

Leveraging Microblogs for Spatiotemporal Music Information Retrieval

Markus Schedl

Department of Computational Perception
Johannes Kepler University
Linz, Austria
<http://www.cp.jku.at>

Abstract. We present results of text data mining experiments for music retrieval, analyzing microblogs gathered from November 2011 to September 2012 to infer *music listening patterns* all around the world. We assess *relationships between particular music preferences and spatial properties*, such as month, weekday, and country, and the *temporal stability of listening activities*. The findings of our study will help improve music retrieval and recommendation systems in that it will allow to incorporate geospatial and cultural information into models for music retrieval, which has not been looked into before.

1 Introduction

Exploiting social media to enrich retrieval methods by including user-generated data is a quite recent strategy. In particular in Music Information Retrieval (MIR), work that leverages social media data is almost non-existent, except for publications that make use of the music service `Last.fm`¹. On the other hand, retrieval methods that take into account cultural differences in the perception and consumption of music are highly desired [6]. Aiming to narrow this gap, we first determine music-related microblogs, extract from them information about location and music items (Section 2), annotate each item, and perform an analysis of the resulting annotated data collection that will eventually lead to *spatiotemporal music retrieval* methods. In particular, we investigate the *relationship between music preference and spatial properties* (Section 3) and the *temporal stability of listening patterns* (Section 4). Related work is sketched in Section 5; conclusions are drawn in Section 6. To foster reproducibility and further experimentation, the data collection named `MusicMicro 11.11-09.12` can be downloaded².

2 Determining Music Listening Patterns

We monitored the `Twitter` streaming API from November 2011 to September 2012, using the `Spritzer` feed³ which contains a random selection of 1–2% of all tweets. To determine tweets related to music, we filtered the stream

¹ <http://last.fm>

² <http://www.cp.jku.at/people/schedl/data/MusicMicro/musicmicro.html>

³ <http://gnip.com/twitter/spritzer>

with keywords typically used to communicate music listening activities, such as `#nowplaying`, `#np`, or `#itunes`. We further excluded tweets that did not contain a location. The resulting set still contains a lot of irrelevant microblogs⁴. For this reason, we extracted track and artist lists from `Musicbrainz`⁵ and applied a pattern-based, multistage entity detection technique. Those tweets that could be identified as pointing to a music artist and/or song were retained. We furthermore excluded all tweets posted by users having the substring “radio” in their user names to suppress radio stations as they may distort the results. After these rigorous filtering steps, a total of 594,306 microblogs by 136,866 users, in which we identified 19,529 unique artist names, remained for subsequent investigation. Using the `Yahoo! PlaceFinder` API⁶, we were able to connect the tweets to 20,722 different cities in 180 countries.

Since we ultimately aim at geospatial, semantic music search, we annotated each of the microblogs identified as described above with semantic tags. To this end, we gathered a set of 288 moods from `Allmusic`⁷. This set was then used to index collaborative tag lists extracted from `Last.fm`⁸ for each artist, allowing to project the artist/tweet space to a semantic tag space. This projection (i) effectively reduces computational complexity and (ii) allows to retrieve music items by semantic labels, which is important when the user does not know the music item she is searching for.

3 Relation: Listening Preferences – Spatial Properties

To obtain the overall distribution of music listening activity, we compute a normalized, worldwide *tag distribution vector* \mathbf{T} as $\frac{\sum_{c \in C} \sum_{u \in U(c)} \frac{\mathbf{T}(c,u)}{|A(u)|}}{\sum_{c \in C} |U(c)|}$, where C is the set of countries, $U(c)$ are the unique users in country c , and $\mathbf{T}(c,u)$ is the tag distribution vector of user u in country c , i.e., a vector containing aggregated tag occurrence counts for $A(u)$ (multiset of artists user u has listened to). The normalized tag distribution vector $\mathbf{T}(c)$ of a particular country c is likewise defined as $\mathbf{T}(c) = \sum_{u \in U(c)} \frac{\mathbf{T}(c,u)}{|A(u)|}$.

Figure 1 depicts the most frequent 20 tags and the five countries with highest tweeting activities. We see that preference for particular *music mood varies considerably between countries*, a fact that will allow future music retrieval approaches to personalize results by incorporating spatial information. The bar chart shows the relative difference between $\mathbf{T}(c)$ and \mathbf{T} as the normalized Manhattan distance: $\frac{L_1(\mathbf{T}(c), \mathbf{T})}{\mathbf{T}}$. For instance, “lyrical” music is listened to in the USA 116% more frequently than the worldwide average suggests. “Smooth” music is listened to in Brazil 53% less frequently than the worldwide average.

Quantitative analysis of listening correlation between countries (independent of time), similar to the approach presented in the previous section, shows indeed a relatively low mean correlation and a high standard deviation: $\bar{\rho} = 0.7382 (\pm 0.2012)$, $\min(\rho) = -0.0238$, $\max(\rho) = 1.0$.

⁴ For instance, `#np` is also used to indicate playing video games.

⁵ <http://musicbrainz.org>

⁶ <http://developer.yahoo.com/geo/placefinder>

⁷ <http://allmusic.com>

⁸ <http://last.fm/api>

4 Temporal Stability of Listening Activities

To investigate whether listening patterns are consistent over time (independent of the country), we compute Pearson’s correlation coefficient between the normalized tag vectors for each pair of months: $\rho(\overline{\mathbf{T}(m_i)}, \overline{\mathbf{T}(m_j)})$. The mean ρ value over all pairs of months is $\bar{\rho} = 0.9974 (\pm 0.0021)$; the maximum is 0.998 (January vs. February 2012); the minimum is 0.9902 (November 2011 vs. September 2012). We thus conclude that *music listening patterns are highly independent of month*.

Analyzing analogously the correlation between tag vectors aggregated at the level of weekdays, we find that listening patterns are highly correlated between weekdays. In fact, the mean correlation between aggregated tag vectors among workdays only and among weekends only is 0.9999. The corresponding mean correlation between workdays and weekends is 0.9993. Nevertheless, due to the large sample size, the difference between the two correlation values is significant, according to Fisher’s r-to-z transformation. There is hence a *significant difference in listening behavior between workdays and weekends*, which may be explained by different music preferences during working and partying hours.

5 Related Work

There meanwhile exists extensive literature on the topics of social media mining (SMM) and –retrieval. For instance, Alhadi et al. recently presented an approach to predict interesting tweets, based on the retweet activity of users [1]. A comprehensive overview of related methods can be found in [4].

When it comes to geospatial analysis for music retrieval, in contrast, only very recently researchers have looked into culture-specific music creation and listening [6]. Serra’s work so far focused on musical properties of non-western music. Bridging SMM and MIR research, Zangerle et al. present an approach to music recommendation based on microblog co-occurrences of artist and track names [7]. Schedl and Hauger present an approach to extract music genre patterns for different regions of the world [5] and a user interface to explore these patterns [3]. Unlike in the paper at hand, previous work typically focused on the aspect of music genre, which is known to be an ill-defined concept [2]; whereas we present a more general approach and further consider spatiotemporal aspects.

6 Conclusion and Outlook

Based on a microblog collection covering 11 months, we presented an approach to annotate tweets with music artist/track names, semantic tags, and geographic data. We then use these annotations to *infer music listening activities* and relate them to spatiotemporal properties. We found that *music listening is independent of month*. There is a statistically significant *difference between workdays and weekends and between countries* (irrespective of time), though.

As part of future work, we will exploit more specific information – tracks to describe listening activities and cities to describe corresponding locations. Incorporating the findings of this study, we will elaborate music retrieval systems providing serendipitous experiences [8]. We will further look into domains other than music, e.g., movies, politicians, or shares.

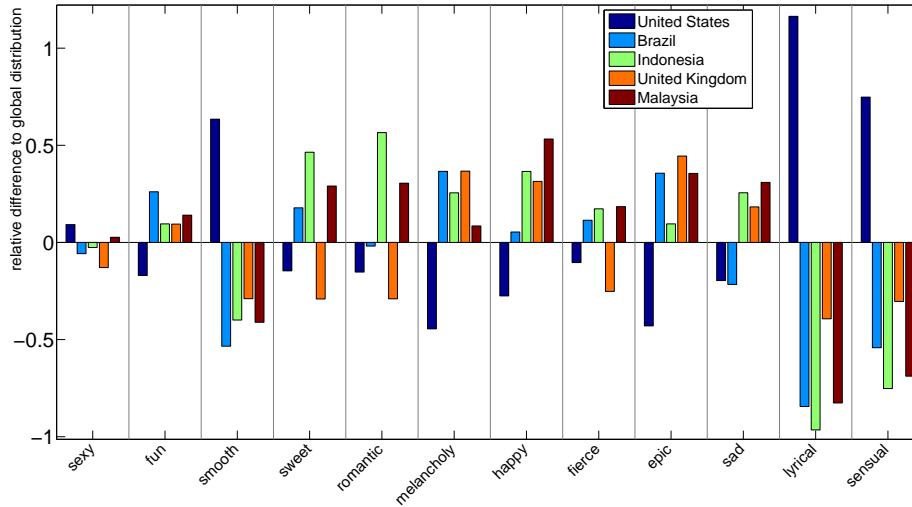


Fig. 1. Music mood distribution for top countries and top tags.

7 Acknowledgments

This research is supported by the Austrian Science Funds (FWF): P22856-N23.

References

1. A. C. Alhadi, T. Gottron, J. Kunegis, and N. Naveed. LiveTweet: Monitoring and Predicting Interesting Microblog Posts. In *Proc. ECIR*, Apr 2012.
2. J.-J. Aucouturier and F. Pachet. Representing Musical Genre: A State of the Art. *Journal of New Music Research*, 32(1):83–93, 2003.
3. D. Hauger and M. Schedl. Exploring Geospatial Music Listening Patterns in Microblog Data. In *Proc. AMR*, Oct 2012.
4. N. Ramzan, R. van Zwol, J.-S. Lee, K. Clüver, and X.-S. Hua, editors. *Social Media Retrieval*. Springer, Nov 2012.
5. M. Schedl and D. Hauger. Mining Microblogs to Infer Music Artist Similarity and Cultural Listening Patterns. In *Proc. WWW Workshop: AdMIRE*, Apr 2012.
6. X. Serra. Data Gathering for a Culture Specific Approach in MIR. In *Proc. WWW Workshop: AdMIRE*, Apr 2012.
7. E. Zangerle, W. Gassler, and G. Specht. Exploiting Twitter’s Collective Knowledge for Music Recommendations. In *Proc. WWW Workshop: #MSM*, Apr 2012.
8. Y. C. Zhang, D. O. Seaghdha, D. Quercia, and T. Jambor. Auralist: Introducing Serendipity into Music Recommendation. In *Proc. WSDM*, Feb 2012.