

INTERACTIVELY EXPLORING SIMILARITIES BETWEEN MUSIC GENRES BASED ON USER-GENERATED TAGS

Stefan Brandl

Johannes Kepler University (JKU)
Linz, Austria
stefan.brandl@jku.at

Markus Schedl

Johannes Kepler University (JKU) and
Linz Institute of Technology (LIT)
Linz, Austria
markus.schedl@jku.at

ABSTRACT

Searching for music by genre is one of the most common strategies. Knowledge about similarities between (sub-)genres likewise facilitates discovery of new music. However, given the often very fine-grained genre taxonomies used by major music providers (e.g., Spotify organizes their collection according to more than 5,000 micro-genres), grasping the meaning of those genre names is impossible for most users. Addressing this issue, we present Genre Similarity Explorer (GSE), an interactive exploration tool for pairwise genre similarity. Genre similarity is quantified based on co-occurrences of genre tags in a collection of user-generated song annotations.

1. MOTIVATION

While music genre is a debated concept [1,2], music listeners still commonly use genre for search, e.g., in catalogs of music streaming services, and as anchor points to explore a collection. However, genre assignments have to be exclusive by nature to best encapsulate what a song sounds like. Inadvertently, similar sounding songs can therefore be associated with different genres. Therefore, a quantitative measure for genre similarity is useful to better understand the relation of genres to each other, and to discover new genres using a known one as anchor.

2. COMPUTING GENRE SIMILARITIES

Our work assumes that similarities between genres can — to some extent — be expressed by the weighted co-occurrences of genre tags in songs these tags appear in. To given an example, our approach assumes that genres "rap" and "hip hop" are more similar than "rap" and "techno" because the former combination appears in 1.6206% of songs in our dataset while the latter only in 0.0799%, with the caveat that user-generated tags may not reflect musicologically grounded characteristics of song.

For this purpose, we gathered tag annotations for every song in the LFM-2b dataset [3], which was listened to at least 10 times. In total, this results in over 4 million songs for which we gathered tag annotations from Last.fm. Tags are provided in a json-like format where each tag is given a score from 0 to 100, measured as rounded ratio of the number of users who assigned the respective tag, where a value of 100 indicates the most frequently assigned tag.

Since Last.fm's tags are user-created there is a lot of variance and many highly specific tags that occur exactly once. In total there are over 1 million unique tags.

We filtered tags dropping every tag not contained in Everynoise's list of over 5,000 genres¹ and dropping every song with no tags remaining. After this step, 1.6 million songs and 2,723 distinct genres remained, while highly specific tags had been dropped.

All resulting <genre, song, weight> triples are then represented in a matrix of shape $s \times g$, where $s = 1,637,385$ songs and $g = 2,723$ genres. Note that this matrix is highly sparse (99.8885%). Each s -dimensional column vector represents a genre profile, in which each value is in $[1, 100]$. We calculate cosine similarity between all pairs of column/genre vectors a and b as follows:

$$\text{sim}_{\cos}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \cdot \|\mathbf{b}\|} = \frac{\sum_i a_i \cdot b_i}{\sqrt{\sum_i (a_i)^2} \cdot \sqrt{\sum_i (b_i)^2}} \quad (1)$$

3. EXPLORING GENRE SIMILARITIES

We present Genre Similarity Explorer (GSE) as a way to interactively explore genre similarities. GSE is available at <http://www.cp.jku.at/projects/GSE/>. Genre similarities are presented as a matrix, where 1 means highest (cosine) similarity and 0 means lowest similarity. For these matrices, we consider only the top k genres, meaning genres with the highest number of appearances in unique songs. The top 20 most common genres according to <https://www.allmusic.com> are highlighted in bold text.

In Figure 1, which visualizes the similarity matrix for the top 50 genres, we can see that no two genres are very



¹ <https://everynoise.com/everynoiseld.cgi?scope=all>




Figure 1. Similarity matrix for the top 50 genres, i.e. those with the highest number of appearances in unique songs

similar to each other. We measure the highest similarity between "hip hop" and "rap" at 0.3853. Genre-pairs that we expected to have high similarity are "rock" and "alternative rock" (0.3238), "hardcore" and "metal core" (0.3015), "classic rock" and "rock" (0.2765) and "metal" and "trash metal" (0.2707). Unexpectedly, "downtempo" and "electronica" (0.1408) as well as "singer-songwriter" and "rock" (0.1247) scored higher than expected. We can explain this by investigating co-occurrences: 48.31% of "singer-songwriter" songs are also tagged "rock", while 36.67% of all "downtempo" songs are tagged "electronica".

On the other hand, combinations such as "electro" and "electronica" (0.1526), "country" and "folk" (0.0737) or even genres that sound very entangled by their names alone such as "punk" and "post-punk" (0.0655) scored lower than expected.

4. FUTURE WORK

We contemplate several extensions of GSE. For instance, we plan to extend our approach to audio features by aggregating those features on a per-genre level, enabling an exploration of genre similarities via audio characteristics.

Another avenue for a further extension of GSE is to use genre taxonomies instead of folksonomies, to allow for hierarchical sub-genre exploration and comparison [4].

Currently, we are also working on a graph-based layout that allows music playback of the most prototypical examples of songs for each genre and pairs of genres, making the exploration more exciting.

5. ACKNOWLEDGMENTS

This research received support by the Austrian Science Fund (FWF): P33526.

6. REFERENCES

- [1] B. L. Sturm, "Classification accuracy is not enough - on the evaluation of music genre recognition systems," *J. Intell. Inf. Syst.*, vol. 41, no. 3, pp. 371–406, 2013. [Online]. Available: <https://doi.org/10.1007/s10844-013-0250-y>
- [2] M. Schedl, A. Flexer, and J. Urbano, "The neglected user in music information retrieval research," *J. Intell. Inf. Syst.*, vol. 41, no. 3, pp. 523–539, 2013. [Online]. Available: <https://doi.org/10.1007/s10844-013-0247-6>
- [3] A. B. Melchiorre, N. Rekabsaz, E. Parada-Cabaleiro, S. Brandl, O. Lesota, and M. Schedl, "Investigating gender fairness of recommendation algorithms in the music domain," *Information Processing & Management*, vol. 58, no. 5, p. 102666, 2021.
- [4] B. L. Sturm and F. Gouyon, "Revisiting inter-genre similarity," *IEEE Signal Process. Lett.*, vol. 20, no. 11, pp. 1050–1053, 2013. [Online]. Available: <https://doi.org/10.1109/LSP.2013.2280031>