



# Enhancing Sequential Music Recommendation with Negative Feedback-informed Contrastive Learning

Pavan Seshadri

pseshadri9@gatech.edu  
Georgia Institute of Technology  
Music Informatics Group  
Atlanta, Georgia, USA

Shahrzad Shashaani

shahrzad.shashaani@tuwien.ac.at  
TU Wien  
Faculty of Informatics  
Vienna, Austria

Peter Knees

peter.knees@tuwien.ac.at  
TU Wien  
Faculty of Informatics  
Vienna, Austria

## ABSTRACT

Modern music streaming services are heavily based on recommendation engines to serve content to users. Sequential recommendation—continuously providing new items within a single session in a contextually coherent manner—has been an emerging topic in current literature. User feedback—a positive or negative response to the item presented—is used to drive content recommendations by learning user preferences. We extend this idea to session-based recommendation to provide context-coherent music recommendations by modelling negative user feedback, i.e., skips, in the loss function.

We propose a sequence-aware contrastive sub-task to structure item embeddings in session-based music recommendation, such that true next-positive items (ignoring skipped items) are structured closer in the session embedding space, while skipped tracks are structured farther away from all items in the session. This directly affects item rankings using a K-nearest-neighbors search for next-item recommendations, while also promoting the rank of the true next item. Experiments incorporating this task into SoTA methods for sequential item recommendation show consistent performance gains in terms of next-item hit rate, item ranking, and skip down-ranking on three music recommendation datasets, strongly benefiting from the increasing presence of user feedback.

## CCS CONCEPTS

• Information systems → Recommender systems; Music retrieval.

## KEYWORDS

Sequential Recommendation, Music Recommendation, Negative Feedback, Contrastive Learning

## ACM Reference Format:

Pavan Seshadri, Shahrzad Shashaani, and Peter Knees. 2024. Enhancing Sequential Music Recommendation with Negative Feedback-informed Contrastive Learning. In *18th ACM Conference on Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3640457.3688188>



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

RecSys '24, October 14–18, 2024, Bari, Italy  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0505-2/24/10  
<https://doi.org/10.1145/3640457.3688188>

## 1 INTRODUCTION

The advent of music recommender systems (MRS) [21] has revolutionized the music industry, evidenced by their widespread integration into music streaming platforms. Sequential music recommendation, a core task within MRS, focuses on extending user sessions by suggesting the next track. Traditional methods rely on user history to learn preferences through techniques like collaborative filtering [13], often facing the challenge of the cold start problem for new users or tracks [22]. The model must guess preferences until the user and/or track has interacted with the system enough to learn a profile [7]. Sequential recommendation addresses this by learning session-level relationships, enabling more accurate predictions based on recent interactions.

This study aims to leverage implicit and explicit within-context preference signals present in listening sessions to learn robust profiles for sequential recommendations. We investigate learning session-level information and the effects of incorporating user feedback through a contrastive learning task. We consider CNN-, RNN-, and transformer-based architectures, influenced by SoTA methods for sequential retail recommendation, and show a consistent improvement for all architectures through our extension. To our best knowledge, learning from negative signals/user feedback has not been explored thoroughly for sequential music recommendation, most likely due to a lack of public data containing user feedback, especially time-aware feedback. In fact, many public music recommendation datasets were collected before the streaming boom, where logged listening histories would primarily be sourced from user selection, leading to a low source of negative signals [21].

To account for the change of listening behavior, in this work, we investigate three frequently used datasets with different characteristics across different periods, ranging from personal-collection-sourced listening profiles with few skips, to a large volume of streaming connected logs with increased skipping events, to a dedicated skip prediction dataset. We show that an increasing amount of negative samples can be progressively exploited to learn effective session-level representations.

## 2 RELATED WORK

Sequential recommendation systems can generally be divided into two types: *session-aware* systems leverage session-level history from identifiable users, while *session-based* systems ignore user-labels and aim to build user-agnostic representations using solely discrete sessions [18]. In this study, we investigate a *session-based* system that implicitly learns contextual information through anonymous listening sessions.

Several session-based approaches have been proposed for retail recommendation tasks. *CASER* [25] and *NextItNet* [29] employ convolutional filters to learn effective representations. *GRU4Rec* [8] is a Recurrent Neural Network (RNN) recommender system to generate recommendations for short session-based data by modeling entire user sessions. *SLi-Rec* [28] combines both long-term and short-term user modeling and introduces time-aware and content-aware controllers into the RNN structure. Other session-based methods make use of attention mechanisms. *BERT4Rec* [24] uses the bidirectional attention mechanism from BERT [5] to learn a robust item embedding space for sequential recommendation. *SASRec* [11] is a self-attention based sequential model to handle sparse datasets and capture longer-term semantics.

Sequential approaches have also been proposed for music recommendation tasks. These include *CoSERNN* [6], which learns from contextual information such as device type and timestamps to learn embeddings, and Online Learning to Rank for Sequential Music Recommendation [17], which learns from music tags describing content information for online learning to rank.

Early work on incorporating negative user feedback for music recommendation has considered ad-hoc adjustments based on content and context similarity, e.g. [1, 10, 15]. In more recent work, performing a study close to our task, Wen et al. [27] investigate leveraging implicit user feedback, immediately after clicking for video and music recommendations to improve WRMF [9] and BPR [19]. Park and Lee [16] exploit item, document, or track negative feedback, and negative preference, for contrastive learning.

In contrast to existing approaches, our method focuses on modeling context within sequences rather than relying completely on user taste or long-term preference profiles. Since we formulate the core idea of our method as a learning objective via a loss term, it can be incorporated into a variety of existing recommendation algorithms.

## 3 METHOD

### 3.1 Problem Statement

Given an item history sequence  $S$ , standard feedback-agnostic approaches to sequential recommendation aim to predict the next item  $i$  at time  $t + 1$  for each  $i \in S$ . In other words, if  $F(S)$  is a standard sequential recommender, operating on  $S = \{i_0, i_1, \dots, i_t\}$ , then  $F(S) = \{i_1, i_2, \dots, i_{t+1}\}$ . As this approach is feedback-agnostic, it is assumed that each item in  $S$  is a valid recommendation.

Including user feedback, however, induces the possibility that certain items in  $S$  may not be valid recommendations, and thus solely predicting the item at the next time step is not optimal. To combat this, we separate the items in  $S$  with positive response  $P$  and negative response  $N$  into disjoint sets. As in [23], We re-formulate the problem to instead desire the next-positive item for each  $i_t$ .

If  $j \in T_p$  denotes the set of time steps of positive items in  $S$ , then for each  $i_t$ , we aim to predict the next-positive item in  $S$ ,  $i_m$ , such that  $m = \min_j \{j \in T_p \mid j > t\}$ . The difference  $m - t$  denotes the number of negative items between  $i_t$  and its next positive item  $i_m$ . To learn  $F(S)$ , we model a probability distribution,  $p(t_{i+1} = i \mid S_t), \forall i \in I$ , which effectively provides a ranking of probabilities to select the desired next item. For each  $i_t \in S$ , we aim to raise the predicted probability of positive item  $i_m$ , while penalizing the probability of all negative items,  $N = \{i_{t+n} \mid 0 < n < m - t\}$ .

### 3.2 Incorporating Negative Feedback

For our core task of sequential recommendation, we target learning an inner product space, such that  $\hat{\mathbf{E}} \cdot \mathbf{M}^T$  produces distances which can be mapped to the probability space described in Section 3.1. Thus, these distances correspond to the ranking of each item during inference. To incorporate negative feedback, we propose a sequence-aware contrastive task used to affect these distances by regularizing the learned item embeddings, which in turn affects the ranking of positive items against negative items.

While feedback-agnostic systems essentially model a similarity function where  $\hat{\mathbf{E}}(i_t) \sim \mathbf{E}(i_{t+1})$ , we aim to regularize the item embedding vectors to encourage our system to instead produce  $\hat{\mathbf{E}}(i_t) \sim \mathbf{E}(i_{t+m})$ , where  $i_m$  is the next-positive example from  $i_t$ . To more robustly learn this embedding space, we also wish to treat all skipped tracks in a sequence as negative examples.

To incorporate this into our model, we simply add an additional contrastive loss term to our learning objective. We employ InfoNCE noise-contrastive estimation [26] to essentially model a classifier that is optimized to distinguish the next-positive item in the input sequence against all negative items in the sequence via a discriminative function,  $f_k$ . Utilizing cosine similarity as the discriminative function in the contrastive term directly augments the learned inner product space to promote the ranking of the true next-positive item against the rankings of all negative items. Formally, the loss term per item is shown below:

$$\mathcal{L}_{NCE} = -\mathbb{E}_X \left[ \log \frac{f_k(\mathbf{i}_m, \hat{\mathbf{i}}_{t+1})}{f_k(\mathbf{i}_m, \hat{\mathbf{i}}_{t+1}) + \sum_{n_j \in N} f_k(n_j, \hat{\mathbf{i}}_{t+1})} \right] \quad (1)$$

where  $\mathbf{i}_m$  is the embedding of the next-positive sample, and  $\hat{\mathbf{i}}_{t+1}$  is the predicted embedding of the next item.

## 4 EXPERIMENTS

### 4.1 Models and Training

We employ the following baseline models used in sequential item recommendation for our task: GRU4Rec [8], CASER [25], SASRec [11], and BERT4Rec [24]. We use these specifically for the variety in which sequential patterns are learned: by recurrent networks, convolutional filters unidirectional, and bidirectional attention, respectively. We aim to demonstrate that our loss function incorporating negative feedback can be generalized into a variety of architectures for sequential music recommendation.

**4.1.1 Training Objective and Procedure.** We use equivalent implementations as laid out by the previous authors, except in sampling during training: empirically we find that using a sampled softmax with negative log likelihood improves training stability with a large vocabulary of items to recommend (~300-500K in this study). For each mini-batch for each session, we uniformly sample 1,000 unseen items and rank the target tracks alongside these. These 1,000 items are re-sampled each epoch, such that as training continues, the model continually learns to “rank” the target items with an increasing subset of the total tracks, as the number of unique tracks sampled for comparison increases.

For our loss function, in line with other contrastive recommendation systems, we append the negative log-likelihood loss term

for the sequential recommendation, with the contrastive loss term, shown below, where  $\alpha$  and  $\beta$  are constants:  $\mathcal{L} = \alpha\mathcal{L}_{NLL} + \beta\mathcal{L}_{NCE}$ .

For all models except BERT4Rec, we train them to predict the next item at each time step. BERT4Rec uses bidirectional training with a masked language model, and a token is added at the end of each sequence to predict the next track based on the overall context. The final and second-to-last items in each session are used as testing and validation sets, respectively.

**4.1.2 Hyperparameters and Implementation.** We select a maximum sequence length of 20 for all datasets, breaking down longer sessions into multiple discrete sessions. For all models, the embedding and hidden dimensions are both set to 128. We refer to the base implementations as per the respective authors for the configuration of the GRU4Rec and CASER models. For both SASRec and BERT4Rec, we stack 2 self-attentive encoder blocks with 8 attention heads. We set the masking proportion for BERT4Rec to  $p = 0.2$ . All parameters are initialized via truncated normal sampling with  $\mu = 0$ ,  $\sigma = 1$  in range  $[-0.02, 0.02]$ . We tune the optimal  $\alpha, \beta \in [0.1, 0.2, 0.5, 0.75, 1]$  using the validation set. For the MSSD, we select  $\alpha = 1$ ,  $\beta = 0.5$ , and for both LFM-1K and LFM-2B, we select  $\alpha = 1$ ,  $\beta = 0.2$ . We use the ADAM optimizer [12] with a learning rate of 0.005, selected after tuning through the validation set with  $lr \in [0.0001, 0.0005, 0.001, 0.005]$ .

**4.1.3 Non-sequential baseline competitors.** We also compare the models to non-sequential session-based baselines, namely Weighted Regularized Matrix Factorization (WRMF) [9] and Bayesian Personalized Ranking (BPR) [19], and their negative feedback informed extensions proposed by Wen et al. [27]. We evaluate both extensions: -BL uses post-click feedback as corrections for preference labels; -NR probabilistically samples items across all types of feedback.

## 4.2 Datasets

For this study we use three real-world datasets that differ in terms of represented listening tasks, skipping activity, and years of listening activity contained, namely the Music Streaming Sessions Dataset (MSSD) [2] using data from Spotify, the LFM-2B dataset [20], and the LFM-1K dataset [3, 4], both using data from Last.fm.<sup>1</sup>

MSSD contains ~160 million Spotify user sessions predominantly from the year 2018. We sample ~600K discrete sessions containing ~9M item interactions with ~600K total unique tracks, containing roughly 50% explicitly annotated skipped tracks.

The two Last.fm datasets contain user listening histories with timestamp information, which we split into sessions wherever timestamp differences over 20 mins occur. We interpret a track to have been skipped if the next event starts within 30 secs after the track has been started. We remove any session with less than 5 events. LFM-1K contains 19 million discrete listening events from 2005 to 2009 (i.e., before streaming was the dominant way of listening) with only a 0.8% skip rate, insufficient to learn from negative feedback. Thus, we sample sessions such that roughly half will contain a skip event. This provides a set of ~80K sessions spanning ~1.4M item interactions, containing ~300K tracks with a skip rate of ~11%. LFM-2B contains more than 2 billion listening events from 2005 until 2020. We specifically use the provided 2020 subset containing ~7.6M listening events. After processing, the used subset

contains ~14% skipped tracks. We assume that this subset already has a stronger connection to streaming-based listening behavior than LFM-1K. In all datasets, we discard all present user information and treat each session individually, i.e., as a new user.

## 4.3 Evaluation

For evaluating the core recommendation task of next predicting a desired item, we measure Hit Rate @{1, 5, 10, 20} and Mean Average Precision @10 to emphasize the top recommendations. This common evaluation strategy alone has the limitation the next-positive-item is assumed to be the only true desired output of a recommender. Moreover, in the case of discrete sessions, the final (and penultimate, etc.) items in the sequence may not have a next-positive example recorded in the dataset. Second, evaluation should reflect whether the penalization of negative items relative to positive items improves the learning of session-context overall to better predict unobserved items. We, therefore, evaluate an additional skip down-ranking task to quantify the reduction in the ranking of skipped tracks utilizing Mean Reciprocal Rank @10. A lowered MRR value for skips is considered an improvement and serves as a proxy for a better learned representation of context.

## 5 RESULTS AND DISCUSSION

### 5.1 Core-Recommendation Task

Table 1 shows the results for all models, comparing the baselines with negative feedback informed extensions (ours for the sequential approaches, [27] for the non-sequential) on all three datasets. We find that our feedback-informed loss term unanimously increases performance across datasets. Additionally, for the non-sequential baselines (WRMF, BPR) we find improvements consistent with those reported in [27].

We observe the greatest absolute performance across all models on the MSSD, followed by the LFM-2B and LFM-1K, respectively. With regards to sequential music recommendation, the specific type of listening behavior and context could possibly affect the model’s ability to model the interaction sequences effectively: the MSSD (collected 2019) is fully streaming based and contains a high proportion of algorithmically driven content and performs best, while the LFM-1K (collected 2009) contains mostly user curation and performs worst. LFM-2B (collected 2020) contains a mix of both and performs in the middle. Performance gains appear to correlate with skip percentage in the dataset as expected: With more available negative feedback, the feedback-informed loss term can better regularize embeddings.

### 5.2 Skip down-ranking

Table 2 details the results of the skip down-ranking experiment. We observe a consistent decrease in scores (lower is better) after learning from negative feedback, with a notable exception in BERT4Rec on MSSD, where we observe an increase. Additionally, this dataset contains the highest proportion of skipped tracks in both the testing and training sets (roughly 50% for both), and this particular combination (BERT4Rec on MSSD) displayed the highest performance increase in the core-recommendation performance (18-74% increase on hit rate).

<sup>1</sup>Spotify: <https://open.spotify.com>, Last.fm: <https://last.fm>

Data	Metric	SASRec		BERT4Rec		GRU4Rec		CASER		WRMF			BPR		
		orig.	ours	orig.	ours	orig.	ours	orig.	ours	orig.	-BL	-NR	orig.	-BL	-NR
MSSD	HR@1	.377	<b>.410 (9%)</b>	.204	.355 (74%)	.210	.235 (12%)	.223	.251 (13%)	.203	.207	.211	.208	.213	.216
	HR@5	.615	<b>.628 (2%)</b>	.450	.608 (35%)	.398	.431 (8%)	.412	.455 (10%)	.400	.406	.411	.406	.413	.421
	HR@10	.696	<b>.706 (1%)</b>	.553	.693 (25%)	.491	.534 (9%)	.518	.530 (2%)	.488	.498	.505	.502	.510	.517
	HR@20	.767	<b>.774 (1%)</b>	.648	.767 (18%)	.600	.627 (5%)	.616	.649 (5%)	.597	.603	.608	.605	.609	.612
	MAP@10	.397	<b>.417 (5%)</b>	.185	.369 (99%)	.176	.193 (10%)	.195	.221 (13%)	.149	.156	.164	.160	.168	.173
LFM-2B	HR@1	.190	<b>.221 (16%)</b>	.101	.117 (16%)	.096	.102 (6%)	.102	.112 (10%)	.097	.098	.102	.098	.105	.107
	HR@5	.371	<b>.400 (8%)</b>	.227	.248 (9%)	.203	.221 (9%)	.197	.208 (6%)	.138	.142	.147	.143	.152	.169
	HR@10	.452	<b>.477 (6%)</b>	.292	.320 (10%)	.273	.291 (7%)	.281	.302 (7%)	.269	.273	.279	.276	.286	.294
	HR@20	.532	<b>.553 (4%)</b>	.366	.394 (8%)	.311	.342 (10%)	.326	.354 (9%)	.305	.311	.316	.314	.324	.348
	MAP@10	.188	<b>.219 (16%)</b>	.098	.110 (12%)	.078	.085 (9%)	.081	.088 (9%)	.062	.065	.067	.066	.072	.078
LFM-1K	HR@1	.152	<b>.181 (19%)</b>	.069	.086 (25%)	.048	.059 (23%)	.052	.071 (37%)	.042	.044	.047	.043	.050	.052
	HR@5	.301	<b>.330 (10%)</b>	.207	.230 (11%)	.182	.200 (10%)	.188	.198 (5%)	.139	.146	.177	.150	.153	.159
	HR@10	.392	<b>.421 (7%)</b>	.299	.320 (7%)	.261	.289 (11%)	.269	.293 (9%)	.270	.285	.294	.292	.298	.301
	HR@20	.478	<b>.491 (3%)</b>	.413	.433 (5%)	.388	.397 (2%)	.390	.404 (4%)	.346	.368	.375	.369	.374	.376
	MAP@10	.092	<b>.107 (16%)</b>	.049	.064 (31%)	.034	.042 (24%)	.037	.040 (8%)	.028	.031	.034	.032	.036	.038

Table 1: Hit Ratio @ [1, 5, 10, 20] and Mean Average Precision @ 10 for the sequential models (SASRec, BERT4Rec, GRU4Rec, CASER) and the non-sequential baselines (WRMF, BPR) on the three datasets. Non “orig.” models incorporate negative feedback. They are compared to their “orig.” baselines which do not model negative feedback. Numbers in parentheses show the relative increase in the percentage of the approach over the respective baseline; bold entries mark the better performing approach between the baseline and negative feedback-informed approach. The overall best performance is also highlighted in bold (i.e., SASRec-ours).

Metric		SASRec		BERT4Rec	
		orig.	ours	orig.	ours
MRR@10	MSSD	.960	.950 (-1%)	.840	.950 (13%)
MRR@10	LFM-2B	.969	.911 (-6%)	.953	.950 (-0.3%)
MRR@10	LFM-1K	.731	.731 (0%)	.540	.460 (-15%)

Table 2: Mean Reciprocal Rank @ 10 on skip targets for the best sequential models (SASRec, BERT4Rec) on the three datasets. Lower values are better.

### 5.3 Model Comparison

Through our picked baselines, we observe multiple methods of learning from sequential interactions. SASRec and BERT4Rec both employ self-attentive architectures in uni- and bidirectional representations, resp. GRU4Rec uses recurrent networks, while CASER uses convolutional filters. We generally find superior performance with the self-attentive architectures, with SASRec consistently performing best across all datasets. For all datasets, we observe similar relative improvements in learning from negative feedback across all sequential models, except for BERT4Rec on the MSSD.

While we observe the largest improvements in terms of HR and MAP, as per the results of the skip down-ranking, we also observe that skipped testing targets raise in rank. There are a couple of hypotheses that could support this: first, BERT4Rec is the only bidirectional approach of the 4 frontend sequential models. As we observe higher absolute performance consistently using SASRec (a similar architecture, differing in training objective for causal inference), the “forward” direction may be simply more impactful

than the “backward” direction for sequential music recommendation. As our feedback-informed loss function explicitly models the forward direction, it could simply be enforcing this while training, leading to greater performance. Additionally, as the MSSD contains a high amount of algorithm driven content with user feedback, this particular combination may be learning the underlying algorithm rather than user intent. This could lead to significantly greater performance in predicting the end tracks, but also lead to making the same mistakes as the underlying algorithm (i.e., recommending tracks that will be skipped).

## 6 CONCLUSIONS

We have presented a contrastive learning objective that incorporates sequential negative feedback (i.e., skips) for next-track music recommendation. We showed a consistent positive impact across three real-world datasets encompassing different listening behaviors, and across four SoTA sequential recommendation architectures. We additionally identify the ability of our proposed feedback-aware systems on reliably down-ranking unseen skipped tracks within the testing set.

Direct future work would include the incorporation of explicit feedback (e.g. thumbs up/down) and hard negative/positive samples based on historical co-occurrences. Beyond the identified utility of incorporating negative feedback, our work points to open challenges to be addressed in future work, foremost a better understanding of user intents and different streaming behaviors, cf. [14], leading to more considerate models, both of data and in recommenders, and connected to that more specific evaluation strategies beyond single retrieval metrics focusing only on one aspect of systems.

## ACKNOWLEDGMENTS

This research was funded in whole or in part by the Austrian Science Fund (FWF) (10.55776/P33526). For open access purposes, the author has applied a CC BY public copyright license to any author-accepted manuscript version arising from this submission.

## REFERENCES

- [1] Klaas Bosteele and Etienne E. Kerre. 2009. A Fuzzy Framework for Defining Dynamic Playlist Generation Heuristics. *Fuzzy Sets and Systems* 160, 23 (2009), 3342–3358. <https://doi.org/10.1016/j.fss.2009.05.013>
- [2] Brian Brost, Rishabh Mehrotra, and Tristan Jehan. 2019. The Music Streaming Sessions Dataset. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 2594–2600. <https://doi.org/10.1145/3308558.3313641>
- [3] Óscar Celma. 2010. *Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space* (1st ed.). Springer. <https://doi.org/10.1007/978-3-642-13287-2>
- [4] Óscar Celma. 2022. *lastfm Music Recommendation Dataset*. <https://doi.org/10.5281/zenodo.6090214>
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, MN, USA, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [6] Casper Hansen, Christian Hansen, Lucas Maystre, Rishabh Mehrotra, Brian Brost, Federico Tomasi, and Mounia Lalmas. 2020. Contextual and Sequential User Embeddings for Large-Scale Music Recommendation. In *Proceedings of the 14th ACM Conference on Recommender Systems* (Virtual Event, Brazil) (RecSys '20). Association for Computing Machinery, New York, NY, USA, 53–62. <https://doi.org/10.1145/3383313.3412248>
- [7] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. *CoRR* abs/1708.05031 (2017), arXiv:1708.05031 <http://arxiv.org/abs/1708.05031>
- [8] 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* (San Juan, Puerto Rico), Yoshua Bengio and Yann LeCun (Eds.). <https://doi.org/10.48550/arXiv.1511.06939>
- [9] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining, ICDM, 2008*. IEEE, 263–272. <https://doi.org/10.1109/ICDM.2008.22>
- [10] Yajie Hu and Mitsunori Ogihara. 2011. NextOne Player: A Music Recommendation System Based on User Behavior. In *Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR, Miami, United States, 103–108*. <https://doi.org/10.5281/zenodo.1418301>
- [11] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *Proceedings of the 2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206. <https://doi.org/10.1109/ICDM.2018.00035>
- [12] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015* (San Diego, CA, USA), Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [13] Peter Knees and Markus Schedl. 2013. A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 10, 1, Article 2 (Dec 2013), 21 pages. <https://doi.org/10.1145/2542205.2542206>
- [14] Francesco Meggetto, Crawford Revie, John Levine, and Yashar Moshfeghi. 2021. On Skipping Behaviour Types in Music Streaming Sessions. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (Virtual Event, Queensland, Australia) (CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 3333–3337. <https://doi.org/10.1145/3459637.3482123>
- [15] Elias Pampalk, Tim Pohle, and Gerhard Widmer. 2005. Dynamic Playlist Generation Based on Skipping Behavior. In *Proceedings of the 6th International Conference on Music Information Retrieval, ISMIR, London, United Kingdom, 634–637*. <https://doi.org/10.5281/zenodo.1414932>
- [16] Minju Park and Kyogu Lee. 2022. Exploiting Negative Preference in Content-based Music Recommendation with Contrastive Learning. In *Proceedings of the 16th ACM Conference on Recommender Systems* (Seattle, WA, USA) (RecSys '22). Association for Computing Machinery, New York, NY, USA, 229–236. <https://doi.org/10.1145/3523227.3546768>
- [17] Bruno L. Pereira, Alberto Ueda, Gustavo Penha, Rodrygo L. T. Santos, and Nivio Ziviani. 2019. Online Learning to Rank for Sequential Music Recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) (RecSys '19). Association for Computing Machinery, New York, NY, USA, 237–245. <https://doi.org/10.1145/3298689.3347019>
- [18] Massimo Quadrona, Alexandros Karatzoglou, Balázs Hidasi, and Paolo Cremonesi. 2017. Personalizing Session-Based Recommendations with Hierarchical Recurrent Neural Networks. In *Proceedings of the Eleventh ACM Conference on Recommender Systems* (Como, Italy) (RecSys '17). Association for Computing Machinery, New York, NY, USA, 130–137. <https://doi.org/10.1145/3109859.3109896>
- [19] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) (UAI '09). AUAI Press, Arlington, Virginia, USA, 452–461. <https://doi.org/10.48550/arXiv.1205.2618>
- [20] Markus Schedl, Stefan Brandl, Oleg Lesota, Emilia Parada-Cabaleiro, David Penz, and Navid Rekabsaz. 2022. LFM-2b: A Dataset of Enriched Music Listening Events for Recommender Systems Research and Fairness Analysis. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (Regensburg, Germany) (CHIIR '22). Association for Computing Machinery, New York, NY, USA, 337–341. <https://doi.org/10.1145/3498366.3505791>
- [21] Markus Schedl, Peter Knees, Brian McFee, and Dmitry Bogdanov. 2022. *Music Recommendation Systems: Techniques, Use Cases, and Challenges*. Springer US, New York, NY, 927–971. [https://doi.org/10.1007/978-1-0716-2197-4\\_24](https://doi.org/10.1007/978-1-0716-2197-4_24)
- [22] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskis. 2015. *Music Recommender Systems*. Springer US, Boston, MA, 453–492. [https://doi.org/10.1007/978-1-4899-7637-6\\_13](https://doi.org/10.1007/978-1-4899-7637-6_13)
- [23] Pavan Seshadri and Peter Knees. 2023. Leveraging Negative Signals with Self-Attention for Sequential Music Recommendation. In *Proceedings of the 1st Workshop on Music Recommender Systems, co-located with the 17th ACM conference on Recommender Systems*.
- [24] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (Beijing, China) (CIKM '19). Association for Computing Machinery, New York, NY, USA, 1441–1450. <https://doi.org/10.1145/3357384.3357895>
- [25] Jiayi Tang and Ke Wang. 2018. Personalized Top-N Sequential Recommendation via Convolutional Sequence Embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) (WSDM '18). Association for Computing Machinery, New York, NY, USA, 565–573. <https://doi.org/10.1145/3159652.3159656>
- [26] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR* abs/1807.03748 (2018), arXiv:1807.03748 <http://arxiv.org/abs/1807.03748>
- [27] Hongyi Wen, Longqi Yang, and Deborah Estrin. 2019. Leveraging Post-Click Feedback for Content Recommendations. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) (RecSys '19). Association for Computing Machinery, New York, NY, USA, 278–286. <https://doi.org/10.1145/3298689.3347037>
- [28] Zeping Yu, Jianxun Lian, Ahmad Mahmood, Gongshen Liu, and Xing Xie. 2019. Adaptive User Modeling with Long and Short-Term Preferences for Personalized Recommendation. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 4213–4219. <https://doi.org/10.24963/ijcai.2019/585>
- [29] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M. Jose, and Xiangnan He. 2019. A Simple Convolutional Generative Network for Next Item Recommendation. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (Melbourne VIC, Australia) (WSDM '19). Association for Computing Machinery, New York, NY, USA, 582–590. <https://doi.org/10.1145/3289600.3290975>