TAGNN: Target Attentive Graph Neural Networks for Session-based Recommendation

Feng Yu^{1,2,*}, Yanqiao Zhu^{1,2,*}, Qiang Liu^{3,4}, Shu Wu^{1,2,†}, Liang Wang^{1,2}, and Tieniu Tan^{1,2}

¹Center for Research on Intelligent Perception and Computing, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³RealAI ⁴Tsinghua University

{feng.yu,shu.wu,wangliang,tnt}@nlpr.ia.ac.cn,yanqiao.zhu@cripac.ia.ac.cn,qiang.liu@realai.ai

ABSTRACT

Session-based recommendation nowadays plays a vital role in many websites, which aims to predict users' actions based on anonymous sessions. There have emerged many studies that model a session as a sequence or a graph via investigating temporal transitions of items in a session. However, these methods compress a session into one fixed representation vector without considering the target items to be predicted. The fixed vector will restrict the representation ability of the recommender model, considering the diversity of target items and users' interests. In this paper, we propose a novel target attentive graph neural network (TAGNN) model for session-based recommendation. In TAGNN, target-aware attention adaptively activates different user interests with respect to varied target items. The learned interest representation vector varies with different target items, greatly improving the expressiveness of the model. Moreover, TAGNN harnesses the power of graph neural networks to capture rich item transitions in sessions. Comprehensive experiments conducted on real-world datasets demonstrate its superiority over state-of-the-art methods.

CCS CONCEPTS

• Information systems → Recommender systems; • Computing methodologies → Machine learning; Neural networks; • Applied computing → Online shopping.

ACM Reference Format:

Feng Yu, Yanqiao Zhu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2020. TAGNN: Target Attentive Graph Neural Networks for Session-based Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20), July 25–30, 2020, Virtual Event, China.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3397271.3401319

SIGIR '20, July 25-30, 2020, Virtual Event, China

1 INTRODUCTION

Recommender systems are one of the most successful applications in the fields of data mining and machine learning research. Most previous work studies approaches that personalize the recommendation according to constantly recorded user profiles. In many real-world applications, however, such long-term profiles, even the users' identities may not exist. As an emerging method aiming to solve this problem, session-based recommendation predicts the next action (e.g., which item to click) of a user, given her previous behaviors within the ongoing session.

Considering the high practical values of session-based recommendation, many approaches have been proposed so far. Traditionally, Markov-chain-based methods [8] predict the user's next action solely based on the previous action. Such a strong independence assumption suffers from noisy data and thus restricts its use in session-based recommendation scenarios. Recent trends in recommender systems have led to a proliferation of studies using deep neural network techniques. Models based on recurrent neural networks (RNNs) have achieved promising performance. For example, Hidasi et al. propose a RNN-based method GRU4Rec [1] to model short-term preferences with gated recurrent units (GRUs). Recently, NARM [4] proposes two RNN-based subsystems to capture users' local and global preference respectively. Similar to NARM, STAMP [6] extracts users' potential interests using a simple multilayer perception model and an attentive network.

Despite their effectiveness, we would argue that those methods are still in their infancy. Previous work highlights that complex user behavioral patterns are of great significance for session-based recommendation [4, 6]. However, these sequence-based methods only model sequential transitions between consecutive items, with complex transitions neglected. Take repeated purchases as an example, which is one of the most prominent behaviors in e-shopping scenarios. Suppose a session for a user is $s = v_1 \rightarrow v_2 \rightarrow v_1 \rightarrow v_3$. Then, it is hard for these sequence-based methods to capture such a to-and-fro relationship between items. Specifically, they will be confused about the relationship between item v_1 and items (v_2, v_3) . In this paper, we propose to discover the complex transitional patterns underneath sessions through session graphs [2, 11, 12]. By modeling items in sessions as session graphs, this natural means of encoding the abundant temporal patterns within sessions produces more accurate representation for each item.

Moreover, candidate items are usually abundant and users' interests are usually diverse. Previous work [4, 6, 11] represents one session using one embedding vector. That fixed-size vector represents all interests of a single user, which cannot express diverse user

^{*}The first two authors contributed equally to this work.

[†]To whom correspondence should be addressed.

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 the author(s) 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.

^{© 2020} Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-8016-4/20/07...\$15.00 https://doi.org/10.1145/3397271.3401319

interests, and thereby limits the expressiveness of the recommender model. As a brute-force solution, we may consider enlarging the dimension of that fixed-length vector, which will in turn expose the risk of overfitting and deteriorate the model performance. However, we observe that it is not necessary to embed all user interests into one vector when making prediction for a specific candidate item. For example, suppose that a customer has a historical session of (swimming suits, purse, milk, frying pan). If we want to recommend a handbag for her, we focus on her interests in the purse rather than the frying pan. That is to say, the interests of a user with rich behaviors can be specifically activated given a target item [13]. In this paper, we refine the proposed graph-based model through a novel target attention module. We term the resulting model as Target Attentive Graph Neural Networks for session-based recommendation, TAGNN¹ for brevity. The proposed target attention module aims to adaptively activate user interests by considering the relevance of historical behaviors given a target item. By introducing a local target attentive unit, specific user interests related to a target item in the current session are activated, which will benefit downstream session representations as a result.

Figure 1 gives an overview of the TAGNN method. We first construct session graphs using items in historical sessions. After that, we obtain corresponding embeddings using graph neural networks to capture complex item transitions based on session graphs. Given the item embeddings, we employ a target-aware attentive network to activate specific user interests with respect to a target item. Following that, we construct session embeddings. At last, for each session, we can infer the user's next action based on item embeddings and the session embedding.

The main contribution of this work is threefold. Firstly, we model items in sessions as session graphs to capture complex item transitions within sessions. Then, we employ graph neural networks to obtain item embeddings. Secondly, to adaptively activate users' diverse interests in sessions, we propose a novel target attentive network. The proposed target attentive module can reveal the relevance of historical actions given a certain target item, which further improves session representations. Finally, we conduct extensive experiments on real-world datasets. The empirical studies show that our method achieves state-of-the-art performance.

2 THE PROPOSED METHOD: TAGNN

2.1 Preliminaries

In session-based recommendation, an anonymous session can be represented by a list $s = [v_{s,i}]_{i=1}^{s_n}$ ordered by timestamps and we denote $V = \{v_i\}_{i=1}^m$ as the set consisting of all unique items (e.g., user clicks) involved in sessions. Session-based recommendation aims to predict the next action v_{s,s_n+1} given session *s*. Our model produces a ranking list of probabilities for all candidate items and items with top-*k* probability values will be selected for recommendation.

In our model, we represent each session *s* as a directed *session* graph $\mathcal{G}_s = (\mathcal{V}_s, \mathcal{E}_s, \mathcal{A}_s)$, where $\mathcal{V}_s, \mathcal{E}_s, \mathcal{A}_s$ are the node set, the edge set, and the adjacency matrix, respectively. In this graph \mathcal{G}_s , each node represents an item $v_{s,i} \in V$ and each edge $(v_{s,i-1}, v_{s,i}) \in \mathcal{E}_s$ represents a user visits item $v_{s,i-1}$ and $v_{s,i}$ consecutively. Here



Figure 1: An overview of the proposed TAGNN method. Session graphs are constructed based on sessions at first. Then, graph neural networks capture rich item transitions in sessions. Last, from one session embedding vector, target-aware attention adaptively activates different user interests concerning varied target items to be predicted.



Figure 2: An example of a session graph \mathcal{G}_s and the corresponding adjacency matrix $A_s = [A_s^{(in)}; A_s^{(out)}]$.

we define A_s as the concatenation of two adjacency matrices $A_s^{(out)}$ and $A_s^{(in)}$ to reflect the bidirectional relationship between items in sessions, where $A_s^{(out)}$ and $A_s^{(in)}$ represents weighted connections of outgoing and incoming edges respectively. Considering items may appear in the session repeatedly, we conduct row normalization for the two adjacency matrices. For example, consider a session $s = [v_1, v_2, v_3, v_1, v_3, v_4]$, we construct a session graph \mathcal{G}_s as shown in Figure 2. Here we use the same strategy of constructing session graphs as SR-GNN [11]. However, it is flexible to adopt different mechanisms of constructing the session graph for different sessionbased recommendation scenarios. For clarity, we omit subscript *s* for referring item embeddings hereafter.

2.2 Learning Item Embeddings

After constructed session graphs, we transform every node $v_i \in V$ into a unified embedding space. The resulting vector $v_i \in \mathbb{R}^d$ is a *d*-dimensional representation of item v_i obtained using graph neural networks. Then, we can represent each session *s* using item embeddings. The graph neural network (GNN) [3, 5, 10] is a class of widely used deep learning models. GNNs generate node representations

¹Code available at https://github.com/CRIPAC-DIG/TAGNN

on top of graph topology, which models complex item connections. Therefore, they are particularly suitable for session-based recommendation. In this paper, we employ gated graph neural networks (GGNNs) [5], a variant of GNN, to learn the node vectors. Formally, for node $v_{s,i}$ of graph G_s , its update rules are:

$$\boldsymbol{a}_{s,i}^{(t)} = \boldsymbol{A}_{s,i} \left[\boldsymbol{v}_1^{(t-1)}, \dots, \boldsymbol{v}_{s_n}^{(t-1)} \right]^\top \boldsymbol{H} + \boldsymbol{b}, \tag{1}$$

$$\boldsymbol{z}_{s,i}^{(t)} = \sigma \left(\boldsymbol{W}_{z} \boldsymbol{a}_{s,i}^{(t)} + \boldsymbol{U}_{z} \boldsymbol{v}_{i}^{(t-1)} \right),$$
(2)

$$\mathbf{r}_{s,i}^{(t)} = \sigma \left(\mathbf{W}_r \mathbf{a}_{s,i}^{(t)} + \mathbf{U}_r \mathbf{v}_i^{(t-1)} \right), \tag{3}$$

$$\boldsymbol{v}_{i}^{(t)} = \tanh\left(\boldsymbol{W}_{o}\boldsymbol{a}_{s,i}^{(t)} + \boldsymbol{U}_{o}\left(\boldsymbol{r}_{s,i}^{(t)} \odot \boldsymbol{v}_{i}^{(t-1)}\right)\right), \tag{4}$$

$$\boldsymbol{v}_{i}^{(t)} = \left(1 - \boldsymbol{z}_{\mathrm{s},i}^{(t)}\right) \odot \boldsymbol{v}_{i}^{(t-1)} + \boldsymbol{z}_{\mathrm{s},i}^{(t)} \odot \boldsymbol{v}_{i}^{(t)}, \tag{5}$$

where *t* is the training step, $A_{s,i:} \in \mathbb{R}^{1 \times 2n}$ is the *i*-th row in matrix A_s corresponding to node $v_{s,i}$, $H \in \mathbb{R}^{d \times 2d}$ and $b \in \mathbb{R}^d$ are weight and bias parameter respectively, $\left[v_1^{(t-1)}, \ldots, v_{s_n}^{(t-1)}\right]$ is the list of node vectors in session $s, z_{s,i} \in \mathbb{R}^{d \times d}$ and $r_{s,i} \in \mathbb{R}^{d \times d}$ are the reset and update gates respectively, $\sigma(\cdot)$ is the sigmoid function, and \odot denotes element-wise multiplication. For each session graph \mathcal{G}_s , the GGNN model propagates information between neighboring nodes. The update and reset gate decides what information to be preserved and discarded respectively.

2.3 Constructing Target-Aware Embeddings

Previous work captures users' interests only using intra-session item representations. In our model, once we obtained the node vector for each item, we proceed to construct target embeddings, to adaptively consider the relevance of historical behaviors concerning target items. Here we define the target items as all candidate items to predict. Usually, a user's action given a recommended item only matches a part of her interests. To model this process, we design a novel target attention mechanism to calculate soft attention scores over all items in the session with respect to each target item.

In this section, we introduce a local target attentive module to calculate attention scores between all items v_i in session s and each target item $v_t \in V$. Firstly, a shared non-linear transformation parameterized by a weight matrix $W \in \mathbb{R}^{d \times d}$ is applied to every node-target pair. Then, we normalize the self-attention scores using the softmax function:

$$\beta_{i,t} = \operatorname{softmax}(e_{i,t}) = \frac{\exp\left(\boldsymbol{v}_t^\top \boldsymbol{W} \boldsymbol{v}_i\right)}{\sum_{j=1}^m \exp\left(\boldsymbol{v}_t^\top \boldsymbol{W} \boldsymbol{v}_j\right)}.$$
(6)

Finally, for each session *s*, the users' interests towards a target item v_t is represented by $s_{\text{target}}^t \in \mathbb{R}^d$, as given below:

$$s_{\text{target}}^{t} = \sum_{i=1}^{s_{n}} \beta_{i,t} \boldsymbol{v}_{i}.$$
 (7)

The obtained *target embedding* for representing users' interests varies with different target items.

2.4 Generating Session Embeddings

In this section, we further exploit users' short- and long-term preference exhibited in the current session *s* using node representations involved in session *s*. The resulting two representations along with the user's target embedding will be further concatenated to generate better session embeddings.

Local embedding. As the user's final action is usually determined by her last action, we simply represent the user's short-term preference as a *local embedding* $s_{\text{local}} \in \mathbb{R}^d$ as the representation of the last-visited item v_{s,s_n} .

Global embedding. Then, we represent the user's long-term preference as a *global embedding* $s_{global} \in \mathbb{R}^d$ by aggregating all involved node vectors. We adopt another soft-attention mechanism to draw dependencies between the last-visited item and each item involved in the session:

$$\alpha_i = \boldsymbol{q}^\top \, \boldsymbol{\sigma}(\boldsymbol{W}_1 \boldsymbol{v}_{s_n} + \boldsymbol{W}_2 \boldsymbol{v}_i + \boldsymbol{c}), \tag{8}$$

$$s_{\text{global}} = \sum_{i=1}^{s_n} \alpha_i v_i, \tag{9}$$

where $q, c \in \mathbb{R}^d$ and $W_1, W_2 \in \mathbb{R}^{d \times d}$ are weight parameters.

Session embedding. Finally, we generate the session embedding *s* of session *s* by taking linear transformation over the concatenation of the local and global embeddings and the target embedding:

$$\mathbf{s}_t = \mathbf{W}_3[\mathbf{s}_{\text{target}}^t; \mathbf{s}_{\text{local}}; \mathbf{s}_{\text{global}}], \tag{10}$$

where $W_3 \in \mathbb{R}^{d \times 3d}$ projects the three vectors into one embedding space \mathbb{R}^d . Please kindly note that we generate different session embeddings for each target item.

2.5 Making Recommendation

After obtaining all item embeddings and session embeddings, we compute the recommendation score $\hat{z_t}$ for each target item $v_t \in V$ by taking inner-product of item embedding v_t and session representation s. Following that, we use the softmax function over all unnormalized scores z for all target items and get the final output vector:

$$\hat{z_t} = \boldsymbol{s}_t^\top \boldsymbol{v}_t, \tag{11}$$

$$\hat{\boldsymbol{y}} = \operatorname{softmax}\left(\hat{\boldsymbol{z}}\right).$$
 (12)

Here $\hat{y} \in \mathbb{R}^m$ denotes the probabilities of nodes being the next action in *s*. The items with the top-*k* probabilities in \hat{y} will be selected as recommended items.

For training the model, we define the loss function as the crossentropy of the prediction and the ground truth:

$$\mathcal{L}(\hat{\boldsymbol{y}}) = -\sum_{i=1}^{m} \boldsymbol{y}_i \log\left(\hat{\boldsymbol{y}}_i\right) + (1 - \boldsymbol{y}_i) \log\left(1 - \hat{\boldsymbol{y}}_i\right), \quad (13)$$

where \boldsymbol{y} denotes the one-hot encoding vector of the ground truth items. We use the back-propagation through time (BPTT) algorithm to train the proposed model.

3 EXPERIMENTS

In this section, we aim to answer the following two questions:

RQ1. Does the proposed TAGNN achieve state-of-the-art performance compared with existing representative baseline algorithms?

RQ2. How do different schemes for representing user interests affect the model performance?

Table 1: Statistics of datasets used in the experiments

Statistics	Yoochoose 1/64	Diginetica
# Clicks	557,248	982,961
# Training sessions	369,859	719,470
# Test sessions	55,898	60,858
# Unique items	17,377	43,097
Average length	6.16	5.12

3.1 Experimental Configurations

Datasets. We evaluate the proposed method using two widelyused real-world datasets Yoochoose² and Diginetica³, obtained from two contests in data mining conference RecSys 2015 and CIKM 2016 respectively. For fair comparison, we closely follow the same data preprocessing scheme as Li et al. [4], Liu et al. [6], Wu et al. [11]. Specifically, we drop items appearing less than 5 times and sessions consisting of less than 2 items. For generating training and test sets, sessions of last days are used as the test set for Yoochoose, and sessions of last weeks as the test set for Diginetica. For an existing session $s = [v_{s,1}, v_{s,2}, \dots, v_{s,s_n}]$, we generate a series of input session sequences and corresponding labels as $([v_{s,1}], v_{s,2}), ([v_{s,1}, v_{s,2}], v_{s,3}), \dots, ([v_{s,1}, v_{s,2}, \dots, v_{s,s_{n-1}}], v_{s,s_n})$. Since the Yoochoose dataset is too large, we only use its the most recent 1/64 fractions of the training sessions, denoted as Yoochoose 1/64. The statistics of datasets used throughout experiments are summarized in Table 1.

Baselines. To evaluate the performance of the proposed method, we comprehensively compare TAGNN with representative baselines. The traditional baselines include (a) frequency-based methods POP and S-POP, (b) similarity-based method Item-KNN [9], and (c) factorization-based methods Bayesian personalized ranking (BPR-MF) [7] and factorizing personalized Markov chain model (FPMC) [8]. We also consider deep learning baselines, including RNN-based recommender model GRU4REC [1], neural attentive recommender model (NARM) [5], short-term attention/memory priority model (STAMP) [6], and GNN-based recommender model SR-GNN [11].

Evaluation Metrics. We adopt two commonly-used metrics for evaluation, including Recall@20 and MRR@20. The former one evaluates the proportion of correct recommendation in an unranked list, while the latter one further considers the position of correct recommended items in a ranked list.

Hyperparameter Setup. Following previous methods [6, 11], we set d = 100 for hidden dimensionality in all experiments. We tune other hyperparameters based on a random 10% validation set. The initial learning rate for Adam is set to 0.001 and will decay by 0.1 after every 3 training epochs. The batch size is set to 100 for both datasets and the t_2 penalty is set to 10^{-5} .

3.2 Comparison with Baseline Methods (RQ1)

To evaluate the performance of the proposed method, we firstly compare it with existing representative baselines (RQ1). The overall

Table 2:	The j	performa	nce of	TAGNN	compared	with	other
baseline	meth	iods using	g two d	latasets.			

	Digin	etica	Yoochoose 1/64		
Method	Recall@20	Recall@20 MRR@20 Re		MRR@20	
POP	0.89	0.20	6.71	1.65	
S-POP	21.06	13.68	30.44	18.35	
Item-KNN	35.75	11.57	51.60	21.81	
BPR-MF	5.24	1.98	31.31	12.08	
FPMC	26.53	6.95	45.62	15.01	
GRU4REC	29.45	8.33	60.64	22.89	
NARM	49.70	16.17	68.32	28.63	
STAMP	45.64	14.32	68.74	29.67	
SR-GNN	50.73	17.59	70.57	30.94	
TAGNN	51.31	18.03	71.02	31.12	
Improv.(%)	1.14	2.50	0.64	0.58	

performance in terms of Recall@20 and MRR@20 is summarized in Table 2, with the highest performance highlighted in boldface.

In all, TAGNN aggregates session items into session graphs and further considers modeling user preference through target-aware attention. It is apparent from the table that the proposed TAGNN model achieves state-of-the-art performance on all datasets in terms of Recall@20 and MRR@20, which confirms the effectiveness of the proposed method.

This table is quite revealing in several ways. Firstly, traditional methods including POP and S-POP achieve poor performance. They mainly emphasize items with high co-occurrence, which is oversimplified in session-based recommendation. Interestingly, the simple method Item-KNN still shows favorable performance, compared with POP, BPR-MF, and FPMC. Without knowing the sequential information, Item-KNN only recommends items with high similarity. This may be explained by the fact that latent factors representing user preference play a key role in generating accurate recommendation. It can also be seen that Item-KNN surpasses most Markov-chain-based methods, such as BPR-MF and FPMC, which demonstrates that modeling limited dependencies in session sequence is not realistic in session-based recommendation scenarios.

Secondly, there is a clear trend that deep learning methods greatly outperform conventional models. These methods have a stronger capability to capture complex user behavior, leading to superior performance over traditional ones. Sequential models such as GRU4REC and NARM only considers single-way transitions between successive item. Compared with SR-GNN, which further models session as graphs and is able to capture more implicit connections between user clicks, these methods neglect complex item transitional patterns. However, the performance of these models is still inferior to that of the proposed method. TAGNN elaborates graph-based models by further considering user interests with target-aware attentions. This mechanism specifically activates diverse user interests given different target items, which improves the expressiveness of the recommender model. In summary, these results demonstrate the efficacy of the proposed TAGNN method.

²http://2015.recsyschallenge.com/challenge.html

³http://cikm2016.cs.iupui.edu/cikm-cup



Figure 3: The performance of different session representations.

3.3 Ablation Studies (RQ2)

We conduct ablation studies on session embedding strategies in this section. We design 4 model variants to analyze how different representations for user preference affect the model performance (RQ2): (a) local embedding only (TAGNN-L), (b) global embedding using average pooling only (TAGNN-Avg), (c) attentive global embedding (TAGNN-Att), and (d) local embedding plus attentive global embedding (TAGNN-L+Att). Figure 3 presents experimental results using different session representations.

It is apparent from this figure that the hybrid embedding strategy used by TAGNN achieves the best performance on all datasets, which verifies the necessity of explicitly incorporating target-aware attention for better representing user interests. Also, it is clear that the mixed use of attentive global embedding with local embedding outperforms other opponents, demonstrating the necessity to capture both long- and short-term preference exhibited within sessions. Please note that the performance of TAGNN-L, which merely uses the last item as the session representation stands out in this figure. It indicates that the last item has a great impact on a user's final action. Moreover, we can observe that using average pooling over items in session achieves bad performance. This phenomenon may be explained by the diversity of user behavior in the session, which further highlights the importance of using the proposed target-aware attention to capture diverse user interests.

4 CONCLUSION

In this paper, we have developed a novel target attentive graph neural network model for session-based recommendation. By incorporating graph modeling and a target-aware attention module, the proposed TAGNN jointly considers user interests given a certain target item as well as complex item transitions in sessions. We have conducted thorough empirical evaluation to investigate TAGNN. Extensive experiments on real-world benchmark datasets demonstrate the effectiveness of our model.

ACKNOWLEDGMENTS

This work is jointly supported by National Key Research and Development Program (2018YFB1402600, 2016YFB1001000) and National Natural Science Foundation of China (U19B2038, 61772528).

REFERENCES

- Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. In *ICLR*.
- [2] Fenyu Hu, Yanqiao Zhu, Shu Wu, Weiran Huang, Liang Wang, and Tieniu Tan. 2019. GraphAIR: Graph Representation Learning with Neighborhood Aggregation and Interaction. arXiv.org (2019).
- [3] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In ICLR.
- [4] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural Attentive Session-based Recommendation. In CIKM.
- [5] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. 2016. Gated Graph Sequence Neural Networks. In *ICLR*.
- [6] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: Short-Term Attention/Memory Priority Model for Session-based Recommendation. In KDD.
- [7] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In UAI.
- [8] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing Personalized Markov Chains for Next-basket Recommendation. In WWW.
- [9] Badrul Munir Sarwar, George Karypis, Joseph A. Konstan, and John Riedl. 2001. Item-based Collaborative Filtering Recommendation Algorithms. In WWW.
- [10] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. The Graph Neural Network Model. IEEE T-NN 20, 1 (2009).
- [11] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based Recommendation with Graph Neural Networks. In AAAI.
- [12] Yufeng Zhang, Xueli Yu, Zeyu Cui, Shu Wu, Zhongzhen Wen, and Liang Wang. 2020. Every Document Owns Its Structure: Inductive Text Classification via Graph Neural Networks. In ACL.
- [13] Guorui Zhou, Xiaoqiang Zhu, Chengru 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. In KDD.