MEANTIME: Mixture of Attention Mechanisms with Multi-temporal Embeddings for Sequential Recommendation

Sung Min Cho
Seoul National University
tjdals4565@gmail.com

Eunhyeok Park
POSTECH
canusglow@gmail.com

Sungjoo Yoo
Neural Processing Research Center
Seoul National University
sungjoo.yoo@gmail.com

ABSTRACT

Recently, self-attention based models have achieved state-of-the-art performance in sequential recommendation task. Following the custom from language processing, most of these models rely on a simple positional embedding to exploit the sequential nature of the user's history. However, there are some limitations regarding the current approaches. First, sequential recommendation is different from language processing in that timestamp information is available. Previous models have not made good use of it to extract additional contextual information. Second, using a simple embedding scheme can lead to information bottleneck since the same embedding has to represent all possible contextual biases. Third, since previous models use the same positional embedding in each attention head, they can wastefully learn overlapping patterns. To address these limitations, we propose MEANTIME (Mixture of Attention mechanisms with Multi-temporal Embeddings) which employs multiple types of temporal embeddings designed to capture various patterns from the user's behavior sequence, and an attention structure that fully leverages such diversity. Experiments on real-world data show that our proposed method outperforms current state-of-the-art sequential recommendation methods, and we provide an extensive ablation study to analyze how the model gains from the diverse positional information.

CCS CONCEPTS

• Information systems → Retrieval models and ranking.

KEYWORDS

Sequential Recommendation, Self-attention, Temporal Embedding, BERT

ACM Reference Format:

Sung Min Cho, Eunhyeok Park, and Sungjoo Yoo. 2020. MEANTIME: Mixture of Attention Mechanisms with Multi-temporal Embeddings for Sequential Recommendation. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, September 22–26, 2020, Virtual Event, Brazil. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3383313.3412216>

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.

RecSys '20, September 22–26, 2020, Virtual Event, Brazil

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7583-2/20/09...\$15.00

<https://doi.org/10.1145/3383313.3412216>

1 INTRODUCTION

Capturing users' preferences from their history is essential for making effective recommendations, because users' preferences and items' characteristics are both dynamic. Furthermore, users' preferences depend heavily on the context (e.g., one is likely to be interested in buying a keyboard after purchasing a desktop). Therefore, sequential recommendation aims to predict the next set of items that users are likely to prefer by exploiting their history.

Many algorithms have been proposed to better understand the sequential history of users [5, 8, 10, 18–21, 36]. Despite their excellent performance, most of them ignore the interactions' timestamp values. While recent works such as TiSASRec [14] successfully incorporated time information, their usage of time was also limited to a single embedding scheme. Since the timestamp values are rife with information, it would be beneficial to explore many forms of temporal embeddings that can fully extract diverse patterns, and model architectures that can fully leverage such diversity.

In this paper we propose MEANTIME (Mixture of Attention mechanisms with Multi-temporal Embeddings) which introduces multiple temporal embeddings to better encode both absolute and relative positions of a user-item interaction in a sequence. Moreover, we employ multiple self-attention heads that handle each positional embedding separately. This is profitable because each head can play a role of expert that focuses on a particular pattern of user behaviors. Extensive experiments show that MEANTIME outperforms state-of-the-art algorithms on real-world datasets. Our contributions are as follows:

- We propose to diversify the embedding schemes that are used to encode the positions of user-item interactions in a sequence. This is done by applying multiple kernels to positions/values of timestamps in a sequence to create unique embedding matrices.
- We also propose a novel model architecture that operates multiple self-attention heads simultaneously, where each head specializes in extracting certain patterns from the users' behaviors by utilizing one of the positional embeddings we propose above.
- We conducted thorough experiments to show the effectiveness of our method on real-world datasets. We also provide a comprehensive ablation study that helps understand the effect of various components in our model.

2 RELATED WORK

2.1 Temporal Recommendation

Since temporal information hold contextual information, they can be very crucial for recommendation performance. TimeSVD++ and

Figure 1: Model Architecture

BPTF [12, 27] adopted time factor into the matrix factorization method, and TimeSVD++ was one of the main contributions for the winning of Netflix Grand Prize [11]. Time-LSTM [37] equipped LSTMs with several forms of time gates to better model the time intervals in user’s interaction sequence. Recently, CTA [25] used multiple parametrized kernel functions on temporal information to calibrate the self-attention mechanism.

2.2 Sequential Recommendation

Sequential recommendation aims to suggest relevant items based on the user’s sequential history. Markov-chain based methods [4–7] assume that users’ behaviors are affected only by their last few behaviors. FPMC [18] merges Markov-chains with matrix factorization method for next-basket recommendation. Since the first suggestion by GRU4Rec [9], many RNN-based methods [8, 17, 23, 24, 31, 34, 35] brought the success of RNN into item sequence understanding. To overcome the strong order constraint of RNN models, CNN-based methods [21, 30, 33] were proposed. Some works adopted graph neural network (GNN) to understand user’s session as a graph [26].

As for attention mechanisms, NARM [13] and STAMP [15] incorporated vanilla attention mechanism with RNN. More recently, SASRec [10] successfully employed self-attention mechanism, which was a huge success in NLP areas [1, 2, 22, 29]. BERT4Rec [19] improved SASRec by adopting Transformer [22] and Cloze-task based training method. TiSASRec [14] enhanced SASRec by merging timestamp information into self-attention operations. Our work seeks to further expand this trend by suggesting a novel model architecture that uses multiple types of temporal embeddings and self-attention operations to better capture the diverse patterns in user’s behaviors.

3 PROPOSED METHOD

Several self-attention based recommendation models have been proposed recently [10, 19]. However, the positional factors have been rather overlooked in these models for many reasons. First, they don’t utilize timestamp values which hold important contextual information. While TiSASRec [14] successfully addressed this issue,

they also used a simple embedding scheme for temporal values. This can possibly lead to an information bottleneck since the same embedding has to represent all possible positional biases. Furthermore, since all attention heads use the same positional information, they might wastefully learn the same patterns. As a way to mitigate these problems, we devise a novel model architecture that injects various temporal embeddings for each attention head as shown in Figure 1.

3.1 Problem Formulation

Let U be a set of users, and V a set of items. For each user $u \in U$, we have a sequence of items $V^u = [v_1^u, \dots, v_k^u, \dots, v_{|V^u|}^u | v_k^u \in V]$ that the user previously interacted with in chronological order and the corresponding time sequence of the interaction $T^u = [t_1^u, \dots, t_k^u, \dots, t_{|V^u|}^u | t_k^u \in \mathbb{N}]$ that stores the absolute timestamp values. Our goal is to predict the next item v_{next}^u that the user u is likely to interact with at the target timestamp t_{next}^u based on the given history (V^u, T^u) .

Our model is based on a fixed-length model with length N as in previous works [10, 19]. Given the history (V_u, T_u) of user u , we feed $\mathbf{v} = [v_{|V_u|-N+2}^u, \dots, v_{|V_u|}^u, [\text{MASK}]]$ and $\mathbf{t} = [t_{|V_u|-N+2}^u, \dots, t_{|V_u|}^u, t_{next}^u]$ to the model. Then the model estimates v_{next} as output. If the history is shorter than $N - 1$, it is padded by a special token [PAD].

3.2 Input Embedding

The embedding layers convert $\mathbf{v}$ and $\mathbf{t}$ to hidden features that are provided to the attention module. In the case of item array $\mathbf{v}$, each element has an item index. This array is embedded as $E^{item} \in \mathbb{R}^{N \times h}$ using the conventional look-up operation with a learnable item embedding table $M^I \in \mathbb{R}^{(|V|+1) \times h}$, where h is the hidden feature size. M^I holds embeddings for $|V|$ items and [MASK] token.

In the case of $\mathbf{t}$, each element has a timestamp value. In previous studies, these values were embedded in a single embedding scheme, if they were used at all. However, given that user’s behaviors are

Figure 2: Brief overview of single attention head in multi-headed attention structure. (a) absolute self-attention in Transformer (b) relative self-attention in Transformer-XL and MEANTIME (c) absolute self-attention in MEANTIME.

mixed with various patterns, a single embedding might not suffice. Moreover, using different encoding functions (instead of plain learnable values) for each embedding is beneficial since we can train each of them to learn a unique pattern. With these in mind, we propose six methods to embed $\mathbf{t}$: three absolute methods (Figure 1(a)), and three relative methods (Figure 1(b)).

Absolute embeddings encode each position in the sequence as E^{Day}, E^{Pos} , or $E^{Con} \in \mathbb{R}^{N \times h}$. The first is **Day**-embedding which converts the day of each timestamp into an embedding vector to get E^{Day} . We keep a learnable embedding matrix $M^D \in \mathbb{R}^{|D| \times h}$ where $|D|$ is the number of possible days in the span of the dataset. In our analysis, **Day** is a reasonable choice of time length considering the trade-off between matrix size and the effectiveness. **Pos**-embedding is analogous to the learnable positional embedding used in [2, 19], where each position has a corresponding embedding in $M^P \in \mathbb{R}^{N \times h}$. We convert each position to get E^{Pos} . Note that in this case, the embeddings depend only on the indices of time array and not on the timestamp values. **Con**-embedding is similar to **Pos**-embedding, except that all position vectors are shared ($M^C \in \mathbb{R}^{1 \times h}$) to remove positional bias. Hence, E^{Con} is a stack of repeated vectors.

Relative embeddings encode the relationship between each interaction pair in the sequence as E^{Sin}, E^{Exp} , or $E^{Log} \in \mathbb{R}^{N \times N \times h}$ by utilizing the temporal difference information. First, we define a matrix of temporal differences $\mathbf{D} \in \mathbb{R}^{N \times N}$ whose element is defined as $d_{ab} = (\mathbf{t}_a - \mathbf{t}_b)/\tau$, where τ is an adjustable unit time difference. We introduce three encoding functions for $\mathbf{D}$. First, **Sin** encoder converts the difference d_{ab} to a hidden vector $\vec{\theta}_{ab} \in \mathbb{R}^{1 \times h}$, by the following equation:

$$\vec{\theta}_{ab,2c} = \sin\left(\frac{d_{ab}}{freq^{\frac{2c}{h}}}\right) \quad \vec{\theta}_{ab,2c+1} = \cos\left(\frac{d_{ab}}{freq^{\frac{2c}{h}}}\right) \quad (1)$$

where $\vec{\theta}_{ab,c}$ is the c^{th} value of the vector $\vec{\theta}_{ab}$ and $freq$ is an adjustable parameter. Likewise, **Exp** encoder and **Log** encoder converts d_{ab} to $\vec{e}_{ab}$ and $\vec{l}_{ab}$ respectively by applying the following equations:

$$\vec{e}_{ab,c} = \exp\left(\frac{-|d_{ab}|}{freq^{\frac{c}{h}}}\right) \quad \vec{l}_{ab,c} = \log\left(1 + \frac{|d_{ab}|}{freq^{\frac{c}{h}}}\right) \quad (2)$$

Stacking these vectors gives us E^{Sin}, E^{Exp} and E^{Log} respectively.

Each embedding offers a distinctive view on temporal data. For example, **Sin** captures periodic occurrences. Larger time gap is either quickly decayed to zero in **Exp** or manageably increased

in **Log**. **Day** can confine the attention scope within the same (or similar) day. **Pos** can pick up the patterns learnt by the previous models. Finally, **Con** removes positional bias altogether so that the attention can focus only on the relationship between items. These distinctive views enable our self-attention heads (Figure 1(c)) to attend to different characteristics of the sequence.

3.3 Self-Attention Structure

3.3.1 Attention Architecture. Multi-head self-attention proposed in Transformer [22] shows excellent performance in recommendation tasks [19]. MEANTIME is also based on this structure. Traditionally, self-attention based models [2, 10, 19] employ a simple absolute positional encoding by feeding the sum of E^{Item} and E^{Pos} as input as shown in Figure 2(a). Advanced transformers [1, 29] enhanced this as shown in Figure 2(b). E^{Pos} is replaced by the relative positional embedding E^R that only depends on the difference between the positional indices of the array, and positional information is no longer used at key and value encoder. Parameters b_I and b_R serve as global content/positional bias respectively. For absolute embedding, we also adopt the previous structure. In addition, we propose separating the item and positional embedding for absolute attention as shown in Figure 2(c).

As explained in section 3.2, we can diversify the positional information that is fed to the attention module. MEANTIME utilizes 3 absolute embeddings and 3 relative embeddings. When fed a relative embedding, the head will operate as in Figure 2(b) with either E^{Sin}, E^{Exp} or E^{Log} as input to K_R . When fed an absolute embedding, the head will operate as in Figure 2(c) with either E^{Day}, E^{Pos} or E^{Con} as input to Q_A and K_A . As in previous works, the dimension of each head is $\frac{h}{n}$. The choice of n and embedding types are adjustable with regards to the characteristics of the dataset.

3.3.2 Stacking Layers. After self-attention, MEANTIME operates similar to [2, 19]. We apply Position-wise Feed Forward Network (FFN) to the result of self-attention to complete a single layer:

$$\text{FFN}(x) = \text{GELU}(xW^1 + b^1)W^2 + b^2 \quad (3)$$

where $W^1 \in \mathbb{R}^{h \times 4h}$, $b^1 \in \mathbb{R}^{4h}$, $W^2 \in \mathbb{R}^{4h \times h}$ and $b^2 \in \mathbb{R}^h$ are learnable parameters. Then we stack L such layers. As in previous works [10, 19, 22], we apply a residual connection for each sublayer

Table 1: Performance Comparison.

Datasets	Metrics	MARANK	SAS	TISAS	BERT	MEANTIME	Improvement
ML-1M	Recall@5	0.4427	0.5708	<u>0.5869</u>	0.5862	0.6415	9.31%
	Recall@10	0.5993	0.6966	<u>0.7102</u>	0.6987	0.7412	4.36%
	NDCG@5	0.3042	0.4142	0.4307	<u>0.4420</u>	0.4932	11.58%
	NDCG@10	0.3548	0.4550	0.4708	<u>0.4786</u>	0.5255	9.81%
ML-20M	Recall@5	0.3938	0.5274	0.5412	<u>0.5910</u>	0.6225	5.34%
	Recall@10	0.5503	0.6887	0.7021	<u>0.7176</u>	0.7491	4.39%
	NDCG@5	0.2665	0.3700	0.3814	<u>0.4469</u>	0.4739	6.04%
	NDCG@10	0.3170	0.4223	0.4336	<u>0.4879</u>	0.5150	5.55%
Beauty	Recall@5	0.1577	0.2011	0.2081	<u>0.2271</u>	0.2470	8.78%
	Recall@10	0.2302	0.2714	0.2805	<u>0.3095</u>	0.3330	7.58%
	NDCG@5	0.1085	0.1464	0.1512	<u>0.1633</u>	0.1764	8.05%
	NDCG@10	0.1318	0.1689	0.1745	<u>0.1898</u>	0.2041	7.52%
Game	Recall@5	0.2846	0.3925	0.3857	<u>0.4211</u>	0.4752	12.85%
	Recall@10	0.4217	0.5201	0.5085	<u>0.5590</u>	0.6100	9.12%
	NDCG@5	0.1893	0.2735	0.2735	<u>0.2970</u>	0.3446	16.03%
	NDCG@10	0.2334	0.3147	0.3133	<u>0.3415</u>	0.3882	13.67%

to facilitate training:

$$\begin{aligned} y &= x + \text{Dropout}(\text{Attention}(\text{LayerNorm}(x))) \\ z &= y + \text{Dropout}(\text{FFN}(\text{LayerNorm}(y))) \end{aligned} \quad (4)$$

Note that while some works [19, 22] apply LayerNorm at the very end, we empirically chose to apply LayerNorm in the front. This is also in accordance with the recent report on the order of layer normalization in Transformers [28].

3.3.3 Prediction Layer. Given the outputs $[o^1, \dots, o^N] \in \mathbb{R}^{N \times h}$ from the last layer, we obtain the item score distribution for each position by:

$$P(V|\mathbf{v}, \mathbf{t}) = \text{softmax}(\text{GELU}(o^i W^P + b^P) E^T + b^O) \quad (5)$$

where $W^P \in \mathbb{R}^{h \times h}$, $b^P \in \mathbb{R}^h$, and $b^O \in \mathbb{R}^{|V|}$ are learnable parameters and $E^T \in \mathbb{R}^{h \times |V|}$ is an item embedding table. Note that we use shared item embeddings for both input and output prediction. $P(v_i = v|\mathbf{v}, \mathbf{t})$ represents the probability that an item at position i is item v .

3.4 Model Training

To train our model, we adopt the existing technique [19]. We sample a subarray of items $\mathbf{v} = [v_k^u, \dots, v_{k+N-1}^u]$ and timestamps $\mathbf{t} = [t_k^u, \dots, t_{k+N-1}^u]$ from (V^u, T^u) , and convert $\mathbf{v}$ to $\mathbf{v}'$ by randomly masking a portion of it with [MASK] by some probability ρ . Then we feed $\mathbf{v}'$ and $\mathbf{t}$ to our model to get $P(V|\mathbf{v}', \mathbf{t})$. Next, we calculate the loss as follows:

$$L = - \sum_{\mathbf{v}'_i \text{ is masked}} \log P(v_i = \mathbf{v}_i | \mathbf{v}', \mathbf{t}) \quad (6)$$

4 EXPERIMENTS

4.1 Setup

4.1.1 Datasets. We evaluate our model on four real-word datasets with various domains and sparsity: MovieLens 1M [3], MovieLens 20M [3], Amazon Beauty [16] and Amazon Game [16]. We follow the common data preprocessing procedure from [5, 10, 14, 19]. We

convert each dataset into an implicit dataset by treating each numerical rating or review as a presence of a user-item interaction. We group the interactions by user ids and then sort them by timestamps to form a sequence for each user. Following the custom practice [5, 10, 14, 19], we discard users and items with less than five interactions to ensure the quality of the dataset.

4.1.2 Evaluation. To evaluate the performance of each model, we adopt the widely used *leave-one-out* evaluation task. For each user's item sequence, we hold out the last item for test, and the second last item for validation. We use the rest for training. We follow the common practice [10, 14, 19] of letting the model rank the ground truth item together with 100 randomly sampled negative items which haven't yet been interacted by the user. As in [19], we sample the negative items according to their popularity to make the task more realistic. To score the ranked list, we use two common metrics: *Recall*, and *Normalized Discounted Cumulative Gain* (NDCG). Both Recall@K and NDCG@K are designed to have larger values when the ground truth item is ranked higher up in the top-k list. We report with $K = 5$ and 10.

4.2 Baselines Models & Implementation Details

In order to validate the effectiveness of our method, we compare it with the state-of-the-art baselines: MARank [32], SASRec [10], TiSASRec [14], and BERT4Rec [19]. These models are evaluated based on the code provided by the authors. We implemented MEANTIME¹ with PyTorch. All models are fully tuned through an extensive grid-search on hyperparameters: hidden dimension $h \in \{16, 32, 64, 128\}$, number of heads $n \in \{2, 4\}$, dropout ratio $\in \{0.2, 0.5\}$, and weight decay $wd \in \{0, 0.00001\}$. We used the same maximum sequence length ($N = 200$ for MovieLens, $N = 50$ for Amazon) and the number of layers ($L = 2$) for all transformer-based models to ensure fairness. We train all models until their validation accuracy doesn't improve for 20 epochs, and then report with the test set. We report the optimal result for each model. As for embedding types in MEANTIME, we used the best combination of embedding types among

¹Our code is available at <https://github.com/SungMinCho/MEANTIME>

Figure 3: NDCG@10 Comparison with All Possible Combinations of Positional Embeddings on ML-1M (left) and Game (right) Dataset. C, D, P, S, E and L represent Con, Day, Pos, Sin, Exp and Log, respectively. We used $h = 60$ for all experiments.

all possible combinations with $n = 2, 4$: **Day+Pos+Sin+Log** for MovieLens 1M and 20M, **Day+Sin+Exp+Log** for Amazon Beauty, and **Day+Pos+Sin+Exp** for Amazon Game.

4.3 Results

Table 1 shows the performance of all models on four datasets. It also shows the relative improvements of our model compared to the strongest baselines. Our model gives the best performance in all metrics and datasets. On average, MEANTIME achieves **9.07%**, **6.36%**, **10.43%**, and **9.14%** improvements over the strongest baselines in Recall@5, Recall@10, NDCG@5 and NDCG@10, respectively. Given that we tune all models under the same set of hyperparameters (including the number of attention heads), we can deduce that the performance boost comes from the novel multi-temporal embeddings and attention operations that we proposed. Especially, although TiSASRec uses additional temporal information, it does not show impressive improvement over SASRec and BERT4Rec. We argue that using a single embedding scheme throughout all attention heads was a hindrance to their potential improvement. In contrast, each attention head takes different temporal embedding to extract specialized patterns in our design. This joint design of diverse temporal embeddings and attention mechanism was the key to our improvement.

4.4 Effect of Multi-temporal Embeddings

Figure 3 shows the comparison of all possible combinations of temporal embeddings on ML-1M and Amazon Game dataset. We fixed all other hyperparameters (e.g., $h = 60$). Interestingly, the importance of each embedding seems to differ in each dataset. In ML-1M, relative position embeddings (especially **Log**) seem to play an important role, whereas in Amazon Game, absolute position embeddings (especially **Day**) seem to be essential. Therefore, the optimal combination is also different for each dataset. Further analysis with different h shows us that while optimal combination differs slightly for each h , the overall tendency stays the same for each dataset. The same applies to other datasets. This suggests that datasets with distinct characteristics require carefully chosen temporal embeddings that provide the right views on the context of user-item interactions. Note that while we only used $n = 2$ and 4 for Table 1 to be fair with other baselines, using all possible n can further improve our performance.

5 CONCLUSION

In this paper we presented MEANTIME, which uses time information and multiple encoding functions to provide distinctive positional embeddings to self-attention modules. Experiments on four real-world datasets show that our model outperforms state-of-the-art baselines in sequential recommendation. Through extensive ablation study, we also showed that different datasets require different positional embeddings to gain optimal performance. This suggests that one must carefully tune the positional factor in their model considering the characteristics of their own dataset.

ACKNOWLEDGMENTS

This work was supported by Samsung Electronics and the National Research Foundation of Korea (NRF) grant (PF Class Heterogeneous High Performance Computer Development, NRF-2016M3C4A7952587) funded by the Ministry of Science, ICT & Future Planning (No. 2013R1A3A2003664).

REFERENCES

- [1] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. *Association for Computational Linguistics (ACL)* (2019).
- [2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)* (2019).
- [3] F. Maxwell Harper and Joseph A. Konstan. 2016. The MovieLens Datasets: History and Context. *ACM Transactions on Interactive Intelligent Systems (tiis)* (2016).
- [4] Ruining He, Chen Fang, Zhaowen Wang, and Julian J. McAuley. 2016. Vista: A Visually, Socially, and Temporally-aware Model for Artistic Recommendation. *Recommender Systems (RecSys)* (2016).
- [5] Ruining He, Wang-Cheng Kang, and Julian J. McAuley. 2017. Translation-based Recommendation. *Recommender Systems (RecSys)* (2017).
- [6] Ruining He and Julian McAuley. 2016. Fusing similarity models with markov chains for sparse sequential recommendation. *International Conference on Data Mining (ICDM)* (2016).
- [7] Ruining He and Julian McAuley. 2016. Fusing similarity models with markov chains for sparse sequential recommendation. *International Conference on Data Mining (ICDM)* (2016).
- [8] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent neural networks with top-k gains for session-based recommendations. *International Conference on Information and Knowledge Management (CIKM)* (2018).
- [9] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. *International Conference on Learning Representations (ICLR)* (2016).
- [10] W. Kang and J. McAuley. 2018. Self-Attentive Sequential Recommendation. *International Conference on Data Mining (ICDM)* (2018).
- [11] Yehuda Koren. 2009. The bellkor solution to the netflix grand prize. *Netflix prize documentation* (2009).
- [12] Yehuda Koren. 2009. Collaborative filtering with temporal dynamics. *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2009).

- [13] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural attentive session-based recommendation. *International Conference on Information and Knowledge Management (CIKM)* (2017).
- [14] Jiacheng Li, Yujie Wang, and Julian McAuley. 2020. Time Interval Aware Self-Attention for Sequential Recommendation. *Web Search and Data Mining (WSDM)* (2020).
- [15] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2018).
- [16] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. *ACM SIGIR Conference on Research and Development in Information Retrieval* (2015).
- [17] Massimo Quadroni, Alexandros Karatzoglou, Balázs Hidasi, and Paolo Cremonesi. 2017. Personalizing session-based recommendations with hierarchical recurrent neural networks. *Recommender Systems (RecSys)* (2017).
- [18] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized markov chains for next-basket recommendation. *International Conference on World Wide Web (WWW)* (2010).
- [19] 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. *International Conference on Information and Knowledge Management (CIKM)* (2019).
- [20] Jiaxi Tang, Francois Belletti, Sagar Jain, Minmin Chen, Alex Beutel, Can Xu, and Ed H. Chi. 2019. Towards neural mixture recommender for long range dependent user sequences. *The World Wide Web Conference* (2019).
- [21] Jiaxi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. *Web Search and Data Mining (WSDM)* (2018).
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* (2017).
- [23] Bo Wu, Wen-Huang Cheng, Yongdong Zhang, Qiushi Huang, Jintao Li, and Tao Mei. 2017. Sequential Prediction of Social Media Popularity with Deep Temporal Context Networks. *International Joint Conference on Artificial Intelligence (IJCAI)* (2017).
- [24] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J Smola, and How Jing. 2017. Recurrent recommender networks. *Web Search and Data Mining (WSDM)* (2017).
- [25] Jibang Wu, Renqin Cai, and Hongning Wang. 2020. Déjà vu: A Contextualized Temporal Attention Mechanism for Sequential Recommendation. *Proceedings of The Web Conference* (2020).
- [26] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. *AAAI Conference on Artificial Intelligence* (2019).
- [27] Liang Xiong, Xi Chen, Tzu-Kuo Huang, Jeff Schneider, and Jaime G Carbonell. 2010. Temporal collaborative filtering with bayesian probabilistic tensor factorization. *International Conference on Data Mining (ICDM)* (2010).
- [28] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. On layer normalization in the transformer architecture. *arXiv preprint arXiv:2002.04745* (2020).
- [29] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in Neural Information Processing Systems (NeurIPS)* (2019).
- [30] Jiaxuan You, Yichen Wang, Aditya Pal, Pong Eksombatchai, Chuck Rosenberg, and Jure Leskovec. 2019. Hierarchical temporal convolutional networks for dynamic recommender systems. *International Conference on World Wide Web (WWW)* (2019).
- [31] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. A dynamic recurrent model for next basket recommendation. *ACM SIGIR conference on Research and Development in Information Retrieval* (2016).
- [32] Lu Yu, Chuxu Zhang, Shangsong Liang, and Xiangliang Zhang. 2019. Multi-order attentive ranking model for sequential recommendation. *AAAI Conference on Artificial Intelligence* (2019).
- [33] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M Jose, and Xiangnan He. 2019. A simple convolutional generative network for next item recommendation. *Web Search and Data Mining (WSDM)* (2019).
- [34] Yuyu Zhang, Hanjun Dai, Chang Xu, Jun Feng, Taifeng Wang, Jiang Bian, Bin Wang, and Tie-Yan Liu. 2014. Sequential click prediction for sponsored search with recurrent neural networks. *AAAI Conference on Artificial Intelligence (AAAI)* (2014).
- [35] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. *AAAI Conference on Artificial Intelligence (AAAI)* (2019).
- [36] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2018).
- [37] Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. 2017. What to Do Next: Modeling User Behaviors by Time-LSTM. *International Joint Conference on Artificial Intelligence (IJCAI)* (2017).