

# Sequences of Diverse Song Recommendations

Nava Tintarev  
Delft University of Technology  
Delft, The Netherlands

Christoph Lofi  
Delft University of Technology  
Delft, The Netherlands

Cynthia C. S. Liem  
Delft University of Technology  
Delft, The Netherlands

## ACM Reference Format:

Nava Tintarev, Christoph Lofi, and Cynthia C. S. Liem. 2017. Sequences of Diverse Song Recommendations. In *Dutch Information Retrieval*. ACM, New York, NY, USA, 1 page. [https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)

## 1 RELEVANCE TO DIR

Recommender systems suggest users items to consume, try, or buy, by learning from our past interactions, inferring our interests, and making predictions. A criticism of recommender systems has been that they “over-personalize”: censoring user choices over time and effectively polarising users’ preferences [2, 7]. One approach that addresses this relative uniformity of recommendation is diversification of items within a list. However, not only algorithmic diversity, but also user perceptions of diversification are expected to play a role: Wilhelmssen et al. found that people were able to detect different degrees of diversification [10], while Ferwerda et al. found that diversification only improved recommendation attractiveness if it also lead to a perceived sense of discovery [4].

In music, there has been progress on automatically generating playlists [1, 3, 5, 9], but to the best of our knowledge this work did not also consider ordering and diversity, or whether the improvements in accuracy were beneficial from a user perspective. This work is part of a larger agenda of understanding user requirements for constructing *sequences of items*. This enables the community work on diversification algorithms that consider user perceptions of diversity.

## 2 RESEARCH SUMMARY

The paper in UMAP’17 [8] raised the question of what happens in domains, like music, where recommended items are usually consumed in sequence. A sequence not only affords a recommender system more chances to make accurate recommendations, but also to mix familiar and unfamiliar items. However, this situation also creates new challenges for recommendation quality and user satisfaction (i.e., two given items are good recommendations when considered in isolation, but create a poor experience when consumed in sequence). In contrast, recommendation lists are commonly seen as top-N recommendations where the user selects one or only a handful of items from a list of N items. The paper presented an exploratory study with users of a live recommender systems which presents playlists, focusing on the perspective of users, addressing:

- How users perceive diversity: Which items should be kept familiar, which should be diversified?
- How users perceive the role of ordering: What expectations do they have?

In this study three-quarters of participants felt that the list needed to be diverse, and almost all felt that the songs had to be novel to them. However, they also preferred specific types of familiarity: two-thirds expected to see familiar genres, 20% familiar artists, and 5% familiar titles. One surprising result was that *participants did not expect the recommended list to be in a particular order, nor did they find this property to be important*. Users do not seem to be consciously seeking a sequential narrative (i.e., they might be using the system for background listening), even if songs can (and are likely to) be consumed in a sequence.

Consequently, this pilot study enables us to formulate new hypotheses about user perceptions of ordering in diverse sequences. Firstly, when users consider ordering they may be looking for something in particular, like an increase in tempo across songs, or thematic progression (found to be important in user generated playlists [6]). It may also be more important to identify constraints on ordering than to seek an optimal ordering. There are orderings that would be unsuitable for most listeners, like playing songs in very different genres, or highly different tempos in direct succession. Our current work studies which types of ordering choices are perceptible to users, and influence satisfaction. It investigates the role of varying the distance of songs, and using intermediate songs to create a “smooth transition”. This work involves a service-neutral environment, and studies active listening behavior.

## REFERENCES

- [1] Natalie Aizenberg, Yehuda Koren, and Oren Somekh. 2012. Build Your Own Music Recommender by Modeling Internet Radio Streams. In *WWW*. 1–10.
- [2] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348 (2015), 1130–1132.
- [3] Shuo Chen, Joshua Moore, Thorsten Joachims, and Douglas Turnbull. 2012. Playlist Prediction via Metric Embedding. In *SIGKDD*. 714–722.
- [4] Bruce Ferwerda, Mark P Gaus, Andreu Vall, Marko Tkalcić, and Markus Schedl. 2017. How item discovery enabled by diversity leads to increased recommendation list attractiveness. In *Proceedings of the Symposium on Applied Computing*. ACM, 1693–1696.
- [5] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In *International Conference on Learning Representations (ICLR)*.
- [6] Brian McFee and Gert RG Lanckriet. 2011. The Natural Language of Playlists. In *International Conference on Music Information Retrieval (ISMIR)*. 537–542.
- [7] Dimitar Nikolov, Diego FM Oliveira, Alessandro Flammini, and Filippo Menczer. 2015. Measuring online social bubbles. *PeerJ Computer Science* 1 (2015), e38.
- [8] Nava Tintarev, Christoph Lofi, and Cynthia Liem. 2017. Sequences of Diverse Song Recommendations: An exploratory study in a commercial system. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*. ACM, 391–392.
- [9] Roberto Turrin, Andrea Condorelli, Paolo Cremonesi, Roberto Pagano, and Massimo Quadrana. 2015. Large scale music recommendation. In *Workshop on Large-Scale Recommender Systems (LSRS 2015) at ACM RecSys*.
- [10] Martijn C Willemsen, Bart P Knijnenburg, Mark P Gaus, LC Velter-Bremmers, and Kai Fu. 2011. Using latent features diversification to reduce choice difficulty in recommendation lists. *RecSys* 11 (2011), 14–20.

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).

DIR’17, November 2017, Hilversum

© 2017 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06.

[https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)