

A Review-Driven Neural Model for Sequential Recommendation

Chenliang Li¹, Xichuan Niu², Xiangyang Luo^{3*}, Zhenzhong Chen² and Cong Quan⁴

¹School of Cyber Science and Engineering, Wuhan University, China

²School of Remote Sensing and Information Engineering, Wuhan University, China

³State Key Lab of Mathematical Engineering and Advanced Computing, Zhengzhou, China

⁴School of Computer Science, Wuhan University, China

{cllee, niuxichuan}@whu.edu.cn, xiangyangluo@126.com, {zzchen, quancong}@whu.edu.cn

Abstract

Writing review for a purchased item is a unique channel to express a user's opinion in E-Commerce. Recently, many deep learning based solutions have been proposed by exploiting user reviews for rating prediction. In contrast, there has been few attempt to enlist the semantic signals covered by user reviews for the task of collaborative filtering. In this paper, we propose a novel review-driven neural sequential recommendation model (named RNS) by considering users' intrinsic preference (long-term) and sequential patterns (short-term). In detail, RNS is devised to encode each user or item with the aspect-aware representations extracted from the reviews. Given a sequence of historical purchased items for a user, we devise a novel hierarchical attention over attention mechanism to capture sequential patterns at both union-level and individual-level. Extensive experiments on three real-world datasets of different domains demonstrate that RNS obtains significant performance improvement over uptodate state-of-the-art sequential recommendation models.

1 Introduction

Recommender systems are now an indispensable asset in many E-Commerce platforms to enhance their productivity. It is well recognized that these systems largely drive sales and improve user satisfaction by automatically pushing what a user is looking for [Keiningham *et al.*, 2007]. The next purchase decision made by a user is often influenced by her recent behaviors. For example, after buying a SLR camera, the user would be highly interested in camera lenses. Recently, sequential recommendation has drawn increasing attention from both academic and industrial circles. The task is to identify the next item that a user will purchase by considering her temporal preference as a sequence of purchased items.

The key challenge of sequential recommendation is to dynamically approximate the current preference of a user by

considering both the general preferences and sequential patterns between items. As being the seminal work, FPMC extends the latent factor learning model by encoding the sequential transition patterns between items [Rendle *et al.*, 2010]. The prediction made by FPMC is a linear combination of the user general preference and item sequential association which both are calculated by conducting inner-product between the corresponding latent vectors. Following this line, many variants are proposed to include more complex patterns [Wang *et al.*, 2015; He and McAuley, 2016]. Recent advances in deep learning techniques have bred many neural solutions for sequential recommendation. Both convolutional neural networks (CNN) and recurrent neural networks (RNN) based models are proposed to capture a user's dynamic preferences [Hidasi *et al.*, 2015; Yu *et al.*, 2016; Chen *et al.*, 2018; Tang and Wang, 2018], leading to state-of-the-art recommendation performance. However, these existing methods do not consider user reviews to enhance sequential recommendation. The semantic information contained in user reviews can reveal different features of items and also the preferences of users. It is widely recognized that exploiting review information would largely improve the rating prediction accuracy [Wang and Blei, 2011; Bao *et al.*, 2014; Zheng *et al.*, 2017; Wu *et al.*, 2018; Wu *et al.*, 2019; Li *et al.*, 2019].

To this end, we propose a review-driven neural model for sequential recommendation, named RNS. The incorporation of review text brings more rich semantic information and strengthens model's expressive ability. Specifically, we first embed users and items into low-dimensional dense spaces through aspect-aware convolutional processing of the review documents, at the same time, user's general preference is extracted. Then a hierarchical attention over attention mechanism is employed to capture user sequential pattern at both union-level and individual-level. After that, we combine user's intrinsic (long-term) preference and temporal (short-term) preference to get a hybrid representation and make more accurate next items recommendation. On three real-world datasets from different domains, RNS outperforms the existing state-of-the-art methods remarkably. Overall, the key contributions of this work are summarized as below:

- We propose a novel review-driven neural model that exploits reviews for sequential recommendation. To the best of our knowledge, this is the first attempt to harness

*Xiangyang Luo is the corresponding author.

the rich semantics from reviews for this task.

- We introduce an aspect-based convolutional network to identify user general preference from review document and utilize a hierarchical attention over attention mechanism to model user sequential preference at different granularities.
- On three real-world datasets with diverse characteristics, our results demonstrate the superiority of the proposed RNS in terms of four metrics.

2 Related Work

In the following, we review the existing literatures that are highly relevant to our work: namely review-based recommendation and sequential recommendation.

2.1 Review-based Recommendation

Conventionally, many recommender systems are mainly developed in the paradigm of collaborative filtering (CF). Within this line of literatures, great effort is made under the framework of matrix factorization or the probabilistic counterpart [Salakhutdinov and Mnih, 2007; Koren *et al.*, 2009]. The idea is to represent each user or item as a low dimension latent vector. The binary interaction for a user-item pair is then estimated by performing inner-product of the corresponding latent vectors. However, the user-item interaction data is very sparse, which is an inherent obstacle for latent factor learning based solutions. Currently, to tackle the data sparsity problem, many techniques have been developed by utilizing the semantic signals provided by reviews. These advances have delivered state-of-the-art recommendation performance [Wang and Blei, 2011; Bao *et al.*, 2014; Zheng *et al.*, 2017; Catherine and Cohen, 2017; Wu *et al.*, 2018; Wu *et al.*, 2019; Li *et al.*, 2019].

Many works propose to distill semantic latent factors from reviews with topic modeling techniques [Blei *et al.*, 2003]. These solutions incorporate the topical factors learnt from reviews into latent factor learning framework [Wang and Blei, 2011; Bao *et al.*, 2014]. However, given that the topic modeling represents a document as a bag-of-words (BOW), much semantic context information encoded with the word order is inevitably lost. The recent surge of deep learning has drawn remarkable attention from the community. Both convolutional neural networks (CNN) [Kim, 2014] and recurrent neural network (RNN) [Mikolov *et al.*, 2010] are widely adopted to extract semantic representation from reviews for rating prediction [Zheng *et al.*, 2017; Catherine and Cohen, 2017; Tay *et al.*, 2018; Wu *et al.*, 2018; Wu *et al.*, 2019; Li *et al.*, 2019]. Empowered by dense representation and neural composition ability, local contextual information in reviews is well preserved and easily composited to form high-level features. These neural solutions have delivered significant performance gain over the existing BOW-based alternatives.

Many works also propose to derive aspect-level features from reviews. Several approaches utilize the external toolkit for aspect extraction [Zhang *et al.*, 2014; He *et al.*, 2015]. Recently, several deep learning based solutions have been proposed. These works utilize attention mechanism to extract the aspect and derive the aspect-level representations [Cheng

et al., 2018; Chin *et al.*, 2018]. By highlighting the important words or aspect, the above mentioned solutions enable better rating prediction and facilitate semantic explainability. However, the aforementioned techniques are mainly devised for the task of rating prediction. Neither of these works focuses on enhancing sequential recommendation with user reviews.

2.2 Sequential Recommendation

Sequential recommendation can be considered as a special kind of CF with implicit feedback. The effective learning of sequential patterns is widely verified to be a critical issue for sequential recommendation [Liu *et al.*, 2009; Rendle *et al.*, 2010; Wang *et al.*, 2015; Hidasi *et al.*, 2015; He and McAuley, 2016; Yu *et al.*, 2016; Tang and Wang, 2018; Chen *et al.*, 2018; Kang and McAuley, 2018]. Earlier methods aim to extract sequential rules based on statistic co-occurrence pattern [Agrawal and Srikant, 1995; Liu *et al.*, 2009]. FPMC [Rendle *et al.*, 2010] is the first work that automatically encodes the sequential relations for items with latent vectors. Afterwards, several variants are proposed to incorporate auxiliary information. Fossil integrates high-order Markov chains and item similarity together for next item recommendation [He and McAuley, 2016]. HRM proposes a hierarchical architecture to capture both sequential patterns and users' general preference [Wang *et al.*, 2015]. Recent efforts focus on enhancing sequential recommendation by using deep learning techniques, leading to the state-of-the-art performance. GRU4Rec and DREAM utilize a RNN network to capture sequential patterns [Hidasi *et al.*, 2015; Yu *et al.*, 2016]. RUM utilizes the attention and memory mechanisms to express and update the users' preference in a dynamic fashion [Chen *et al.*, 2018]. The dynamic preference is modeled in both item and feature levels. Caser proposes to utilize both horizontal and vertical convolutional operations to harness the union-level and point-level sequential patterns [Tang and Wang, 2018]. They also propose a skip behavior to consider the items in the next few actions for training. SASRec builds a self-attention model for sequential recommendation [Kang and McAuley, 2018]. It can adaptively put different weights to historical items at each time step. Though these deep learning based solutions significantly enhance the performance of sequential recommendation. However, semantic information covered by user reviews are mainly overlooked. The proposed RNS in this work is the first attempt to exploit textual reviews in this field.

3 The Proposed Algorithm

In this section, we firstly give a formal formulation of review-driven sequential recommendation problem and then present the details of our proposed RNS.

3.1 Problem Formulation

Let \mathcal{U} and \mathcal{I} represent the user and item set, respectively. Each user $u \in \mathcal{U}$ is associated with a sequence of interacted items arranged chronologically as: $S^u = (S_1^u, S_2^u, \dots, S_{|S^u|}^u)$, where $S_t^u \in \mathcal{I}$ denotes the item purchased by user u at time step t . In this work, we exploit user reviews to enhance




Figure 1: The network architecture of RNS.

sequential recommendation with implicit feedback. Consequently, we merge the set of reviews written by user u and the set of reviews written for item i to form user document D_u and item document D_i respectively. Given user u , her recently purchased L items and their corresponding review documents, the goal is to rank candidate items in terms of the likelihood that the user will purchase in the next.

3.2 The Architecture of RNS

Figure 1 demonstrates the network architecture of RNS. The objective of RNS is to derive the item representation, the user’s intrinsic preference and temporal preference at union-level and individual-level from the corresponding reviews. Specifically, for each user u , we learn a long-term user representation \mathbf{p}_u^l by utilizing an aspect-aware convolutional network (A-CNN) over user document D_u . Likewise, each item i is encoded as a representation vector \mathbf{q}_i via another parallel A-CNN over item document D_i .

To capture user’s temporal preference denoted by \mathbf{p}_u^s , we utilize a hierarchical attention over attention mechanism to exploit the recent L items (*i.e.*, $\mathbf{q}_1^u, \dots, \mathbf{q}_L^u$) she has purchased at both union-level and individual-level. Then, the user’s current preference is formed through a linear fusion:

$$\mathbf{p}_u = \mathbf{p}_u^l + \alpha \cdot \mathbf{p}_u^s \quad (1)$$

where α controls the importance of the temporal preference. The preference score (*i.e.*, purchasing likelihood) of a candidate item j is computed:

$$s_{uj} = \sigma(\mathbf{p}_u^\top \mathbf{q}_j) \quad (2)$$

Here, we choose *sigmoid* function $\sigma(x) = 1/(1 + e^{-x})$ to meet the binary constraint (*i.e.*, 0/1).

3.3 Aspect-Aware CNN (A-CNN)

Since the A-CNN architectures for user and item documents are identical (but with different parameters), we only describe

the extraction process of user documents for simplicity. We first map each word in each user document into the corresponding embedding of d dimensions. Thus, the user document is transformed to an embedding matrix $\mathbf{M}_u \in \mathbb{R}^{l \times d}$, where l is the document length. However, even for the same word, the semantics or sentimental polarity could be totally different for two different aspects in the same domain. For instance, word “low” in two sentences “The price is very low” and “This computer is low resolution” convey contrary sentiments towards aspects *price* and *resolution* respectively. So all words share the same d dimensional vector across all aspects is unreasonable. Hence, we introduce aspect-specific embedding transformation matrices: $\mathbf{T}_a \in \mathbb{R}^{d \times d}$ and multiply it with original word embedding matrices: $\mathbf{M}_u^a = \mathbf{M}_u \mathbf{T}_a$ where \mathbf{M}_u^a is the embedding matrix for D_u for aspect a . Therefore, we can represent D_u as a tensor $\mathbf{M}_u^* \in \mathbb{R}^{l \times d \times K}$ for K aspects. We then utilize a multi-channel convolutional operation analogous to colored images for feature extraction. Specifically, A-CNN has n filters $\mathbf{F}^k \in \mathbb{R}^{h \times d \times K}$, $1 \leq k \leq n$, where h is the filter height size that can have multiple values. For example, if $n = 10$, $h = \{1, 3, 5, 7, 9\}$, then there will be two filters for each size. The k^{th} filter derives its feature as follows:

$$\mathbf{z}^k[i] = \text{ReLU}(\mathbf{M}_u^*[i : i + h - 1] \odot \mathbf{F}^k + b^k) \quad (3)$$

where $1 \leq i \leq l - h + 1$ is the start point of sliding window in the user document, b^k is the bias, \odot is the convolution operator and *ReLU* is chosen as the activation function. The user’s long-term preference is encoded as follows:

$$\mathbf{p}_u^l = [\max(\mathbf{z}^1), \max(\mathbf{z}^2), \dots, \max(\mathbf{z}^n)] \quad (4)$$

The item representation \mathbf{q}_i is extracted with the same procedure.

Intuitively, the order of purchased items in user behaviors is pivotal for the sequential recommendation. Given a user u and her recent L purchased items $\mathbf{q}_1^u, \mathbf{q}_2^u, \dots, \mathbf{q}_L^u$ arranged

in chronological order, the latest item \mathbf{q}_L^u is more likely to reflect her temporal preference, while \mathbf{q}_1^u has minimal impact. Accordingly, we encode this temporal order information by including position embeddings to the item representation: $\mathbf{q}_m^u = \mathbf{q}_m^u + \mathbf{o}_m$ where \mathbf{o}_m is embedding for m -th position, and $1 \leq m \leq L$. Note that the sinusoid version of position embedding [Vaswani *et al.*, 2017] is also investigated, but it leads to unfavorable performance, which is in line with the observation in [Kang and McAuley, 2018].

3.4 Hierarchical Attention over Attention.

We further focus on deriving the user short-term preference \mathbf{p}_u^s from her recent L items. A user’s sequential preference could dynamically evolve over time. Moreover, different historical items could have different impacts on the user’s next purchase decision. To accommodate these characteristics, we resort to the attention mechanism that has achieved great success in relevant works [Tay *et al.*, 2018; Wu *et al.*, 2019]. Specifically, for each of user u ’s L purchased items \mathbf{q}_m^u ($1 \leq m \leq L$) and each candidate item \mathbf{q}_j , we first calculate the weight of each \mathbf{q}_m^u via an attention mechanism:

$$w_m^u = \text{softmax}(\mathbf{q}_j^\top \mathbf{q}_m^u) = \frac{\exp(\mathbf{q}_j^\top \mathbf{q}_m^u)}{\sum_{i=1}^L \exp(\mathbf{q}_j^\top \mathbf{q}_i^u)} \quad (5)$$

Note that \mathbf{q}_m^u is a fusion of item representation and position embedding. This design enables RNS to easily incorporate temporal order for attention calculation. Then, the short-term preference of user u with respect to candidate item \mathbf{q}_j is computed as a weighted sum of the item embeddings as follows:

$$\mathbf{p}_u^{s_1} = \sum_{m=1}^L w_m^u \cdot \mathbf{q}_m^u \quad (6)$$

Here, $\mathbf{p}_u^{s_1}$ is user sequential preference representation at union-level. That is, $\mathbf{p}_u^{s_1}$ is jointly encoded by all previous L purchased items.

We can simply take $\mathbf{p}_u^{s_1}$ as the final short-term preference derived for user u , but it only captures sequential pattern at union-level, *i.e.*, every item of past L ones has contribution for next item prediction. However, a user would buy an Apple earphone just because he bought an iPhone recently. In this case, other purchases except iPhone are just noises for making recommendation. Hence, we further explore the previous purchase records’ influence on future action at individual level, *i.e.*, identify the most relevant item. Inspired by the pointer mechanism [Tay *et al.*, 2018], we choose the item with maximum attention weight:

$$m_u = \arg \max_m (w_m^u) \quad (7)$$

$$\mathbf{p}_u^{s_2} = \mathbf{q}_{m_u}^u \quad (8)$$

where $\mathbf{p}_u^{s_2}$ is user sequential pattern representation at individual level. Then, we further utilize attention mechanism to discriminate which short-term preference is more important: union-level or individual-level. The final short-term user

	IV	PS	THI
# users	1,372	7,417	10,076
# items	7,957	33,798	66,710
# interactions	23,181	117,385	169,245
sparsity	99.79%	99.95%	99.97%
avg. actions per user	16.9	15.8	16.8

Table 1: Statistics of datasets

preference is calculated as follows:

$$\beta_n^u = \text{softmax}(\mathbf{q}_j^\top \mathbf{p}_u^{s_n}) = \frac{\exp(\mathbf{q}_j^\top \mathbf{p}_u^{s_n})}{\sum_{i=1}^2 \exp(\mathbf{q}_j^\top \mathbf{p}_u^{s_i})} \quad (9)$$

$$\mathbf{p}_u^s = \sum_{n=1}^2 \beta_n^u \cdot \mathbf{p}_u^{s_n} \quad (10)$$

The preference score is then calculated according to Equation 1-2.

3.5 Model Inference and Optimization

The model parameters of RNS include word embeddings, position embeddings, aspect transformation matrices, and two A-CNN networks. For each user u , we extract each L successive items and their next item as the target from the user’s sequence \mathcal{S}^u to form a training instance. Following previous work [Tang and Wang, 2018], for each training instance t_j with target item j , we randomly sample x negative items, denoted as $\mathcal{N}(j)$. Let \mathcal{C}^u be the collection of user u ’s all training instances. The objective function is defined as binary cross-entropy loss with L_2 norm regularization:

$$\mathcal{L} = \sum_u \sum_{t_j \in \mathcal{C}^u} (-\log(s_{uj}) + \sum_{i \in \mathcal{N}(j)} -\log(1 - s_{ui})) + \lambda \|\Theta\|^2 \quad (11)$$

where λ is coefficient for the regularization and Θ denotes all model parameters. The model is trained via Adma optimizer.

4 Experiments

In this section, we conduct experiments on three real-world datasets to evaluate our proposed RNS against up-to-date state-of-the-art methods. We then perform the ablation study to investigate the validity of each design choice made in RNS. Finally, we demonstrate the superiority of the proposed hierarchical attention over attention mechanism with case study.

4.1 Experimental Setup

Datasets. We perform our experiments on the Amazon dataset¹. This dataset contains product purchase history from Amazon ranging from May 1996 to July 2014. We conduct experiments on three categories: *Instant Video (IV)*, *Pet Supplies (PS)* and *Tools and Home Improvement (THI)*. Following previous works [Chen *et al.*, 2018], the users that have purchased less than 10 items are removed. We then take explicit ratings as purchase behavior. We hold the first 70% of

¹<http://jmcauley.ucsd.edu/data/amazon/>

Datasets	Instant Video				Pet Supplies				Tools and Home Improvement			
	Precision	Recall	NDCG	HR	Precision	Recall	NDCG	HR	Precision	Recall	NDCG	HR
Measures@5	0.0783	0.0876	0.0873	0.2648	0.0615	0.0622	0.0750	0.2376	0.0542	0.0575	0.0587	0.2105
Pop	0.1067	0.1214	0.1187	0.3900	0.0985	0.1009	0.1075	0.3761	0.0785	0.0828	0.0911	0.3047
FPMC	0.1036	0.1166	0.1137	0.3804	0.0943	0.0940	0.1015	0.3646	0.0712	0.0747	0.0821	0.2915
GRU4Rec	0.1160	0.1258	0.1245	0.3921	0.1013	0.1005	0.1086	0.3790	0.0790	0.0853	0.0936	0.3166
RUM(I)	0.1142	0.1274	0.1274	0.3928	<u>0.1041</u>	0.1047	0.1146	0.3973	0.0803	0.0855	0.0925	0.3199
RUM(F)	0.1152	0.1321	0.1456	0.4060	0.1038	0.1136	0.1252	0.3975	0.0819	0.0873	0.1029	0.3378
Caser	<u>0.1183</u>	0.1295	0.1436	0.4054	0.1026	0.1105	0.1241	0.3922	0.0820	0.0875	0.1036	0.3384
SASRec	0.1329	0.1531	0.1648	0.4446	0.1146	0.1252	0.1388	0.4362	0.0894	0.0943	0.1120	0.3614
RNS	12.4%	15.9%	13.2%	9.5%	10.0%	10.2%	10.9%	9.7%	9.0%	7.8%	8.1%	6.8%
Improvement												

Table 2: Performance comparison for baselines and RNS. The best and second best results are highlighted in boldface and underlined respectively. Improvement over the best baseline are shown in the last row.

items in each user’s sequence as the training set and the rest are used for testing. After preprocessing, basic statistics of the datasets are listed in Table 1.

Metrics. Following previous works [Kang and McAuley, 2018], for each test item, we randomly sample 100 negative items, and rank these items with the items in the testing set together. We adopt 4 widely used metrics for performance evaluation: *Precision@N*, *Recall@N*, *NDCG@N* and *HR@N*, where *N* indicates top-*N* ranked items. The first two metrics regard the recommendation task as classification problem and evaluate the Top-*N* classification results. *NDCG@N* (Normalized Discounted Cumulative Gain) evaluates ranking quality by taking the positions of ground-truth items in the ranking list into consideration. *HR@N* (Hit Ratio) gives the ratio of testing instances that can obtain at least one correct recommendation in the top-*N*.

Baseline Methods. We compare RNS with the following state-of-the-art baselines including sequential learning based models and neural network based models: (1) the popularity based recommender, **Pop**; (2) the factor based sequential pattern learning, **FPMC** [Rendle *et al.*, 2010]; (3) RNN based sequential recommender, **GRU4Rec** [Hidasi *et al.*, 2015]; (4) attention and memory network based sequential recommender, item-level **RUM(I)** and feature-level **RUM(F)** [Chen *et al.*, 2018]; (5) CNN based sequential recommender, **Caser** [Tang and Wang, 2018]; (6) self-attention based sequential recommender, **SASRec** [Kang and McAuley, 2018].

Parameter Settings. As to the baselines, we utilize the recommended setting by their original work. For the proposed RNS, we set $L = 5$ which is a common setting in most relevant works. The number of aspects is set to be 5^2 , and embedding size is set to 25, $n = 10$ and $h = \{1, 3, 5, 7, 9\}$. The learning rate is set to 0.001, $x = 3$, $\alpha = 0.1$ and $\lambda = 0.0001$. We set $N = 5$ for performance evaluation.

4.2 Results and Discussion

The overall performance of different methods is presented in Table 2. Several observations can be made:

First, the non-personalized Pop method gives the worst performance on all metrics across three datasets, which is followed by RNN based solution GRU4Rec. Also, The non-

²We observe that K is optimal at 5 on all datasets. Detailed results are not included due to page limitation.

Datasets	IV		PS		THI	
	Recall	HR	Recall	HR	Recall	HR
Measures@5						
RNS	0.1531	0.4446	0.1252	0.4362	0.0943	0.3614
RNS-i	0.1455	0.4264	0.1214	0.4277	0.0904	0.3487
RNS-u	0.1480	0.4308	0.1219	0.4297	0.0885	0.3446
RNS-pe	0.1516	0.4359	0.1220	0.4289	0.0905	0.3489
RNS-at	0.1468	0.4176	0.1214	0.4293	0.0926	0.3544

Table 3: Ablation study of RNS.


neural method FPMC surpasses GRU4Rec significantly. Here FPMC models user preference through integrating matrix factorization (MF) with sequential information, while GRU4Rec utilizes RNNs based on the hypothesis that the latter state depends on the previous one. This result might highlight that using RNN along is not sufficient to learn users’ dynamic preference and long-term user preference modeled by MF in FPMC is also beneficial.

Second, RUM(I) and RUM(F) achieve comparable results on three datasets, indicating the memory mechanism and updating operation in RUM can provide expressive ability to model user dynamic preference. Caser and SASRec outperform the other baselines on all settings except *Precision@5* in *Pet Supplies*. Caser generally achieves better recommendation performance than other baselines in both *Instant Video* and *Pet Supplies*. In contrast, SASRec dominates in *Tools and Home Improvement*.

Third, our proposed RNS consistently outperforms all baselines by a wide margin across all datasets. Specifically, the average performance gains of all metrics on three datasets are 12.8%, 10.2%, 7.9% in *Instant Video*, *Pet Supplies* and *Tools and Home Improvement* respectively. The substantial improvement of our model over the baselines could be credited to the following reasons: (1) we model user’s long-term preference through aspect-aware convolutional processing of the review document. The rich semantics provided by the reviews help us better characterize user’s general preference; (2) we capture sequential patterns at both union-level and individual-level based on a novel hierarchical attention over attention mechanism. These traits are confirmed by further analysis of RNS in the following.

4.3 Ablation Study

As there are many components in RNS architecture, we evaluate their effects via ablation study, as shown in Table 3.


 Figure 2: Parameter sensitivity of RNS: L and α

Due to space limitation, we only list the results with respect to *Recall* and *HR*. Similar performance patterns are also observed for *Precision* and *NDCG*. RNS-u means that the short-term user preference is modeled by excluding union-level, *i.e.*, $\mathbf{p}_u^s = \mathbf{p}_u^{s_2}$, and RNS-i for excluding individual level, *i.e.*, $\mathbf{p}_u^s = \mathbf{p}_u^{s_1}$. RNS-pe means that position embeddings are excluded, while RNS-at means that only one aspect is considered, *i.e.*, without applying aspect transformation.

From Table 3, we can observe that keeping only union-level or individual-level preference will reduce the recommendation accuracy, suggesting that the hierarchical attention mechanism is beneficial to identify the user’s dynamic preference. On both *Instant Video* and *Pet Supplies* datasets, RNS-u outperforms RNS-i on two metrics, while on *Tools and Home Improvement* dataset RNS-i is better. This observation suggests that the importance of union-level and individual-level depends on dataset. Without the position embeddings, RNS-pe achieves worse performance than RNS on three datasets. This suggests that temporal information in forms of purchase order is a useful signal for sequential recommendation. Also, without utilizing aspect-aware feature extraction, RNS experiences the performance degradation to some extent. This suggests that aspect-aware features extracted from reviews are more discriminative.

4.4 Parameter Sensitivity

We further study the effects of parameters L and α on model performance. The results of *Precision@5* on three datasets are shown. We omit the similar patterns observed for other metrics due to space limitation. Figure 2(a) plots the performance pattern by varying L values. We observe that RNS consistently achieves better performance at $L = 5$ across datasets. When L becomes smaller or larger, RNS experiences performance degradation to some extent. This is reasonable since a small L produces less useful historical signals and a large L would inevitably introduce much noise. Hence, we set $L = 5$ in our experiments.

Parameter α in Equation 1 controls the importance of sequential patterns (*i.e.*, union-level and individual-level preferences) in RNS. Figure 2(b) plots the performance pattern by varying α values. When $\alpha = 0$, RNS degrades to a variant of DeepCoNN proposed for rating prediction [Zheng *et al.*, 2017]. We observe that RNS performs much worse in this setting. This is consistent with the existing finding that sequential patterns are crucial for sequential recommendation [Rendle *et al.*, 2010]. The optimal performance is consistently achieved when $\alpha = 0.1$ in all the three datasets. When




Figure 3: Illustration of attention weights hierarchically.

α becomes increasingly larger, the performance of RNS further deteriorates to a larger extent. These results further suggest that both short-term and long-term user preferences are complementary to each other for sequential recommendation. Accordingly, we set $\alpha = 0.1$ in the experiments.

4.5 Case Study

To illustrate the intuition of the hierarchical attention over attention at both union-level and individual-level, We present the case study about two users sampled in *Tools and Home Improvement*. The attention weights in both union-level and individual-level are displayed for interpretation.

Figure 3(a) demonstrates a case where individual level is more crucial than union-level, as the user wants to buy an air conditioner cover just because he bought an air conditioner not long ago. Figure 3(b) demonstrates a completely opposite example, where all recently purchased 5 items contribute nearly equal for next purchase decision. The user needs wall-paper adhesive since he bought several kinds of wallpapers and related tools before. These observations confirm that the proposed hierarchical attention over attention mechanism in RNS is effective to capture users’ dynamic preference.

5 Conclusion

In this paper, we propose a novel neural model for sequential recommendation with reviews, named RNS. Our model incorporates both a user’s long-term intrinsic preference and short-term preference to predict her next actions. It utilizes aspect-aware CNN network to extract user and item representations from the corresponding review documents. A novel hierarchical attention over attention mechanism is proposed to capture sequential patterns at both union-level and individual-level. Experimental results show that RNS significantly outperforms strong state-of-the-art methods under different metrics. We will extend RNS to facilitate fine-grained textual explanation for sequential recommendation.

Acknowledgements

This research was supported by National Natural Science Foundation of China (No. 61872278, No. U1636219, No. U1804263), National Key R&D Program of China (No. 2016YFB0801303, No. 2016QY01W0105), Science and Technology Innovation Talent Project of Henan Province (No. 184200510018).

References

- [Agrawal and Srikant, 1995] Rakesh Agrawal and Ramakrishnan Srikant. Mining sequential patterns. In *ICDE*, pages 3–14, 1995.
- [Bao *et al.*, 2014] Yang Bao, Hui Fang, and Jie Zhang. Topicmf: Simultaneously exploiting ratings and reviews for recommendation. In *AAAI*, pages 2–8, 2014.
- [Blei *et al.*, 2003] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [Catherine and Cohen, 2017] Rose Catherine and William W. Cohen. Transnets: Learning to transform for recommendation. In *RecSys*, pages 288–296, 2017.
- [Chen *et al.*, 2018] Xu Chen, Hongteng Xu, Yongfeng Zhang, Jiaxi Tang, Yixin Cao, Zheng Qin, and Hongyuan Zha. Sequential recommendation with user memory networks. In *WSDM*, pages 108–116, 2018.
- [Cheng *et al.*, 2018] Zhiyong Cheng, Ying Ding, Lei Zhu, and Mohan S. Kankanhalli. Aspect-aware latent factor model: Rating prediction with ratings and reviews. In *WWW*, pages 639–648, 2018.
- [Chin *et al.*, 2018] Jin Yao Chin, Kaiqi Zhao, Shafiq Joty, and Gao Cong. Anr: Aspect-based neural recommender. In *CIKM*, pages 147–156, 2018.
- [He and McAuley, 2016] Ruining He and Julian McAuley. Fusing similarity models with markov chains for sparse sequential recommendation. In *ICDM*, pages 191–200, 2016.
- [He *et al.*, 2015] Xiangnan He, Tao Chen, Min-Yen Kan, and Xiao Chen. Trirank: Review-aware explainable recommendation by modeling aspects. In *CIKM*, pages 1661–1670, 2015.
- [Hidasi *et al.*, 2015] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. Session-based recommendations with recurrent neural networks. *CoRR*, abs/1511.06939, 2015.
- [Kang and McAuley, 2018] Wang-Cheng Kang and Julian McAuley. Self-attentive sequential recommendation. In *ICDM*, pages 197–206, 2018.
- [Keiningham *et al.*, 2007] Timothy L. Keiningham, Bruce Cooil, Lerzan Aksoy, Tor W. Andreassen, and Jay Weiner. The value of different customer satisfaction and loyalty metrics in predicting customer retention, recommendation, and share-of-wallet. *Managing Service Quality: An International Journal*, 17(4):361–384, 2007.
- [Kim, 2014] Yoon Kim. Convolutional neural networks for sentence classification. In *EMNLP*, pages 1746–1751, 2014.
- [Koren *et al.*, 2009] Yehuda Koren, Robert M. Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *IEEE Computer*, 42(8):30–37, 2009.
- [Li *et al.*, 2019] Chenliang Li, Cong Quan, Peng Li, Yunwei Qi, Yuming Deng, and Libing Wu. A capsule network for recommendation and explaining what you like and dislike. In *SIGIR*, 2019.
- [Liu *et al.*, 2009] Duen-Ren Liu, Chin-Hui Lai, and Wang-Jung Lee. A hybrid of sequential rules and collaborative filtering for product recommendation. *Inf. Sci.*, 179(20):3505–3519, 2009.
- [Mikolov *et al.*, 2010] Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Cernocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Proc. of INTERSPEECH*, pages 1045–1048, 2010.
- [Rendle *et al.*, 2010] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. Factorizing personalized markov chains for next-basket recommendation. In *WWW*, pages 811–820, 2010.
- [Salakhutdinov and Mnih, 2007] Ruslan Salakhutdinov and Andriy Mnih. Probabilistic matrix factorization. In *NIPS*, pages 1257–1264, 2007.
- [Tang and Wang, 2018] Jiaxi Tang and Ke Wang. Personalized top-n sequential recommendation via convolutional sequence embedding. In *WSDM*, pages 565–573, 2018.
- [Tay *et al.*, 2018] Yi Tay, Anh Tuan Luu, and Siu Cheung Hui. Multi-pointer co-attention networks for recommendation. In *KDD*, pages 2309–2318, 2018.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pages 6000–6010, 2017.
- [Wang and Blei, 2011] Chong Wang and David M. Blei. Collaborative topic modeling for recommending scientific articles. In *KDD*, pages 448–456, 2011.
- [Wang *et al.*, 2015] Pengfei Wang, Jiafeng Guo, Yanyan Lan, Jun Xu, Shengxian Wan, and Xueqi Cheng. Learning hierarchical representation model for nextbasket recommendation. In *SIGIR*, pages 403–412, 2015.
- [Wu *et al.*, 2018] Libing Wu, Cong Quan, Chenliang Li, and Donghong Ji. PARL: let strangers speak out what you like. In *CIKM*, pages 677–686, 2018.
- [Wu *et al.*, 2019] Libing Wu, Cong Quan, Chenliang Li, Qian Wang, Bolong Zheng, and Xiangyang Luo. A context-aware user-item representation learning for item recommendation. *ACM Trans. Inf. Syst.*, 37(2):22:1–22:29, 2019.
- [Yu *et al.*, 2016] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. A dynamic recurrent model for next basket recommendation. In *SIGIR*, pages 729–732, 2016.
- [Zhang *et al.*, 2014] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *SIGIR*, pages 83–92, 2014.
- [Zheng *et al.*, 2017] Lei Zheng, Vahid Noroozi, and Philip S. Yu. Joint deep modeling of users and items using reviews for recommendation. In *WSDM*, pages 425–434, 2017.