

# Session-based Sequential Skip Prediction via Recurrent Neural Networks

Lin Zhu

Ctrip Travel Network Technology Co., Limited.  
Shanghai, P.R. China  
lizhonyx@163.com

Yihong Chen

Ctrip Travel Network Technology Co., Limited.  
Shanghai, P.R. China  
yihongchen@Ctrip.com

## ABSTRACT

The focus of WSDM cup 2019 is session-based sequential skip prediction, i.e. predicting whether users will skip tracks, given their immediately preceding interactions in their listening session. This paper provides the solution of our team **ekffar** to this challenge. We focus on recurrent-neural-network-based deep learning approaches which have previously been shown to perform well on session-based recommendation problems. We show that by choosing an appropriate recurrent architecture that properly accounts for the given information such as user interaction features and song metadata, a single neural network could achieve a Mean Average Accuracy (AA) score of 0.648 on the withheld test data. Meanwhile, by ensembling several variants of the core model, the overall recommendation accuracy can be improved even further. By using the proposed approach, our team was able to attain the 1st place in the competition. We have open-sourced our implementation at GitHub<sup>1</sup>.

## KEYWORDS

Recurrent Neural Network, Skip Prediction, Session Modeling

### ACM Reference format:

Lin Zhu and Yihong Chen. 2019. Session-based Sequential Skip Prediction via Recurrent Neural Networks. In *Proceedings of Proceedings of the ACM WSDM Cup 2019, Melbourne, Australia, Feb 11, 2019 (WSDM Cup '19)*, 4 pages. DOI: 10.1145/nmnnnnn.nmnnnnn

## 1 INTRODUCTION

The accurate prediction/modeling of user action has long been a topic of vital importance in large-scale web and mobile applications, the quality of the model or accuracy of prediction plays a vital role in various revenue generating applications, such as content recommendation [6], advertisement displaying [16], search result ranking [21], etc. Traditional

<sup>1</sup><https://github.com/linzhu123455/spotify-skip-prediction-top-1-solution>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*WSDM Cup '19, Melbourne, Australia*

© 2019 ACM. 978-x-xxxx-xxxx-x/YY/MM...\$15.00  
DOI: 10.1145/nmnnnnn.nmnnnnn

approaches, such as matrix Factorization [18] and its variants, seek to model global interactions, e.g. by learning low-dimensional user and item embeddings and predicting their interactions in the embedding space. While these algorithms are able to effectively model static user preferences, they fail to account for sequential and dynamical nature of user behavior. This limitation motivates the study of *session-based* user action modeling problem, where the terminology ‘session’ refers to a group of interactions that take place within a given time interval. A session usually has a goal, such as finding a good restaurant in a city, or listening to music of a certain style or mood.

In this year, the WSDM Cup, a competition held annually as part of the prestigious ACM International Conference on Web Search and Data Mining (WSDM), challenges competitors from all over the world to help and build a better user behavior modeling system [2]. The challenge dataset which consists of roughly 130 million music listening sessions with associated user interactions, was kindly produced by Spotify, a major online music streaming service with online music streaming service with over 190 million active users interacting with a library of over 40 million tracks. The task is to predict whether individual tracks encountered in the second half of a listening session will be skipped by a particular user.

To tackle this challenge, we focus on recurrent-neural-network-based deep learning approaches since they are able to deal with large datasets in an efficient and effective way, and have previously been shown to perform well on session-based recommendation problems. We show that by choosing an appropriate recurrent architecture that properly accounts for the given information such as user interaction features and song metadata, a single neural network could achieve a Mean Average Accuracy (AA) score of 0.648 on the withheld test data, achieving top 1 position on the challenge leaderboard. Meanwhile, by ensembling several variants of the core model, the overall recommendation accuracy can be improved even further.

The paper is organized as follows. A overview of the WSDM-Spotify challenge is presented in Section 2, Section 3 discusses the components of our approach, and the paper concludes in Section 4.

## 2 THE CHALLENGE DESCRIPTION

### 2.1 Dataset

The public part of the dataset consists of roughly 130 million listening sessions with associated user interactions and

context features on the Spotify service. The user interaction with each listened song is described by a set of features, such as the skip behavior, as well as the number of times the user did a seek forward or seek back within track (“hist\_user\_behavior\_n\_seekfwd” and “hist\_user\_behavior\_n\_seekback”, respectively), on the other hand, the context features include the hour of day that the listening happened (“hour\_of\_day”) and the type of context the playback occurred within (“context\_type”), etc. In addition to the public part of the dataset, approximately 30 million listening sessions are used for the challenge leaderboard. For these leaderboard sessions the participant is provided all the user interaction features for the first half of the session, but only the track ids for the second half. Meanwhile, the meta information of the tracks that users interacted with during these sessions are also provided, including numerous acoustic features and metadata (such as duration and release year).

## 2.2 Challenge

The task is to predict whether individual tracks encountered in a listening session will be skipped by a particular user. In order to do this, complete information about the first half of a users listening session is provided, while the prediction is to be carried out on the second half. Participants have access to metadata, as well as acoustic descriptors, for all the tracks encountered in listening sessions. The output of a prediction is a binary variable for each track in the second half of the session indicating if it was skipped or not.

## 2.3 Evaluation Metric

Although it would also be useful to predict the skip behavior across the entire session, as skip prediction can be used to avoid recommendation of a potential track to the user based on the user’s immediately preceding interactions, accurate prediction of whether the next immediate track is going to be skipped is therefore most important. Based on these considerations, Mean Average Accuracy is adopted as the primary evaluation metric for the challenge, with the average accuracy defined by

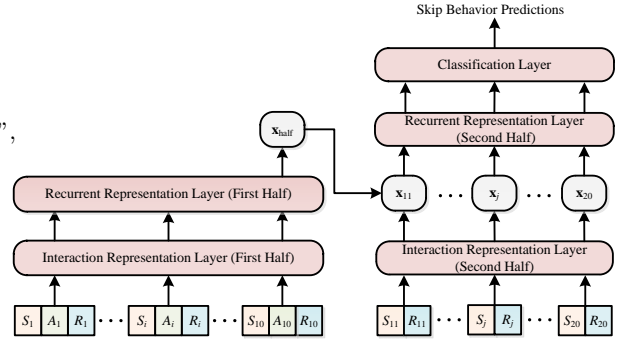
$$AA = \frac{\sum_{i=1}^T A(i) L(i)}{T}, \quad (1)$$

where  $T$  is the number of tracks to be predicted for the given session,  $A(i)$  is the accuracy at position  $i$  of the sequence,  $L(i)$  is the boolean indicator for if the  $i$ th prediction was correct.

# 3 OUR SOLUTION

## 3.1 Problem Formulation

The given information about the first half of a session be summarized as a series of triplets  $\{(S_i, A_i, R_i)\}_{1 \leq i \leq 10}$ , where  $S_i$ ,  $A_i$ , and  $R_i$  denote the song, the user interaction associated with the song, and the position of song listening event within the session, on the other hand, information provided about the second half of a session can be summarized as a



**Figure 1: The overall architecture of our skip prediction model.**

series of doublets  $\{(S_i, R_i)\}_{11 \leq i \leq 20}$ , with the user interaction information held out. The objective of the competition is to predict the user skip behavior associated with songs in the second half of the session. To integrate the provided information in a principled way, we parametrize the desired prediction model as a neural architecture that consists of three main components, as illustrated in Figure 1:

- **Interaction Representation Layers** that transforms triplets in the first half of a session and doublets in the second half of a session into feature representations;
- **Recurrent Representation Layers** that utilize recurrent neural network models to refine the triplet and doublet representations by exploiting the sequential nature of user sessions;
- A **Classification Layer** that aggregate the outputs of previous layers and calculate skip probabilities for songs in the second half of the session.

## 3.2 Interaction Representation Layers

As mentioned in Section 2, features for user interactions and the tracks that users interacted with during these sessions are provided, which can be straightforwardly used as their representations. Given the sensitivity of neural networks to input scaling [6, 19], numerical features are normalized to the interval  $[0, 1]$ , and categorical features such as ‘context\_type’ are mapped to dense vectors via embeddings, which are learned jointly with all other model parameters through back-propagation.

Additionally, we borrowed ideas from [10, 11], and treated each track as a word and each listening session as a sentence, Glove [23] was then used to train on these “sentences” to obtain additional 150-dimensional embeddings for songs.

As for the session position information, we simply incorporate it as an additional numerical feature. While more sophisticated positional encoding schemes [26] are readily applicable, we found that such a simple option already performs well on the challenge dataset.

All of the aforementioned features are accordingly concatenated to obtain vector representations for triplets  $\{(S_i, A_i, R_i)\}_{1 \leq i \leq 10}$

in the first half of the session, as well as doublets  $\{(S_i, R_i)\}_{11 \leq i \leq 20}$  in the second half of the session.

### 3.3 Recurrent Neural Networks for Session Modeling

Motivated by the tremendous success of deep neural networks (DNN) in a number of tasks such as image and natural language processing (NLP), various approaches have been proposed in recent years to adapt DNN algorithms for data mining tasks [3, 6, 13]. In particular, session-based user action modeling problem shares some similarities with NLP-related problems [14] as they both deal with sequences. As a result, the Recurrent Neural Networks (RNNs), which have proven to be suitable for handling various sequence modeling tasks [9, 12, 22, 25], have also been successfully applied to session-based user action modeling, achieving remarkable results of outperforming traditional recommendation methods by 15% to 30% in terms of various evaluation metrics [1, 7, 14, 20, 24].

The main difference between RNNs and conventional feed-forward neural network models is the existence of an internal hidden state in the units that recurrently process and memorize sequences of inputs. Traditional RNN models usually suffer from the vanishing gradient problem when the models need to process long sequences. The Gated Recurrent Unit (GRU) network is a technique proposed by Cho et al. [4] which aims to solve the problem of gradient vanishing in RNN. Its basic idea is to control the update of network states with an update gate and a reset gate. Formally, let  $\mathbf{x}_t \in \mathbb{R}^\alpha$  be the input vector in the  $t$  step and  $\alpha$  is the dimensionality of  $\mathbf{x}_t$ . The output vector (also the activation vector in GRU)  $\mathbf{o}_t \in \mathbb{R}^\alpha$  and the network state  $\mathbf{s}_t \in \mathbb{R}^\alpha$  is computed as:

$$\begin{aligned} \mathbf{o}_t &= (1 - \mathbf{u}_t) \odot \mathbf{o}_{t-1} + \mathbf{u}_t \odot \mathbf{s}_t, \\ \mathbf{u}_t &= \sigma(\mathbf{W}_u^x \cdot \mathbf{x}_t + \mathbf{W}_u^s \cdot \mathbf{s}_t), \\ \mathbf{s}_t &= \tanh(\mathbf{W}^x \cdot \mathbf{x}_t + \mathbf{W}^s \cdot (\mathbf{r}_t \odot \mathbf{o}_{t-1})), \\ \mathbf{r}_t &= \sigma(\mathbf{W}_r^x \cdot \mathbf{x}_t + \mathbf{W}_r^s \cdot \mathbf{s}_t), \end{aligned} \quad (2)$$

Based on the above considerations, we adopted GRUs to model user sessions. Concretely, for the first half of the session, we used triplet representations described in the previous section as input, and passed them through 2 stacked GRU networks, the final states of which were then concatenated as a vector  $\mathbf{x}_{\text{half}}$  to represent the entire first half session. Meanwhile, for the second half of the session, let the doublet representation for session position  $i$  be denoted as  $\mathbf{x}_i$ , we enrich it with  $\mathbf{x}_{\text{half}}$  as:

$$\mathbf{x}_i = \begin{bmatrix} \mathbf{x}_i \\ \mathbf{x}_{\text{half}} \\ \mathbf{x}_{\text{half}} \odot \text{ReLU}(\mathbf{x}_i) \end{bmatrix}, \quad (3)$$

where  $\odot$  is the Hadamard product, ReLU is a dense layer with Rectified Linear Units (ReLU) activations [8] to match the dimensions of  $\mathbf{x}_i$  and  $\mathbf{x}_{\text{half}}$ .  $\mathbf{x}_i$  returned by (3) can be considered as an enriched representation as it encodes the contextual information from the first and the second half of the session, as well as the information associated with the specific session position  $i$ , and all of these information can be beneficial for the considered skip prediction task.

### 3.4 Classification Layer

Given the enriched representation (3) as input, the classification layer is simply a multilayer perceptron with sigmoid output activation function and 2 layers of ReLU activations. Meanwhile, to inject more information into the model, we further adopt a multi task scenario where model not only predicts the user skip behavior, but also other user interactions such as 'context\_switch', 'no\_pause\_before\_play', and 'short\_pause\_before\_play'.

### 3.5 Implementation Details

The neural network model is implemented using Keras, a diagram of the model drawn using the Keras function "plot\_model" is given in Fig.2. Note that the user sessions can have unequal length, varying between 10 and 20 in the challenge dataset, we resolve this issue by padding short sequences with pre-specified constants. The model was trained with the Adam algorithm [17] with an initial learning rate of 0.0005, batch size of 2200.

### 3.6 Results

Based on the neural network architecture, a single trained model could achieve a AA score of 0.648 on the withheld test dataset, already achieving top 1 position on the competition leaderboard, we also tried several simple variants of the core model, for example by replacing the ReLU activations with exponential linear units (ELUs) [5], or by adding batch normalization layer [15] before the activations, or by varying batch size and number of hidden units. An ensemble of 6 such model variants could further boost the leaderboard accuracy to 0.651.

## 4 CONCLUSION

In this paper, we describe our solution to WSDM cup 2019, we focus on recurrent-neural-network-based deep learning approaches since they have previously been shown to perform well on session-based recommendation problems. We show that by choosing an appropriate recurrent architecture that properly accounts for the given information such as user interaction features and song metadata, a single neural network could achieve top 1 position on the challenge leaderboard. Meanwhile, by ensembling several variants of the core model, the overall recommendation accuracy can be improved even further.

## REFERENCES

- [1] Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, and Ed H Chi. 2018. Latent Cross: Making Use of Context in Recurrent Recommender Systems. In *WSDM*. 46–54.
- [2] Brian Brost, Rishabh Mehrotra, and Tristan Jehan. 2019. The Music Streaming Sessions Dataset. In *Proceedings of the 2019 Web Conference*. ACM.
- [3] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishii Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, and others. 2016. Wide & deep learning for recommender systems. In *DLRS*. 7–10.
- [4] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).

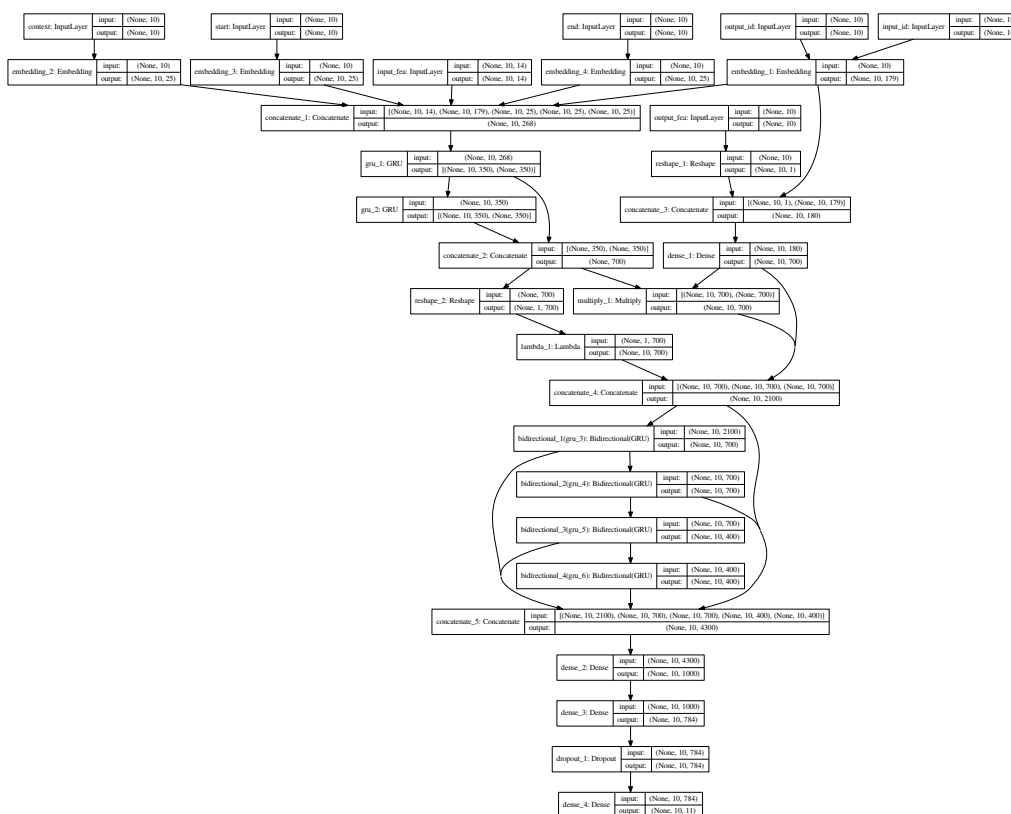


Figure 2: Model diagram.

- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2015. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289* (2015).
- [6] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *RecSys*. 191–198.
- [7] Tim Donkers, Benedikt Loepp, and Jürgen Ziegler. 2017. Sequential user-based recurrent neural network recommendations. In *RecSys*. 152–160.
- [8] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. Deep sparse rectifier neural networks. In *AISTATS*. 315–323.
- [9] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *I-CASSP*. 6645–6649.
- [10] Mihajlo Grbovic and Haibin Cheng. 2018. Real-time Personalization using Embeddings for Search Ranking at Airbnb. In *SIGKDD*. 311–320.
- [11] Mihajlo Grbovic, Vladan Radosavljevic, Nemanja Djuric, Narayan Bhamidipati, Jaikit Savla, Varun Bhagwan, and Doug Sharp. 2015. E-commerce in your inbox: Product recommendations at scale. In *SIGKDD*. 1809–1818.
- [12] David Ha and Douglas Eck. 2017. A neural representation of sketch drawings. In *ICLR*.
- [13] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. 173–182.
- [14] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In *ICLR*.
- [15] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*. 448–456.
- [16] Yuchin Juan, Damien Lefortier, and Olivier Chapelle. 2017. Field-aware factorization machines in a real-world online advertising system. In *WWW*. 680–688.
- [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 8 (2009), 30–37.
- [19] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. 2012. Efficient backprop. In *Neural networks: Tricks of the trade*. Springer, 9–48.
- [20] Jarana Manotumruksa, Craig Macdonald, and Iadh Ounis. 2018. A Contextual Attention Recurrent Architecture for Context-Aware Venue Recommendation. In *SIGIR*. 555–564.
- [21] H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, and others. 2013. Ad click prediction: a view from the trenches. In *SIGKDD*. 1222–1230.
- [22] Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *INTERSPEECH*.
- [23] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*. 1532–1543.
- [24] Massimo Quadrana, Alexandros Karatzoglou, Balázs Hidasi, and Paolo Cremonesi. 2017. Personalizing Session-based Recommendations with Hierarchical Recurrent Neural Networks. *arXiv preprint arXiv:1706.04148* (2017).
- [25] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*. 3104–3112.
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*. 5998–6008.