

Personalizing Session-based Recommendations with Hierarchical Recurrent Neural Networks

Massimo Quadrana
Politecnico di Milano, Milan, Italy
massimo.quadrana@polimi.it

Balázs Hidasi
Gravity R&D, Budapest, Hungary
balazs.hidasi@gravityrd.com

Alexandros Karatzoglou
Telefonica Research, Barcelona, Spain
alexk@tid.es

Paolo Cremonesi
Politecnico di Milano, Milan, Italy
paolo.cremonesi@polimi.it

ABSTRACT

Session-based recommendations are highly relevant in many modern on-line services (e.g. e-commerce, video streaming) and recommendation settings. Recently, Recurrent Neural Networks have been shown to perform very well in session-based settings. While in many session-based recommendation domains user identifiers are hard to come by, there are also domains in which user profiles are readily available. We propose a seamless way to personalize RNN models with cross-session information transfer and devise a Hierarchical RNN model that relays and evolves latent hidden states of the RNNs across user sessions. Results on two industry datasets show large improvements over the session-only RNNs.

CCS CONCEPTS

• Information systems → Recommender systems; • Computing methodologies → Neural networks;

KEYWORDS

recurrent neural networks; personalization; session-based recommendation; session-aware recommendation

1 INTRODUCTION

In many online systems where recommendations are applied, interactions between a user and the system are organized into sessions. A session is a group of interactions that take place within a given time frame. Sessions from a user can occur on the same day, or over several days, weeks, or months. A session usually has a goal, such as finding a good restaurant in a city, or listening to music of a certain style or mood.

Providing recommendations in these domains poses unique challenges that until recently have been mainly tackled by applying conventional recommender algorithms [10] on either the last interaction or the last session (session-based recommenders). Recurrent Neural Networks (RNNs) have been recently used for the purpose of session-based recommendations [7] outperforming item-based

methods by 15% to 30% in terms of ranking metrics. In *session-based recommenders*, recommendations are provided based solely on the interactions in the current user session, as users are assumed to be anonymous. But in many of these systems there are cases where a user might be logged-in (e.g. music streaming services) or some form of user identifier might be present (cookie or other identifier). In these cases it is reasonable to assume that the user behavior in past sessions might provide valuable information for providing recommendations in the next session.

A simple way of incorporating past user session information in session-based algorithm would be to simply concatenate past and current user sessions. While this seems like a reasonable approach, we will see in the experimental section that this does not yield the best results.

In this work we describe a novel algorithm based on RNNs that can deal with both cases: (i) *session-aware recommenders*, when user identifiers are present and propagate information from the previous user session to the next, thus improving the recommendation accuracy, and (ii) *session-based recommenders*, when there are no past sessions (i.e., no user identifiers). The algorithm is based on a Hierarchical RNN where the hidden state of a lower-level RNN at the end of one user session is passed as an input to a higher-level RNN which aims at predicting a good initialization (i.e., a good context vector) for the hidden state of the lower RNN for the next session of the user.

We evaluate the Hierarchical RNNs on two datasets from industry comparing them to the plain session-based RNN and to item-based collaborative filtering. Hierarchical RNNs outperform both alternatives by a healthy margin.

2 RELATED WORK

Session-based recommendations. Classical CF methods (e.g. matrix factorization) break down in the session-based setting when no user profile can be constructed from past user behavior. A natural solution to this problem is the item-to-item recommendation approach [11, 16]. In this setting an item-to-item similarity matrix is precomputed from the available session data, items that are often clicked together in sessions are deemed to be similar. These similarities are then used to create recommendations. While simple, this method has been proven to be effective and is widely employed. Though, these methods only take into account the last click of the user, in effect ignoring the information of the previous clicks.

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'17, August 27–31, 2017, Como, Italy.

© 2017 ACM. 978-1-4503-4652-8/17/08...\$15.00

DOI: <http://dx.doi.org/10.1145/3109859.3109896>

Recurrent Neural Models. RNNs are the deep models of choice when dealing with sequential data [12]. RNNs have been used in image and video captioning, time series prediction, natural language processing and much more. Long Short-Term Memory (LSTM) [9] networks are a type of RNNs that have been shown to work particularly well, it includes additional gates that regulate when and how much to take the input into account and when to reset the hidden state. This helps with the vanishing gradient problem that often plagues the standard RNN models. Slightly simplified version of LSTM – that still maintains all their properties – are Gated Recurrent Units (GRUs) [3] which we use in this work.

RNNs were first used to model session data in [7]. The recurrent neural network is trained with a ranking loss on a one-hot representation of the session (clicked) item-IDs. The RNN is then used to provide recommendations after each click for new sessions. This work only focused on the clicked item-IDs in the current session while here we aim at modeling the user behavior across sessions as well. RNNs were also used to jointly model the content or features of items together with click-sequence interactions [8]. By including item features extracted for example, from the thumbnail image of videos or the textual description of a product, the so-called parallel-RNN model provided superior recommendation quality wrt. ‘plain’ RNNs. In [19], authors proposed data augmentation techniques to improve the performance of the RNN for session-based recommendations, these techniques have though the side effect of increasing training times as a single session is split into several sub-sessions for training. RNNs have also been used in more standard user-item collaborative filtering settings where the aim is to model the evolution of the user and items factors [20] [4] where the results are though less impressive, with the proposed methods barely outperforming standard matrix factorization methods. Finally a sequence to sequence model with a version of Hierarchical Recurrent Neural Networks was used for generative context-aware query suggestion in [17].

3 MODEL

In this section we describe the proposed Hierarchical RNN (*HRNN* henceforth) model for personalized session-based recommendation.

3.1 Session-based Recurrent Neural Network

Our model is based on the session-based Recurrent Neural Network (RNN henceforth) model presented in [7]. RNN is based on a single Gated Recurrent Unit (GRU) layer that models the interactions of the user within a session. The RNN takes as input the current item ID in the session and outputs a score for each item representing the likelihood of being the next item in the session.

Formally, for each session $S_m = \{i_{m,1}, i_{m,2}, \dots, i_{m,N_m}\}$, RNN computes the following session-level representation

$$s_{m,n} = GRU_{ses}(i_{m,n}, s_{m,n-1}), n = 1, \dots, N_m - 1 \quad (1)$$

where GRU_{ses} is the *session-level* GRU and $s_{m,n}$ its hidden state at step n , being $s_{m,0} = 0$ (the null vector), and $i_{m,n}$ is the one-hot vector of the current item ID ¹.

The output of the RNN is a score $\hat{r}_{m,n}$ for every item in the catalog indicating the likelihood of being the next item in the session

(or, equivalently, its relevance for the next step in the session)

$$\hat{r}_{m,n} = g(s_{m,n}), n = 1, \dots, N_m - 1 \quad (2)$$

where $g(\cdot)$ is a non-linear function like softmax or tanh depending on the loss function. During training, scores are compared to a one-hot vector of the next item ID in the session to compute the loss. The network can be trained with several ranking loss functions such as cross-entropy, BPR [14] and TOP1 [7]. In this work, the TOP1 loss always outperformed other ranking losses, so we consider only it in the rest of the paper. The **TOP1** loss is the regularized approximation of the relative rank of the relevant item. The relative rank of the relevant item is given by $\frac{1}{N_S} \cdot \sum_{j=1}^{N_S} I\{\hat{r}_{s,j} > \hat{r}_{s,i}\}$ where $r_{s,j}$ is the score of a sampled ‘irrelevant item’. $I\{\cdot\}$ is approximated with a sigmoid. To force the scores of negative examples (‘irrelevant items’) towards zero a regularization term is added to the loss. The final loss function is as follows: $L_s = \frac{1}{N_S} \cdot \sum_{j=1}^{N_S} \sigma(\hat{r}_{s,j} - \hat{r}_{s,i}) + \sigma(\hat{r}_{s,j}^2)$

RNN is trained efficiently with session-parallel mini-batches. At each training step, the input to GRU_{ses} is the stacked one-hot representation of the current item ID of a batch of sessions. The session-parallel mechanism keeps the pointers to the current item of every session in the mini-batch and resets the hidden state of the RNN when sessions end. To further reduce the computational complexity, the loss is computed over the current item IDs and a sample of negative items. Specifically, the current item ID of each session is used as positive item and the IDs of the remaining sessions in the mini-batch as negative items when computing the loss. This makes explicit negative item sampling unnecessary and enables popularity-based sampling. However, since user-identifiers are unknown in pure session-based scenarios, there are good chances that negative samples will be ‘contaminated’ by positive items the user interacts with in other sessions.

3.2 Personalized Session-based Hierarchical Recurrent Neural Network

Our *HRNN* model builds on top of RNN by: (i) adding an additional GRU layer to model information across user sessions and to track the evolution of the user interests over time; (ii) using a powerful user-parallel mini-batch mechanism for efficient training.

3.2.1 Architecture. Beside the session-level GRU, our *HRNN* model adds one *user-level* GRU (GRU_{usr}) to model the user activity across sessions.

Figure 1 shows a graphical representation of *HRNN*. At each time step, recommendations are generated by GRU_{ses} , as in RNN. However, when a session ends, the user representation is updated. When a new session starts, the hidden state of GRU_{usr} is used to initialize GRU_{ses} and, optionally, propagated in input to GRU_{ses} .

Formally, for each user u with sessions $C^u = \{S_1^u, S_2^u, \dots, S_{M_u}^u\}$, the user-level GRU takes as input the session-level representations $s_1^u, s_2^u, \dots, s_{M_u}^u$, being $s_m^u = s_{m,N_m-1}^u$ the last hidden state of GRU_{ses} of each user session S_m^u , and uses them to update the user-level representation c_m^u . Henceforth we drop the user superscript u to unclutter notation. The user-level representation c_m is updated as

$$c_m = GRU_{usr}(s_m, c_{m-1}), m = 1, \dots, M_u \quad (3)$$

where $c_0 = 0$ (the null vector). The input to the user-level GRU is connected to the last hidden state of the session-level GRU. In this

¹To simplify the explanation we use a notation similar to [17].

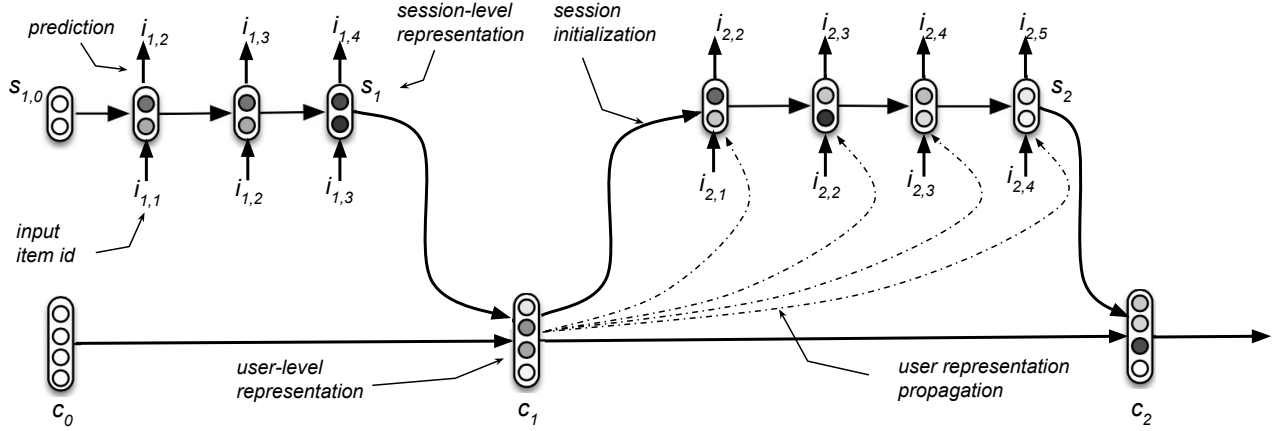


Figure 1: Graphical representation of the proposed Hierarchical RNN model for personalized session-based recommendation. The model is composed of an hierarchy of two GRUs, the session-level GRU (GRU_{ses}) and the user-level GRU (GRU_{usr}). The session-level GRU models the user activity within sessions and generates recommendations. The user-level GRU models the evolution of the user across sessions and provides personalization capabilities to the session-level GRU by initializing its hidden state and, optionally, by propagating the user representation in input.

way, the user-level GRU can track the evolution of the user across sessions and, in turn, model the dynamics user interests seamlessly. Notice that the user-level representation is kept *fixed* throughout the session and it is updated only when the session ends.

The user-level representation is then used to initialize the hidden state of the session-level GRU. Given c_m , the initial hidden state $s_{m+1,0}$ of the session-level GRU for the following session is set to

$$s_{m+1,0} = \tanh(W_{init}c_m + b_{init}) \quad (4)$$

where W_{init} and b_{init} are the initialization weights and biases respectively. In this way, the information relative to the preferences expressed by the user in the previous sessions is transferred to the session-level. Session-level representations are then updated as follows

$$s_{m+1,n} = GRU_{ses}(i_{m+1,n}, s_{m+1,n-1} [, c_m]), \quad n = 1, \dots, N_{m+1} - 1 \quad (5)$$

where the square brackets indicate that c_m can be *optionally* propagated in input to the session-level GRU.

The model is trained end-to-end using back-propagation [15]. The weights of GRU_{usr} are updated only between sessions, i.e. when a session ends and when the forthcoming session starts. However, when the user representation is propagated in input to GRU_{ses} , the weights of GRU_{usr} are updated also within sessions even if c_m is kept fixed. We also tried with propagating the user-level representation to the final prediction layer (i.e., by adding term c_m in Equation 2) but we always incurred into severe degradation of the performances, even wrt. simple session-based RNN. We therefore discarded this setting from this discussion.

Note here that the GRU_{usr} does not simply pass on the hidden state of the previous user session to the next but also learns (during

training) how user sessions evolve during time. We will see in the experimental section that this is crucial in achieving increased performance. In effect GRU_{usr} computes and evolves a user profile that is based on the previous user sessions, thus in effect personalizing the GRU_{ses} . In the original RNN, users who had clicked/interacted with the same sequence of items in a session would get the same recommendations; in *HRNN* this is not anymore the case, recommendations will be influenced by the the users past sessions as well.

In summary, we considered the following two different *HRNN* settings, depending on whether the user representation c_m is considered in Equation 5:

- *HRNN Init*, in which c_m is used only to initialize the representation of the next session.
- *HRNN All*, in which c_m is used for initialization and propagated in input at each step of the next session.

In *HRNN Init*, the session-level GRU can exploit the historical preferences along with the session-level dynamics of the user interest. *HRNN All* instead enforces the usage of the user representation at session-level at the expense of a slightly greater model complexity. As we will see, this can lead to substantially different results depending on the recommendation scenario.

3.2.2 Learning. For the sake of efficiency in training, we have edited the session-parallel mini-batch mechanism described in [8] to account for user identifiers during training (see Figure 2). We first group sessions by user and then sort session events within each group by time-stamp. We then order users at random. At the first iteration, the first item of the first session of the first B users constitute the input to the *HRNN*; the second item in each session




Figure 2: User-parallel mini-batches for mini-batch size 2.

constitute its output. The output is then used as input for the next iteration, and so on. When a session in the mini-batch ends, Equation 3 is used to update the hidden state of GRU_{usr} and Equation 4 to initialize the hidden state of GRU_{ses} for the forthcoming session, if any. When a user has been processed completely, the hidden states of both GRU_{usr} and GRU_{ses} are reset and the next user is put in its place in the mini-batch.

With user-parallel mini-batches we can train $HRNN$ s efficiently over users having different number of sessions and sessions of different length. Moreover, this mechanism allows to sample negative items in a user-independent fashion, hence reducing the chances of ‘contamination’ of the negative samples with actual positive items. The sampling procedure is still popularity-based, since the likelihood for an item to appear in the mini-batch is proportional to its popularity. Both properties are known to be beneficial for pairwise learning with implicit user feedback [13].

4 EXPERIMENTS

In this section we describe our experimental setup and provide an in-depth discussion of the obtained results.

4.1 Datasets

We used two datasets for our experiments. The first is the XING² Recsys Challenge 2016 dataset [1] that contains interactions on job postings for 770k users over a 80-days period. User interactions come with time-stamps and interaction type (click, bookmark, reply and delete). We named this dataset XING. The second dataset is a proprietary dataset from a Youtube-like video-on-demand web site. The dataset tracks the videos watched by 13k users over a 2-months period. Viewing events lasting less than a fixed threshold are not tracked. We named this dataset VIDEO.

We manually partitioned the interaction data into sessions by using a 30-minute idle threshold. For the XING dataset we discarded interactions having type ‘delete’. We also discarded repeated interactions of the same type within sessions to reduce noise (e.g. repeated clicks on the same job posting within a session). We then preprocessed both datasets as follows. We removed items with support less than 20 for XING and 10 for VIDEO since items with low support are not optimal for modeling. We removed sessions having

²<https://www.xing.com/en>

Dataset	XING	VIDEO
Users	11,479	13,717
Items	59,297	19,218
Sessions	89,591	133,165
Events	546,862	825,449
Events per item [†]	6/9.3/14.4	9/43.0/683.7
Events per session [†]	4/6.1/5.5	4/6.2/8.2
Sessions per user [†]	6/7.8/4.8	7/9.8/7.8
Training - Events	488,576	745,482
Training - Sessions	78,276	120,160
Test - Events	58,286	79,967
Test - Sessions	11,315	13,605

Table 1: Main properties of the datasets ([†] median/mean/std).

< 3 interactions to filter too short and poorly informative sessions, and kept users having ≥ 5 sessions to have sufficient cross-session information for proper modeling of returning users.

The test set is build with the last session of each user. The remaining sessions form the training set. We also filtered items in the test set that do not belong to the training set. This partitioning allows to run the evaluation over users having different amounts of historical sessions in their profiles, hence to measure the recommendation quality over users having different degrees of activity with the system (see Section 4.3.1 for an in-depth analysis). We further partitioned the training set with the same procedure to tune the hyper-parameters of the algorithms. The characteristics of the datasets are summarized in Table 1.

4.2 Baselines and parameter tuning

We compared our $HRNN$ model against several baselines, namely Personal Pop (PPOP), Item-KNN, RNN and RNN Concat:

- Personal Pop (PPOP) recommends the item with the largest number of interactions by the user.
- Item-KNN computes an item-to-item cosine similarity based on the co-occurrence of items within sessions.
- RNN adopts the same model described in [7]. This model uses the basic GRU with a TOP1 loss function and session-parallel minibatching: sessions from the same user are fed to the RNN independently from each other.
- RNN Concat is the same as RNN, but sessions from the same user are concatenated into a single session³.

We optimize the neural models for TOP1 loss using AdaGrad [5] with momentum for 10 epochs. Increasing the number of epochs did not significantly improve the loss in all models. We used dropout regularization [18] on the hidden states of RNN and $HRNN$. We applied dropout also to GRU_{ses} initialization for $HRNN$ (Equation 4). We used single-layer GRU networks in both levels in the hierarchy, as using multiple layers did not improve the performance. In order to assess how the capacity of the network impacts the recommendation quality, in our experiments we considered both *small* networks

³We experimented with a variant of this baseline that adds a special ‘separator item’ to enforce the separation between sessions; however, it did not show better performance in our experiments.

Model	Batch size	Dropout	Learning rate/ Momentum
RNN (<i>small</i>) [†]	100	0.2	0.1/0.1
RNN (<i>large</i>) [†]	100	0.2	0.1/0.0
HRNN All (<i>small</i>)	100	0.1/0.1/0.2	0.1/0.2
HRNN All (<i>large</i>)	50	0.0/0.2/0.2	0.05/0.3
HRNN Init (<i>small</i>)	50	0.0/0.1/0.0	0.1/0.0
HRNN Init (<i>large</i>)	100	0.1/0.2/0.2	0.1/0.1

Table 2: Network hyper-parameters on XING ([†] RNN and RNN Concat share the same values).

with 100 hidden units per GRU layer and *large* networks with 500 hidden units per GRU layer.

We tuned the hyper-parameters of each model (baselines included) on the validation set using random search [2]. To help the reproducibility of our experiments, we report the hyper-parameters used in our experiments on the XING dataset in Table 2⁴. Dropout probabilities for HRNNs are relative to the user-level GRU, session-level GRU and initialization in this order. The optimal neighborhood size for Item-KNN is 300 for both datasets.

Neural models were trained on Nvidia K80 GPUs equipped with 12GB of GPU memory⁵. Training times vary from ~5 minutes for the small RNN model on XING to ~30 minutes for the large HRNN All on VIDEO. Evaluation took no longer than 2 minutes for all the experiments. We want to highlight that training times do not significantly differ between RNN and HRNNs, with HRNN All being the most computationally expensive model due to the higher complexity of its architecture (see Figure 1).

4.3 Results

We evaluate wrt. the sequential next-item prediction task, i.e. given an event of the user session, we evaluate how well the algorithm predicts the following event. All RNN-based models are fed with events in the session one after the other, and we check the rank of the item selected by the user in the next event. In addition, HRNN models and RNN Concat are *bootstrapped* with all the sessions in the user history that precede the testing one in their original order. This step slows down the evaluation but it is necessary to properly set the internal representations of the personalized models (e.g., the user-level representation for HRNNs) before evaluation starts. Notice that the evaluation metrics are still computed only over events in the test set, so evaluation remains fair. In addition, we discarded the first prediction computed by the RNN Concat baseline in each test session, since it is the only method capable of recommending the first event in the user sessions.

As recommender systems can suggest only few items at once, the relevant item should be amongst the first few items in the recommendation list. We therefore evaluate the recommendation quality in terms of Recall@5, Precision@5 and Mean Reciprocal Rank (MRR@5). In sequential next-item prediction, Recall@5 is equivalent to the hit-rate metric, and it measures the proportion of cases out of all test cases in which the relevant item is amongst

the top-5 items. This is an accurate model for certain practical scenarios where no recommendation is highlighted and their absolute order does not matter, and strongly correlates with important KPIs such as CTR [6]. Precision@5 measures the fraction of correct recommendations in the top-5 positions of each recommendation list. MRR@5 is the reciprocal rank of the relevant item, where the reciprocal rank is manually set to zero if the rank is greater than 5. MRR takes the rank of the items into account, which is important in cases where the order of recommendations matters.

Table 3 summarizes the results for both XING and VIDEO datasets. We trained each neural model for 10 times with different random seeds and report the average results⁶. We used Wilcoxon signed-rank test to assess significance of the difference between the proposed HRNN models and the state-of-the-art session-based RNN and the naïve personalization strategy used by RNN Concat.

Results on XING. On this dataset, the simple personalized popularity baseline is a very competitive method, capable of outperforming the more sophisticated Item-KNN baseline by large margins. As prior studies on the dataset have already shown, users’ activity within and across sessions has a high degree of repetitiveness. This makes the generation of ‘non-trivial’ personalized recommendations in this scenario very challenging [1].

This is further highlighted by the poor performance of session-based RNN that is always significantly worse than PPOP independently of the capacity of the network. Nevertheless, personalized session-based recommendation can overcome its limitations and achieve superior performance in terms of Recall and Precision with both small and large networks. HRNNs significantly outperform RNN Concat in terms of Recall and Precision (up to +3%/+1% with small/large networks), and provide significantly better MRR with large networks (up to +5.4% with HRNN All). Moreover, HRNNs significantly outperform the strong PPOP baseline of ~11% in Recall and Precision, while obtaining comparable MRR. This is a significant result in a domain where more trivial personalization strategies are so effective.

The comparison between the two HRNN variants does not highlight significant differences, excluded a small (~2%) advantage of HRNN All over HRNN Init in MRR. The differences in terms of Recall and Precision are not statistically significant. Having established the superiority of HRNNs over session-based recommendation and trivial concatenation, we resort to the VIDEO dataset to shed further light on the differences between the proposed personalized session-based recommendation solutions.

Results on VIDEO. The experiments on this dataset exhibit drastically different results from the XING dataset. Item-KNN baseline significantly outperforms PPOP, and session-based RNN can outperform both baselines by large margins. This is in line with past results over similar datasets [7, 8].

RNN Concat has comparable Recall and Precision with respect to session-based RNNs and interestingly significantly better MRR

⁴The source code is available at <https://github.com/mquadr/hgru4rec>.

⁵We employed Amazon EC2 *p2.xlarge* spot instances in our experiments.

⁶The random seed controls the initialization of the parameters of the network, which in turn can lead substantially different results in absolute terms. Even though, we have not observed substantial differences in the relative performances of neural models when using different random seeds.

		XING			VIDEO		
		Recall@5	MRR@5	Precision@5	Recall@5	MRR@5	Precision@5
small	Item-KNN	0.0697	0.0406	0.0139	0.4192	0.2916	0.0838
	PPOP	0.1326	0.0939	0.0265	0.3887	0.3031	0.0777
	RNN	0.1292	0.0799	0.0258	0.4639	0.3366	0.0928
	RNN Concat	<i>0.1358</i>	<i>0.0844</i>	<i>0.0272</i>	0.4682	0.3459	0.0936
	<i>HRNN All</i>	<u>0.1334[†]</u>	<u>0.0842</u>	<u>0.0267[†]</u>	<u>0.5272</u>	<u>0.3663</u>	<u>0.1054</u>
<i>HRNN Init</i>	<u>0.1337[†]</u>	<u>0.0832</u>	<u>0.0267[†]</u>	<u>0.5421</u>	<u>0.4119</u>	<u>0.1084</u>	
large	RNN	0.1317	0.0796	0.0263	0.5551	0.3886	0.1110
	RNN Concat	<i>0.1467</i>	<i>0.0878</i>	<i>0.0293</i>	0.5582	<i>0.4333</i>	<i>0.1116</i>
	<i>HRNN All</i>	<u>0.1482[†]</u>	<u>0.0925</u>	<u>0.0296[†]</u>	0.5191	0.3877	0.1038
	<i>HRNN Init</i>	<u>0.1473[†]</u>	<u>0.0901</u>	<u>0.0295[†]</u>	<u>0.5947</u>	<u>0.4433</u>	<u>0.1189</u>

Table 3: Results of Recall, MRR and Precision for $N = 5$ on the XING and VIDEO datasets. *small* networks have 100 hidden units for RNNs and 100+100 for HRNNs; *large* networks have 500 hidden units for RNNs and 500+500 for HRNNs. All the networks have statistically significant different (ssd.) results from the baselines (Wilcoxon signed-rank test $p < 0.01$). Networks ssd. from RNN are in italic; HRNNs ssd. from RNN Concat are underlined. Best values are in bold. All differences between HRNNs are significant excluded the values marked with [†] superscript.

when large networks are used. This suggest that straight concatenation does not enhance the retrieval capabilities of the RNN recommender but strengthens its ability in ranking items correctly.

However, *HRNN Init* has significantly better performance than all baselines. It significantly outperforms all baselines and RNNs (up to 6.5% better Recall and 2.3% MRR wrt. RNN Concat with large networks). In other words, also in this scenario the more complex cross-session dynamics modeled by *HRNN* provide significant advantages in the overall recommendation quality. We will investigate on the possible reasons for these results in the following sections. It is worth noting that *HRNN All* performs poorly in this scenario. We impute the context-enforcing policy used in this setting for the severe degradation of the recommendation quality. One possible explanation could be that the consumption of multimedia content (videos in our case) is a strongly session-based scenario, much stronger than the job search scenario represented in XING. Users may follow general community trends and have long-term interests that are captured by the user-level GRU. However, the user activity within a session can be totally disconnected from her more recent sessions, and even from her general interests (for example, users having a strong general interest over extreme-sport videos may occasionally watch cartoon movie trailers). *HRNN Init* models the user taste dynamics and lets the session-level GRU free to exploit them according to the actual evolution of the user interests within session. Its greater flexibility leads to superior recommendation quality.

4.3.1 Analysis on the user history length. We investigate deeper the behavior of Hierarchical RNN models. Since we expect the length of the user history to have an impact on the recommendation quality, we breakdown the evaluation by the number of sessions in the history of the user. This serves as a proxy for the ‘freshness’ of the user within the system, and allows to evaluate recommenders under different amounts of historical information about the user modeled by the user-level GRU before a session begins. To this purpose, we partitioned user histories into two groups: ‘Short’ user histories having ≤ 6 sessions and ‘Long’ user histories having 6

History length	XING	VIDEO
Short (≤ 6)	67.00%	53.69%
Long (> 6)	33.00%	46.31%

Table 4: Percentage of sessions in each history length group.

or more. The statistics on the fraction of sessions belonging to each group for both datasets are reported in Table 4. Since our goal is to measure the impact of the complex cross-session dynamics used in *HRNN* wrt. traditional RNN, we restrict these analysis to RNN-based recommenders in the *large* configuration. For each algorithm, we compute the average Recall@5 and MRR@5 per test session grouped by history length. The analysis on Precision@5 returns similar results to Recall@5, so we omit it here also for space reasons. To enhance the robustness of the experimental results, we run the evaluation 10 times with different random seeds and report the median value per algorithm. Figure 3 shows the results on XING. As the length of the user history grows, we can notice that Recall slightly increases and MRR slightly decreases in MRR for all methods, session-based RNN included. The relative performance between methods does not changes significantly between short and long user histories, with *HRNN All* being the best performing model with 12% better Recall@5 and 14-16% better MRR@5 wrt. session-based RNN. *HRNN Init* has performance comparable to RNN Concat and *HRNN All* in accordance to our previous findings.

Figure 4 shows the results on VIDEO. Recall of all methods – session-based RNN included – improves with the length of the user history. MRR instead improves with history length only for RNN Concat and *HRNN Init*. This highlights the need for effective personalization strategies to obtain superior recommendation quality at session-level for users that heavily utilize the system. Moreover, the performance gain of *HRNN Init* wrt. session-based RNN grows from 5%/12% (short) to 7%/19% (long) in Recall@5/MRR@5, further highlighting the quality of our personalization strategy. Coherently with our previous findings, *HRNN All* does not perform well in




Figure 3: Recall@5 and MRR@5 on XING grouped by user history length.




Figure 4: Recall@5 and MRR@5 on VIDEO grouped by user history length.

this scenario, and its performances are steady (or even decrease) between the two groups.

In summary, the length of the user history has a significant impact on the recommendation quality as expected. In loosely session-bounded domains like XING, in which the user activity is highly repetitive and less diverse across sessions, enforcing the user representation in input at session-level provides slightly better performance over the simpler initialization-only approach. However, in the more strongly session-based scenario, in which the user activity across sessions has a higher degree of variability and may significantly diverge from the user historical interest and tastes, the simpler and more efficient *HRNN Init* variant has significantly better recommendation quality.

4.3.2 Analysis within sessions. Here we breakdown by number of events within the session in order to measure the impact of personalization within the user session. We limit the analysis to sessions having length ≥ 5 (6,736 sessions for XING and 8,254 for VIDEO). We compute the average values of each metric grouped by position within the session (Beginning, Middle and End). Beginning refers to the first 2 events of the session, Middle to the 3rd and 4th, and End to any event after the 4th. As in the previous analysis, we focus on RNN-based models in the *large* configuration and report the median of the averages values of each metric computed over 10 runs with different random seeds. Results for Recall@5 and MRR@5 are shown in Figure 5 and Figure 6 for XING and VIDEO respectively.

On XING, the performance of all methods increases with the number of previous items in the session, suggesting that the user context at session-level is properly leveraged by all RNN-based




Figure 5: Recall@5 and MRR@5 on XING for different positions within session.




Figure 6: Recall@5 and MRR@5 on VIDEO for different positions within session.

models. However, there is a wide margin between personalized and ‘pure’ session-based models. Both *HRNNs* have similar Recall@5 and are comparable to RNN Concat. Interestingly, the gain in MRR@5 of *HRNN All* wrt. to both RNN and RNN Concat grows with the number of items processed, meaning that in this scenario the historical user information becomes more useful as the user session continues. *HRNN Init* has constantly better MRR than RNN Concat, with wider margins at the beginning and at the end of the session.

On VIDEO, the behavior within session is different. We can notice that both Recall and MRR increase between the beginning and end of a session as expected. *HRNN Init* exhibits a large improvement over RNN and RNN Concat at the beginning of the session (up to 10% better Recall and 25% better MRR). This conforms with the intuition that past user activity can be effectively used to predict the first actions of the user in the forthcoming session with greater accuracy. After the first few events the gain in Recall of personalized over pure session-based models reduces, while the gain in MRR stays stable. In other words, after a few events, the session-level dynamics start to prevail over longer-term user interest dynamics, making personalization strategies less effective. However, personalization still provides superior ranking quality all over the session, as testified by the higher MRR of both *HRNN Init* and RNN Concat over RNN. It is important to notice that better recommendations at beginning of a session are more impactful than later in the session because they are more likely to increase the chances of retaining the user. Finally, *HRNN All* is always the worse method, further underpinning the superiority of the *HRNN Init* variant.

	VIDEOXXL		
	Recall@5	MRR@5	Precision@5
RNN	0.3415	0.2314	0.0683
RNN Concat	0.3459	0.2368	0.0692
<i>HRNN All</i>	0.3621	0.2658	0.0724
<i>HRNN Init</i>	0.4362	0.3261	0.0872

Table 5: Results on the VIDEOXXL dataset for small networks. Best values are in bold.

4.3.3 *Experiments on a large-scale dataset.* We validated *HRNNs* over a larger version of the VIDEO dataset used in the previous experiments. This dataset is composed by the interactions of 810k users on 380k videos over the same 2 months periods, for a total of 33M events and 8.5M sessions. We then applied the same preprocessing steps used for the VIDEO dataset. We named this dataset VIDEOXXL. Due to our limited computational resources, we could only test small networks (100 hidden units for RNN and for 100+100 *HRNNs*) on this large-scale dataset. We run all RNNs and *HRNNs* once using the same hyper-parameters learned on the small VIDEO dataset. Although not optimal, this approach provides a first approximation on the applicability of our solution under a more general setting. For the same reason, we do not provide an exhaustive analysis on the experimental results as done for the smaller datasets. To speed up evaluation, we computed the rank of the relevant item compared to the 50,000 most supported items, as done in [8].

Results are summarized in Table 5 and confirm our previous findings on the small VIDEO dataset. RNN Concat is not effective and performs similarly to session-based RNN. On the other hand, Hierarchical RNNs outperform session-based RNN by healthy margins. In particular, *HRNN Init* outperforms session-based RNN by ~28% in Recall@5 and by ~41% in MRR@5. These results further confirm the effectiveness of the *HRNN* models presented in this paper, and further underpin the superiority of *HRNN Init* over the alternative approaches for personalized session-based recommendation.

5 CONCLUSIONS AND FUTURE WORK

In this paper we addressed the problem of personalizing session-based recommendation by proposing a model based Hierarchical RNN, that extends previous RNN-based session modeling with one additional GRU level that models the user activity across sessions and the evolution of her interests over time. *HRNNs* provide a seamless way of transferring the knowledge acquired on the long-term dynamics of the user interest to session-level, and hence to provide personalized session-based recommendations to returning users. The proposed *HRNNs* model significantly outperform both state-of-the-art session-based RNNs and the other basic personalization strategies for session-based recommendation on two real-world datasets having different nature. In particular, we noticed that the simpler approach that only initializes the session-level representation with the evolving representation of the user (*HRNN Init*) gives the best results. We delved into the dynamics of session-based RNN models within and across sessions, providing extensive evidences of the superiority of the proposed *HRNN Init* approach and setting

new state-of-the-art performances for personalized session-based recommendation.

As future works, we plan to study attention models and the usage of item and user features as ways of refining user representations and improve session-based recommendation even further. We also plan to investigate personalized session-based models in other domains, such as music recommendation, e-commerce and online advertisement.

REFERENCES

- [1] *RecSys Challenge '16: Proceedings of the Recommender Systems Challenge*, New York, NY, USA, 2016. ACM.
- [2] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *J. Mach. Learn. Res.*, 13:281–305, February 2012.
- [3] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder–decoder approaches. In *SSST-8: 8th Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, 2014.
- [4] Robin Devooght and Hugues Bersini. Collaborative filtering with recurrent neural networks. *arXiv preprint arXiv:1608.07400*, 2016.
- [5] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [6] B. Hidasi and D. Tikk. Fast ALS-based tensor factorization for context-aware recommendation from implicit feedback. In *ECML-PKDD'12, Part II*, number 7524 in LNCS, pages 67–82. Springer, 2012.
- [7] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. Session-based recommendations with recurrent neural networks. *CoRR*, abs/1511.06939, 2015.
- [8] Balázs Hidasi, Massimo Quadrana, Alexandros Karatzoglou, and Domonkos Tikk. Parallel recurrent neural network architectures for feature-rich session-based recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, RecSys '16, pages 241–248, New York, NY, USA, 2016. ACM.
- [9] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [10] Noam Koenigstein and Yehuda Koren. Towards scalable and accurate item-oriented recommendations. In *Proceedings of the 7th ACM Conference on Recommender Systems*, RecSys '13, pages 419–422, New York, NY, USA, 2013. ACM.
- [11] G. Linden, B. Smith, and J. York. Amazon.com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, 7(1):76–80, 2003.
- [12] Zachary C Lipton, John Berkowitz, and Charles Elkan. A critical review of recurrent neural networks for sequence learning. *CoRR*, 1506.00019, 2015.
- [13] Steffen Rendle and Christoph Freudenthaler. Improving pairwise learning for item recommendation from implicit feedback. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*, WSDM '14, pages 273–282, New York, NY, USA, 2014. ACM.
- [14] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI '09, pages 452–461, Arlington, Virginia, United States, 2009. AUAI Press.
- [15] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. *Neurocomputing: Foundations of research*. chapter Learning Internal Representations by Error Propagation, pages 673–695. MIT Press, Cambridge, MA, USA, 1988.
- [16] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *WWW'01: 10th Int. Conf. on World Wide Web*, pages 285–295, 2001.
- [17] Alessandro Sordani, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, pages 553–562. ACM, 2015.
- [18] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, January 2014.
- [19] Yong Kiam Tan, Xinxing Xu, and Yong Liu. Improved recurrent neural networks for session-based recommendations. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, DLRS 2016, pages 17–22, New York, NY, USA, 2016. ACM.
- [20] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J. Smola, and How Jing. Recurrent recommender networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, WSDM '17, pages 495–503, New York, NY, USA, 2017. ACM.