Mass Media Musical Meaning: Opportunities from the Collaborative Web

Cynthia C. S. Liem*

Multimedia Computing Group, Delft University of Technology
c.c.s.liem@tudelft.nl

Abstract. In the digital domain, music is usually studied from a positivist viewpoint, focusing on general 'objective' music descriptors. In this work, we strive to put music in a more social and cultural context, looking into ways to unify data analysis methods with thoughts from the humanities on musical meaning and significance. More specifically, we investigate whether information in collaborative web resources on movie plot narratives and folksonomic song tags is capable of revealing common associations between these two. Reported initial findings suggest this is indeed the case, which opens opportunities for further work in this area, cross-disciplinary collaborations, and novel contextually oriented music information retrieval application scenarios.

Keywords: music information retrieval, cultural context, narrative elements, collaborative web resources, data science, text retrieval, mass media, musicology

1 Introduction

Music hardly occurs without contextualization. For different types of social occasions, different choices of music are considered appropriate to be played. Our music taste reflects our cultural and social background and sense of belonging. Socially and culturally established connotative associations also strongly are used when music is intended to support a multimedia narrative or persuasive message, such as a movie or a commercial.

The topic of socially and culturally established meaning of music would naturally be more closely related to the humanities than to computer science, and indeed most strongly has been studied under the first discipline. However, if music information retrieval techniques are meant to serve the organization and use of digital music data in real-world use scenarios, this topic cannot be ignored.

While at first, it appears hard to study contextual music usage at scale in the digital domain, in fact some interesting collaborative web resources exist. Conjecturing that contextual, socially established meaning of music is (partially)

* The research leading to these results has received funding from the European Union Seventh Framework Programme FP7 / 2007 - 2013 through the PHENICX project under Grant Agreement no. 601166.

encoded in the way music was used in mass media, in this paper we study how information from the Internet Movie Database (IMDb) and `last.fm` can be used to obtain associations from folksonomic music tags to movie plot narratives, and the other way around. By applying automated text analysis techniques and word cloud visualizations, we will show initial outcomes suggesting that interesting associations between music and non-musical narrative elements can indeed be found in collaborative resources. This opens doors to more sophisticated analysis methods, cross-disciplinary collaboration opportunities, and potential incorporation in novel search engine scenarios.

2 Related work

In musicological study, in the beginning of the 20th century, music was mostly seen as a positivist phenomenon. Under this view, music could be studied in an absolute and independent sense, and was considered to be fundamentally represented by scores. However, from the 1980s onwards, a new stream of thinking emphasizing subjectivity, criticism and value judgements emerged, in which the contextual and social surroundings of music became important, and 'the music itself' even became a taboo concept [9]. As such, culture, context and identity have become major topics in modern musicology research.

In a certain sense, similar developments occurred in the music information retrieval field [6] [3] [4]. Apart from studying aspects of the isolated musical object, the roles of context, usage scenarios and relations to other domains and modalities have become increasingly important. Furthermore, the social web has increasingly been studied as a potential source of information for getting information on context and typical user-entered labels of music objects.

In musicology and psychology, many studies exist on the way music has been used in the context of movies and other forms of multimedia (e.g. [8] [2] [1]). It has been conjectured, and even shown in small-scale studies, that music can influence the perception of a moving image, and the other way around. A particularly interesting categorisation of functions of music in relation to moving images, though not backed up from a data perspective, was made by Lissa [7], listing that music can be used to indicate movement, stylize real sounds, represent space, represent time, communicate meaning through deformation of sound, provide a commentary external to the narrative, serve as music within the narrative, indicate psychology of actors, provide a source of empathy, function as a symbol, anticipate action, and serve as a formal unifying factor. This was later recategorized by Tagg and Clarida [8], who proposed a sign typology for music involving sonic, kinetic and tactile anaphones (sounds that were metaphorical for sonic, kinetic and tactile sensory qualities), genre synecdoche, in which a small glimpse of an exotic genre will evoke the full exotic context, episodic markers announcing musical episode transitions, and style indicators.

Partially inspired by the works mentioned above, in previous work, we performed a crowdsourcing study [5] investigating what categories of narrative elements were commonly associated to production music fragments, and found

that generalization was possible in terms of event structure. In the current work, we do not explicitly solicit conceptions of the crowd, but rather study common connections between narrative (thematic) elements and music song descriptors, expressed in folksonomic terms.

3 Methodology

In this work, we are interested in discovering if collaborative web resources are capable of revealing connections between musical vocabularies and non-musical narrative elements, and vice versa. For this, we consider two major resources: the *Internet Movie Database (IMDb)* and *last.fm*. The IMDb (`http://www.imdb.com`) is a large user-contributed resource on movies and TV series, including listings of many types of metadata attributes, including actors and actresses, but also soundtracks and descriptions of plots. `last.fm` (`http://www.last.fm`) is a social web resource in which users can log and share their music playing behavior, and describe the songs they listen to in tags.

First, we crawled the IMDb for movies which also had soundtrack listings associated to them. For each of these soundtracks, we queried `last.fm` for taggings of soundtrack songs, and kept those movies that had at least one soundtrack with a `last.fm` tag associated to it. 22,357 unique IMDb movies with plot descriptions had at least one soundtrack song with a `last.fm` tag associated to it; in total, considering the soundtracks of all these movies, 265,376 song tags could be found. Subsequently, we addressed two questions:

- *Do different narrative elements emerge for different song tags?* For this, a mapping was made from song tags to all movie plots which had a soundtrack to which this song tag was associated. For example, the song tag 'guitar' will be mapped to all plot descriptions in the corpus of movies which had at least one soundtrack song on which 'guitar' was used as a `last.fm` tag.
- The other way around, *do different song tags emerge for different narrative elements?* For this, we built an Apache Lucene[1] search index, framing our data in an automated text retrieval scenario. Using the map from song tags to movie plots, we treat the song tag as the document key, and all corresponding movie plot stories as document text. Subsequently, we can issue natural language queries to the search index, which will surface the 'documents' matching the query best (so, cases in which movie plot descriptions cause a close match to the issued query), and then consider the corresponding document keys, which consist of song tags. To reduce noise, we only index song tags which were used at least 1000 times in `last.fm`.

4 Initial outcomes

In this section, we report initial outcomes resulting from exploration of the data connections which were established as explained above. Unfortunately, no standard metrics or ground truth rules exist on the relation between music song

[1] `lucene.apache.org`

tags and narrative elements. Therefore, as a first way to still display emergent patterns in the data, in discussing outcomes we will use Wordle[2] word cloud visualizations, which apply common automatic statistical text analysis methods to visualize the most important words in text corpora.

4.1 Do different narrative elements emerge for different song tags?

We take the mapping from song tags to movie plots as described in the previous section, and then examine what kinds of words occur in the aggregated collection of movie plots associated to a certain song tag. In Figure 1, word clouds are displayed based on several song tag queries reflecting different music genres.

Like any other collaborative web resource, both the `last.fm` and IMDb corpora are noisy. Furthermore, especially when many movie plots are found for a song tag, while the text analysis provided by Wordle filters out common stopwords and highly frequent words, certain universal elements which are no stopword, but still occur in almost all movie plot narratives ('one', 'life', 'man', 'woman') stand out. More sophisticated language analysis would be needed to filter those out.

Still, some distinctive characteristics already emerge for different genres. For example, while 'family' occurs in each of the word clouds, the movie plots associated to 'opera' have a stronger connection to family-oriented themes like ' marriage' and 'Christmas'. Movie plots associated to 'salsa' have an explicit relation to dancing. Movie plots associated to 'rap' suggest that younger main characters are involved ('school' is more prevalent than in other tags) and seem to suggest slightly stronger male connotations than the other word clouds ('man' and 'father' are relatively large; 'brother' is equally sized to 'girl', but 'woman', 'girlfriend' and 'mother' are clearly smaller; finally, any given names emerging in the word cloud are male names).

4.2 Do different song tags emerge for different narrative elements?

In the reversed scenario of looking up collections of song tags for non-musical narrative elements, further interesting socially established aspects can be found. Figure 2 shows what song tags are returned by querying the Lucene search index using several names of cities. Again, certain very popular song tags occur for all word clouds (in particular 'rock', the most frequently used tag on `last.fm`). However, it is clear that beyond this, differences occur between the word clouds, revealing connections between geographical locations and typically associated music styles (e.g. 'chanson' for Paris, 'blues' for Chicago, 'anime' for Tokyo). In Figure 3, we show further interesting results for various queries which are non-musical, but do express a narrative context. Again, the broadly popular 'rock' song tag strongly occurs in all results, but next to this, we notice different nuances which indeed represent typical connotations for the given queries (e.g. some 'rougher' genres like metal and punk for 'car chase', genres associated to

[2] `www.wordle.net`

warmer regions like Jamaica, Italy and Brazil for 'beach holiday', and some less rough genres like instrumental, atmospheric, blues and classical for 'candlelight dinner').

5 Conclusion and future opportunities

In this paper, we described early efforts in investigating opportunities to retrieve contextual meaning associations for music from large collaborative web resources. Information on narrative movie plot elements and soundtrack use was obtained from the IMDb, while information on descriptive song tags for soundtrack songs was obtained from `last.fm`. The information from the two data sources was connected by building mappings and text search indices associating song tags to movie plots and vice versa. Initial data visualizations suggest that connotative associations between music tags and narrative plot elements indeed are reflected in these collaborative web resources.

The work reported in this paper is highly novel work-in-progress, and current results thus may still appear somewhat anecdotical. However, we believe this work opens doors to a lot of opportunities and future work. First of all, more sophisticated text analysis methods should be applied to further reveal the potential of the data, and devise more sophisticated models and algorithms to infer cultural views on musical meaning from data. In parallel, we believe concrete possibilities for cross-disciplinary cooperations between computer (data) science and the humanities can be investigated, as well as novel querying paradigms for music information retrieval systems based on cultural context.

References

1. Cohen, A.J.: How music influences the interpretation of film and video: Approaches from experimental psychology. In: Kendall, R., Savage, R.W. (eds.) Selected Reports in Ethnomusicology: Perspectives in Systematic Musicology, vol. 12, pp. 15–36. Department of Ethnomusicology, University of California, Los Angeles (2005)
2. Cook, N.: Analysing musical multimedia. Oxford University Press, New York, USA (1998)
3. Kaminskas, M., Ricci, F.: Contextual music information retrieval: State of the art and challenges. Computer Science Review (2012)
4. Knees, P., Schedl, M.: A survey of music similarity and recommendation from music context data. ACM TOMM 10 (December 2013)
5. Liem, C.C.S., Larson, M.A., Hanjalic, A.: When music makes a scene - Characterizing music in multimedia contexts via user scene descriptions. International Journal of Multimedia Information Retrieval 2, 15–30 (2013)
6. Liem, C.C.S., Müller, M., Eck, D., Tzanetakis, G., Hanjalic, A.: The Need for Music Information Retrieval with User-Centered and Multimodal Strategies. In: Proc. 1st Int. ACM MIRUM workshop. pp. 1–6 (November 2011)
7. Lissa, Z.: Ästhetik der Filmmusik. Henschelverlag, Berlin, Germany (1965)
8. Tagg, P., Clarida, B.: Ten Little Title Tunes — Towards a Musicology of the Mass Media. The Mass Media Scholar's Press, New York, USA & Montreal, Canada (2003)
9. Wiering, F., Volk, A.: Musicology. Tutorial slides, ISMIR 2011

(a) ‘opera’

(b) ‘salsa’

(c) ‘rap’

Fig. 1: Wordle word clouds of aggregated movie plots for several genre-related music song tags.

(a) ‘Paris’

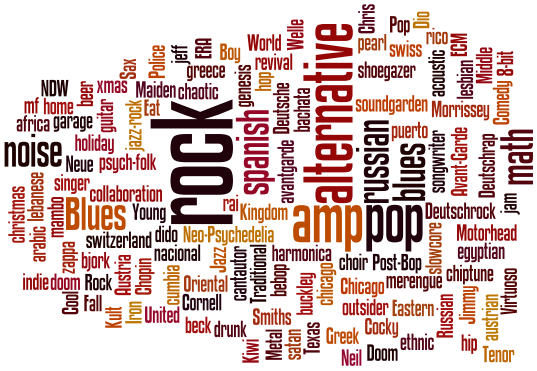

(b) ‘Chicago’

(c) ‘Tokyo’

Fig. 2: Wordle word clouds for song tags retrieved for city name queries on the Lucene search index.

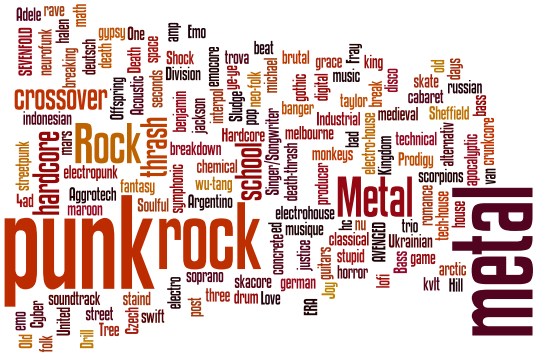

(a) ‘car chase’

(b) ‘beach holiday’

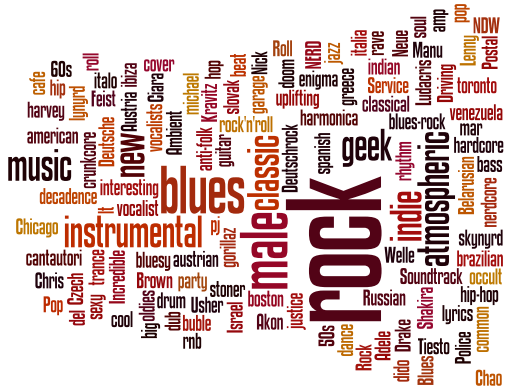

(c) ‘candlelight dinner’

Fig. 3: Wordle word clouds for song tags retrieved for various other non-musical narrative context queries on the Lucene search index.