

# Exploring Longitudinal Effects of Session-based Recommendations

ANDRES FERRARO, Music Technology Group - Universitat Pompeu Fabra, Spain

DIETMAR JANNACH, AAU Klagenfurt, Austria

XAVIER SERRA, Music Technology Group - Universitat Pompeu Fabra, Spain

Session-based recommendation is a problem setting where the task of a recommender system is to make suitable item suggestions based only on a few observed user interactions in an ongoing session. The lack of long-term preference information about individual users in such settings usually results in a limited level of personalization, where a small set of popular items may be recommended to many users. This repeated exposure of such a subset of the items through the recommendations may in turn lead to a reinforcement effect over time, and to a system which is not able to help users discover new content anymore to the desirable extent.

In this work, we investigate such potential longitudinal effects of session-based recommendations in a simulation-based approach. Specifically, we analyze to what extent algorithms of different types may lead to concentration effects over time. Our experiments in the music domain reveal that all investigated algorithms—both neural and heuristic ones—may lead to lower item coverage and to a higher concentration on a subset of the items. Additional simulation experiments however also indicate that relatively simple re-ranking strategies, e.g., by avoiding too many repeated recommendations in the music domain, may help to deal with this problem.

CCS Concepts: • **Information systems** → **Recommender systems**.

Additional Key Words and Phrases: Session-based Recommendation, Longitudinal Effects, Bias

## ACM Reference Format:

Andres Ferraro, Dietmar Jannach, and Xavier Serra. 2020. Exploring Longitudinal Effects of Session-based Recommendations. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20), September 22–26, 2020, Virtual Event, Brazil*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3383313.3412213>

## 1 INTRODUCTION

Recommender systems (RS) more and more determine which content we see in the online world. The exposure of selected items through recommendations aims at achieving certain desired effects like the discovery of new content. However, depending on the underlying algorithm, the repeated exposure of some items may also lead to undesired effects, like a too strong emphasis on already popular items. Such “blockbuster effects”, which may result in decreased sales diversity in e-commerce, were previously investigated for traditional recommendation approaches based on long-term preference profiles [7, 11], and they were also observed in field tests [13]. Understanding such undesired effects is, however, important in many domains. On news websites or on social media sites, for example, an over-emphasis on certain types of content may lead to a biased distribution of information, and to effects like filter bubbles [17].

The problem of such biases might be particularly pronounced in the context of *session-based* recommendation scenarios, where we have to deal with anonymous or first-time users [9, 18]. In such situations, the recommendations by the system can be based only on the few observed interactions in the ongoing session. The level of personalization

---

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).

© 2020 Copyright held by the owner/author(s).

Manuscript submitted to ACM

may therefore be low and lead to the effect that many of the provided recommendations consist mainly of generally popular items, eventually resulting in a reinforcement effect.

This work is a first step to investigate such effects for the problem of session-based recommendation in a simulation-based approach. Starting from real-world data of recorded user interactions from the music domain, we first generate session-based recommendations from a selected set of seed tracks with different algorithms. We then assume that a certain fraction of the recommendations are listened to by the users, and we correspondingly extend the dataset with these new interactions. This process is repeated in the simulation and from time to time we re-train the underlying models. After these retraining steps, we measure if the recommendations have changed on a system-wide level. In particular, we analyze if the recommendations have become more concentrated on a small subset of the items or not.

Our experiments show that all investigated types of algorithms, both recent neural ones and heuristic-based ones, may lead to undesired concentration effects over time. Furthermore, we find that even though the prediction accuracy of some algorithms is often comparable, they may exhibit largely different concentration tendencies and, as a result, recommend very different sets of items to users in the end. This observation is particularly important from a practical point of view, since such differences might go unnoticed when an algorithmic comparison is solely based on accuracy measures. In practice, we are generally interested in a system that makes highly accurate predictions but does not exhibit undesired reinforcement tendencies. In an additional simulation experiment, we therefore investigated the effects when applying a re-ranking strategy to avoid too many repeated recommendations. The experiment provides indications that relatively simple strategies can help to counteract the undesired effects without a loss in accuracy.

## 2 PREVIOUS WORKS

*Catalog coverage* is a traditional quality factor in recommender systems and measures how many of the available items are presented to users in their top-n lists [8]. In our work, we adopt the definition from [8] and additionally measure how coverage develops over time. In other works, catalog coverage is often referred to as *aggregate diversity*, e.g., in [2]. One assumption is that higher aggregate diversity will ultimately lead to higher *sales diversity*, as investigated for example in [13]. In our work, we also assume that higher aggregate diversity (i.e., coverage) has positive effects. Besides aggregate diversity, researchers also frequently investigate diversity at the individual level, see [4]. Aggregate and individual diversity are, however, not necessarily correlated [13]. One can, for example, recommend the same set of highly diverse items to everyone, which does however not lead to high aggregate diversity [21].

Various strategies were proposed to increase catalog coverage, typically based on re-ranking the top-n items returned by an accuracy-optimized algorithm [2, 21]. Such strategies have the advantage that the resulting re-ranked lists remain highly accurate, i.e., only small compromises on accuracy have to be made. In our work, we investigate the effects of a comparably simple re-ranking strategy in a long-term perspective. Alternative model-based strategies for counteracting in particular popularity biases for traditional recommendation scenarios were proposed, e.g., in [1, 11].

Using simulation-based techniques as a research method, e.g., in the form of agent-based modeling, has a long tradition in various fields, e.g., in managerial science [20]. Simulation-based research is however comparably rare in the field of RS. Recently, Zhang et al. [23] used agent-based simulation to analyze longitudinal effects of recommender systems. Among other aspects, their simulations revealed a phenomenon called *performance paradox*, where it turned out that a strong reliance by users on the recommendations may lead to suboptimal performance development over time. It was also found that RS can concentrate on a small set of items. Such concentration biases, also measured in terms of the Gini index, were previously explored in [11]. Here, the authors as in [23] observed in their simulation approach that some algorithms may lead to a concentration over time. However, some algorithms were also suited to

increase catalog coverage and to decrease the Gini index over time. Our work continues the lines of research presented in [5, 6, 11, 23]. However, differently from previous works we consider session-based recommendation scenarios where no long-term preferences are available.

### 3 METHODOLOGY

Our general research methodology is based on simulation principles that were also adopted in [11, 23]. As a starting point, we use datasets containing real user interactions. We focus on the music domain and use two public datasets that include session information. One contains listening histories from the online music service *last.fm* [19] called *30Music*; the other one, called *#nowplaying*, was extracted from music-related Twitter messages in [22]. Regarding dataset characteristics, the *#nowplaying* dataset comprises 1.2M listening events in 146k sessions for 61k items. The *30Music* dataset is even larger, with 2.8M events 166k sessions for 446k items, i.e., sessions here are generally longer as well. Note that both datasets contain user IDs and that long-term listening histories are available. Since we focus on session-based recommendation problems, we do not take these long-term models into account when recommending.

The main simulation procedure is as follows:

- (1) For each session in the dataset, we select one track as a seed for a new listening session.<sup>1</sup>
- (2) We generate playlists from the seed track using different session-based algorithms and measure the characteristics of the recommendations at a system-wide level.
- (3) We assume that a certain fraction of the tracks in the playlist are listened to by the users and we add these simulated listening events to the dataset.
- (4) We update the models at defined intervals and continue with Step (1).

In our experiments, we used the following configuration:

- As seeds in **Step 1**, we randomly selected one of the tracks played in each session. We also made experiments with other seeds, e.g, the most frequently played track in a session. The obtained results were similar to the random seeds in terms of the general characteristics, which is why we omit them from this paper.
- In **Step 2**, we created recommendations of list length 30. If all 30 tracks are listened, this roughly corresponds to a two hours music experience in case of pop songs.
- In **Step 3**, we assume that the users on average consume about one third of the recommendations, i.e., 10 tracks, based on the observations regarding adoption rates in the music domain in [6, 12]. The selection of the tracks was done randomly. In general, modifying this parameter would result in slower or faster changes in the behavior of the recommender, but it would not impact the general characteristics of the emerging phenomena.
- We retrain the models in **Step 4** after having generated artificial playlists 3 times. Assuming that models in real-world deployments are retrained over night, there would be 3 sessions per day before the models are updated.

Measurements are taken in Step 2 after each re-training step. To see how recommendations change over time for a given set of items, we repeatedly took the following measures for the tracks that were used in the first simulation round:

- the *Gini* index to assess the concentration of the recommendations on certain items. Higher values mean higher concentration [11, 23];
- catalog *coverage* in terms of the absolute number of different items appearing in the top-n lists;
- the average item *popularity* in terms of the number of times a recommended track appears in the dataset;
- the information retrieval measures *precision*, *recall*, and F1 at list length 10 as accuracy measures.

<sup>1</sup>Alternative strategies with longer seed sessions are possible as well, but beyond the scope of this paper.

In our work, we seek to understand differences between algorithms in terms of their longitudinal effects. Depending on the application domain, the choice of the algorithm can then be based on these observations. We consider algorithms from different families in our simulations, as shown in Table 1. The hyper-parameters of the algorithms were optimized for accuracy on the training data. To ensure reproducibility we share all code and data used in the experiments online, including the configuration for splitting the data and the parameters used for each algorithm.<sup>2</sup>

Table 1. Algorithms used in the Comparison

GRU4REC	The first widely-used neural approach to session-based recommendation, based on RNNS [9].
NARM	An attention-based neural method [14], often leading to competitive results [16].
SKNN	A nearest neighbor technique that shows competitive results in a number of domains [15].
CAGH	A simple yet often effective baseline proposed in [3], which recommends the greatest hits of artists that are similar to those appearing in the seed tracks.

## 4 RESULTS

We report results for three experiments, discussed in Section 4.1 to Section 4.3.

### 4.1 Experiment 1: Analysis of Initial Recommendations

In the first measurement, we determine precision and recall for the different algorithms for the first round of recommendations, i.e., on the original data, and we additionally measure the Gini index, coverage, and the average popularity of the recommended items.

Table 2. Results for first simulation round for the #nowplaying dataset.

Algorithm	F1	Precision	Recall	Gini	Popularity (abs.)	Popularity (rel.)	Coverage
SKNN	<b>0.1550</b>	0.1482	<b>0.1624</b>	0.4782	57,7683	39,8138	<b>61,161</b>
NARM	0.1481	<b>0.1490</b>	0.1472	0.5982	65,7828	48,0066	59,578
GRU4REC	0.1227	0.1175	0.1283	<b>0.4169</b>	<b>22,7044</b>	<b>4,7920</b>	61,119
CAGH	0.0002	0.0005	0.0001	0.9301	171.7474	153.6176	24,718

The results for the #nowplaying dataset are shown in Table 2. In terms of accuracy measures, we find that SKNN and NARM are working best in this experiment. GRU4REC works slightly worse in this setup, where we only use the first element of a session to create the playlist. CAGH, although often competitive in alternative setups, does not work in a satisfactory way in such a cold-start setup. Regarding the other metrics, we find that among the well-performing techniques, NARM has both the highest concentration bias and the strongest tendency to recommend popular items (see column “Popularity (abs.)”). This is important because it shows that deep learning based techniques such as GRU4REC and NARM, despite comparable performance in accuracy, can recommend largely different items to users in their top-10 lists. SKNN represents the middle ground here, but still leans quite strongly to recommend popular items. The column “Popularity (rel.)” shows the difference between the average popularity of the recommendations and the seed track. All algorithms recommend tracks that are more popular than the seed tracks, with GRU4REC being the best in terms of retaining the original popularity level. CAGH, by design, is worst here, as it only recommends greatest hits of artists.

<sup>2</sup><https://github.com/andrebola/session-rec-effect>

The differences in terms of *coverage* among SKNN, NARM, and GRU4REC are low and all of them recommend almost all of the about 61k different tracks at least once. The results for the 30Music datasets shown in Table 3 are similar in terms of the accuracy measures. On this dataset, however, SKNN is also favorable in terms of the beyond-accuracy measures.

Table 3. Results for the first simulation round for the 30Music dataset.

Algorithm	F1	Precision	Recall	Gini	Popularity (abs.)	Popularity (rel.)	Coverage
SKNN	<b>0.1988</b>	<b>0.1802</b>	0.2218	<b>0.5629</b>	<b>21,6084</b>	<b>15,5684</b>	<b>429,338</b>
NARM	0.1955	0.1697	<b>0.2306</b>	0.7116	23.9657	17.9276	365,736
GRU4REC	0.1537	0.1318	0.1844	0.6547	24,0323	18,0633	397,470
CAGH	0.0000	0.0000	0.0000	0.9340	83.2055	77.0883	141,257

#### 4.2 Experiment 2: Longitudinal Analysis of Concentration, Coverage, and Popularity Effects

In this experiment, we repeatedly generated playlists by randomly choosing seed tracks for each training session. In total, we made 30 simulation rounds. After each round we added the simulated interactions to the dataset and we re-trained the models in every third round, leading to the 10 iterations shown in Figure 1.

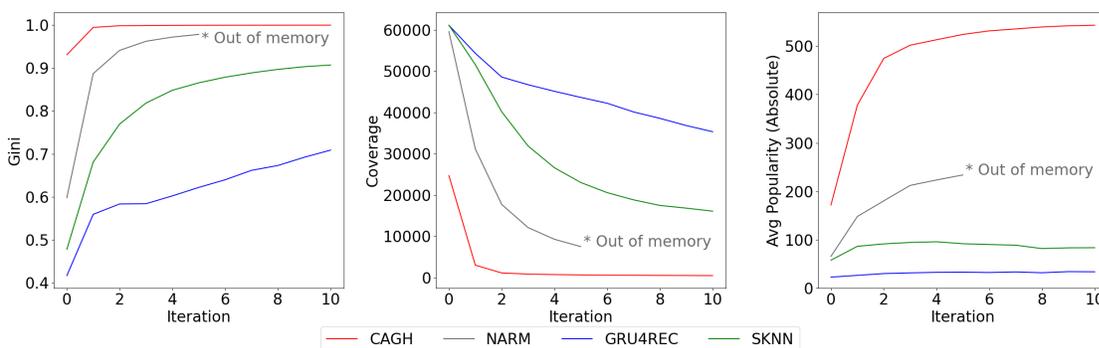


Fig. 1. Simulation Results for the #nowplaying Dataset. NARM ran out of memory (>64 GB) after 5 iterations as we add more data to the training set. Additional simulations (not shown here) in which we created playlists for only 20% of the data in each round confirmed the trends observed for the full datasets.

Looking at the Gini index (Figure 1), we observe that all algorithms in this comparison lead to an increased concentration effect over time. As expected from the results after the initial recommendations, the concentration increases most strongly when using NARM (excluding again CAGH), and it increases more slowly for GRU4REC. The development of the coverage metric follows this trend, as shown in Figure 1, i.e., the coverage of all algorithms goes down steadily during the simulation, with the strongest effect observed for NARM and the weakest for GRU4REC. Interestingly, the average popularity level remains mostly constant for SKNN and GRU4REC, with GRU4REC generally having a lower popularity bias than SKNN. For NARM, the popularity bias however increases over time.

The simulation results for 30Music are shown in Figure 2. The general observations are similar to those obtained for the #nowplaying dataset, i.e., all algorithms lead to a concentration over time and to a reduced coverage. As expected from the results of the initial recommendations (Table 2 and Table 3), we can however see that SKNN behaves favorably on this dataset in terms of concentration and coverage.

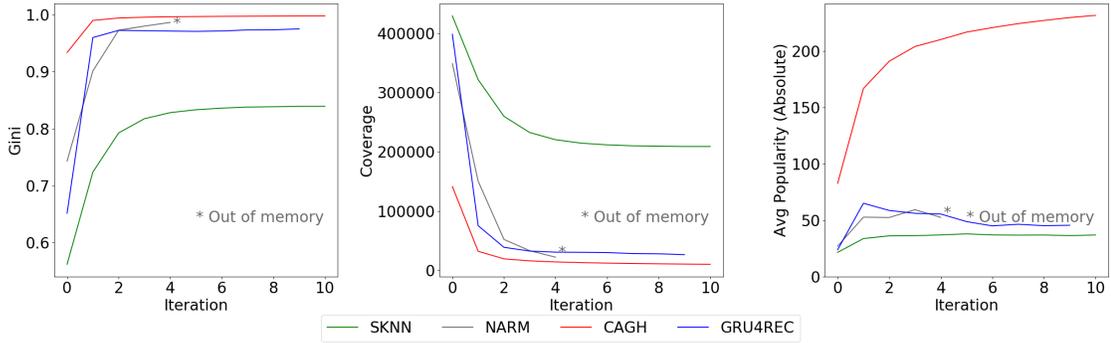


Fig. 2. Simulation Results for the 30Music Dataset. NARM again ran out of memory after a few iterations.

### 4.3 Experiment 3: Longitudinal Effects of Using Reranking as a Countermeasure

In a final set of experiments our goal was to investigate the effectiveness of applying reranking strategies to avoid concentration effects as was done, e.g., [2]. We analyzed the effects of two basic heuristics:

- In *Reranking Strategy 1*, our goal was to avoid recommending tracks too often that were recommended frequently in previous rounds to all users. Technically, in each round we count the number of times an item  $i$  was recommended, denoted as  $previous\_recs(i)$ . In the following round, we penalize frequently recommended tracks by moving them back in the recommendation lists. The penalty  $p$  in terms of the number positions to move the item  $i$  back is computed heuristically as  $p = 10 * \log(previous\_recs(i))$ .
- *Reranking Strategy 2* is personalized, and it tries to avoid recommendations that an individual user has consumed previously in the same session. Technically, for each user  $u$ , we count the number of times it has consumed an item  $i$  in the session, denoted as  $previous\_consumptions(i, u)$ . The penalty  $p$  for user  $u$  and item  $i$  is consumed as  $p = 10 \times previous\_consumptions(i, u)$ . For example, if a user has listened to a track two times before, the track will be moved back 20 positions in the new recommendation list.

The results show that both re-ranking strategies are effective, but in slightly different ways. When avoiding to repeatedly recommend the same tracks to everyone (*Reranking Strategy 1*), we can observe that the increase in concentration can be slowed down for all algorithms except again for CAGH (compare Figure 3a). For the case of individualized reranking *Reranking Strategy 2*, the increase of the concentration bias can be stopped at a certain level for SKNN and GRU4REC. For NARM, however, which exhibited a strong concentration bias already at the beginning, this personalized reranking does not seem to be very effective. This phenomenon can be attributed to the limited level of personalization of NARM, as shown in Table 2. Overall, however, the results also indicate that already simple reranking strategies can be effective countermeasures to avoid undesired concentration effects.

Furthermore, the reranking strategies do also not lead to a loss in accuracy. Looking at the precision and recall values obtained in our simulation experiment, we see that the accuracy of GRU4REC for both reranking strategies remains almost constant; for SKNN, the performance is even slightly increased, as observed previously for the music domain in [10]. These results are consistent for both datasets. Specifically, if we average precision and recall over all iterations for GRU4REC without reranking on the #nowplaying dataset, both precision and recall are at about 0.11. Applying

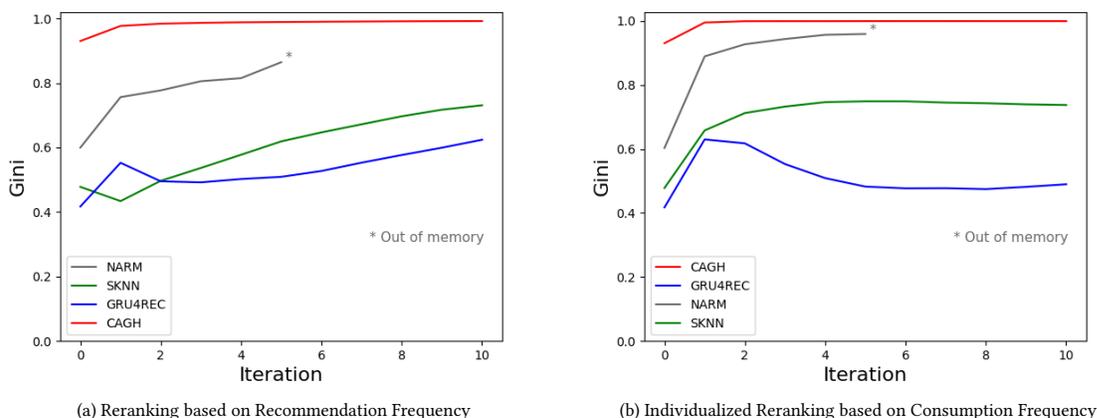


Fig. 3. Simulation Results (Reranking) for the #nowplaying Dataset.

either reranking strategy only leads to changes at the third place after the decimal. For SKNN, precision and recall even go slightly up with both strategies from about 0.13 up to 0.16.

Finally, additional measurements, which we omit for space reasons here, show that coverage also ceases to go down for SKNN and GRU4REC when a personalized reranking strategy is applied, and that the popularity bias continues to remain stable. The same measurements were furthermore made for the 30Music dataset (except for NARM, again due to high computational costs). The results are again generally in line with what was observed for the #nowplaying dataset.

## 5 CONCLUSION

In this work we analyzed through a simulation-based approach to what extent modern approaches to session-based recommendation exhibit certain potentially undesired biases, e.g., recommending the same set of items to everyone, and if these biases are reinforced over time. Our analysis in the music domain shows that, unlike to the findings in [11], all examined approaches can have such tendencies, although to a different extent. For practitioners, our work implies that such effects should be considered in the choice of algorithms, in particular because accuracy measures alone do not reveal such important differences between algorithms. Furthermore, our first experiments with personalized and non-personalized reranking strategies showed that already simple heuristics can be helpful countermeasures. On a methodological level, we see our work as another step to move beyond today’s “single-snapshot” analysis of recommendation algorithms, which does not allow us to investigate or simulate longitudinal effects of such systems.

Our work so far is limited to one particular domain and a specific set of assumptions used in the simulation, e.g., that the observed interactions all come from the recommendations, which leads to an amplification of the observed effects. Our future works include the investigation of other domains like e-commerce and the consideration of scenarios where long-term preference information about the users can be leveraged to diversify and de-bias the recommendations.

## ACKNOWLEDGMENTS

This research has been partially supported by Kakao Corp.

## REFERENCES

- [1] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Controlling Popularity Bias in Learning-to-Rank Recommendation. In *Proceedings of the 11th ACM Conference on Recommender Systems (RecSys 2017)*. 42–46.
- [2] Gediminas Adomavicius and YoungOk Kwon. 2012. Improving Aggregate Recommendation Diversity Using Ranking-Based Techniques. *IEEE Trans. on Knowl. and Data Eng.* 24, 5 (2012), 896–911.
- [3] Geoffray Bonnin and Dietmar Jannach. 2014. Automated generation of music playlists: Survey and experiments. *Comput. Surveys* 47, 2 (2014), 1–35.
- [4] Pablo Castells, Neil J. Hurley, and Saul Vargas. 2015. Novelty and Diversity in Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). 881–918.
- [5] Allison JB Chaney, Brandon M Stewart, and Barbara E Engelhardt. 2018. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys 2018)*. 224–232.
- [6] Andres Ferraro, Dmitry Bogdanov, Xavier Serra, and Jsson Yoon. 2019. Artist and style exposure bias in collaborative filtering based music recommendations. In *Proceedings of the 1st Workshop on Designing Human-Centric Music Information Research Systems*. 8–10.
- [7] Daniel Fleder and Kartik Hosanagar. 2009. Blockbuster Culture’s Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science* 55, 5 (2009), 697–712.
- [8] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. 2004. Evaluating Collaborative Filtering Recommender Systems. *ACM Trans. Inf. Syst.* 22, 1 (2004), 5–53.
- [9] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In *4th International Conference on Learning Representations, ICLR 2016*.
- [10] Dietmar Jannach, Lukas Lerche, and Iman Kamehkhosh. 2015. Beyond “Hitting the Hits” - Generating Coherent Music Playlist Continuations with the Right Tracks. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys 2015)*. 187–194.
- [11] Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. 2015. What recommenders recommend: an analysis of recommendation biases and possible countermeasures. *User Modeling and User-Adapted Interaction* 25, 5 (2015), 427–491.
- [12] Iman Kamehkhosh, Geoffray Bonnin, and Dietmar Jannach. 2020. Effects of Recommendations on the Playlist Creation Behavior of Users. *User Modeling and User-Adapted Interaction* 30 (2020), 285–322.
- [13] Dokyun Lee and Kartik Hosanagar. 2019. How Do Recommender Systems Affect Sales Diversity? A Cross-Category Investigation via Randomized Field Experiment. *Information Systems Research* 30, 1 (2019), 239–259.
- [14] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural Attentive Session-based Recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17)*. 1419–1428.
- [15] Malte Ludewig and Dietmar Jannach. 2018. Evaluation of Session-based Recommendation Algorithms. *User-Modeling and User-Adapted Interaction* 28, 4–5 (2018), 331–390.
- [16] Malte Ludewig, Noemi Mauro, Sara Latifi, and Dietmar Jannach. 2019. Performance Comparison of Neural and Non-Neural Approaches to Session-based Recommendation. In *Proceedings of the 2019 ACM Conference on Recommender Systems (RecSys 2019)*.
- [17] Eli Pariser. 2012. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Books, London.
- [18] Massimo Quadrana, Paolo Cremonesi, and Dietmar Jannach. 2018. Sequence-Aware Recommender Systems. *Comput. Surveys* 54 (2018), 1–36. Issue 1.
- [19] Roberto Turrin, Massimo Quadrana, Andrea Condorelli, Roberto Pagano, and Paolo Cremonesi. 2015. 30Music Listening and Playlists Dataset. In *RecSys 2015 Poster Proceedings*. Dataset: <http://recsys.deib.polimi.it/datasets/>.
- [20] Friederike Wall. 2016. Agent-based modeling in managerial science: an illustrative survey and study. *Review of Managerial Science* 10, 1 (2016), 135–193.
- [21] Qi Wang, Jijun Yu, and Weiwei Deng. 2019. An adjustable re-ranking approach for improving the individual and aggregate diversities of product recommendations. *Electron. Commer. Res.* 19, 1 (2019), 59–79.
- [22] Eva Zangerle, Martin Pichl, Wolfgang Gassler, and Günther Specht. 2014. #nowplaying Music Dataset: Extracting Listening Behavior from Twitter. In *Proceedings of the First International Workshop on Internet-Scale Multimedia Management (WISMM '14)*. 21–26.
- [23] Jingjing Zhang, Gediminas Adomavicius, Alok Gupta, and Wolfgang Ketter. 2020. Consumption and Performance: Understanding Longitudinal Dynamics of Recommender Systems via an Agent-Based Simulation Framework. *Information Systems Research* forthcoming (2020).