
Learning to Rehearse in Long Sequence Memorization

Zhu Zhang^{*1,2} Chang Zhou^{*2} Jianxin Ma² Zhijie Lin¹ Jingren Zhou² Hongxia Yang² Zhou Zhao¹

Abstract

Existing reasoning tasks often have an important assumption that the input contents can be always accessed while reasoning, requiring unlimited storage resources and suffering from severe time delay on long sequences. To achieve efficient reasoning on long sequences with limited storage resources, memory augmented neural networks introduce a human-like write-read memory to compress and memorize the long input sequence in one pass, trying to answer subsequent queries only based on the memory. But they have two serious drawbacks: 1) they continually update the memory from current information and inevitably forget the early contents; 2) they do not distinguish what information is important and treat all contents equally. In this paper, we propose the Rehearsal Memory (RM) to enhance long-sequence memorization by self-supervised rehearsal with a history sampler. To alleviate the gradual forgetting of early information, we design self-supervised rehearsal training with recollection and familiarity tasks. Further, we design a history sampler to select informative fragments for rehearsal training, making the memory focus on the crucial information. We evaluate the performance of our rehearsal memory by the synthetic bAbI task and several downstream tasks, including text/video question answering and recommendation on long sequences.

1. Introduction

In recent years, the tremendous progress of neural networks has enabled machines to perform reasoning given the input contents X and a query Q , e.g., infer the answer of given questions from the text/video stream in text/video question answering (Seo et al., 2016; Jin et al., 2019a; Le et al.,

2020b), or predict whether a user will click the given item based on the user behavior sequence in recommender systems (Ren et al., 2019; Pi et al., 2019; Zhang et al., 2021). Studies that achieve top performances at such reasoning tasks usually have an important assumption that the raw input contents X can be always accessed while answering the query Q . In this setting, the complex interaction between X and Q can be designed to extract query-relevant information from X with little loss, such as co-attention interaction (Xiong et al., 2016; Jin et al., 2019b). Though these methods (Seo et al., 2016; Le et al., 2020b) can effectively handle these reasoning tasks, they require unlimited storage resources to hold the original input X . Further, they have to encode the whole input contents and develop the elaborate interaction from scratch, which are time-consuming. This is not acceptable for online services that require instant response such as recommender systems, as the input sequence becomes extremely long (Ren et al., 2019).

To achieve efficient reasoning on long sequences with limited storage resources, memory augmented neural networks (MANNs) (Graves et al., 2014; 2016) introduce a write-read memory M with fixed-size capacity (size much smaller than $|X|$) to compress and remember the input contents X . In the inference phase, they can capture query-relevant clues directly from the memory M , i.e., the raw input X is *not* needed at the time of answering Q . This procedure is very similar to the daily situation of our human beings, i.e., we may not know the tasks Q that we will answer in the future when we are experiencing current events, but we have the instincts to continually memorize our experiences within the limited memory capacity, from which we can rapidly recall and draw upon past events to guide our behaviors given the present tasks (Moscovitch et al., 2016; Baddeley, 1992). Such human-like memory-based methods bring three benefits for long-sequence reasoning: 1) storage efficiency: we only need to maintain the limited memory M rather than X ; 2) reasoning efficiency: inference over M and Q is more lightweight than inference over X and Q from scratch; 3) high reusability: the maintained memory M can be reused for any query Q .

However, existing MANNs have two serious drawbacks for memory-based long-sequence reasoning. First, these approaches ignore the long-term memorization ability of the memory. They learn how to maintain the memory M only

^{*}Equal contribution ¹Zhejiang University, China ²DAMO Academy, Alibaba Group, China. Correspondence to: Zhou Zhao <zhaozhou@zju.edu.cn>.

by back-propagated losses to the final answer and do not design any specific training target for long-term memorization, which inevitably lead to the gradual forgetting of early contents (Le et al., 2019a). That is, when dealing with the long input sequences, these approaches may fail to answer the query relevant to early contents due to the lack of long-term memorization. Second, determining what to remember in the memory with limited capacity is crucial to retain sufficient clues for subsequent Q . This is especially challenging since the information compression procedure in M is totally not aware of Q . But existing MANNs do not distinguish what information is important and treat all contents equally. Thus, due to lack of information discrimination, these approaches may store too much meaningless information but lose vital evidence for subsequent reasoning.

In this paper, we propose the Rehearsal Memory (RM) to enhance long-sequence memorization by self-supervised rehearsal with a history sampler. To overcome gradual forgetting of early information and increase the generalization ability of the memorization technique, we develop two extra self-supervised rehearsal tasks to recall the recorded history contents from the memory. The two tasks are inspired by the observation that human beings can recall details nearby some specific events and distinguish whether a series of events happened in the history, which respectively correspond to two different memory processes revealed in cognitive, neuropsychological, and neuroimaging studies, i.e., *recollection* and *familiarity* (Yonelinas, 2002; Moscovitch et al., 2016). Concretely, the recollection task aims to predict the masked items in history fragments H , which are sampled from the original input stream and parts of items are masked as the prediction target. This task tries to endow the memory with the *recollection* ability that enables one to relive past episodes. And the familiarity task tries to distinguish whether a historical fragment H ever appears in the input stream, where we directly sample positive fragments from the input stream and replace parts of the items in positive ones as negative fragments. This task resembles the *familiarity* process that recognizes experienced events or stimulus as familiar.

To make the rehearsal memory have the ability of remembering the crucial information, we further train an independent history sampler to select informative fragments H for self-supervised rehearsal training. Similar to the teacher-student architecture in knowledge distillation (Hinton et al., 2015), we expect the history sampler (i.e. the teacher) to capture the characteristic of important fragments in the current environment and guide the rehearsal memory (i.e. the student) to remember task-relevant clues. Concretely, we independently train a conventional reasoning model that can access raw contents X while answering the query Q as the history sampler. The model contains the attention interaction between history fragments H and the query Q , where the attention

weight can be regarded as the importance of each fragment. After training, the history sampler can select the vital fragments based on the attention weights for self-supervised rehearsal training. This is similar to the procedure where human beings learn to memorize meaningful experiences, i.e., we have gone through a lot of tasks to slowly understand which information is likely to be used in future tasks and pay more attention to them during memorization (Moscovitch et al., 2016).

In conclusion, we propose the self-supervised memory rehearsal to enhance the long-sequence memorization for subsequent reasoning. We design the self-supervised recollection and familiarity tasks to solve *how to rehearse*, which can alleviate the gradual forgetting of early information. Further, we adopt a history sampler to decide *what to rehearse*, which guides the memory to remember critical information. We illustrate the ability of our rehearsal memory via the synthetic bAbI task and several downstream tasks, including text/video question answering and recommendation on long sequences.

2. Related Works

Memory augmented neural networks (MANNs) introduce the external memory to store and access the past contents by differentiable write-read operators. Neural Turing Machine (NTM) (Graves et al., 2014) and Differentiable Neural Computer (DNC) (Graves et al., 2016) are the typical MANNs for human-like memorization and reasoning, whose inferences rely only on the memory with limited capacity rather than starting from the original input. In this line of research, Rae et al. (2016) adopt the sparse memory accessing to reduce computational cost. Csordás & Schmidhuber (2019) introduce the key/value separation problem of content-based addressing and adopt a mask for memory operations as a solution. Le et al. (2019b) manipulate both data and programs stored in memory to perform universal computations. And Santoro et al. (2018); Le et al. (2020a) consider the complex relational reasoning with the information they remember.

However, these works exploit MANNs mainly to help capture complex dependencies in dealing with input sequences, but do not explore the potential of MANNs in the field of memory-based long-sequence reasoning. They learn how to maintain the memory only by back-propagated losses to the final answer but do not design specific training target for long-term memorization, inevitably incurring gradual forgetting of early contents during memorizing long sequences (Le et al., 2019a). Recently, there are a few works trying to alleviate this problem. Le et al. (2019a) propose to measure “remember” ability by the final gradient on the early input, and adopt a uniform writing operation on the memory to balance between maximizing memorization and forgetting. Munkhdalai et al. (2019) design the meta-learned

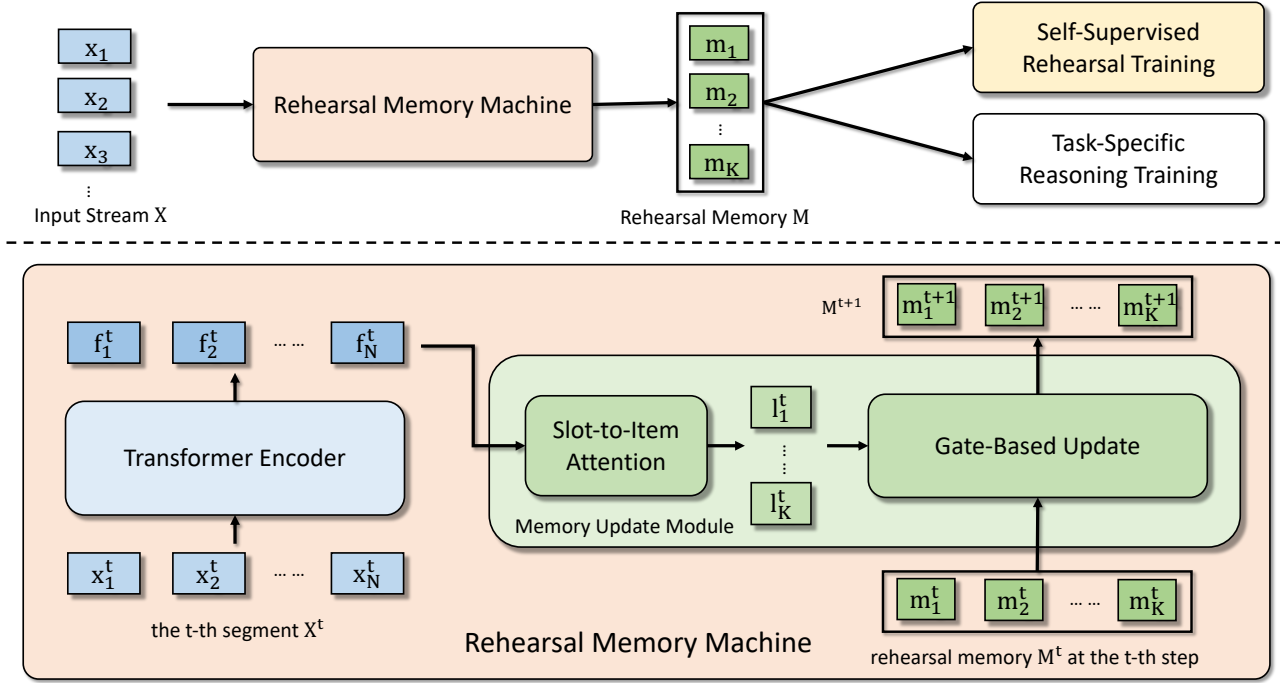


Figure 1. The Framework of Rehearsal Memory.

neural memory instead of the conventional array-structured memory and memorize the current and past information by reconstructing the written values via the memory function. Besides, Compressive Transformer (Rae et al., 2019) maps the past memory to a smaller compressed memory for long-range sequence learning, where the compressed memory preserves much original information by a high compression rate. But considering the compressed memory is implemented by a FIFO queue, it will completely forget the contents beyond a fixed range.

Our approach is different and parallel to these techniques, we try to enhance long-sequence memorization by self-supervised memory rehearsal, i.e., recall the recorded history contents from the memory to overcome gradual forgetting of early information. We design the recollection task to enable the memory to relive past episodes and adopt the familiarity task to make the memory recognize experienced events. A recent work (Park et al., 2020) also introduces a self-supervised memory loss to ensure how well the current input is written to the memory, but it only focuses on remembering the current information and ignoring the long-term memorization. Further, compared to previous techniques that have no assumptions on what behavior will be remembered the most, we propose a history sampler to distinguish the characteristic of important fragments in the current environment and guide the memory rehearsal to remember task-relevant clues.

3. Rehearsal Memory

3.1. Problem Formulation

Given the input stream $\mathbf{X} = \{x_1, x_2, \dots\}$ and a query \mathbf{Q} , the directly reasoning methods (Seo et al., 2016; Le et al., 2020b) learn the model $\mathcal{T}(\mathbf{X}, \mathbf{Q})$ to predict the answer \mathbf{A} . These is an important assumption that the input stream \mathbf{X} can be always accessed while reasoning. And complex interaction between \mathbf{X} and \mathbf{Q} can be designed to extract query-relevant information in $\mathcal{T}(\mathbf{X}, \mathbf{Q})$. Obviously, these methods have to store the original input \mathbf{X} and infer the answer \mathbf{A} from scratch when the query \mathbf{Q} is known. In this paper, we explore the human-like memory-based reasoning on long sequences, where we compress the input stream \mathbf{X} into a fixed-size memory $\mathbf{M} = \{\mathbf{m}_k\}_{k=1}^K$ with K memory slots and then infer the answer \mathbf{A} for any relevant query \mathbf{Q} by $\mathbf{A} = \mathcal{R}(\mathbf{M}, \mathbf{Q})$. Here we only need to store the compressed memory \mathbf{M} , which can be updated in real-time and reused for a series of queries. Since the slot number K in the memory is irrelevant to the input length $|\mathbf{X}|$, this setting only requires $O(1)$ storage space rather than $O(|\mathbf{X}|)$ in directly reasoning model $\mathcal{T}(\mathbf{X}, \mathbf{Q})$.

As shown in Figure 1, we apply a rehearsal memory machine $\mathcal{G}_\Theta(\mathbf{X})$ to compress the input stream \mathbf{X} into rehearsal memory $\mathbf{M} = \{\mathbf{m}_k\}_{k=1}^K$ with K memory slots. During the training stage, we simultaneously develop self-supervised rehearsal training and task-specific reasoning training based on the memory \mathbf{M} . For self-supervised rehearsal training,

we develop a rehearsal model $\mathcal{H}_\xi(\mathbf{M}, \mathbf{H})$ to reconstruct the masked history fragments (recollection task) and distinguish positive history fragments from negative ones (familiarity task), where \mathbf{H} means the critical history fragments that are selected by the history sampler $\mathcal{S}_\Psi(\mathbf{Q}, \mathbf{X})$. For task-specific reasoning training, we develop the task-specific reason model $\mathcal{R}_\Omega(\mathbf{M}, \mathbf{Q})$ to answer the given query \mathbf{Q} . *During the testing stage*, we maintain the rehearsal memory $\mathbf{M} = \mathcal{G}_\Theta(\mathbf{X})$ from the stream \mathbf{X} and then infer the answer \mathbf{A} for any relevant query \mathbf{Q} by $\mathbf{A} = \mathcal{R}_\Omega(\mathbf{M}, \mathbf{Q})$, where the rehearsal model $\mathcal{H}_\xi(\mathbf{M}, \mathbf{H})$ and the history sampler $\mathcal{S}_\Psi(\mathbf{Q}, \mathbf{X})$ are no longer needed.

3.2. Rehearsal Memory Machine

We deal with the input stream \mathbf{X} from the segment level rather than item level, i.e., we cut the input sequence into fixed-length segments and memorize them into the rehearsal memory segment-by-segment. Compared to existing MANNs (Graves et al., 2014; 2016), which store the input stream item-by-item orderly with a RNN-based controller, our segment-level memorization can further capture the bi-directional context of each item and improve the modeling efficiency. We denote the t -th segment as $\mathbf{X}^t = \{\mathbf{x}_n^t\}_{n=1}^N$ with N items and the current memory as $\mathbf{M}^t = \{\mathbf{m}_k^t\}_{k=1}^K$, where we have recorded $t-1$ segments in \mathbf{M}^t . The \mathbf{x}_n^t and \mathbf{m}_k^t have the same dimension d_x .

We first model the t -th segment by a Transformer encoder (Vaswani et al., 2017) and obtain the sequence features $\mathbf{F}^t = \{\mathbf{f}_n^t\}_{n=1}^N$ with dimension d_x . After it, we apply a memory update module to write \mathbf{F}^t into \mathbf{M}^t . We apply a slot-to-item attention to align the sequence features to slot features in the current memory \mathbf{M}^t , and then develop the gate-based update. Concretely, we first calculate the slot-to-item attention matrix where each element means the relevance of a slot-item pair, and then learn aligned features $\mathbf{L}^t = \{\mathbf{l}_k^t\}_{k=1}^K$ for each slot, given by

$$\begin{aligned} \alpha_{kn}^t &= \mathbf{w}_a^\top \tanh(\mathbf{W}_1^a \mathbf{m}_k^t + \mathbf{W}_2^a \mathbf{f}_n^t + \mathbf{b}^a), \\ \hat{\alpha}_{kn}^t &= \frac{\exp(\alpha_{kn}^t)}{\sum_{j=1}^K \exp(\alpha_{jn}^t)}, \mathbf{l}_k^t = \sum_{n=1}^N \hat{\alpha}_{kn}^t \mathbf{f}_n^t, \end{aligned} \quad (1)$$

where $\mathbf{W}_1^a \in \mathbb{R}^{d_{model} \times d_x}$, $\mathbf{W}_2^a \in \mathbb{R}^{d_{model} \times d_x}$ and $\mathbf{b}^a \in \mathbb{R}^{d_{model}}$ are the projection matrices and bias. \mathbf{w}_a^\top is the row vector. Next, the k -th slot feature \mathbf{m}_k^t is updated with its aligned feature \mathbf{l}_k^t based on a GRU unit with d_x -d hidden states, given by

$$\mathbf{m}_k^{t+1} = \text{GRU}(\mathbf{m}_k^t, \mathbf{l}_k^t), \quad (2)$$

where \mathbf{l}_k^t is the current input of the GRU unit and \mathbf{m}_k^t is the hidden state at the t -th step. And \mathbf{m}_k^{t+1} is the new slot feature after the gate-based update. After memorizing T segments, we can obtain rehearsal memory \mathbf{M}^{T+1} and we denote it by \mathbf{M} for convenience.

3.3. Self-Supervised Rehearsal Training

As shown in Figure 2, based on the maintained memory \mathbf{M} , we apply the memory rehearsal technique to enhance long-sequence memorization. We first design the self-supervised recollection and familiarity tasks to solve *how to rehearse*, which can alleviate the gradual forgetting of early information. We next adopt a history sampler to decide *what to rehearse*, which guides the memory to remember critical task-relevant clues.

3.3.1. HOW TO REHEARSE: SELF-SUPERVISED RECOLLECTION AND FAMILIARITY TASKS

We design the rehearsal model $\mathcal{H}_\xi(\mathbf{M}, \mathbf{H})$ with recollection and familiarity tasks. The recollection task reconstructs the masked positive history fragments to enable the memory to relive past episodes. And the familiarity task tries to distinguish positive history fragments from negative ones for making the memory recognize experienced events.

First, we apply an independent history sampler to select the B segments from the input stream as the history fragment set $\mathbf{H} = \{H^b\}_{b=1}^B$, which is illustrated in the next section. Each fragment $H^b = \{h_1^b, h_2^b, \dots, h_N^b\}$ contains N items and each item h_* corresponds to a feature \mathbf{x}_* . For the b -th fragment, we randomly mask 50% of items in the fragment and add an especial item [cls] at the beginning to obtain the masked positive history fragment $H^{b+} = \{h_{[cls]}^{b+}, h_1^{b+}, h_{[m_1]}^{b+}, \dots, h_N^{b+}\}$, where $h_{[m_1]}^{b+}$ means the first masked item. In order to guarantee that the model $\mathcal{H}_\xi(\mathbf{M}, \mathbf{H})$ reconstructs the masked fragment by utilizing the maintained memory \mathbf{M} rather than only relying on fragment context, we set the mask ratio to 50% instead of 15% in BERT (Devlin et al., 2019). Moreover, we construct the masked negative history fragment $H^{b-} = \{h_{[cls]}^{b-}, h_1^{b-}, h_{[m_1]}^{b-}, \dots, h_N^{b-}\}$ by replacing 50% of unmasked items in the positive fragment, where the replacement items are sampled from other input stream to make the negative fragment distinguishable. Here we construct the positive fragment set $\mathbf{H}^+ = \{H^{b+}\}_{b=1}^B$ and corresponding negative fragment set $\mathbf{H}^- = \{H^{b-}\}_{b=1}^B$ from the original fragment set \mathbf{H} . Next, we adopt a bidirectional Transformer decoder (Vaswani et al., 2017) without the future masking to model each history fragment H^{b+}/H^{b-} . In the decoder, each history item can interact with all other items in the fragment. The rehearsal memory \mathbf{M} is input to the ‘‘encoder-decoder multi-head attention sub-layer’’ in each decoder layer, where the queries come from the previous decoder layer and the memory slots are regarded as the keys and values. This allows each item in the decoder to attend over all slot features in the memory \mathbf{M} . Finally, we obtain the features $\{\mathbf{r}_{[cls]}^{b+/b-}, \mathbf{r}_1^{b+/b-}, \mathbf{r}_{[m_1]}^{b+/b-}, \dots, \mathbf{r}_N^{b+/b-}\}$ where each $\mathbf{r}_*^{b+/b-}$ has the dimension d_x .

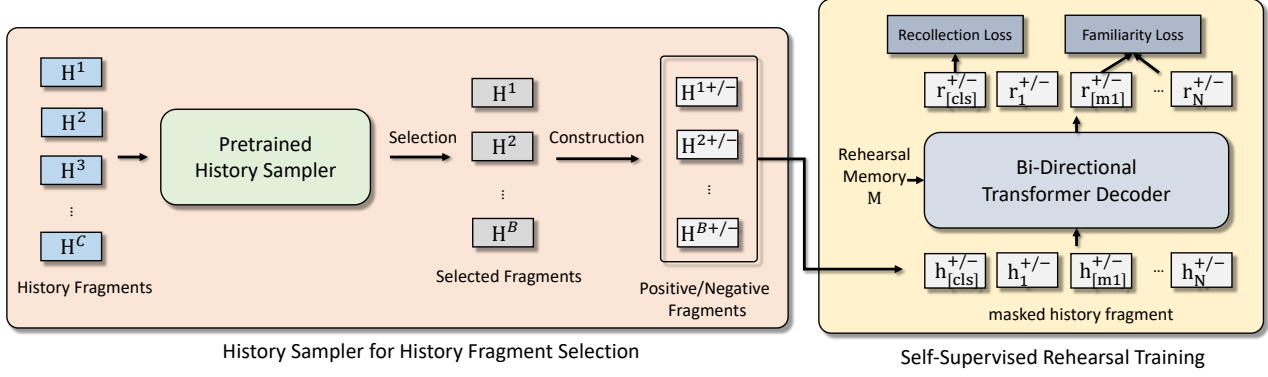


Figure 2. Self-Supervised Rehearsal Training with the History Sampler.

Recollection Task. We first predict the masked items of positive history fragments to build the item-level reconstruction for the recollection task. Considering there are too many item types, we apply the contrastive training (He et al., 2020; Chen et al., 2020; Zhang et al., 2020) based on the ground truth and other sampled items. For the $N/2$ masked items, we compute the recollection loss for the b -th fragment by

$$\mathcal{L}_i^b = \log \frac{\exp(\mathbf{r}_{[m_i]}^{b+} \cdot \mathbf{y}_i)}{\exp(\mathbf{r}_{[m_i]}^{b+} \cdot \mathbf{y}_i) + \sum_{j=1}^J \exp(\mathbf{r}_{[m_i]}^{b+} \cdot \mathbf{y}_j)}, \quad (3)$$

$$\mathcal{L}_{rec}^b = -\frac{2}{N} \sum_{i=1}^{N/2} \mathcal{L}_i^b$$

where $\mathbf{y}_i \in \mathbb{R}^{d_x}$ is the feature of ground truth of the i -th masked item, $\mathbf{y}_j \in \mathbb{R}^{d_x}$ is the feature of sampled items and $\mathbf{r}_{[m_i]}^{b+} \cdot \mathbf{y}_*$ is the inner product of two features.

Familiarity Task. Next, we predict whether the masked history fragment ever appears in the current input stream, i.e. distinguish positive history fragments from negative ones. This training objective makes the memory learn the ability of recognizing experienced events. Concretely, we project each feature $\mathbf{r}_{[cls]}^{b+/b-}$ into a confident score $s^{b+/b-} \in (0, 1)$ by a linear layer with the sigmoid activation, and calculate the familiarity loss by

$$\mathcal{L}_{fam}^b = -\log(s^{b+}) + \log(1 - s^{b-}), \quad (4)$$

where \mathcal{L}_{fam}^b is the familiarity loss for the b -th pair of positive and negative fragments in $\mathbf{H}^+/\mathbf{H}^-$.

3.3.2. WHAT TO REHEARSE: HISTORY SAMPLER

Existing MANNs (Graves et al., 2014; 2016) often have no assumptions on what information needs to be remembered the most. But due to the limited capacity of the memory, it is crucial to distinguish what contents are important for

subsequent inference and pay more attention to them during memorization. Thus, we further train a history sampler $\mathcal{S}_\Psi(\mathbf{Q}, \mathbf{X})$ to select informative history fragments \mathbf{H} for self-supervised rehearsal training, which is independent to the memory machine $\mathbf{M} = \mathcal{G}_\Theta(\mathbf{X})$. In knowledge distillation (Hinton et al., 2015), the teacher model can access the privileged information and transfer the knowledge to the student model. Similar to it, our history sampler, which is the teacher and can access raw contents \mathbf{X} while answering the query \mathbf{Q} , learns the ability of distinguishing task-relevant important fragments and guides the rehearsal memory (i.e., the student) to remember critical clues.

Concretely, we independently train a directly reasoning model as the history sampler $\mathcal{S}_\Psi(\mathbf{Q}, \mathbf{X})$, which can access raw contents \mathbf{X} while answering the query \mathbf{Q} . We first cut the entire input \mathbf{X} into C history fragments $\{H^c\}_{c=1}^C$ just like the rehearsal memory machine, where each fragment contains N items. Next, we obtain the fragment features $\{\mathbf{h}^c\}_{c=1}^C$ by averaging the item features in each fragment. After it, we develop the attention-based reasoning for the query \mathbf{Q} on these fragment features. The query feature $\mathbf{q} \in \mathbb{R}^{d_{model}}$ is modeled by the task-specific encoder in different downstream tasks, which is introduced in Section A of the supplementary material. Given the query feature \mathbf{q} and fragment features $\{\mathbf{h}^c\}_{c=1}^C$, we conduct the attention method to aggregate query-relevant clues from fragments, given by

$$\beta^c = \mathbf{w}_h^\top \tanh(\mathbf{W}_1^h \mathbf{q} + \mathbf{W}_2^h \mathbf{h}^c + \mathbf{b}^h),$$

$$\hat{\beta}^c = \frac{\exp(\beta^c)}{\sum_{j=1}^C \exp(\beta^j)}, \quad \mathbf{e} = \sum_{c=1}^C \hat{\beta}^c \mathbf{h}^c, \quad (5)$$

where $\mathbf{W}_1^h \in \mathbb{R}^{d_{model} \times d_{model}}$, $\mathbf{W}_2^h \in \mathbb{R}^{d_{model} \times d_x}$ and $\mathbf{b}^h \in \mathbb{R}^{d_{model}}$ are the projection matrices and bias. And \mathbf{w}_h^\top is the row vector. We then obtain the reasoning feature $\mathbf{a} = [\mathbf{e}; \mathbf{q}]$ by concatenating the query and query-relevant fragment features, and design the final reasoning layer for different tasks, shown in Section A of the supplementary

material. After independent training, the attention weights $\{\beta^c\}_{c=1}^C$ can be regarded as the importance score of each fragment for the query \mathbf{Q} . Thus, we can sample the vital history fragments with high attention weights for each (\mathbf{X}, \mathbf{Q}) pair. Specifically, to guarantee the sampled fragments appear in the entire input stream, we select $B/2$ fragments from $\{H^c\}_{c=1}^{C/2}$ with large weights and choose another $B/2$ fragments from $\{H^c\}_{c=C/2}^C$ to constitute the fragment set $\mathbf{H} = \{H^b\}_{b=1}^B$. The final recollection loss \mathcal{L}_{rec} and familiarity loss \mathcal{L}_{fam} are computed by

$$\mathcal{L}_{rec} = \mathbb{E}_{b \sim \mathcal{S}_\Psi} [\mathcal{L}_{rec}^b], \quad \mathcal{L}_{fam} = \mathbb{E}_{b \sim \mathcal{S}_\Psi} [\mathcal{L}_{fam}^b]. \quad (6)$$

3.4. Task-Specific Reasoning Training

Besides self-supervised rehearsal training, we simultaneously develop task-specific reasoning training. For several downstream tasks, we propose different task-specific reason model $\mathcal{R}_\Omega(\mathbf{M}, \mathbf{Q})$ based on the memory \mathbf{M} . Here we adopt the simple and mature components in the reason model for a fair comparison. The details are introduced in Section A of the supplementary material. Briefly, we first learn the query representation \mathbf{q} by a task-specific encoder and then perform the multi-hop attention-based reasoning. Finally, we obtain the reason loss \mathcal{L}_r from $\mathcal{R}_\Omega(\mathbf{M}, \mathbf{Q})$.

Eventually, we combine the rehearsal and reason losses to train our model, given by

$$\mathcal{L}_{rm} = \lambda_1 \mathcal{L}_{rec} + \lambda_2 \mathcal{L}_{fam} + \lambda_3 \mathcal{L}_r, \quad (7)$$

where λ_1 , λ_2 and λ_3 are applied to adjust the balance of three losses.

4. Experiments

In this section, we first verify our rehearsal memory on the widely-used short-sequence reasoning task bAbI. Next, we mainly compare our approach with diverse baselines on several long-sequence reasoning tasks. We then perform ablation studies on the memory rehearsal techniques and analyze the impact of crucial hyper-parameters.

4.1. Experiment Setting

Model Setting. We first introduce the common model settings for all downstream tasks. We set the layer number of the Transformer encoder and bi-directional Transformer decoder to 3. The head number in Multi-Head Attention is set to 4. We set λ_1 , λ_2 and λ_3 to 1.0, 0.5 and 1.0, respectively. The number B of history fragments is set to 6. During training, we apply an Adam optimizer (Duchi et al., 2011) to minimize the multi-task loss \mathcal{L}_{rm} , where the initial learning rate is set to 0.001.

Baseline. We compare our rehearsal memory with the directly reasoning methods and the memory-based reasoning

Table 1. Performance Comparisons for Synthetical bAbI Task: mean \pm std. and best error over 10 runs.

Method	mean \pm std error	best error
DNC	16.7 \pm 7.6	3.8
NUTM	5.6 \pm 1.9	3.3
DMSDNC	1.53 \pm 1.33	0.16
STM	0.39 \pm 0.18	0.15
CT	0.81 \pm 0.26	0.34
RM	0.33 \pm 0.15	0.12

approaches. The directly reasoning baselines are different in downstream tasks and the memory-based baselines mainly are DNC (Graves et al., 2016), NUTM (Le et al., 2019b), DMSDNC (Park et al., 2020), STM (Le et al., 2020a) and Compressive Transformer (CT) (Rae et al., 2019). For a fair comparison, we modify the reasoning module of memory-based baselines to be consistent with our rehearsal memory, i.e. we conduct multi-hop attention-based reasoning based on the built memory. And the number of memory slots in these baselines is also set to K . Besides, we set the core number of NUTM to 4, the query number of STM to 8 and the memory block number of DMSDNC to 2. As for CT, the layer number of the Transformer is set to 3 as our rehearsal memory and the compression rate is set 5.

4.2. Rehearsal Memory on Short-Sequence Reasoning

The bAbI dataset (Weston et al., 2015) is a synthetic text question answering benchmark and widely applied to evaluate the memorization and reasoning performance of MANNs. This dataset contains 20 reasoning tasks and requires to be solved with one common model. Although most of these tasks only give short-sequence text input (less than 100 words) and existing methods (Park et al., 2020; Le et al., 2020a) have solved these tasks well, we still compare our rehearsal memory with other memory-based baselines to verify the short-sequence reasoning performance. We set the d_x and d_{model} to 128. The number K of memory slots is set to 20. And we naturally take each sentence in input texts as a segment and the maximum length N of segments is set to 15. Due to limited word types in this dataset, we sample all other words as negative items in \mathcal{L}_{rec} .

The results are summarized in Table 1. The RM model solves these bAbI tasks with near-zero error and outperforms existing baselines at the mean and best error rate of 10 runs, verifying that our RM method can conduct effective memorization and reasoning on short-sequence tasks. In these methods, DNC, NUTM and DMSDNC model the input contents word-by-word. STM processes input texts as a sentence-level sequence. And CT and RM methods cut the input texts into sentences for modeling. From the results,

Table 2. Performance Comparisons for Long-Sequence Text Question Answering on NarrativeQA.

Method	Setting	Val MRR	Test MRR
AS Reader	Directly	26.9	25.9
E2E-MN	Directly	29.1	28.6
DNC	Memory-Based	25.8	25.2
NUTM	Memory-Based	27.7	27.2
DMSDNC	Memory-Based	28.1	27.5
STM	Memory-Based	27.2	26.7
CT	Memory-Based	28.7	28.3
RM	Memory-Based	29.4	28.7

we can find STM, CT and RM achieve lower error rates than other baselines, suggesting the importance of sentence-level modeling. Considering the bAbI tasks are close to being solved, we design another difficult synthetic task in Section B of the supplementary material to evaluate our RM model.

4.3. Rehearsal Memory on Long-Sequence Reasoning

We then compare our approach with diverse baselines on several long-sequence reasoning tasks.

4.3.1. LONG-SEQUENCE TEXT QUESTION ANSWERING

We apply the NarrativeQA dataset (Kočíšký et al., 2018) with long input contents for long-sequence text question answering. This dataset contains 1,572 stories and corresponding summaries generated by humans, where each summary contains more than 600 tokens on average. And there are 46,765 questions in total. We adopt the multi-choice form to answer the given question based on a summary, where other answers for questions associated with the same summary are regarded as answer candidates. We compute the mean reciprocal rank (MRR) as the metric, i.e., the rank of the correct answer among candidates. Besides the memory-based methods, we adopt directly reasoning model **AS Reader** (Kadlec et al., 2016) and **E2E-MN** (Sukhbaatar et al., 2015) as baselines. The AS Reader applies a pointer network to generate the answer and E2E-MN employs the end-to-end memory network to conduct multi-hop reasoning. For our rehearsal memory, we set the d_x and d_{model} to 256. The number K of memory slots is set to 20. We naturally take each sentence in summaries as a segment and the maximum length N of segments is set to 20. And we sample all other words as negative items in \mathcal{L}_{rec} .

We report the results in Table 2. Our RM method obtains the best performance among memory-based approaches, which demonstrates our self-supervised rehearsal training with the history sampler can effectively enhance long-sequence memorization and reasoning. Further, the RM model slightly

Table 3. Performance Comparisons for Long-Term Video Question Answering on ActivityNet-QA.

Method	Setting	Accuracy
E-VQA	Directly	25.2
E-MN	Directly	27.9
E-SA	Directly	31.8
HCRN	Directly	37.6
DNC	Memory-Based	30.3
NUTM	Memory-Based	33.1
DMSDNC	Memory-Based	32.4
STM	Memory-Based	33.7
CT	Memory-Based	35.4
RM	Memory-Based	36.3

outperforms early directly reasoning methods AS Reader and E2E-MN, showing the ability of efficient reasoning on long sequences with limited storage resources.

4.3.2. LONG-SEQUENCE VIDEO QUESTION ANSWERING

The ActivityNet-QA dataset (Yu et al., 2019) contains 5,800 videos from the ActivityNet (Caba Heilbron et al., 2015). The average video duration of this dataset is about 180s and is the longest in VQA datasets. We compare our method with four directly reasoning baselines, including three basic models **E-VQA**, **E-MN**, **E-SA** from (Yu et al., 2019) and the SOTA model **HCRN** (Le et al., 2020b). For our rehearsal memory, we set the d_x and d_{model} to 256. The number K of memory slots and length N of segments are both set to 20. And in \mathcal{L}_{rec} , we select 30 other frame features from the video as the sampled items.

As shown in Table 3, the RM method obtains a better performance than other memory-based baselines. Compared to the best baseline CT, our RM model further achieves the 0.9% absolute improvement, showing the effectiveness of our model designs and self-supervised rehearsal training. Moreover, the RM method outperforms the basic directly reasoning baselines E-VQA, E-MN and E-SA, but slightly worse than the SOTA method HCRN. This suggests our rehearsal memory can reduce the gap between memory-based and directly reasoning paradigms.

4.3.3. LIFELONG SEQUENCE RECOMMENDATION

The lifelong sequence recommendation (Ren et al., 2019) aims to predict whether the user will click a given item based on long sequences, thus it can be regarded as a long-sequence reasoning task. The XLong dataset (Ren et al., 2019) is sampled from the click logs on Alibaba. The length of historical behavior sequences in this dataset is 1000. We compare our method with four directly reasoning methods

Table 4. Performance Comparisons for Lifelong Sequence Recommendation on XLong.

Method	Setting	AUC
GRU4REC	Directly	0.8702
Caser	Directly	0.8390
RUM	Directly	0.8649
DIEN	Directly	0.8793
HPMN	Memory-Based	0.8645
MIMN	Memory-Based	0.8731
RM	Memory-Based	0.8817

Table 5. Ablation Results about the Rehearsal Losses and History Sampler.

Method	NarrativeQA		ActNet-QA	XLong
	Val	Test	Acc.	AUC
w/o. rehearsal	27.9	27.5	24.6	0.8745
only \mathcal{L}_{rec}	29.2	28.6	36.0	0.8802
only \mathcal{L}_{fam}	28.6	28.1	35.4	0.8776
random sampler	28.7	28.3	35.7	0.8813
Full	29.4	28.7	36.3	0.8817

GRU4REC (Hidasi et al., 2015), **Caser** (Tang & Wang, 2018), **DIEN** (Zhou et al., 2019), **RUM** (Chen et al., 2018) and two memory-based methods **HPMN** (Ren et al., 2019) and **MIMN** (Pi et al., 2019), where the HPMN method builds the memory by hierarchical RNNs and the MIMN method introduces a write-read memory as in (Graves et al., 2014). For our rehearsal memory, we set the d_x and d_{model} to 64. The number K of memory slots and length N of segments are both set to 20. And in \mathcal{L}_{rec} , we select 200 items from the large item set as the sampled items.

The results are shown in Table 4. our RM method not only outperforms other memory-based approaches, but also achieves better performance than directly reasoning baselines. This is because our rehearsal memory can aggregate and organize the long-term interests from user behavior sequences and these interests can be activated during next-item prediction. But the directly reasoning approaches may fail to learn such informative interest representations.

4.4. Ablation Study for Self-Supervised Rehearsal

We next perform ablation studies on the self-supervised rehearsal losses and history sampler. Concretely, we first completely discard the self-supervised rehearsal training to produce the ablation model RM (w/o. rehearsal). We then remove the recollection or familiarity loss to produce two ablation models RM (only \mathcal{L}_{fam}) and RM (only \mathcal{L}_{rec}),

Table 6. Ablation Results about the Transformer Encoder and Mask Ratio.

Method	NarrativeQA		ActNet-QA	XLong
	Val	Test	Acc.	AUC
GRU Encoder	29.1	28.3	35.9	0.8789
15% Mask Ratio	29.0	28.5	36.0	0.8809
Full	29.4	28.7	36.3	0.8817

where the history sampler is still retained. Next, we replace the independent history sampler with a random sampler to generate the ablation model RM (random sampler), which randomly selects B history fragments from the input stream for rehearsal training.

We conduct the ablation experiments on NarrativeQA, ActivityNet-QA and XLong datasets. The results are reported in Table 5. We can find the full model outperforms the model RM (w/o. rehearsal), demonstrating the self-supervised rehearsal training with the history sampler can further boost the long-sequence memorization and reasoning ability of rehearsal memory. Further, the full model has better performance than RM (only \mathcal{L}_{fam}) and RM (only \mathcal{L}_{rec}) on all metrics, which illustrates two rehearsal tasks are both helpful for alleviating the issue of gradual forgetting. And RM (only \mathcal{L}_{rec}) achieves better results than RM (only \mathcal{L}_{fam}), showing the recollection task that enables the memory to relive past episode is more important for rehearsal training. Moreover, the ablation model RM (random sampler) has the performance degradation than the full model. This fact indicates it is critical to select informative history fragments for rehearsal training. By the guidance of the history sampler, the rehearsal memory can remember task-relevant clues for subsequent reasoning.

4.5. Ablation Study for Model Setting

In this section, we conduct ablation study about the model settings. Existing MANNs often use a RNN as their controller, but we apply a Transformer encoder to improve the sequential modeling ability of RM. Thus, we replace the Transformer encoder in the rehearsal memory machine with a bi-directional GRU encoder. As shown in Table 6, the full model achieves better performance than the model with the GRU Encoder, verifying the effectiveness of the Transformer encoder.

For the masked history fragments, we set the mask ratio to 50% instead of 15% in BERT. We compare the results of two mask ratios in Table 6. We can find the full model with the 50% mask ratio outperforms the one with the 15% mask ratio. This fact suggests that the large mask ratio makes the rehearsal model $\mathcal{H}_\xi(\mathbf{M}, \mathbf{H})$ utilize the maintained memory




Figure 3. Effect of the Memory Slot Number K .




Figure 4. Effect of the Segment Length N on the ActivityNet-QA dataset.

M than only relying on fragment context, and is beneficial for the self-supervised rehearsal training of RM.

4.6. Hyper-Parameters Analysis

We then explore the effect of two crucial hyper-parameters: the memory slot number K and the segment length N . We first set the slot number K to [10, 15, 20, 25] and compare our RM method with two baselines STM and CT on the NarrativeQA and ActivityNet-QA datasets. We display the results in Figure 3. We note that the performance of all three methods gradually improves with the increase of slot number and slowly reaches the bottleneck. When the number of memory slots exceeds 20, more slots can not bring much improvement. By comparison, we can find our RM method achieves the best performance on different slot numbers, verifying the effectiveness and stability of our rehearsal memory. Moreover, the performance of CT is terrible when the slot number is too few. This is because the CT method implements the compressed memory by a FIFO queue and completely forgets the contents beyond the queue, i.e., its memorization range severely depends on the slot number.

We then set the segment length N to [10, 15, 20, 25, 30] and report the results on the ActivityNet-QA dataset in Figure 4. We can find that when the segment length is set to 10, the RM method achieves poor results and the performance is relatively stable when the segment length changes between 20 and 30. This is because when the segment is too short, important evidence may be scattered in different segments,

and the model cannot effectively capture the evidence and infer the answer.

5. Conclusions

In this paper, we propose the self-supervised rehearsal to enhance long-sequence memorization. We design the self-supervised recollection and familiarity tasks to alleviate the gradual forgetting of early information. Further, we adopt a history sampler to guide the memory to remember critical information. Extensive experiments on a series of downstream tasks verify the performance of our method. For future work, we will further explore the property of rehearsal memory.

Acknowledgments

This work is supported by the National Key R&D Program of China under Grant No. 2018AAA0100603. This research is supported by the National Natural Science Foundation of China under Grant No.61836002 and No.62072397, and the Zhejiang Natural Science Foundation LR19F020006.

References

- Baddeley, A. Working memory. *Science*, 255(5044):556–559, 1992.
- Caba Heilbron, F., Escorcia, V., Ghanem, B., and Carlos Niebles, J. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 961–970, 2015.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- Chen, X., Xu, H., Zhang, Y., Tang, J., Cao, Y., Qin, Z., and Zha, H. Sequential recommendation with user memory networks. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pp. 108–116, 2018.
- Csordás, R. and Schmidhuber, J. Improving differentiable neural computers through memory masking, deallocation, and link distribution sharpness control. *arXiv preprint arXiv:1904.10278*, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference on The North American Chapter of the Association for Computational Linguistics*, 2019.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization.

- Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- Graves, A., Wayne, G., and Danihelka, I. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., Colmenarejo, S. G., Grefenstette, E., Ramalho, T., Agapiou, J., et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471–476, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- Hidasi, B., Karatzoglou, A., Baltrunas, L., and Tikk, D. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*, 2015.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Jin, W., Zhao, Z., Gu, M., Yu, J., Xiao, J., and Zhuang, Y. Multi-interaction network with object relation for video question answering. In *Proceedings of the ACM International Conference on Multimedia*, pp. 1193–1201, 2019a.
- Jin, W., Zhao, Z., Gu, M., Yu, J., Xiao, J., and Zhuang, Y. Video dialog via multi-grained convolutional self-attention context networks. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 465–474, 2019b.
- Kadlec, R., Schmid, M., Bajgar, O., and Kleindienst, J. Text understanding with the attention sum reader network. *arXiv preprint arXiv:1603.01547*, 2016.
- Kočíšký, T., Schwarz, J., Blunsom, P., Dyer, C., Hermann, K. M., Melis, G., and Grefenstette, E. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018.
- Le, H., Tran, T., and Venkatesh, S. Learning to remember more with less memorization. *arXiv preprint arXiv:1901.01347*, 2019a.
- Le, H., Tran, T., and Venkatesh, S. Neural stored-program memory. *arXiv preprint arXiv:1906.08862*, 2019b.
- Le, H., Tran, T., and Venkatesh, S. Self-attentive associative memory. *arXiv preprint arXiv:2002.03519*, 2020a.
- Le, T. M., Le, V., Venkatesh, S., and Tran, T. Hierarchical conditional relation networks for video question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9972–9981, 2020b.
- Moscovitch, M., Cabeza, R., Winocur, G., and Nadel, L. Episodic memory and beyond: the hippocampus and neocortex in transformation. *Annual review of psychology*, 67:105–134, 2016.
- Munkhdalai, T., Sordoni, A., Wang, T., and Trischler, A. Metalearned neural memory. In *Advances in Neural Information Processing Systems*, pp. 13331–13342, 2019.
- Park, T., Choi, I., and Lee, M. Distributed memory based self-supervised differentiable neural computer. *arXiv preprint arXiv:2007.10637*, 2020.
- Pi, Q., Bian, W., Zhou, G., Zhu, X., and Gai, K. Practice on long sequential user behavior modeling for click-through rate prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2671–2679, 2019.
- Rae, J., Hunt, J. J., Danihelka, I., Harley, T., Senior, A. W., Wayne, G., Graves, A., and Lillicrap, T. Scaling memory-augmented neural networks with sparse reads and writes. In *Advances in Neural Information Processing Systems*, pp. 3621–3629, 2016.
- Rae, J. W., Potapenko, A., Jayakumar, S. M., and Lillicrap, T. P. Compressive transformers for long-range sequence modelling. *arXiv preprint arXiv:1911.05507*, 2019.
- Ren, K., Qin, J., Fang, Y., Zhang, W., Zheng, L., Bian, W., Zhou, G., Xu, J., Yu, Y., Zhu, X., et al. Lifelong sequential modeling with personalized memorization for user response prediction. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 565–574, 2019.
- Santoro, A., Faulkner, R., Raposo, D., Rae, J., Chrzanowski, M., Weber, T., Wierstra, D., Vinyals, O., Pascanu, R., and Lillicrap, T. Relational recurrent neural networks. In *Advances in neural information processing systems*, pp. 7299–7310, 2018.
- Seo, M., Kembhavi, A., Farhadi, A., and Hajishirzi, H. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*, 2016.
- Sukhbaatar, S., Weston, J., Fergus, R., et al. End-to-end memory networks. In *Advances in Neural Information Processing Systems*, pp. 2440–2448, 2015.
- Tang, J. and Wang, K. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pp. 565–573, 2018.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- Weston, J., Bordes, A., Chopra, S., Rush, A. M., van Merriënboer, B., Joulin, A., and Mikolov, T. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*, 2015.
- Xiong, C., Zhong, V., and Socher, R. Dynamic coattention networks for question answering. *arXiv preprint arXiv:1611.01604*, 2016.
- Yonelinas, A. P. The nature of recollection and familiarity: A review of 30 years of research. *Journal of memory and language*, 46(3):441–517, 2002.
- Yu, Z., Xu, D., Yu, J., Yu, T., Zhao, Z., Zhuang, Y., and Tao, D. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *Proceedings of the American Association for Artificial Intelligence*, volume 33, pp. 9127–9134, 2019.
- Zhang, S., Yao, D., Zhao, Z., Chua, T.-S., and Wu, F. Causerec: Counterfactual user sequence synthesis for sequential recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021.
- Zhang, Z., Zhao, Z., Lin, Z., He, X., et al. Counterfactual contrastive learning for weakly-supervised vision-language grounding. *Advances in Neural Information Processing Systems*, 33:18123–18134, 2020.
- Zhou, G., Mou, N., Fan, Y., Pi, Q., Bian, W., Zhou, C., Zhu, X., and Gai, K. Deep interest evolution network for click-through rate prediction. In *Proceedings of the American Association for Artificial Intelligence*, volume 33, pp. 5941–5948, 2019.