The Potential of the Confluence of Theoretical and Algorithmic Modeling in Music Recommendation

Christine Bauer
Johannes Kepler University Linz
Linz, Austria
christine.bauer@jku.at

ABSTRACT

The task of a music recommender system is to predict what music item a particular user would like to listen to next. This position paper discusses the main challenges of the music preference prediction task: the lack of information on the many contextual factors influencing a user's music preferences in existing open datasets, the lack of clarity of what the right choice of music is and whether a right choice exists at all; the multitude of criteria (beyond accuracy) that have to be met for a "good" music item recommendation; and the need for explanations on relationships to identify (and potentially counteract) unwanted biases in recommendation approaches.

The paper substantiates the position that the confluence of theoretical modeling (which seeks to explain behaviors) and algorithmic modeling (which seeks to predict behaviors) seems to be an effective avenue to take in computational modeling for music recommender systems.

KEYWORDS

computational modeling; music recommendation; preference prediction; human-computer interaction

CHI'19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
https://doi.org/10.1145/nnnnnnn.nnnnnnn

Examples of differences associated with music items, users, and their consumption behaviour include the following[34]:

- very low consumption time in the dimension of minutes, whereas a book or a travel are consumed during days or weeks;
- consumption in sequences (e.g., playlists);
- music often consumed passively (e.g., while jogging, travelling, working);
- consumption is highly driven by situational context;
- users are likely to appreciate the re-recommendation of the same item while a user is less likely to read the same news article over and over again; and
- music evokes strong emotions.

Sidebar 1: The specialties of the music domain

INTRODUCTION

Before the era of the Internet, access to music content (e.g., music recordings) was restricted to local availability of their physical representations (e.g., vinyl). Thereby, the selection and aggregation of content had traditionally been exposed to human control [25]. For instance, a small group of Artist&Repertoire managers working for the major music labels scouted new artists and developed them commercially.

Nowadays,—owing to the development of the Social Web that allows for easy distribution of user-generated content—the intermediary level of experts (e.g., the Artist&Repertoire managers at music labels) that traditionally "prefiltered" content before it reached potential consumers is bypassed. This results in the situation that users currently face: music content is abundantly available online and the amount of overall available content increases tremendously on a daily basis.

However, the opportunity to access a large amount of content frequently leads to information overload [8] or choice overload [19], because people do not find the content that they are interested in or do not know what to choose. Assisting users in searching, sorting, and filtering the massive amount of online content [24], recommender systems have become important tools in people's everyday life and do not only facilitate the interaction with music content [15], but also support versatile activities such as shopping [26], consuming news [25], or finding persons for any kind of social matching [23].

Recommender systems are computer systems that provide suggestions for items that are deemed interesting to a particular target user, assisting that particular user in various decision-making processes (e.g., relating to what music to listen to) [27]. The general term used to denote to what the system recommends to users is "item" [27]; in case of music recommender systems (MRS) it is the music item (e.g., musical work, artist, genre).

There are universally valid principles for designing recommender systems, such as that a recommender system typically consist of three key components (i.e., user, item, and matching mechanism) [3]. Still, a recommender system needs to be put into context because there are product- and sector-specific characteristics that a recommender system needs to consider (be customized to) to provide useful and effective recommendations for the specific type of item [27, 33]. Sidebar 1 presents the specialties of the music domain compared to other domains deploying recommender systems.

RATIONALE

An ideal MRS proposes "the right music, to the right user, at the right moment" [21]. However, this is a complex task because various factors influence a user's music preferences in a given situation [6]. Many studies have investigated the relationships between music preferences and various person-related characteristics (e.g., demographics [17], personality traits [28], social influences [12]. Besides person-related characteristics, also situation-related factors (e.g., temporal aspects [18], or weather [13]) influence a user's music preferences. The task of an MRS is to predict what a particular user would like to listen to next. Basically, there are two computational modeling approaches to build upon for this music preference prediction task:

- *Theoretical modeling* seeks to explain users' listening behavior. For advancing MRS, the first step would be to observe a user's listening behavior and perform analyses to explain where a user's listening behavior results from (e.g., from person-related characteristics or situational factors, and from which of these in particular). Then, building on these findings (e.g., knowing that Finnish listeners are more likely to prefer heavy metal than Italian listeners [30]), future user models may be created for predictions.
- *Algorithmic modeling* seeks to predict users' listening behavior. Algorithmic modeling may rely on approaches that are capable of identifying listening patterns within a user's listening history or across users without necessarily delivering descriptions that help *explaining* the relationships of the identified patterns. For instance, approaches such as deep neural networks frequently leave us with "black boxes" [20] because the resulting models are complex and frequently they do not produce an intelligible description of the results produced in each case. Still, the resulting models may be apt to deliver remarkably accurate predictions. In other words, algorithmic modeling may recommend music to the user what he or she will indeed like in the very moment without understanding whether it was indeed the "right" choice—and if—why it was "right".

CHALLENGES

One challenge for music preference prediction is that it is (almost) impossible to say what is *the* right choice for a particular user in the particular moment; it is typically a set of items that is *right* or *okay*.

Another challenge of algorithmic modeling is that—currently—we can only model based on data that we have available. For MRS, several open datasets exist, such as the Million Song Dataset [10], the LFM-1b dataset [29], or the recently released Music Streaming Sessions Dataset [14]. However, there are many factors influencing a user's music preferences for which we do not have (sufficient) data available (yet) to exploit for algorithmic modeling. Theoretical modeling—thus, the "explaining approach"—may help here to advance MRS. It is also a viable basis to provide an informed route what kind of data should be collected so that algorithmic modeling may come into play here to use its powerful mechanisms to exploit the additional data to make even better predictions.

A further challenge relates to evaluation of MRS: What does it mean if an MRS recommends a music item to a user and the user indeed listens to the item? Potentially, it is the user's most favorite song and so the user enjoyed listening to it. Maybe, though, the user listens to the item because the algorithm provided it as the next one to listen to in the playlist, but the user was distracted at the very moment because of receiving a phone call (or was not present in the room for some minutes). In such cases, the recommendation was maybe not a "bad" one because the user did not hear it anyways, but was it a good prediction then?

With respect to biases as inherent in recommendation systems (e.g., the popularity bias phenomenon [16] suggesting that over time the most popular music items tend to get more and more attention, while music items in the long tail get less and less attention [22]), the ability to understand and explain models seems to be a crucial prerequisite to uncover such bias and develop and take effective measures to counteract unwanted bias.

PREVIOUS AND ONGOING RESEARCH, AND INTERESTS

A major part of my previous and ongoing research is aimed at integrating contextual information into (user) modeling. Basically, my work on context modeling takes a conceptual viewpoint (e.g., [2, 4]). It points towards the various potentially relevant contextual factors that we tend to "forget" in modeling (for various reasons such as, for instance, the non-availability of useful datasets including such contextual information).

With the main objective at improving MRS, some part of my research on MRS is geared towards identifying relationships between various aspects (such as age [32], user connections [5, 7], user country [6], real-world events [35], mainstreaminess [31]) and music preferences or listening behavior. Findings are then used to improve MRS performance (for instance in [6, 31, 32]).

To a considerable extent, ideas on the (contextual) components that could improve MRS are based on literature from various disciplines such as cognitive science (e.g., [36]), social psychology (e.g., [11]), and computer science [21]. In addition, ideas emanate from my own experience of many years in the music domain—which is a significant knowledge source that is not available to every researcher.

PROSPECTS

Overall, recommender systems research has predominantly focused on improving the prediction accuracy of algorithms based on existing datasets (reflecting users' historic item ratings or consumption behavior) [9]. However, to date, comprehensive contextual information about users and the specific situational settings in which those consume the items is rarely available in existing datasets [1]—and is especially true for music-related datasets.

The confluence of theoretical and algorithmic modeling seems to be an effective avenue to take in computational modeling for MRS.

ACKNOWLEDGMENTS

This research is supported by the Austrian Science Fund V579.

REFERENCES

[1] Gediminas Adomavicius and Alexander Tuzhilin. 2015. Context-aware recommender systems. In *Recommender Systems Handbook* (2nd ed.), Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer, New York, NY, 191–226. https://doi.org/10.1007/978-1-4899-7637-6_6

[2] Christine Bauer. 2014. A framework for conceptualizing context for intelligent systems (CCFIS). *Journal of Ambient Intelligence and Smart Environments* 6, 4 (2014), 403–417. https://doi.org/10.3233/AIS-140269

[3] Christine Bauer, Marta Kholodylo, and Christine Strauss. 2017. Music Recommender Systems: Challenges and Opportunities for Non-Superstar Artists. In *30th Bled eConference*. 21–32. https://doi.org/10.18690/978-961-286-043-1.3

[4] Christine Bauer and Alexander Novotny. 2017. A consolidated view of context for intelligent systems. *Journal of Ambient Intelligence and Smart Environments* 9, 4 (2017), 377–393. https://doi.org/10.3233/ais-170445

[5] Christine Bauer and Markus Schedl. 2018. Investigating cross-country relationship between users' social ties and music mainstreaminess. In *19th International Society for Music Information Retrieval Conference (ISMIR'18)*. ISMIR, 678–686.

[6] Christine Bauer and Markus Schedl. 2018. On the Importance of Considering Country-specific Aspects on the Online-Market: An Example of Music Recommendation Considering Country-Specific Mainstream. In *51st Hawaii International Conference on System Sciences (HICSS'18)*. 3647–3656. http://hdl.handle.net/10125/50349

[7] Christine Bauer and Markus Schedl. 2019. A cross-country investigation of user connection patterns in online social networks. In *52nd Hawaii International Conference on System Sciences (HICSS'19)*. 2166–2175. http://hdl.handle.net/10125/59655

[8] David Bawden and Lyn Robinson. 2009. The dark side of information: overload, anxiety and other paradoxes and pathologies. *Journal of Information Science* 35, 2 (2009), 180–191. https://doi.org/10.1177/0165551508095781

[9] Joeran Beel, Marcel Genzmehr, Stefan Langer, Andreas Nürnberger, and Bela Gipp. 2013. A Comparative Analysis of Offline and Online Evaluations and Discussion of Research Paper Recommender System Evaluation. In *International Workshop on Reproducibility and Replication in Recommender Systems Evaluation (RepSys'13)*. ACM, 7–14. https://doi.org/10.1145/2532508.2532511

[10] Thierry Bertin-Mahieux, Daniel P.W. Ellis, Brian Whitman, and Paul Lamere. 2011. The Million Song Dataset. In *12th International Conference on Music Information Retrieval (ISMIR'11)*. ISMIR.

[11] Arielle Bonneville-Roussy, Peter J. Rentfrow, Man K. Xu, and Jeff Potter. 2013. Music through the ages: Trends in musical engagement and preferences from adolescence through middle adulthood. *Journal of Personality and Social Psychology* 105, 4 (2013), 703–717. https://doi.org/10.1037/a0033770

[12] Arielle Bonneville-Roussy and John Rust. 2018. Age trends in musical preferences in adulthood: 2. Sources of social influences as determinants of preferences. *Musicae Scientiae* 22, 2 (2018), 175–195. https://doi.org/10.1177/1029864917704016

[13] Matthias Braunhofer, Marius Kaminskas, and Francesco Ricci. 2011. Recommending music for places of interest in a mobile travel guide. In *5th ACM Conference on Recommender Systems (RecSys'11)*. ACM, 253–256.

[14] Brian Brost, Rishabh Mehrotra, and Tristan Jehan. 2019. The Music Streaming Sessions Dataset. In *The Web Conference 2019*. ACM.

[15] Òscar Celma. 2010. *Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space*. Springer, Berlin, Heidelberg, Germany.

[16] Òscar Celma and Pedro Cano. 2008. From Hits to Niches?: Or How Popular Artists Can Bias Music Recommendation and Discovery. In *2Nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition (NETFLIX'08)*. ACM, New York, NY, Article 5, 8 pages. https://doi.org/10.1145/1722149.1722154

[17] Zhiyong Cheng, Jialie Shen, Liqiang Nie, Tat-Seng Chua, and Mohan Kankanhalli. 2017. Exploring User-Specific Information in Music Retrieval. In *40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'17)*. ACM, New York, NY, 655–664. https://doi.org/10.1145/3077136.3080772

[18] Ricardo Dias and Manuel J Fonseca. 2013. Improving music recommendation in session-based collaborative filtering by using temporal context. In *IEEE 25th International Conference on Tools with Artificial Intelligence (ICTAI'13)*. IEEE, 783–788. https://doi.org/10.1109/ICTAI.2013.120

[19] Gerald Häubl and Valerie Trifts. 2000. Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science* 19, 1 (2000), 4–21.

[20] Pang Wei Koh and Percy Liang. 2017. Understanding Black-box Predictions via Influence Functions. In *34th International Conference on Machine Learning (ICML'17)*. 1885–1894.

[21] Audrey Laplante. 2014. Improving music recommender systems: what can we learn from research on music tags?. In *15th International Society for Music Information Retrieval Conference (ISMIR'14)*. ISMIR, 451–456.

[22] Mark Levy and Klaas Bosteels. 2010. Music recommendation and the long tail. In *1st Workshop On Music Recommendation And Discovery (WOMRAD'10)*.

[23] Julia M. Mayer, Quentin Jones, and Starr Roxanne Hiltz. 2015. Identifying Opportunities for Valuable Encounters: Toward Context-Aware Social Matching Systems. *ACM Transactions on Information Systems* 34, 1, Article 1 (July 2015), 32 pages. https://doi.org/10.1145/2751557

[24] Miquel Montaner, Beatriz López, and Josep Lluís de la Rosa. 2003. A taxonomy of recommender agents on the Internet. *Artificial Intelligence Review* 19, 4 (2003), 285–330. https://doi.org/10.1023/A:1022850703159

[25] Oliver Oechslein and Thomas Hess. 2014. The Value of a Recommendation: The Role of Social Ties in Social Recommender Systems. In *47th Hawaii International Conference on System Science (HICSS'14)*. IEEE, 1864–1873. https://doi.org/10.1109/HICSS.2014.235

[26] Gal Oestreicher-Singer and Arun Sundararajan. 2012. Recommendation Networks and the Long Tail of Electronic Commerce. *MIS Quarterly* 36, 1 (2012), 65–83.

[27] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. *Recommender Systems Handbook* (2nd ed.). Springer, New York, NY. https://doi.org/10.1007/978-1-4899-7637-6

[28] Thomas Schäfer and Claudia Mehlhorn. 2017. Can personality traits predict musical style preferences? A meta-analysis. *Personality and Individual Differences* 116 (2017), 265–273.

[29] Markus Schedl. 2016. The LFM-1b Dataset for Music Retrieval and Recommendation. In *ACM on International Conference on Multimedia Retrieval (ICMR'16)*. ACM, New York, NY, 103–110. https://doi.org/10.1145/2911996.2912004

[30] Markus Schedl. 2017. Investigating country-specific music preferences and music recommendation algorithms with the LFM-1b dataset. *International Journal of Multimedia Information Retrieval* 6, 1 (2017), 71–84. https://doi.org/10.1007/s13735-017-0118-y

[31] Markus Schedl and Christine Bauer. 2017. Introducing Global and Regional Mainstreaminess for Improving Personalized Music Recommendation. In *15th International Conference on Advances in Mobile Computing & Multimedia (MoMM'17)*. ACM, New York, NY, 74–81. https://doi.org/10.1145/3151848.3151849

[32] Markus Schedl and Christine Bauer. 2017. Online Music Listening Culture of Kids and Adolescents: Listening Analysis and Music Recommendation Tailored to the Young. In *1st International Workshop on Children and Recommender Systems (KidRec'17)*. ACM, New York, NY.

[33] 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.

[34] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskas. 2015. Music Recommender Systems. In *Recommender Systems Handbook* (2nd ed.), Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer, New York, NY, 453–492. https://doi.org/10.1007/978-1-4899-7637-6_13

[35] Markus Schedl, Eelco Wiechert, and Christine Bauer. [n. d.]. The effects of real-world events on music listening behavior: an intervention time series analysis. In *WWW '18 Companion: The 2018 Web Conference Companion (WWW'18)*. ACM, 75–76. https://doi.org/10.1145/3184558.3186936

[36] Catherine J. Stevens. 2012. Music perception and cognition: a review of recent cross-cultural research. *Topics in Cognitive Science* 4, 4 (2012), 653–667. https://doi.org/10.1111/j.1756-8765.2012.01215.x