

OVERVIEW AND NEW CHALLENGES OF MUSIC RECOMMENDATION RESEARCH IN 2018

Markus Schedl
JKU Linz, Austria
markus.schedl@jku.at

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

Fabien Gouyon
Pandora, USA
fgouyon@pandora.com
(corresponding author)

1. MOTIVATION AND FORMAT

Music recommendation is a central topic at the ISMIR conference. Every edition of the conference comes with a significant number of papers related to that topic, yet the last ISMIR tutorial specialized on that topic took place eleven years ago [1].

Given the dramatic, fast-paced evolution of music recommendation research in the last decade, and the relevance of this topic to the whole MIR community, we believe that a tutorial on the topic would be timely.

We propose to organize an *introductory tutorial* providing an overview of music recommendation research, as well as the challenges it faces today. We therefore do not require particular prerequisite knowledge or skills from the audience, other than the basic understanding of the main MIR concepts that are to be expected from the general ISMIR audience. We will incorporate both *academic and industrial* points of view into the tutorial. Accompanying the tutorial, we will *publish online a comprehensive set of slides*, including references to state-of-the-art work, links to open datasets and open implementations of several of the presented techniques.

2. 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. An ISMIR tutorial in 2007 [1] (and its updated 2011 version at RecSys [2]) provided an overview of the progress at the time towards alleviating this issue.

However, things have evolved since then, as we are now 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, 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 music recommendation technologies, and for the related research communities.

The authors of this proposal recently gave a successful tutorial addressing these opportunities and challenges to the Recommender Systems community [3]. We think a similar tutorial, updated and specifically tailored to ISMIR, would be very relevant and timely for the MIR community.

3. OUTLINE OF THE TUTORIAL

1. Introduction to music recommendation

The tutorial starts with a description of the particular domain of music for recommendation and an overview of data sources and algorithms.

- (a) *Evolution of the music industry and music discovery*: We highlight the rapid development and scaling of the digital music industry using recent figures of music sales
- (b) *What makes music recommendation special?*: We discuss differences between music and other domains, such as movies (e.g., consumption duration, repetition, etc.)
- (c) *Data and algorithms*: We continue by distinguishing music content, purpose and context on the one hand, and listener background, intent and context on the other, as different dimensions of data sources. Which we then connect to algorithms typically applied.

2. It's all about the use case

In this segment, we develop the notion that the current evolution of the music industry results in an explosion of use cases for MIR research. Some of which we are going to describe in detail in the following segments.

3. Use case 1: Playlist generation

This segment describes the use case of playlist generation, highlighting its specificities in terms of requirements and methods.

- (a) *What makes a good playlist?*
- (b) *Exploration vs. exploitation*



- (c) *Giving it a focus*: e.g., a genre, (re)discovery, newly released music, etc.
- (d) *Methods*: curatorial, collaborative filtering, content-based, personalization
- (e) *Evaluation*: offline vs. A/B testing

4. Use case 2: Context-aware music recommendation

Regarding the music recommendation task as dependent on the context of the listener, in this segment, we review context categories used in music recommendation and present examples of context-aware MRS.

- (a) *Context categories*: We categorize various dimensions of the user context, e.g., time, location, activity, weather, social context, etc.
- (b) *Cultural/regional specificities*: We summarize findings about country-specific differences in music preferences and in requirements for MRS.
- (c) *Methods*: We outline the most frequently adopted approaches in context-aware MRS.
- (d) *Evaluation*: We highlight particular challenges in evaluating context-aware MRS.

5. Use case 3: Recommendation in the creative processes of music making

In this segment, we discuss the particular aspects of recommendation in music making, based on experiences/publications from recent and ongoing projects in this area, with a focus on user-centric requirement identification and theoretic future systems.

- (a) *Recommendation to facilitate collaboration*: We present a real-world use case from an online jam community, where an RS should bring together compatible and complementary musicians and recorded tracks that match an ongoing session, representing a rather traditional RS scenario.
- (b) *Asking expert users*: We present topics emerging from interviews with international music producers on the role of recommendation in their work. These represent less traditional and defined, hence more challenging scenarios.
- (c) *Implications and theoretical systems*: We discuss requirements that differ from those of consumer-based systems, i.e., stronger need for serendipity, preservation of originality, systems embodying opposition.
- (d) *Evaluation*: Discussing the difficulties of quantitative evaluation, we argue for qualitative and user-centric evaluation strategies.

6. What's next?

In this last segment, we briefly mention other use cases, with a focus on innovative ones. We also reflect on social implications of music recommendation growth and on challenges ahead.

- (a) *Innovative use cases*: e.g. user-centric MRS (e.g., emotion- and personality-aware)
- (b) *With the fun comes responsibility*: filter bubble, ethics, influence on popular culture, etc.
- (c) *Challenges* from academic and industry perspectives

4. PRESENTERS

Dr. Markus Schedl is an Associate Professor at the Department of Computational Perception of the Johannes Kepler University Linz, Austria.¹ His main research interests include recommender systems, user modeling, web and social media mining, multimedia, and music information retrieval. He (co-)authored more than 150 refereed journal articles and conference papers on these topics. Markus already gave several tutorials at the most important multimedia, information retrieval, and recommender systems conferences, including ISMIR, ACM Multimedia, ACM SIGIR, ACM ICMR, and ACM Recommender Systems.

Dr. Peter Knees is an Assistant Professor of the Faculty of Informatics, TU Wien, Austria. For over a decade he has been an active member of the ISMIR community, reaching out to the related areas of multimedia, text IR, and recommender systems. Apart from serving on the program committees of major conferences in the field, he has organized several workshops on topics of media retrieval. He is an experienced teacher of graduate-level courses on recommender systems and information retrieval and has given tutorials on music information retrieval at RecSys, SIGIR, ECIR, RuSSIR, and the Indonesian Summer School on MIR.²

Dr. Fabien Gouyon is Principal Scientist at the music streaming service Pandora, where he does applied research on personalized music recommendation and works with the Music Genome Project. Before joining Pandora, he received a PhD in Computer Science while working in the Music Technology Group in the University Pompeu Fabra in Barcelona, and was a co-founder of Barcelona Music and Audio Technologies (BMAT), worked in the Austrian Research Institute for Artificial Intelligence in Vienna, and started and led the Sound and Music Computing Group while teaching at the University of Porto. He was President of ISMIR in 2016-2017.³

5. SPECIAL REQUIREMENTS

For the presentation of the tutorial, we will need a projector and an audio system (compatible with a typical notebook's 3.5mm stereo jack), capable of playing music sufficiently loud to be perceivable by the entire audience (to allow playback of videos and sound material for demonstration of the discussed approaches).

¹ <http://www.cp.jku.at/people/schedl>

² <https://www.ifs.tuwien.ac.at/~knees>

³ <http://www.fabiengouyon.org>

6. REFERENCES

- [1] O. Celma and P. Lamere. Music recommendation tutorial. In *Proc. ISMIR*, Vienna, Austria, 2007.
- [2] O. Celma and P. Lamere. Music recommendation and discovery revisited. In *Proc. RecSys*, Chicago, USA, 2011.
- [3] M. Schedl, P. Knees, and F. Gouyon. New paths in music recommendation systems research. In *Proc. RecSys*, Como, Italy, 2017.