

Non-local Recurrent Neural Memory for Supervised Sequence Modeling

Canmiao Fu^{1,2,*}, Wenjie Pei^{2,*}, Qiong Cao², Chaopeng Zhang¹,
Yong Zhao^{1,†}, Xiaoyong Shen² and Yu-Wing Tai^{2,†}

¹School of ECE, Peking University ²Tencent

fcm@pku.edu.cn, wenjiecoder@outlook.com, freyaqcao@tencent.com, cpz@pku.edu.cn,
yongzhao@pku.edu.cn, goodshenxy@gmail.com, yuwingtai@tencent.com

Abstract

Typical methods for supervised sequence modeling are built upon the recurrent neural networks to capture temporal dependencies. One potential limitation of these methods is that they only model explicitly information interactions between adjacent time steps in a sequence, hence the high-order interactions between nonadjacent time steps are not fully exploited. It greatly limits the capability of modeling the long-range temporal dependencies since one-order interactions cannot be maintained for a long term due to information dilution and gradient vanishing. To tackle this limitation, we propose the Non-local Recurrent Neural Memory (NRNM) for supervised sequence modeling, which performs non-local operations to learn full-order interactions within a sliding temporal block and models global interactions between blocks in a gated recurrent manner. Consequently, our model is able to capture the long-range dependencies. Besides, the latent high-level features contained in high-order interactions can be distilled by our model. We demonstrate the merits of our NRNM on two different tasks: action recognition and sentiment analysis.

1. Introduction

Supervised sequence modeling aims to build models to extract effective features from variety of sequence data such as video data or text data via supervised learning. It has extensive applications ranging from computer vision [26, 35] to natural language processing [8, 42]. The key challenge in supervised sequence modeling is to capture the long-range temporal dependencies, which are used to further learn the high-level feature for the whole sequence.

Most state-of-the-art methods for supervised sequence modeling are built upon the recurrent neural networks (RNN) [32], which has been validated its effectiveness [33, 52]. One crucial limitation of the vanilla-RNNs is the gradient-vanishing problem along the temporal domain, which results in the inability to model long-term dependen-

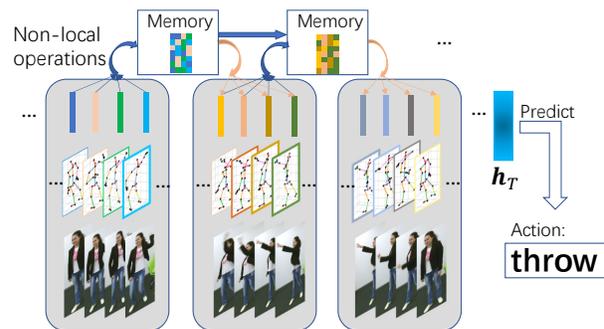


Figure 1. Given a video sample for action recognition, our proposed model (NRNM) performs non-local operations within each memory block to learn high-order interactions between hidden states of different time steps. Meanwhile, the global interactions between memory blocks are modeled in a gated recurrent manner. The learned memory states are in turn leveraged to refine the hidden states in future time steps. Thus, the long-range dependencies can be captured. The model predicts the action based on the hidden state (h_T) in the last time step.

cies. This limitation is then substantially mitigated by gated recurrent networks such as GRU [4] and LSTM [11], which employ learnable gates to selectively retain information in the memory or hidden states. The memory-based methods for sequence modeling [34, 39, 45] are further proposed to address the issue of limited memory of recurrent networks. However, a potential drawback of these methods is that they only model explicitly the information interactions between adjacent time steps in the sequence, hence the high-order interactions between nonadjacent time steps are not fully exploited. This drawback gives rise to two negative consequences: 1) the high-level features contained in the interactions between nonadjacent time steps cannot be distilled; 2) it greatly limits the modeling of long-range temporal dependencies since one-order interaction information cannot be maintained in a long term due to the information dilution and gradient vanishing along with recurrent operations.

Inspired by non-local methods [3, 44] which aim to explore potential interactions between all pairs of feature portions, we propose to perform non-local operations to model

*Both authors contributed equally.

†Corresponding authors.

the high-order interactions between non-adjacent time steps in a sequence. The captured high-order interactions are able to not only help to distill latent high-level features which are hardly learned by typical sequence modeling methods focusing on one-order interactions, but also contribute to modeling long-range temporal dependencies since the non-local operations strength the latent feature propagation and thus substantially alleviate the vanishing-gradient problem. Since exploring full-order interactions between all time steps for a long sequence is computationally expensive and also not necessary due to information redundancy, we model full-order interactions by non-local operations within a temporal block (a segment of sequence) and slide the block to recurrently update the extracted information. More specifically, we propose the Non-local Recurrent Neural Memory (*NRNM*) to perform the blockwise non-local operations to learn full-order interactions within each memory block and capture the local but high-resolution temporal dependencies. Meanwhile, the global interactions between adjacent blocks are captured by updating the memory states in a gated recurrent manner when sliding the memory cell. Consequently, the long-range dependencies can be maintained. Figure 1 illustrates our method by an example of action recognition.

Compared to typical supervised sequence modeling methods, especially recurrent networks with memory mechanism, our *NRNM* benefits from following advantages:

- It is able to model 1) the local full-order interactions between all time steps within a segment of sequence (memory block) and 2) the global interactions between memory blocks. Thus, It can capture much longer temporal dependencies.
- The proposed *NRNM* is able to learn latent high-level features contained in high-order interactions between non-adjacent time steps, which may be missed by conventional methods.
- The *NRNM* cell can be seamlessly integrated into any existing sequence models with recurrent structure to enhance the power of sequence modeling. The integrated model can be trained in an end-to-end manner.

2. Related work

Graphical sequence models. The conventional graphical models for sequence modeling can be roughly divided into two categories: generative and discriminative models. A well-known example of generative model is Hidden Markov Model (HMM) [30], which models sequence data in a chain of latent k -nomial features. Discriminative graphical models model the distribution over all class labels conditioned on the input data. Conditional Random Fields (CRF) [19] is a discriminative model for sequential predictions by modeling the linear mapping between observations and labels. To tackle its limitation of linear mapping, many nonlinear

CRF-variants are proposed [25, 27, 28, 41]. The disadvantages of graphical model compared to recurrent networks lie in the hard optimization and limited capability of temporal modeling. Our model is designed based on recurrent networks.

Recurrent Networks. Recurrent Neural Network [32] learns a hidden representation for each time step by taking into account both current and previous information. Benefited from its advantages such as easy training and temporal modeling, it has been successfully applied to, amongst others, handwriting recognition [2] and speech recognition [33]. However, the key limitation of vanilla-RNN is the gradient vanishing problem during training [10] and thus cannot model long-range temporal dependencies. This limitation is alleviated by gated recurrent networks such as Long Shot-Term Memory (LSTM) [11] and Gate Recurrent Unit (GRU) [4], which selectively retain information by learnable gates. Nevertheless, a potential limitation of these models is that they only model explicitly one-order interactions between adjacent time steps, hence the high-order interactions between nonadjacent time steps are not fully captured. Our model is proposed to circumvent this drawback by employing non-local operations to model full-order interactions in a block-wise manner. Meanwhile, the global interactions between blocks are modeled by a gated recurrent mechanism. Thus, our model is able to model long-range temporal dependencies and distill high-level features that are contained in high-order interactions.

Memory-based recurrent networks. Memory networks are first proposed to rectify the drawback of limited memory of recurrent networks [39, 45], which are then extended for various tasks, especially in natural language processing. Most of these models build external memory units upon a basis model to augment its memory [9, 34, 39, 45]. In particular, attention mechanism [1] is employed to filter the information flow from memory [8, 18, 39, 47]. The primary difference between these memory-based recurrent networks and our model is that these models focus on augmenting the memory size to memorize more information for reference while our model aims to model high-order interactions between different time steps in a sequence, which is not concerned by existing memory-based networks.

3. Non-local Recurrent Neural Memory

Given as input a sequence, our Non-local Recurrent Neural Memory (*NRNM*) is designed as a memory module to capture the long-range temporal dependencies in a non-local manner. It can be seamlessly integrated into any existing sequential models with recurrent structure to enhance sequence modeling. As illustrated in Figure 2, we build our *NRNM* upon a LSTM backbone as an instantiation. We will first elaborate on the cell structure of our *NRNM* and then describe how *NRNM* and the LSTM backbone perform sequence modeling collaboratively.

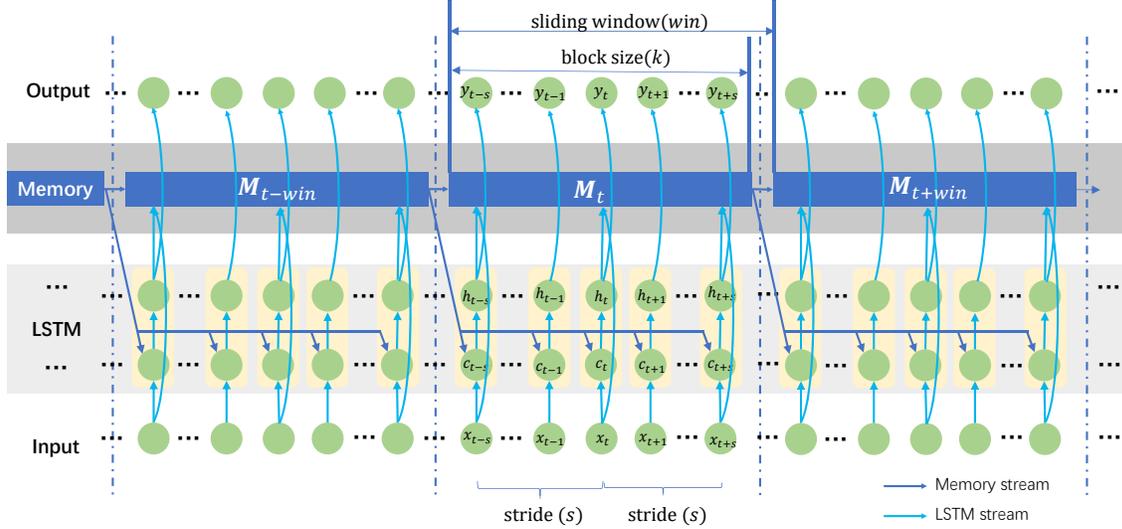


Figure 2. The architecture of our method. Our proposed *NRNM* is built upon the LSTM backbone to learn high-order interactions between LSTM hidden states of different time steps within each memory block. Meanwhile, the global interactions between memory blocks are modeled in a gated recurrent manner. The learned memory states are in turn used to refine the LSTM hidden states in future time steps.

3.1. NRNM Cell

Residing on a basis sequence model (a standard LSTM backbone), the proposed Non-local Recurrent Neural Memory (*NRNM*) aims to maintain a memory cell along the temporal dimension which can not only distill the underlying information contained in past time steps of the input sequence but also capture the temporal dependencies, both operating in the long temporal range. To this end, the *NRNM* cell performs non-local operations on a segment of the sequence (termed as non-local memory block) sliding along the temporal domain, as shown in Figure 2. This blocking design is analogous to DenseNet [12] which performs dense connection (a form of non-local operation) in blocks. The obtained memory embeddings are further leveraged to refine the hidden embeddings which are used for the final prediction. The memory state is updated recurrently when sliding the memory block, which is consistent with update of the basis LSTM backbone.

Consider an input sequence $\mathbf{x}_{1,\dots,T} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ of length T in which $\mathbf{x}_t \in \mathbb{R}^D$ denotes the observation at the t -th time step. The hidden representation of the input sequence at time step t learned by the LSTM backbone is denoted as \mathbf{h}_t . *NRNM* learns the memory embedding $\widetilde{\mathbf{M}}_t$ for a block (a segment of time steps) with blocking size k covering the temporal interval $[t - k + 1, t]$ by refining the underlying information contained in this time interval. Specifically, we consider two types of source information for *NRNM* cell: 1) the learned hidden representations in this time interval $[\mathbf{h}_{t-k+1}, \dots, \mathbf{h}_t]$ by the LSTM backbone; 2) the original input features $[\mathbf{x}_{t-k+1}, \dots, \mathbf{x}_t]$. Hence the memory embedding $\widetilde{\mathbf{M}}_t$ at time step t is formulated as:

$$\widetilde{\mathbf{M}}_t = f([\mathbf{h}_{t-k+1}, \dots, \mathbf{h}_t], [\mathbf{x}_{t-k+1}, \dots, \mathbf{x}_t]), \quad (1)$$

where f is the nonlinear transformation function performed by *NRNM* cell. Here we incorporate the input feature \mathbf{x} which is already assimilated in the hidden representation \mathbf{h} of the basis LSTM backbone since we aim to explore the latent interactions between hidden representations and input features in the current block (i.e., the interval $[t - k + 1, t]$).

Next we elaborate on the transformation function f of *NRNM* cell presented in Figure 3. To distill information in current block that is worth to retain in memory, we apply Self-Attention mechanism implemented with multi-head attention [42] to model latent full-order interactions among source informations: original input features and the hidden representations by LSTM in the current block:

$$\begin{aligned} \mathbf{C} &= \text{Concat}([\mathbf{h}_{t-k+1}, \dots, \mathbf{h}_t], [\mathbf{x}_{t-k+1}, \dots, \mathbf{x}_t]), \\ \mathbf{Q}, \mathbf{K}, \mathbf{V} &= (\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v)\mathbf{C}, \\ \mathbf{W}_{att} &= \text{softmax}(\mathbf{Q}\mathbf{K}^\top / \sqrt{m}), \\ \mathbf{M}_{att} &= \mathbf{W}_{att}\mathbf{V}. \end{aligned} \quad (2)$$

Herein, $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ are queries, keys and values of Self-Attention transformed by parameters $\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v$ from the source information \mathbf{C} respectively. \mathbf{W}_{att} is the derived attention weights calculated by dot-product attention scheme scaled by the memory hidden size m . The obtained attention embeddings \mathbf{M}_{att} is then fed into two skip-connection layers and one fully-connected layer to achieve the memory embedding $\widetilde{\mathbf{M}}_t$.

The physical interpretation of this design is that the source information is composed of $2k$ information units: k LSTM hidden states and k input features. Each information unit of the obtained memory embedding $\widetilde{\mathbf{M}}_t$ is constructed by attending into each of these $2k$ source information units while the size of memory embedding $\widetilde{\mathbf{M}}_t$ can be

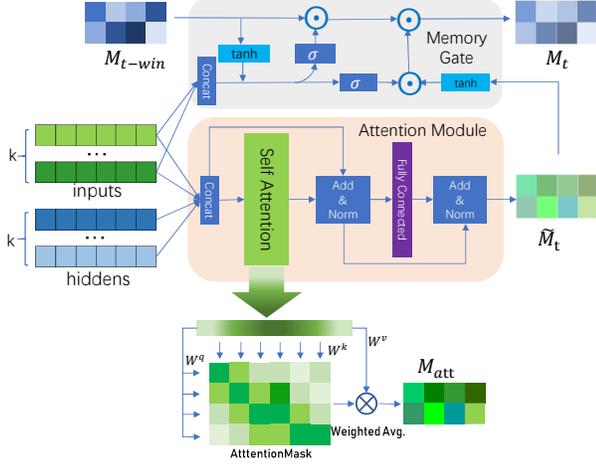


Figure 3. The structure of *NRNM* cell.

customized via the parametric transformation. As such, the full-order latent interactions between the source information units are explored in a non-local way. Another benefit of such non-local operations is that it strengthens latent feature propagation and thus alleviates the vanishing-gradient problem, which is always suffered by recurrent networks.

Since LSTM hidden states already contain history information by recurrent structure, in practice we use a striding scheme to select hidden states as the source information for *NRNM* cell to avoid potential information redundancy and improve the modeling efficiency. For instance, we pick hidden states every s time steps in the temporal interval $[t - k + 1, t]$ for the source information, given $\text{stride} = s$.

Gated recurrent update of memory state. \tilde{M}_t only contains information within current temporal block ($[t - k + 1, t]$). To model the temporal consistency between adjacent memory blocks, we also update the memory state in a gated recurrent manner, which is similar to the recurrent scheme of LSTM. Specifically, the final memory state M_t for current memory block is obtained by:

$$M_t = G_i \odot \tanh(\tilde{M}_t) + G_f \odot M_{t-win}, \quad (3)$$

where win is the sliding window size of *NRNM* cell which controls the updating frequency of memory state. G_i and G_f are input gate and forget gate respectively to balance the memory information flow from current time step \tilde{M}_t and previous memory state M_{t-win} . They are modeled by measuring the compatibility between current input feature and previous memory state:

$$\begin{aligned} G_i &= \text{sigmoid}(\mathbf{W}_{im} \cdot [\mathbf{x}_{t-k+1}, \dots, \mathbf{x}_t, M_{t-win}] + \mathbf{B}_{im}), \\ G_f &= \text{sigmoid}(\mathbf{W}_{fm} \cdot [\mathbf{x}_{t-k+1}, \dots, \mathbf{x}_t, M_{t-win}] + \mathbf{B}_{fm}), \end{aligned} \quad (4)$$

wherein, \mathbf{W}_{im} and \mathbf{W}_{fm} are transformation matrices while \mathbf{B}_{im} and \mathbf{B}_{fm} are bias terms.

Modeling long-range dependencies. We aim to capture underlying long-range dependencies in a sequence by a two-pronged strategy:

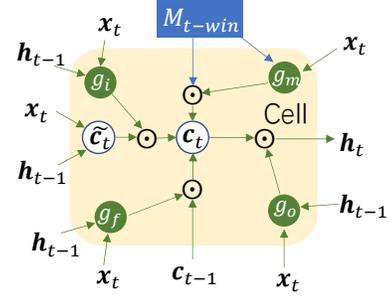


Figure 4. The LSTM cell is updated by incorporating the memory state.

- We perform non-local operations within each temporal block by *NRNM* cell to capture the full-order interactions locally between different time steps and distill the high-quality memory state. Hence, the local but high-resolution temporal information can be captured.
- We update the memory state in a gated recurrent manner smoothly when sliding the window of memory block along the temporal domain. It is designed to capture the global temporal dependencies between memory blocks in low resolution considering the potential information redundancy and computational efficiency.

3.2. Sequence Modeling

Our *NRNM* can be seamlessly integrated into the LSTM backbone to enhance the power of sequence modeling. Specifically, we incorporate the obtained memory state into the recurrent update of LSTM cell states to help refine its quality as shown in Figure 4:

$$\mathbf{v}_m = \text{flatten}(M_{t-win}) \quad (5)$$

$$\mathbf{C}_t = \mathbf{g}_f \odot \mathbf{C}_{t-1} + \mathbf{g}_i \odot \tilde{\mathbf{C}}_t + \mathbf{g}_m \odot \mathbf{v}_m,$$

where \mathbf{C}_{t-1} , \mathbf{C}_t and $\tilde{\mathbf{C}}_t$ are previous LSTM cell state, current cell state and candidate cell state respectively. \mathbf{v}_m is the vector flattened from the memory state M_{t-win} . \mathbf{g}_f and \mathbf{g}_i are the routine forget gate and input gate of LSTM cell to balance the information flow between the current time step and previous step. All $\tilde{\mathbf{C}}_t$, \mathbf{g}_f and \mathbf{g}_i are modeled in a similar nonlinear way as a function of input feature \mathbf{x}_t and previous hidden state \mathbf{h}_{t-1} . For instance, the input gate \mathbf{g}_i is modeled as:

$$\mathbf{g}_i = \text{sigmoid}(\mathbf{W}_i \cdot \mathbf{x}_t + \mathbf{U}_i \cdot \mathbf{h}_{t-1} + \mathbf{b}_i). \quad (6)$$

In Equation 5, we enrich the modeling of the LSTM cell state \mathbf{C}_t by incorporating our *NRNM* cell state M_t via a memory gate \mathbf{g}_m . The memory gate is constructed as a matrix to control the information flow from the memory state M_t , which is derived by measuring the relevance (compatibility) between current input feature and the memory state:

$$\mathbf{g}_m = \text{sigmoid}(\mathbf{W}_m \cdot \mathbf{x}_t + \mathbf{U}_m \cdot \text{flatten}(M_{t-win}) + \mathbf{b}_m), \quad (7)$$

where \mathbf{W}_m and \mathbf{U}_m are transformation matrices and \mathbf{b}_m is the bias term.

The newly constructed cell state \mathbf{C}_t are further used to derive the hidden state \mathbf{h}_t of the whole sequence model prepared for the final prediction:

$$\mathbf{h}_t = \mathbf{g}_o \odot \tanh(\mathbf{C}_t), \quad (8)$$

where \mathbf{g}_o is the output gate which is modeled in a similar way to the input gate in Equation 6.

3.3. End-to-end Parameter Learning

The memory state of our *NRNM* for current block is learned based on the hidden states within this block of the LSTM backbone while the obtained memory state is in turn leveraged to refine the hidden states in future time steps. Hence, our *NRNM* and the LSTM backbone are integrated seamlessly and refine each other alternately.

The learned hidden representations $\{\mathbf{h}_t\}_{t=1,\dots,T}$ in Equation 8 for a sequence with length T can be used for any sequence prediction task such as step-wise prediction (like language modeling) or sequence classification (like action classification). In subsequent experiments, we validate our model in two tasks of sequence classification with different modalities: action recognition and sentiment analysis. Below we present the loss function for training our model for sequence classification, but it is straightforward to substitute the loss function to adapt our model to tasks of step-wise prediction.

Given a training set $\mathcal{D} = \{\mathbf{x}_{1,\dots,T^n}^n, y^n\}_{n=1,\dots,N}$ containing N sequences of length T^n and their associated labels y^n . We learn our *NRNM* and the LSTM backbone jointly in an end-to-end manner by minimizing the conditional negative log-likelihood of the training data with respect to the parameters:

$$\mathcal{L} = - \sum_{n=1}^N \log P(y^n | \mathbf{x}_{1,\dots,T^n}^n), \quad (9)$$

where the probability of the predicted label y^n among K classes is calculated by the hidden state in the last time step:

$$P(y^n | \mathbf{x}_{1,\dots,T^n}^n) = \frac{\exp(\mathbf{W}_{y^n}^\top \mathbf{h}_{T^n} + \mathbf{b})}{\sum_{i=1}^K \exp(\mathbf{W}_i^\top \mathbf{h}_{T^n} + \mathbf{b}_i)}. \quad (10)$$

Herein, \mathbf{W}^\top and \mathbf{b} is the parameters for linear transformation and bias term.

4. Experiments on Action Recognition

To evaluate the performance of our proposed *NRNM* model, we first consider the task of action recognition in which the temporal dependencies between frames in a video are the most discriminative cues.

4.1. Dataset and Evaluation Protocol

We evaluate our method on the NTU dataset[35] which is currently the largest action recognition dataset. It is a RGB+D-based dataset containing 56,880 video sequences and 4 million frames collected from 40 distinct subjects.

The dataset includes 60 action categories. 3D skeleton data (i.e. 3D coordinates of 25 body joints) is provided using Microsoft Kinect.

In our experiments, we opt for NTU dataset using only 3D skeleton joint information rather than Kinetics [15] based on RGB information for action recognition since single-frame RGB information already provides much implication for action recognition and weakens the importance of temporal dependencies [29]. Dropping RGB-D information enforces our model to recognize actions relying on temporal information of joints.

Two standard evaluation metrics are provided in [35]: Cross-Subject (CS) and Cross-View (CV). CS evaluation splits 40 subjects equally into training and test sets consisting of 40,320 and 16,560 samples respectively. In CV evaluation, samples of camera 1 are used for testing and samples from cameras 2 and 3 for training. We report both metrics for performance evaluation.

4.2. Implementation

Our *NRNM* is built on a 3-layer LSTM backbone. The number of hidden units of all recurrent networks mentioned in this work (vanila-RNN, GRU, LSTM) is tuned on a validation set by selecting the best configuration from the option set $\{128, 256, 512\}$. We employ 4-head attention scheme in practice. The size of memory state is set to be same as the combined size of input hidden states, i.e., the dimensions are $[\text{block size } (k) / \text{stride } (s), \dim(\mathbf{h}_t)]$. Following Tu et al. [40], Zoneout [17] is employed for network regularization. The dropout value is set to 0.5 to prevent potential overfitting. Adam [16] is used with the initial learning rate of 0.001 for gradient descent optimization.

4.3. Investigation on NRNM

We first perform experiments to investigate our proposed *NRNM* systematically.

Effect of the block size k . We first conduct experiments on NTU dataset to investigate the performance of *NRNM* as a function of the block size. Concretely, we evaluate our method using an increasing number of block sizes: 4, 6, 8, 10, and 12 while fixing the other hyper-parameters.

Figure 5(a) shows that the accuracy initially increases as the increase of the block size, which is reasonable since larger block size allows *NRNM* to incorporate information of more time steps in memory and thus enables *NRNM* to capture longer temporal dependencies. As the block size increases further after the saturated state at the block size of 8, the performance starts to decrease. We surmise that the non-local operations on a long block of sequence result in overfitting on the training data and information redundancy.

Effect of the integrated location of NRNM on the LSTM backbone. We next study the effect of integrating the *NRNM* into different layers of the 3-layer LSTM backbone.

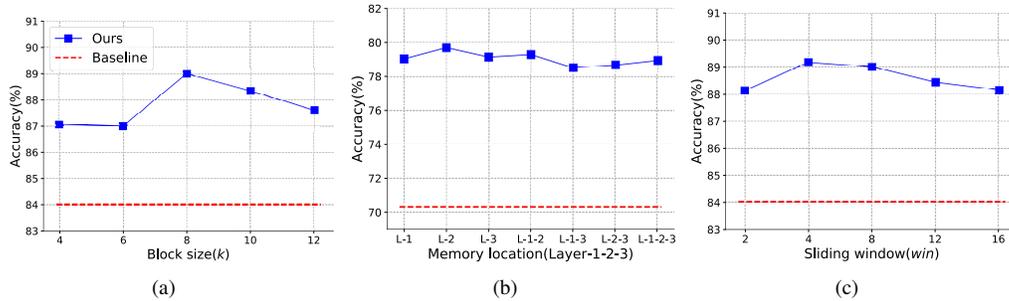


Figure 5. Ablation study of *NRNM* on NTU dataset by exploring the effect of (a) the block size k , (b) the integrated location of *NRNM* on the LSTM backbone and (c) the sliding window size win . The performance of the baseline (a standard LSTM) is presented for reference.

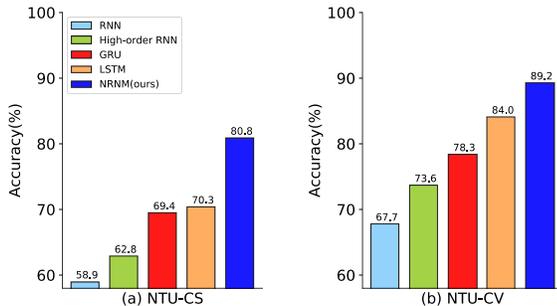


Figure 6. Comparison of our model with other basic recurrent models in terms of classification accuracy (%) on NTU dataset in both Cross-Subject (CS) and Cross-View (CV) metrics.

Figure 5(b) presents the results, from which we can conclude: 1) Integrating *NRNM* at any layer of LSTM outperforms the standard LSTM; 2) Only integrating *NRNM* once at one layer performs better than applying *NRNM* at multiple layers which would lead to information redundancy and overfitting; 3) Integrating *NRNM* at the middle layer achieves the best performance, which is probably because the layer-2 hidden states of LSTM are more suitable for *NRNM* to distill information than the low-level and high-level features learned by layer-1 and layer-3 hidden states.

Effect of sliding window size win . Then we investigate the effect of sliding window size, which is used to control the updating frequency of memory state. Theoretically, too small sliding window size implies much overlap between two adjacent memory blocks and thus tends to lead to information redundancy. On the other hand, too large sliding window size results in large non-accessed temporal interval between two adjacent memory blocks and would potentially miss information in the interval.

In this set of experiments, we set block size to 8 time steps, and consider different sliding window size. Figure 5(c) reports that the model performs well when the sliding window is around 4 to 8 while the performance decreases at other values, which validates our analysis.

Comparison with LSTM baselines. To investigate the effectiveness of our *NRNM*, we compare our model to the basic recurrent models including vanilla-RNN, GRU, LSTM

and high-order RNN on NTU dataset in two evaluation metrics: Cross-Subject (CS) and Cross-View (CV). Figure 6 shows 1) all RNNs with memory or gated structure outperforms vanilla-RNN and high-order RNN by a large margin, which indicates the advantages of memory and gated structure for controlling information flow; 2) high-order RNN performs better than vanilla-RNN which implies the necessity of the non-local operations since high-order connections can be considered as a simple non-local operation in a local area. It is also consistent with the existing conclusions [37, 50]; 3) our *NRNM* outperforms LSTM significantly which demonstrates the superiority of our model over standard LSTM.

4.4. Comparison with State-of-the-arts

In this set of experiments, we compare our model with the state-of-the-art methods for action recognition on NTU dataset in both Cross-Subject (CS) and Cross-View (CV) metrics. It should be noted that we do not compare with methods which employ extra information or prior knowledge such as joint connections for each part of body or human body structure modeling [36, 48].

Table 1 reports the experimental results. Our model achieves the best performance in both CS and CV metrics, which demonstrates the superiority of our model over other recurrent networks, especially those with memory or gated structures. While our model outperforms the standard LSTM model substantially, the methods based on LSTM [38, 52] boost the performance over LSTM by introducing extra attention mechanisms.

Analysis on model complexity. To compare the model complexity between our model and other recurrent baselines and investigate whether the performance gain of our model is boosted by the augmented model complexity, we evaluate the performance of the recurrent baselines with different model complexities (configurations) in Table 2. Our model substantially outperforms other baselines under optimized configurations, which demonstrates that the performance superiority of our model is not resulted from the increased capacity by the extra parameters.

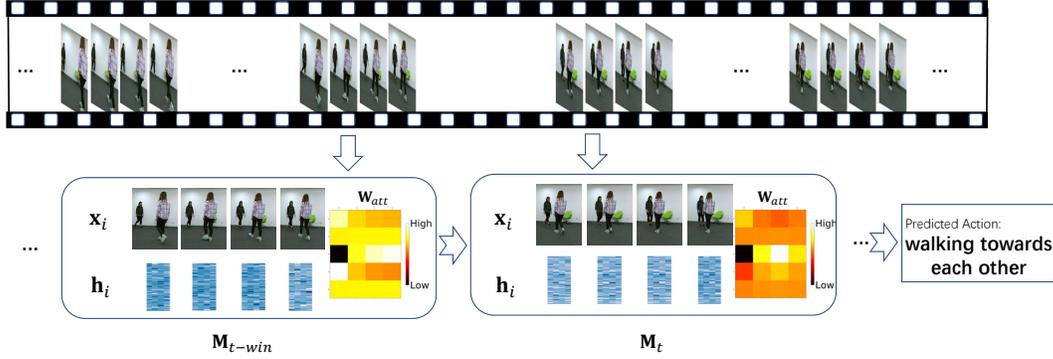


Figure 7. Visualization of an example with labeled action “walking towards each other”. Our model is able to correctly recognize it while LSTM misclassifies it as “punching/slapping other person”. The temporal variations of relative distance between two persons are key to recognize the action. Our model can successfully capture it while LSTM fails. Two blocks of memory states and the attention weights W_{att} in Equation 2 are visualized.

	CS	CV
HBRNN-L[7]	59.1	64.0
Part-aware LSTM[35]	62.9	70.3
Trust Gate ST-LSTM[21]	69.2	77.7
Two-stream RNN[43]	71.3	79.5
Ensemble TS-LSTM[20]	74.6	81.3
VA-LSTM[51]	79.4	87.6
STA-LSTM[38]	73.4	81.2
EleAtt-LSTM[52]	78.4	85.0
EleAtt-GRU[52]	79.8	87.1
LSTM (baseline)	70.3	84.0
<i>NRNM</i> (ours)	80.8	89.2

Table 1. Classification accuracy (%) on NTU by different methods in both Cross-Subject (CS) and Cross-View (CV) metrics.

	CV(%)	#Parameters
3-LSTM (256)	83.9	1.5M
3-LSTM (512)	84.0	5.6M
5-LSTM (512)	83.1	9.8M
3-ElAtt-LSTM (256)	85.5	1.8M
6-ElAtt-LSTM (256)	82.7	3.8M
4-ElAtt-LSTM (512)	83.4	8.9M
3-ElAtt-GRU (100)	87.1	0.3M
3-ElAtt-GRU (256)	85.4	1.4M
5-ElAtt-GRU (256)	85.0	2.5M
<i>NRNM</i> (ours)	89.2	3.6M

Table 2. Classification accuracy (%) on NTU by different methods with different model complexities in Cross-View (CV) metrics. Here 3-LSTM (256) denotes the LSTM equipped with 3 hidden layers comprising 256 hidden units. Note that all results are reported from our implementations.

4.5. Qualitative Analysis

To qualitatively illustrate the advantages of the proposed *NRNM*, figure 7 presents a concrete video example with the action label “walking towards each other” (groundtruth). In this example, it is quite challenging to recognize the action since it can only be inferred by the temporal variations

of the relative distance between two persons in the scene. Hence, capturing the long-range dependencies is crucial to recognize it. The standard LSTM misclassifies it as “punching/slapping other person” while our model is able to correctly classify it due to the capability to model long-range temporal information by our designed *NRNM*.

Figure 7 visualizes two blocks of memory states, each of which is learned by *NRNM* cell via incorporating information of multiple frames including input features x_i and the hidden states h_i of LSTM backbone. To obtain more insights into the non-local operations of *NRNM*, we visualize the attention weights W_{att} in Equation 2 to show that each unit of memory state is calculated by attending to all units of source information (x_i and h_i).

5. Experiments on Sentiment Analysis

Next we perform experiments on task of sentiment analysis to evaluate our model on the text modality. Specifically, we aim to identify online movie reviews as positive or negative, which is a sequence classification problem.

5.1. Dataset and Evaluation Protocol

We use the IMDB Review dataset [22] which is a standard benchmark for sentiment analysis. It contains 50,000 labeled reviews among which 25,000 samples are used for training and the rest for testing. The average length of reviews is 241 words and the maximum length is 2526 words[5]. Note that the IMDB dataset also provides additional 50,000 unlabeled reviews, which are used by several customized semi-supervised learning methods [5, 6, 14, 24, 31]. Since we only use labeled data for supervised training, we compare our methods with those methods based on supervised learning using the same set of training data for a fair comparison.

The torchtext¹ is used for data preprocessing. Following the training strategy in Dai et al. [5], we pretrain a language model for extracting word embeddings.

¹<https://github.com/pytorch/text>

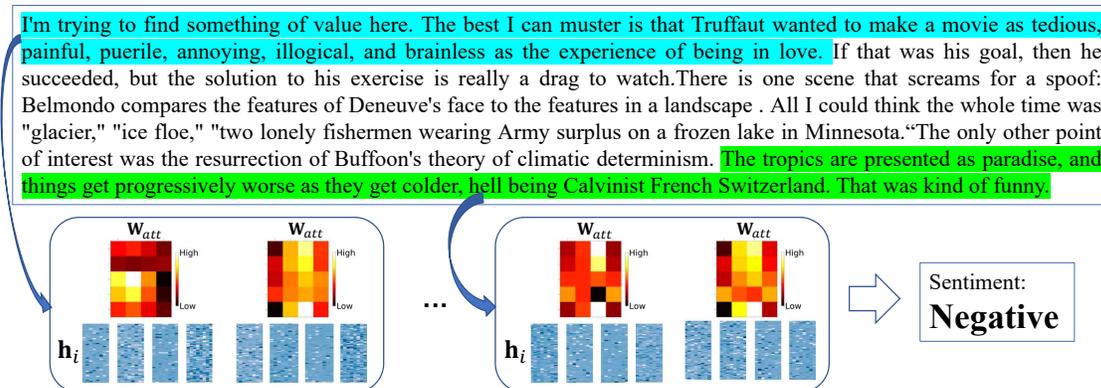


Figure 8. Visualization of an example of movie review with the groundtruth label “negative”. Our model is able to correctly classify it while LSTM fails. The last sentence (in green color) which seems positive tends to misguide models. The first sentence is an important cue for negative sentiment, which is hardly captured by LSTM since it is easily forgotten by the hidden state h_T in the last time step.

5.2. Comparison with LSTM Baselines

We first conduct a set of experiments to compare our model to the basic recurrent networks including vanilla-RNN, GRU, LSTM and high-order RNN. Figure 9 shows that our model outperforms all other baselines significantly which reveals the remarkable advantages of our *NRNM*. Besides, while LSTM and GRU perform much better than vanilla-RNN, high-order RNN also boosts the performance by a large margin compared to vanilla-RNN. It again demonstrates the benefits of high-order connections which are a simple form of non-local operations in local area.

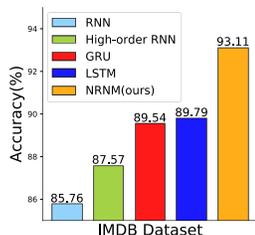


Figure 9. Comparison of our model with other basic recurrent models in terms of classification accuracy (%) on IMDB dataset.

5.3. Comparison with the State-of-the-arts

Next we compare our *NRNM* with the state-of-the-art methods including LSTM[46], oh-CNN [13] and oh-2LSTMp[14] which learn the word embeddings by customized CNN or LSTM instead of using existing pretrained word embedding vocabulary, DSL [46] and MLDL[46] which perform a dual learning between language modeling and sentiment analysis, GLoMo[49] which is a transfer learning framework, and BCN+Char+CoVe[23] which trains a machine translation model to encode the word embeddings to improve the performance of sentiment analysis.

Table 3 shows that our model achieves the best performance among all methods. It is worth mentioning that our model even performs better than GLoMo[49] and BCN+Char+CoVe[23], which employ additional data for either transfer learning or training a individual machine translation model.

Methods	Accuracy
LSTM [46]	89.9
MLDL [46]	92.6
GLoMo [49]	89.2
oh-2LSTMp [14]	91.9
DSL [46]	90.8
oh-CNN [13]	91.6
BCN+Char+CoVe [23]	92.1
LSTM (baseline)	89.8
<i>NRNM</i> (ours)	93.1

Table 3. Classification accuracy (%) on IMDB dataset by different methods.

5.4. Qualitative Analysis

Figure 8 illustrates an example of sentiment analysis from IMDB dataset. This example of movie review is fairly challenging since the last sentence of the review seems to be positive which is prone to misguide models, especially when we use the hidden state of last time step h_T for prediction. Our model is able to correctly classify it as “negative” while LSTM fails. We also visualize the attention weights of non-local operations (W_{att} Equation 2) in two blocks of *NRNM* states to show the attendance of each information units of source information for calculating the *NRNM* states. The first memory block corresponds to the first sentence which is an important cue of negative sentiment while the second memory block corresponds to the last sentence.

6. Conclusion

In this work, we have presented the Non-local Recurrent Neural Memory (*NRNM*) for supervised sequence modeling. We perform non-local operations within each memory block to model full-order interactions between non-adjacent time steps and model the global interactions between memory blocks in a gated recurrent manner. Thus, the long-range temporal dependencies are captured. Our method achieves the state-of-the-art performance for tasks of action recognition and sentiment analysis.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- [2] R. Bertolami, H. Bunke, S. Fernandez, A. Graves, M. Liwicki, and J. Schmidhuber. A novel connectionist system for improved unconstrained handwriting recognition. *IEEE T-PAMI*, 31(5), 2009.
- [3] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, 2005.
- [4] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [5] A. M. Dai and Q. V. Le. Semi-supervised sequence learning. In *NeurIPS*, 2015.
- [6] A. B. Dieng, C. Wang, J. Gao, and J. Paisley. Topicrnn: A recurrent neural network with long-range semantic dependency. In *ICLR*, 2017.
- [7] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *CVPR*, 2015.
- [8] E. Grave, A. Joulin, and N. Usunier. Improving neural language models with a continuous cache. *ICLR*, 2017.
- [9] A. Graves, G. Wayne, and I. Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [10] S. Hochreiter, Y. Bengio, P. Frasconi, J. Schmidhuber, et al. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- [11] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [12] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- [13] R. Johnson and T. Zhang. Effective use of word order for text categorization with convolutional neural networks. *arXiv preprint arXiv:1412.1058*, 2014.
- [14] R. Johnson and T. Zhang. Supervised and semi-supervised text categorization using lstm for region embeddings. In *ICML*, 2016.
- [15] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [17] D. Krueger, T. Maharaj, J. Kramár, M. Pezeshki, N. Ballas, N. R. Ke, A. Goyal, Y. Bengio, A. Courville, and C. Pal. Zoneout: Regularizing rnns by randomly preserving hidden activations. *ICLR*, 2017.
- [18] A. Kumar, O. Irsoy, P. Ondruska, M. Iyyer, J. Bradbury, I. Gulrajani, V. Zhong, R. Paulus, and R. Socher. Ask me anything: Dynamic memory networks for natural language processing. In *ICML*, 2016.
- [19] J. Lafferty, A. McCallum, and F. C. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001.
- [20] I. Lee, D. Kim, S. Kang, and S. Lee. Ensemble deep learning for skeleton-based action recognition using temporal sliding lstm networks. In *ICCV*, 2017.
- [21] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *ECCV*, 2016.
- [22] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*. Association for Computational Linguistics, 2011.
- [23] B. McCann, J. Bradbury, C. Xiong, and R. Socher. Learned in translation: Contextualized word vectors. In *NeurIPS*, 2017.
- [24] T. Miyato, A. M. Dai, and I. Goodfellow. Adversarial training methods for semi-supervised text classification. *ICLR*, 2017.
- [25] L.-P. Morency, A. Quattoni, and T. Darrell. Latent-dynamic discriminative models for continuous gesture recognition. In *CVPR*, 2007.
- [26] W. Pei, T. Baltrusaitis, D. M. Tax, and L.-P. Morency. Temporal attention-gated model for robust sequence classification. In *CVPR*, 2017.
- [27] W. Pei, H. Dibeklioglu, D. M. Tax, and L. van der Maaten. Multivariate time-series classification using the hidden-unit logistic model. *IEEE transactions on neural networks and learning systems*, 29(4):920–931, 2018.
- [28] J. Peng, L. Bo, and J. Xu. Conditional neural fields. In *NIPS*, 2009.
- [29] Z. Qiu and et al. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, 2017.
- [30] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [31] A. Radford, R. Jozefowicz, and I. Sutskever. Learning to generate reviews and discovering sentiment. *arXiv preprint arXiv:1704.01444*, 2017.
- [32] D. E. Rumelhart, G. E. Hinton, R. J. Williams, et al. Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1, 1988.
- [33] H. Sak, A. Senior, and F. Beaufays. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *Fifteenth annual conference of the international speech communication association*, 2014.
- [34] A. Santoro, R. Faulkner, D. Raposo, J. Rae, M. Chrzanowski, T. Weber, D. Wierstra, O. Vinyals, R. Pascanu, and T. Lillicrap. Relational recurrent neural networks. In *NeurIPS*, 2018.
- [35] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016.
- [36] C. Si, Y. Jing, W. Wang, L. Wang, and T. Tan. Skeleton-based action recognition with spatial reasoning and temporal stack learning. In *ECCV*, 2018.
- [37] R. Soltani and H. Jiang. Higher order recurrent neural networks. *arXiv preprint arXiv:1605.00064*, 2016.

- [38] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu. Spatio-temporal attention-based lstm networks for 3d action recognition and detection. *IEEE TIP*, 27(7):3459–3471, 2018.
- [39] S. Sukhbaatar, J. Weston, R. Fergus, et al. End-to-end memory networks. In *NeurIPS*, 2015.
- [40] J. Tu, H. Liu, F. Meng, M. Liu, and R. Ding. Spatial-temporal data augmentation based on lstm autoencoder network for skeleton-based human action recognition. In *ICIP*, 2018.
- [41] L. Van Der Maaten, M. Welling, and L. Saul. Hidden-unit conditional random fields. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011.
- [42] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [43] H. Wang and L. Wang. Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks. In *CVPR*, 2017.
- [44] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *CVPR*, 2018.
- [45] J. Weston, S. Chopra, and A. Bordes. Memory networks. In *ICLR*, 2015.
- [46] Y. Xia, X. Tan, F. Tian, T. Qin, N. Yu, and T.-Y. Liu. Model-level dual learning. In *ICML*, 2018.
- [47] C. Xiong, S. Merity, and R. Socher. Dynamic memory networks for visual and textual question answering. In *ICML*, 2016.
- [48] S. Yan, Y. Xiong, and D. Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018.
- [49] Z. Yang, B. Dhingra, K. He, W. W. Cohen, R. Salakhutdinov, Y. LeCun, et al. Glomo: Unsupervisedly learned relational graphs as transferable representations. In *NeurIPS*, 2018.
- [50] C. Zhang and P. C. Woodland. High order recurrent neural networks for acoustic modelling. In *ICASSP*, 2018.
- [51] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *ICCV*, 2017.
- [52] P. Zhang, J. Xue, C. Lan, W. Zeng, Z. Gao, and N. Zheng. Adding attentiveness to the neurons in recurrent neural networks. In *ECCV*, 2018.