Exploring Music Diversity Needs Across Countries

Bruce Ferwerda
Johannes Kepler University
Altenberger Str. 69
A-4040 Linz, Austria
bruce.ferwerda@jku.at

Andreu Vall
Johannes Kepler University
Altenberger Str. 69
A-4040 Linz, Austria
andreu.vall@jku.at

Marko Tkalcic
Free University of Bolzano
Piazza Domenicani 3
Bolzano, Italy
marko.tkalcic@unibz.it

Markus Schedl
Johannes Kepler University
Altenberger Str. 69
A-4040 Linz, Austria
markus.schedl@jku.at

ABSTRACT

Providing diversity in recommendations has shown to positively influence the user's subjective evaluations such as satisfaction. However, it is often unknown how much diversity a recommendation set needs to consist of. In this work, we explored how music users of Last.fm apply diversity in their listening behavior. We analyzed a dataset with the music listening history of 53,309 Last.fm users capturing their *total* listening events until August 2014. We complemented this dataset with The Echo Nest features and Hofstede's cultural dimensions to explore how music diversity is applied across countries. Between 47 countries, we found distinct relationships between the cultural dimensions and music diversity variables. These results suggest that different country-based diversity measurements should be considered when applied to a recommendation set in order to maximize the user's subjective evaluations. The country-based relationships also provide opportunities for recommender systems to personalize experiences when user data is limited by being able to rely on the user's demographics.

CCS Concepts

•Human-centered computing → User models; •Social and professional topics → Cultural characteristics;

Keywords

Music Recommendations; Diversity; Cultural Differences

1. INTRODUCTION

Providing recommendation diversity to users has become an important feature for recommender systems. It has been acknowledged that focusing only on recommendation accuracy (i.e., recommending the most relevant items to the user) is not enough [11]. Recommending items with high accuracy often result in a set of recommendations that are too similar to each other, and thereby not covering the full spectrum of the user's interest. Giving in on accuracy by introducing diversity can positively influence the subjective evaluation of the recommendations, such as, user satisfaction [3].

A common method to introduce recommendation diversity is to employ a *bounded greedy* algorithm [19]. This algorithm proposes the most relevant items that are at the same time maximally diverse from each other. However, are maximally diverse items always desired? Prior research has identified that individual characteristics (e.g., expertise) play a role in how much diversity is desired by the user [3]. The problem that persist is that often self-report measures are used to identify these influential individual characteristics. Whereas registration and login processes are becoming easier (e.g., single sign-on buttons) [1], asking additional questions may become a bothersome process for (new) users.

The implicit acquisition (i.e., without the use of questionnaires) of individual differences is still a challenging task. Especially for new users there is not enough behavioral data yet to make inferences about user preferences. In this work we take a country level perspective. Country information often already consists in a standard user profile and is therefore easy to acquire. We look at music listening behavior of Last.fm [2] users from 47 countries and explore how they apply diversity to their playlists. We computed a diversity measure based on the unique listening events (i.e., of artists and genres) of users per country. Additionally, by using The Echo Nest features we were able to look at artist characteristics (i.e., how known, familiar, or popular the artists are that have been listened to by users). This provides insights of diversity patterns based on listening characteristics (i.e., artist or genre level) as well as on artist characteristics.

2. RELATED WORK

Recommender systems intend to create a personalized set of items that are most relevant to the user. However, highly relevant items often appear too similar to each other. A set of items showing too much similarities (e.g., too many highly

[1]Buttons implemented in applications that allow users to register and login with their social media account. For example, "login with your Facebook account."
[2]http://www.last.fm

WOODSTOCK '97 El Paso, Texas USA

DOI: 10.475/123_4

relevant items) can, in turn, cause choice overload [17]. Bollen et al. [3] and Willemsen et al. [21] investigated the influence of diversity on movie recommendations and found that diversity has a positive effect on the attractiveness of the recommendation set, the difficulty to make a choice, and eventually on the choice satisfaction. Bollen et al. [3] additionally identified individual differences. For example, they showed that increased expertise has positive effects on perceived item variety and attractiveness.

Besides individual characteristics, research has shown that culture consists of useful cues as well. General behavior and preferences have shown to be rooted and embodied in culture [13], hence looking at behavior on a country level may provide useful information for the desired recommendation diversity. In a comprehensive study, Hofstede et al. [12] describe national cultures among six dimensions: power distance, individualism, masculinity, uncertainty avoidance, long-term orientation and indulgence, which could help to explain diversity differences between countries in this study.

3. METHOD

A Last.fm dataset was used consisting of 53,309 users of 47 countries with the *total* listening history of each user until August 2014 [16]. The dataset consists of users' listening (i.e., user ID, time-stamp, artist name, and track name) and profile information (i.e., gender, age, country). A diversity measure was created by aggregating each user's listening history by artist and genre to identify the unique instances of each respectively. [3] For example, a user has a listening history of 90 events originating from two artists/genres; the diversity would be two here. Each diversity measure was normalized ($r \in [0,1]$) due to the unequal number of users between countries, and aggregated into a mean diversity value for each country (Table 1). The dataset was complemented with Hofstede's cultural dimensions and The Echo Nest features to better understand cultural diversity patterns.

3.1 Hofstede's Cultural Dimensions

The most comprehensive framework for national cultures is considered to be Hofstede's cultural dimensions. They defined six dimensions to identify cultures [12]:

Power distance: defines the extent to which power is distributed unequally by less powerful members of institutions (e.g., family). High power distance indicates that a hierarchy is clearly established and executed in society. Low power distance indicates that authority is questioned and attempted to distribute power equally.

Individualism: defines the degree of integration of people into societal groups. High individualism is defined by loose social ties. The main emphasis is on the "I" instead of the "we," while opposite for low individualistic cultures.

Masculinity: defines a society's preference for achievement, heroism, assertiveness and material rewards for success (countries scoring high in this dimension). Whereas low masculinity represents a preference for cooperation, modesty, caring for the weak and quality of life.

[3]Genre was obtained through The Echo Nest API. Additionally, in order to maintain a manageable dataset, we decided not to focus on listening events on a track level in this study.

Uncertainty avoidance: defines a society's tolerance for ambiguity. High scoring countries in this scale are more inclined to opt for stiff codes of behavior, guidelines, laws. Whereas more acceptance of differing thoughts and/or ideas are accepting for those scoring low in this dimension.

Long-term orientation: is associated with the connection of the past with the current and future actions and/or challenges. Lower scoring countries tend to believe that traditions are honored and kept, and value steadfastness. High scoring countries believe more that adaptation and circumstantial, pragmatic problem-solving are necessary.

Indulgence: defines in general the happiness of a country. Countries scoring high in this dimension are related to a society that allows relatively free gratification of basic and natural human desires related to enjoying life and having fun (e.g., be in control of their own life and emotions). Whereas low scoring countries show more controlled gratification of needs and regulate it by means of strict social norms.

3.2 The Echo Nest Features

We complemented the dataset with The Echo Nest features through their API. [4] The Echo Nest is a common service to gather additional information about artists [2]. We added the following features: genre, hotness (i.e., how much attention an artist is getting. For example, in the media), familiarity (i.e., how likely it is that someone has heard of the artist), and discovery (i.e., how unexpectedly popular an artist is). As the features (i.e., hotness, familiarity, and discovery) follow a normal distribution, we calculated per country the mean values for each feature to represent the central tendency of the artists listened to by users in a country. Measurements represent a value $r \in [0,1]$ (Table 1).

4. RESULTS

The analyses we conducted are divided into two parts. In the first part we investigated the relationship between the cultural dimensions and diversity based on listening characteristics (i.e., whether diversity occurs on an artist or genre level). In the second part we take a deeper look into the relationship between cultural dimensions and artist characteristics (i.e., hotness, familiarity, and discovery). We performed correlation analyses to explore the relationship between cultural dimensions and different diversity measures, and report Pearson's correlation ($r \in [-1,1]$) to indicate the linear relationship.

4.1 Listening Characteristics

Table 2 shows the correlation results between the cultural dimensions and listening characteristics. A negative correlation represents the degree of diversity, whereas a positive correlation indicates homogeneity. Positive correlations were found between a culture's power distance and artist (r=.279, p=.043) and genre (r=.329, p=.027) diversity. This indicates that users in countries scoring high on this dimension tend to apply less diversity by artist as well as by genre. Negative correlations were found between the individualism dimensions and artist (r=-.373, p=.012) and genre (r=-.265, p=.048) diversity, which indicates that users in individualistic countries tend to apply music diversity on an artist as

[4]http://developer.echonest.com/

Table 1: Number of users per country. Hofstede's cultural dimensions ($r \in [0,120]$ with 60 as a midlevel): power distance (PDI), individualism (IDV), masculinity (MAS), uncertainty avoidance (UAI), long-term orientation (LTO), and indulgence (IND). Unique listening events ($r \in [0,1]$ where 0 means divers and 1 means homogeneous): artist and genre . The Echo Nest features ($r \in [0,1]$): hotness, familiarity, and discovery.

Country	#Users	PDI	IDV	MAS	UAI	LTO	IND	Artist	Genre	Hotness	Familiarity	Discovery
U.S.A.	10255	40	91	62	46	26	68	0.0354	0.1192	0.5284	0.5188	0.3907
Russia	5024	93	39	36	95	81	20	0.0324	0.1622	0.5023	0.4700	0.3830
Germany	4578	35	67	66	65	83	40	0.0137	0.0614	0.5194	0.4969	0.3950
Great Britain	4534	35	89	66	35	51	69	0.0429	0.2030	0.5276	0.5175	0.3906
Poland	4408	68	60	64	93	38	29	0.0372	0.1475	0.5306	0.5122	0.3905
Brazil	3886	69	38	49	76	44	59	0.0104	0.0405	0.5491	0.5424	0.3892
Finland	1409	33	63	26	59	38	57	0.0308	0.0906	0.5234	0.4919	0.3961
Netherlands	1375	38	80	14	53	67	68	0.0361	0.1285	0.5270	0.5027	0.3961
Spain	1243	57	51	42	86	48	44	0.0302	0.0824	0.5304	0.5084	0.3976
Sweden	1231	31	71	5	29	53	78	0.0627	0.1981	0.5278	0.4988	0.3998
Ukraine	1143	N/A	N/A	N/A	N/A	86	14	0.0672	0.2041	0.5088	0.4787	0.3848
Canada	1077	39	80	52	48	36	68	0.0155	0.0578	0.5248	0.5154	0.3868
France	1055	68	71	43	86	63	48	0.0335	0.0993	0.5190	0.4954	0.3937
Australia	976	38	90	61	51	21	71	0.0261	0.0509	0.5297	0.5197	0.3889
Italy	974	50	76	70	75	61	30	0.0500	0.1261	0.5295	0.5198	0.3894
Japan	806	54	46	95	92	88	42	0.0535	0.1318	0.4870	0.4531	0.3755
Norway	750	31	69	8	50	35	55	0.0537	0.1517	0.5317	0.5054	0.3986
Mexico	705	81	30	69	82	24	97	0.0594	0.1424	0.5337	0.5154	0.3954
Czech Rep	632	57	58	57	74	70	29	0.0275	0.0862	0.5186	0.4995	0.3888
Belarus	558	N/A	N/A	N/A	N/A	81	15	0.0154	0.0537	0.5018	0.4648	0.3851
Belgium	513	65	75	54	94	82	57	0.0403	0.0793	0.5208	0.5003	0.3928
Indonesia	484	78	14	46	48	62	38	0.1210	0.1858	0.5481	0.5357	0.3893
Turkey	479	66	37	45	85	46	49	0.1242	0.2475	0.5264	0.5117	0.3928
Chile	425	63	23	28	86	31	68	0.0416	0.0931	0.5338	0.5255	0.3857
Croatia	372	73	33	40	80	58	33	0.0794	0.1445	0.5185	0.5072	0.3842
Portugal	291	63	27	31	104	28	33	0.1412	0.2392	0.5232	0.5106	0.3870
Argentina	282	49	46	56	86	20	62	0.0232	0.0379	0.5310	0.5210	0.3879
Switzerland	277	34	68	70	58	74	66	0.0324	0.0798	0.5202	0.5001	0.3944
Austria	276	11	55	79	70	60	63	0.0373	0.0702	0.5228	0.5043	0.3915
Denmark	272	18	74	16	23	35	70	0.0614	0.1023	0.5322	0.5101	0.3970
Hungary	272	46	80	88	82	58	31	0.0801	0.1513	0.5057	0.4803	0.3863
Serbia	253	86	25	43	92	52	28	0.1450	0.2320	0.5059	0.4895	0.3811
Romania	237	90	30	42	90	52	20	0.1119	0.2200	0.5090	0.4850	0.3878
Bulgeria	236	70	30	40	85	69	16	0.0832	0.1182	0.5152	0.4979	0.3866
Ireland	220	28	70	68	35	24	65	0.1220	0.1819	0.5400	0.5383	0.3891
Lithuania	202	42	60	19	65	82	16	0.1017	0.1743	0.5122	0.4910	0.3848
Slovak Rep	192	104	52	110	51	77	28	0.2369	0.3394	0.5127	0.4954	0.3861
Greece	175	60	35	57	112	45	50	0.1003	0.1675	0.5129	0.4982	0.3856
Latvia	165	44	70	9	63	69	13	0.1095	0.1818	0.5150	0.4922	0.3871
New Zealand	164	22	79	58	49	33	75	0.1699	0.2499	0.5327	0.5223	0.3905
China	162	80	20	66	30	87	24	0.0573	0.1155	0.5128	0.4748	0.3895
Colombia	159	67	13	64	80	13	83	0.0606	0.1071	0.5359	0.5208	0.3947
Iran	135	58	41	43	59	14	40	0.2379	0.2471	0.5320	0.5156	0.3892
India	122	77	48	56	40	51	26	0.2127	0.2542	0.5437	0.5393	0.3857
Venezuela	118	81	12	73	76	16	100	0.5122	0.3238	0.5431	0.5319	0.3896
Estonia	107	40	60	30	60	82	16	0.1155	0.1372	0.5229	0.5018	0.3888
Israel	100	13	54	47	81	38	N/A	0.1740	0.2177	0.5227	0.5075	0.3877

Table 2: Correlation results between unique listening events (artist and genre) and Hofstede's cultural dimensions: power distance (PDI), individualism (IDV), masculinity (MAS), uncertainty avoidance (UAI), long-term orientation (LTO), and indulgence (IND).

	PDI	IDV	MAS	UAI	LTO	IND
Artist	.279*	-.373**	.155	-.020	-.259*	.080
Genre	.329*	-.265*	.074	-.051	-.108	-.113

Note. *p<.05, **p<.01

Table 3: Correlation results between The Echo Nest features (hotness, familiarity, and discovery) and Hofstede's cultural dimensions: power distance (PDI), individualism (IDV), masculinity (MAS), uncertainty avoidance (UAI), long-term orientation (LTO), and indulgence (IND).

	PDI	IDV	MAS	UAI	LTO	IND
Hot.	-.131	-.0.39	-.135	-.359*	-.641**	.557**
Fam.	-.100	-.229	.009	-.255	-.677**	.520**
Disc.	-.367**	-.294*	-.311*	-.366*	-.274*	.517**

Note. *p<.05, **p<.01

well as genre level. Finally, a negative correlation was found between long-term orientation and artist diversity (r=-.259, p=.048). Users in countries scoring high on this dimension tend to listen to more diverse artists.

4.2 Artist Characteristics

The Echo Nest features allows us to gain more insights of the diversity characteristics of the artists. A negative correlation indicates that users of a country involve artists that score low on the respective feature of The Echo Nest when scoring high in the correlated cultural dimension (Table 3).

Hotness. Hotness was found to be negatively correlated with uncertainty avoidance (r=-.359, p=.015) and long-term orientation (r=-.641, p=.000), while a positive correlation was found with indulgence (r=.557, p=.000).

Familiarity. Familiarity was found to be negatively correlated with long-term orientation (r=-.677, p=.000), but positively correlated with indulgence (r=.520, p=.000).

Discovery. Discovery was found correlated negatively with five out of six cultural dimensions: power distance (r=-.367, p=.013), individualism (r=-.294, p=.050), masculinity (r=-.311, p=.038), uncertainty avoidance (r=-.366, p=.013), and long-term orientation (r=-.274, p=.042). A positive correlation was found with indulgence (r=.517, p=.000).

5. CONCLUSION & IMPLICATIONS

We show with our results that different diversity patterns exist and that they are related to cultural dimensions. When looking at the relationship between listening characteristics (i.e., artist and genre) and cultural dimensions, distinct correlations were found. Users in countries scoring high on power distance appear to apply less diversity in artist and genre to their playlists. Power distance refers to the extent inequality of power distributions within a group is accepted, and is related to obedience [12]. This could be an explanation why less diversity is applied.

The individualism dimension is found to be related to more diversity on an artist as well as on a genre level. One explanation could be that in individualistic countries, people are relying on loose social ties, and are expected to look after themselves [12]. Therefore, they may be inclined to explore and more willing to try different artists and genres.

A negative correlation was found of only artist diversity and long-term orientation. Users in low scoring countries in this dimension believe that traditions are honored and kept. Moreover they value steadfastness [12], hence they would be less inclined to try diverse artists.

The artist characteristics (i.e., The Echo Nest features) provided a deeper insight into the characteristics of the artists. Hotness and discovery both correlate negatively with uncertainty avoidance. This means that users in countries scoring high in this dimension are less prone to listen to artists that are considered "hot" (e.g., getting a lot of media attention) or "discoverable" (e.g., "one-hit wonders"). These countries are characterized by having less tolerance to ambiguity [12], whereas media attention (e.g., conflicting information) or one-hit wonders may increase ambiguity and uncertainty.

Negative correlations were found between all features and long-term orientation. Users in countries scoring high on long-term orientation listen less to artists that are hot, unknown, or unexpectedly popular. As mentioned earlier, these countries value traditions and steadfastness [12], which could explain the relationship found.

All the features show a positive correlation with the indulgence dimension. This means that users in indulgent countries are more prone to listen artist that are hot, unknown, or unexpectedly popular. This would be in line with the explanation that indulgent countries are characterized of being in control of their own life [12], and therefore are more open to all kind of the different artist characteristics (i.e., hotness, familiarity, and discovery).

Approaching diversity on a country level enables the creation of proxy measures for personalization when data is limited, such as the "cold-start problem" in recommender systems. Users' personality has gained interest to make inferences for personalization (e.g., [5, 9, 20]) to solve this problem. One way to extract personality is facilitated by the increased connectedness of applications and social media (e.g., single sign-on buttons). This allows exploitation of social media data for personality acquisition (e.g., Facebook [1, 4, 14], Twitter [10, 15], Instagram [6, 7, 8, 18]). However, a connection with the user's social media account is still needed. Our results could be used to make inferences about the user's diversity needs based on their country, which is often available through the user's account information.

In this work we showed that there are differences between cultural dimensions and diversity on an artist and genre levels. To keep the dataset manageable in this study, we disregarded track diversity. This is a limitation of this work, but also a future direction that we would want to pursue.

6. ACKNOWLEDGMENT

This research is supported by the Austrian Science Fund (FWF): P25655; and by the EU FP7/2013-2016 through the PHENICX project under grant agreement no.601166.

7. REFERENCES

[1] M. D. Back, J. M. Stopfer, S. Vazire, S. Gaddis, S. C. Schmukle, B. Egloff, and S. D. Gosling. Facebook profiles reflect actual personality, not self-idealization. *Psychological science*, 2010.

[2] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere. The million song dataset. In *ISMIR 2011*, pages 591–596. Miami, Florida, 2011.

[3] D. Bollen, B. P. Knijnenburg, M. C. Willemsen, and M. Graus. Understanding choice overload in recommender systems. In *Proceedings of the fourth ACM RecSys*, pages 63–70. ACM, 2010.

[4] F. Celli, E. Bruni, and B. Lepri. Automatic personality and interaction style recognition from facebook profile pictures. In *Proceedings of the ACM MM*, pages 1101–1104. ACM, 2014.

[5] B. Ferwerda and M. Schedl. Enhancing music recommender systems with personality information and emotional states: A proposal. In *Proc. EMPIRE workshop*, 2014.

[6] B. Ferwerda, M. Schedl, and M. Tkalcic. Predicting personality traits with instagram pictures. In *3rd Workshop on EMPIRE*, 2015.

[7] B. Ferwerda, M. Schedl, and M. Tkalcic. Personality traits and the relationship with (non-) disclosure behavior on facebook. *WWW*, 2016.

[8] B. Ferwerda, M. Schedl, and M. Tkalcic. Using instagram picture features to predict users' personality. In *MultiMedia Modeling*, pages 850–861, 2016.

[9] B. Ferwerda, E. Yang, M. Schedl, and M. Tkalcic. Personality traits predict music taxonomy preferences. In *CHI'15 Ext. Abstracts*, 2015.

[10] J. Golbeck, C. Robles, M. Edmondson, and K. Turner. Predicting personality from twitter. In *SocialCom*, pages 149–156, 2011.

[11] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.

[12] G. Hofstede, G. J. Hofstede, and M. Minkov. *Cultures and organizations: Software of the mind*, volume 2. Citeseer, 1991.

[13] S. Kitayama and H. Park. Cultural shaping of self, emotion, and well-being: How does it work? *Soc. and Pers. Psych. Compass*, 2007.

[14] G. Park, H. A. Schwartz, J. C. Eichstaedt, M. L. Kern, M. Kosinski, D. J. Stillwell, L. H. Ungar, and M. E. Seligman. Automatic personality assessment through social media language. 2014.

[15] D. Quercia, M. Kosinski, D. Stillwell, and J. Crowcroft. Our twitter profiles, our selves: Predicting personality with twitter. In *SocialCom*.

[16] M. Schedl. The lfm-1b dataset for music retrieval and recommendation. *ICMR (submitted)*, 2016.

[17] B. Scheibehenne, R. Greifeneder, and P. M. Todd. Can there ever be too many options? a meta-analytic review of choice overload. *Journal of Consumer Research*, 37(3):409–425, 2010.

[18] M. Skowron, B. Ferwerda, M. Tkalcic, and M. Schedl. Fusing social media cues: Personality prediction from twitter and instagram. *WWW*, 2016.

[19] B. Smyth and P. McClave. Similarity vs. diversity. In *Case-Based Reasoning Research and Development*, pages 347–361. Springer, 2001.

[20] M. Tkalcic, B. Ferwerda, D. Hauger, and M. Schedl. Personality correlates for digital concert program notes. 2015.

[21] M. C. Willemsen, B. P. Knijnenburg, M. P. Graus, L. C. Velter-Bremmers, and K. Fu. Using latent features diversification to reduce choice difficulty in recommendation lists. *RecSys*, 2011.