/COE12>, <https://doi.org/10.55776/P36413>. This research is funded by the Know-Center within the COMET – Competence Centers for Excellent Technologies Programme, funded by bmvit, bmdw, FFG, and SFG.

REFERENCES

- [1] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, Bamshad Mobasher, and Edward C. Malthouse. 2021. User-centered Evaluation of Popularity Bias in Recommender Systems. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization, UMAP 2021, Utrecht, The Netherlands, June, 21-25, 2021*. 119–129.
- [2] Allison JB Chaney, Brandon M Stewart, and Barbara E Engelhardt. 2018. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM conference on recommender systems*. 224–232.
- [3] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, David Sculley, and Yoni Halpern. 2020. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 525–534.
- [4] Mukund Deshpande and George Karypis. 2004. Item-based top-N recommendation algorithms. *ACM Transactions on Information Systems - TOIS* 22, 1 (jan 2004), 143–177. <https://doi.org/10.1145/963770.963776>
- [5] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, YongDong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, China) (SIGIR '20)*. Association for Computing Machinery, New York, NY, USA, 639–648. <https://doi.org/10.1145/3397271.3401063>
- [6] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web (Perth, Australia) (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. <https://doi.org/10.1145/3038912.3052569>
- [7] Dominik Kowald and Emanuel Lacic. 2022. Popularity bias in collaborative filtering-based multimedia recommender systems. In *International Workshop on Algorithmic Bias in Search and Recommendation*. Springer, 1–11.
- [8] Dominik Kowald, Markus Schedl, and Elisabeth Lex. 2020. The Unfairness of Popularity Bias in Music Recommendation: A Reproducibility Study. In *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 12036)*. 35–42.
- [9] Emanuel Lacic, Leon Fadljevic, Franz Weissenboeck, Stefanie Lindstaedt, and Dominik Kowald. 2022. What drives readership? an online study on user interface types and popularity bias mitigation in news article recommendations. In *European Conference on Information Retrieval*. Springer, 172–179.
- [10] Oleg Lesota, Stefan Brandl, Matthias Wenzel, Alessandro B. Melchiorre, Elisabeth Lex, Navid Rekabsaz, and Markus Schedl. 2022. Exploring Cross-group Discrepancies in Calibrated Popularity for Accuracy/Fairness Trade-off Optimization. In *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)*, Himan Abdollahpouri, Shaghayegh Sahebi, Mehdi Elahi, Masoud Mansoury, Babak Loni, Zahra Nazari, and Maria Dimakopoulou (Eds.). CEUR-WS.org. <https://ceur-ws.org/Vol-3268/paper5.pdf>
- [11] Oleg Lesota, Alessandro B. Melchiorre, Navid Rekabsaz, Stefan Brandl, Dominik Kowald, Elisabeth Lex, and Markus Schedl. 2021. Analyzing Item Popularity Bias of Music Recommender Systems: Are Different Genders Equally Affected?. In *RecSys '21: Fifteenth ACM Conference on Recommender Systems, Amsterdam, The Netherlands, 27 September 2021 - 1 October 2021*. ACM, 601–606. <https://doi.org/10.1145/3460231.3478843>
- [12] Oleg Lesota, Emilia Parada-Cabaleiro, Stefan Brandl, Elisabeth Lex, Navid Rekabsaz, and Markus Schedl. 2022. Traces of Globalization in Online Music Consumption Patterns and Results of Recommendation Algorithms. In *Proceedings of the 23rd International Society for Music Information Retrieval Conference, ISMIR 2022, Bengaluru, India, December 4-8, 2022*, Preeti Rao, Hema A. Murthy, Ajay Srinivasamurthy, Rachel M. Bittner, Rafael Caro Repetto, Masataka Goto, Xavier Serra, and Marius Miron (Eds.). 291–297. <https://archives.ismir.net/ismir2022/paper/000034.pdf>
- [13] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback Loop and Bias Amplification in Recommender Systems. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, Mathieu d'Aquin, Stefan Dietze, Claudia Hauff, Edward Curry, and Philippe Cudré-Mauroux (Eds.). ACM, 2145–2148. <https://doi.org/10.1145/3340531.3412152>
- [14] Alessandro B. Melchiorre, Eva Zangerle, and Markus Schedl. 2020. Personality Bias of Music Recommendation Algorithms. In *RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22-26, 2020*, Rodrygo L. T. Santos, Leandro Balby Marinho, Elizabeth M. Daly, Li Chen, Kim Falk, Noam Koenigstein, and Edleno Silva de Moura (Eds.). ACM, 533–538. <https://doi.org/10.1145/3383313.3412223>
- [15] Peter Muellner, Dominik Kowald, and Elisabeth Lex. 2021. Robustness of meta matrix factorization against strict privacy constraints. In *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28-April 1, 2021, Proceedings, Part II* 43. Springer, 107–119.
- [16] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (Montreal, Quebec, Canada) (UAI '09)*. AUAI Press, Arlington, Virginia, USA, 452–461.
- [17] Markus Schedl, Stefan Brandl, Oleg Lesota, Emilia Parada-Cabaleiro, David Penz, and Navid Rekabsaz. 2022. LFM-2b: A Dataset of Enriched Music Listening Events for Recommender Systems Research and Fairness Analysis. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval, Regensburg, Germany, March 14 - 18, 2022*. 337–341.
- [18] Sebastian Scher, Simone Kopeinik, Andreas Trügler, and Dominik Kowald. 2023. Modelling the long-term fairness dynamics of data-driven targeted help on job seekers. *Scientific Reports* 13, 1 (2023), 1727.
- [19] Harald Steck. 2018. Calibrated recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018*, Sole Pera, Michael D. Ekstrand, Xavier Amatriain, and John O'Donovan (Eds.). ACM, 154–162. <https://doi.org/10.1145/3240323.3240372>
- [20] Jie Xu, Yazhou Ren, Huayi Tang, Xiaorong Pu, Xiaofeng Zhu, Ming Zeng, and Lifang He. 2021. Multi-VAE: Learning Disentangled View-common and View-peculiar Visual Representations for Multi-view Clustering. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 9214–9223. <https://doi.org/10.1109/ICCV48922.2021.00910>
- [21] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Yushuo Chen, Xingyu Pan, Kaiyuan Li, Yujie Lu, Hui Wang, Changxin Tian, Yingqian Min, Zhichao Feng, Xinyan Fan, Xu Chen, Pengfei Wang, Wendi Ji, Yaliang Li, Xiaoling Wang, and Ji-Rong Wen. 2021. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (Virtual Event, Queensland, Australia) (CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 4653–4664. <https://doi.org/10.1145/3459637.3482016>