User Behavior Modeling with Deep Learning for Recommendation: Recent Advances

NOAH'S ARK LAB

Presenter Bio

Yong Liu
Researcher
Huawei Noah's Ark Lab

Weiwen Liu
Researcher
Huawei Noah's Ark Lab

Wei Guo
Researcher
Huawei Noah's Ark Lab

Kefan Wang
M.D. Student
University of Science
and Technology of China

NOAH'S ARK LAB

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

User Behavior Modeling

- In the digital age, **user-generated data** is rapidly accumulating, providing valuable information resources for businesses, research institutions, and government agencies.

- The proliferation of user-generated data has given rise to a new challenge – **Information Overload**.

- Understanding and modeling user behavior has become a critical challenge and opportunity in the fields of **information science and business**.

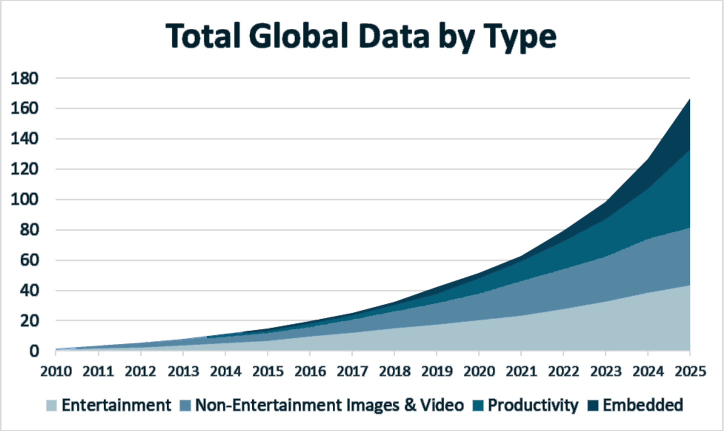

Data Accumulation

User Behavior Modeling

- In the digital age, **user-generated data** is rapidly accumulating, providing valuable information resources for businesses, research institutions, and government agencies.

- The proliferation of user-generated data has given rise to a new challenge – **Information Overload**.

- Understanding and modeling user behavior has become a critical challenge and opportunity in the fields of **information science and business**.

information overload

User Behavior Modeling

- In the digital age, **user-generated data** is rapidly accumulating, providing valuable information resources for businesses, research institutions, and government agencies.

- The proliferation of user-generated data has given rise to a new challenge – **Information Overload**.

- Understanding and modeling user behavior has become a critical challenge and opportunity in the fields of **information science and business**.

User Modeling

User Behavior Modeling

- **User behavior modeling (UBM)** extracts personalized interests from user behavior history, which is key to recommender systems.
- In recent years, a series of advanced techniques have been increasingly employed in user behavior modeling, including RNN, Attention, GNN and CNN.

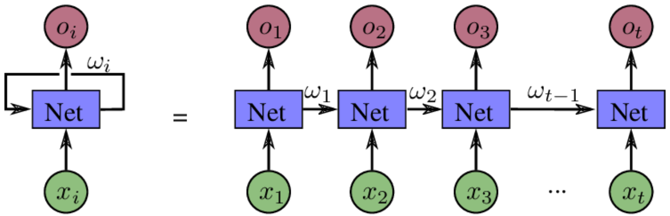

RNN

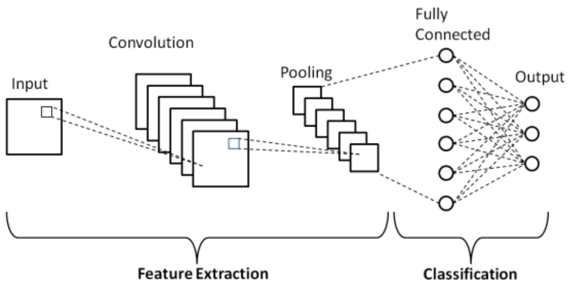

CNN

User Behavior Modeling

- The goal of UBM is to learn a function $F_{\Theta}^{UBM}(\cdot)$ for predicting how likely a user u will be interested in an item i.

  $p(u,i) = F_{\Theta}^{UBM}(u,i,H_u,f_u,f_i,f_c), \forall u \in U, i \in I.$

- Each behavior record $b_{u,k} = \{v_k, t_k, f_k\}, \forall b_{u,k} \in H_u$, consists of interacted item $v_k \in U$, the timestamp t_k, and the related features f_k.

- Three development trends in recent years
  - **Longer length**
  - **Increased multiplicity**
  - **Growing heterogeneity**

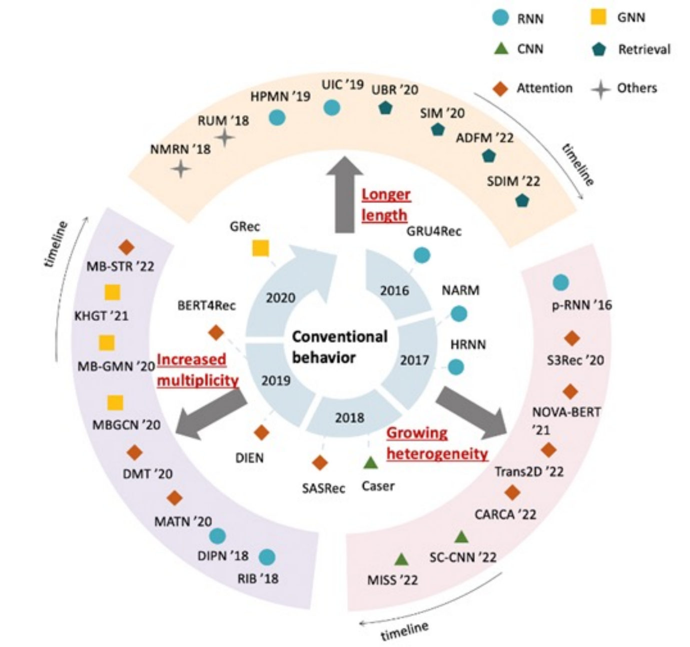

Development trends

User Behavior Modeling

- **Taxonomy**

- Conventional UBM
- Long-Sequence UBM
- Multi-Type UBM
- UBM with Side Information
- Industrial Practices
- Future Prospects

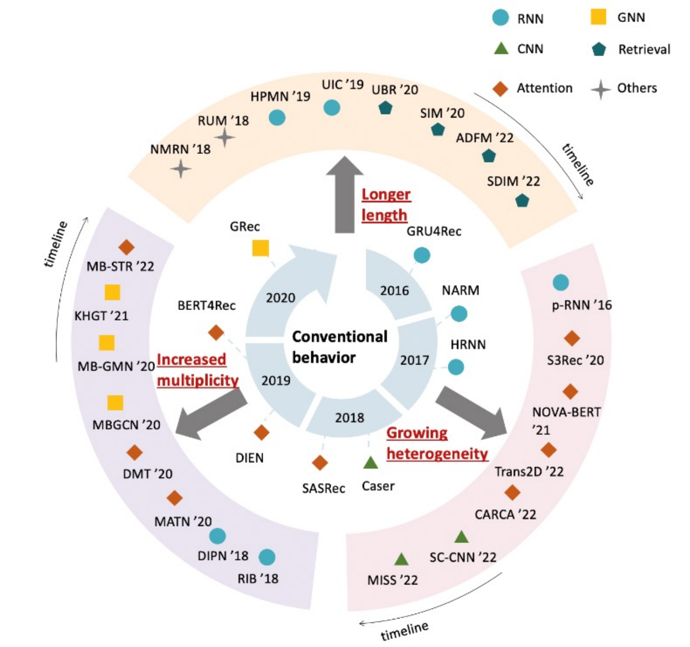

Development trends

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention, GNN
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Conventional UBM

- Learn user interests from **simple** historical behavior sequences, usually with **limited length**.
- Formulated as

$$p(u,i) = F_{\Theta}^{UBM}(u, i, H_u^S), \forall u \in U, i \in I.$$

- Various deep network structures have been adopted to learn from H_u.

RNN-based Methods

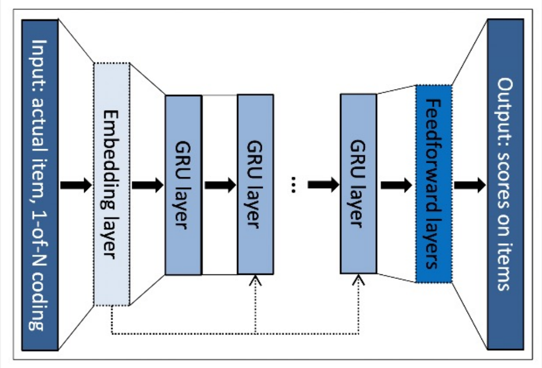

GRU4Rec [ICLR '16]

CNN-based Methods

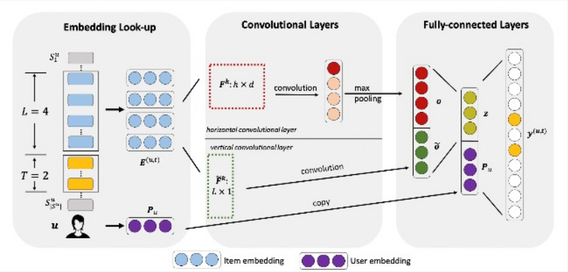

Caser [WSDM '18]

Attention-based Methods

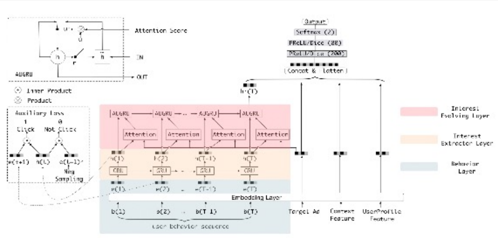

DIN [AAAI '19]

Hidasi, Balázs, et al. "Session-based recommendations with recurrent neural networks." *arXiv preprint arXiv:1511.06939* (2015).
Tang, Jiaxi, and Ke Wang. "Personalized top-n sequential recommendation via convolutional sequence embedding." *Proceedings of the eleventh ACM international conference on web search and data mining*. 2018.
Zhou, Guorui, et al. "Deep interest network for click-through rate prediction." *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2018.

Conventional UBM

- **RNN-based Methods**

- Capture information from previous time steps and utilize it in subsequent time steps.
- Be employed to capture users' long and short-term interests.
- Limitation: Complex gating mechanisms increase the computational load.

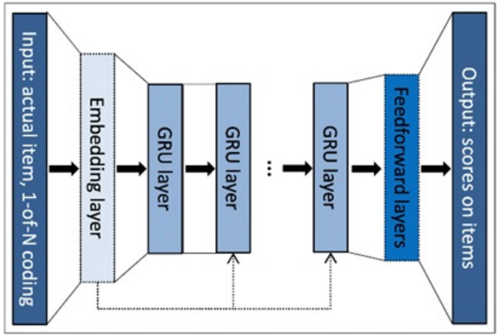

GRU4Rec [ICLR '16]

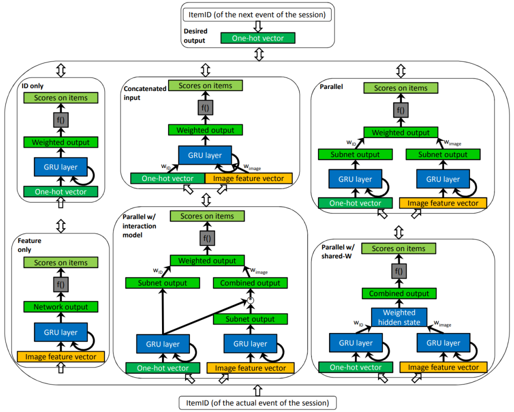

p-RNN [Recsys '16]

Hidasi, Balázs, et al. "Session-based recommendations with recurrent neural networks." *arXiv preprint arXiv:1511.06939* (2015).
Hidasi, Balázs, et al. "Parallel recurrent neural network architectures for feature-rich session-based recommendations." *Proceedings of the 10th ACM conference on recommender systems*. 2016.

Conventional UBM

▪ CNN-based Methods

- How to capture the skip behaviors well, where the next step is influenced by the behaviors a few steps earlier.
- Limitation: Because of the size limitations of CNN's filters, it is difficult for CNN to capture global information and long–term dependencies.

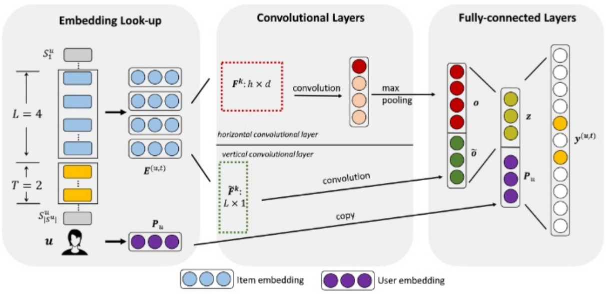

Caser [WSDM ‘18]

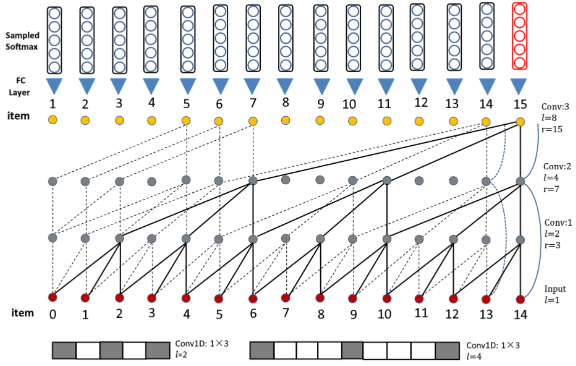

NextItNet [WSDM ‘19]

Tang, Jiaxi, and Ke Wang. "Personalized top-n sequential recommendation via convolutional sequence embedding." *Proceedings of the eleventh ACM international conference on web search and data mining*. 2018.

Yuan, Fajie, et al. "A simple convolutional generative network for next item recommendation." *Proceedings of the twelfth ACM international conference on web search and data mining*. 2019.

Conventional UBM

▪ Attention-based Methods

- Modeling interactions between any pair of behaviors, without degradation over the encoding.
- Avoid performance degradation caused by a large distance.

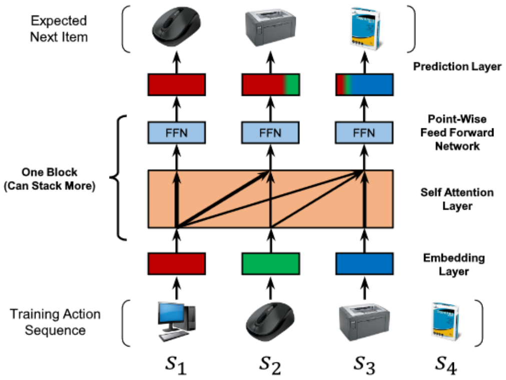

SASRec [ICDM '18]

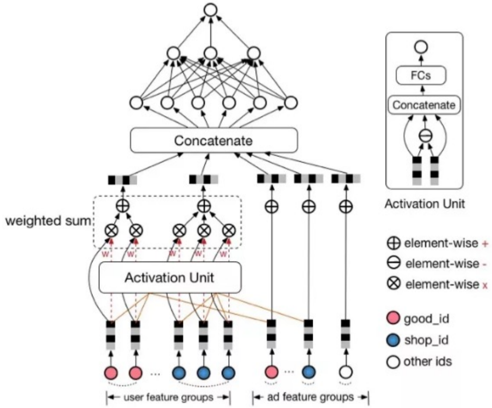

DIN [KDD '18]

Kang, Wang-Cheng, and Julian McAuley. "Self-attentive sequential recommendation." *2018 IEEE international conference on data mining (ICDM)*. IEEE, 2018.
Zhou, Guorui, et al. "Deep interest network for click-through rate prediction." *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2018.

Conventional UBM

- **GNN-based Methods**
- Using the properties of graphs to learn the structural characteristics of interactions.
- By introducing graph contrastive learning, we can enhance the data from a global perspective.

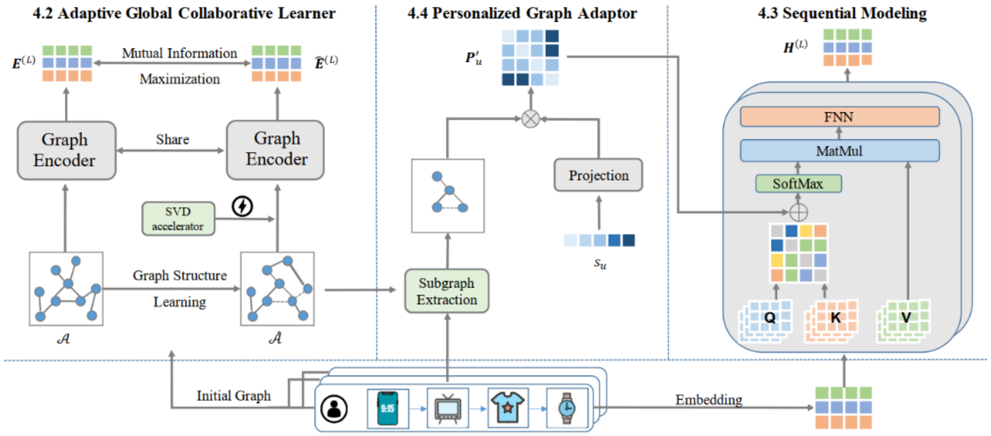

APGL4SR

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Long-Sequence UBM

- Learn user interests from **long** historical behavior sequences (as least in thousands).
- Formulated as

$$p(u,i) = F_{\Theta}^{UBM}(u,i,H_u^L),\forall u \in U, i \in I.$$

- Enables to take advantage of **long-term** behavior dependencies and the **periodicity** of user behaviors.
- Longer sequences may contain more noise and require longer inference time.

Memory-augmented Methods

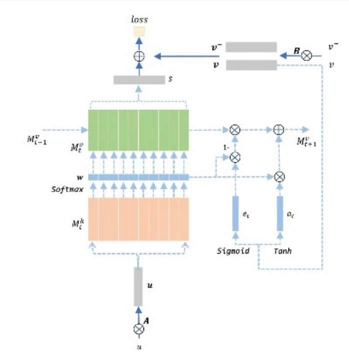

NARM [KDD '18]

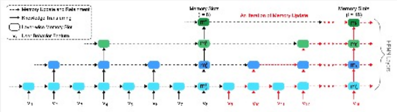

HPMN [SIGIR '19]

Retrieval Methods

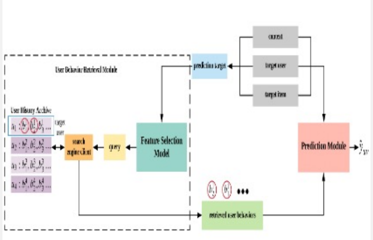

UBR [SIGIR '20]

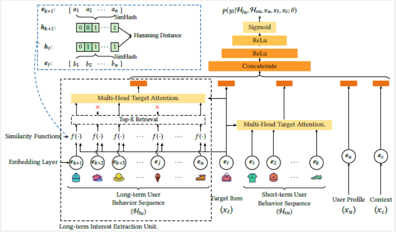

ETA [Arxiv '21]

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Memory-augmented Methods

- **Hierarchical Periodical Memory Network, HPMN**

➢ User behavior sequences vary in length and there exist extremely long sequences.
➢ Dynamic and multi-facet user interests.

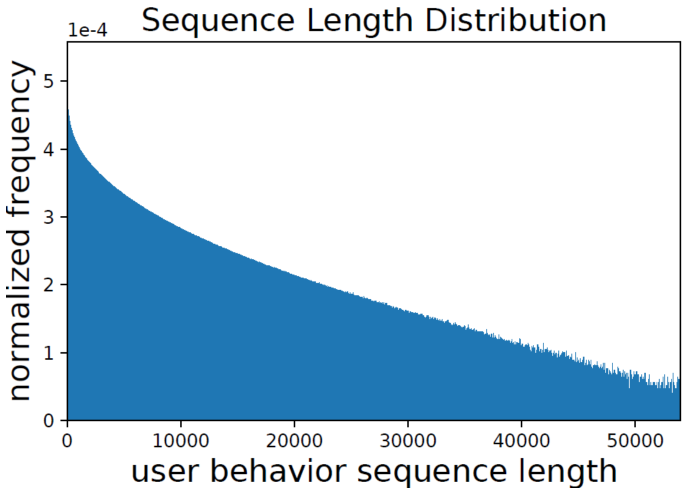

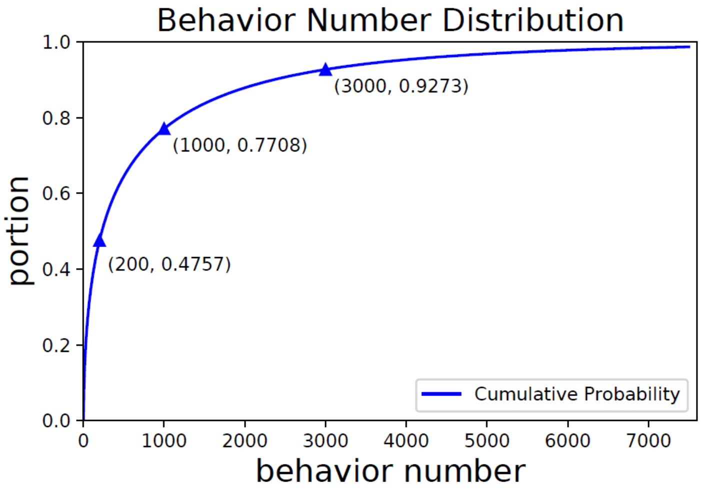

Ren, Kan, et al. "Lifelong sequential modeling with personalized memorization for user response prediction." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

Memory-augmented Methods

- **Hierarchical Periodical Memory Network, HPMN**

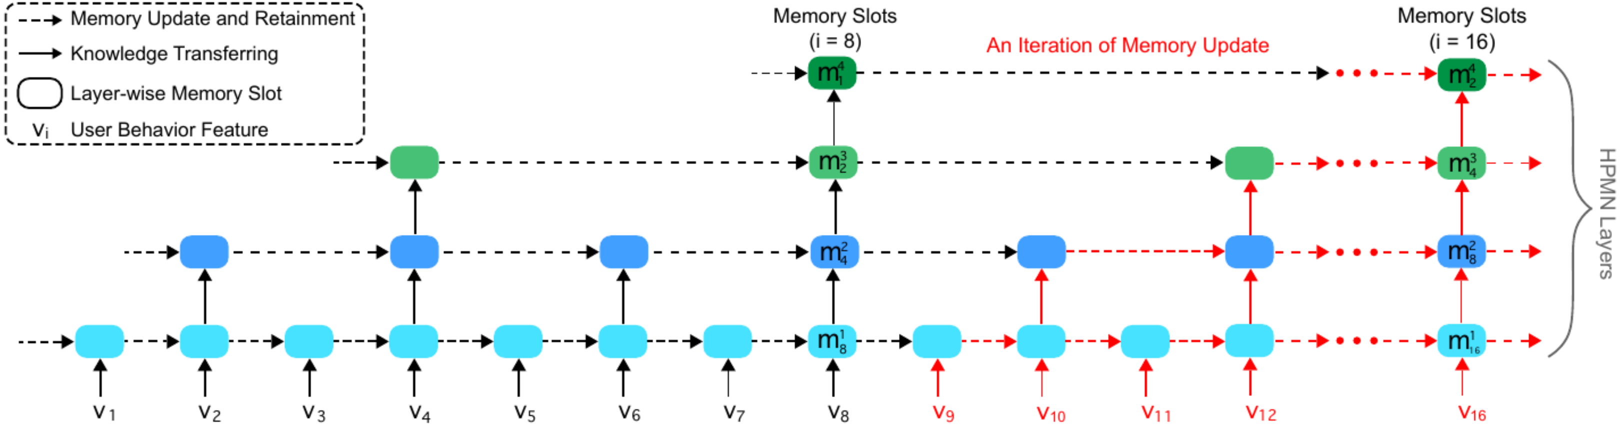

Figure 3: The framework of HPMN model with four layers maintaining user memory in four ($D = 4$) memory slots. The update period t^j of j-th layer follows an exponential sequence $\{2^{j-1}\}_{j=1}^{D}$ as an example. The red part means the incremental updating mechanism; the dotted line means the periodic memorization and forgetting.

Ren, Kan, et al. "Lifelong sequential modeling with personalized memorization for user response prediction." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

Memory-augmented Methods

- **Hierarchical Periodical Memory Network, HPMN**
  - The content in the j-th memory slot at step i
    - $\{\boldsymbol{m}_i^j\}_{j=1}^D$
  - Memory query and attentional reading
    - Given the query vector of the target item $\boldsymbol{v}$
    - Calculate the attention weight $w^j = E(\boldsymbol{m}^j, \boldsymbol{v})$ for each j-th memory slot
    - User representation $\boldsymbol{r} = \sum_j^D w^j \cdot \boldsymbol{m}^j$ at step i
  - Periodical and gate-based (soft) writing

$$\boldsymbol{m}_i^j = \begin{cases} g^j\left(\boldsymbol{m}_i^{j-1}, \boldsymbol{m}_{i-1}^j\right) & \text{if } i \bmod t^j = 0\,, \\ \boldsymbol{m}_{i-1}^j & \text{otherwise}\,, \end{cases}$$

Ren, Kan, et al. "Lifelong sequential modeling with personalized memorization for user response prediction." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

▪ Hierarchical Periodical Memory Network, HPMN

- Real-time query only on the maintained user memory
  - w/o inference over the whole user behavior sequence online

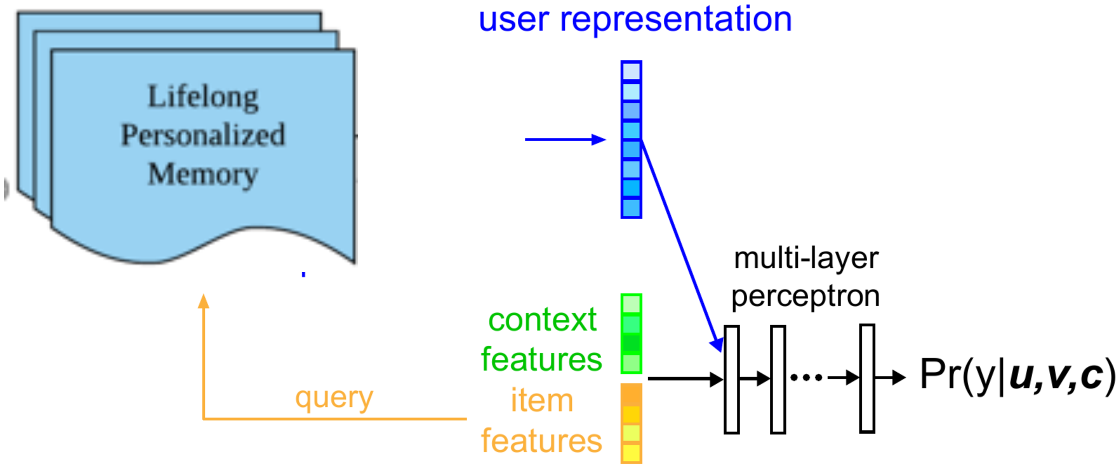

Figure 4: The overall user response prediction.

Ren, Kan, et al. "Lifelong sequential modeling with personalized memorization for user response prediction." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

▪ Hierarchical Periodical Memory Network, HPMN

- Overall Framework

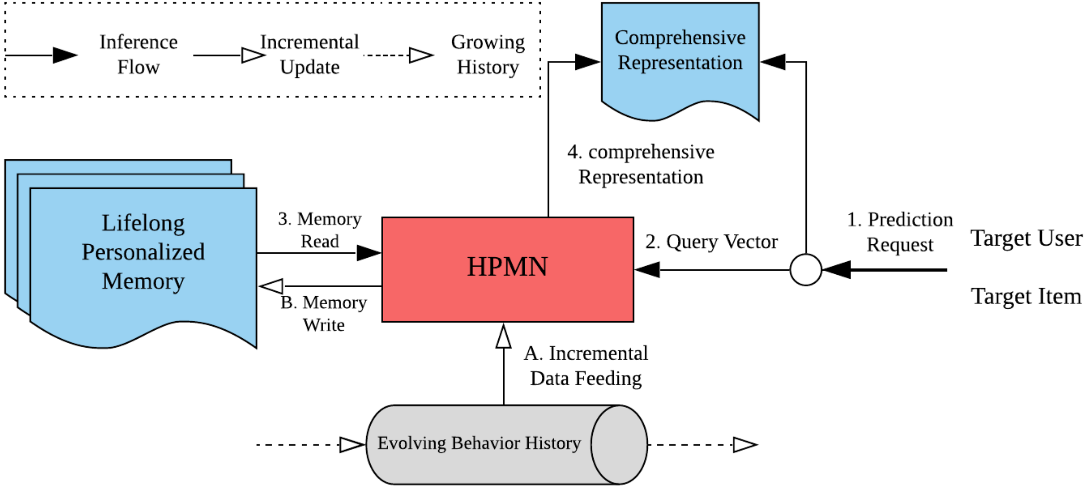

Figure 2: The LSM framework.

Ren, Kan, et al. "Lifelong sequential modeling with personalized memorization for user response prediction." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

▪ Hierarchical Periodical Memory Network, HPMN

Experimental results:

Table 4: Performance Comparison. (* indicates p-value < 10^{-6} in the significance test. ↑ and ↓ indicates the *performance* over lifelong sequences (with length T) is better or worse than the same model over short sequences (with length s). AUC: the higher, the better; Log-loss: the lower, the better. The second best performance of each metric is underlined.)

Model Group	Model	Len.	AUC Amazon	AUC Taobao	AUC XLong	Log-loss Amazon	Log-loss Taobao	Log-loss XLong
Group 2	GRU4Rec	s	0.7669	0.8431	0.8716	0.5650	0.4867	0.4583
	Caser	s	0.7509	0.8260	0.8467	0.5795	0.5094	0.4955
	DIEN	s	0.7725	0.8914	0.8725	0.5604	0.4184	0.4515
	RUM	s	0.7434	0.8327	0.8512	0.5819	0.5400	0.4931
Group 1	DNN	T	0.7546	0.7460	0.8152	0.6869	0.5681	0.5365
	SVD++	T	0.7155	0.8371	0.8008	0.6216	0.8371	1.7054
Group 2	GRU4Rec	T	0.7760 ↑	0.8471 ↑	0.8702 ↓	0.5569 ↑	0.4827 ↑	0.4630 ↓
	Caser	T	0.7582 ↑	0.8745 ↑	0.8390 ↓	0.5704 ↑	0.4550 ↑	0.5050 ↓
	DIEN	T	0.7770 ↑	0.8934 ↑	0.8716 ↓	0.5564 ↑	0.4155 ↑	0.4559 ↓
	RUM	T	0.7464 ↑	0.8370 ↑	0.8649 ↑	0.6301 ↓	0.4966 ↑	0.4620 ↑
Group 3	LSTM	T	0.7765	0.8681	0.8686	0.5612	0.4603	0.4570
	SHAN	T	0.7763	0.8828	0.8369	0.5595	0.4318	0.5000
	HPMN	T	**0.7809***	**0.9240***	**0.8929***	**0.5535***	**0.3487***	**0.4150***

Visualized Analysis:

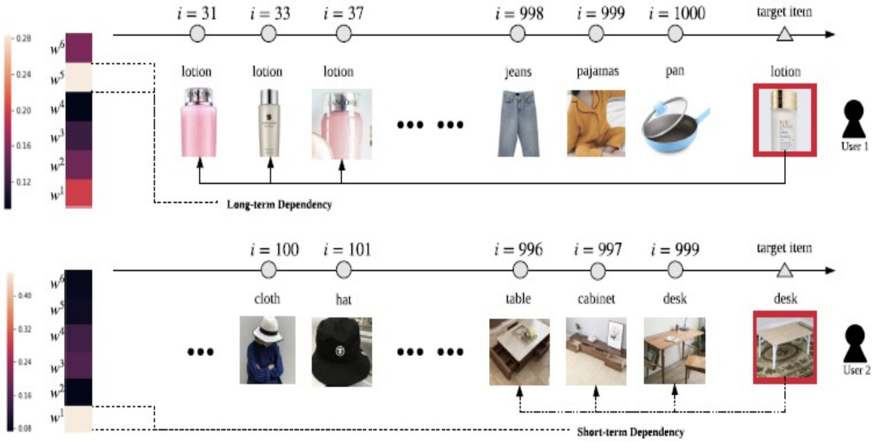

Ren, Kan, et al. "Lifelong sequential modeling with personalized memorization for user response prediction." *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019.

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Retrieval Methods

- **User Behavior Retrieval, UBR**
  - Most of existing sequential CTR works use the most recent N behaviors, if N is large
    - Heavy burden on system overhead (storage and latency), unable to model ultra-long historical behavior sequences
    - Longer sequences have a lot of noise
    - Each prediction uses exactly the same recent N behaviors, which may be not suitable for different target items
  - Instead of designing more complex model, retrieval methods turn to the data perspective
    - For each prediction, retrieve the most useful N behaviors (N is not large) from all the user's log
    - Use search engine technique to retrieve relevant behaviors

Qin, Jiarui, et al. "User behavior retrieval for click-through rate prediction." *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2020.

Retrieval Methods

- User Behavior Retrieval, UBR

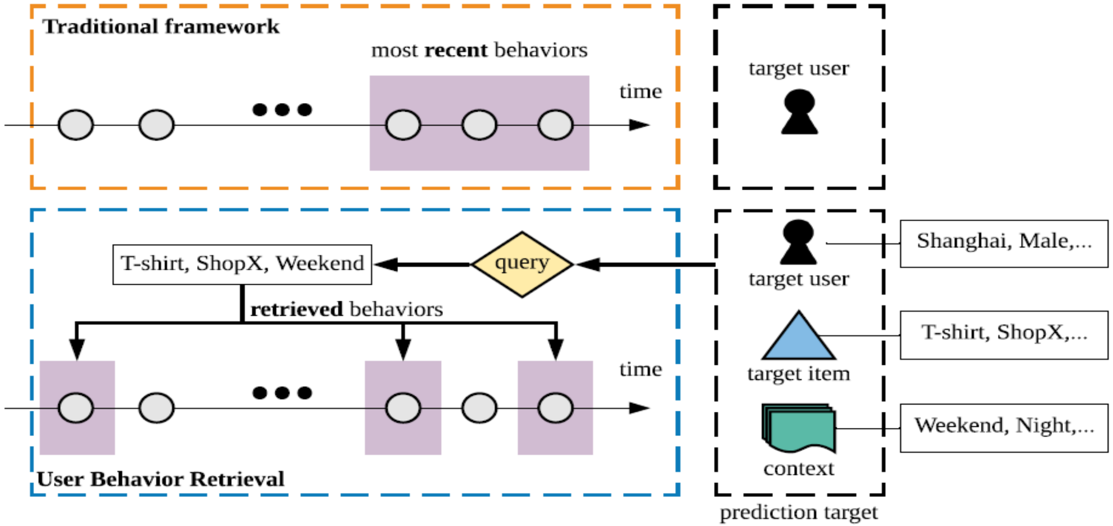

Qin, Jiarui, et al. "User behavior retrieval for click-through rate prediction." *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2020.

Retrieval Methods

▪ User Behavior Retrieval, UBR

- For the same user, when predicting interest in different items, different historical behaviors of the user can be retrieved for modeling.

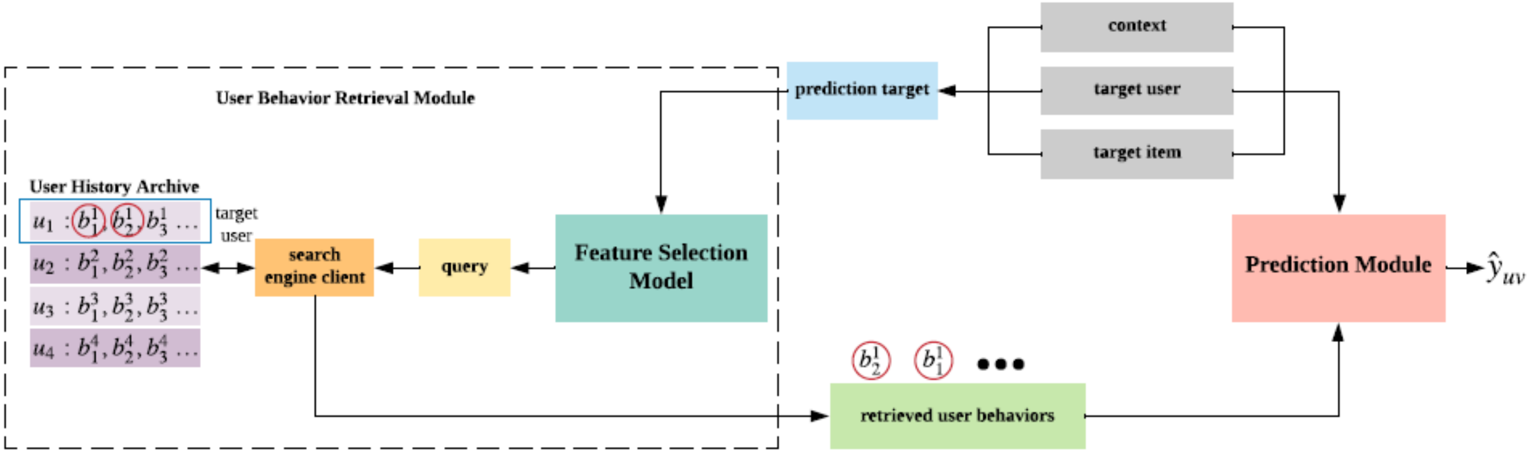

Qin, Jiarui, et al. "User behavior retrieval for click-through rate prediction." *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2020.

▪ User Behavior Retrieval, UBR

- Feature selection module: Use self-attention to select important features and combine them with the user ID to form a query.
- Behavior search module: searches for the most important historical behavior of the current user based on the query.

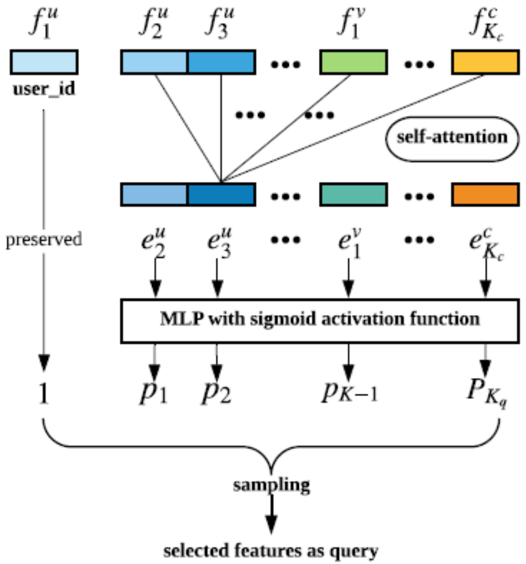

Behavior storage: feature based inverted index

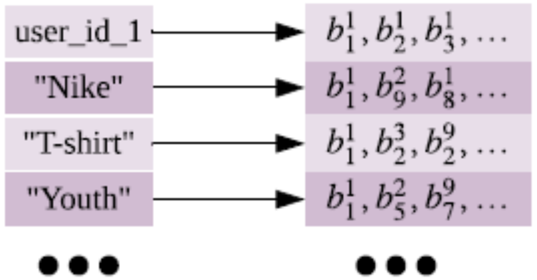

The query is formulated as

$$f_1^u \; AND \; (f_1 \; OR \; f_2 \; OR ... OR \; f_n),$$

Qin, Jiarui, et al. "User behavior retrieval for click-through rate prediction." *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2020.

User Behavior Retrieval, UBR

Experimental results:

Group1: Same length of user sequences

Model	Tmall		Taobao		Alipay	
	AUC	LL	AUC	LL	AUC	LL
GRU4Rec	0.762	0.585	0.677	0.661	0.6131	0.699
Caser	0.762	0.579	0.673	0.657	0.655	0.676
SASRec	0.755	0.586	0.670	0.658	0.648	0.679
HPMN	0.763	0.579	0.668	0.660	0.615	0.703
MIMN	0.753	0.591	0.662	0.686	0.664	0.675
DIN	0.766	0.576	0.678	0.649	0.732	0.616
DIEN	0.775	0.567	0.677	0.659	0.730	0.616
UBR4CTR	**0.807**	**0.516**	**0.752**	**0.571**	**0.895**	**0.417**
Imprv.	4.1%	9.0%	10.9%	12.0%	22.3%	32.3%

Experimental results:

Group2: Longer sequences for baselines

Model	Tmall		Taobao		Alipay	
	AUC	LL	AUC	LL	AUC	LL
GRU4Rec	0.781	0.560	0.677	0.660	0.639	0.684
Caser	0.774	0.566	0.645	0.659	0.705	0.631
SASRec	0.769	0.578	0.669	0.654	0.711	0.637
HPMN	0.767	0.579	0.655	0.664	0.703	0.643
MIMN	0.759	0.590	0.659	0.659	0.719	0.634
DIN	0.791	0.546	0.605	0.679	0.856	0.506
DIEN	0.805	0.538	0.704	0.656	0.843	0.491
UBR4CTR	**0.807**	**0.516**	**0.752**	**0.571**	**0.895**	**0.417**
Imprv.	0.2%	4.1%	6.8%	12.7%	4.6%	15.1%

Qin, Jiarui, et al. "User behavior retrieval for click-through rate prediction." *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2020.

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Multi-Type UBM

- Explicitly model different behavior types (e.g., view, click, purchase)
- Formulated as

$$p(u, i, b) = F_{\Theta}^{UBM}(u, i, H_u^{MB}), \forall u \in U, i \in I.$$

- Different behavior types have different characteristics, but are also correlated.

Late Fusion

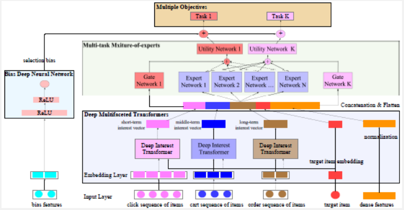

DMT [CIKM '20]

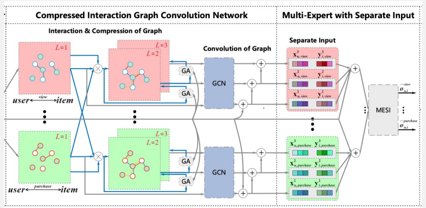

CIGF [WWW '23]

Early Fusion

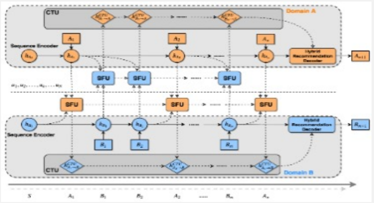

π-Net [SIGIR '19]

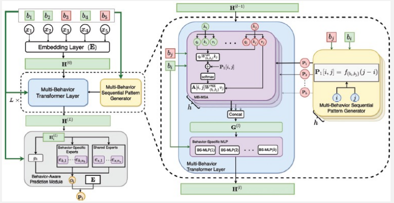

MBSTR [SIGIR '22]

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Late Fusion

- **Compressed Interaction Graph based Framework, CIGF**
  - Users interact with items through different behaviors, which can be treated “as features” for multi-behavior relation learning or “as labels” for multi-task supervised learning.
  - Inadequate modeling of high-order relations when treating multi-behavior data “as features”.
  - Potential gradient conflict when treating multi-behavior data “as labels”.

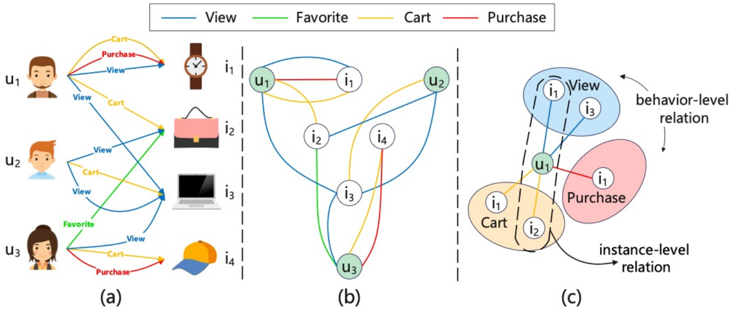

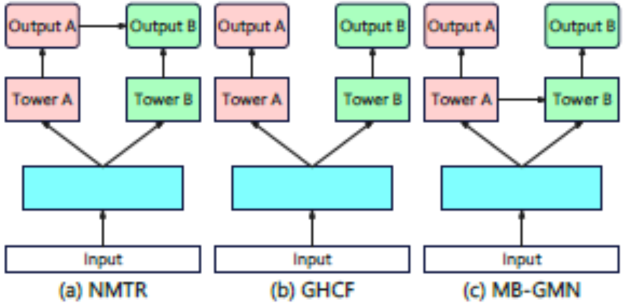

Guo, Wei, et al. "Compressed Interaction Graph based Framework for Multi-behavior Recommendation." *arXiv preprint arXiv:2303.02418* (2023).

Late Fusion

▪ Compressed Interaction Graph Convolution Network, CIGCN

➢ Interaction & Compression of Graph

- Three technical details: subgraph partitioning, instance-level interaction, and multi-head attention mechanism for compression.

➢ Graph Convolution

- The information of neighbor nodes is aggregated by a message passing mechanism.
- Residual operation is set to alleviate the over-smoothing phenomenon caused by increasing layers and excavate the higher-order behavior correlation.

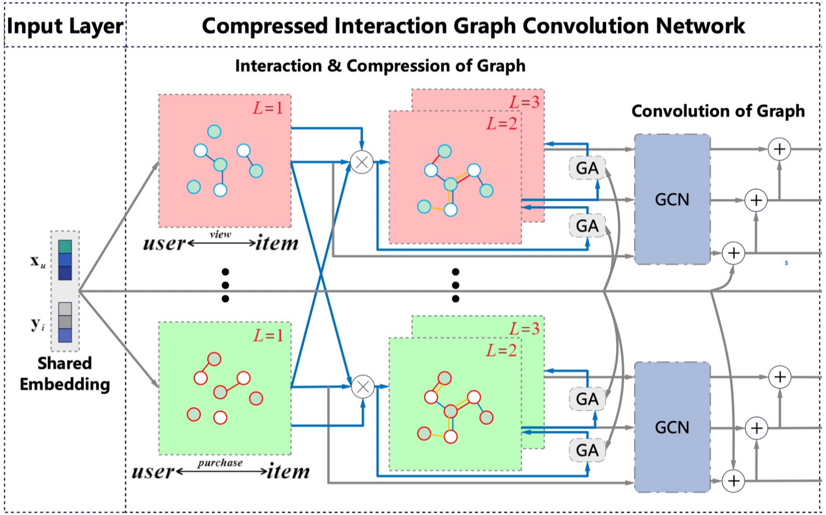

Guo, Wei, et al. "Compressed Interaction Graph based Framework for Multi-behavior Recommendation." *arXiv preprint arXiv:2303.02418* (2023).

Late Fusion

- **Multi-Expert with Separate Input, MESI**

➢ Multi-Expert with Separate Input
  - Different expert networks were used to extract perceptual information represented by decoupled inputs.
  - Decoupled inputs are used to generate different task perception gates to converge expert subsets of different behavioral perceptions.

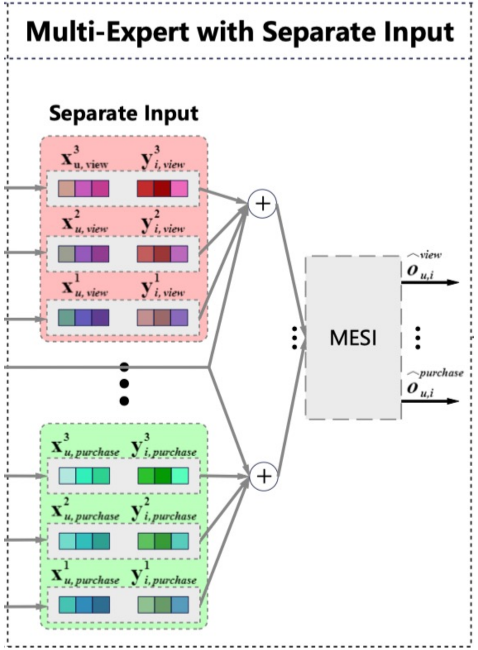

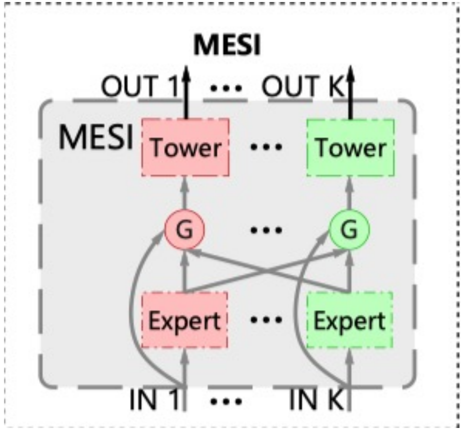

Guo, Wei, et al. "Compressed Interaction Graph based Framework for Multi-behavior Recommendation." *arXiv preprint arXiv:2303.02418* (2023).

▪ Compressed Interaction Graph based Framework, CIGF

- CIGF consistently yields superior performance on all three datasets.
- It is important to simultaneously consider multi-behavior and high-order relations.

Dataset	Beibei		Taobao		IJCAI	
Model	HR	NDCG	HR	NDCG	HR	NDCG
DMF	0.597	0.336	0.305	0.189	0.392	0.250
AutoRec	0.607	0.341	0.313	0.190	0.448	0.287
NGCF	0.611	0.375	0.302	0.185	0.461	0.292
LightGCN	0.643	0.378	0.373	0.235	0.443	0.283
NMTR	0.613	0.349	0.332	0.179	0.481	0.304
DIPN	0.631	0.394	0.317	0.178	0.475	0.296
MATN	0.626	0.385	0.354	0.209	0.489	0.309
$NGCF_M$	0.634	0.372	0.374	0.221	0.481	0.307
$LightGCN_M$	0.651	0.391	0.391	0.243	0.486	0.317
GHCF	0.608	0.378	0.415	0.241	-	-
MBGCN	0.642	0.376	0.369	0.222	0.463	0.277
MB-GMN	0.691	0.410	0.491	0.300	0.532	0.345
CIGF	**0.700**⋆	**0.443**⋆	**0.592**⋆	**0.383**⋆	**0.601**⋆	**0.400**⋆
%Improv	1.30%	8.05%	20.57%	27.67%	12.97%	15.94%

Guo, Wei, et al. "Compressed Interaction Graph based Framework for Multi-behavior Recommendation." *arXiv preprint arXiv:2303.02418* (2023).

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Early Fusion

- **Multi-Behavior Sequential Transformer Recommender, MBSTR**

- In real-world online service platforms, user interactions are intrinsically Sequential and Multi-behavioral.
- Sequential: Interest dynamics and evolution.
- Multi-behavioral: Fine-grained interest patterns, Solution for data sparsity.

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

Early Fusion

- **Multi-Behavior Sequential Transformer Recommender, MBSTR**

➢ Complex multi-behavior dependencies
  - Capture fine-grained heterogeneous dependencies.

➢ Diverse multi-behavior sequential patterns
  - Explicitly model the diverse sequential patterns

➢ Sparse target behavior data
  - An appropriate training strategy with a dedicated prediction module

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

Early Fusion

▪ Multi-Behavior Sequential Transformer Recommender, MBSTR

MB-Transformer Layer

- Capture heterogeneous dependencies

MB-sequential pattern generator

- Model diverse sequential patterns

Behavior-aware masked prediction

- Give better predictions on target behavior

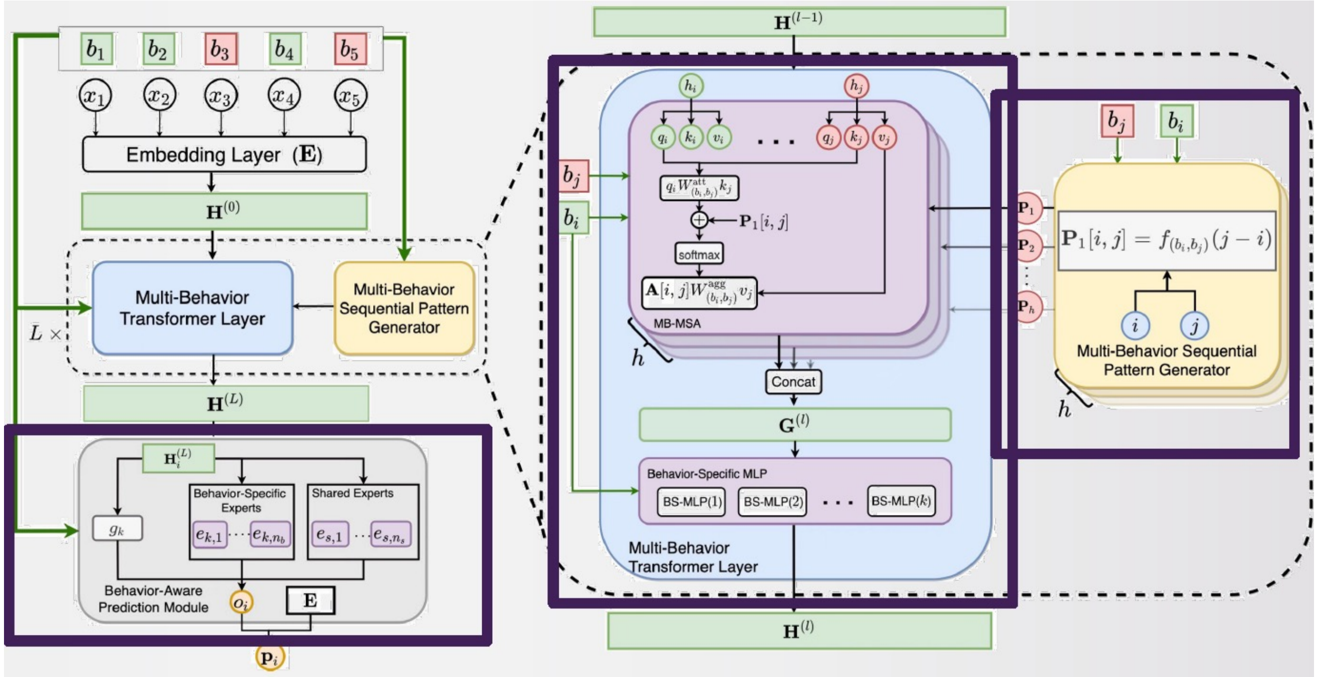

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

Early Fusion

▪ Multi-Behavior Sequential Transformer Recommender, MBSTR

- Traditionally: homogeneous
  - Words, image patches, Audio segments…
- Heterogeneous dependencies in multi-behavior settings
- Two components
  - MB-MSA: multi-behavior integration.
  - BS-MLP: behavior-specific transformation.

Algorithm 1: Multi-Behavior Multi-head Self-Attention

Input: $H^{(l-1)} \in \mathbb{R}^{n\times d}, b \in \mathcal{B}^n, P^{(l)} \in \mathbb{R}^{h\times n\times n}$
Output: $G^{(l)} \in \mathbb{R}^{n\times d}$

```
for head m = 1 to h do
    /* Step 1. Behavior-specific projection.                        */
    Q_m ← f_{Q_m}(H^{(l-1)}, b)
    K_m ← f_{K_m}(H^{(l-1)}, b)
    V_m ← f_{V_m}(H^{(l-1)}, b)
    /* Step 2. Cross behavior similarity.                           */
    for i = 1 to n, j = 1 to n do
        A_m[i, j] = Q_m[i] W^{att}_{(b[i],b[j])} K_m[j] / √d
    end
    /* Step 3. Sequential pattern injection and softmax.            */
    A_m ← softmax(A_m + P[m])
    /* Step 4. Cross behavior information aggregation.              */
    for i = 1 to n do
        G^{(l)}_m[i] ← Σ_j A_m[i, j] · W^{agg}_{(b[i],b[j])} · V_m[j]
    end
end
G^{(l)} ← Concat(G^{(l)}_1, ..., G^{(l)}_h)
```

4.1.2 Behavior-Specific Multi-Layer Perceptron. To model behavior semantics and perform feature transformation, we apply Behavior-Specific MLPs (BS-MLP). Specifically, we specify a distinct MLP layer for each type of behavior. The BS-MLP is defined as follows:

$$H^{(l)} = \text{BS-MLP}(G^{(l)}, b), \tag{4}$$

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

Early Fusion

- **Multi-Behavior Sequential Transformer Recommender, MBSTR**

➢ Diverse sequential patterns

- E.g. clicks: short-term interests vs purchases: long-term interests

➢ Encode sequential patterns:

$$\underbrace{\mathbf{P}[i,j]}_{\substack{\text{Multi-behavior}\\ \text{relative positional}\\ \text{encoding}}} = f_{\underbrace{(\mathbf{b}[i],\mathbf{b}[j])}_{\text{Behavior types}}} \underbrace{(j-i)}_{\text{Relative position}}$$

➢ Add P to attention matrix

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

Early Fusion

▪ Multi-Behavior Sequential Transformer Recommender, MBSTR

- How to effectively train the model with multi-behavior data?
  - Multi-behavior data as inputs and supervision signals

- Behavior-Aware Masked Item Prediction

$\mathbf{X} = (\mathbf{x}, \mathbf{b})$ corrupted sequence: $\hat{\mathbf{x}}$ masked items: $\bar{\mathbf{x}}$

$$\min_{\theta} -\log p_{\theta}(\bar{\mathbf{x}} \mid \hat{\mathbf{x}}, \mathbf{b}) = -\sum_{i=1}^{n} m_i \log p_{\theta}\left(\mathbf{x}[i] \mid \hat{\mathbf{x}}, \mathbf{b}\right),$$

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

▪ Multi-Behavior Sequential Transformer Recommender, MBSTR

- The effectiveness of the MB-STR model.
- Both sequential and multi-behavioral information bring benefits to model performance.

Dataset			Yelp		Taobao		IJCAI	
Metrics			HR	NDCG	HR	NDCG	HR	NDCG
O	NS	MF	0.755	0.481	0.262	0.153	0.285	0.185
		DMF	0.756	0.485	0.305	0.189	0.392	0.250
		NGCF	0.789	0.500	0.302	0.185	0.461	0.292
		LightGCN	0.810	0.513	0.373	0.235	0.443	0.283
	S	SASRec	0.796	0.504	0.372	0.221	0.597	0.406
		BERT4Rec	0.816	0.531	0.385	0.234	0.605	0.431
M	NS	NGCF_M	0.793	0.492	0.374	0.221	0.481	0.307
		LightGCN_M	<u>0.872</u>	<u>0.585</u>	0.391	0.243	0.486	0.317
		NMTR	0.790	0.478	0.332	0.179	0.481	0.304
		MATN	0.826	0.530	0.354	0.209	0.489	0.309
		MBGCN	0.796	0.502	0.369	0.222	0.463	0.277
		MB-GMN	0.87	0.582	0.491	0.300	0.532	0.345
	S	DIPN	0.791	0.500	0.317	0.178	0.475	0.296
		SASRec_M	0.819	0.531	0.637	0.442	0.795	0.611
		BERT4Rec_M	0.838	0.558	<u>0.675</u>	<u>0.476</u>	<u>0.816</u>	<u>0.632</u>
		DMT	0.652	0.515	0.666	0.415	0.682	0.513
Our MB-STR			$\mathbf{0.882}^*$	$\mathbf{0.624}^*$	$\mathbf{0.768}^*$	$\mathbf{0.608}^*$	$\mathbf{0.879}^*$	$\mathbf{0.713}^*$
Rela, Improv.			1.15%	6.67%	13.78%	27.73%	7.72%	12.82%

Yuan, Enming, et al. "Multi-behavior sequential transformer recommender." *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2022.

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Behavior type definition
  - Multi-behavior type fusion and prediction
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

UBM with Side Information: Time Information

- In the past, researchers usually only considered the relative time of items as the ranking basis.

- A simple example: **Positional Encoding**

$$PE(pos, 2i) = \sin\left(\frac{pos}{10000^{2i/d_{\text{model}}}}\right)$$

$$PE(pos, 2i+1) = \cos\left(\frac{pos}{10000^{2i/d_{\text{model}}}}\right)$$

- However, the interaction time of each item actually contains a lot of information.

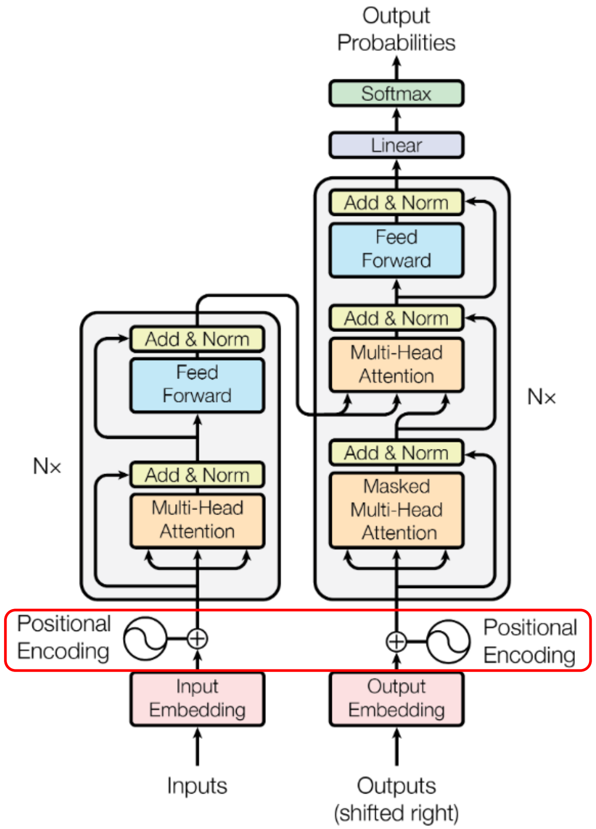

Vaswani, Ashish, et al. "Attention is all you need." *Advances in neural information processing systems* 30 (2017).

UBM with Side Information: Time Information

- There are some other methods using time information:

- In shared account recommendation, the interaction time can be used to distinguish the interests of different users. (π-Net)
- By perceiving interaction intervals, it helps to extract users' emerging preferences. (TIEN)
- To better model user short–term interests, you can also use dynamic time intervals to split sessions. (TiSSA)

Ma, Muyang, et al. "π-net: A parallel information-sharing network for shared-account cross-domain sequential recommendations." *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval. 2019.*

Li, Xiang, et al. "Deep time-aware item evolution network for click-through rate prediction." *Proceedings of the 29th ACM International Conference on Information & Knowledge Management. 2020.*

Lei, Chenyi, Shouling Ji, and Zhao Li. "Tissa: A time slice self-attention approach for modeling sequential user behaviors." *The World Wide Web Conference. 2019.*

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Behavior type definition
  - Multi-behavior type fusion and prediction
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

UBM with Side Information: Item Attribute [FDSA, IJCAI '19]

- Not only the user interaction sequence contains information, but also the attribute information contains a lot of user interests.

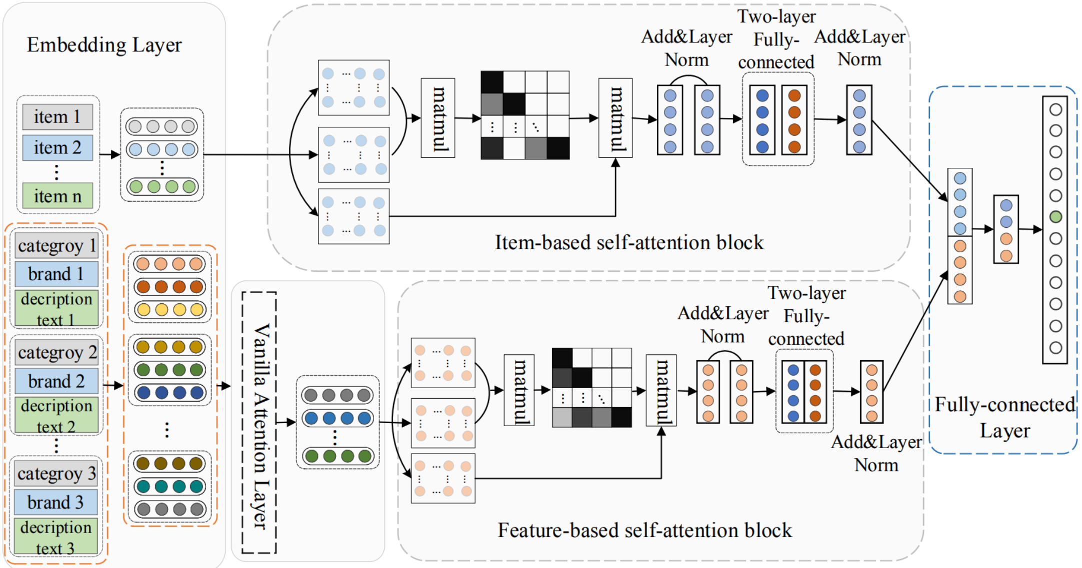

Zhang, Tingting, et al. "Feature-level Deeper Self-Attention Network for Sequential Recommendation." *IJCAI*. 2019.

UBM with Side Information: Item Attribute [FDSA, IJCAI '19]

- Through the fusion of interaction information and attribute information, the modeling effect can be significantly improved.

Dataset	Method	@5		@10	
		Hit	NDCG	Hit	NDCG
Tmall	PopRec	0.1532	0.0988	0.2397	0.1267
	BPR	0.1749	0.1129	0.2647	0.1418
	FPMC	0.2731	0.2034	0.3680	0.2339
	TransRec	0.2652	0.1854	0.3773	0.2214
	GRU4Rec	0.1674	0.1217	0.2446	0.1465
	CSAN	0.3481	0.2440	0.4787	0.2863
	SASRec	0.3572	0.2531	0.4840	0.2940
	SASRec+	0.3427	0.2415	0.4714	0.2829
	SASRec++	0.3550	0.2534	0.4785	0.2932
	CFSA	0.3836	0.2724	0.5152	0.3149
	FDSA	**0.3940**	**0.2820**	**0.5197**	**0.3226**
Toys and Games	PopRec	0.1952	0.1287	0.3058	0.1643
	BPR	0.2096	0.1394	0.3219	0.1756
	FPMC	0.2983	0.2261	0.3833	0.2535
	TransRec	0.3135	0.2255	0.4206	0.2600
	GRU4Rec	0.2039	0.1359	0.3118	0.1705
	CSAN	0.2327	0.1601	0.3404	0.1947
	SASRec	0.3292	0.2334	0.4441	0.2705
	SASRec+	0.3367	0.2410	0.4510	0.2776
	SASRec++	0.3394	0.2428	0.4544	0.2799
	CFSA	0.3391	0.2411	0.4538	0.2782
	FDSA	**0.3571**	**0.2572**	**0.4738**	**0.2949**

Zhang, Tingting, et al. "Feature-level Deeper Self-Attention Network for Sequential Recommendation." *IJCAI*. 2019.

Item Attribute: Transferring in KCG [TechCD, SIGIR ’23]

- This is a cross–domain zero–shot cognitive diagnosis task.
- The key point of information transferring is to find an intermediary.
- Attributes of users and items can be transferred through KCGs.

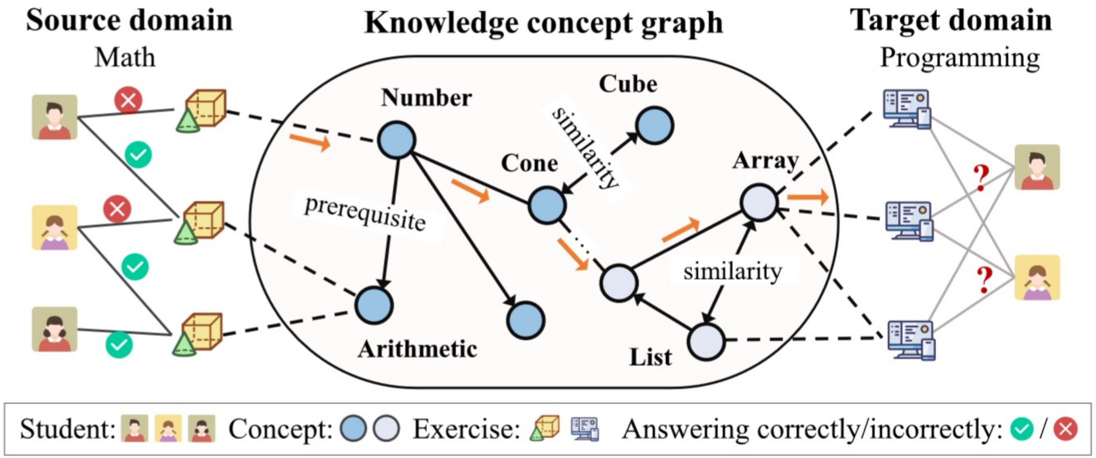

“↑” demonstrates a linking from the mature domain to the new domain.

Gao, Weibo, et al. "Leveraging Transferable Knowledge Concept Graph Embedding for Cold-Start Cognitive Diagnosis." *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2023.

Item Attribute: Transferring in KCG [TechCD, SIGIR ’23]

- Employing a pedagogical KCG as the intermediary to facilitate the sharing of student cognitive states across different domains.
- Using GCN to aggregate the information in KCG.
- Consider the high–order output of GCN as the general interest.

$$z_i^{(l)} = \sum_{r\in\mathcal{R}_i} \frac{1}{|\mathcal{P}_i^r|} \sum_{(e_j,r,e_i)\in\mathcal{P}_i^r} \mathcal{W}_r z_j^{(l-1)},$$

$$z_i^{\#} = \frac{1}{L-\lambda+1}\sum_{l=\lambda}^{L} z_i^{(l)}, z_i^{*} = \frac{1}{L+1}\sum_{l=0}^{L} z_i^{(l)}.$$

$$\boldsymbol{h}_u = \frac{1}{|\mathcal{H}_u^S|}\sum_{v\in\mathcal{H}_u^S} z_v^{\#},$$

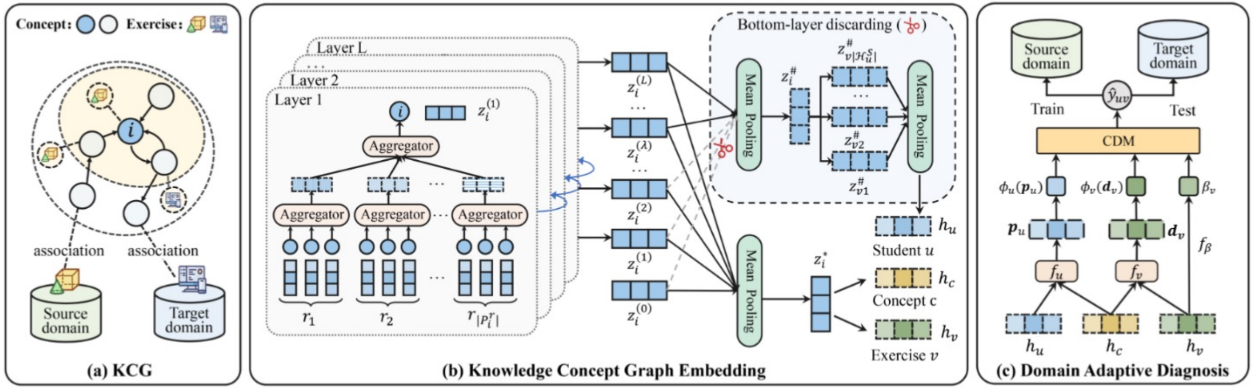

TechCD [SIGIR ’23]

Gao, Weibo, et al. "Leveraging Transferable Knowledge Concept Graph Embedding for Cold-Start Cognitive Diagnosis." *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2023.

Item Attribute: Transferring in KCG [TechCD, SIGIR ’23]

- This work aims to make a better cognitive diagnosis for students, leading to better exercise recommendation results, etc.
- We can observe a significant improvement in the model's prediction of student performance.

Dataset	Metric	IRT				MIRT				NeuralCD				Random
		Oracle	NLP	GCN	TechCD	Oracle	NLP	GCN	TechCD	Oracle	NLP	GCN	TechCD	
S-CM T-AM	ACC (%) ↑	77.89*	59.84	56.72	**63.45**	73.83*	56.44	56.74	**64.73**	74.65*	56.44	57.05	**57.06**	50.13
	AUC (%) ↑	84.98*	65.32	56.62	**67.42**	79.26*	65.52	56.60	**68.90**	81.07*	57.09	**57.44**	53.68	50.14
	RMSE (%) ↓	38.91*	47.98	50.75	**47.59**	48.40*	48.30	50.79	**47.06**	41.17*	49.69	50.72	**49.49**	57.70
S-AM T-CM	ACC (%) ↑	77.67*	55.88	56.92	**57.72**	74.07*	55.88	56.92	**57.78**	74.34*	55.88	56.80	**56.99**	49.91
	AUC (%) ↑	85.50*	50.68	56.62	**58.99**	81.16*	**60.56**	56.62	59.02	81.61*	53.67	**57.55**	52.40	49.89
	RMSE (%) ↓	39.08*	53.21	54.46	**52.85**	47.93*	**48.52**	50.50	52.85	41.52*	49.87	50.72	**49.57**	57.78

Gao, Weibo, et al. "Leveraging Transferable Knowledge Concept Graph Embedding for Cold-Start Cognitive Diagnosis." *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2023.

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Behavior type definition
  - Multi-behavior type fusion and prediction
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

UBM with Side Information: Multi-modal Information

- Multi–modal information, such as text, pictures, audio, etc., contains users' interests.
- It's vital but hard to make use of multi–modal information, for the different input form of each modality.

- Information of different modality is usually embedded into different representation spaces.
- The key point is Multimodal Fusion, that is, fusing the representations of different embedding space.

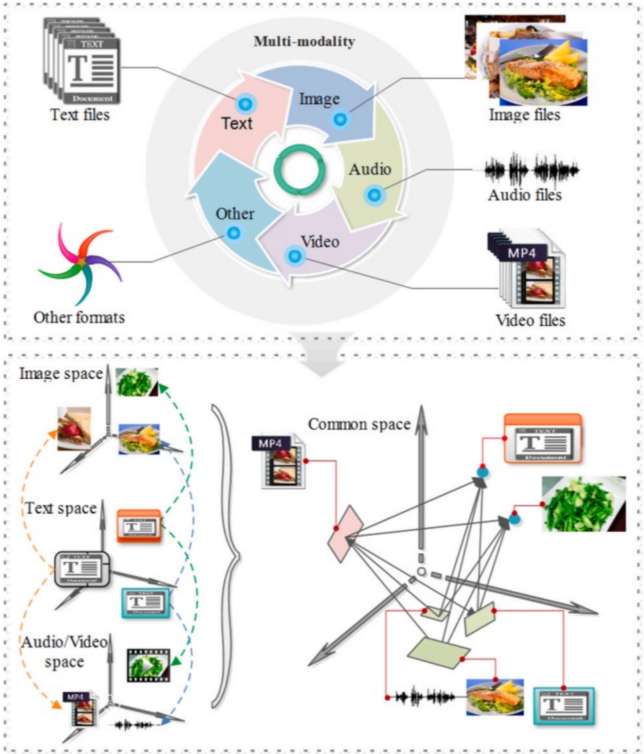

Lei, Zhenfeng, et al. "Is the suggested food your desired?: Multi-modal recipe recommendation with demand-based knowledge graph." *Expert Systems with Applications* 186 (2021): 115708.

Multi-modal Information: Fusing Multiple Modes [p-RNN, Recsys '16]

- Items often come with rich feature representation such as detailed text description or images.
- An ordinary model using RNN to fuse the multi-modal data.

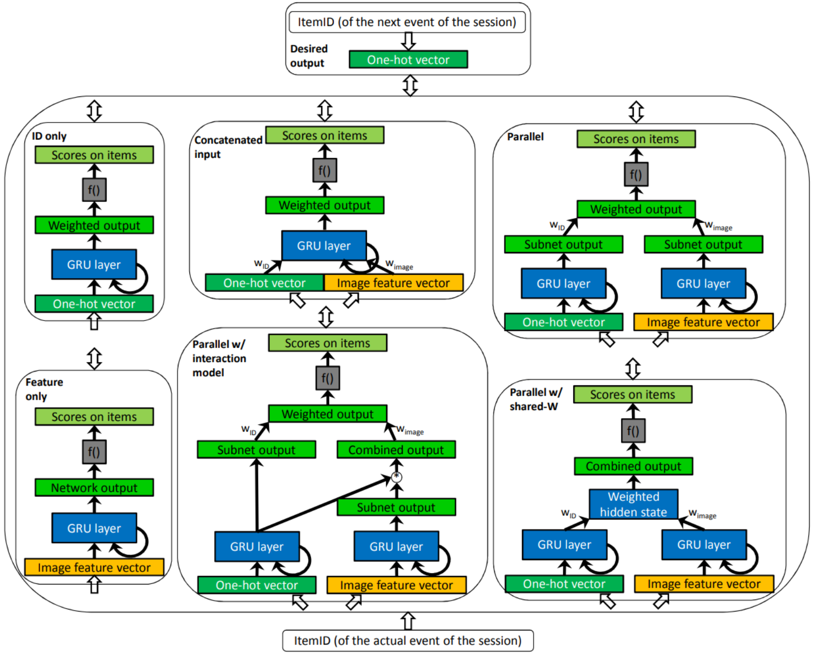

Hidasi, Balázs, et al. "Parallel recurrent neural network architectures for feature-rich session-based recommendations." *Proceedings of the 10th ACM conference on recommender systems*. 2016.

Multi-modal Information: Fusing Multiple Modes [p-RNN, Recsys ’16]

- The left table shows that the introduction of multimodal information helps the model achieve better results with a smaller number of parameters.
- The right figure shows that even if the GRU is simply used to integrate the multi–mode embedding, the effect can be significantly

Method	Recall@20	MRR@20
Item-kNN	0.6263	0.3740
ID only	0.6831 (+9.07%)	0.3847 (+2.85%)
ID only (200)	0.6963 (+11.17%)	0.3881 (+3.77%)
Feature only	0.5367 (−14.30%)	0.3065 (−18.05%)
Concatenated	0.6766 (+8.03%)	0.3850 (+2.94%)
Parallel (sim)	0.6765 (+8.01%)	0.4014 (+7.34%)
Parallel (alt)	0.6874 (+9.76%)	0.4331 (+15.81%)
Parallel (res)	0.7028 (+12.21%)	**0.4440** (+18.72%)
Parallel (int)	**0.7040** (+12.41%)	0.4361 (+16.60%)
Shared-W (sim)	0.6681 (+6.66%)	0.4007 (+7.13%)
Shared-W (alt)	0.6804 (+8.63%)	0.4035 (+7.88%)
Shared-W (res)	0.6425 (+2.58%)	0.3541 (−5.31%)
Shared-W (int)	0.6658 (+6.31%)	0.3715 (−0.66%)
Int. model (sim)	0.6751 (+7.78%)	0.3998 (+6.90%)
Int. model (alt)	0.6847 (+9.32%)	0.4104 (+9.74%)
Int. model (res)	0.6749 (+7.76%)	0.4098 (+9.56%)
Int. model (int)	0.6843 (+9.25%)	0.4040 (+8.02%)

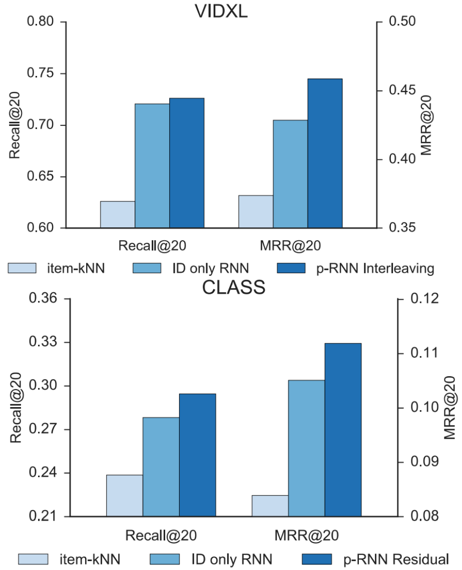

Hidasi, Balázs, et al. "Parallel recurrent neural network architectures for feature-rich session-based recommendations." *Proceedings of the 10th ACM conference on recommender systems*. 2016.

Multi-modal Information: Fine-tune LLM [LLaMA-Adapter, arXiv ’23]

- The ordinary large language model(LLM) can only deal with text, so it’s necessary to adopt some other training strategies to introduce multi-modal information.
- LLMs can make more efficient use of multimodal information by using the extensive background knowledge they store.

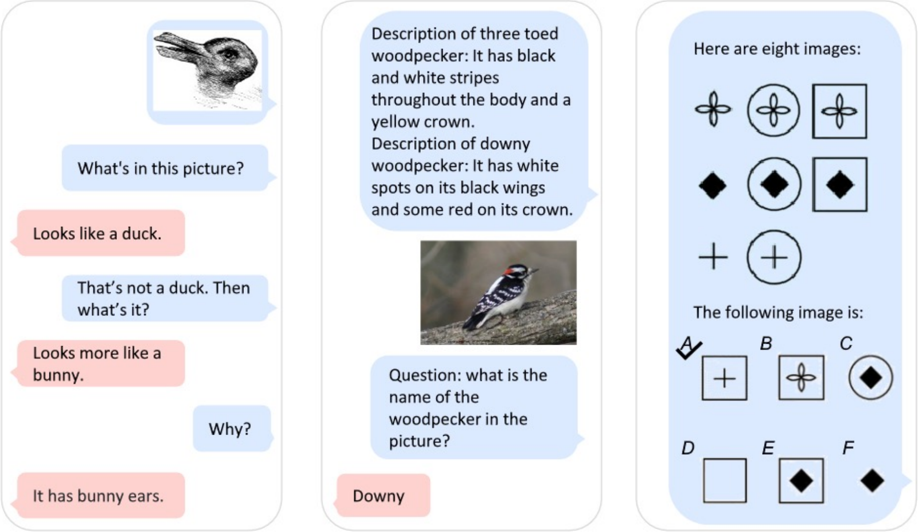

Huang, Shaohan, et al. "Language is not all you need: Aligning perception with language models." *arXiv preprint arXiv:2302.14045* (2023).

Multi-modal Information: Fine-tune LLM [LLaMA-Adapter, arXiv '23]

- With adapter, a lightweight but efficient additional network, LLM can be fine-tuned efficiently.
- By freezing the core architecture and making targeted adjustments solely to the adapter, the model can be effectively tuned , and the model can also be adapted to a particular domain.

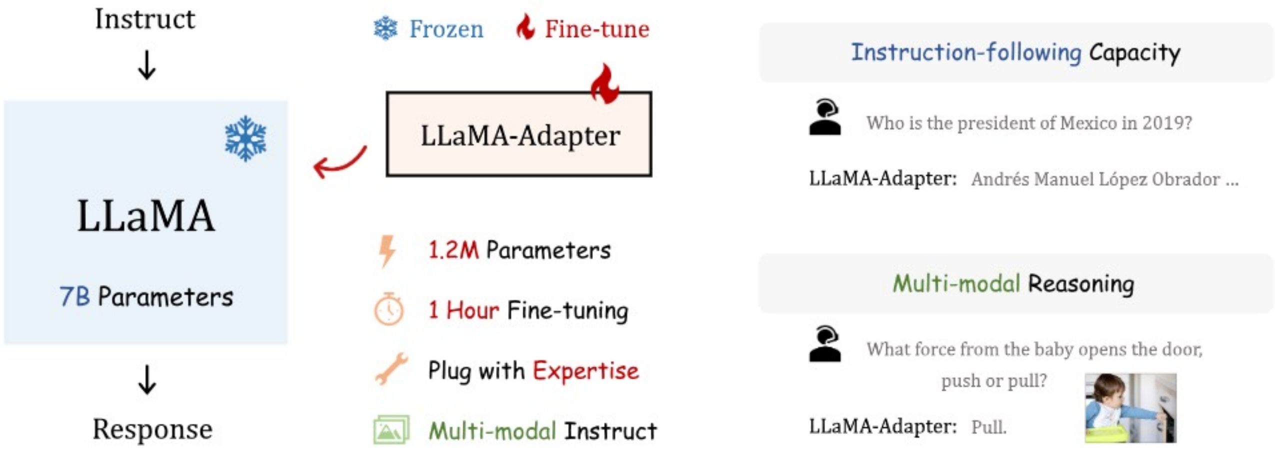

LLaMA-Adapter [arXiv '23]

Zhang, Renrui, et al. "Llama-adapter: Efficient fine-tuning of language models with zero-init attention." *arXiv preprint arXiv:2303.16199* (2023).

Multi-modal Information: Fine-tune LLM [LLaMA-Adapter, arXiv '23]

- This work uses the adapter to input multimodal information, giving the LLM the ability to process both images and texts.
- With the input of multi–modal information, the LLM has more comprehensive modeling and perception capabilities.

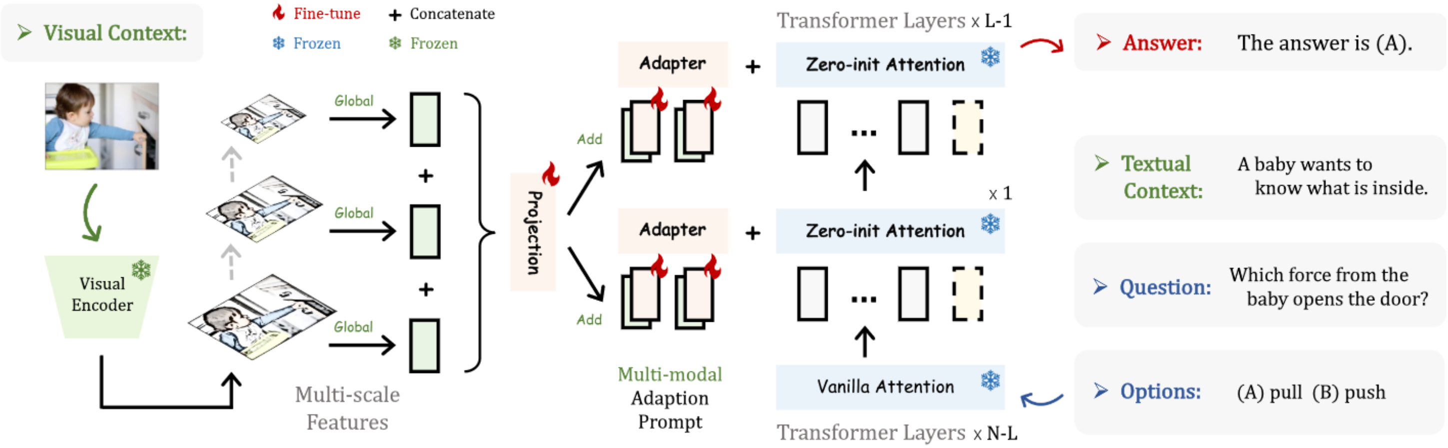

LLaMA-Adapter [arXiv '23]

Zhang, Renrui, et al. "Llama-adapter: Efficient fine-tuning of language models with zero-init attention." *arXiv preprint arXiv:2303.16199* (2023).

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

Industrial Practices

- UBM has been deployed in e-commerce, app store, and coupon allocation.
- The consideration of long-range behaviors, multi-type behaviors, and side information all achieve improvements over conventional UBM.
- Hashing or sampling techniques can improve the efficiency of UBM.
- GNN-based methods may have difficulty in deploying online.

Model	Application Scenario	Baseline	Gains & Costs
DIN	Online advertising	Embedding&MLP	+10.0% CTR, +3.8% RPM
DIEN	Online advertising	Embedding&MLP DIN	+20.7% CTR, +17.1% eCPM, -3.0% PPC +11.8% CTR, +10.4% eCPM, -1.0% PPC
GRU4Rec+	Online video	Strategy	+5% Watch time, +5% Video play, +4% Click
BST	E-commerce	DIN	+3.02% CTR, +4ms RT
(a) Conventional UBM			
UIC	Online advertising	DIEN	+7.5% CTR, +6% RPM
SIM	Online advertising	UIC DIEN	+7.1% CTR, +4.4% RPM, +53 times MSL +2.1 $d_{category}$
UBR	App store	w/o UBR	+6.6% eCPM, +11.1% CTR
ETA	E-commerce	SIM	+1.8% CTR, +3.1% GMV, -2ms IT
SDIM	Online search	w/o Long sequence	+2.98% CTR, +2.69% VBR, +1ms IT
ADFM	Online advertising	SIM	+4.7% CTR, +3.1% RPM, -70.8% Storage
(b) Long-sequence UBM			
DMT	E-commerce	DIEN	+4.5% CTR, +4.6% CVR, +6.0% GMV
ZEUS	E-commerce	DMT	+6.0% CTR, +9.7% CVR, 11.7% GMV
DIPN	Coupon allocation	Strategy	+41.1% Usage Rate, +39.8% GMV
(c) Multi-type UBM			
NOVA-BERT	App store	BERT	+0.192 $\times 10^9$ FLOPs, +7.1 Mb Model Size
SEMI	E-commerce	BST	+9.32% NBV, +10.45% DT, +12.10% CWR
TiSSA	E-commerce	w/o TiSSA	+1.56% CTR, +2.09% CVR, +3.66% GMV
(d) UBM with side information			

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

New Trends and Techniques: Reinforcement Learning

- Maximize **long-term reward**
- Model user preferences in an **interactive way**
  - Recommendation results influence users' current state (e.g., current preferences and interests) $\Rightarrow s$
  - Interacting with users in multiple turns $\Rightarrow state\ transition$
  - Acting in recommending users' favored items/item lists $\Rightarrow a$
  - Receive feedback from users $\Rightarrow r$

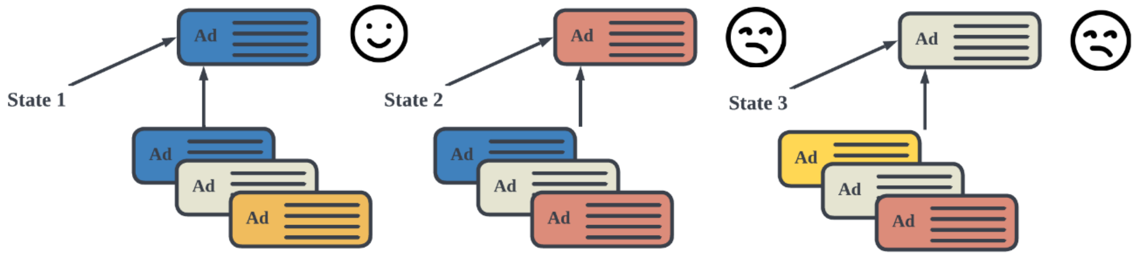

New Trends and Techniques: Reinforcement Learning

- Suppose we have a set of users $\mathcal{U}$ and a set of items $\mathcal{V}$. Then for each user, at timestamp t,
  - The recommender system observe a sequence of users' historical behaviors H_t **(state representation)**
  - Recommend an item or a list of items A_t to the user according to a recommendation policy π_t
  - The user observes the recommendation results and provides feedback f_{A_t}
  - The recommender system updates the policy π_t according to the feedback f_{A_t}
- **State representation**
  - Should summarize past information (user, item, context) such that all relevant information is not missed

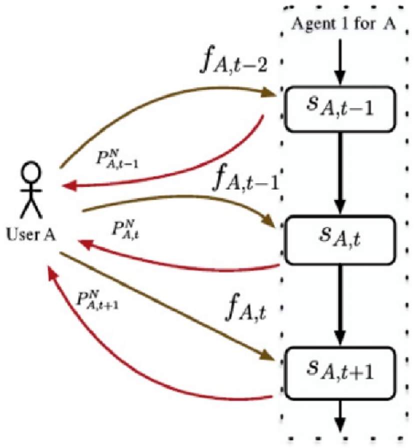

Reinforcement Learning: State Representation [DEERS, KDD'18]

- Positive behaviors: purchase, click…
- Negative behaviors: skip, leave…
- Indicate user preferences from the opposite aspects
- **Advantage:**
  - Avoid bad recommendation cases
- **Challenges:**
  - Negative feedback could bury the positive ones
  - May not come from users' dislike

Zhao, Xiangyu, et al. "Recommendations with negative feedback via pairwise deep reinforcement learning." KDD. 2018.

Reinforcement Learning: State Representation [DEERS, KDD'18]

- **Intuition:**
  - Recommend an item that is similar to the positive items
  - While dissimlar to the negative items
- GRU to capture users' sequential preferences
- Users tend to click one item while skip the other items in the same category

$$L(\theta) = \mathbb{E}_{s,a,r,s'}\left[\left(y - Q(s_+, s_-, a; \theta)\right)^2 - \alpha\left(Q(s_+, s_-, a; \theta) - Q(s_+, s_-, a^C; \theta)\right)^2\right]$$

Time	State	Item	Category	Feedback
1	s_1	a_1	A	skip
2	s_2	a_2	B	click
3	s_3	a_3	A	click
4	s_4	a_4	C	skip
5	s_5	a_5	B	skip
6	s_6	a_6	A	skip
7	s_7	a_7	C	order

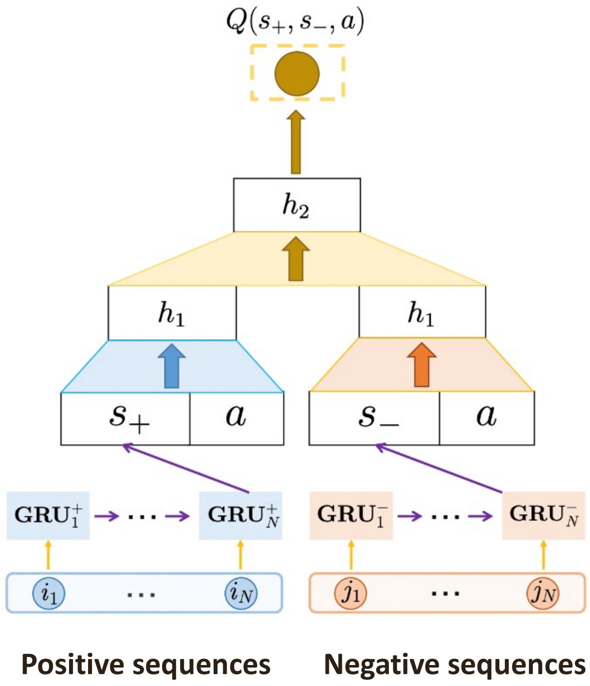

Zhao, Xiangyu, et al. "Recommendations with negative feedback via pairwise deep reinforcement learning." KDD. 2018.

Reinforcement Learning: State Representation [DEERS, KDD’18]

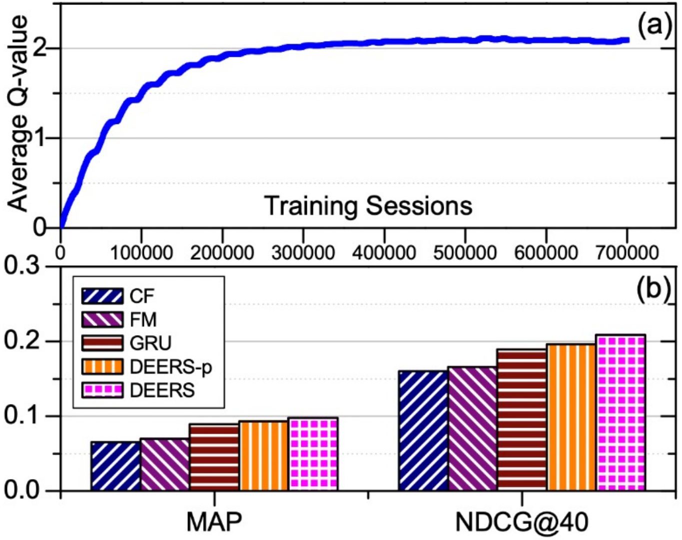

Offline evaluation

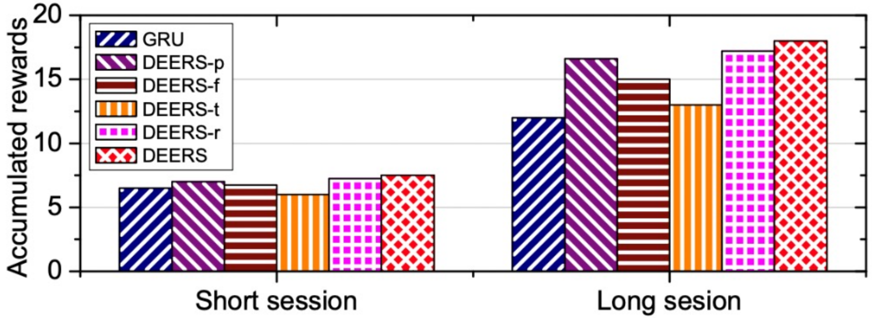

Online simulated evaluation

Zhao, Xiangyu, et al. "Recommendations with negative feedback via pairwise deep reinforcement learning." KDD. 2018.

Reinforcement Learning: State Representation [DRR, KBS'20]

- State representation: user current preferences and interests
- The state representation module
  - **DRR-p:** utilize a product operator to capture the pairwise local dependency between items.
  - **DRR-u:** add the pairwise interactions of user-item.
  - **DRR-ave:** eliminate the position effects
  - **DRR-att:** apply the attention mechanism

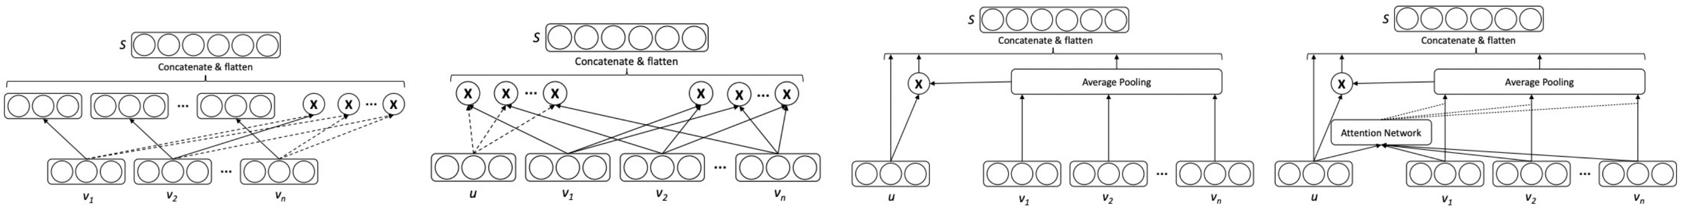

(a) DRR-p (b) DRR-u (c) DRR-ave (d) DRR-att

Liu, Feng, et al. "Deep reinforcement learning based recommendation with explicit user-item interactions modeling." *arXiv preprint arXiv:1810.12027* (2018).

Reinforcement Learning: State Representation [DRR, KBS'20]

- DRR outperforms existing representative SL and RL baselines

Overall Ranking performance on the **ML (100k)** and **Yahoo! Music** datasets.

Model	**ML (100k)**			**Yahoo! Music**		
	Precision@20	NDCG@20	MAP	Precision@20	NDCG@20	MAP
Popularity	0.5685	0.8720	0.6017	0.3424	0.8715	0.3928
PMF	0.5845	0.8849	0.6446	0.3657	0.8763	0.4235
SVD++	0.5876	0.8866	0.6461	0.3662	0.8789	0.4386
AFM	0.6325	0.8914	0.7038	0.3722	0.8809	0.4487
DeepFM	0.6362	0.8941	0.7097	0.3745	0.8818	0.4506
DQN	0.6076	0.8815	0.6704	0.3647	0.8812	0.4405
DDPG	0.6052	0.8870	0.6713	0.3664	0.8805	0.4412
DEERS	0.6481	0.8933	0.7226	0.3761	0.8849	0.4495
DRR-p	0.6112	0.8889	0.6924	0.3673	0.8825	0.4436
DRR-u	0.6244	0.8907	0.7075	0.3690	0.8831	0.4483
DRR-ave	0.6564	0.8982	0.7425	0.3763	0.8846	0.4685
DRR-att	**0.6784**[a]	**0.9025**[a]	**0.7768**[a]	**0.3802**[a]	**0.8913**[a]	**0.4750**[a]

[a]Indicates the statistically significant improvements (i.e., two-sided t-test with $p < 5e{-}5$) over the best baseline.

Offline experiments

The average rewards on the four datasets.

Model	**ML (100k)**	**Yahoo! Music**	**ML (1M)**	**Jester**
LinUCB	0.4266	0.0989	0.4996	0.2391
HLinUCB	0.3214	0.1062	0.5428	0.2488
DQN	0.5752	0.1207	0.6002	0.2758
DDPG	0.5783	0.1149	0.5937	0.2805
DEERS	0.7035	0.1625	0.6635	0.3274
DRR-p	0.6338	0.1337	0.6114	0.2932
DRR-u	0.6522	0.1385	0.6273	0.3046
DRR-ave	0.7105	0.1633	0.6746	0.3315
DRR-att	**0.7466**[a]	**0.1936**[a]	**0.6824**[a]	**0.3437**[a]

[a]Indicates the statistically significant improvements (i.e., two-sided t-test with $p < 1e{-}5$) over the best baseline.

Online simulation: Total rewards

Liu, Feng, et al. "Deep reinforcement learning based recommendation with explicit user-item interactions modeling." *arXiv preprint arXiv:1810.12027* (2018).

OUTLINE

- **01 INTRODUCTION**
- **02 CONVENTIONAL UBM**
  - Network structures: RNN, CNN, Attention
- **03 LONG-SEQUENCE UBM**
  - Memory-augmented methods
  - User behavior retrieval methods
- **04 MULTI-TYPE UBM**
  - Late fusion methods
  - Early fusion methods
- **05 UBM WITH SIDE INFORMATION**
  - Time information
  - Item attribute
  - Multi-modal information
- **06 INDUSTRIAL PRACTICES**
- **07 NEW TRENDS AND TECHNIQUES**
  - Reinforcement learning
  - Large language models

New Trends and Techniques: Large Language Models

- Large language models contains open-world knowledge and reasoning ability

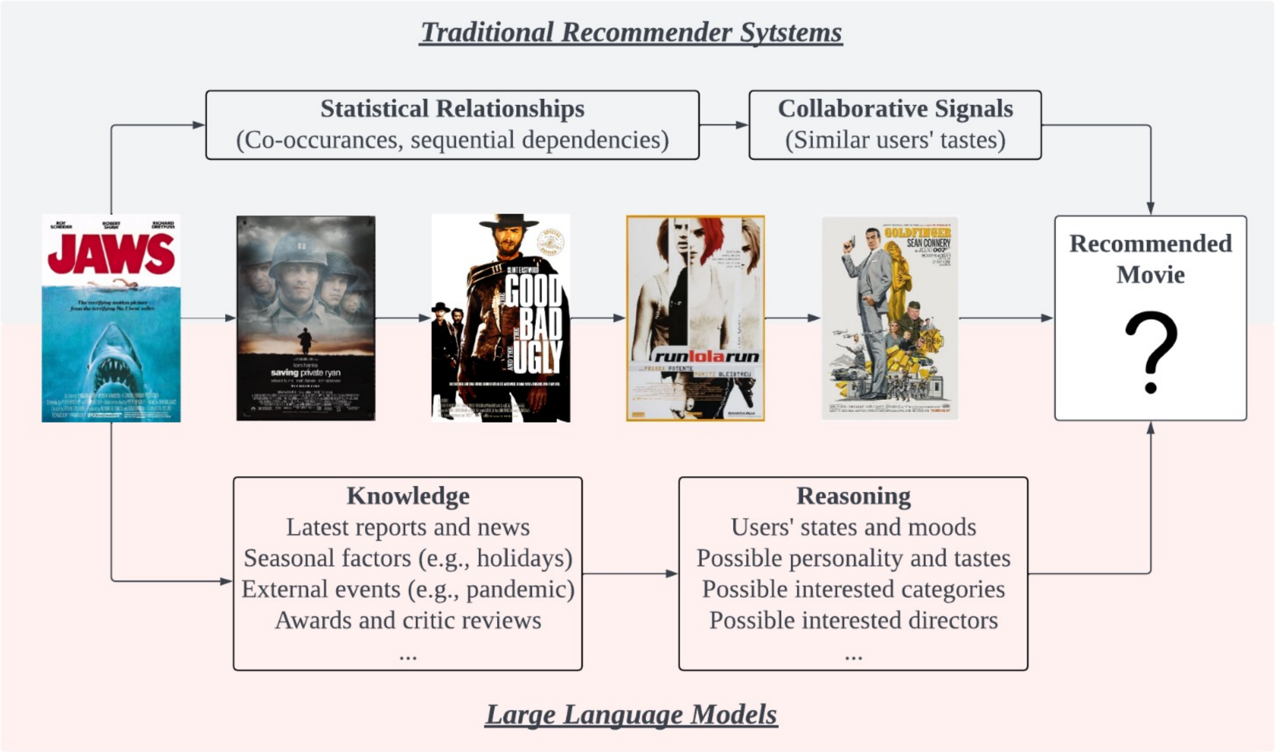

Large Language Models: Open-world Knowledge [Kar, DLP ’23]

- Large language models can perform logical reasoning
- Understand underlying preferences and motives that drive user behaviors

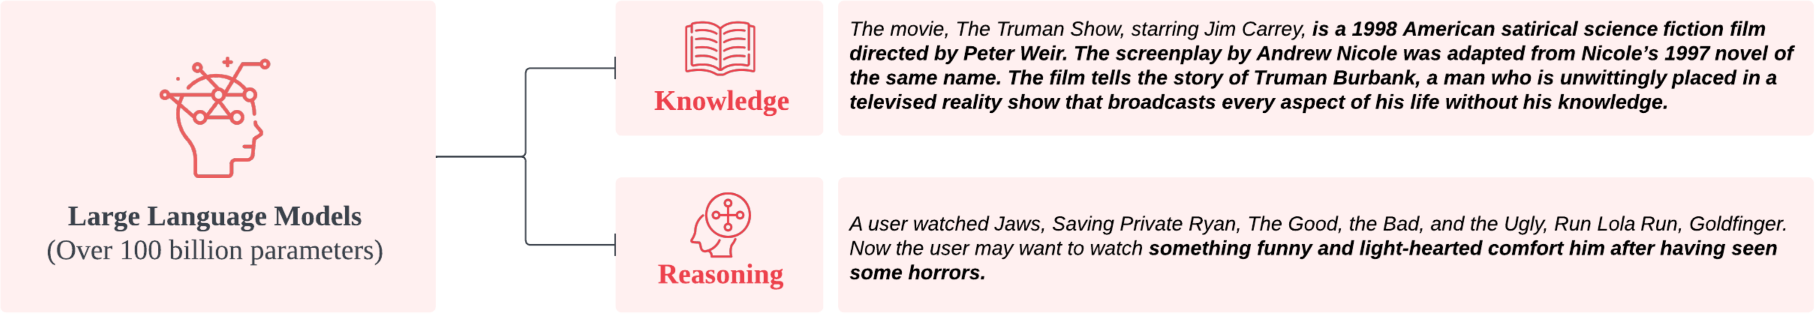

Xi, Yunjia, et al. "Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models." *arXiv preprint arXiv:2306.10933* (2023).

Large Language Models: Open-world Knowledge [Kar, DLP '23]

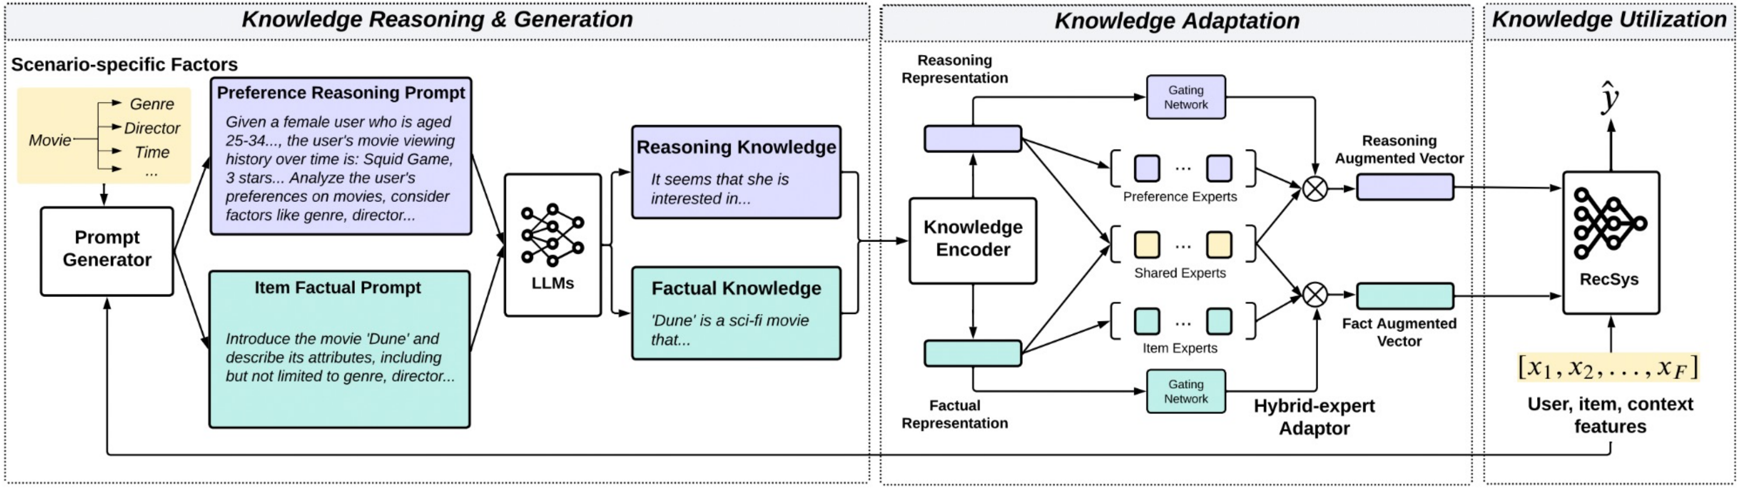

Xi, Yunjia, et al. "Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models." *arXiv preprint arXiv:2306.10933* (2023).

Large Language Models: Open-world Knowledge [Kar, DLP '23]

Backbone model	MovieLens-1M				Amazon-Books			
	AUC		LL		AUC		LL	
	base	KAR	base	KAR	base	KAR	base	KAR
DCNv2	0.7924	**0.8049***	0.5451	**0.5315***	0.8269	**0.8350***	0.4973	**0.4865***
DCNv1	0.7929	**0.8044***	0.5457	**0.5319***	0.8268	**0.8348***	0.4973	**0.4869***
DeepFM	0.7928	**0.8041***	0.5462	**0.5321***	0.8269	**0.8347***	0.4969	**0.4873***
FiBiNet	0.7925	**0.8051***	0.5450	**0.5310***	0.8269	**0.8351***	0.4973	**0.4870***
AutoInt	0.7934	**0.8060***	0.5440	**0.5297***	0.8262	**0.8357***	0.4981	**0.4863***
FiGNN	0.7944	**0.8054***	0.5424	**0.5307***	0.8270	**0.8352***	0.4977	**0.4870***
xDeepFM	0.7942	**0.8041***	0.5457	**0.5317***	0.8271	**0.8351***	0.4971	**0.4866***
DIEN	0.7960	**0.8059***	0.5469	**0.5298***	0.8307	**0.8391***	0.4926	**0.4812***
DIN	0.7975	**0.8066***	0.5387	**0.5304***	0.8304	**0.8418***	0.4937	**0.4801***

$*$ denotes statistically significant improvement (t-test with p-value < 0.05) over the backbone model.

Model	MovieLens-1M		Amazon-Books	
	AUC	LL	AUC	LL
UnisRec	0.7891	0.5496	0.8196	0.5063
VQ-Rec	0.7914	0.5456	0.8226	0.5025
base(DIN)	<u>0.7975</u>	<u>0.5387</u>	<u>0.8304</u>	<u>0.4937</u>
KAR(DIN)	**0.8066***	**0.5304***	**0.8418***	**0.4801***

$*$ denotes statistically significant improvement (t-test with p-value < 0.05) over the baseline/backbone models.

Xi, Yunjia, et al. "Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models." *arXiv preprint arXiv:2306.10933* (2023).

Large Language Models: Knowledge Prompt-tuning [KP4SR, MM '23]

- Limitation of adapting LLM to recommendation:
  - Language model lack domain knowledge
  - Struggle to capture users' fine-grained preferences
- Introduce an external knowledge graph to LM for behavior modeling
- Two further challenges:
  - How to convert structured knowledge graphs into text sequences
  - How to deal with the noise caused by irrelevant entities and relationships

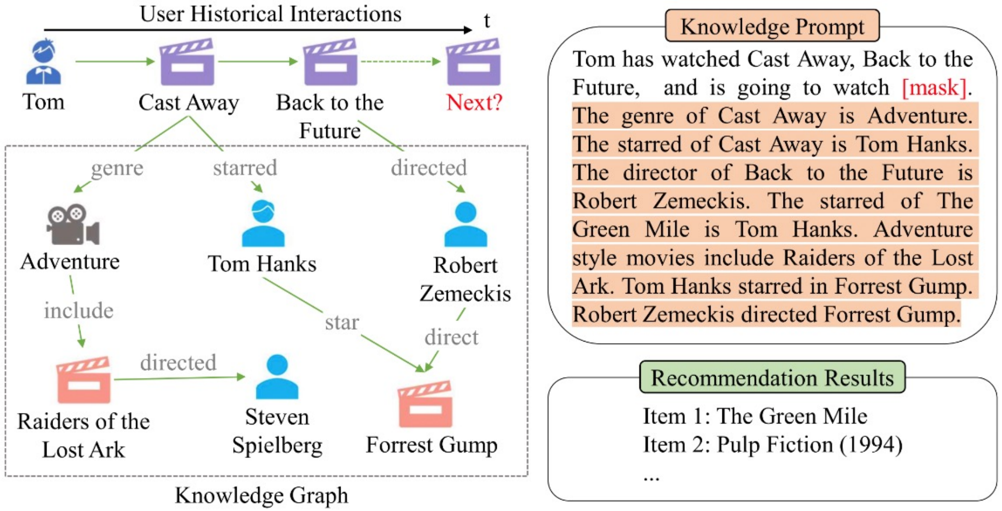

Zhai, Jianyang, et al. "Knowledge Prompt-tuning for Sequential Recommendation." *ACM MM.* (2023).

Large Language Models: Knowledge Prompt-tuning [KP4SR, MM '23]

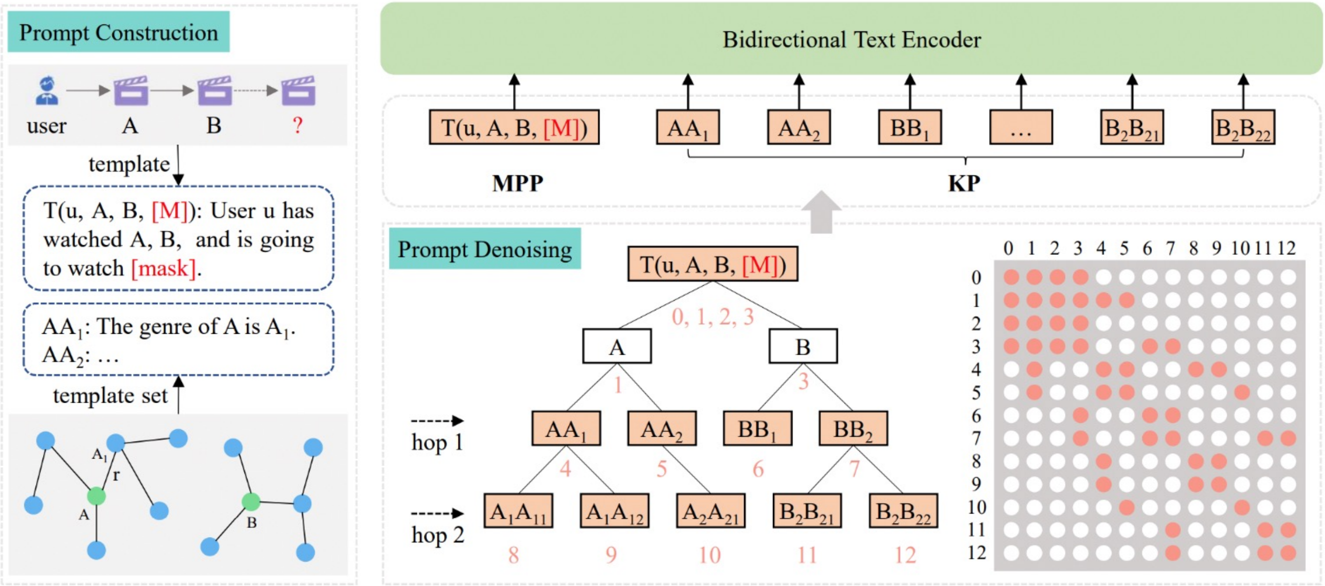

- Prompt construction
  - Convert user behaviors into masked personalized prompts
  - Convert triples in KGs to knowledge prompts
- Prompt denoising
  - Control the attention between input tokens
  - Implemented with a tree-structure

Zhai, Jianyang, et al. "Knowledge Prompt-tuning for Sequential Recommendation." *ACM MM.* (2023).

Large Language Models: Knowledge Prompt-tuning [KP4SR, MM ’23]

- LM-based methods surpassed traditional UBM methods
- KP4SR improves the baseline by a large margin

Methods	books				music				movies			
	NDCG@5	NDCG@10	HR@5	HR@10	NDCG@5	NDCG@10	HR@5	HR@10	NDCG@5	NDCG@10	HR@5	HR@10
Caser	0.0220	0.0294	0.0356	0.0587	0.0165	0.0232	0.0271	0.0477	0.0309	0.0462	0.052	0.0999
GRU4Rec	0.0235	0.0317	0.0380	0.0635	0.0222	0.0317	0.0374	0.0668	0.0378	0.0554	0.0643	0.1191
BERT4Rec	0.0204	0.0282	0.0323	0.0567	0.0242	0.0356	0.0426	0.0781	0.0328	0.0488	0.056	0.1062
SASRec	0.0254	0.0362	0.0466	0.0803	0.0327	0.047	0.0634	0.1078	0.0312	0.0459	0.0538	0.0994
GRU4RecF	0.0240	0.0321	0.0381	0.0633	0.0266	0.0377	0.0441	0.0788	0.0372	0.0551	0.0644	0.1204
GRU4RecKG	0.0233	0.0314	0.0373	0.0625	0.0222	0.0313	0.0380	0.0664	0.0374	0.0562	0.0650	0.1237
KSR	0.0240	0.0317	0.0383	0.0623	0.0330	0.0411	0.0504	0.0757	0.0394	0.0574	0.0679	0.1242
FDSA	0.0221	0.0309	0.0355	0.0631	0.0185	0.0261	0.0304	0.0539	0.0354	0.0523	0.0604	0.1132
SASRecF	0.0238	0.0319	0.0379	0.0631	0.0310	0.0418	0.0503	0.0839	0.0294	0.0441	0.0503	0.0964
S^3Rec	0.0249	0.0356	0.0452	0.0783	0.0301	0.0443	0.0524	0.0968	0.0306	0.0461	0.0536	0.1019
DIF-SR	0.0298	0.0416	0.0584	<u>0.0948</u>	0.0573	0.0678	**0.1110**	**0.1433**	0.0492	0.0689	0.0875	**0.1489**
P5	<u>0.0433</u>	<u>0.0501</u>	<u>0.0604</u>	0.0813	<u>0.0815</u>	<u>0.0879</u>	0.0994	0.1193	<u>0.0618</u>	<u>0.0738</u>	<u>0.0888</u>	0.1261
KP4SR	**0.0609**	**0.0691**	**0.0824**	**0.1077**	**0.0906**	**0.0975**	<u>0.1108</u>	<u>0.1319</u>	**0.0755**	**0.0891**	**0.1058**	<u>0.1481</u>

Zhai, Jianyang, et al. "Knowledge Prompt-tuning for Sequential Recommendation." *ACM MM.* (2023).

Large Language Models: Lifelong Behavior Comprehension [RELLA ArXiv ’23]

- LLMs fail to extract the useful information from long user behavior sequence

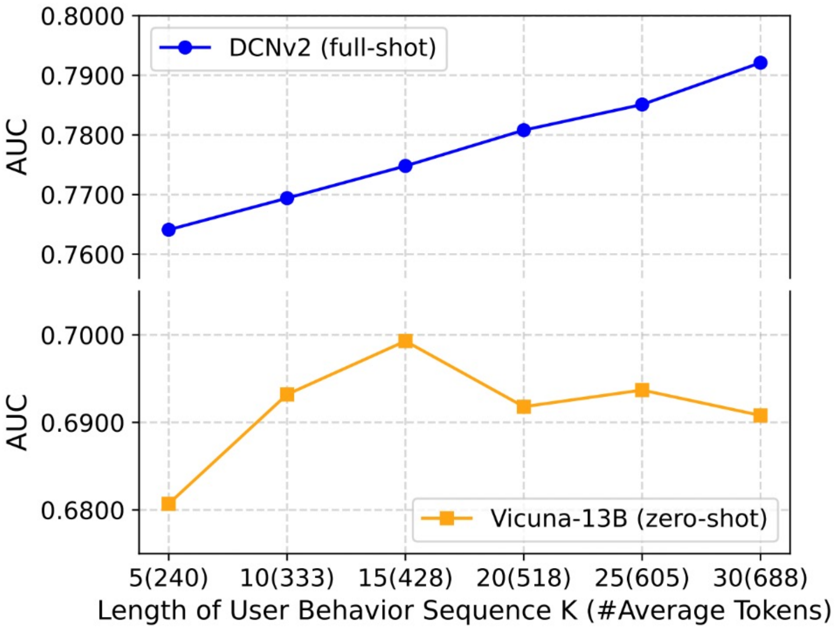

Lin, Jianghao, et al. "ReLLa: Retrieval-enhanced Large Language Models for Lifelong Sequential Behavior Comprehension in Recommendation." *arXiv preprint arXiv:2308.11131* (2023).

Large Language Models: Lifelong Behavior Comprehension [RELLA ArXiv ’23]

- ReLLa:
  - Apply LLM to obtain the semantic vectors
  - Retrieve top-K behaviors w.r.t. the target item

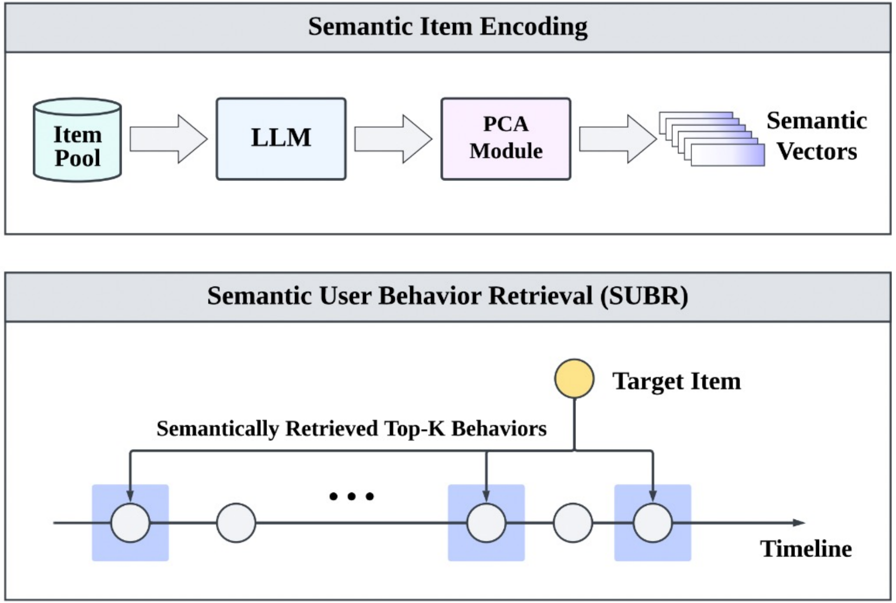

Lin, Jianghao, et al. "ReLLa: Retrieval-enhanced Large Language Models for Lifelong Sequential Behavior Comprehension in Recommendation." *arXiv preprint arXiv:2308.11131* (2023).

Large Language Models: Lifelong Behavior Comprehension [RELLA ArXiv ’23]

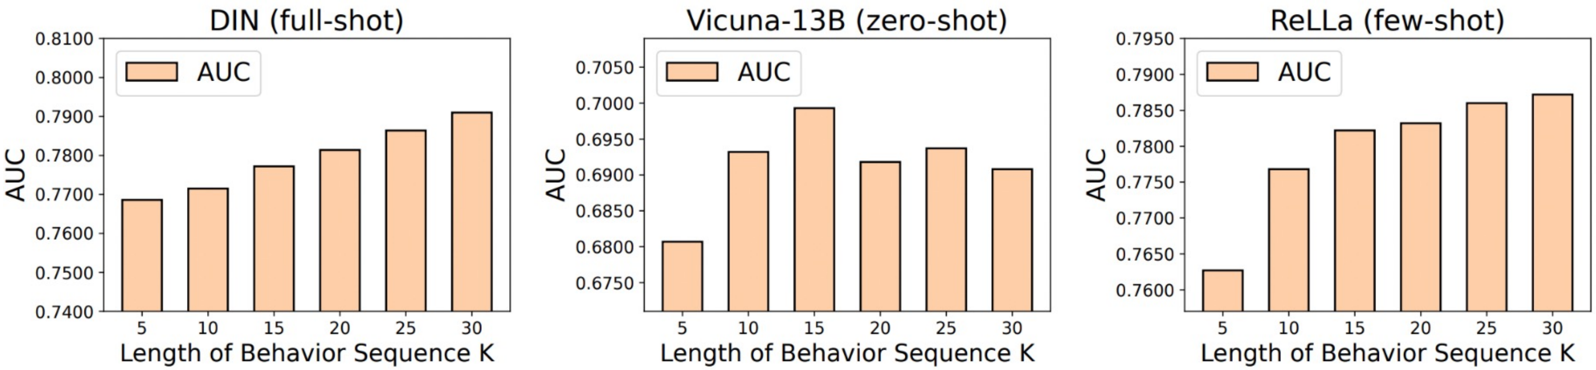

Model		MovieLens-1M AUC	Log Loss	ACC	Rel.Impr
Zero-shot	Vicuna-7B	0.6739	0.9510	0.5644	4.06%
	Vicuna-13B	0.6993	0.6291	0.6493	0.29%
	ReLLa (Ours)	**0.7013***	**0.6250***	**0.6507***	-
Few-shot	DCNv2	0.7265	0.6237	0.6630	8.36%
	DIN	0.7269	0.6153	0.6600	8.30%
	TALLRec	0.7596	0.5859	0.6963	3.63%
	ReLLa (Ours)	**0.7872***	**0.5625***	**0.7059***	-
Full-shot	DIN	0.7962	0.5425	0.7252	-

Lin, Jianghao, et al. "ReLLa: Retrieval-enhanced Large Language Models for Lifelong Sequential Behavior Comprehension in Recommendation." *arXiv preprint arXiv:2308.11131* (2023).

Conclusion

- **Conventional UBM:** RNN, CNN, and Attention-based techniques
- **Longer length:** memory-augmented methods, and retrieval methods
- **Increase multiplicity:** click, purchase, skip…
- **Growing heterogeneity:** item attribute, multi-modal information, time information.
- **New trends and techniques:** reinforcement learning and large language models.

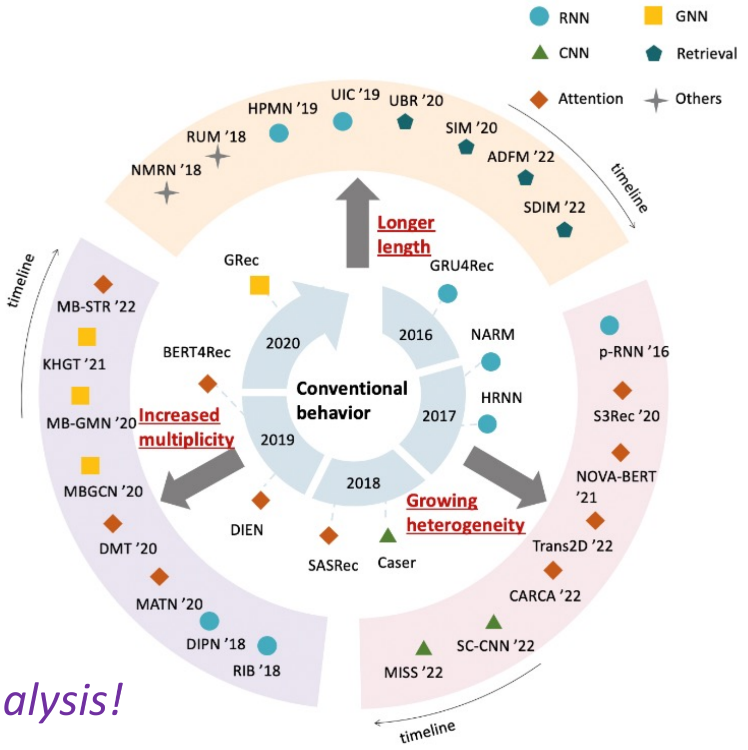

Towards complicated, complete, human-like analysis!

Future Prospects

- **Deeper information fusion:** combine long and multi-type behavior modeling.
- **More efficient learning method:** manage the trade-off between effectiveness and efficiency for online serving.
- **More interpretable user representations:** improve interpretability.
- **More advanced techniques:** build pre-trained unified models.

Thank you!