Journal of Mobile Multimedia, Vol. 0, No. 0 (2005) 000–000

An Analysis of Global and Regional Mainstreaminess for Personalized Music Recommender Systems

MARKUS SCHEDL
Department of Computational Perception, Johannes Kepler University Linz
Altenberger Straße 69, A-4040 Linz, Austria
markus.schedl@jku.at

CHRISTINE BAUER
Department of Computational Perception, Johannes Kepler University Linz
Altenberger Straße 69, A-4040 Linz, Austria
christine.bauer@jku.at

Received (received date)
Revised (revised date)

The *music mainstreaminess of a listener* reflects how strong a person's listening preferences correspond to those of the larger population. Considering that music mainstream may be defined from different perspectives, we show country-specific differences and study how taking into account music mainstreaminess influences the quality of music recommendations.

In this paper, we first propose 11 novel mainstreaminess measures characterizing music listeners, considering both a global and a country-specific basis for mainstreaminess. To this end, we model *preference profiles* (as a vector over artists) for users, countries, and globally, incorporating artist frequency, listener frequency, and a newly proposed TF-IDF-inspired weighting function, which we call artist frequency–inverse listener frequency (AF-ILF). The resulting preference profile for each user u is then related to the respective country-specific and global preference profile using fraction-based approaches, symmetrized Kullback-Leibler divergence, and Kendall's τ rank correlation, in order to quantify u's mainstreaminess. Second, we detail country-specific peculiarities concerning what defines the countries' mainstream and discuss the proposed mainstreaminess definitions. Third, we show that incorporating the proposed global and country-specific mainstreaminess measures into the music recommendation process can notably improve accuracy of rating prediction.

Keywords: music mainstreaminess, music recommender systems, artist frequency-inverse listener frequency, popularity, country-specific differences
Communicated by: to be filled by the Editorial

1 Introduction

In the era of digitalization, music has become easier to access than ever: a tremendous number of musical recordings are readily available to consume on online platforms such as YouTube, Spotify, or iTunes. This opportunity to access a large number of musical works, though, results in information overload [8], which requires new tools to assist users in choosing from the huge amount of musical content" [39]. Music recommender systems (MRS) have, thus, become a significant research topic over the past few years [11, 43, 6] and current online music platforms typically use some sort of MRS.

In general, the idea behind recommender systems is to assist users in searching, sorting, and filtering the vast amount of information available [29]. MRS are specifically built to assist users in navigating through the myriad of available musical recordings and provide them with music suggestions that would fit the respective user's interest or, respectively, automatically generate consecutive recommendations that build a personalized playlist [43]. The challenge is "to propose the right music, to the right user, at the right moment" [24].

Various automatic approaches to music recommendation have been proposed [45]. As summarized in the review by Schedl et al. [45], most MRS rely mainly on some sort of content-based filtering [5] or collaborative filtering [26]. Content-based MRS may, for instance, consider acoustic similarity information on the song level [49], or use the song's music genre, or the performing artist of the music item to quantify similarities [27]. MRS employing collaborative filtering do not require exogenous information about neither users nor music items. Instead, a user is suggested music listened to by users with similar preferences or listening patterns [34].

Another variant, popularity-based recommendation approaches, resemble a primitive form of collaborative filtering, where items are recommended to users based on how popular those items are overall among other users. Such approaches are built on the assumption that the target user is more likely to like a very popular item than one of the far less popular items [11, 44]. Popularity-based recommendation approaches are particularly applicable in hit-driven domains—such as in the music industry. Accordingly, popularity-based MRS approaches are widely adopted to complement other approaches in cold start situations, when there is limited information about new users and/or items available in the system [13, 50].

One approach for considering popularity in the music domain is to describe music listeners "in terms of the degree to which they prefer music items that are currently popular or rather ignore such trends" [38]. Harnessing music mainstreaminess in combination with collaborative filtering techniques tends to deliver better results with respect to music recommendation accuracy and rating prediction error than pure collaborative filtering approaches alone [16, 44, 48, 41].

However, a limitation of existing work on quantifying a user's music mainstreaminess is that music mainstream is viewed from a global perspective. There exist regional peculiarities to mainstream, though [7]. For instance, music consumption behavior is affected by culturally influenced music preferences, market regulations, local radio airplay, etc. (e.g., [47, 20, 10, 35]). In other words, regional aspects shape users' music preferences and music consumption behavior. Accordingly, we can assume country-specific differences concerning which artists are popular.

With respect to the music recommendation research domain, the definition of specific measures that can capture a user's mainstreaminess (i) on both, a global and a country-specific level, and (ii) in ways that can easily be operationalized in music recommendation is a new target of research (e.g., [41, 7]). Calling on this, the main contributions of this paper are three-fold: (i) the definition of several novel measures for user mainstreaminess, considering both a global and a regional, country-specific basis, (ii) the illustration of country-specific peculiarities of these mainstreaminess definitions, and (iii) an analysis of the performance of the proposed mainstreaminess measures for personalized music recommendation.

The remainder of the paper is organized as follows. In Section 2, we provide a brief overview over existing work on mainstreaminess and popularity in music recommendation, and introduce the dataset on which we conduct our experiments. We then detail the proposed mainstreaminess measures in Section 3 and provide examples that show their value to distill the regional mainstream, in addition to a global one. In Section 4, we discuss for a few prototype countries the relationship between their regional mainstream in comparison to the global mainstream. Section 5 shows how to exploit the proposed mainstreaminess measures in collaborative filtering recommendation and highlights the additional values of doing so.Eventually, we round off the paper in Section 6 with a conclusion and directions for future research.

2 Conceptual Foundations and Related Work

2.1 Music Popularity and Mainstreaminess

In the context of recommender systems, popularity-based approaches are widely adopted in numerous domains, including music [13, 23, 50], news [51], or product recommendation in electronic commerce in general [1]. Popularity is thereby typically constructed as a general consensus of a group's attitude about entities [23].

While various ways exist to define and measure popularity (for instance, in terms of sales figures, media coverage, etc.), in the field of MRS, music popularity is frequently characterized by using the total playcounts of a music item—i.e., the number of listening events the music item realizes by all listeners in total cf. [11]. With respect to music popularity by using playcounts, the *long tail concept* as described in [2] is specifically applicable to the (online) music industry [12]; on online music platforms there is a concentration of playcounts on the most popular music items (the head), and then there is a long tail of less popular items [11, 9].

A more general concept to popularity concentration is referred to as *mainstream*. Although literature in the field of popular music studies and popular music cultures references to *mainstream* frequently, the term itself remains rather poorly defined, cf. e.g., [4]. According to the Oxford Dictionaries, *mainstream* is defined as "The ideas, attitudes, or activities that are shared by most people and regarded as normal or conventional". Due to the strong connection of the concepts, the terms *mainstream* and *long tail* are often used interchangeably. The mainstream is thereby frequently also referred to with other terms and phrases (e.g., *hits* [11], *the head* [15]) to circumscribe the phenomenon; the overall concept is also called, for instance, the hit-driven paradigm [11], the long-tail concept [11, 2], etc.

In MRS research, the user feature *music mainstreaminess of a user* [16, 44] essentially describes whether and how strong a user's music listening preferences correspond to those of the overall population. While other listening-centric features, for instance, serendipity [52] or novelty [14], are frequently exploited when modeling a user's music consumption behavior and providing music recommendations, music mainstreaminess is a rather new target of research [16, 44, 48]. Thereby, the mainstreaminess feature is used to analyze a user's ranking of music items and compare it with the overall ranking of artists, albums, or tracks [48].

2.2 Related Work on the Quantification of Music Mainstreaminess

Formal definitions to measure the level of music mainstreaminess of a user are scarce in literature (e.g., [44, 48, 41]). Most existing approaches quantify music mainstreaminess as fractions of the target user's playcounts among the playcounts of the overall population. A limitation of this approach is that it disproportionately privileges the absolute top hits [41], which is problematic for long-tail distributions, which are present for music item popularity on online music platforms. There is a high concentration of demands on the most popular items and a long tail of less popular items. Privileging the top hits leads to low performance of fraction-based user models of mainstreaminess in collaborative filtering approaches [41].

To overcome this limitation, Schedl and Bauer [41] proposed measurement approaches based on *rank-order correlation* and *Kullback-Leibler (KL) divergence*. However, also their work shares with existing fraction-based approaches to quantify mainstreaminess that music mainstream is viewed from a global perspective and does not take regional peculiarities of music mainstream into account.

2.3 Cultural and Regional Aspects Influencing Music Mainstreaminess

As human preferences and behavior are rooted and embodied in culture [22], also music preferences and music consumption behavior are affected by cultural aspects [17, 20, 47]. For instance, music perceptions vary across cultures [25, 30, 46, 47] and music preferences are shaped by cultural aspects [3]. For example, in the European countries, pop music preferences disconverge rather than converge [10].

Still, not only cultural aspects, but also regional (e.g., country-specific) mechanisms affect music consumption; particularly important are national market structures—including distribution channels, legislation, subsidizing, and local radio airplay—that vary across countries [33, 35, 19]. In other words, regional aspects shape users' music preferences and music consumption behavior. Being aware that culture does not equate nation [21, 28], we emphasize that cultural aspects as well as national market structures contribute to users' music consumption preferences and behavior. Accordingly, we can assume country-specific differences concerning the popularity of artists. Against this background, we focus on country-specific differences in the paper at hand.

Closest to our work is the study presented in [48], which analyzes the recommendation performance of mainstreaminess (spelled "mainstreamness") and a user's country, among other features. Our work significantly differs from [48] in various regards: First, we use an open dataset to allow for replication. Second, [48] propose only one global mainstreaminess measure that compares a user's preferences to the overall dataset (global population), while we define mainstreaminess in various ways (based on fractional, divergence, and rank correlation functions) and at various levels (global and country-specific). Third, we also propose a novel weighting approach based on "inverse listening frequency" that highlights artists popular in a specific country, thus, contributing to its mainstream, but not necessarily on a global level.

2.4 Data Preparation

For our experiments, we deploy the LFM-1b dataset [39], which covers 1,088,161,692 listening events of 120,322 unique users, who listened to 32,291,134 unique tracks by 3,190,371 unique artists. The core component of the dataset is the cleaned user-artist-playcount matrix (UAM) containing the number of listening events of 120,175 users to 585,095 unique artists. The distribution of listening events of the Last.fm data corresponds to a typical long-tail distribution [11]. As 65,132 user profiles do not contain any country information, we exclude those from our experiments since they do not contribute to defining a country's mainstreaminess.

3 Formalizing Mainstreaminess

When describing how well a user's listening preferences reflect those of an overall population, e.g., globally or within a country, what is considered *mainstream* depends on the selection of a population; this is a phenomenon which we will also show in our analysis. Consequently, we propose several quantitative measures for user mainstreaminess, both on a global and on a country-specific level, depending on the selection of the population against which the target user is compared. Our approach is inspired by the well-established monotonicity assumptions in text processing and information retrieval [37]: the TF-IDF (term frequency–inverse document frequency) weighting. Based on this assumption, our proposed mainstreaminess measures rely on the concepts of *artist frequency (AF)*, *listener frequency (LF)*, and *artist frequency–inverse listener frequency (AF-ILF)*.

We define AF_{a,U_1} as the sum of the number of tracks by artist a listened to by a set of users U_1. Note that U_1 may be a single user u, all users in a country c, or the entirety of users in the collection (i.e., the global population g). Accordingly, we define LF_{a,U_2} as the number of listeners of artist a within a user population U_2. And we eventually define $AF{\cdot}ILF_{a,U_1,U_2}$ as in Equation 1. We set $AF{\cdot}ILF_{a,U_1,U_2} = 0$ iff $LF_{a,U_2} = 0$.

$$AF{\cdot}ILF_{a,U_1,U_2} = \log\left(1 + AF_{a,U_1}\right) \cdot \log\left(1 + \frac{|U_2|}{LF_{a,U_2}}\right) \quad (1)$$

Note that U_1 and U_2 may represent a single user, all users in the same country, or all users in the dataset (cf. Subsection 2.4). Therefore, this definition allows us to easily formalize both the global and the regional definitions of mainstreaminess, by varying U_1 and U_2. The ILF weighting term can be integrated when computing the *preference profile* for a user or for a country, e.g., $AF{\cdot}ILF_{a,u,c}$, where U_1 contains only the user u and U_2 all users in country c (to which u belongs), or $AF{\cdot}ILF_{a,c,g}$, where U_1 is composed of all users in country c (to which u belongs) and U_2 of all users in the dataset. Using ILF is motivated by the fact that, when determined by $AF_{a,c}$ or $LF_{a,c}$, the top artists in each country c are often identical or very similar to the global top artists (cf. Tables 1, 2, 3, and 4). In order to uncover the respective country-specific mainstream, we therefore use $ILF_{a,g}$ to penalize globally popular artists.

Artist	**AF**	**Artist**	**LF**
The Beatles	2,985,509	Radiohead	24,829
Radiohead	2,579,453	Nirvana	24,249
Pink Floyd	2,351,436	Coldplay	23,714
Metallica	1,970,569	Daft Punk	23,661
Muse	1,896,941	Red Hot Chili Peppers	22,609
Arctic Monkeys	1,803,975	Muse	22,429
Daft Punk	1,787,739	Queen	21,778
Coldplay	1,755,333	The Beatles	21,738
Linkin Park	1,691,122	Pink Floyd	21,129
Red Hot Chili Peppers	1,627,851	David Bowie	20,602

Table 1. Global top artists in the LFM-1b dataset, according to artist frequency (AF) and listener frequency (LF), considering the 53,258 users with country information.

Artist	**AF**
Stam1na	105,633
In Flames	97,645
CMX	90,032
Kotiteollisuus	82,309
Turmion Kätilöt	78,722
Amorphis	78,159
Nightwish	75,742
Mokoma	73,453
Muse	69,507
Metallica	69,499
Artist	**LF**
Metallica	703
Nightwish	695
Muse	693
Daft Punk	675
Queen	671
System of a Down	663
Coldplay	634
Nirvana	614
Pendulum	613
Iron Maiden	609
Artist	**AF-ILF**
St. Hood	70.526
The Sun Sawed in 1/2	67.490
tiko-μ	66.546
Worth the Pain	66.058
Cutdown	65.247
Katariina Hänninen	64.955
Game Music Finland	64.835
Daisuke Ishiwatari	63.565
Altis	63.235
Redrum-187	62.428

Table 2. Top artists for Finland (1,407 users), according to artist frequency (AF), listener frequency (LF), and artist frequency–inverse listener frequency (AF-ILF).

Tables 2, 3, and 4 illustrate the effect of this weighting. It shows the top artists for Finland, Italy, and Turkey, in terms of $AF_{a,c}$, $LF_{a,c}$, and $AF{\cdot}ILF_{a,c,g}$, i.e., AF computed on the country level, ILF on the global level. As can be seen, the AF and even more the LF measures are not suited well to distill the essential mainstream of a country, except maybe for countries such as Finland that show a very specific music taste far away from the global taste [40]. In contrast, AF-ILF is capable of identifying those artists that are popular in a specific country, but not worldwide.

Based on the above definitions, we compute preference profiles globally (PP_g), for a country (PP_c), and for a user (PP_u). Given the LFM-1b dataset [39], these profiles are 585,095-dimensional vectors containing the AF, LF, or AF-ILF scores over all artists in the dataset. Figure 1 provides an example by visualizing the preference profiles for Finland, a country

Artist	**AF**
Radiohead	68,160
The Beatles	65,498
Pink Floyd	60,558
Fabrizio De André	53,928
Muse	48,168
Depeche Mode	42,586
Afterhours	42,473
Verdena	42,338
Sigur Rós	41,748
Arctic Monkeys	39,755
Artist	**LF**
Radiohead	556
Pink Floyd	539
The Beatles	505
David Bowie	500
Muse	500
Nirvana	497
Coldplay	475
The Cure	466
Depeche Mode	459
Daft Punk	457
Artist	**AF-ILF**
CaneSecco	68.451
DSA Commando	66.049
Veronica Marchi	65.864
Train To Roots	65.459
Alessandro Raina	64.228
Machete Empire	63.915
Danti	62.958
Dargen D'Amico	62.453
宝塚歌劇団・宙組	62.228
Aquefrigide	61.663

Table 3. Top artists for Italy (972 users), according to artist frequency (AF), listener frequency (LF), and artist frequency–inverse listener frequency (AF-ILF).

that does particularly not correspond to the global music mainstream. Please note that artist IDs (on the x-axis) are sorted with respect to their *global* popularity in regards to the respective measure (AF, LF, or AF-ILF). As can be seen, while the distributions of the AF- and LF-based preference profiles follow a similar trend, the AL-ILF weighting considerably increases the importance of globally less popular, but country-wise more popular artists (also see Tables 2, 3, and 4).

Exploiting the profiles, we propose three categories of mainstreaminess measures on the user level: fraction-based (F), symmetrized Kullback-Leibler divergence (D), and rank-order correlation according to Kendall's τ (C). The adoption of fraction-based measures is motivated by their easy interpretability (due to the share of overlap between a user's and the global or a country's preference profiles). Kullback-Leibler divergence is a well-established method

Artist	**AF**
Pink Floyd	68,887
Metallica	42,784
Daft Punk	42,020
Iron Maiden	34,174
Radiohead	31,390
Massive Attack	30,669
The Beatles	27,951
Opeth	25,744
Depeche Mode	25,075
Dream Theater	24,286
Artist	**LF**
Pink Floyd	292
Radiohead	289
Metallica	268
Coldplay	261
Nirvana	251
Massive Attack	249
The Beatles	240
Red Hot Chili Peppers	240
Queen	238
Led Zeppelin	236
Artist	**AF-ILF**
Cüneyt Ergün	64.473
Floyd Red Crow Westerman	61.955
Fırat Tanış	58.666
Acil Servis	58.439
Taste (Rory Gallager)	58.366
Mezarkabul	57.799
Rachmaninoff Sergey	57.733
Mabel Matiz	57.619
Grup Yorum	56.855
Yüzyüzeyken Konuşuruz	56.748

Table 4. Top artists for Turkey (479 users), according to artist frequency (AF), listener frequency (LF), and artist frequency–inverse listener frequency (AF-ILF).

to compare distributions (discrete preference profiles in our case). Employing rank-order correlation is motivated by the fact that conversion of feature values to ranks has already been proven successful for music similarity tasks [32].

We provide formulas for the specific measures in Table 5, where $\widehat{X}$ denotes the sum-to-unity normalized vector X and $ranks(PP_U^W)$ represents the real-valued preference profile converted to ranks, i.e. the vector containing all normalized item frequencies of user u, with respect to the frequency weighting approach W (AF or LF). When using $AF \cdot ILF$, $ranks\left(PP_u^W\right)$ is extended to $ranks\left(PP_{u,c}^{AF \cdot ILF}\right)$, i.e. AF computed for user u, ILF on country c, or $ranks\left(PP_{c,g}^{AF \cdot ILF}\right)$, i.e. AF computed on country c, ILF globally. Note that we invert the results of the fraction-based formulations and the symmetrized KL-divergences in order to be consistent in that higher values always indicate closer to the mainstream, while lower

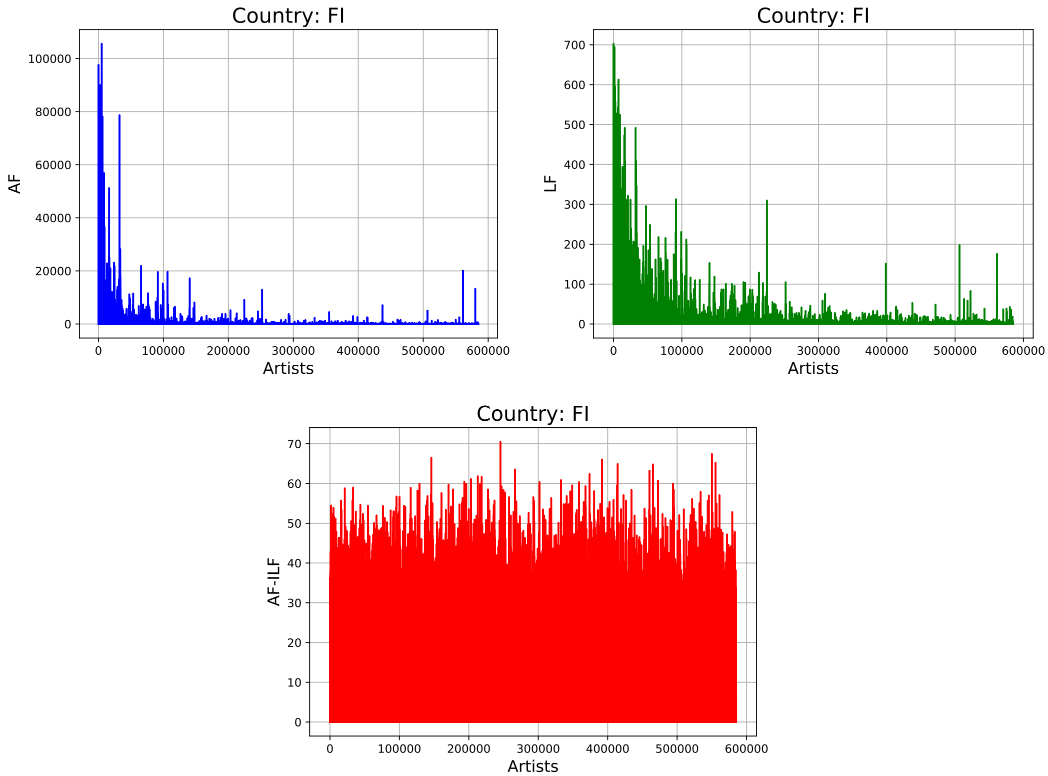

Fig. 1. Artist frequency (AF), listener frequency (LF), and artist frequency–inverse listener frequency (AF-ILF) for Finland. Artist IDs (x-axis) are sorted by global AF, LF, or AF-ILF values, respectively.

ones indicate farther away from the mainstream.

4 Analysis of Global Versus Country-Specific Mainstream

In order to identify archetypal countries for mainstreaminess distributions, we investigate these distributions for the 47 countries in the dataset (cf. Subsection 2.4) that contain at least 100 listeners. Figure 2 illustrates four different examples, showing the country-specific listener frequency for the global top 50,000 artists, for the countries United States (US), Finland (FI), Brazil (BR), and Japan (JP). In all four plots, artists are sorted with respect to their global popularity in decreasing order along the x-axis. The black curve indicates the global trend, adjusted to the listener frequency in the respective country. Looking at the United States, we see that—except for some jitter—the distribution of listener frequencies among artists quite closely follows the global distribution (black curve). For Brazil, and even more for Finland, in contrast, a second trend curve becomes visible, indicating that in addition to the global trend (evidenced by a substantial amount of items along the black curve), certain artists within the countries are much more popular than expected from a global perspective. In Finland and Brazil, these country-specific popular artists follow approximately the same pattern as the global trend curve. In contrast, Japan does not reveal a clear secondary trend curve; there

Abbr.	**Formula**
$F_{g:AF,u:AF}$	$1-\frac{1}{\lvert A\rvert}\cdot\sum_{a\in A}\frac{\lvert\widehat{AF}_{a,u}-\widehat{AF}_{a,g}\rvert}{\max\left(\widehat{AF}_{a,u},\widehat{AF}_{a,g}\right)}$
$F_{g:AF,u:AF\cdot ILF}$	$1-\frac{1}{\lvert A\rvert}\cdot\sum_{a\in A}\frac{\lvert\widehat{AF\cdot ILF}_{a,u,g}-\widehat{AF}_{a,g}\rvert}{\max\left(\widehat{AF\cdot ILF}_{a,u,g},\widehat{AF}_{a,g}\right)}$
$F_{g:AF\cdot ILF,u:AF\cdot ILF}$	$1-\frac{1}{\lvert A\rvert}\cdot\sum_{a\in A}\frac{\lvert\widehat{AF\cdot ILF}_{a,u,g}-\widehat{AF\cdot ILF}_{a,g,g}\rvert}{\max\left(\widehat{AF\cdot ILF}_{a,u,g},\widehat{AF\cdot ILF}_{a,g,g}\right)}$
$F_{c:AF,u:AF}$	$1-\frac{1}{\lvert A\rvert}\cdot\sum_{a\in A}\frac{\lvert\widehat{AF}_{a,u}-\widehat{AF}_{a,c}\rvert}{\max\left(\widehat{AF}_{a,u},\widehat{AF}_{a,c}\right)}$
$F_{c:AF\cdot ILF,u:AF\cdot ILF}$	$1-\frac{1}{\lvert A\rvert}\cdot\sum_{a\in A}\frac{\lvert\widehat{AF\cdot ILF}_{a,u,c}-\widehat{AF\cdot ILF}_{a,c,g}\rvert}{\max\left(\widehat{AF\cdot ILF}_{a,u,c},\widehat{AF\cdot ILF}_{a,c,g}\right)}$
$D_{g:AF,u:AF}$	$\frac{1}{2}\cdot\left(\sum_{a\in A}\widehat{AF}_{a,u}\cdot\log\frac{\widehat{AF}_{a,u}}{\widehat{AF}_{a,g}}+\sum_{a\in A}\widehat{AF}_{a,g}\cdot\log\frac{\widehat{AF}_{a,g}}{\widehat{AF}_{a,u}}\right)^{-1}$
$D_{c:AF,u:AF}$	$\frac{1}{2}\cdot\left(\sum_{a\in A}\widehat{AF}_{a,u}\cdot\log\frac{\widehat{AF}_{a,u}}{\widehat{AF}_{a,c}}+\sum_{a\in A}\widehat{AF}_{a,c}\cdot\log\frac{\widehat{AF}_{a,c}}{\widehat{AF}_{a,u}}\right)^{-1}$
$D_{c:AF\cdot ILF,u:AF\cdot ILF}$	$\frac{1}{2}\cdot\left(\sum_{a\in A}\widehat{AF\cdot ILF}_{a,u,g}\cdot\log\frac{\widehat{AF\cdot ILF}_{a,u,g}}{\widehat{AF\cdot ILF}_{a,c,g}}+\sum_{a\in A}\widehat{AF\cdot ILF}_{a,c,g}\cdot\log\frac{\widehat{AF\cdot ILF}_{a,c,g}}{\widehat{AF\cdot ILF}_{a,u,g}}\right)^{-1}$
$C_{g:AF,u:AF}$	$\tau\left(ranks\left(PP_g^{AF}\right),ranks\left(PP_u^{AF}\right)\right)$
$C_{c:AF,u:AF}$	$\tau\left(ranks\left(PP_c^{AF}\right),ranks\left(PP_u^{AF}\right)\right)$
$C_{c:AF\cdot ILF,u:AF\cdot ILF}$	$\tau\left(ranks\left(PP_{u,c}^{AF\cdot ILF}\right),ranks\left(PP_{c,g}^{AF\cdot ILF}\right)\right)$

Table 5. Proposed music mainstreaminess measures on the user level. Terms denote the following: F stands for the fraction-based approach, D refers to the symmetrized Kullback-Leibler divergence approach, and C is used as abbreviation for the approaches based on rank-order correlation according to Kendall's τ. A is a list of all artists; $\widehat{AF}$ denotes the sum-to-unity normalized AF value; $ranks(PP_u^W)$ represents the real-valued preference profile converted to ranks, i.e. the vector containing all normalized item frequencies of user u, with respect to the frequency weighting approach W (AF or LF); in case of $AF\cdot ILF$, $ranks\left(PP_u^W\right)$ is extended to $ranks\left(PP_{u,c}^{AF\cdot ILF}\right)$, i.e. AF computed for user u, ILF on country c, or $ranks\left(PP_{c,g}^{AF\cdot ILF}\right)$, i.e. AF computed on country c, ILF globally. Note that we invert the values of some measures (F and D) in order to ensure that higher values always indicate closer to the mainstream.

are rather many individual outliers that do not seem to follow a particular pattern.

To quantitatively identify and analyze the country-specific outliers that deviate from the global trend, we next use a sliding window of 5 artists, which we run over the top 1,000 AF, LF, and AF-ILF values of artists, sorted in the same way as in Figure 2, i.e., in decreasing order of global popularity, again for the top 47 countries in the dataset. We compute the mean AF, LF, and AF-ILF value within each window and relate it to the corresponding value of the first artist in the window. If this fraction exceeds a certain threshold, we consider the corresponding artist an outlier. For our experiments that we present in the following, we set that threshold to 100%, meaning that an outlier's value must be at least twice as large as the mean value in its window (in case of a positive outlier); or at most 50% of the value of the

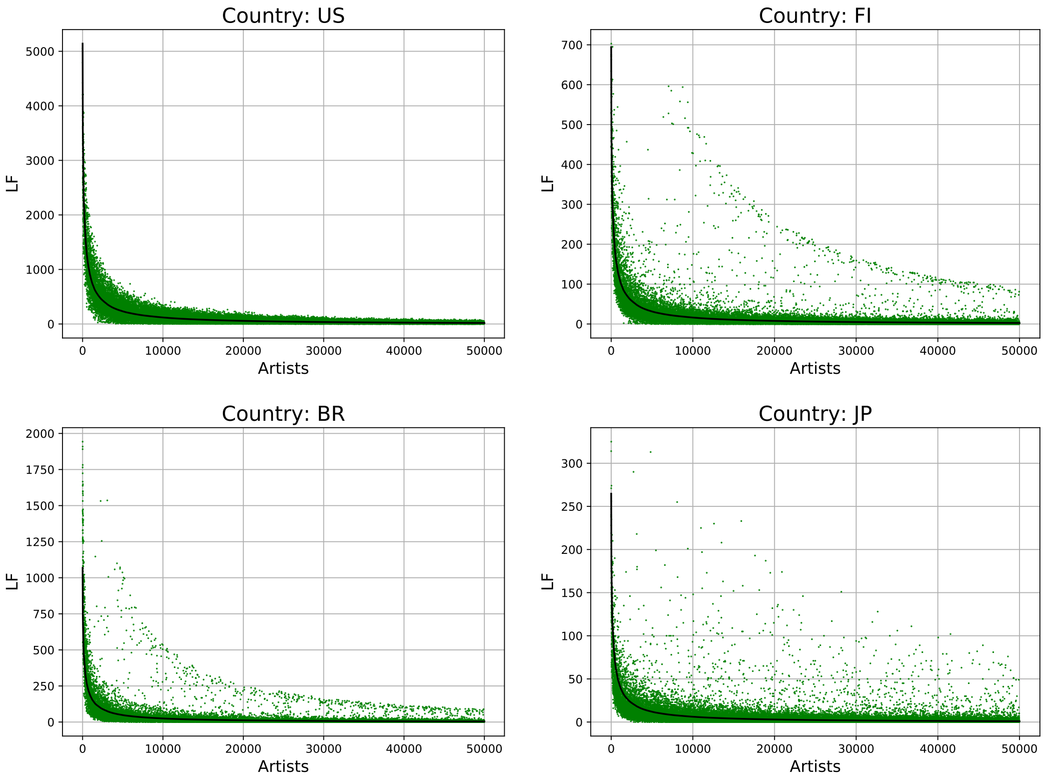

Fig. 2. Country-specific listener frequency (LF) for global top 50,000 artists, for the United States (US), Finland (FI), Brazil (BR), and Japan (JP). In all four plots, artists are sorted with respect to their global popularity in decreasing order. The black curve indicates the global trend, adjusted to the LF in the respective country.

mean value in its window (in case of a negative outlier).

In doing so, we identify country-specific outliers that do not correspond to the global trend, meaning that the identified artists are particularly more (if positive) or particularly less popular in the respective country. Table 6 shows examples of positive AF outliers for Finland. Among the most salient outliers, we find the Finnish metal band “Amorphis”, but also metal bands from neighboring countries such as “Soilwork” from Sweden.

Table 7 shows the top country-specific positive outliers for Germany. The artist with the highest AF difference to the expected AF values in its neighborhood (window) is “Die Ärzte”, a German punk rock band. Also other German bands rank high (e.g., “Rammstein”, “Volbeat”, and “In Extremo”).

To exemplify also negative outliers, Table 8 shows for the United States, the first (highest global position) positive and negative outliers that appear along the trend when using the AF measure. Among the negative outliers, we find mostly hard rock and metal bands, which corroborates previous findings that these genres are underrepresented in the United States compared to the global mean [42].

Artist	Rank	Difference
In Flames	25	+162.74%
Katatonia	73	+112.78%
Amon Amarth	90	+102.17%
Pendulum	99	+124.77%
Children of Bodom	122	+120.17%
Sonata Arctica	134	+146.35%
Bullet for My Valentine	138	+105.89%
HIM	154	+103.20%
Lamb of God	169	+136.27%
Sabaton	195	+168.01%
Amorphis	**203**	**+229.48%**
Infected Mushroom	220	+101.34%
Kamelot	248	+110.62%
Gojira	255	+128.40%
Dimmu Borgir	275	+140.08%
Soilwork	288	+220.73%
Burzum	305	+105.12%
Finntroll	314	+165.20%
Fear Factory	328	+122.30%
Biffy Clyro	365	+140.82%

Table 6. Results of outlier analysis for artist–frequency (AF) values in Finland. The first 20 positive outliers are shown together with their global rank and the difference between their AF values and the mean AF values in a window of size 5, succeeding the artist.

Artist	Rank	Difference
Rammstein	13	+115.87%
Rise Against	59	+128.29%
Mumford & Sons	85	+100.64%
Amon Amarth	90	+122.67%
Enter Shikari	179	+128.08%
Grateful Dead	261	+266.76%
Volbeat	287	+138.91%
3 Doors Down	298	+112.16%
Finntroll	314	+105.71%
Machine Head	325	+115.04%
The Gaslight Anthem	352	+102.57%
Biffy Clyro	365	+142.99%
Flogging Molly	395	+102.68%
Die Ärzte	**437**	**+310.54%**
Simple Plan	462	+158.99%
Heaven Shall Burn	505	+173.12%
La Dispute	541	+132.26%
Emilie Autumn	543	+116.91%
In Extremo	563	+194.80%
Combichrist	565	+121.34%

Table 7. Results of outlier analysis for artist–frequency (AF) values in Germany. The first 20 positive outliers are shown together with their global rank and the difference between their AF values and the mean AF values in a window of size 5, succeeding the artist.

5 Music Recommendation Tailored to User Mainstreaminess

To evaluate the proposed mainstreaminess measures (cf. Section 3) with respect to their ability to improve performance in music recommendation, we conduct rating prediction experiments,

Artist	**Rank**	**Difference**
Radiohead	1	+101.42%
Rammstein	13	-60.13%
Nine Inch Nails	20	+101.68%
Nightwish	23	-54.26%
In Flames	25	-54.56%
AC/DC	36	-53.89%
Korn	39	-53.46%
Marilyn Manson	52	-56.09%
The White Stripes	70	+112.77%
Katatonia	73	-60.63%
Within Temptation	74	-63.20%
30 Seconds to Mars	81	-56.39%
Guns N' Roses	82	-63.45%
Amon Amarth	90	-55.56%
Anathema	97	-54.23%
Avenged Sevenfold	101	-64.63%
Modest Mouse	105	+142.16%
Bring Me the Horizon	106	-54.01%
Limp Bizkit	116	-73.35%
Blur	129	-54.05%

Table 8. Results of outlier analysis for artist–frequency (AF) values in the United States. The first 20 positive and negative outliers are shown together with their global rank and the difference between their AF values and the mean AF values in a window of size 5, succeeding the artist.

which is a common approach to recommender systems evaluation. For this evaluation, we use again the LFM-1b dataset of user-generated listening events from Last.fm [39], as discussed in Subsection 2.4.

5.1 Experimental Setup

While we are aware that a truly user-centric evaluation would be beneficial for this kind of research, conducting a user study on tens of thousands of users (or even only a representative subset of the users) is beyond the scope of this paper. We therefore stick to the common approach of quantifying the performance of a recommender system by conducting a rating prediction task. To this end, we normalize and scale the playcount values in the UAM to the range [0, 1000] for each user individually, assuming that higher numbers of playcounts indicate higher user preference for an artist.

We apply the common singular value decomposition (SVD) method according to [36] to factorize the UAM and in turn effect rating prediction. In 5-fold cross-validation experiments, we use root mean square error (RMSE) and mean absolute error (MAE) as performance measures.

To obtain a baseline, we first run the rating prediction experiment on the global group of 65,132 users and report results of the error measures in the first row of Table 9. To study the influence of both, the different mainstreaminess *definitions* and mainstreaminess *levels* on recommendation performance, we then create subsets of users for each combination of mainstreaminess measure and country with at least 1,000 users.[a] To this end, we split the users

[a]The restriction to countries with at least 1,000 users was made to allow for a meaningful analysis, as performed in [40].

Mainstreaminess	**user set**	**w.RMSE**	**w.MAE**
Baseline (global UAM)		29.105	25.202
$F_{g:AF,u:AF}$	*all*	26.377	24.050
	high	3.714	1.308
	mid	12.574	9.887
	low	14.186	11.625
$F_{g:AF,u:AF \cdot ILF}$	*all*	21.137	18.617
	high	3.681	1.299
	mid	11.035	8.191
	low	14.426	11.868
$F_{g:AF \cdot ILF,u:AF \cdot ILF}$	*all*	19.140	16.769
	high	11.777	9.121
	mid	13.396	10.833
	low	8.708	5.806
$F_{c:AF,u:AF}$	*all*	14.465	11.958
	high	3.723	1.309
	mid	8.681	6.112
	low	12.706	9.952
$F_{c:AF \cdot ILF,u:AF \cdot ILF}$	*all*	17.615	15.301
	high	9.237	6.648
	mid	3.686	1.305
	low	10.122	7.610
$D_{g:AF,u:AF}$	*all*	24.026	21.705
	high	10.561	8.024
	mid	9.854	7.299
	low	5.365	2.909
$D_{c:AF,u:AF}$	*all*	28.021	25.746
	high	5.365	2.912
	mid	13.510	10.840
	low	25.923	22.621
$D_{c:AF \cdot ILF,u:AF \cdot ILF}$	*all*	14.628	11.624
	high	3.656	1.281
	mid	7.035	4.515
	low	8.589	5.670
$C_{g:AF,u:AF}$	*all*	15.906	13.525
	high	3.680	1.291
	mid	7.443	4.472
	low	19.183	16.373
$C_{c:AF,u:AF}$	*all*	14.349	12.032
	high	3.687	1.290
	mid	4.270	1.833
	low	3.692	1.308
$C_{c:AF \cdot ILF,u:AF \cdot ILF}$	*all*	30.827	28.535
	high	7.680	5.187
	mid	4.825	2.340
	low	10.785	8.1084

Table 9. Weighted root mean square error (RMSE) and weighted mean absolute error (MAE) for various mainstreaminess definitions and levels, i.e. user sets. Rating values are scaled to [0, 1000]. Experiments are conducted on the country level (except for first row using the complete UAM with random item selection in each fold, irrespective of country) and error measures are averaged (arithmetic mean) over all countries with more than 1,000 users and weighted by number of users in the respective country. In the individual experiments, *all* refers to the group of all users in each considered country, *low* only to the users in the lower 3-quantile (tertile) w.r.t. the respective mainstreaminess definition, *mid* and *high* defined analogously.

in each country into three (almost) equally sized subsets according to their mainstreaminess value: *low* corresponds to users in the lower 3-quantile (tertile) w.r.t. the respective mainstreaminess definition, *mid* and *high*, respectively, to the mid and upper tertile. In the individual experiments, *all* refers to the group of all users in each considered country, *low* only to the users in the lower 3-quantile (tertile) w.r.t. the respective mainstreaminess definition, *mid* and *high* defined analogously. Further, conducting the same experiment on all users in each country (user set *all*) allows for a comparison of a pure mainstreaminess filtering approach versus a combination of mainstreaminess filtering and demographic (country) filtering.

5.2 Results and Discussion

Table 9 shows the error measures (RMSE and MAE) for different *definitions* and *levels* of mainstreaminess, averaged over all considered countries (cf. Subsection 2.4), RMSE and MAE weighted by the number of users in the respective country. In the following discussion, we concentrate on RMSE since it is more common and considers larger differences between predicted and true ratings disproportionately more severe than smaller ones.

As a general finding, our results show that tailoring the recommendations to a user's mainstreaminess level (*low*, *mid*, *high*) leads to substantial error reductions, irrespective of the applied mainstreaminess measure. More specifically, $C_{c:AF,u:AF}$ outperforms the other measures in four regards: First, it leads to the lowest overall RMSE of 14.349 (*all*). Second, the errors realized by $C_{c:AF,u:AF}$ are also the lowest for each of the three user sets (*low*, *mid*, *high*). If better performance is achieved on a set with another measure, the difference is just in the third position after the decimal point. Third, $C_{c:AF,u:AF}$ performs on each of the three user sets (*low*, *mid*, *high*) in a balanced way (weighted RMSE amounts to respectively 3.692, 4.270, and 3.687), whereas the other mainstreaminess measures yield a rather unbalanced picture since each of them performs on at least one set far worse than on the other(s), e.g., $C_{g:AF,u:AF}$ with 19.183, 7.443, and 3.681, respectively, for *low*, *mid*, and *high*. Fourth, $C_{c:AF,u:AF}$ performs well also on the low mainstreaminess user set (*low*), which is a user segment that is typically difficult to satisfy.

The fraction-based approaches $F_{g:AF,u:AF}$, $F_{c:AF,u:AF}$, and $F_{g:AF,u:AF \cdot ILF}$ have in common that they perform far better in the high mainstreaminess segment than in the mid and the low one. This could indicate that these measures still privilege globally popular items too much and, thus, produce more errors in the mid and low segments.

Interestingly, the approaches based on symmetrized Kullback-Leibler divergence (D) perform worse when tailored towards a user's country ($D_{c:AF,u:AF}$), compared to their application on a global level ($D_{g:AF,u:AF}$). Combining the country-specific tailoring with the AF-ILF weighting allows for better results compared to applying both separately.

While our results do not suggest a general superiority of mainstreaminess measures that incorporate AF-ILF, first results of our deeper analysis on the country level indicate that these measures seem to perform particularly well for countries far from the global mainstream, such as Finland (RMSE of $D_{c:AF \cdot ILF,u:AF \cdot ILF}$ for *all*=5.985, *high*=1.346, *mid*=1.365, *low*=1.418), but worse for high mainstream countries, such as the USA (RMSE of $D_{c:AF \cdot ILF,u:AF \cdot ILF}$ for *all*=57.489, *high*=4.071, *mid*=4.077, *low*=55.968). In the presented example, the low mainstream country Finland is small, and the respective weighted error measures in Table 9 do not reflect this country's users to the same extent as the large and high mainstream United States. As part of our ongoing large-scale analysis, delving into detail on country-specific aspects, we will investigate as a next step what factors influence the performance differences between countries for a given mainstreaminess measure.

A direct comparison of the RMSE achieved by our approach with the RMSE reported in [48], the work closest to ours, is unfortunately impossible since Vigliensoni and Fujinaga quantized playcounts into a 5-point Likert rating scale: [1, 5]. Still, in a rough estimation, our results suggest that the accuracy of our best $C_{c:AF,u:AF}$ approach delivers a new benchmark in the combination of demographic (country) filtering and mainstreaminess filtering, with a RMSE of 14.3 on a [0, 1000] scale. The best RMSE reported in [48] when considering mainstreamness and country information is approximately 0.9 on the much narrower [1, 5] scale (cf. approach *u.c.m.* in Figure 2 of [48]).

6 Conclusions and Outlook

The *music mainstreaminess of a listener* reflects how strong a person's listening preferences correspond to those of the larger population. We consider that music mainstream may be defined from different perspectives. In this paper, we took into account that there are regional differences of what is considered mainstream, due to cultural characteristics and different market structures across countries.

The main contributions of this paper are three-fold: First, we proposed 11 novel measures to quantify the music mainstreaminess of a user, a country, and an entire population. Those are based on fractional (F), divergence (D), and rank correlation (C) functions.

Second, we illustrated country-specific peculiarities of music preferences and *country-specific mainstream* employing the LFM-1b dataset [39]. We identified archetypal countries: (i) those countries where the mainstream of the country corresponds to the global trend (e.g., the United States), (ii) those countries with a distinct country-specific mainstream in addition to the global mainstream (e.g., Finland), and (iii) those countries roughly following the global mainstream trend without a clear secondary trend curve, but showing various country-specific outliers over the whole global artist popularity range (e.g., Brazil and Japan).

Third, we studied the performance of the proposed mainstreaminess measures for personalized music recommendation. Considering that music mainstream may be defined from a global but also a country-specific perspective, we particularly studied how the combination of a user's mainstreaminess and demographic (country) filtering influences the quality of music recommendations. Based on the LFM-1b dataset [39], we investigated the performance of the proposed measures in a rating prediction task, employing probabilistic matrix factorization. To quantify performance, we computed country-averaged, weighted RMSE and MAE figures for all mainstreaminess definitions and various mainstreaminess levels, and compared these with a global baseline. Overall, our results suggest that incorporating any kind of mainstreaminess information outperforms the baseline. Our best approach combines demographic filtering (based on a user profile's country) and mainstreaminess filtering based on Kendall's τ (variant $C_{c:AF,u:AF}$) and outperforms applying these filtering approaches separately. While our results do not hint at a general superiority of mainstreaminess measures that incorporate AF-ILF, they do show that such measures perform much better than others for countries whose preference profiles are far away from the global taste (e.g., Finland).

As part of future work, we will take an in-depth look at the differences between countries, i.e. analyze in which countries which mainstreaminess functions perform particularly well or poorly. Additionally, we plan to analyze how well our results generalize to other datasets providing demographic user information, e.g., the Million Musical Tweets Dataset [18], a playlist dataset crawled from Spotify users [31], or on a larger scale Spotify's official Million Playlist Dataset,[b]released as part of the ACM Recommender Systems Challenge 2018 on automatic playlist continuation. We further plan user studies to investigate with qualitative methods whether incorporating mainstreaminess information improves users' perceived satisfaction with recommendations.

7 Acknowledgments

This research is supported by the Austrian Science Fund (FWF): V579.

References

1. Hyung Jun Ahn. Utilizing popularity characteristics for product recommendation. *International Journal of Electronic Commerce*, 11(2):59–80, 2006.
2. Christian Anderson. *The long tail: Why the future of business is selling less of more*. Hachette Books, New York, NY, USA, 2006.
3. Young Min Baek. Relationship between cultural distance and cross-cultural music video consumption on youtube. *Social Science Computer Review*, 33(6):730–748, 2015.
4. Sarah Baker, Andy Bennett, and Jodie Taylor. *Redefining mainstream popular music*. Routledge, 2013.
5. Chumki Basu, Haym Hirsh, and William Cohen. Recommendation as classification: Using social and content-based information in recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 714–720. American Association Intelligence, 1998.
6. Christine Bauer, Marta Kholodylo, and Christine Strauss. Music recommender systems: Challenges and opportunities for non-superstar artists. In Andreja Pucihar, Mirjana Kljajić Borštnar, Christian Kittl, Pascal Ravesteijn, Roger Clarke, and Roger Bons, editors, *Proceedings of 30th Bled eConference*, pages 21–32, Bled, Slovenia, 2017. University of Maribor Press.
7. Christine Bauer and Markus Schedl. On the importance of considering country-specific aspects on the online-market: An example of music recommendation considering country-specific mainstream. In *51st Hawaii International Conference on System Sciences (HICSS 2018)*, pages 3647–3656.
8. David Bawden and Lyn Robinson. The dark side of information: overload, anxiety and other paradoxes and pathologies. *Journal of Information Science*, 35(2):180–191, 2009.
9. Erik Brynjolfsson, Yu Hu, and Duncan Simester. Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, 57(8):1373–1386, 2011.
10. Oliver Budzinski and Julia Pannicke. Do preferences for pop music converge across countries?:empirical evidence from the eurovision song contest. *Creative Industries Journal*, pages 1–20, 2017.
11. Òscar Celma. *Music Recommendation and Discovery*. Springer, Berlin, Heidelberg, Germany, 2010.
12. Òscar Celma and Pedro Cano. From hits to niches?: Or how popular artists can bias music recommendation and discovery. In *Proceedings of the v2Nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition*, NETFLIX 2008, pages 5:1–5:8, New York, NY, USA, 2008. ACM.

[b]`https://recsys-challenge.spotify.com/details`

13. Zhiyong Cheng and Jialie Shen. Just-for-me: An adaptive personalization system for location-aware social music recommendation. In *Proceedings of the International Conference on Multimedia Retrieval*, ICMR 2014, pages 185:185–185:192, New York, NY, USA, 2014. ACM.
14. Charles L.A. Clarke, Maheedhar Kolla, Gordon V. Cormack, Olga Vechtomova, Azin Ashkan, Stefan Büttcher, and Ian MacKinnon. Novelty and diversity in information retrieval evaluation. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2008, pages 659–666, New York, NY, USA, July 20-24 2008. ACM.
15. Paolo Cremonesi, Franca Garzotto, Roberto Pagano, and Massimo Quadrana. Recommending without short head. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW 2014, pages 245–246. ACM, 2014.
16. Katayoun Farrahi, Markus Schedl, Andreu Vall, David Hauger, and Marko Tkalčič. Impact of listening behavior on music recommendation. In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, ISMIR 2014, pages 483–488, 2014.
17. Bruce Ferwerda. *Improving the User Experience of Music Recommender Systems Through Personality and Cultural Information*. Phd, Johannes Kepler University Linz, Linz, Austria, 2016.
18. David Hauger, Markus Schedl, Andrej Košir, and Marko Tkalčič. The million musical tweets dataset: What can we learn from microblogs. In *Proceedings of the 14th International Society for Music Information Retrieval Conference*, ISMIR 2013, 2013.
19. Brian J Hracs, Michael Seman, and Tarek E Virani. *The production and consumption of music in the digital age*, volume 58. Routledge, 2016.
20. Xiao Hu, Jin Ha Lee, Kahyun Choi, and J.S. Downie. A cross-cultural study of mood in k-pop songs. In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, ISMIR 2014, pages 217–238, 2014.
21. Michael L Jones. Hofstede-culturally questionable? 2007.
22. Shinobu Kitayama and Hyekyung Park. Cultural shaping of self, emotion, and well-being: How does it work? *Social and Personality Psychology Compass*, 1(1):202–222, 2007.
23. Rajeev Kumar, B.K. Verma, and Shyam Sunder Rastogi. Social popularity based svd++ recommender system. *International Journal of Computer Applications*, 87(14):33–37, 2014.
24. Audrey Laplante. Improving music recommender systems: what can we learn from research on music tags? In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, ISMIR 2014, pages 451–456, 2014.
25. Jin Ha Lee and Xiao Hu. Cross-cultural similarities and differences in music mood perception. *iConference 2014 Proceedings*, 2014.
26. Greg Linden, Brent Smith, and Jeremy York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 4(1):76–80, 2003.
27. Brian McFee, Luke Barrington, and Gert Lanckriet. Learning content similarity for music recommendation. *IEEE transactions on audio, speech, and language processing*, 20(8):2207–2218, 2012.
28. Brendan McSweeney. Hofstede's model of national cultural differences and their consequences: A triumph of faith-a failure of analysis. *Human relations*, 55(1):89–118, 2002.
29. Miquel Montaner, Beatriz López, and Josep Lluís de la Rosa. A taxonomy of recommender agents on the internet. *Artificial Intelligence Review*, 19(4):285–330, 2003.
30. Steven J Morrison and Steven M Demorest. Cultural constraints on music perception and cognition. *Progress in brain research*, 178:67–77, 2009.
31. Martin Pichl, Eva Zangerle, and Günther Specht. Towards a context-aware music recommendation approach: What is hidden in the playlist name? In *Proceedings of the 15th IEEE International Conference on Data Mining Workshop*, ICDM 15, pages 1360–1365, Piscataway, NJ, USA, November 2015. IEEE.
32. Tim Pohle, Markus Schedl, Peter Knees, and Gerhard Widmer. Automatically adapting the structure of audio similarity spaces. In *Proceedings of the 1st Workshop on Learning the Semantics of Audio Signals*, LSAS 2006, December 2006.

33. Dominic Power And and Daniel Hallencreutz. Competitiveness, local production systems and global commodity chains in the music industry: entering the us market. *Regional Studies*, 41(3):377–389, 2007.
34. Francesco Ricci, Lior Rokach, and Bracha Shapira. *Recommender Systems Handbook*. Springer, New York, NY, 2nd edition, 2015.
35. Paul Rutten. Local popular music on the national and international markets. *Cultural Studies*, 5(3):294–305, 1991.
36. Ruslan Salakhutdinov and Andriy Mnih. Probabilistic matrix factorization. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, NIPS 2007, pages 1257–1264. Curran Associates Inc., 2007.
37. Gerard Salton, Andrew Wong, and Chung-Shu Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.
38. Markus Schedl. Ameliorating music recommendation: Integrating music content, music context, and user context for improved music retrieval and recommendation. In *Proceedings of International Conference on Advances in Mobile Computing & Multimedia (MoMM)*, MoMM '13, pages 3:3–3:9, New York, NY, USA, 2013. ACM.
39. Markus Schedl. The lfm-1b dataset for music retrieval and recommendation. In *Proceedings of the 17th ACM International Conference on Multimedia Retrieval*, ICMR 2016, New York, NY, USA, April 2016.
40. Markus Schedl. Investigating country-specific music preferences and music recommendation algorithms with the lfm-1b dataset. *International Journal of Multimedia Information Retrieval*, 6(1):71–84, 2017.
41. Markus Schedl and Christine Bauer. Distance- and rank-based music mainstreaminess measurement. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization*, UMAP 2017, pages 364–367, New York, NY, USA, 2017. ACM.
42. Markus Schedl and Bruce Ferwerda. Large-Scale Analysis of Group-Specific Music Genre Taste from Collaborative Tags. In *Proceedings of the 19th IEEE International Symposium on Multimedia (ISM 2017)*, Taichung, Taiwan, December 2017.
43. Markus Schedl, Emilia Gómez, and Julián Urbano. Music information retrieval: Recent developments and applications. *Foundations and Trends in Information Retrieval*, 8(2-3):127–261, 2014.
44. Markus Schedl and David Hauger. Tailoring music recommendations to users by considering diversity, mainstreaminess, and novelty. In *Proceedings of the 38th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2015, August 2015.
45. Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskas. *Recommender Systems Handbook*, chapter Music Recommender Systems, pages 453–492. Springer, New York, NY, USA, 2nd edition, 2015.
46. Abhishek Singhi and Daniel G Brown. On cultural, textual and experiential aspects of music mood. In *ISMIR*, pages 3–8, 2014.
47. Catherine J. Stevens. Music perception and cognition: A review of recent cross-cultural research. *Topics in Cognitive Science*, 4(4):653–667, 2012.
48. Gabriel Vigliensoni and Ichiro Fujinaga. Automatic music recommendation systems: do demographic, profiling, and contextual features improve their performance? In *Proceedings of the 17th International Society for Music Information Retrieval Conference*, ISMIR 2016, pages 94–100, 2016.
49. Linxing Xiao, Lie Lu, Frank Seide, and Jie Zhou. Learning a music similarity measure on automatic annotations with application to playlist generation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, ICASSP 2009, pages 1885–1888. IEEE, 2009.
50. Yan Yan, Tianlong Liu, and Zhenyu Wang. A music recommendation algorithm based on hybrid collaborative filtering technique. In Xichun Zhang, Maosong Sun, Zhenyu Wang, and Xuanjing Huang, editors, *4th National Conference on Social Media Processing*, SMP 2015, pages 233–240,

Singapore, 2015. Springer.

51. JungAe Yang. Effects of popularity-based news recommendations ("most-viewed") on users' exposure to online news. *Media Psychology*, 19(2):243–271, 2016.
52. Yuan Cao Zhang, Diarmuid Ó. Seaghdha, Daniele Quercia, and Tamas Jambor. Auralist: Introducing serendipity into music recommendation. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining*, WSDM 2012, Seattle, WA, USA, February 8-12 2012.