

Integrating Deep Event-Level and Script-Level Information for Script Event Prediction

Long Bai, Saiping Guan, Jiafeng Guo, Zixuan Li, Xiaolong Jin, Xueqi Cheng

CAS Key Lab of Network Data Science and Technology, Institute of Computing Technology, Chinese Academy of Sciences (CAS);

School of Computer Science and Technology, University of Chinese Academy of Sciences

{bailong18b, guansaiping, guojiafeng, lizixuan, jinxiaolong, cxq}@ict.ac.cn

Abstract

Scripts are structured sequences of events together with the participants, which are extracted from the texts. Script event prediction aims to predict the subsequent event given the historical events in the script. Two kinds of information facilitate this task, namely, the event-level information and the script-level information. At the event level, existing studies view an event as a verb with its participants, while neglecting other useful properties, such as the state of the participants. At the script level, most existing studies only consider a single event sequence corresponding to one common protagonist. In this paper, we propose a Transformer-based model, called MCPredictor, which integrates deep event-level and script-level information for script event prediction. At the event level, MCPredictor utilizes the rich information in the text to obtain more comprehensive event semantic representations. At the script-level, it considers multiple event sequences corresponding to different participants of the subsequent event. The experimental results on the widely-used New York Times corpus demonstrate the effectiveness and superiority of the proposed model.

1 Introduction

Scripts, consisting of structured sequences of events, are a kind of knowledge that describes everyday scenarios (Abelson and Schank, 1977). A typical script is the *restaurant script*, which describes the scenario of a person going to a restaurant. In this script, “customer enter restaurant”, “customer order food”, “customer eat food” and “customer pay bill” happen successively. This structured knowledge is helpful to downstream natural language processing tasks, such as anaphora resolution (Bean and Riloff, 2004) and story generation (Chaturvedi et al., 2017).

Script event prediction aims to predict the subsequent event based on the historical events in the

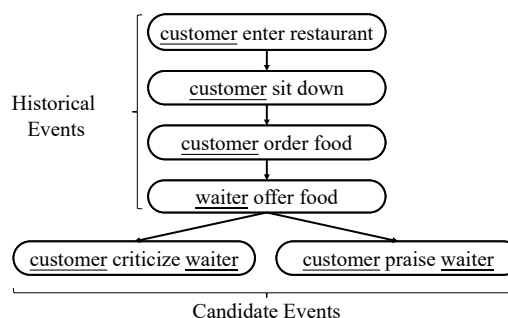


Figure 1: An example of the script event prediction task, where two essential participants are underlined.

script. What is vital to this task is to understand the historical events comprehensively. Therefore, two categories of information are essential, namely, the event-level information and the script-level information. The event-level information contains necessary elements to describe the events, such as the verbs and their participants, while the script-level information describes how the events are connected and structured, such as via the temporal order or a common participant.

The existing studies only consider the verb with its participants (usually the headwords) at the event level. However, this event representation method suffers from a lack of necessary information to derive a more accurate prediction. There exists other important properties of the events in the original texts, such as the intention and the state of the participants. For instance, as shown in Figure 1, if the waiter’s service is friendly, the customer will be more likely to praise the waiter. If the customer is irritable, he/she will be more likely to criticize the waiter. Unfortunately, the current formulation does not consider these features. From the aspect of the script-level information, existing studies only model the events that share a specific protagonist. These events are organized into a sequence by temporal order, which is thus called the narra-

tive event chain. However, an event may contain multiple participants, each of which has its own influence on the occurrence of the event. As shown in Figure 1, the two participants, i.e., the customer and the waiter, take their own actions and jointly determine whether the customer will criticize or praise the waiter.

Motivated by these observations, at the event level, we trace the events back to their original texts and consider all the constituents in the texts that describe the events. Then, we utilize the informative constituents to obtain more comprehensive event semantic representations. With respect to the script level, we view the script as multiple narrative event chains derived from the participants of the subsequent event. Via modeling the narrative event chains of multiple participants, we are able to capture their behavior trends and predict the occurrence of the subsequent event. Thus, by integrating the above deep event-level and script-level information, we propose MCPredictor, a Transformer-based script event prediction model. At the event level, MCPredictor contains an event encoding component and a text encoding component. Via aggregating the output of the two components, it obtains more comprehensive event semantic representations. At the script level, MCPredictor contains a chain modeling component and a scoring component. The former integrates the temporal order information into event representations. The latter aggregates the influence of multiple narrative event chains through an attention mechanism to predict the subsequent event. In general, this paper has the following main contributions:

- We emphasize the importance of both the event-level information and the script-level information for script event prediction, and go deep into it;
- We propose the MCPredictor model to integrate both kinds of information for script event prediction. It introduces rich information from the original texts to enhance the event-level information and learns the script-level information by aggregating the influence of multiple narrative event chains on the subsequent event;
- The proposed model achieves state-of-the-art performance on the widely-used benchmark dataset. The event-level and the script-level

information introduced by us are also testified to be effective.

2 Related Work

Recent studies on script event prediction start from (Chambers and Jurafsky, 2008), which proposes the basic structure of the narrative event chain with a specific participant (called the protagonist). Each event is represented as a tuple of the verb and the dependency relation between the verb and the protagonist, i.e., $\langle verb, dependency \rangle$. Then, it uses Pointwise Mutual Information (PMI) to measure the score of two events to be in the same narrative event chain. Finally, it aggregates the pairwise scores to infer a candidate’s probability of being the subsequent event of the narrative event chain.

Balasubramanian et al. (2013) pointed out that the event representation method mentioned above may lose the co-occurrence information between a subject and its object. Therefore, they represented an event as a $\langle subject, verb, object \rangle$ triple. Pichotta and Mooney (2014); Granroth-Wilding and Clark (2016) further extended the event representation method by taking the indirect object into consideration. Following studies on script event prediction are mainly based on this event representation method. The above symbolic representation methods may cause the sparsity problem. Therefore, distributed event representation method is applied in more recent studies (Modi and Titov, 2014; Granroth-Wilding and Clark, 2016). In addition, early studies aggregate the scores of each event in the narrative event chain and the candidate event, ignoring the temporal order of the events. Therefore, Pichotta and Mooney (2016); Wang et al. (2017); Lv et al. (2019) introduced LSTM to integrate temporal order information¹. Li et al. (2018) further converted the narrative event chain to a narrative event evolutionary graph and used graph neural networks to model it.

Conventionally, an event contains a verb and several participants, where the headwords are used to represent the participants. This event representation method suffers from a lack of information since only a few words are considered. Therefore, Ding et al. (2019) used if-then commonsense knowledge to enrich the event representation. Lee and Goldwasser (2019); Zheng et al. (2020); Lee et al. (2020) labeled extra discourse relations be-

¹Following (Wang et al., 2017), we use the term “temporal order”. Precisely, it is the “narrative order”.

tween the events according to the conjunctions between them. Zheng et al. (2020); Lv et al. (2020); Lee et al. (2020) introduced pre-trained language models, such as BERT (Devlin et al., 2019), to script event prediction and achieved excellent performance. Still, these studies lose some of the informative constituents in the texts that directly describe the events.

The above models all focus on a single narrative event chain. The only study that considers the multiple chains is (Chambers and Jurafsky, 2009). However, our study is different from it in the following aspects: (1) when representing an event, it only considers the verb and the dependency relation between the verb and the protagonist, while we keep all the participants in an event; (2) it adopts the symbolic event representation method and a pair-based model (PMI), while we use distributed event representation method and consider the temporal order information.

3 Problem Statement

In this paper, an event $e = \langle v, a_0, a_1, a_2, t \rangle$ consists of a verb v , three participants² (i.e., subject a_0 , object a_1 , and indirect object a_2), and its original sentence t where the event appears. Here, $t = \{w_0, w_1, \dots\}$ consists of a sequence of words.

The script event prediction task aims to predict the most probable subsequent event e^* given the candidate event set \mathcal{S} and the historical events \mathcal{H} . Here, $\mathcal{S} = \{e_{c_0}, e_{c_1}, \dots, e_{c_{m-1}}\}$ consists of the m candidate events and $\mathcal{H} = \{e_0, e_1, \dots\}$ consists of the happened events. Note that, since the candidate events have not happened yet, there are no corresponding sentences for them.

4 The MCPredictor Model

In this section, we will describe our MCPredictor model. As show in Figure 2, it consists of four main components: (1) event encoding, (2) text encoding, (3) chain modeling, and (4) scoring. We will describe the model in single-chain mode and then introduce how the results from the multiple chains (derived from different participants of the candidate event) are aggregated. Finally, we will list the variants of the scoring component.

² Usually, an event only contains one or two participants, However, to uniformly characterize all possible conditions, we pad the missing participants with null.

4.1 Event Encoding Component

The event encoding component aims to map the events into low-dimensional vectors. The embeddings \mathbf{e}_i ³ for events e_i ($i \in \{0, \dots, n-1\}$, n is the number of historical events) is calculated via mapping the verbs v_i and their participants $a_{i,0}, a_{i,1}, a_{i,2}$ into the same vector space \mathbb{R}^{d_w} :

$$\mathbf{e}_i = \tanh(W_v^T \mathbf{v}_i + W_{a_0}^T \mathbf{a}_{i,0} + W_{a_1}^T \mathbf{a}_{i,1} + W_{a_2}^T \mathbf{a}_{i,2} + \mathbf{b}_e), \quad (1)$$

where $W_v, W_{a_0}, W_{a_1}, W_{a_2} \in \mathbb{R}^{d_w \times d_e}$ are mapping matrices; $\mathbf{b}_e \in \mathbb{R}^{d_e}$ is the bias vector; d_w is the word embedding size and d_e is the event embedding size. The candidate event e_c is encoded similarly.

4.2 Text Encoding Component

The text encoding component is to encode the sentence where each event appears into a vector. Since pre-trained language models, such as BERT (Devlin et al., 2019), have shown great power in encoding natural language, we apply BERT-tiny in this component and use the output embedding of the “[CLS]” tag as the sentence embedding.

Specifically, to avoid information leakage, we replace the other events in the same narrative event chain with the “[UNK]” tag when encoding the sentence of the focused event in this component. In addition, we add role tags before and after the verb to specify dependency relation (“[subj]”, “[obj]” or “[iobj]”) between the verb and the protagonist. For instance, the sentence “*He entered the restaurant and asked the waiter for the menu*” will be converted into “*He [subj] entered [subj] the restaurant and [UNK]*”, if we focus on the event corresponding to the verb “entered”.

Then, the embeddings of events e_i that integrate sentence information are:

$$\mathbf{e}'_i = \mathbf{e}_i + \text{BERT}(t_i), \quad (2)$$

where BERT is the BERT-tiny network and t_i are the converted sentence corresponding to events e_i .

4.3 Chain Modeling Component

In the above components, the events are separately encoded without temporal information. Therefore, we use stacked Transformer (Vaswani et al., 2017) layers to deeply integrate the temporal order information into the event representations. Following

³ In what follows, following the convention, the embeddings are denoted with the same letters but in boldface.

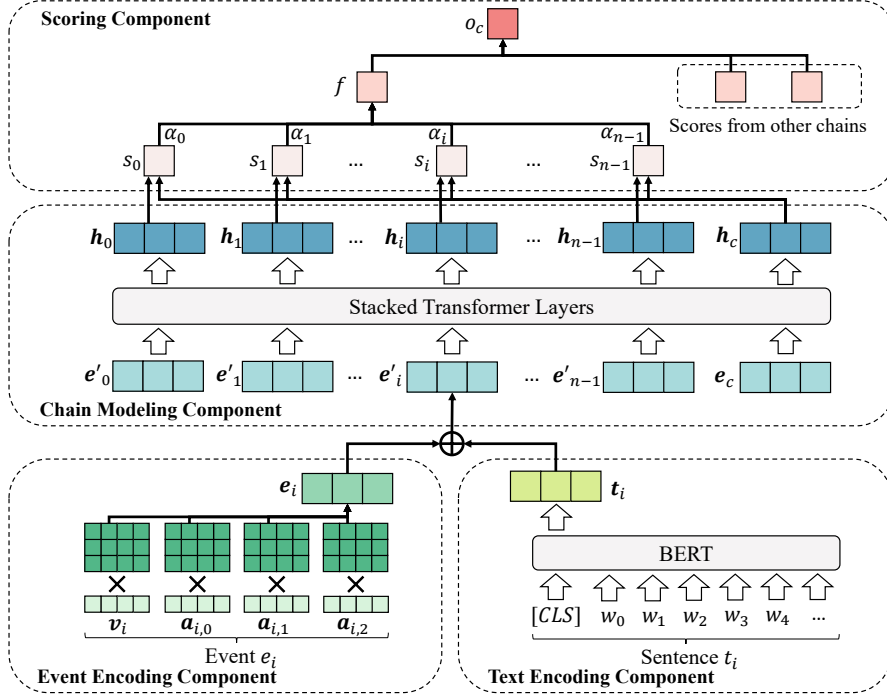


Figure 2: An illustrative diagram of the proposed MCPredictor model.

the convention (Wang et al., 2017; Li et al., 2018), to model the interactions between the historical events (i.e., the historical narrative event chain) and the candidate event, we append the candidate event to the end of the historical narrative event chain. As mentioned in Section 3, there is no corresponding sentence for the candidate event, thus, we only use its event embedding calculated by Equation 1. Additionally, the positional embeddings (Vaswani et al., 2017) are added to the event embeddings to specify their positions in the chain. Finally, the last Transformer layer outputs the hidden vectors \mathbf{h}_i and \mathbf{h}_c for historical events e_i and the candidate event e_c , respectively.

4.4 Scoring Component

The scoring component aims to calculate the score of each candidate event. This component contains two steps: event-level scoring and script-level aggregation. The former calculates the similarity score between each historical event and the candidate event. The latter aggregates these scores to derive the similarity score between the script and the candidate event. Besides, we will describe the variants of this component.

4.4.1 Event-Level Scoring

Aggregating the event-level similarity scores after modeling the temporal order is better than only

considering the event-pair similarity scores or only using the last hidden vector of the chain for prediction, as verified by (Wang et al., 2017). Therefore, before evaluating the script-level similarity scores, we calculate the pairwise similarity scores s_i between the candidate event e_c and historical events e_i using the hidden vectors from the chain modeling component in Section 4.3. Here, we use the negative Euclidean distance (denoted as E-Score) as the similarity scores:

$$s_i = -\|\mathbf{h}_i - \mathbf{h}_c\|_2. \quad (3)$$

4.4.2 Script-Level Aggregation

This step is to aggregate the event-level scores from multiple narrative event chains. Let us begin with a single narrative event chain.

We use an attention mechanism to derive the similarity score f between the historical narrative event chain and the candidate event e_c :

$$f = \sum_{i=0}^{n-1} \alpha_i s_i, \quad (4)$$

where the attention weight α_i of each event e_i is calculated by:

$$\alpha_i = \frac{\exp(u_i)}{\sum_{k=0}^{n-1} \exp(u_k)}, \quad (5)$$

where we use the scaled-dot product attention (Vaswani et al., 2017):

$$u_i = \frac{1}{\sqrt{d_e}} \mathbf{h}_i^T \mathbf{h}_c. \quad (6)$$

The similarity score is calculated by aggregating the similarity scores from multiple chains:

$$o_{c_i} = \sum_{j=0}^2 f_{i,j}, \quad (7)$$

where $f_{i,j}$ is the similarity score between the candidate event e_{c_i} and the historical narrative event chain corresponding to its j -th participant.

Then, the probability of each candidate event $e_{c_i} \in \mathcal{S}$ to be the correct subsequent event is calculated by:

$$\Pr(e_{c_i}|\mathcal{H}) = \frac{\exp(o_{c_i})}{\sum_{k=0}^{m-1} \exp(o_{c_k})}, \quad (8)$$

where o_{c_i} is the score of each candidate event.

Note that, in previous studies, only a single narrative event chain of the given protagonist is considered in the historical events \mathcal{H} . Therefore, the score of each candidate event o_{c_i} is equal to f_i , where f_i represents the similarity score of the historical narrative event chain and the candidate event e_{c_i} calculated via Equation 4.

Finally, we select the most possible candidate event as the subsequent event e^* as follows:

$$e^* = \arg \max_{e_{c_i} \in \mathcal{S}} \Pr(e_{c_i}|\mathcal{H}). \quad (9)$$

4.4.3 Variants

We try three other similarity score functions and three other attention functions on the proposed model. The score functions are:

- **C-Score** is the cosine similarity:

$$\text{C-Score}(\mathbf{h}_i, \mathbf{h}_c) = \frac{\mathbf{h}_i^T \mathbf{h}_c}{\|\mathbf{h}_i\| \|\mathbf{h}_c\|}. \quad (10)$$

- **M-Score** is the negative Manhattan distance:

$$\text{M-Score}(\mathbf{h}_i, \mathbf{h}_c) = -\|\mathbf{h}_i - \mathbf{h}_c\|_1. \quad (11)$$

- **L-Score** is the linear transformation score:

$$\text{L-Score}(\mathbf{h}_i, \mathbf{h}_c) = \mathbf{w}_{se}^T \mathbf{h}_i + \mathbf{w}_{sc}^T \mathbf{h}_c + b_s, \quad (12)$$

where $\mathbf{w}_{se}, \mathbf{w}_{sc} \in \mathbb{R}^{d_e}$ are the weight vectors and $b_s \in \mathbb{R}$ is the bias.

The attention functions are:

- **Average attention** simply averages the event-level scores;

- **Dot product attention** is calculated by:

$$u_i = \mathbf{h}_i^T \mathbf{h}_c. \quad (13)$$

- **Additive attention** is calculated by:

$$u_i = \mathbf{w}_{ae}^T \mathbf{h}_i + \mathbf{w}_{ac}^T \mathbf{h}_c + b_a, \quad (14)$$

where $\mathbf{w}_{ae}, \mathbf{w}_{ac} \in \mathbb{R}^{d_e}$ are the weight vectors and $b_a \in \mathbb{R}$ is the bias.

4.5 Training Details

The training objective is to minimize the cross-entropy loss:

$$L(\Theta) = -\frac{1}{N} \sum_{i=1}^N \log \Pr(e_i^*|\mathcal{H}_i) + \lambda L_2(\Theta), \quad (15)$$

where e_i^* and \mathcal{H}_i denote the correct answer and the historical events for the i -th training sample, respectively; N is the number of training samples; Θ denotes all the model parameters; $L_2(\cdot)$ denotes the L_2 regularization; λ denotes the regularization factor. We then train the model using Adam (Kingma and Ba, 2015) algorithm with 100-size mini-batch.

5 Experiments

In this section, we compare MCPredictor with a number of baselines to validate its effectiveness. In addition, we investigate the variants of MCPredictor and the importance of different constituents. Finally we conduct case studies on the model.

5.1 Dataset

Following (Granroth-Wilding and Clark, 2016), we extract events from the widely-used New York Time portion of the Gigaword corpus (Graff et al., 2003). The C&C tool (Curran et al., 2007) is used for POS tagging and dependency parsing, and OpenNLP⁴ for coreference resolution. For the participants that are coreference entities, we select the events they participate in to construct the event chain. For the non-coreference entities, we select the events by matching their headwords. The short chains will be padded with null events (verb and arguments are all null) to the max sequence length.

⁴<http://opennlp.apache.org>

# train documents	830,645
# development documents	103,583
# test documents	103,805
# train samples	1,440,295
# development samples	10,000
# test samples	10,000

Table 1: Detailed information about the dataset.

Since the development and test sets are unavailable, we generate them using the codes from [Granroth-Wilding and Clark \(2016\)](#) and their development and test document split. Thus, the results of the baseline models may be different from what they are reported in previous studies. Detailed information about the dataset is shown in Table 1, where #X denotes the size of X. Each sample in the training set, development set, and test set contains five candidate events, where only one choice is correct.

5.2 Experiment Settings

Following the convention, event sequence length n is set to 8; word embedding size d_w is set to 300; event embedding size d_e is set to 128; the number of Transformer layers is selected from $\{1, \underline{2}\}$; the dimension of the feed-forward network in Transformer is selected from $\{512, \underline{1024}\}$; the dropout rate is set to 0.1; the learning rate is set to $1e-4$ except that the learning rate of BERT-tiny is set to $1e-5$; the regularization factor λ is set to $1e-6$. All the hyper-parameters are tuned on the development set, and the best settings are underlined. All the experiments are run under Tesla V100.

5.3 Baselines

- **PMI** ([Chambers and Jurafsky, 2008](#)) uses PMI to measure the pairwise similarity of events.
- **Event-Comp** ([Granroth-Wilding and Clark, 2016](#)) uses multi-layer perceptron to encode events and calculates their pairwise similarity.
- **PairLSTM** ([Wang et al., 2017](#)) uses LSTM to integrate temporal order information into event embeddings.
- **SGNN** ([Li et al., 2018](#)) merges all the narrative event chains into a narrative event evolution graph. The scaled graph neural network is used to obtain graph information in event embeddings.

Method	Accuracy(%)
PMI*	30.52
Event-Comp*	49.57
PairLSTM*	50.83
SGNN*	52.45
SAM-Net*	54.48
Lv2020*	58.66
NG*	63.59
PMI	30.28
Event-Comp	50.19
PairLSTM	50.32
SGNN	52.30
SAM-Net	55.60
SCPredictor-s	58.28
MCPredictor-s	59.24
SCPredictor	66.24
MCPredictor	67.05

Table 2: Model accuracy on the test set. “*” denotes the performance from previous studies.

- **SAM-Net** ([Lv et al., 2019](#)) uses an LSTM network and self-attention mechanism to capture semantic features.
- **NG** ([Lee et al., 2020](#)) builds a narrative graph to enrich the event semantic representation via introducing discourse relations.
- **Lv2020** ([Lv et al., 2020](#)) introduces external commonsense knowledge and uses a BERT-based model to predict the subsequent event.
- **SCPredictor**, an ablation of MCPredictor, removes the scores from the other narrative event chains at the script level.
- **SCPredictor-s and MCPredictor-s**, ablations of SCPredictor and MCPredictor, respectively, remove the sentence information at the event level.

For PMI, we use the version implemented by ([Granroth-Wilding and Clark, 2016](#)). For PairLSTM, we adopt the version implemented by ([Li et al., 2018](#)). For others, we use the versions provided in their original papers. For comparison, accuracy(%) of predicting the correct subsequent events is used as the evaluation metric.

5.4 Results and Analyses

Final experimental results on the script event prediction task are shown in Table 2. From the results, we have the following observations:

- SCPredictor and MCPredictor outperform SCPredictor-s and MCPredictor-s respectively by more than 7.81%. This significant improvement indicates that introducing sentence information at the event level successfully enhances the event representations;
- MCPredictor-s and MCPredictor outperform SCPredictor-s and SCPredictor respectively, which demonstrates the effectiveness of introducing multiple narrative event chains at the script level. In addition, the improvement of multiple chains decreases slightly when the sentence information is introduced (from 0.96% to 0.81%). This is probably because extra events brought by other narrative event chains may partially covered by the sentences corresponding to existing events. In the development set, only about 13% of the samples contain the extra events that are not covered by the sentences corresponding to existing events when introducing other narrative event chains. Still, the multiple narrative event chains contribute to the model;
- SCPredictor-s outperforms existing models (i.e., PMI, Event-Comp, PairLSTM, SGNN, and SAM-Net) by more than 2.68% under the same input (only the verb and the headwords of the participants are used). This improvement indicates that the Transformer networks can model the script better than existing network structures.

5.5 Human Evaluation

In addition to automatic evaluation, we also conduct human evaluation to further study the performance of the proposed MCPredictor model. Specifically, we randomly select 100 samples in the development set. Without loss of generality, we select SAM-Net model to compare with MCPredictor manually under these samples. MCPredictor correctly answers 72 out of 100, while SAM-Net only correctly answers 60. These results demonstrate that the MCPredictor generally derives more plausible answers.

5.6 Comparative Studies on Variants

To further study the effects of different scoring functions and attention functions, we conduct comparative studies on the development set.

The comparative results on different scoring functions are listed in Table 3, where Δ means the

Method	Accuracy(%)	Δ
E-Score	67.35	/
M-Score	65.36	-1.99
L-Score	65.05	-2.30
C-Score	60.68	-6.63

Table 3: Performance comparison of MCPredictor over different scoring functions.

Method	Accuracy(%)	Δ
scaled-dot	67.35	/
dot	66.21	-1.14
additive	65.84	-1.51
avg	64.74	-2.61

Table 4: Performance comparison of MCPredictor over different attention functions.

performance difference between the best scoring function and each alternative scoring function. As presented in Table 3, E-Score achieves the best performance among the four scoring functions, which is consistent with the result in (Li et al., 2018). The reason for C-Score underperforming the other three scoring functions may be that it only measures the angle between the two vectors, however, ignores their lengths.

Similarly, the comparative results on different attention functions are listed in Table 4. In this table, scaled-dot is the scaled-dot product attention in Equation 6, dot indicates the dot product attention in Equation 13, additive is the additive attention in Equation 14, and avg denotes the average attention. Similar to Table 3, Δ means the performance difference. As shown in Table 4, the scaled-dot product attention outperforms the others. The average attention underperforms the other attention functions, which shows the effectiveness of the attention mechanism. The reason for dot product attention underperforming the scaled-dot attention function may be that dot product attention tends to give a single event too much weight, which may suppress the contributions of other related events.

5.7 Detailed Analyses on Constituents

What is the contribution of each constituent from the sentences? To go deep into it, we use a masking mechanism to hide some constituents. The results on the development set are listed in Table 5, where “All” denotes that all constituents are used; “V” denotes the verbs; “N” denotes the nouns; “J” denotes the adjectives; “R” denotes the adverbs; and “-X”

Protagonist A: First Virtual	Original Sentence
set+up(executive, company)	First Virtual, which Unger mann <i>set up</i> in ...
demonstrate(executive, product)	First Virtual ... publicly <i>demonstrated</i> its products ...
aim(company, null, segment)	First Virtual is <i>aiming</i> at an emerging segment ...
be(company, able)	First Virtual <i>is</i> able to <i>offer</i> its lower prices by
offer(company, price)	<i>limiting</i> its ATM network speeds to 25 million bits
limit(company, speed, bit)	a second ...
sneak(company, null, organization)	it could <i>sneak</i> into organizations ...
develop(company, software)	First Virtual has also <i>developed</i> software ...
protagonist B: Ralph Unger mann	Original Sentence
consider(father, executive)	Ralph Unger mann ... <i>is considered</i> a founding father ...
set+out(executive, problem)	Ralph Unger mann ... has <i>set out</i> to solve his newest company ...
set+up(executive, company)	First Virtual, which Unger mann <i>set up</i> in ...
predict(executive)	Unger mann <i>predicts</i> that ...
contend(executive)	Unger mann <i>contends</i> that ...
correct choice: sell(executive, company); wrong choice: launch(company, attack)	

Table 6: Case study of MCPredictor. The historical events used by SCPredictor are in boldface.

Constituents	Accuracy(%)	Δ
All	67.35	/
-V	60.54	-6.81
-N	64.81	-2.54
-J	66.89	-0.46
-R	67.03	-0.32
-V(self)	67.29	-0.06
-V(others)	60.90	-6.45
-V&N	55.36	-11.99
-V&R	60.23	-7.12
-N&J	64.10	-3.25
-R&J	66.55	-0.80

Table 5: Performance comparison of MCPredictor over different constituents.

denotes that X is masked. In addition, we study the influence of their pairwise combinations X and Y (denoted as “X&Y”). Considering the modification relationship between the constituents, we only conduct the experiments on the combinations of “V&R”, “V&N”, “N&J”, and “R&J”, respectively.

Masking the verbs brings 6.81% decrease to the performance. To further study the influences of different verbs, we mask the focused verb (denoted as “V(self)”) and the verbs excepting the focused one (denoted as “V(others)”) separately. The results show that masking the focused verb almost do no harm to the performance, while masking others brings 6.45% decrease to the performance. This is because the event encoding component still obtains the information about the focused event despite the

text encoding component masking it. In addition, these results show that the implicit relevance of the other verbs and the focused verb is important for script event prediction.

Masking the nouns brings 2.54% decrease to the performance. The nouns are usually the participants which influence the word sense of the verbs. Thus, the nouns are also important for prediction.

Masking the adjectives brings 0.46% decrease to the performance when comparing “All” and “-J”, while it brings 0.71% decrease when comparing “-N” and “-N&J”. This phenomenon shows that using the combination of the nouns and the adjectives brings richer semantic information than using them separately. The same phenomenon also appears when comparing “All”, “-V”, and “-V&R”.

5.8 Case Studies

To have a better understanding of the effects of multiple narrative event chains, we study the cases in the development set where the MCPredictor model selects the correct choice while the SCPredictor model selects the wrong one. As presented in Table 6, protagonist A is “company” in the events, and protagonist B is “executive” in the events. When only the narrative event chain of protagonist A is provided, the correct candidate event is not very convincing, since protagonist B is rarely mentioned in this chain. In addition, the events in this chain are common, and many events can be the subsequent event. Thus, SCPredictor is likely to select the wrong candidate event as the prediction result.

On the contrary, MCPredictor integrates evidence from both narrative event chains. It is thus able to predict the correct candidate event. Moreover, the events contain less information comparing to the original sentences, which shows the necessity to introduce the sentence information for them.

6 Conclusion and Future Work

In this paper, we proposed the MCPredictor model to handle the script event prediction task, which integrates deep event-level and script-level information. At the event level, MCPredictor additionally utilizes the original sentence to obtain more informative event representation. At the script level, it considers multiple event sequences corresponding to different participants of the subsequent event. Experimental results demonstrate its merits and superiority.

However, currently, we still need to know the participants of the candidate events to extract multiple chains. When participants are unknown, we have to enumerate all possible combinations of entities in the script, which is time-consuming. Moreover, extracting informative constituents of events besides the verb and their participants is still a challenge. In the future, we will study these problems.

Acknowledgment

The work is supported by the National Natural Science Foundation of China under grants U1911401, 62002341, 61772501, U1836206 and 61722211, the GFKJ Innovation Program, Beijing Academy of Artificial Intelligence under grant BAAI2019ZD0306, the Lenovo-CAS Joint Lab Youth Scientist Project, and the Foundation and Frontier Research Key Program of Chongqing Science and Technology Commission (No. c-stc2017jcyjBX0059). Thanks to Yixing Fan, Huawei Shen, Jinhua Gao, Yunqi Qiu, Xiao Ding, and the reviewers for the constructive discussions and suggestions.

References

Robert Abelson and Roger C Schank. 1977. Scripts, plans, goals and understanding. *An inquiry into human knowledge structures New Jersey*, 10.

Niranjan Balasubramanian, Stephen Soderland, Mausam, and Oren Etzioni. 2013. [Generating coherent event schemas at scale](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1721–1731, Seattle,

Washington, USA. Association for Computational Linguistics.

- David Bean and Ellen Riloff. 2004. [Unsupervised learning of contextual role knowledge for coreference resolution](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 297–304, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Nathanael Chambers and Dan Jurafsky. 2008. [Unsupervised learning of narrative event chains](#). In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.
- Nathanael Chambers and Dan Jurafsky. 2009. [Unsupervised learning of narrative schemas and their participants](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610, Suntec, Singapore. Association for Computational Linguistics.
- Snigdha Chaturvedi, Haoruo Peng, and Dan Roth. 2017. [Story comprehension for predicting what happens next](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1603–1614, Copenhagen, Denmark. Association for Computational Linguistics.
- James Curran, Stephen Clark, and Johan Bos. 2007. [Linguistically motivated large-scale NLP with c&c and boxer](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 33–36, Prague, Czech Republic. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xiao Ding, Kuo Liao, Ting Liu, Zhongyang Li, and Junwen Duan. 2019. [Event representation learning enhanced with external commonsense knowledge](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4893–4902, Hong Kong, China. Association for Computational Linguistics.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.

- Mark Granroth-Wilding and Stephen Clark. 2016. [What happens next? event prediction using a compositional neural network model.](#)
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization.](#) In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.*
- I-Ta Lee and Dan Goldwasser. 2019. [Multi-relational script learning for discourse relations.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4214–4226, Florence, Italy. Association for Computational Linguistics.
- I-Ta Lee, Maria Leonor Pacheco, and Dan Goldwasser. 2020. [Weakly-supervised modeling of contextualized event embedding for discourse relations.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4962–4972, Online. Association for Computational Linguistics.
- Zhongyang Li, Xiao Ding, and Ting Liu. 2018. [Constructing narrative event evolutionary graph for script event prediction.](#) In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4201–4207. International Joint Conferences on Artificial Intelligence Organization.
- Shangwen Lv, Wanhui Qian, Longtao Huang, Jizhong Han, and Songlin Hu. 2019. [Sam-net: Integrating event-level and chain-level attentions to predict what happens next.](#) In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6802–6809.
- Shangwen Lv, Fuqing Zhu, and Songlin Hu. 2020. [Integrating external event knowledge for script learning.](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 306–315, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Ashutosh Modi and Ivan Titov. 2014. [Learning semantic script knowledge with event embeddings.](#) In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings.*
- Karl Pichotta and Raymond Mooney. 2014. [Statistical script learning with multi-argument events.](#) In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 220–229, Gothenburg, Sweden. Association for Computational Linguistics.
- Karl Pichotta and Raymond Mooney. 2016. [Learning statistical scripts with lstm recurrent neural networks.](#)
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need.](#) In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Zhongqing Wang, Yue Zhang, and Ching-Yun Chang. 2017. [Integrating order information and event relation for script event prediction.](#) In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 57–67, Copenhagen, Denmark. Association for Computational Linguistics.
- Jianming Zheng, Fei Cai, and Honghui Chen. 2020. [Incorporating scenario knowledge into a unified fine-tuning architecture for event representation.](#) In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 249–258.