

Music Similarity and Retrieval

Peter Knees
Department of Computational Perception
Johannes Kepler University Linz
Altenberger Str. 69, 4040 Linz, Austria
peter.knees@jku.at

Markus Schedl
Department of Computational Perception
Johannes Kepler University Linz
Altenberger Str. 69, 4040 Linz, Austria
markus.schedl@jku.at

ABSTRACT

This tutorial serves as an introductory course to the field of and state-of-the-art in music information retrieval (MIR) and in particular to music similarity estimation which is an essential component of music retrieval. Apart from explaining approaches that estimate similarity based on acoustic properties of an audio signal, we review methods that exploit (mostly textual) meta-data from the Web to build representations of music then used for similarity calculation. Additionally, topics such as (large-scale) music indexing, information extraction for music, personalization in music retrieval, and evaluation of MIR systems are addressed.

Categories and Subject Descriptors

H.5.5 [Sound and Music Computing]: Methodologies and techniques; H.3.3 [Information Search and Retrieval]

Keywords

music information retrieval, music similarity, music retrieval, content-based MIR, context-based MIR

MOTIVATION AND CONTENT

Music is omnipresent in today's society, especially on the Web and in social media, and the amount of music available via streaming services, online stores, and platforms like YouTube has skyrocketed over the last couple of years. Music information retrieval (MIR) is a research field that aims – among other things – at making the information contained in ever-growing digital music repositories accessible in an intelligent manner by automatically extracting semantically meaningful information from various representations linked to music entities, such as digital audio files, band Web pages, song lyrics, or tweets on listening activities.

A key approach in MIR is to describe music via computational features, which can be broadly categorized into three classes: *music content*, *music context*, and *user context*. While music content-based features are derived directly

from the audio signal of the music file [1], music context refers to pieces of information that are not encoded in the actual audio file, nevertheless play an important role in human perception of music, such aspects include the meaning of song lyrics, the background of an artist, or even the cover of an album [2]. The user context, in contrast to the other two, includes environmental aspects as well as physical and mental activities of the music listener [3].

The aim of the tutorial is to give a sound and comprehensive, nevertheless easy-to-understand, introduction to MIR. First, we review and discuss the ideas behind the three categories of computational features (content, music context, and user context). Then, we focus on approaches for music similarity estimation, in particular approaches that estimate similarity based on acoustic properties of the actual musical signal and approaches that exploit meta-data from the Web. This includes an introduction to the field of Web-based MIR and a detailed description and comparison of data sources of music context (e.g., Web pages, blogs, micro-blogs, social networks, tags, lyrics, playlists). Additionally, topics such as automatic (large-scale) music indexing, information extraction for music, personalization and user adaptation in music retrieval, as well as the challenging task of evaluating MIR systems beyond the traditional IR-related measures and the difficulties entailed by the need for objective quantification are addressed. We further demonstrate applications such as automatic playlist generators, music search engines, music recommender systems, and intelligent interfaces to music, that utilize the methods presented.

The material presented in the tutorial is available online at <http://www.cp.jku.at/tutorials/sigir2013.html>

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union Seventh Framework Programme FP7 / 2007–2013 through PHENICX project under grant agreement no. 601166

REFERENCES

- [1] M. Casey, R. Veltkamp, M. Goto, M. Leman, C. Rhodes, and M. Slaney. Content-Based Music Information Retrieval: Current Directions and Future Challenges. *Proc IEEE*, 96:668–696, April 2008.
- [2] P. Knees and M. Schedl. A Survey of Music Similarity and Recommendation from Music Context Data. *ACM TOMCCAP*, 2013. Accepted for publication.
- [3] M. Schedl and P. Knees. Personalization in Multimodal Music Retrieval. In *Proc Workshop AMR*, 2011.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR'13, July 28–August 1, 2013, Dublin, Ireland.
ACM 978-1-4503-2034-4/13/07.