中国科学院深圳先进技术研究院
SHENZHEN INSTITUTES OF ADVANCED TECHNOLOGY
CHINESE ACADEMY OF SCIENCES

AAAI 2021

A User-Adaptive Layer Selection Framework for Very Deep Sequential Recommender Models

Lei Chen, Fajie Yuan, Jiaxi Yang, Xiang Ao, Chengming Li, Min Yang

◆ Contents

➢ Abstract

- **SRS:** Sequential Recommender Systems

- Very deep recommender models

  →**Obstacle:** high network latency

- **Argument:** Treating all users equally during the inference phase is inefficient in running time, as well as sub-optimal in accuracy.

Abstract

- **SkipRec:** An adaptive inference framework by learning to skip inactive hidden layers on a per-user basis.
  - ✓ Devising a **policy network** to automatically determine which layers should be retained and which layers should be skipped.
  - ✓ So as to achieve user-specific decisions.
- **To derive the optimal skipping policy:** Using **gumbel softmax** and **reinforcement learning** to solve the non-differentiable problem during backpropagation.

- Extensive experimental results show that SkipRec attains comparable or better accuracy with much less inference time.

➢ Introduction

◆ **Standard CF vs Sequential recommender models**

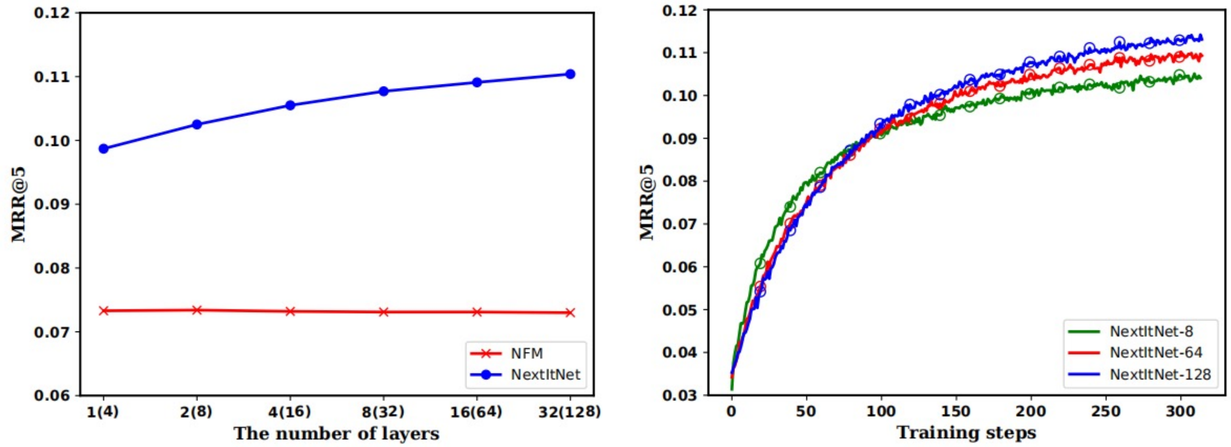

(a) (b)

◆ Sequential recommender models enjoy excellent model expressivity by stacking **very deep** layers.

◆ **A real problem arising:** The inference cost largely increases, leading to an inevitable time delay for online service.

➢ Introduction

- Users in recommender systems are **unique** (i.e., **personalized**) !
  - Passing all of them through hidden layers of the same depth is non-optimal and computationally inefficient.

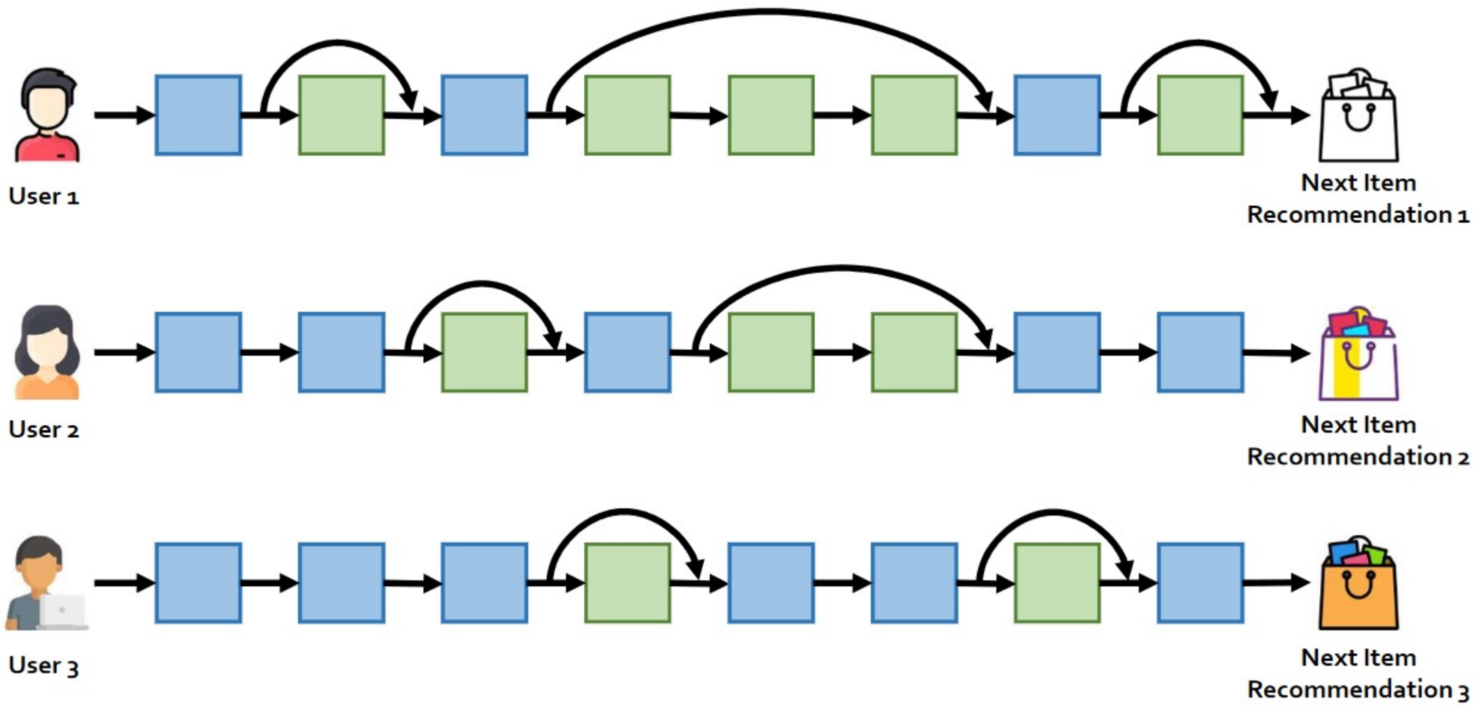

➢ Introduction

- **SkipRec:** An adaptive inference framework for deep sequential recommender models, which defines the network structure adaptively on a per-user basis.
  - Devising a **policy network** to automatically determine which layers in the backbone network should be retained and which layers are skipped, so as to obtain the user-specific decision in SRS.

- **SkipRec** is a **general network depth selection framework** which directly applies to a broad range of deep recommendation models.

➢ Introduction

◆ Contributions:

- ✓ We are the first to emphasize **the unnecessity in executing the same number of hidden layers for all users** in a deep recommender model.
- ✓ We propose **SkipRec**, **a user-specific depth selection framework** where the number of network layers can be selected on a per-user basis. SkipRec enables each user to have their own skipping policies, which is the first personalized depth selection method for the recommendation task.

➢ Introduction

◆ Contributions:

- ✓ We propose **SkipRec-Gumbel** and **SkipRec-RL** to derive the optimal skipping policy (skipping or retaining) without suffering from the non-differentiable problem during backpropagation.
- ✓ Extensive experiments show that the proposed **SkipRec** attains **competitive or better performance with less inference time** in three real-world SRS datasets.

Related work

- **Conventional SRS**
  - Markov Chain (MC)
  - Factorization models
  - ……

- **Deep learning (DL) based SRS**
  - **RNN-based:** GRU4Rec
  - **CNN-based:** Caser, NextItNet
  - **Self-attention based:** SASRec, BERT4Rec
  - ……

➢ Methodology

◆ Problem definition

- **Given:** a sequence of historical user behaviors $X^{\mathrm{u}} = [x_1^u, x_2^u, \dots, x_t^u]$.
- **Output:** predicted item x_{t+1}^u that the user will interact with at next time.

- **Next-item recommendation**
- **Top-N recommendation**

➢ Methodology

◆ SkipRec: The overall architecture

- ➢ Backbone network
- ➢ Policy network

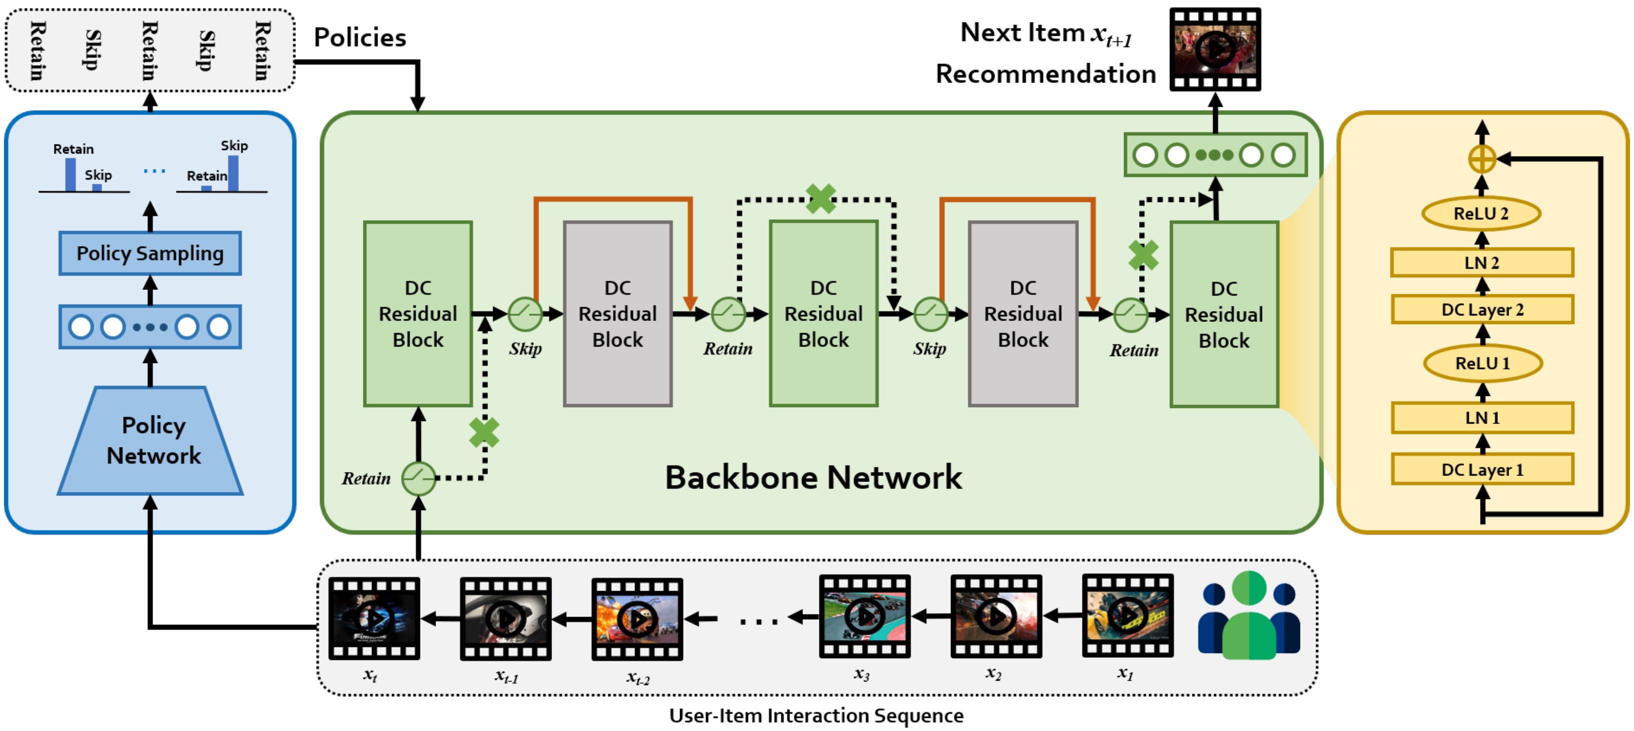

➢ Methodology

◆ Backbone network: NextItNet

- a stack of **dilated convolutional layers**, which are wrapped by a **residual block** structure every two layers.

$$\mathbf{E}_l^u = \lambda \times \mathcal{F}_l(\mathbf{E}_{l-1}^u) + \mathbf{E}_{l-1}^u$$

$$\mathcal{F}_l(\mathbf{E}_{l-1}^u) = \sigma\left(\mathbf{LN}_2\left(\psi_2\left(\sigma\left(\mathbf{LN}_1\left(\psi_1(\mathbf{E}_{l-1}^u)\right)\right)\right)\right)\right)$$

$$p(x_{t+1}^u | x_{1:t}^u) = \mathrm{softmax}(\mathbf{W}\mathbf{E}_l^u + \mathbf{b})$$

$$p\left(X^u; \Omega\right) = \prod_{i=2}^{t} p\left(x_i^u | x_{1:i-1}^u, \theta\right) p\left(x_1^u\right)$$

➢ Methodology

◆ Policy network: Lightweight NextItNet model

- ➢ Without any restrictions, the policy network can also be implemented with **any deep neural networks**, e.g., a **RNN** model.

$$\mathbf{E}_{l\ adaptive}^{u} = \mathbf{I}_l(\mathbf{E}^u)\mathbf{E}_l^u + (1 - \mathbf{I}_l(\mathbf{E}^u))\,\mathbf{E}_{l-1\ adaptive}^{u}$$

- ➢ $\mathbf{I}_l(\mathbf{E}^u)$ is a binary variable indicating whether the residual block $\mathbf{E}_l^u$ could be retained or skipped based on user sequence $\mathbf{E}^u$.

➢ Methodology

◆ Policy network

- ➢ **Gumbeling Softmax Sampling**
- ➢ **Reinforcement learning**

➢ SkipRec-Gumbel

$$\mathbf{z} = \text{one\_hot}\left(\arg\max_i \left[g_i + \log \pi_i\right]\right)$$

$$\alpha_i = \frac{\exp\left(\left(\log \pi_i + g_i\right)/\tau\right)}{\sum_{j=1}^{k} \exp\left(\left(\log \pi_j + g_j\right)/\tau\right)} \quad \text{for } i = 1, \ldots, k$$

➢ Methodology

➢ SkipRec-RL

☐ Reward function

$$\mathbf{R}(\mathcal{A}) = \begin{cases} 1 - \left(\sum_l \mathbb{1}(\mathcal{A}_l)/N\right)^2 & \text{if correct} \\ -\gamma & \text{otherwise} \end{cases}$$

☐ Self-critical sequence training (SCST)

$$\mathcal{L}_{RL} = -\sum_{l=1}^{N} \log p(\mathcal{A}_l^s)\left(\mathbf{R}(\mathcal{A}_l^s) - \mathbf{R}(\hat{\mathcal{A}}_l)\right)$$

$$\mathcal{L} = \mathcal{L}_{CE} + \beta \mathcal{L}_{RL}$$

- **Training Procedure: One-stage and Two-stage.**
  - **One-stage training:** the parameters of the policy network and the backbone network are initialized randomly and trained jointly.
  - **Two-stage training:** we first pre-train the NextItNet model with training data and initialize the backbone network with pre-trained parameters, which can help the backbone network obtain better feature representation ability at the beginning and make the training of the policy network more conducive.

➢ Experiments

◆ Experimental setup

- Datasets: ML20, ML100 and Weishi.

Dataset	#items	#interactions	#sequences	Length t
Weishi	66K	10M	1,048,575	10
ML20	54K	27.7M	1,491,478	20
ML100	54K	27.7M	457,350	100

- Baselines: NFM, GRU4Rec, Caser, NextItNet and NextItNet+.

- Evaluation metrics: HR@5, HR@20, MRR@5, MRR@20, Inference speedup.

◆ Experimental results

➢ Quantitative Evaluation

Model	Weishi					ML20					ML100				
	MRR@5	MRR@20	HR@5	HR@20	Speedup	MRR@5	MRR@20	HR@5	HR@20	Speedup	MRR@5	MRR@20	HR@5	HR@20	Speedup
MostPop	0.0050	0.0121	0.0187	0.0940	\	0.0044	0.0076	0.0134	0.0485	\	0.0040	0.0068	0.0124	0.0433	\
NFM	0.0734	0.0876	0.1307	0.2774	\	0.0483	0.0585	0.0856	0.1933	\	0.0341	0.0402	0.0573	0.1221	\
GRU4Rec	0.1001	0.1148	0.1649	0.3216	\	0.0938	0.1082	0.1577	0.3115	\	0.0973	0.1120	0.1611	0.3117	\
Caser	0.0911	0.1051	0.1498	0.2973	\	0.0915	0.1045	0.1509	0.2857	\	0.0927	0.1058	0.1505	0.2858	\
NextItNet	0.1025	0.1175	0.1669	0.3234	\	0.1019	0.1173	0.1686	0.3285	\	0.1073	0.1226	0.1752	0.3339	\
NextItNet+	0.1091	0.1243	0.1767	0.3354	1.00×	0.1067	0.1234	0.1784	0.3393	1.00×	0.1117	0.1276	0.1819	0.3465	1.00×
SkipRec-Gumbel	0.1105	0.1264	0.1796	0.3442	1.66×	0.1091	0.1249	0.1783	0.3424	1.92×	0.1122	0.1281	0.1827	0.3474	1.35×
Gumbel w/ pre	**0.1121**	**0.1279**	**0.1826**	0.3453	1.74×	**0.1104**	**0.1261**	**0.1800**	**0.3431**	1.72×	**0.1147**	**0.1306**	**0.1862**	**0.3511**	1.30×
SkipRec-RL	0.1098	0.1255	0.1788	0.3426	2.00×	0.1087	0.1243	0.1773	0.3402	2.28×	0.1111	0.1270	0.1799	0.3455	2.08×
RL w/ pre	0.1117	0.1275	0.1813	**0.3460**	2.01×	0.1101	0.1258	0.1793	0.3413	2.30×	0.1136	0.1292	0.1852	0.3488	2.03×

- ☐ **NextItNet+** and **NextItNet** perform better than NFM, GRU4Rec and Caser with substantial improvements in recommendation accuracy.
- ☐ **SkipRec-Gumbel** and **SkipRec-RL** attain comparable or better recommendation accuracy and notable inference speedup especially when we pre-train the backbone network in advance.

➢ Ablation Studies on the Policy Network

□ Random Policy

Data	Model	MRR@5	MRR@20	HR@5	HR@20
Weishi	NextItNet+	0.1091	0.1243	0.1767	0.3354
	SkipRec-Gumbel	**0.1121**	**0.1279**	**0.1826**	0.3453
	SkipRec-RL	0.1117	0.1275	0.1813	**0.3460**
	Random policy	0.1056	0.1212	0.1727	0.3346
ML20	NextItNet+	0.1067	0.1234	0.1784	0.3393
	Skip-Gumbel	**0.1104**	**0.1261**	**0.1800**	**0.3431**
	SkipRec-RL	0.1101	0.1258	0.1793	0.3413
	Random policy	0.1018	0.1172	0.1676	0.3277

□ A well-optimized policy largely performs better than a random policy.

➢ Experiments

➢ Ablation Studies on the Policy Network

◻ Policy Network with GRU

Data	Model	MRR@5	MRR@20	HR@5	HR@20
Weishi	NextItNet+	0.1091	0.1243	0.1767	0.3354
	SkipRec-Gumbel	0.1121	0.1279	**0.1826**	0.3453
	SkipRec-Gumbel-GRU	**0.1123**	**0.1284**	0.1816	**0.3473**
	SkipRec-RL	0.1117	0.1275	0.1813	0.3460
	SkipRec-RL-GRU	0.1112	0.1267	0.1801	0.3413
ML20	NextItNet+	0.1067	0.1234	0.1784	0.3393
	SkipRec-Gumbel	**0.1104**	**0.1261**	0.1800	**0.3431**
	SkipRec-Gumbel-GRU	0.1100	0.1255	**0.1818**	0.3426
	SkipRec-RL	0.1101	0.1258	0.1793	0.3413
	SkipRec-RL-GRU	0.1096	0.1254	0.1780	0.3414

◻ Designing the policy network by a Gated Recurrent Unit (GRU) can also achieve comparable performance.

➢ Experiments

➢ Ablation Studies on the Policy Network

◻ Visualization of Policies

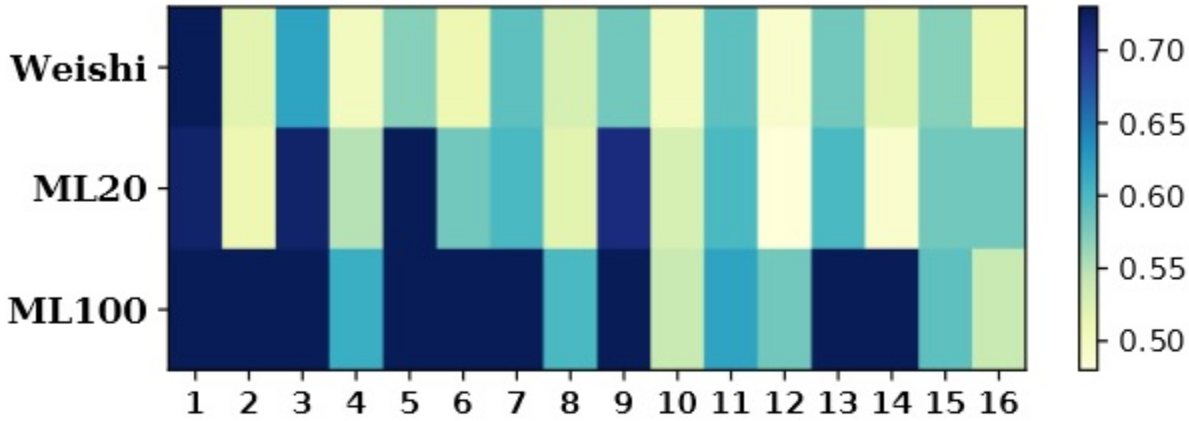

- ◻ Different datasets have different skipping policies. SkipRec allows NextItNet+ to automatically identify the right policy in determining which layers in the backbone network should be retained and which layers should be skipped on a per-user basis.

➢ Experiments

➢ Convergence Behavior Analysis

- SkipRec-Gumbel and SkipRec-RL prevent overfitting better than the NextItNet+.
- NextItNet+ converges faster than SkipRec given that more parameters are updated in each round of backpropagation.

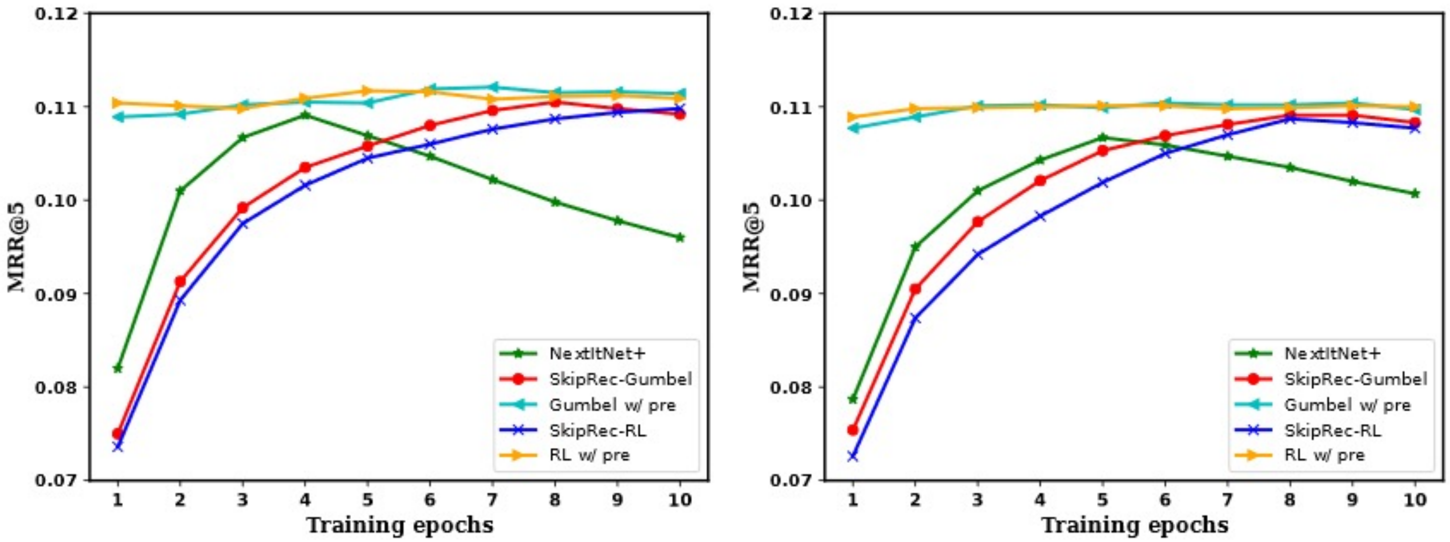

Experiments

Adaptability Experiment

- To verify the generality of SkipRec, we specify it with SASRec. Similar conclusions can be made as specified with NextItNet.

Data	Model	MRR@5	MRR@20	HR@5	HR@20	Speedup
Weishi	SASRec+	0.1076	0.1225	0.1749	0.3298	1.00×
	SkipRec-Gumbel	0.1089	0.1246	0.1770	0.3397	1.52×
	Gumbel w/ pre	**0.1102**	**0.1258**	**0.1778**	**0.3411**	1.37×
	SkipRec-RL	0.1086	0.1240	0.1761	0.3384	1.46×
	RL w/ pre	0.1094	0.1250	0.1768	0.3388	1.58×
ML20	SASRec+	0.1154	0.1310	0.1890	0.3498	1.00×
	SkipRec-Gumbel	0.1151	0.1307	0.1872	0.3500	1.57×
	Gumbel w/ pre	**0.1189**	**0.1346**	**0.1910**	**0.3538**	1.31×
	SkipRec-RL	0.1147	0.1304	0.1861	0.3493	1.69×
	RL w/ pre	0.1175	0.1333	0.1879	0.3524	1.53×

➢ Conclusion

◆ Conclusion

- ✓ We have proposed **an adaptive layer selection framework (SkipRec)**, whereby the number of network layers can be selected on a per-user basis.
- ✓ We devise a **policy network** to automatically determine which layers should be retained and which layers should be skipped in the backbone network.
- ✓ We expect our studies will inspire new research in exploring **deep, effective and efficient recommender models**.

◆ Future work

- A small drawback of SkipRec is that it **converges slightly slower** than the backbone network due to fewer parameters are trained in each round.
- In the future, we would explore advanced techniques to speedup the training of such deep and large recommender models so as to free up more computational resources.

Thank You

中国科学院深圳先进技术研究院

SHENZHEN INSTITUTES OF ADVANCED TECHNOLOGY
CHINESE ACADEMY OF SCIENCES