<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *The 19th IEEE International Symposium on Multimedia (ISM2017), Taichung, Taiwan, December 11-13, 2017..*

Citation for the original published paper:

Schedl, M., Ferwerda, B. (2017)

Large-scale Analysis of Group-specific Music Genre Taste From Collaborative Tags
In: *Proceedings - 2017 IEEE International Symposium on Multimedia, ISM 2017*,
Code 134021 (pp. 479-482).

<https://doi.org/10.1109/ISM.2017.95>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:hj:diva-37614>

Large-scale Analysis of Group-specific Music Genre Taste From Collaborative Tags

Markus Schedl

Department of Computational Perception
Johannes Kepler University Linz, Austria
Email: markus.schedl@jku.at

Bruce Ferwerda

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

Abstract—In this paper, we describe the *LFM-1b User Genre Profile* dataset. It provides detailed information on musical genre preferences for more than 120,000 listeners and links to the LFM-1b dataset. We created the dataset by exploiting social tags, indexing them using two genre term sets, and aggregating the resulting annotated listening events on the user level. We foresee several applications of the dataset in music retrieval and recommendation tasks, among others to build and evaluate decent user models, to alleviate cold-start situations in music recommender systems, and to increase their performance using the additional abstraction layer of genre. We further present results of statistical analyses of the dataset, regarding genre preferences and their consistencies. We do so for the entire user population and for user groups defined by demographic similarities. Moreover, we report interesting insights about correlations between musical preferences on the genre level.

I. INTRODUCTION AND CONTEXT

The importance of considering user characteristics in recommender systems has been highlighted many times [1]. In the music domain, relying only on listening histories or user ratings is nevertheless still the most widely adopted approach to build collaborative filtering algorithms, even though recent work shows that integrating additional listener or listening information is beneficial [2], [3].

Unlike in the text IR domain, in music IR and recommendation, we observe a lack of standardized music datasets, in particular sets that include detailed listener characteristics, going beyond rating or playcount information or social tags. One notable exception is the recently presented LFM-1b dataset [4], which aggregates information about more than one billion listening events gathered from more than 120,000 Last.fm users. Its remarkable feature, in addition to its size, is additional listener-centric information, such as mainstreamness and diversity of a listener’s taste, next to standard demographic data. However, while the dataset offers information on the track, album, and artist level, it lacks aggregate information on a higher level, i.e. genre. Having access to such information would allow to create detailed genre profiles on the user level and to conduct comprehensive analysis, retrieval, and recommendation experiments on a large scale.

In this resource paper, we therefore propose an extension to the LFM-1b set, referred to as *LFM-1b User Genre Profile* dataset. Please note that the LFM-1b User Genre Profile is considered derivative work according to paragraph 4.1 of

Last.fm’s API Terms of Service and can therefore be “published, distributed or otherwise communicated to the public in any media known now”.¹ In the following, we detail data acquisition, creation, and content of the dataset (Section II), provide insights gained through statistical analyses of the genre profiles (Section III), and conclude with a discussion of the dataset’s limitations and possible extensions (Section IV).

II. DATASET DESCRIPTION

A. Data Acquisition and Processing

Taking as input the list of artists in the LFM-1b dataset, we exploit the Last.fm API endpoint `artist.getTopTags` to fetch the most important user-generated tags for each artist, together with their weights (in the range [0,100]).² Subsequently, we create two index term sets, one comprising 20 main genres from Allmusic (rnb, rap, electronic, rock, new age, classical, reggae, blues, country, world, folk, easy listening, jazz, vocal, children’s, punk, alternative, spoken word, pop, and heavy metal), the other consisting of 1,998 genres and styles from Freebase, which are partly very specific (e.g., visual kei, hoedown, or technical death metal). We casefold tags and index terms and describe each artist by a weighted bag-of-words representation of genres. Thereafter, we consider each user’s playcount vector over artists and compute, for each index term set, two variants of a *genre profile*: an unweighted and a playcount-weighted. The former treats artists irrespective of their playcounts, i.e., the genre tags of an artist listened to once contribute to the user’s genre profile in the same way as those listened to many times. In the playcount-weighted variant, in contrast, each artist’s genre occurrence is multiplied with the respective playcount value of the user for that artist. This procedure results in a genre profile for each user, which we represent as k dimensional feature vector over the k genres in the corresponding index term set. We refrain from normalizing the genre profiles in the dataset to avoid losing information about the playcounts, but do so on the user level (sum-to-1) for the statistical analyses we report below.

¹<http://www.last.fm/api/tos>

²We decided against targeting the track level, because much fewer tags are available on that level than on the artist level.

File	Content
LFM-1b_artist_genres_allmusic.txt	genre annotators for artists, using Allmusic genres as index terms; format: artist \t [genre-id \t]*, where genre-id maps to line number in genres_allmusic.txt, the first line indexed by 0
LFM-1b_artist_genres_freebase.txt	genre annotators for artists, using Freebase genres and styles as index terms; format: same as above
LFM-1b_UGP_noPC_allmusic.txt	actual user genre profiles based on the Allmusic index terms, without weighing w.r.t. playcounts; format: user-id \t $occ(g_1, u)$ \t ... \t $occ(g_{ G }, u)$, where $occ(g, u)$ denotes the number of occurrences of genre $g \in G$ aggregated over all artists listened to by user u
LFM-1b_UGP_noPC_freebase.txt	same as above, but using Freebase index terms
LFM-1b_UGP_weightedPC_allmusic.txt	same as LFM-1b_UGP_noPC_allmusic.txt, but genre occurrences are weighted with the respective playcount value of the artist, listened to by user u
LFM-1b_UGP_weightedPC_freebase.txt	same as above, but using Freebase index terms
genres_allmusic.txt	index terms of 20 Allmusic genres; format: genre
genres_freebase.txt	index terms of 1,998 Freebase genres and styles; format: genre or style
user_sets_min[100,200,500,1000]/A-(s,e)_G-[m,f]_C-[country].txt	files containing the user-ids of various user groups, created based on age (A), gender (G), and country (C) information: s=start, e=end, m=male, f=female, country=code according to ISO 3166-1 alpha-2 country codes; organized in directories specifying the minimum number of users in a group for the group to be included; format: user-id \t country \t age \t gender

TABLE I: Description of the files constituting the dataset.

B. Content Description

The dataset can be downloaded as compressed file from <http://www.cp.jku.at/datasets/LFM-1b/> and comprises the components given in Table I. The table also presents a detailed description of the structure of the included files and should be self-explaining. As described above, we use two index term sets (from Allmusic and from Freebase) and provide unweighted and playcount-weighted user genre profiles. Furthermore, in order to enable experiments on different groups of users, we create and include subsets of users with respect to similar age, same gender, and same country. To this end, we analyze all combinations of age, gender, and country, but only include groups with a minimum of 100, 200, 500, and 1,000 users. For instance, the file `user_sets_min500/A-(18,21)_G-m_C-UK.txt` contains all user-ids and demographic information for male users at least 18 and at most 21 years old who indicate to live in the UK. This group consists of 597 listeners, which exceeds the threshold of 500, reflected in the directory name.

In addition, we include a Python script which shows how to load the dataset and, based on demographics extracted from the main LFM-1b dataset, performs the statistical experiments the results reported in the next section are based on.

III. STATISTICAL ANALYSIS

In the following, we present results of statistical analyses conducted to obtain insight into the distribution of listening events over genres per user group, the consistency of genre

preferences within groups (measuring agreement by Krippendorff's α [5]), and correlations between genres (measured by Pearson's correlation coefficient). We focus on the data obtained using the Allmusic genre index and exclude from the detailed analysis genres whose overall share among all users' listening events falls below 3%. These are easy listening (2.29%), world (1.76%), classical (1.41%), country (1.40%), reggae (1.25%), vocal (1.18%), new age (0.84%), spoken word (0.24%), and children's (0.02%).

A. Genre Profiles for User Groups

We define a user group as a subset of users with same country or gender, or similar age.³ Table II shows the genre profiles for all users and user groups with more than 1,000 members. Per user group, the mean share of listening events over genres in percent is given. In addition, the last column contains agreement score, i.e. Krippendorff's α values.

1) *Overall genre preferences and agreement:* The first row in Table II, which shows the overall genre distribution, reveals that the top genres listened to by the sample of Last.fm users in LFM-1b are rock (18.27%), alternative (16.75%), and pop (13.64%). Furthermore, with an overall agreement score of $\alpha = 0.493$, moderate agreement in genre preferences can be observed, according to [6]. This overall agreement is substantially lower than all agreements within user groups, except for the age group (41,50) where it is only slightly above

³We also consider combinations of groups (e.g., users in same country and with similar age). Due to space limitations, the respective results are only available on request by mail, though.

the overall α . We hence conclude that genre preferences are indeed more homogeneous for people in the same country, with the same gender, or similar age.

2) *Intra-group genre preferences and agreement*: To analyze the results on the genre level, we highlight for each genre the user group with highest and with lowest share of listening events, cf. blue and red values in Table II, respectively. We do so for each category of user groups (defined by country, age, or gender). Focusing on the *country*, we can for instance see that rnb is almost twice as popular in the US than in Russia, metal about three times as popular in Finland than in the US. Highest agreement in genre preferences is found in the UK, Brazil, and Sweden ($\alpha \geq 0.58$), lowest in Germany and Finland ($\alpha \leq 0.52$).

Looking at the different *age* groups, we observe a continuous preference increase from young to old users for the genres blues and jazz, while a steady decrease for the genres rap, rock, punk, alternative, and metal is revealed. Agreement in listening preference is quite stable for different age groups ($\alpha \approx 0.55$), except for the group (41,50) for which it is much lower ($\alpha = 0.50$). This group therefore shows a higher diversity in their music taste.

Differences can also be made out with respect to *gender*. Most of them range below one percentage point though, except for metal where a clear preference of males is evident (2.28 pp higher share for males) and pop which is clearly preferred by women (2.54 pp higher share for females). Notably, the highest agreement in preference over all user groups can be found among female listeners ($\alpha = 0.626$).

3) *Consistency of music preferences*: To quantify the consistency of the genre profiles for each genre within user groups, we separately show in Table III the standard deviations as well as the fraction between standard deviations and means (σ/μ), the latter to more easily compare the standard deviations between genres. Considering the entire user population, we observe that profiles for alternative, rock, and pop are the most consistent ones with their average standard deviation staying below half of the corresponding mean ($\sigma/\mu = 0.33$, 0.35, and 0.43, respectively). On the other hand, metal (1.76), rnb (1.42), and rap (1.44) are least consistent, overall.

Due to space limitations, we cannot discuss all consistency results here, but would like to highlight some interesting observations. Note that we always compare within-group consistencies to the overall genre consistencies given in the first data row of Table III. Therefore, for the individual user groups, the table shows in parentheses the difference between the group-specific σ/μ values and the overall σ/μ values, to more easily compare the results. We denote this difference in the following by $\Delta \frac{\sigma}{\mu}$. Negative values therefore indicate higher consistency within the respective group than overall. Positive values indicate lower consistencies.

While metal shows the highest variation in genre profiles of the entire population ($\sigma/\mu = 1.76$), Fins ($\Delta \frac{\sigma}{\mu} = -0.64$), Ukrainians (-0.52), Russians (-0.47), and Poles (-0.46) have a quite stable share in their listening profile. On the other hand, this genre's stability is lowest among US-Americans

(0.15) and the British (0.10). For pop, only Finland has a relatively stable share (0.11). Russians have highly diverse preferences for rnb (0.27) and folk (0.18), but rather stable ones for electronic (-0.11). Brazilians are consistent in their preferences for rap (-0.15) and blues (-0.17). Polish listeners are quite diverse with respect to rap preferences (0.23).

With regard to age, we observe higher than overall preference stability over all age groups for the majority of genres, except for rnb, where it is considerably lower over all groups and pop, where it is lower for the eldest. Interestingly, for jazz, and to a smaller extent for blues and folk, consistency increases with age. An inverse trend is revealed for metal, a genre for which the younger have more stable preferences.

In comparison to the overall genre consistency, female listeners' taste is particularly stable for rap, electronic, blues, and punk ($\Delta \frac{\sigma}{\mu} < -0.10$), also not substantially less stable for any other genre. On the other hand, males are much less consistent in their genre preferences. They particularly disagree in their preference for rnb (0.16). Only for metal (-0.28) a substantially higher agreement than overall is evident.

B. Correlations Between Genres

To investigate which genres users tend to have a joint preference for, we compute Pearson's correlation coefficient between the genre profiles of all users in the dataset. We again exclude genres whose overall share among all listening events falls below 3%. Table IV shows the pairwise correlations and highlights the highest and lowest values in each row, i.e. genre. Overall, the highest correlations, both positive and negative, are found for metal, respectively with rock (0.505) and pop (-0.518). Another correlation almost as high is found between rnb and rap (0.485). Several others are around 0.4 (punk—alternative, jazz—blues) or around -0.4 (rock—jazz, rock—rnb, rock—rap, blues—electronic). All remaining ones are well below an absolute value of 0.4. Please note that all correlations are significant at $p < 10^{-5}$, except for correlation between rnb and blues ($p = 0.0179$).

IV. LIMITATIONS AND FUTURE WORK

Even though we are sure that the proposed dataset is a highly valuable extension to the LFM-1b set, we do not want to conceal its limitations. First of all, the usage of Last.fm data obviously introduces a community bias. For instance, studies showed that listeners of classical music are underrepresented [7], whereas fans of metal and alternative music are overrepresented on the platform [8]. The sample of people present in the dataset at hand (and in general in the LFM-1b set) does therefore not generalize to the population at large. It can nevertheless give an indication of the user composition of a typical music platform. Since the extraction of genre profiles relies on the availability of user-generated tags for the artists in the collection, largely unknown artists may therefore not be represented accurately. However, since these artists are also listened to very infrequently, this fact does not substantially affect the results. Finally, the quality of the index terms, in particular from Freebase, could be

country	age	gender	users	mb	rap	elect.	rock	blues	folk	jazz	punk	altern.	pop	metal	α
-	-	-	120175	3.34	3.41	11.18	18.27	3.28	5.61	3.97	6.19	16.75	13.64	3.98	0.493
US	-	-	10255	3.00	3.22	11.17	18.82	3.07	6.06	3.79	7.53	17.69	13.56	3.29	0.554
RU	-	-	5024	1.55	3.10	14.30	20.60	2.28	4.58	3.03	7.76	18.14	10.58	6.10	0.564
DE	-	-	4578	1.96	3.15	11.90	19.80	2.59	5.67	3.10	7.93	17.26	12.02	6.00	0.510
UK	-	-	4534	2.88	2.76	12.08	18.47	3.10	5.49	4.02	7.32	18.10	13.55	3.35	0.582
PL	-	-	4408	2.18	3.81	11.14	19.45	2.72	4.85	3.49	7.28	19.08	10.96	7.19	0.503
BR	-	-	3886	2.88	1.90	8.29	19.91	3.26	6.05	3.47	7.49	18.72	13.92	5.92	0.586
FI	-	-	1409	1.88	3.40	11.55	21.45	2.20	4.95	2.92	6.56	16.41	11.48	9.85	0.520
NL	-	-	1375	2.64	2.70	11.81	18.18	3.65	6.17	4.20	5.64	17.18	13.37	4.32	0.532
ES	-	-	1243	2.41	2.09	9.86	19.64	3.25	6.07	3.71	6.60	16.95	14.22	5.12	0.560
SE	-	-	1231	2.29	2.60	12.01	19.03	3.07	6.12	3.53	6.15	17.44	14.11	4.82	0.584
UA	-	-	1143	1.69	2.82	13.42	20.86	2.46	4.92	3.13	7.25	18.16	10.56	6.64	0.565
CA	-	-	1077	2.20	2.89	11.76	19.16	2.78	6.37	3.53	7.48	18.26	13.02	4.35	0.575
FR	-	-	1055	2.87	3.44	12.77	17.58	3.25	5.68	4.71	5.55	16.89	12.99	3.73	0.535
-	(6,17)	-	3416	2.88	3.63	10.70	20.27	2.19	4.61	2.45	8.66	19.18	13.26	5.98	0.544
-	(18,21)	-	13784	2.48	3.37	11.49	20.03	2.42	5.17	2.98	8.05	18.69	12.60	5.72	0.554
-	(22,25)	-	13204	2.21	2.83	11.84	19.77	2.65	5.67	3.44	7.33	18.22	12.35	5.60	0.562
-	(26,30)	-	7745	2.19	2.59	12.18	19.23	2.88	5.88	3.90	6.77	17.76	12.36	5.11	0.552
-	(31,40)	-	5113	2.31	2.35	11.84	18.70	3.41	6.04	4.43	6.08	16.99	12.84	4.56	0.535
-	(41,50)	-	1662	2.99	1.75	10.24	18.21	4.60	5.95	5.26	4.69	14.83	13.94	3.51	0.496
-	-	m	39969	2.23	3.08	11.80	19.66	2.93	5.36	3.70	7.25	17.55	11.94	5.84	0.511
-	-	f	15802	2.90	2.40	10.95	18.91	2.81	6.23	3.43	6.89	18.63	14.48	3.56	0.626

TABLE II: Genre profiles (playcount-weighted) of top demographic groups including at least 1,000 users. Shares are given in average percentages over genres. Blue font is used to indicate **highest share** per genre within each category of user groups (country, age, gender), red to indicate **lowest share**. The last column contains agreement scores per user group (Krippendorff's α).

country	age	gender	users	mb	rap	elect.	rock	blues	folk	jazz	punk	altern.	pop	metal
-	-	-	120175	3.34 (1.42)	3.41 (1.44)	11.18 (0.74)	18.27 (0.35)	3.28 (1.01)	5.61 (0.75)	3.97 (0.85)	6.19 (0.89)	16.75 (0.33)	13.64 (0.43)	3.98 (1.76)
US	-	-	10255	4.54 (0.09)	4.40 (-0.07)	7.46 (-0.07)	5.61 (-0.05)	3.16 (0.02)	4.12 (-0.07)	3.12 (-0.03)	6.15 (-0.08)	4.97 (-0.05)	5.43 (-0.03)	6.28 (0.15)
RU	-	-	5024	2.62 (0.27)	4.64 (0.06)	9.11 (-0.11)	6.52 (-0.03)	2.47 (0.08)	4.27 (0.18)	2.84 (0.08)	6.35 (-0.08)	5.17 (-0.04)	5.34 (0.08)	7.87 (-0.47)
DE	-	-	4578	3.10 (0.16)	4.34 (-0.06)	8.72 (-0.01)	6.24 (-0.03)	2.77 (0.06)	5.05 (0.14)	3.04 (0.13)	6.29 (-0.10)	5.38 (-0.02)	5.90 (0.06)	8.57 (-0.33)
UK	-	-	4534	4.04 (-0.02)	3.95 (-0.00)	7.92 (-0.09)	5.37 (-0.06)	2.83 (-0.09)	3.56 (-0.10)	3.13 (-0.07)	5.69 (-0.12)	4.79 (-0.06)	5.00 (-0.06)	6.21 (0.10)
PL	-	-	4408	3.28 (0.09)	6.37 (0.23)	8.25 (-0.00)	6.82 (0.00)	2.60 (-0.05)	3.87 (0.05)	2.97 (-0.00)	5.75 (-0.10)	5.22 (-0.06)	5.70 (0.09)	9.38 (-0.46)
BR	-	-	3886	4.34 (0.09)	2.44 (-0.15)	5.51 (-0.08)	6.18 (-0.04)	2.73 (-0.17)	4.10 (-0.07)	2.96 (0.00)	5.64 (-0.14)	4.74 (-0.08)	5.68 (-0.02)	8.58 (-0.31)
FI	-	-	1409	2.66 (-0.01)	4.84 (-0.01)	8.03 (-0.05)	6.40 (-0.05)	2.45 (0.11)	4.16 (0.09)	2.55 (0.02)	4.99 (-0.13)	4.97 (-0.03)	6.14 (0.11)	11.06 (-0.64)
NL	-	-	1375	3.34 (-0.16)	3.99 (0.04)	8.34 (-0.04)	5.78 (-0.03)	3.05 (-0.17)	4.06 (-0.09)	3.04 (-0.13)	5.10 (0.01)	4.90 (-0.04)	5.32 (-0.03)	7.65 (0.01)
ES	-	-	1243	3.46 (0.02)	3.02 (0.01)	6.04 (-0.13)	5.95 (-0.04)	2.94 (-0.10)	4.05 (-0.08)	3.06 (-0.03)	5.53 (-0.06)	5.19 (-0.02)	6.13 (0.00)	8.60 (-0.08)
SE	-	-	1231	2.79 (-0.20)	3.43 (-0.12)	7.43 (-0.12)	5.67 (-0.05)	2.67 (-0.14)	4.26 (-0.06)	2.55 (-0.13)	4.70 (-0.13)	4.59 (-0.07)	5.73 (-0.02)	7.77 (-0.15)
UA	-	-	1143	2.78 (0.22)	4.25 (0.07)	8.60 (-0.10)	6.61 (-0.03)	2.51 (0.01)	4.55 (0.17)	2.91 (0.08)	5.95 (-0.07)	5.10 (-0.05)	5.23 (0.07)	8.26 (-0.52)
CA	-	-	1077	3.10 (-0.01)	3.82 (-0.11)	7.26 (-0.13)	5.41 (-0.06)	2.51 (-0.11)	4.56 (-0.04)	2.92 (-0.02)	6.15 (-0.07)	4.90 (-0.06)	5.32 (-0.02)	7.58 (-0.02)
FR	-	-	1055	3.80 (-0.10)	4.70 (-0.07)	7.98 (-0.12)	6.04 (-0.00)	2.79 (-0.15)	3.73 (-0.10)	3.23 (-0.17)	4.62 (-0.06)	4.59 (-0.06)	5.25 (-0.02)	6.76 (0.05)
-	(6,17)	-	3416	4.44 (0.12)	5.41 (0.05)	7.30 (-0.06)	6.51 (-0.02)	2.40 (0.09)	4.06 (0.13)	2.52 (0.18)	6.36 (-0.16)	5.29 (-0.05)	6.39 (0.05)	8.92 (-0.27)
-	(18,21)	-	13784	3.81 (0.12)	4.96 (0.03)	7.46 (-0.09)	6.11 (-0.04)	2.52 (0.03)	4.25 (0.07)	2.65 (0.04)	6.18 (-0.13)	4.93 (-0.07)	5.90 (0.04)	8.41 (-0.29)
-	(22,25)	-	13204	3.39 (0.11)	3.97 (-0.03)	7.65 (-0.10)	5.88 (-0.05)	2.52 (-0.06)	4.25 (-0.00)	2.77 (-0.05)	5.78 (-0.11)	4.70 (-0.07)	5.62 (0.03)	8.35 (-0.27)
-	(26,30)	-	7745	3.39 (0.12)	3.79 (0.02)	8.49 (-0.05)	5.95 (-0.04)	2.53 (-0.13)	3.99 (-0.07)	2.93 (-0.10)	5.30 (-0.11)	4.68 (-0.07)	5.47 (0.01)	7.87 (-0.22)
-	(31,40)	-	5113	3.57 (0.12)	3.31 (-0.03)	8.33 (-0.04)	6.00 (-0.02)	3.12 (-0.09)	3.96 (-0.10)	3.18 (-0.13)	4.88 (-0.09)	4.84 (-0.04)	5.47 (-0.00)	7.92 (-0.02)
-	(41,50)	-	1662	4.32 (0.02)	2.89 (0.21)	8.68 (0.10)	6.72 (0.02)	3.70 (-0.20)	4.13 (-0.06)	3.81 (-0.13)	4.73 (0.11)	5.16 (0.02)	5.49 (-0.03)	6.51 (0.09)
-	-	m	39969	3.53 (0.16)	4.66 (0.08)	8.51 (-0.02)	6.43 (-0.02)	3.01 (0.02)	4.15 (0.02)	3.15 (0.00)	6.03 (-0.06)	5.15 (-0.04)	5.73 (0.05)	8.62 (-0.28)
-	-	f	15802	4.12 (0.00)	3.15 (-0.12)	6.26 (-0.17)	5.31 (-0.07)	2.47 (-0.13)	4.25 (-0.07)	2.80 (-0.03)	5.35 (-0.12)	4.75 (-0.07)	5.33 (-0.06)	6.33 (0.02)

TABLE III: Standard deviations of playcount-weighted genre profiles as well as $\frac{\sigma}{\mu}$ for entire population (first data row) and $\Delta \frac{\sigma}{\mu}$ for the individual user groups (in parentheses). Colors are used in the same way as in Table II.

	mb	rap	elect.	rock	blues	folk	jazz	punk	altern.	pop	metal
mb		0.485	-0.036	-0.393	-0.007	-0.252	0.168	-0.308	-0.371	0.260	-0.243
rap	0.485		0.072	-0.351	-0.207	-0.326	0.012	-0.090	-0.178	0.021	-0.134
elect.	-0.036	0.072		-0.373	-0.257	-0.032	-0.177	-0.017	-0.114	-0.183	
rock	-0.393	-0.351	-0.336		-0.056	-0.049	-0.397	0.359	0.356	-0.232	0.505
blues	-0.007	-0.207	-0.373	-0.056		0.253	0.364	-0.224	-0.237	0.024	-0.143
folk	-0.252	-0.326	-0.257	-0.049	0.253		0.051	-0.186	-0.080	0.047	-0.120
jazz	0.168	0.012	-0.032	-0.397	0.364	0.051		-0.373	-0.344	-0.041	-0.275
punk	-0.308	-0.090	-0.177	0.359	-0.224	-0.186	-0.373		0.393	-0.205	0.163
altern.	-0.371	-0.178	-0.017	0.356	-0.237	-0.080	-0.344	0.393		0.020	-0.057
pop	0.260	0.021	-0.114	-0.232	0.024	0.047	-0.041	-0.205	0.020		-0.518
metal	-0.243	-0.134	-0.183	0.505	-0.143	-0.120	-0.275	0.163	-0.057	-0.518	

TABLE IV: Correlations between weighted genre profiles.

improved, e.g. near duplicates (electronic vs. electronica) resolved. However, this is a delicate issue as genre definitions are often subject of discussions. For this reason, standard stemming approaches fail. In future work, we will investigate other genre taxonomies and especially hierarchies to provide information on different, but connected granularity levels. We also plan to complement the dataset with annotations other than genre, e.g., instrumentation, geographic terms, or epochs.

REFERENCES

- [1] G. Adomavicius and A. Tuzhilin, *Recommender Systems Handbook*. Springer, 2011, ch. Context-Aware Recommender Systems, pp. 217–253.
- [2] Y. Shi, M. Larson, and A. Hanjalic, “Collaborative Filtering Beyond the User-Item Matrix: A Survey of the State of the Art and Future Challenges,” *ACM Computing Surveys*, vol. 47, no. 1, pp. 3:1–3:45, May 2014. [Online]. Available: <http://doi.acm.org/10.1145/2556270>
- [3] M. Schedl and D. Hauger, “Tailoring Music Recommendations to Users by Considering Diversity, Mainstreamness, and Novelty,” in *Proc. SIGIR*, 2015.
- [4] M. Schedl, “The LFM-1b Dataset for Music Retrieval and Recommendation,” in *Proc. ICMR*, 2016.
- [5] K. Krippendorff, *Content Analysis – An Introduction to Its Methodology*, 3rd ed. SAGE, 2013.
- [6] J. Landis and G. Koch, “The measurement of observer agreement for categorical data,” *Biometrics*, vol. 33, pp. 159–174, 1977.
- [7] M. Schedl and M. Tkalčić, “Genre-based Analysis of Social Media Data on Music Listening Behavior,” in *Proc. ISMM*, 2014.
- [8] P. Lamere, “Social Tagging and Music Information Retrieval,” *J. New Music Research*, vol. 37, no. 2, pp. 101–114, 2008.