

Indicators of Country Similarity in Terms of Music Taste, Cultural, and Socio-economic Factors

Markus Schedl
Johannes Kepler University (JKU)
Linz, Austria
Email: markus.schedl@jku.at

Florian Lemmerich
Leibniz Institute for the Social Sciences (GESIS)
Cologne, Germany
Email: florian.lemmerich@gesis.org

Bruce Ferwerda
Jönköping University
Jönköping, Sweden
Email: bruce.ferwerda@ju.se

Marcin Skowron
Austrian Research Institute for Artificial Intelligence (OFAI)
Vienna, Austria
Email: marcin.skowron@ofai.at

Peter Knees
Vienna University of Technology (TU Wien)
Vienna, Austria
Email: peter.knees@tuwien.ac.at

Abstract—Considering the cultural background of users is known to improve recommender systems for multimedia items. In this work, we focus on music and analyze user demographics and music listening events in a large corpus (120,000 users, 10^9 events) from Last.fm to investigate whether similarity between countries in terms of cultural and socio-economic factors is reflected in music taste. To this end, we propose a tag-based model to describe the music taste of a country and correlate the resulting music profiles to Hofstede’s cultural dimensions and the *Quality of Government* data. Spearman’s rank-order correlation and Quadratic Assignment Procedure indeed indicate statistically significant weak to medium correlations of music taste and several cultural and socio-economic factors. The results will help elaborating culture-aware models of music listeners and in turn likely yield improved music recommender systems.

I. INTRODUCTION AND CONTEXT

Knowledge about the similarities and differences in music taste between countries and about how these relate to cultural and socio-economic dimensions can improve culture-aware and cross-cultural music retrieval and recommender systems and help mitigate the cold-start problem in cases where only the country of a new user is known, but not his or her music taste. This is a common scenario given the single sign-on approach adopted by many current online services and platforms, including recommender systems. This short paper aims at gaining insights into the aforementioned similarities by exploiting *social media* data to model music preferences on the country level and in turn address the research question whether music taste similarity correlates with cultural and socio-economic similarity.

While recently there has been found evidence that culture is related to online behavior, leading to a connection of anthropological theories with computational models, e.g. [1], [2], the aspect of connecting online music consumption patterns to cultural dimensions yet has not been studied in detail (in contrast to other aspects of digital traces, e.g., by connecting cultural boundaries to food and drink habits [3]). Since smaller-scale musicological-driven offline studies have provided insight that cultural traits are connected to musical listening preferences,

e.g. [4], [5], this work aims at developing a computational approach to culture-specific music consumption behavior, with the vision of supporting music recommender systems.

Research that considers individual, user-specific aspects to improve music retrieval and recommendation algorithms has received substantial attention in the past few years, e.g. [6]–[8]. Existing works predominantly focus on emotion or mood perceived when listening to music, aiming to exploit such knowledge for music retrieval, e.g. [9]–[11]. In contrast, studies investigating cultural differences in perception or consumption of music have not been performed until recently. Hu and Lee [12] found differences in perception of moods between American and Chinese listeners. By analyzing Last.fm music listening behavior of users from 49 countries, Ferwerda et al. [13], [14] found relationships between music listening diversity and Hofstede’s cultural dimensions. Similarly, Skowron et al. [15] use the same dimensions to predict music genre preferences of users with different cultural backgrounds.

II. METHODOLOGY AND DATA

A. Modeling culture and socio-economics

To represent *cultural aspects* on the country level, we rely on Hofstede et al.’s work [16] since it is considered a comprehensive and up to date framework for national cultures.¹ They define six dimensions to describe cultures: power distance, individualism, masculinity, uncertainty avoidance, long-term orientation, and indulgence.

We further investigate a range of *socio-economic indicators* taken from the Quality of Government (QoG) dataset,² which aggregates approximately 2,500 variables from more than 100 data sources. From this dataset, we extract a subset of 181 variables for which all the scores are available for the set of analyzed countries. Examples include GDP, income inequality, agriculture’s share of economy, unemployment rate, and life expectancy. Details on the variables are provided in [17].

¹<https://geert-hofstede.com/countries.html>

²<http://qog.pol.gu.se/data/datadownloads/qogbasicdata>

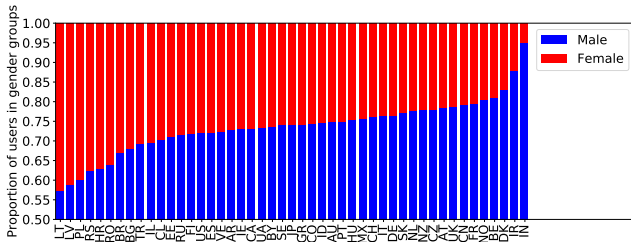


Fig. 1: Distribution of Last.fm users in LFM-1b dataset, over gender groups, sorted increasingly according to share of male users from left to right.




Fig. 2: Age distribution of Last.fm users in LFM-1b dataset, sorted according to median age from young to old.




Fig. 3: Radar plot of genre profiles for Last.fm users in LFM-1b dataset, for top countries.

B. Modeling music preferences

We model music preference on the country level by exploiting the LFM-1b dataset [18], which offers (partly incomplete) demographic information and listening histories for about 120,000 Last.fm users. We consider only countries with at least 100 users and for which all Hofstede’s cultural dimensions are available, which yields about 53,000 users from 44 countries meeting both conditions. Figures 1 and 2 show, respectively, the distribution of gender and age for the 44 countries, which are coded according to the ISO-3166 standard.³ While an initial analysis reveals that the distributions of gender and age in the dataset do not correspond to the population at large, we argue that the data is representative for music lovers who use social media, due to the almost global popularity of Last.fm and the country filtering we apply. This assumption is further supported by a demographic analysis of different social media platforms,⁴ in which Last.fm ranked in the center of all platforms’ distributions.

We define country-specific *genre profiles*, which are used as a proxy for music taste. First, the top tags assigned to each artist in the LFM-1b dataset are gathered from Last.fm.⁵ These tags provide different pieces of information, including instruments (“guitar”), epochs (“80s”), places (“Chicago”), or languages (“Swedish”). We then filter for tags that reflect genre and style, using dictionaries of 20 and 1,998 terms retrieved from Allmusic⁶ and Freebase,⁷ respectively.⁸ The final genre profiles are weight vectors describing the share of each genre among all listening events of the respective country’s population. Formally, the weight of genre g in country c is computed as $w_{c,g} = \frac{\sum_{a \in A_g} l_{c,a}}{\sum_{a \in A} l_{c,a}}$, where A_g is the set of artists tagged with genre g , A is the entire set of artists, and $l_{c,a}$ is the number of listening events to artist a in country c .

To obtain a coarse knowledge about the music preferences, Figure 3 shows a radar plot of the genre profiles according to the Allmusic dictionary, for the countries with at least 100 users in the LFM-1b dataset. Starting with the US, countries are sorted in descending order of users in a counterclockwise manner. To reduce visual clutter and increase readability, we include only the shares of some of the most popular genres from Allmusic. As a general tendency, we observe that the popularity ranking of genres is quite consistent between countries. A few exceptions are, for instance, Japan and China, where the share of pop music is higher than that of alternative. Electronic music is consumed to a disproportionately high amount in Russia, France, Belarus, Hungary, Romania, and Estonia, whereas very little in South American countries (Brazil, Chile, and Argentina), Indonesia, and India. Pop music peaks in Japan, China, and Indonesia; folk in the US, Romania,

³<https://www.iso.org/iso-3166-country-codes.html>

⁴<http://royal.pingdom.com/2012/08/21/report-social-network-demographics-in-2012>

⁵<http://www.last.fm/api/show/artist.getTopTags>

⁶<http://www.allmusic.com>

⁷<http://www.freebase.com>

⁸The full lists of genres and styles can be shared upon request.

TABLE I: Spearman rank-order correlation coefficients (ρ) and corresponding magnitude of p -values according to QAP.

| Aspect | ρ | p -value |
|--|---------|------------|
| Masculinity | 0.2476 | E-17 |
| Power distance | 0.2240 | E-20 |
| Long-term orientation | 0.1791 | E-06 |
| Uncertainty avoidance | 0.1539 | E-01 |
| Indulgence | 0.1484 | E-09 |
| Individualism | 0.1083 | E-01 |
| Islam: total % adherents (arda_isgenpct) | 0.3929 | 0 |
| Ethnic fractionalization (al_ethnic) | 0.3578 | 0 |
| Pct. no schooling, Female 25+ (bl_lu_25f) | 0.3399 | 0 |
| Independent judiciary (bti_ij) | 0.3166 | E-34 |
| Vote fraud (dpi_fraud) | 0.3142 | E-10 |
| Pct. no schooling, male 25+ (bl_lu_25m) | 0.3024 | E-32 |
| Trust in parliament (ess_trparl) | 0.2938 | E-31 |
| Information transparency (diat_iti) | 0.2926 | E-41 |
| Pct. not speaking the official lang. (el_gunn1) | 0.2918 | E-21 |
| Freedom of expression (ess_trparl) | 0.2862 | E-23 |
| Child mortality (epi_chmort) | 0.2824 | E-23 |
| Trust in legal system (ess_trlegal) | 0.2797 | E-20 |
| Corruption commission present in constitution (ccp_cc) | 0.2791 | E-17 |
| Total seats in legislature (dpi_seats) | 0.2786 | E-15 |
| Average schooling years, Female 25+ (bl_asy25f) | 0.2739 | E-22 |
| Associational/Assembly rights (bti_aar) | 0.2690 | E-42 |
| Social safety nets (bti_ssn) | 0.2668 | E-27 |
| Civil society traditions (bti_cst) | 0.2603 | E-18 |
| Hindu: total % adherents (arda_higenpct) | 0.2582 | E-28 |
| Total ecological footprint (ef_ef) | 0.2570 | E-19 |
| Approval of democracy (bti_aod) | 0.2556 | E-11 |
| Geographical distance | -0.1873 | 0 |

Ireland, and Iran. Metal is particularly popular in Finland, Turkey, and Bulgaria.

C. Country Similarity Computation

We estimate proximity of countries in terms of music preferences via cosine similarity over the raw genre playcount vectors, in order to normalize for different amounts of listening events in different countries. To gauge similarity in terms of cultural and socio-economic dimensions, we calculate the Euclidean distance between the respective scores given by Hofstede and QoG data. We further add as another aspect the geographical distance between countries, which is computed as the geodesic distance in kilometers between country capitals [19].

III. EXPERIMENTS AND RESULTS

Since we cannot assume linear relationships between the individual similarity values, we compute Spearman rank-order correlation coefficient (ρ) between the music similarities on the one hand and the similarities for each of the cultural and socio-economic factors on the other. As the similarity scores of one country are dependent with each other, we resort to a Quadratic Assignment Procedure (QAP) to assert the significance of the findings [20]. This method accounts for dependencies in network data by comparing the actual data with a randomization-based null model. We adjust the p -value with Bonferroni correction to control the family-wise error rate. Correlation coefficients and p -values from the QAP are reported in Table I.

The table shows all six cultural dimensions (top rows), the top 20 QoG factors (middle rows), and geographical distance (last row). We see that most correlations are positive and weak to medium ($\rho \in [0.2, 0.4]$). Results for *masculinity*,

power distance, *long-term orientation*, and *indulgence* are statistically significant at low p -values (even after Bonferroni correction). The highest correlations among the socio-economic dimensions, and the only ones with $\rho > 0.3$, are found for factors related to ethnics, religion, and education: *percentage of people who adhere to Islam*, *ethnic fractionalization* (reflects the likelihood that two randomly selected persons from the same country share racial and linguistic characteristic), and *percentage of females aged 25 or older with no schooling*. This gives an indication that populations with similar distribution of religions, races, and languages show similar levels of similarities in terms of music preferences. A bit surprisingly, even though Table I shows that *geographical distance* is negatively correlated with music similarity according to Spearman, i.e., nearby countries share a similar taste, we would have expected a more pronounced correlation.

Looking deeper into the music similarity scores, Figure 4 visualizes the similarities between all pairs of countries according to their genre profiles, computed as indicated in Section II-C. Darker colors indicate higher, brighter colors lower similarity. The figure indeed shows high similarities between the far away countries US, UK, and Australia, which nevertheless share similar culture and values. This geographically high distance between culturally similar countries seems only partly compensated by clusters formed by the geographically and culturally close countries France and Belgium or Russia, Belarus, Ukraine, and the Baltic countries. We therefore conclude that similarity in music preferences is better reflected by cultural and socio-economic similarity than by geographic distance. Furthermore, we can also identify outliers like Japan, China, or Iran, whose music taste is dissimilar from almost all other countries. To investigate whether these empirical observations of country similarities are also reflected when using a dedicated clustering approach, we applied Affinity Propagation [21] on country’s genre profiles, using as input vectors the country-specific normalized number of listeners for each genre, and letting the algorithm decide on the number of clusters.⁹ Table II illustrates the results, which confirm that similarity of music taste is rather a function of shared cultural background and language than geographic proximity. For instance, the Spanish-speaking Latin American countries in the dataset, i.e., Mexico, Chile, Argentina, Colombia, and Venezuela, form a cluster (#6), while Portuguese-speaking Brazil forms a cluster on its own (#2). So do Japan, Indonesia, China, Iran, and India, which can be considered outliers. Again, also the former Soviet countries Russia, Ukraine, and Belarus form a cluster (#1), and most major English-speaking countries do: USA, Canada, Australia, Ireland, and New Zealand (#4).

IV. CONCLUSIONS AND FUTURE WORK

We exploited user-generated data of music listening events to investigate the relationship between music taste and cultural

⁹We used the implementation in scikit-learn (<http://scikit-learn.org>), which identified 11 clusters.

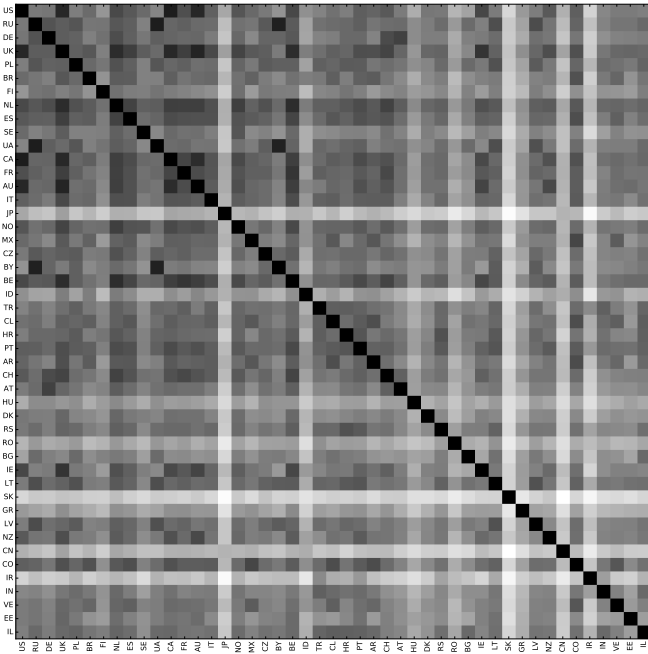


Fig. 4: Similarities between music taste (genre profiles) of Last.fm users in LFM-1b dataset, aggregated on the country level. Darker shades of gray indicate higher similarity.

TABLE II: Country clusters based on Affinity Propagation.

| Clu. | Countries |
|------|--|
| #1 | RU, UA, BY |
| #2 | BR |
| #3 | UK, NL, ES, SE, FR, NO, BE, CH, DK |
| #4 | US, CA, AU, IE, NZ |
| #5 | JP |
| #6 | MX, CL, AR, CO, VE |
| #7 | DE, PL, FI, IT, CZ, TR, HR, PT, AT, HU, RS, RO, BG, LT, SK, GR, LV, EE, IL |
| #8 | ID |
| #9 | CN |
| #10 | IR |
| #11 | IN |

and socio-economic factors, measured on the country level. We found *significant weak to medium correlations* for Hofstede’s *masculinity* and *power distance* as well as for dimensions related to *ethnicity*, *religion*, *education*, and *politics* in the QoG data. The gained insights contribute to a better understanding of the interrelationship between the aforementioned factors and can be applied to alleviate cold-start situations where only country information of a new user is available to the music recommender system, for instance when the system employs a single sign-on approach to register new users.

As part of future work, we will investigate whether the findings also hold for other platforms used to share listening information, by conducting a similar study on music listening datasets mined from Twitter, e.g. [22], [23]. Furthermore, we would like to study the causal relationships between the music taste and the cultural and socio-economic factors. Exploiting information on the individual level, we also plan to investigate to which extent age and gender influence the results.

REFERENCES

- [1] R. Garcia-Gavilanes, D. Quercia, and A. Jaimes, “Cultural dimensions in twitter: Time, individualism and power,” in *Proc. ICWSM*, 2013.
- [2] R. Garcia Gavilanes, “On the quest of discovering cultural trails in social media,” in *Proc. WSDM*, 2013, pp. 747–752.
- [3] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, M. Musolesi, and A. A. F. Loureiro, “You are what you eat (and drink): Identifying cultural boundaries by analyzing food & drink habits in foursquare,” in *Proc. ICWSM*, 2014.
- [4] P. J. Rentfrow and S. D. Gosling, “The do re mi’s of everyday life: the structure and personality correlates of music preferences,” *Journal of personality and social psychology*, vol. 84, no. 6, p. 1236, 2003.
- [5] A. C. North and D. J. Hargreaves, “Lifestyle correlates of musical preference: 1. relationships, living arrangements, beliefs, and crime,” *Psychology of Music*, vol. 35, no. 1, pp. 58–87, 2007.
- [6] Z. Cheng and J. Shen, “Just-for-Me: An Adaptive Personalization System for Location-Aware Social Music Recommendation,” in *Proceedings of the 2014 ACM International Conference on Multimedia Retrieval*, 2014.
- [7] Y. Hu and M. Ogihara, “NextOne Player: A Music Recommendation System Based on User Behavior,” in *Proc. ISMIR*, 2011.
- [8] M. Schedl, S. Stober, E. Gómez, N. Orío, and C. C. Liem, “User-Aware Music Retrieval,” in *Multimodal Music Processing*, M. Müller, M. Goto, and M. Schedl, Eds. Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2012.
- [9] X. Hu and Y. H. Yang, “Cross-dataset and Cross-cultural Music Mood Prediction: A Case on Western and Chinese Pop Songs,” *IEEE Transactions on Affective Computing*, 2016.
- [10] J.-C. Wang, Y.-H. Yang, and H.-M. Wang, “Affective Music Information Retrieval,” in *Emotions and Personality in Personalized Services*. Springer, 2016.
- [11] A. Singhi and D. G. Brown, “On Cultural, Textual and Experiential Aspects of Music Mood,” in *Proc. ISMIR*, 2014.
- [12] X. Hu and J. H. Lee, “A Cross-cultural Study of Music Mood Perception Between American and Chinese Listeners,” in *Proc. ISMIR*, 2012.
- [13] B. Ferwerda, A. Vall, M. Tkalčič, and M. Schedl, “Exploring Music Diversity Needs Across Countries,” in *Proc. UMAP*, 2016.
- [14] B. Ferwerda and M. Schedl, “Investigating the Relationship Between Diversity in Music Consumption Behavior and Cultural Dimensions: A Cross-country Analysis,” in *Workshop on Surprise, Opposition, and Obstruction in Adaptive and Personalized Systems*, July 2016.
- [15] M. Skowron, F. Lemmerich, B. Ferwerda, and M. Schedl, “Predicting genre preferences from cultural and socio-economic factors for music retrieval,” in *Proc. ECIR*, 2017.
- [16] G. Hofstede, G. J. Hofstede, and M. Minkov, *Cultures and Organizations: Software of the Mind*, 3rd ed. McGraw-Hill, USA, 2010.
- [17] J. Teorell, S. Dahlberg, S. Holmberg, B. Rothstein, A. Khomenko, and R. Svensson, *The Quality of Government Standard Dataset, version Jan16*, University of Gothenburg: The Quality of Government Institute, 2016. [Online]. Available: <http://www.qog.pol.gu.se/doi:10.18157/QoGStdJan16>
- [18] M. Schedl, “The LFM-1b Dataset for Music Retrieval and Recommendation,” in *Proc. ACM ICMR*, 2016.
- [19] A. Samoilenko, F. Karimi, D. Edler, J. Kunegis, and M. Strohmaier, “Linguistic neighbourhoods: explaining cultural borders on Wikipedia through multilingual co-editing activity,” *EPJ Data Science*, vol. 5:9, 2016.
- [20] L. Hubert and J. Schultz, “Quadratic assignment as a general data analysis strategy,” *British Journal of Mathematical and Statistical Psychology*, vol. 29, no. 2, pp. 190–241, 1976.
- [21] B. J. Frey and D. Dueck, “Clustering by passing messages between data points,” *Science*, vol. 315, no. 5814, pp. 972–976, 2007. [Online]. Available: <http://science.sciencemag.org/content/315/5814/972>
- [22] E. Zangerle, M. Pichl, W. Gassler, and G. Specht, “#nowplaying music dataset: Extracting listening behavior from twitter,” in *Proc. ACM WISMM*, 2014.
- [23] D. Hauger, M. Schedl, A. Košir, and M. Tkalčič, “The Million Musical Tweets Dataset: What Can We Learn From Microblogs,” in *Proc. ISMIR*, 2013.