



HAL
open science

Do Recommender Systems Promote Local Music? A Reproducibility Study Using Music Streaming Data

Kristina Matrosova, Lilian Marey, Guillaume Salha-Galvan, Thomas Louail, Olivier Bodini, Manuel Moussallam

► **To cite this version:**

Kristina Matrosova, Lilian Marey, Guillaume Salha-Galvan, Thomas Louail, Olivier Bodini, et al.. Do Recommender Systems Promote Local Music? A Reproducibility Study Using Music Streaming Data. RecSys '24: 18th ACM Conference on Recommender Systems, ACM, Oct 2024, Bari (IT), Italy. pp.148-157, <10.1145/3640457.3688065>. <hal-04748808v2>

HAL Id: hal-04748808

<https://hal.science/hal-04748808v2>

Submitted on 28 May 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-SA 4.0 - Attribution - ShareAlike - International License



Do Recommender Systems Promote Local Music? A Reproducibility Study Using Music Streaming Data

Kristina Matrosova*
CNRS, Geographie-Cités
France
LIPN, USPN
France

Lilian Marey
LTCI, Télécom Paris
France
Deezer Research
France

Guillaume Salha-Galvan
Deezer Research
France

Thomas Louail
CNRS, Geographie-Cités
France
PACTE, CNRS, Sciences Po Grenoble
France

Olivier Bodini
LIPN, USPN
France

Manuel Moussallam
Deezer Research
France

ABSTRACT

This paper examines the influence of recommender systems on local music representation, discussing prior findings from an empirical study on the LFM-2b public dataset¹. This prior study argued that different recommender systems exhibit algorithmic biases shifting music consumption either towards or against local content. However, LFM-2b users do not reflect the diverse audience of music streaming services. To assess the robustness of this study's conclusions, we conduct a comparative analysis using proprietary listening data from a global music streaming service, which we publicly release alongside this paper. We observe significant differences in local music consumption patterns between our dataset and LFM-2b, suggesting that caution should be exercised when drawing conclusions on local music based solely on LFM-2b. Moreover, we show that the algorithmic biases exhibited in the original work vary in our dataset, and that several unexplored model parameters can significantly influence these biases and affect the study's conclusion on both datasets. Finally, we discuss the complexity of accurately labeling local music, emphasizing the risk of misleading conclusions due to unreliable, biased, or incomplete labels. To encourage further research and ensure reproducibility, we have publicly shared our dataset and code.

CCS CONCEPTS

• **Information systems** → *Recommender systems; Personalization.*

KEYWORDS

Music Recommendation, Fairness, Algorithmic Bias, Local Music.

¹Previously available at <http://www.cp.jku.at/datasets/LFM-2b/>, the LFM-2b dataset has recently been taken down due to license issues.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys '24, October 14–18, 2024, Bari, Italy

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0505-2/24/10
<https://doi.org/10.1145/3640457.3688065>

ACM Reference Format:

Kristina Matrosova, Lilian Marey, Guillaume Salha-Galvan, Thomas Louail, Olivier Bodini, and Manuel Moussallam. 2024. Do Recommender Systems Promote Local Music? A Reproducibility Study Using Music Streaming Data. In *18th ACM Conference on Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3640457.3688065>

1 INTRODUCTION

Recommender systems are essential for music streaming services like Apple Music, Deezer, and Spotify [9, 26, 38, 39]. They help mitigate information overload problems by showcasing the most relevant content for each user, within large catalogs of millions of songs, albums, and artists [7, 19, 23, 32, 39]. They also assist users in discovering new music they might like on these services [6, 8, 19]. With the rise of streaming as the predominant form of music consumption [25, 34], there has been, however, a noticeable increase in debates about the responsibilities of these systems. Concerns are also growing about their ability to promote a fair and diverse musical landscape and the various biases they might introduce or amplify when recommending music [2, 13, 14, 15, 17, 18, 29, 30, 40].

In particular, Lesota et al. [31] recently argued that some music recommender systems might intensify the predominance of US music consumption in other countries. Specifically, in an empirical study focused on the LFM-2b dataset of listening actions on Last.fm [37], these authors investigated the extent to which standard recommender systems favor US-produced content over local music from the country of origin of each user. Their findings suggest that NeuMF [24], a neural network-based collaborative filtering algorithm, recommends lower proportions of local music than what users from each country actually listen to. In other words, NeuMF exhibits an *algorithmic bias* [15, 22] against local music on LFM-2b. On the contrary, the more classical ItemKNN [12] method yields more calibrated recommendations, and even fosters the consumption of local music in most countries under consideration.

This study undoubtedly raised essential issues regarding the uneven impact of recommender systems on local music consumption. Nonetheless, as big as it might be, the LFM-2b dataset used for evaluation is of a particular nature. Last.fm users tend to be active on the Internet and social media, and are not evenly distributed across countries. Therefore, they might not represent the full spectrum of

music streaming service users. Moreover, as detailed in Section 2, the study did not analyze how important parameters related to model training impact these biases. For these reasons, it remains unclear whether the conclusions of Lesota et al. [31] would hold in other experimental settings and using a different dataset.

In this paper, we address this question by conducting a comprehensive comparative study using proprietary listening data from the global music streaming service Deezer. Our study concentrates on France, Germany, and Brazil – three countries where Deezer is one of the leading market players. Our contributions are as follows:

- Firstly, we show that the Deezer and LFM-2b datasets present different local music consumption patterns. This discrepancy suggests caution when drawing conclusions about local music representation based solely on one dataset like LFM-2b.
- Secondly, we demonstrate that NeuMF and ItemKNN exhibit different algorithmic biases towards local music on Deezer compared to LFM-2b when following the evaluation setup of Lesota et al. [31]. Importantly, we also uncover several factors that significantly influence these biases – including their magnitude but also their direction – thereby affecting this study’s overall conclusions on both datasets. These factors include the number of tracks each model recommends, their training variability, and whether they were trained on data from individual countries or the entire dataset.
- Thirdly, we explain that accurately labeling local music is a complex endeavor, and that the proportion of local music consumed and recommended can vary significantly depending on the source of the labels, their level of completeness, and the various biases introduced by human annotators. Consequently, we recommend prioritizing the development of comprehensive, transparent, and reliable local data labeling in future research. We argue that this foundational step is crucial for studies aiming to understand local music biases, as results based on unreliable labels may be misleading.
- Lastly, along with this paper, we publicly release our Deezer dataset as well as the source code of our experiments. This release aims to ensure full reproducibility of our results and to facilitate future studies on local music recommendation.

The remainder of this paper is organized as follows. In Section 2, we introduce the problem more formally and review the related work in more detail. In Section 3, we introduce our Deezer dataset and compare it to LFM-2b in terms of local music consumption. We report and discuss results from our empirical study on local music recommendation and biases in Section 4, and conclude in Section 5.

2 PRELIMINARIES

We begin this section by formally presenting the problem under consideration, before reviewing the related work.

2.1 Problem Formulation

2.1.1 Notation. In this paper, we consider a set \mathcal{V} of music tracks available in the catalog of a music streaming service, and a set \mathcal{U} of $M \in \mathbb{N}^*$ users on this same service. We denote by $N_{\text{listened}}(u)$ the number of streams performed by each user u over a predefined time period. Moreover, we denote by $N_{\text{local}}(u) \in \{0, \dots, N_{\text{listened}}(u)\}$ the number of these streams that are of music tracks from the

country of origin of u according to some data labeling². We refer to $N_{\text{local}}(u)$ as the number of *local* streams of u . Using this formalism, the proportion of local music listened to by u is:

$$L(u) = \frac{N_{\text{local}}(u)}{N_{\text{listened}}(u)} \in [0, 1]. \quad (1)$$

Additionally, we consider a music recommender system:

$$\text{MRS}_K: \mathcal{U} \rightarrow \mathcal{V}. \quad (2)$$

MRS_K recommends³ K music tracks from \mathcal{V} to each user of the music streaming service, for some fixed value $K < N$. The number of local music tracks recommended to the user u by MRS_K among these K tracks is $N_{\text{local}, \text{MRS}_K}(u) \in \{0, \dots, K\}$. The proportion of local music tracks recommended to u by MRS_K is:

$$L_{\text{MRS}_K}(u) = \frac{N_{\text{local}, \text{MRS}_K}(u)}{K} \in [0, 1]. \quad (3)$$

2.1.2 Objective. Our main goal in this paper is to investigate the impact of MRS_K on local music representation. In line with Lesota et al. [31], our main indicator of interest will be the *algorithmic bias* of MRS_K in favor or against local music, defined as follows:

$$\text{Bias}_{\text{MRS}_K} = \frac{1}{M} \sum_{u \in \mathcal{U}} (L_{\text{MRS}_K}(u) - L(u)), \quad (4)$$

with $\text{Bias}_{\text{MRS}_K} \in [-1, 1]$. In essence, a positive bias (respectively, a negative bias) indicates that, on average, MRS_K recommends more local music (resp., less local music) than what users of the music streaming service organically listen to. The remainder of this paper will analyze this value for different datasets, recommender systems, and settings. We will aim to uncover the various factors that might influence the intensity or even the direction of this bias.

2.2 The study of Lesota et al. [31] on LFM-2b

Analyzing local music algorithmic biases was one of the key objectives of Lesota et al. [31], with a particular emphasis on the predominance of US music consumption in other countries.

2.2.1 Context. Over the past decades, US music has dominated the global music industry, with its cultural influence spreading worldwide. Trends from the US have been widely adopted even in local music productions [20], and the ratio of US music on radio charts has been rapidly increasing⁴ since the 1960s [1]. The consequences of this dominance are mixed [1, 10, 11]. On the one hand, it can potentially stimulate local cultural development through the adaptation to global trends and reinforcement of local identity, a process known as *glocalization* [1, 10]. On the other hand, it is also sometimes perceived as a threat, termed *cultural imperialism*, which could lead to the decline of local cultures [11, 33, 42]. While music has become more centralized with the rise of music streaming services [5, 31], at the time of Lesota et al.’s study [31], limited research

²We note that associating music tracks with specific countries can be an ambiguous task, and that different data labeling rules may yield inconsistent results. The importance of the labeling source will be pointed out and further discussed throughout this paper.

³At this stage, we do not formulate assumptions regarding the specific data or paradigm (e.g., a collaborative filtering or content-based approach [7]) used to recommend tracks.

⁴However, this growth slowed down from the 1990s due to several factors, including the emergence of CDs, which made music production more accessible worldwide, content localization by MTV, and the introduction of laws in countries like France, imposing local music quotas on radio station programming [1].

had focused specifically on the impact of these services and their recommender systems on local music consumption.

2.2.2 Results. In 2022, Lesota et al. [31] published results from their empirical study conducted on a subset⁵ of the LFM-2b public dataset, which includes listening events from users of the Last.fm music website [37]. This study explored the prominence of US cultural imperialism in online music consumption, revealing that while the US maintains a strong position among Last.fm users, its influence varies significantly across countries. The authors also observed varying glocalization patterns depending on countries. The final part of their study, which our reproducibility paper focuses on, investigated whether recommender systems can increase existing predominances and, overall, shift music consumption towards specific countries at the expense of local content. The authors answered this question positively, explaining that the influence is uneven and algorithm-dependent. Their experiments suggest that NeuMF [24], a neural network-based collaborative filtering algorithm, recommends lower proportions of local music than what Last.fm users in each country organically listen to. In other words, NeuMF exhibits a negative $\text{Bias}_{\text{MRS}_K}$ local bias, as computed in Equation (4). Conversely, the more traditional ItemKNN [12] method tends to promote the consumption of local music in most countries, i.e., it is associated with a positive $\text{Bias}_{\text{MRS}_K}$ local music bias.

2.2.3 Limitations of Lesota et al. [31] and Motivations of our Work. This study undoubtedly raises important issues regarding the uneven impact of recommender systems on local music representation. It positions itself within a growing body of scientific research focusing on the fairness of music recommender systems and their biases, not only regarding local music but also other aspects, including gender and music genres [14, 15, 17, 18, 30, 40].

Nonetheless, we believe this study also suffers from limitations that motivate our work. From an algorithmic perspective, the authors investigated biases using a single number of recommended tracks, $K = 10$. The effect of varying K , i.e., allowing their systems to recommend different numbers of tracks, on the magnitude or even the direction of biases is unclear. Additionally, systems like the neural network-based NeuMF include randomness in training [24], yet the robustness of the study’s conclusions to this randomness remains unverified. Thirdly, their systems were trained on all users, but using data solely from users of each specific country might alter the findings. Bauer and Schedl [3, 4], for instance, suggest distinguishing between country-specific and global mainstreamness to ensure a realistic representation of musical preferences in different countries and promote less biased recommendations.

Beyond these algorithmic considerations, replicating this study with a different dataset is also worthwhile. Indeed, Last.fm users tend to be active on the internet and social media and are not evenly distributed across countries [37]. Therefore, they might not reflect the diverse audience of music streaming services. Furthermore, Lesota et al. [31] relied on MusicBrainz, an open music encyclopedia [16], to associate artists with country labels. However, MusicBrainz labels may not only be imprecise but are also

missing for some LFM-2b tracks, which were simply excluded by Lesota et al. [31]. Zanger et al. [44] suggest that this exclusion could lead to measurement errors due to label biases. Indeed, less popular artists or those from locations or genres unrepresented among MusicBrainz human annotators might lack more labels⁶, leading to a distorted representation of local music in some countries. These factors raise the question of whether the findings of Lesota et al. [31] would remain valid with actual listening data from a music streaming service, or by using alternative labeling sources. The remainder of this paper will aim to clarify these aspects.

3 COMPARING LFM-2B TO DEEZER DATA

In this section, we present our dataset of listening events from Deezer and subsequently provide a comparative analysis with LFM-2b.

3.1 The Deezer Dataset

3.1.1 Overview. We examine a proprietary dataset from the global music streaming service Deezer, comprising the listening history of 30 000 randomly selected users on this platform in March 2019. The dataset features an equal distribution of users from the three countries of our study – France, Germany, and Brazil – with 10 000 users from each. It includes approximately 4 million streams across over 565 000 distinct music tracks. To maintain consistency with Lesota et al. [31], we did not filter listening events by streaming context. As a result, the dataset includes both organic streams and recommendations, which we later discuss in Section 4.

3.1.2 Local Music Labeling. The country of each user is determined based on their IP address. Determining an artist’s country may sometimes be ambiguous, for example in the case of artists born in a country but who gained fame in another one. To capture this complexity, we consider three different country labels in our work:

- Our dataset includes the main *country of activity* and *country of origin* of each artist, as provided by Deezer when available, and compiled by this service from public and private sources.
- Moreover, we added the publicly accessible country labels from the MusicBrainz open music encyclopedia [16] when available. We recall that these labels were the ones used by Lesota et al. [31] in their original study on LFM-2b.

3.2 Descriptive Analysis

We now provide a descriptive analysis of our Deezer dataset. We compare it to the subset of LFM-2b⁷ of Lesota et al. [31] and additionally restrict this subset to the three countries of interest in this study. The sample contains 254 users from France, 805 users from Germany, and 1064 users from Brazil, for a total of over 3 million listening events in 2018 and 2019, on around 100 000 music tracks.

3.2.1 LFM-2b vs Deezer. Overall, we observe that Deezer and LFM-2b exhibit quite different patterns of local music consumption.

First of all, by plotting the proportion of local streams in both datasets, using MusicBrainz labels and considering only labeled

⁵The entire LFM-2b dataset includes approximately 2 billion listening events over 15 years from about 120 000 users. Lesota et al. [31] analyzed a subset of 14 million interactions from 2018 and 2019, involving 13 000 users in 20 countries selected for having at least 100 users and artists who had collectively created at least 1 000 tracks.

⁶As an illustration, Jul, a French rapper, has no music genre annotation in MusicBrainz, despite being one of the most-streamed artists on Deezer in France during 2018 and 2019.

⁷We use data kindly provided by the authors via private email communications.




Figure 1: Proportion of local streams by country, according to the LFM-2b and Deezer datasets. All values are computed using MusicBrainz labels, by considering labeled tracks only.

Table 1: Top 10 most streamed music tracks by French users, in the LFM-2b (top) and Deezer (bottom) datasets. All reported country labels originate from MusicBrainz.

Dataset	Artist/Band	Title	Country Label	Singing Language	Release Year	Music Genre
LFM-2b	Portishead	Glory Box	GB	EN	1994	Trip Hop
	Radiohead	Karma Police	GB	EN	1997	Alt. Rock
	The Verve	Bitter Sweet Sym.	GB	EN	1997	Alt. Rock
	Franz Ferdinand	Take Me Out	GB	EN	2004	Indie Rock
	a-ha	Take On Me	NL	EN	1985	Synth Pop
	Angèle	Balance ton quoi	BE	FR	2018	Pop
	The xx	Intro	GB	EN	2009	Trip Hop
	4 Non Blondes	What's Up?	US	EN	1993	Alt. Rock
	Metronomy	The Look	GB	EN	2010	Indie Rock
	Wax Tailor	Que Sera	FR	EN	2004	Trip Hop
Deezer	Ninho	Goutte d'eau	FR	FR	2019	Rap
	Angèle	Tout oublier	BE	FR	2018	Pop
	Lady Gaga	Shallow	US	EN	2018	Folk Pop
	Lompeal	Trop beau	FR	FR	2018	Rap/Pop
	David Guetta	Say My Name	FR	EN	2018	EDM
	Ariana Grande	7 rings	US	EN	2019	Pop
	Alonzo	Assurance vie	FR	FR	2019	Rap
	DJ Snake	Taki Taki	FR	EN	2018	EDM
	Kaaris	Gun salute	FR	FR	2019	Rap
	Booba	PGP	FR	FR	2019	Rap

tracks in both cases (Figure 1), we observe a significantly lower rate of local music in the LFM-2b dataset. For Brazil, for instance, LFM-2b exhibits 2.5 times fewer local streams compared to Deezer. Moreover, we report in the first two columns of Figure 2 histograms of the percentage of local streams per user, here again using MusicBrainz labels. We note that the distributions are different between the two datasets. In the Deezer dataset, across all three countries, users exhibit varying patterns: some do not listen to local music at all, while others listen to local music only, with a large spectrum of behaviors in between. Conversely, the LFM-2b dataset shows a stark contrast, with few users listening to a majority of local music.

To go further, we present in Table 1 the top 10 most streamed music tracks in France, in both datasets. The Deezer dataset not only contains a higher proportion of French music (both in terms of artists and lyrics), but also predominantly features recent releases, from the year prior to, or the same year as the streams. The prevalent genres include pop, rap, and electronic music. On

Table 2: Percentages of (i) labeled streams, (ii) local streams (among the labeled streams) and (iii) local streams (among all streams) in the Deezer dataset, by country and label source. A labeled stream corresponds to a stream of a music track associated with a country label. A local stream corresponds to a stream where the user and the artist have the same country label.

Country	Label Source	Labeled Streams	Local Streams Among Labeled	Local Streams Among All
France	Deezer - Activity	76 %	50 %	38 %
	Deezer - Origin	75 %	34 %	26 %
	MusicBrainz	76 %	38 %	29 %
Germany	Deezer - Activity	60 %	40 %	24 %
	Deezer - Origin	62 %	30 %	18 %
	MusicBrainz	69 %	33 %	23 %
Brazil	Deezer - Activity	41 %	48 %	19 %
	Deezer - Origin	36 %	37 %	13 %
	MusicBrainz	38 %	38 %	14 %

the contrary, the top tracks in the LFM-2b dataset predominantly consist of older releases (dating back one or more decades) with English lyrics and genres such as indie, alternative rock, or trip hop, which are more niche. These differences can be attributed to several characteristics of the Last.fm website, upon which the LFM-2b dataset is built [37]. Firstly, Last.fm caters primarily to music enthusiasts who have a strong inclination towards collecting and organizing their music libraries, potentially resulting in a preference for less mainstream music genres. Secondly, this website's users are predominantly English-speaking persons, which introduces a population bias. While the Deezer dataset may appear to be more reflective of realistic music consumption patterns, it is important to acknowledge the possibility of similar biases existing within it, as well as in data from other streaming services. Hence, it is crucial to proceed with caution when asserting the presence of cultural patterns based on such data. Using multiple data sources for cross-validation becomes imperative to ensure the reliability and accuracy of conclusions.

3.2.2 Impact of Label Sources. Table 2 presents the proportions of labeled streams and local streams (among the labeled ones, and among all streams) in the Deezer dataset, according to the three label sources, i.e., Deezer's country of origin and country of activity, as well as MusicBrainz labels. We observe that none of the labeling sources provides complete coverage. Across the three countries considered, between 64% and 24% of the streams remain unlabeled. Streams in different countries exhibit varying levels of label coverage. For example, the artist's country is identified in 75-76% of streams by French users, depending on the label source, while for streams from Brazil, this coverage drops to only 36-41%. Label coverage varies by country depending on the source. For instance, Deezer's activity labels provide the highest coverage for streams from Brazil, but they offer the least coverage for streams from Germany. Furthermore, the proportions of local consumption strongly vary depending on the label source. For instance, considering only labeled streams, only 38% of French users' streams consist of French tracks according to MusicBrainz labels, whereas Deezer's

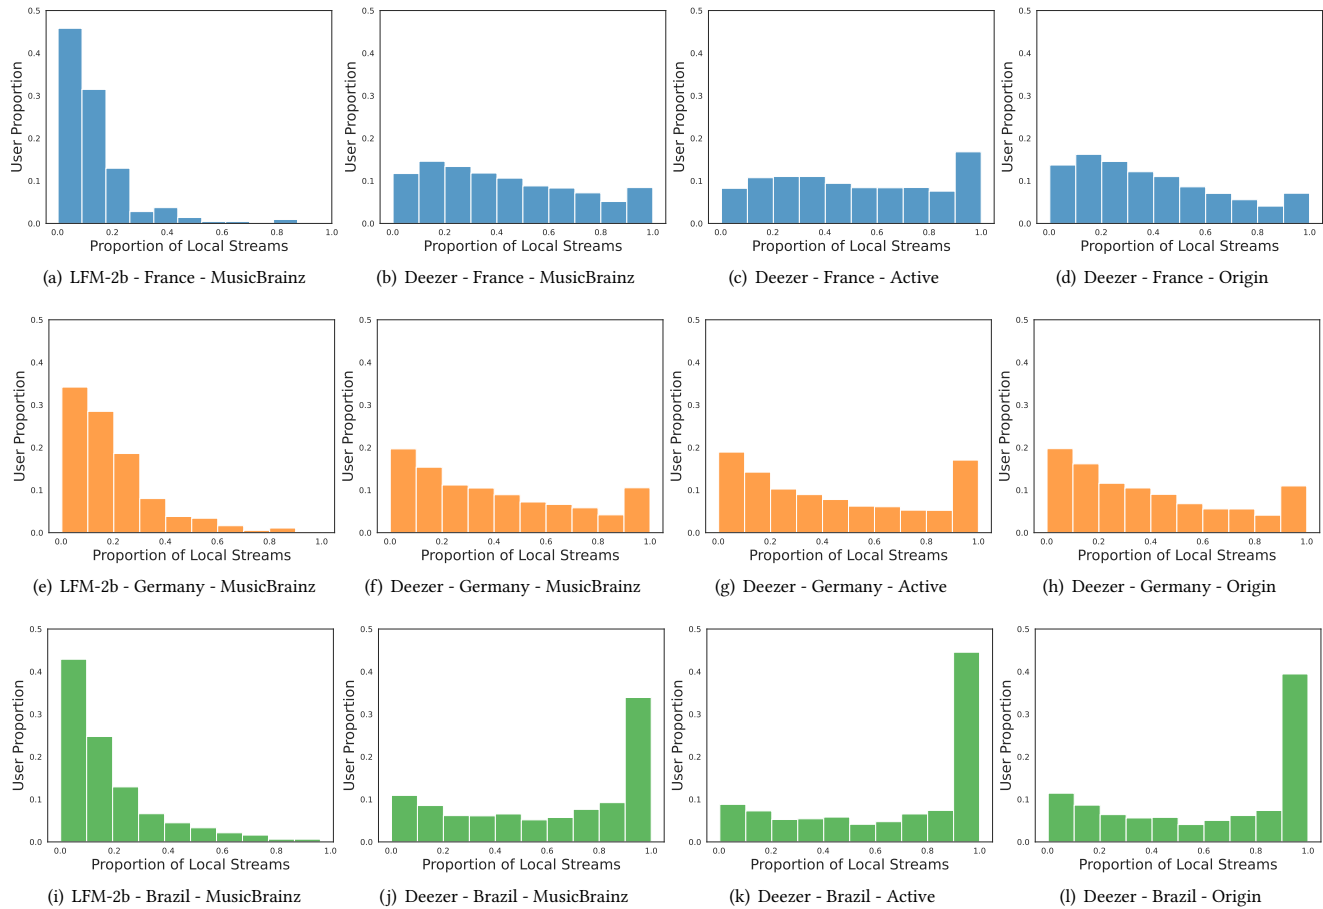


Figure 2: Histograms of the proportion of local streams per user (considering labeled tracks only). Results are split by dataset (i.e., LFM-2b or Deezer), country (i.e., France, Germany, or Brazil), and labeling source (i.e., MusicBrainz labels, Deezer’s country of activity, or Deezer’s country of origin).

activity labels indicate 50%. This difference is notable given that both sources have identical label coverage rates for French streams. Due to incomplete labeling, there’s a significant difference between the proportions of local streams among labeled streams versus all streams. Calculating local streams based solely on labeled data can suggest higher local consumption than what is actually observed across all streams. Moreover, these values aren’t proportional; for instance, when considering only labeled streams, France and Brazil vie for the title of the largest local consumer (depending on the label source). However, when all streams are taken into account, Brazil exhibits the lowest local consumption by a substantial margin.

In summary, obtaining a complete and universally unquestionable labeling of local music proves to be a challenging task. Overall, we advocate against simply filtering out unlabeled tracks, as done in the reference study [31]. Indeed, such an approach may result in removing a majority of the streams from the study, potentially undermining the validity of the study’s conclusion. As outlined in Section 2.2.3, this filtering operation can also introduce label

biases [44], when annotators exhibit preferences for specific countries, music genres, or languages during the annotation process. The extent to which these discrepancies between labels result in inconsistencies regarding the measurement of local music algorithmic biases will be analyzed in Section 4.

4 EXPERIMENTAL ANALYSIS OF LOCAL MUSIC RECOMMENDATION AND BIASES

In this section, we now present our empirical analysis of local music recommendation and biases on the LFM-2b and Deezer datasets. We start by describing the experimental setting. Then, we report and discuss our findings. Notably, we compare them with the main conclusions of the original study of Lesota et al. [31].

4.1 Experimental Setting

4.1.1 Models. We examine the same two collaborative filtering [28] recommender systems analyzed in the reference study:

- NeuMF [24] is a deep learning-based recommender system that integrates traditional matrix factorization (MF) [28] with

neural networks [21]. NeuMF first learns *embedding* vector representations of users and music tracks in the same vector space where proximity reflects similarity, by factorizing a user-track interaction matrix. These embeddings are then processed by a deep neural network architecture trained to predict the most relevant tracks to recommend to each user.

- ItemKNN [12] is a more traditional recommender system based on the nearest neighbor approach [28]. For each user, it assigns similarity scores to tracks by evaluating how similar these tracks are to those the user has previously listened to. This similarity is determined in a collaborative filtering fashion, by analyzing the interactions of other users. When recommending a set of K tracks, ItemKNN selects the top K neighbors with the highest similarity scores. Unlike NeuMF, ItemKNN operates directly on the user-track interaction matrix and does not learn embedding representations.

Lesota et al. [31] trained their ItemKNN and NeuMF models using the entire LFM-2b dataset, including users from all countries. In contrast, our study not only trains these models on the complete LFM-2b and Deezer datasets but also considers country-specific variants. We developed these variants by using only listening data from users within the same country – a realistic setting for a global music streaming service. For instance, to recommend music to a Deezer French user, we would employ the ItemKNN or NeuMF variant trained on Deezer’s listening data from French users.

4.1.2 Task and Implementation Details. For both datasets, we train all models on a top 10 track recommendation task, evaluated using mean reciprocal rank (MRR@10) scores [45] computed on a validation set of 10% randomly selected users, masked during the training phase. Afterwards, we compute the local biases $\text{Bias}_{\text{MRS}_K}$, as defined in Equation (4), for each model in each country, averaged across all users in that country⁸. While in the reference study [31], the authors only reported results for a single number of recommended tracks ($K = 10$), here we consider the more general case of a varying K , with K ranging from 10 to 100 with a step of five tracks.

As in the reference study, we use the implementation of ItemKNN and NeuMF available in RecBole [46], a Python library based on PyTorch [36] that aims to provide a unified framework for developing and reproducing recommendation algorithms. For ItemKNN, we retrieve nearest neighbors to recommend by using cosine similarities computed from the user-track train interaction matrix, with a null value for the shrink parameter [35]. We train all NeuMF models for a maximum of 300 epochs using the Adam optimizer [27], with a learning rate of 0.001, batch sizes of 512 items, a dropout rate of 0.1 [41], and minimizing a binary cross-entropy loss [24]. All NeuMF models learn embedding vectors of dimension 64. For interested readers, we provide exhaustive information on each layer of every neural network in our public GitHub repository (see Section 4.3).

4.2 Results and Discussion

4.2.1 Results on LFM-2b. We begin our analysis with the results obtained on the LFM-2b public dataset. We report in Figure 3 the

local music algorithmic biases of ItemKNN and NeuMF on LFM-2b with MusicBrainz labels, averaged over 20 model runs with standard deviations to assess variability in the training process. Results are split by training variant, i.e., global or country-specific. In particular, Figure 3(a) reports results for the global ItemKNN and NeuMF variants trained on users from all countries, which matches the specific setting of Lesota et al. [31] (with $K = 10$ only in their study). Overall, we reproduce results comparable to those of the original study in this specific setting. NeuMF recommends lower proportions of local music than what users from France, Germany, and Brazil listen to, unveiling negative algorithmic biases. In contrast, ItemKNN tends to foster the consumption of local music in Brazil and Germany, while displaying a negative but relatively small bias in France. Our results are consistent when modifying the number K of recommended tracks.

However, Figure 3(b) reveals that the results change drastically when training ItemKNN and NeuMF in a country-specific fashion, i.e., using data from LFM-2b but selecting users of a single country only, instead of all users. For instance, for $K = 10$, NeuMF now shows a positive bias in Germany, while ItemKNN shows a negative bias in Brazil. Interestingly, increasing the number of recommended tracks K can also reverse the bias direction. For example, NeuMF exhibits a negative bias for $K \in \{10, \dots, 45\}$ but a positive bias for $K \in \{50, \dots, 100\}$. This observation highlights the importance of testing different values of K to draw robust conclusions about the potential biases of each recommender system against local music. Lastly, Figure 3(b) underlines the importance of accounting for variability in the training process, particularly for NeuMF, which shows large ± 1 standard deviation intervals (this variability primarily stems from randomness in the initialization of neural weights, dropout components, and the use of different training splits for each model run). As an illustration, in France and for $K = 10$, NeuMF’s interval overlaps with the “No bias” horizontal dotted line. This emphasizes that NeuMF has shown both positive and negative biases in our experiments, depending on each training instance.

4.2.2 Results on Deezer. We now compare these results with those obtained using the proprietary Deezer dataset, which contains streams from users of the music streaming service Deezer. Figure 4 presents the local music algorithmic biases of ItemKNN and NeuMF on this dataset. Once again, all biases are averaged over 20 model runs with standard deviations. Results are categorized by training variant (i.e., global or country-specific) and label source (i.e., MusicBrainz labels, Deezer’s country of activity, or Deezer’s country of origin). We begin our discussion with an inspection of Figure 4(a), which displays results for the global ItemKNN and NeuMF variants using MusicBrainz labels. This setting is consistent with the one used in the reference study [31] and our earlier Figure 3(a), but applied to the Deezer dataset instead of LFM-2b. We observe that the algorithmic biases exhibited by ItemKNN and NeuMF on LFM-2B vary significantly on Deezer. At $K = 10$, for instance, all models are associated with a positive average bias value, contrary to previous results from Figure 3(a) and Lesota et al. [31] on LFM-2b. We also notice that biases tend to be of higher magnitude on Deezer users. These apparent discrepancies tend to reinforce our discussion from Section 3, highlighting the need for caution when drawing conclusions based solely on one dataset like

⁸We note that one might alternatively compute biases using only users from a test set. However, reporting biases for all users not only aligns with the evaluation protocol of Lesota et al. [31], but also reflects the practical goal of music streaming services, which would typically aim to address local biases across their entire user base.




Figure 3: Local music algorithmic biases of ItemKNN and NeuMF on LFM-2b users in France, Germany, and Brazil, computed for numbers of recommended tracks K varying from 10 to 100 with a step of 5 tracks. Results are split by training variant (“Global” models are trained using listening data from users of all countries, while “Local” models are trained using only listening data from users of the same country). All values are averaged over 20 model runs and reported with ± 1 standard deviation intervals. Values above (respectively, under) the “No bias” 0-level horizontal dotted line indicate that the model exhibits a positive (resp., a negative) algorithmic bias towards local music.

LFM-2b. While this public dataset serves as a convenient starting point for researchers, the observed biases may not be consistent across different datasets with varying listening patterns.

Figure 4(b) reports comparable experiments using our local variants of ItemKNN and NeuMF with MusicBrainz labels. Below, Figure 4(c) and Figure 4(d) present results for global and local models, respectively, but using Deezer’s country of activity labels instead of MusicBrainz labels. Finally, Figure 4(e) and Figure 4(f) show results for global and local models, respectively, but using Deezer’s country of origin. Overall, these figures confirm our previous insights from Section 4.2.1. Training models with data coming from one country only can significantly alter local music biases in both magnitude and direction. Additionally, changing the number of recommended music tracks K and using different user splits or weight initializations can also affect these biases. While these algorithmic factors were not examined in the original work [31], our experiments reveal that they can change the global picture and the study’s conclusion. Properly accounting for these factors is, therefore, crucial to ensure a robust and reliable analysis of local music algorithmic biases.

Figure 4 also highlights that changing the label sources can also substantially affect conclusions. Section 3 had already uncovered that the proportions of local music consumed by users may strongly depend on the label source. Figure 4 further demonstrates that this variation leads to inconsistencies in the measurement of local music algorithmic biases. For instance, the global and local NeuMF models are consistently associated with a negative bias in Germany when relying upon the (country of) “Activity” label (Figure 4(c) and Figure 4(d)), but on the opposite they are associated with a positive bias when using the (country of) “Origin” label (Figure 4(e) and Figure 4(f)). While changing the label source drastically impacts the results, we acknowledge that, nonetheless, some findings remain robust to these changes. For example, Brazil is consistently associated with the highest positive biases across almost all settings in Figure 4. We hypothesize that this consistent behavior may be due to the higher number of Brazilian Deezer users who listen exclusively

or almost exclusively to local music, according to all three label sources (see Figure 2). This aspect might be reflected in our models, although further analysis would be required for confirmation.

4.2.3 Limitations and Future Work. In concluding our discussion, we acknowledge some limitations of our experimental analysis, which also offer opportunities for future research. Firstly, as explained in Section 3.1.1, our Deezer dataset includes both organic streams and recommended streams, to maintain consistency with the original study. However, focusing solely on organic streams for model training could be worthwhile. As Lesota et al. [31] have also pointed out, incorporating recommended tracks may distort the model’s insights about user preferences. Examining the impact of this adjustment on biases towards local music would be an interesting avenue for further investigation. Secondly, Villermet et al. [43] showed that users are very different when it comes to their use of the different features offered by streaming platforms – including algorithmic recommendation – and that only a minority of them primarily relies on algorithmic recommendation to select the music they listened to on streaming platforms. Thus, reproducing the experiment by considering only users who interact with music recommender systems to a certain degree could provide more relevant findings. Thirdly, while we used a specific definition of bias, its perception might actually be subjective, and individual user interviews may be useful to gain additional insights.

At first glance, one might also want to analyze the local music algorithmic biases of numerous other recommender systems, beyond ItemKNN and NeuMF. However, we believe that a more crucial preliminary step for future work will be to improve the accuracy of local music labels. Our study underscores the challenge of accurately labeling local music, and demonstrates how variations in label sources can substantially affect conclusions regarding local music representation and recommendation. Conducting extensive bias analyses on numerous recommender systems with the current state of labeling – where no single label source covers more than




Figure 4: Local music algorithmic biases of ItemKNN and NeuMF on Deezer users in France, Germany, and Brazil, computed for numbers of recommended tracks K varying from 10 to 100 with a step of 5 tracks. Results are split by training variant (“Global” models are trained against listening data from users of all countries, while “Local” models are trained using only listening data from users of the same country), and by label source (i.e., MusicBrainz labels, Deezer’s country of activity, or Deezer’s country of origin). All values are averaged over 20 model runs and reported with ± 1 standard deviation intervals. Values above (resp. under) the “No bias” 0-level dotted line indicate that the model exhibits a positive (resp. negative) algorithmic bias.

80% of streams, and where labels reflect biases from human annotators [44] – could prove fruitless. Indeed, results obtained from misleading labels could render such analyses unreliable. Consequently, we recommend prioritizing the development of comprehensive and reliable local data labeling in future research. We believe that cross-referencing assumptions across multiple labels remains one of the most reliable practices towards achieving this goal.

4.3 Open-Source Code and Data Release

Along with this paper we release two important resources. Firstly, an anonymized version of our Deezer proprietary dataset, containing 4 million listening events from 30 000 Deezer users in France, Germany, and Brazil, along with all three local music labels from our work⁹. The release of this industrial dataset aims to foster future research activities on music recommender systems and local music consumption analysis. In addition, we are open-sourcing the entire Python source code of our experimental analysis, to ensure the

⁹Dataset available at <https://zenodo.org/records/13309698>.

reproducibility of our results. All materials are publicly available on GitHub¹⁰.

5 CONCLUSION

In conclusion, although LFM-2b [37] is publicly available and serves as a convenient starting point for researchers inspecting music recommender systems, caution should be exercised when drawing conclusions about local music consumption based solely on this dataset. Our paper has emphasized significant differences in local music consumption patterns between LFM-2b and a proprietary dataset comprising the listening history of French, German, and Brazilian users of the music streaming service Deezer.

By replicating Lesota et al. [31]’s investigation of algorithmic biases in recommender systems for local music, we have also demonstrated that the two collaborative filtering models they analyzed, NeuMF [24] and ItemKNN [12], display varying biases on Deezer compared to LFM-2b. Moreover, we have identified several factors related to model training that had not been examined in this previous work and can significantly influence these biases, thereby modifying the study’s overall conclusions.

Importantly, we have also explained that the proportion of local music consumed and recommended can vary significantly depending on the label source under consideration, its level of completeness, and biases introduced by human annotators [44]. While obtaining complete and universally accepted local music labels proves to be challenging, we have nonetheless recommended to prioritize the research in this direction. We have argued that this foundational labeling step is crucial for studies aiming to understand local music biases, as results based on unreliable labels may be misleading.

Overall, our work highlights the importance of using multiple model settings and data sources for cross-validation, and to ensure robust conclusions regarding the biases of music recommender systems – not only for local music but also potentially for other aspects such as gender and music genres. As a consequence, we have decided to publicly release our Deezer dataset along with this paper, including listening logs and labels from all three sources used in our experiments. We hope that this release of industrial resources will foster further research. As discussed in Section 4.2.3, our results come with certain limitations that open up interesting avenues for future analyses. Investigating these future directions would undoubtedly contribute to better measuring and enhancing the fairness of recommender systems on music streaming services.

FUNDING INFORMATION

This paper has been realized in the framework of the ‘RECORDS’ grant (ANR-2019-CE38-0013) funded by the ANR (French National Agency of Research).

REFERENCES

- [1] Peter Achterberg, Johan Heilbron, Dick Houtman, and Stef Aupers. 2011. A Cultural Globalization of Popular Music? American, Dutch, French, and German Popular Music Charts (1965 to 2006). *American Behavioral Scientist* 55, 5 (2011), 589–608.
- [2] Christine Bauer, Marta Kholodylo, and Christine Strauss. 2017. Music Recommender Systems Challenges and Opportunities for Non-Superstar Artists. In *Proceedings of the 30th Bled eConference*. 21–32.
- [3] Christine Bauer and Markus Schedl. 2018. On the Importance of Considering Country-Specific Aspects on the Online-Market: an Example of Music Recommendation Considering Country-Specific Mainstream. In *Proceedings of the 51st Hawaii International Conference on System Sciences*. 3647–3656.
- [4] Christine Bauer and Markus Schedl. 2019. Global and Country-Specific Mainstreamness Measures: Definitions, Analysis, and Usage for Improving Personalized Music Recommendation Systems. *PLoS One* 14, 6 (2019), e0217389.
- [5] Pablo Bello and David Garcia. 2021. Cultural Divergence in Popular Music: the Increasing Diversity of Music Consumption on Spotify across Countries. *Humanities and Social Sciences Communications* 8, 1 (2021), 1–8.
- [6] Walid Bendada, Théo Bontempelli, Mathieu Morlon, Benjamin Chapus, Thibault Cador, Thomas Bouabça, and Guillaume Salha-Galvan. 2023. Track Mix Generation on Music Streaming Services using Transformers. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 112–115.
- [7] Jesús Bobadilla, Fernando Ortega, Antonio Hernandez, and Abraham Gutiérrez. 2013. Recommender Systems Survey. *Knowledge-Based Systems* 46 (2013), 109–132.
- [8] Léa Briand, Théo Bontempelli, Walid Bendada, Mathieu Morlon, François Rigaud, Benjamin Chapus, Thomas Bouabça, and Guillaume Salha-Galvan. 2024. Let’s Get It Started: Fostering the Discoverability of New Releases on Deezer. In *Proceedings of the 46th European Conference on Information Retrieval*. Springer, 286–291.
- [9] Léa Briand, Guillaume Salha-Galvan, Walid Bendada, Mathieu Morlon, and Viet-Anh Tran. 2021. A Semi-Personalized System for User Cold Start Recommendation on Music Streaming Apps. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2601–2609.
- [10] Manuel Castells. 2011. *The Power of Identity*. John Wiley & Sons.
- [11] Diana Crane. 2014. Cultural Globalization and the Dominance of the American Film Industry: Cultural Policies, National Film Industries, and Transnational Film. *International Journal of Cultural Policy* 20, 4 (2014), 365–382.
- [12] Mukund Deshpande and George Karypis. 2004. Item-Based Top-N Recommendation Algorithms. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004), 143–177.
- [13] Karljin Dinissen. 2022. Improving Fairness and Transparency for Artists in Music Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3498–3498.
- [14] Karljin Dinissen and Christine Bauer. 2022. Fairness in Music Recommender Systems: A Stakeholder-Centered Mini Review. *Frontiers in Big Data* 5 (2022), 913608.
- [15] Bora Edizel, Francesco Bonchi, Sara Hajian, André Panisson, and Tamir Tassa. 2020. FaiRecSys: Mitigating Algorithmic Bias in Recommender Systems. *International Journal of Data Science and Analytics* 9 (2020), 197–213.
- [16] MusicBrainz: The Open Music Encyclopedia. 2024. <https://musicbrainz.org/>.
- [17] Andres Ferraro. 2019. Music Cold-Start and Long-Tail Recommendation: Bias in Deep Representations. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 586–590.
- [18] Andres Ferraro, Dmitry Bogdanov, Xavier Serra, and Jason Yoon. 2019. Artist and Style Exposure Bias in Collaborative Filtering Based Music Recommendations. In *ISMIR 2019 Workshop on Designing Human-Centric MIR Systems*.
- [19] Andres Ferraro, Peter Knees, Massimo Quadrana, Tao Ye, and Fabien Gouyon. 2023. MuRS: Music Recommender Systems Workshop. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1227–1230.
- [20] Michael Fuhr. 2015. *Globalization and Popular Music in South Korea: Sounding Out K-Pop*. Routledge.
- [21] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT.
- [22] Sara Hajian, Francesco Bonchi, and Carlos Castillo. 2016. Algorithmic Bias: From Discrimination Discovery to Fairness-Aware Data Mining. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2125–2126.
- [23] Casper Hansen, Christian Hansen, Lucas Maystre, Rishabh Mehrotra, Brian Brost, Federico Tomasi, and Mounia Lalmas. 2020. Contextual and Sequential User Embeddings for Large-Scale Music Recommendation. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 53–62.
- [24] 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*. 173–182.
- [25] R Scott Hiller and Jason M Walter. 2017. The Rise of Streaming Music and Implications for Music Production. *Review of Network Economics* 16, 4 (2017), 351–385.
- [26] Kurt Jacobson, Vidhya Murali, Edward Newett, Brian Whitman, and Romain Yon. 2016. Music Personalization at Spotify. *Proceedings of the 10th ACM Conference on Recommender Systems*, 373–373.
- [27] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations*.
- [28] Yehuda Koren and Robert Bell. 2015. Advances in Collaborative Filtering. *Recommender Systems Handbook* (2015), 77–118.
- [29] Dominik Kowald, Markus Schedl, and Elisabeth Lex. 2020. The Unfairness of Popularity Bias in Music Recommendation: A Reproducibility Study. In *Proceedings*

¹⁰Code available at <https://github.com/kmatrosova/FairnessRecsys2024>

- of the 42nd European Conference on Information Retrieval. Springer, 35–42.
- [30] Oleg Lesota, Alessandro Melchiorre, Navid Rekasaz, Stefan Brandl, Dominik Kowald, Elisabeth Lex, and Markus Schedl. 2021. Analyzing Item Popularity Bias of Music Recommender Systems: are Different Genders Equally Affected?. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 601–606.
- [31] Oleg Lesota, Emilia Parada-Cabaleiro, Stefan Brandl, Elisabeth Lex, Navid Rekasaz, 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*. 291–297.
- [32] Yang Li, Kangbo Liu, Ranjan Satapathy, Suhang Wang, and Erik Cambria. 2024. Recent Developments in Recommender Systems: A Survey. *IEEE Computational Intelligence Magazine* 19, 2 (2024), 78–95.
- [33] David Morley. 2006. *Globalisation and Cultural Imperialism Reconsidered*. Routledge, London and New York.
- [34] International Federation of the Phonographic Industry. 2023. *Engaging with Music*. IFPI Technical Report.
- [35] ItemKNN Page on RecBole Documentation. 2024. https://recbole.io/docs/user_guide/model/general/itemknn.html.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Advances in Neural Information Processing Systems* 32 (2019).
- [37] Markus Schedl, Stefan Brandl, Oleg Lesota, Emilia Parada-Cabaleiro, David Penz, and Navid Rekasaz. 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*. 337–341.
- [38] Markus Schedl, Peter Knees, Brian McFee, and Dmitry Bogdanov. 2021. Music Recommendation Systems: Techniques, Use Cases, and Challenges. In *Recommender Systems Handbook*. Springer, 927–971.
- [39] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. 2018. Current Challenges and Visions in Music Recommender Systems Research. *International Journal of Multimedia Information Retrieval* 7, 2 (2018), 95–116.
- [40] Dougal Shakespeare, Lorenzo Porcaro, Emilia Gómez, and Carlos Castillo. 2020. Exploring Artist Gender Bias in Music Recommendation. In *RecSys 2020 Workshop on the Impact of Recommender Systems*.
- [41] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [42] John Tomlinson. 2001. *Cultural Imperialism: A Critical Introduction*. A&C Black.
- [43] Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, and Camille Roth. 2021. Follow the Guides: Disentangling Human and Algorithmic Curation in Online Music Consumption. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 380–389.
- [44] Michael Zanger-Tishler, Julian Nyarko, and Sharad Goel. 2024. Risk Scores, Label Bias, and Everything but the Kitchen Sink. *Science Advances* 10, 13 (2024), eadi8411.
- [45] Eva Zangerle and Christine Bauer. 2022. Evaluating Recommender Systems: Survey and Framework. *ACM Computing Surveys* 55, 8 (2022), 1–38.
- [46] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Yushuo Chen, Xingyu Pan, Kaiyuan Li, Yujie Lu, Hui Wang, Changxin Tian, et al. 2021. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*. 4653–4664.