

New Paths in Music Recommender Systems Research

A Tutorial Held at ACM RecSys 2017

Markus Schedl
Johannes Kepler University Linz
Linz, Austria
markus.schedl@jku.at

Peter Knees
TU Wien
Vienna, Austria
peter.knees@tuwien.ac.at

Fabien Gouyon
Pandora Inc.
USA
fgouyon@pandora.com

ABSTRACT

The particularities of musical data and its multiple modalities make original contributions possible in many core RecSys topics such as content-based and hybrid recommendation, user modeling, interfaces, and context-aware and mobile recommendations. But more urgently, the current revolution in the music industry represents major opportunities and challenges for recommendation systems in general. Recommendation systems are now central to music streaming platforms, which are rapidly increasing in listenership and becoming the top source of revenue for the music industry. It is increasingly more common for a music listener to simply access music than to purchase and own it in a personal collection. In this scenario, recommendation calls no longer for a one-shot recommendation for the purpose of a track or album purchase, but for a recommendation of a listening experience, comprising a very wide range of challenges, such as sequential recommendation, or conversational and contextual recommendations. Recommendation technologies now impact all actors in the rich and complex music industry ecosystem (listeners, labels, music makers and producers, concert halls, advertisers, etc.). To highlight these developments, we focus on three use cases: automatic playlist generation, context-aware music recommendation, and recommendation in the creative process of music making.

ACM Reference format:

Markus Schedl, Peter Knees, and Fabien Gouyon. 2017. New Paths in Music Recommender Systems Research. In *Proceedings of RecSys '17, Como, Italy, August 27-31, 2017*, 2 pages.
DOI: 10.1145/3109859.3109934

MOTIVATION

In the RecSys community, music is too often treated as “just another item”. Yet, the particularities of music data and its multiple modalities open many opportunities, e.g., to leverage content-based audio features or to build comprehensive listener models that go beyond simple user-item interactions. Furthermore, since it is now increasingly more common for a music listener to simply stream music rather than to purchase and own it, today’s music recommenders need to focus on recommending a listening experience. As a consequence, Music Recommender Systems (MRS) [17] research

has to face a wide range of challenges, such as sequential recommendation/automatic radio stationing, sound recommendation to music producers, conversational and contextual recommendation (cf. Google Home, Siri, or Alexa), or even tasks such as recommending *concerts to users or fans to artists*.

While recommender systems in other domains, such as movies, have already been extensively researched, the still underrepresented topic of music recommendation is a highly important research area for its particularities which require specific considerations beyond application of constantly improving collaborative filtering methods. We argue that the music domain should receive increased attention as it provides more diverse use cases than just the typical user-item recommendation scenario that was at the core of previous tutorials on music recommendation [5].

To this end, we focus on three use cases: automatic playlist generation, context-aware music recommendation, and recommendation in the creative process of music making. We incorporate both academic and industrial points of view into the tutorial and echo some of the research directions highlighted in the “Past, Present and Future” session at RecSys 2016, in particular in [1, 3, 7, 13]. The tutorial is accompanied by a set of slides published online.¹

1 BACKGROUND

Since the development of audio compression techniques and mobile music players, the music industry has been in constant turmoil. The first decade of the 2000’s witnessed an explosion of music content availability. The overwhelming choices implied by having most of the world’s music only a few clicks away did set the stage for the birth of music recommendation technologies. The last tutorial on Music Recommendation at RecSys [5], held in 2011, highlighted the progress at the time towards alleviating this issue.

More recently, we are witnessing a revolution in the very way music is consumed. It is increasingly more common to *access* music than to *purchase* and *own* it in a personal collection. Indeed, in the US, the amount of music listening happening on streaming platforms has dramatically increased in the last few years and is now the top source of revenue for the music industry.

This evolution represents major opportunities and challenges for MRS technologies. Recommending music calls no more for a one-shot recommendation for the purpose of a *product* purchase, but instead for an “infinite recommendation loop”, where what is recommended is a listening *experience* and where the recommendation of a specific item to a user never really ends. Further, recommendation technologies can now impact a multitude of complementary use cases, including all actors in the music production and consumption chain. It is therefore in a unique position to contribute significantly

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).

RecSys '17, Como, Italy

© 2017 Copyright held by the owner/author(s). 978-1-4503-4652-8/17/08...\$15.00
DOI: 10.1145/3109859.3109934

¹<http://www.cp.jku.at/tutorials/mrs-recsys.2017>

to recommender systems research as a whole, as it touches many core topics, e.g., *user modeling*, *user interaction and interfaces*, or *context-aware and mobile recommendations*.

2 IT'S ALL ABOUT THE USE CASE

The current evolution of the music industry results in an explosion of use cases for recommender systems, much beyond the traditional item-to-user use case, cf. [13]. This includes scenarios such as automatic playlist/radio generation, context-aware recommendations, and recommendations in the creative process, which we address in more detail.

2.1 Playlist generation

The task of playlist generation is a core recommendation task in the music domain, with particular relevance for the industry. In order to address this important task, we discuss the question of *what makes a good playlist?* and highlight specificities in terms of requirements and methods. This comprises challenges of *exploration vs. exploitation* and *giving a playlist a focus*, e.g., a genre, (re)discovery, or newly released music, together with a discussion of different *methods* for playlist generation, e.g., curatorial, collaborative filtering, content-based, and/or personalized, and practical aspects of *evaluation* (offline vs. A/B testing).

2.2 Context-aware music recommendation

Regarding the music recommendation task as dependent on the context of the listener, we present examples and discuss *methodology* and *evaluation* strategies of context-aware MRS. This is supported by a review of *context categories*, e.g., time, location, activity, weather, and social context, and elaborated in a discussion on *cultural and regional specificities* in which we summarize findings on country-specific differences in music preferences and requirements for MRS.

2.3 Recommendation in the creative processes of music making

To highlight the role of recommendation in the process of music *making*, we review recent and ongoing projects in this area, with a focus on user-centric requirement identification, theoretic future systems, and user-centric evaluation strategies [2, 7, 10]. This comprises a rather traditional role of recommendation to *facilitate collaborations*, presented through a real-world use case from an online jam community, as well as more emerging scenarios originating from interviews with international music producers, often differing in their requirements from consumer-based systems in terms of a *stronger need for serendipity*, *preservation of originality*, and *systems embodying opposition*, cf. [2, 3].

3 WHAT'S NEXT?

Beside a rich variety of other use cases, such as user-centric MRS (e.g., emotion- and personality-aware), recommending fans to artists, artists to labels, concert venues to performers, or tracks to listener groups, we reflect on social implications of music recommendation growth, such as filter bubble, ethics, and influence on popular culture, and on the challenges ahead from both academic and industry perspectives.

THE BAND

Markus Schedl is an Associate Professor at the Department of Computational Perception of the Johannes Kepler University Linz, Austria.² Recent publications on the topic include [9, 12, 15–17, 19].

Peter Knees is an Assistant Professor of the Institute of Software Technology and Interactive Systems of the TU Wien, Austria.³ Recent publications on the topic include [2, 10–12, 17, 19].

Fabien Gouyon is Principal Scientist at the internet radio Pandora.⁴ Recent and less recent publications on the topic include [4, 6, 8, 14, 18].

ACKNOWLEDGMENTS

Supported by the Austrian Science Fund (FWF): P25655 and the Austrian FFG: BRIDGE 1 project *SmarterJam* (858514).

REFERENCES

- [1] Xavier Amatriain and Justin Basilico. 2016. Past, Present, and Future of Recommender Systems: An Industry Perspective. In *Proc. RecSys*. Boston, MA, USA.
- [2] Kristina Andersen and Peter Knees. 2016. Conversations with Expert Users in Music Retrieval and Research Challenges for Creative MIR. In *Proc. ISMIR*. New York, NY, USA.
- [3] Amos Azaria and Jason Hong. 2016. Recommender Systems with Personality. In *Proc. RecSys*. Boston, MA, USA.
- [4] Pedro Cano, Markus Koppenberger, Nicolas Wack, et al. 2005. An Industrial-strength Content-based Music Recommendation System. In *Proc. SIGIR*. Salvador, Brazil.
- [5] Oscar Celma and Paul Lamere. 2011. Music Recommendation and Discovery Revisited. In *Proc. RecSys*. Chicago, IL, USA.
- [6] Marcos Aurélio Domingues, Fabien Gouyon, Alípio Mário Jorge, et al. 2013. Combining usage and content in an online recommendation system for music in the Long Tail. *International Journal of Multimedia Information Retrieval* 2 (2013).
- [7] Michael D. Ekstrand and Martijn C. Willemsen. 2016. Behaviorism is Not Enough: Better Recommendations Through Listening to Users. In *Proc. RecSys*. Boston, MA, USA.
- [8] Fabien Gouyon, Bob L. Sturm, João Lobato Oliveira, Nuno Hespanhol, and Thibault Langlois. 2014. On Evaluation Validity in Music Autotagging. *CoRR* abs/1410.0001 (2014). <http://arxiv.org/abs/1410.0001>
- [9] Marius Kaminskis, Francesco Ricci, and Markus Schedl. 2013. Location-aware Music Recommendation Using Auto-Tagging and Hybrid Matching. In *Proc. RecSys*. Hong Kong, China.
- [10] Peter Knees and Kristina Andersen. 2017. Building Physical Props for Imagining Future Recommender Systems. In *Proc. HUMANIZE*. Limassol, Cyprus.
- [11] Peter Knees, Kristina Andersen, and Marko Tkalić. 2015. "I'd like it to do the opposite": Music-Making Between Recommendation and Obstruction. In *Proc. DMRS*. Bolzano, Italy.
- [12] Peter Knees and Markus Schedl. 2016. *Music Similarity and Retrieval: An Introduction to Audio- and Web-based Strategies*. Springer.
- [13] Tamas Motajsek, Jean-Yves Le Moine, Martha Larson, et al. 2016. Algorithms Aside: Recommendation As The Lens Of Life. In *Proc. RecSys*. Boston, MA, USA.
- [14] Matthew Prockup, Andreas F. Ehmman, Fabien Gouyon, et al. 2015. Modeling Genre with the Music Genome Project: Comparing Human-Labeled Attributes and Audio Features. In *Proc. ISMIR*. Málaga, Spain.
- [15] Markus Schedl. 2015. Listener-aware Music Recommendation from Sensor and Social Media Data. In *Proc. ECML PKDD*. Porto, Portugal.
- [16] Markus Schedl, Emilia Gómez, and Julián Urbano. 2014. Music Information Retrieval: Recent Developments and Applications. *Foundations and Trends in Information Retrieval* 8, 2–3 (2014), 127–261.
- [17] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskis. 2015. Music Recommender Systems. In *Recommender Systems Handbook* (2nd ed.). Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer.
- [18] Mohamed Sordo, Fabien Gouyon, Luis Sarmento, Óscar Celma, and Xavier Serra. 2013. Inferring Semantic Facets of a Music Folksonomy with Wikipedia. *Journal of New Music Research* 42 (December 2013), 346–363. Issue 4.
- [19] Andreu Vall, Marcin Skowron, Peter Knees, and Markus Schedl. 2015. Improving Music Recommendations with a Weighted Factorization of the Tagging Activity. In *Proc. ISMIR*. Málaga, Spain.

²<http://www.cp.jku.at/people/schedl>

³<http://www.ifs.tuwien.ac.at/~knees>

⁴<http://www.fabiengouyon.org>