



Machine Learning Approaches to Hybrid Music Recommender Systems

Andreu Vall¹(✉) and Gerhard Widmer^{1,2}

¹ Institute of Computational Perception, Johannes Kepler University, Linz, Austria
{andreu.vall, gerhard.widmer}@jku.at

² Austrian Research Institute for Artificial Intelligence, Vienna, Austria

Abstract. Music recommender systems have become a key technology supporting the access to increasingly larger music catalogs in on-line music streaming services, on-line music shops, and private collections. The interaction of users with large music catalogs is a complex phenomenon researched from different disciplines. We survey our works investigating the machine learning and data mining aspects of hybrid music recommender systems (i.e., systems that integrate different recommendation techniques). We proposed hybrid music recommender systems robust to the so-called “cold-start problem” for new music items, favoring the discovery of relevant but non-popular music. We thoroughly studied the specific task of music playlist continuation, by analyzing fundamental playlist characteristics, song feature representations, and the relationship between playlists and the songs therein.

Keywords: Music recommender systems
Music playlist continuation · Hybrid recommender systems
Cold-start problem

1 Introduction

Music recommender systems support the interaction of users with large music catalogs. They strongly rely on machine learning and data mining methods to analyze the data describing the users, the music items, the interaction of users with music items, and even the users’ context when interacting with the music items [8, Chap. 7]. Even though the techniques at the core of recommender systems are valid for different item domains (e.g., recommending movies, news, or jobs), the inherent characteristics of the music domain must be considered:

- Music can be recommended at different granularity levels (e.g., songs, albums, artists, or even genres or ready-made playlists).
- Music is often consumed in listening sessions, such as albums or playlists. This defines local session contexts that need to be taken into consideration.
- Music recommender systems should adapt their level of interference to the user needs, ranging from simply supporting the exploration of the music catalog, to providing a full *lean-back* experience.

Music recommender systems have received contributions from different but converging approaches. Research specializing in music information retrieval has often focused on content-based music recommender systems. Music items are represented by features derived from the audio signal, social tags, or web content and recommendations are predicted on the basis of content-wise similarities [5, 7]. On the other hand, research specializing in general recommender systems has usually focused on collaborative filtering techniques, developing statistical models to extract underlying music taste patterns from usage data (e.g., listening logs, or radio stations) [1, 3].

Content-based recommender systems provide fair but limited performance, because the relations derived from content-wise similarities tend to be simple. Collaborative recommender systems capture more abstract music taste patterns, but their performance is heavily affected by the availability of data (users and music items for which few observations are available are poorly represented).

In this paper we survey our recent works on hybrid music recommender systems integrating the strengths of content-based and collaborative recommender systems. We approached two main music recommendation tasks: (1) music artist recommendation, modeling the users' general music preferences over an extended period of time, and (2) music playlist continuation, focusing on next-song recommendations for short listening sessions.

2 Music Artist Recommendation

We proposed in [15] a hybrid extension to a well-established matrix factorization CF approach for implicit feedback datasets [6]. The proposed hybrid CF extension jointly factorizes listening logs and social tags from music streaming services (abundant sources of usage data in the music domain). According to our numerical experiments, the proposed hybrid CF extension yielded more accurate music artist recommendations. Furthermore, we extended the standard evaluation methodology incorporating bootstrap confidence intervals [4] to facilitate the comparison between systems. In the follow-up work [9], we observed that the superior performance of the proposed hybrid CF extension was explained by its robustness to the cold-start problem for new artists (i.e., its ability to better represent music artists for which few observations were available).

3 Music Playlist Continuation

Like previous works on music playlist modeling, we based our research on the exploitation of hand-curated music playlists, which we regard as rich examples from which to learn music compilation patterns.

3.1 Playlist Characteristics and Next-Song Recommendations

Before moving into the design of hybrid music recommender systems for music playlist continuation, we studied which basic playlist characteristics should

be considered to effectively predict next-song recommendations. We studied in [13,14] the importance of the song order, the song context, and the song popularity for next-song recommendations. We compared three existing playlist continuation models of increasing complexity on two datasets of hand-curated music playlists. We observed that considering a longer song context has a positive impact on next-song recommendations. We found that the long-tailed nature of the playlist datasets (common in music collections [2]) makes simple models and highly-expressive models appear to perform comparably. However, further analysis revealed the advantage of using highly-expressive models. Our experiments also suggested either that the song order is not crucial for next-song recommendations, or that even highly-expressive models are unable to exploit it. We also proposed an evaluation approach for next-song recommendations that mitigates known issues in the evaluation of music recommender systems.

3.2 Song Representations and Playlist-Song Membership

Given the results of the experiments described in [13,14], we proposed music recommender systems for playlist continuation able to consider the full playlist song context. We also required the proposed systems to hybridize collaborative and content-based recommender systems to ensure robustness to the cold-start problem.

We identified in [11,12] suitable song-level feature representations for music playlist modeling. We investigated features derived from the audio signal, social tags, and independent listening logs. We found that the features derived from independent listening logs are more expressive than those derived from social tags, which in turn outperform those derived from the audio signal. The combination of features from different modalities outperformed the individual features, suggesting that the different modalities indeed carry complementary information.

We further proposed in [12] a hybrid music recommender system for music playlist continuation robust to the cold-start problem for non-popular songs. However, this approach can only extend playlists for which a profile has been pre-computed at training time. In the follow-up work [10], we proposed another hybrid music recommender system for music playlist continuation that regards playlist-song pairs exclusively in terms of feature vectors. This system learns general “playlist-song” membership relationships, which not only make it robust to the cold-start problem for non-popular songs but also enable the extension of playlists not seen at training time.

4 Lessons Learned and Open Challenges

Additional insights and questions arise from the research conducted in the presented works. Importantly, the evaluation of music recommender systems by means of numerical experiments does not reflect the fact that the usefulness of music recommendations is a highly subjective judgment of the end user. Further research on evaluation metrics for music recommender systems is required.

On a related note, the users' subjectivity on the usefulness of music recommendations makes it challenging to anticipate the actual immediate user needs. An interesting research direction to bridge this gap focuses on the development of interactive interfaces that should let users express their current needs.

Acknowledgments. This research has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 670035 (Con Espressione).

References

1. Aizenberg, N., Koren, Y., Somekh, O.: Build your own music recommender by modeling internet radio streams. In: Proceedings of WWW, pp. 1–10 (2012)
2. Celma, Ó.: Music Recommendation and Discovery. Springer, Heidelberg (2010). <https://doi.org/10.1007/978-3-642-13287-2>
3. Chen, S., Moore, J.L., Turnbull, D., Joachims, T.: Playlist prediction via metric embedding. In: Proceedings of SIGKDD, pp. 714–722 (2012)
4. DiCiccio, T.J., Efron, B.: Bootstrap confidence intervals. *Stat. Sci.* **11**, 189–212 (1996)
5. Flexer, A., Schnitzer, D., Gasser, M., Widmer, G.: Playlist generation using start and end songs. In: Proceedings of ISMIR, pp. 173–178 (2008)
6. Hu, Y., Koren, Y., Volinsky, C.: Collaborative filtering for implicit feedback datasets. In: Proceedings of ICDM, pp. 263–272 (2008)
7. Knees, P., Pohle, T., Schedl, M., Widmer, G.: Combining audio-based similarity with web-based data to accelerate automatic music playlist generation. In: Proceedings of the International Workshop on Multimedia Information Retrieval, pp. 147–154 (2006)
8. Ricci, F., Rokach, L., Shapira, B.: Recommender Systems Handbook, 2nd edn. Springer, Boston (2015). <https://doi.org/10.1007/978-1-4899-7637-6>
9. Vall, A.: Listener-inspired automated music playlist generation. In: Proceedings of RecSys, Vienna, Austria (2015)
10. Vall, A., Dorfer, M., Schedl, M., Widmer, G.: A hybrid approach to music playlist continuation based on playlist-song membership. In: Proceedings of SAC, Pau, France (2018)
11. Vall, A., Eghbal-zadeh, H., Dorfer, M., Schedl, M.: Timbral and semantic features for music playlists. In: Machine Learning for Music Discovery Workshop at ICML, New York, NY, USA (2016)
12. Vall, A., Eghbal-zadeh, H., Dorfer, M., Schedl, M., Widmer, G.: Music playlist continuation by learning from hand-curated examples and song features: Alleviating the cold-start problem for rare and out-of-set songs. In: Proceedings of the Workshop on Deep Learning for Recommender Systems at RecSys, Como, Italy (2017)
13. Vall, A., Quadrana, M., Schedl, M., Widmer, G.: The importance of song context and song order in automated music playlist generation. In: Proceedings of ICMPC-ESCOM, Graz, Austria (2018)
14. Vall, A., Schedl, M., Widmer, G., Quadrana, M., Cremonesi, P.: The importance of song context in music playlists. In: RecSys Poster Proceedings, Como, Italy (2017)
15. Vall, A., Skowron, M., Knees, P., Schedl, M.: Improving music recommendations with a weighted factorization of the tagging activity. In: Proceedings of ISMIR, Málaga, Spain (2015)