

A Generative Approach for Script Event Prediction via Contrastive Fine-Tuning

Fangqi Zhu¹, Jun Gao¹, Changlong Yu, Wei Wang, Chen Xu³,
Xin Mu², Min Yang⁴, Ruifeng Xu^{1,2,5*}

¹ Harbin Institute of Technology, Shenzhen

² Peng Cheng Laboratory

³ Beijing University of Technology

⁴ Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

⁵ Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies

{zhufangqi.hitsz, imgaojun}@gmail.com, xuruifeng@hit.edu.cn

Abstract

Script event prediction aims to predict the subsequent event given the context. This requires the capability to infer the correlations between events. Recent works have attempted to improve event correlation reasoning by using pretrained language models and incorporating external knowledge (e.g., discourse relations). Though promising results have been achieved, some challenges still remain. First, the pretrained language models adopted by current works ignore event-level knowledge, resulting in an inability to capture the correlations between events well. Second, modeling correlations between events with discourse relations is limited because it can only capture explicit correlations between events with discourse markers, and cannot capture many implicit correlations. To this end, we propose a novel generative approach for this task, in which a pretrained language model is fine-tuned with an event-centric pretraining objective and predicts the next event within a generative paradigm. Specifically, we first introduce a novel event-level blank infilling strategy as the learning objective to inject event-level knowledge into the pretrained language model, and then design a likelihood-based contrastive loss for fine-tuning the generative model. Instead of using an additional prediction layer, we perform prediction by using sequence likelihoods generated by the generative model. Our approach models correlations between events in a soft way without any external knowledge. The likelihood-based prediction eliminates the need to use additional networks to make predictions and is somewhat interpretable since it scores each word in the event. Experimental results on the multi-choice narrative cloze (MCNC) task demonstrate that our approach achieves better results than other state-of-the-art baselines. Our code will be available at <https://github.com/zhufq00/mcnc>.

Introduction

A script (Schank and Abelson 2013) refers to a kind of structured knowledge, which involves sequences of events. Script event prediction aims at predicting the subsequent event given an event chain in a script. Figure 1 presents a restaurant script involving sequences of events that occurs when a person enters a restaurant. The goal of the script event prediction task (Granroth-Wilding and Clark 2016) is

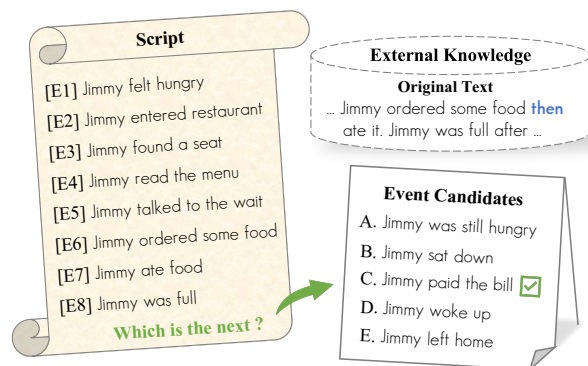


Figure 1: Script event prediction with external knowledge. Giving a script, the task aims to select the subsequent event from event candidates. Events E6~E8 are extracted from the original text where “then” is a discourse marker, which is used by Lee, Pacheco, and Goldwasser (2020) to extract the discourse relation. Note that our method is **free** with this kind of external knowledge (dashed box).

to select the plausible subsequent event from the five event candidates. Solving the script event prediction problem can help us to acquire event knowledge from event chains and is beneficial for many downstream tasks such as story generation (Chaturvedi, Peng, and Roth 2017) and dialogue generation (Danescu-Niculescu-Mizil and Lee 2011).

Understanding events and inferring correlations between events are crucial to acquiring event knowledge. Taking Figure 1 for example, we first need to understand the events (what is happening) and have the commonsense knowledge - “You need to pay after the meal”, and then we can infer the plausible subsequent event - “Jimmy paid the bill” (what happens next). Early works (Granroth-Wilding and Clark 2016; Wang, Zhang, and Chang 2017) used embedding method (e.g., Word2Vec) to learn event representation and then choose the subsequent event by calculating similarities between event candidates and script events. After that, some works (Lv, Zhu, and Hu 2020) attempted to leverage masked language model (MLM) (Liu et al. 2019) to obtain better event representations and they all achieved a substantial improvement. However, some studies (Lee, Pacheco, and Goldwasser 2020) have shown that

*Corresponding author.

these approaches based on pre-trained models still lack an understanding of the relationships between events. Some recent works (Lee, Pacheco, and Goldwasser 2020; Bai et al. 2021) incorporate external knowledge to enhance the model to understand correlations between events, including discourse relations (Bai et al. 2021) and original texts (Lee, Pacheco, and Goldwasser 2020). Lee, Pacheco, and Goldwasser (2020) (NG) extract discourse relations from original texts by a template matching-based method. Bai et al. (2021) use original texts to enhance event representations (See Figure 1 for example).

Despite promising results achieved by these methods, some challenges still remain. First, the pretrained language models adopted by current works ignore event-level knowledge. Specifically, MLMs are pre-trained using token-level learning objectives and thus unable to capture the correlations between events. Second, modeling correlations between events with discourse relations is limited because it can only capture explicit correlations between events with discourse markers, and cannot capture many implicit correlations. For example, in Figure 1, a model can easily recognize a *precedence* relation between E6 event and E7 event based on the discourse marker “then”, while the model struggles to recognize the implicit event relation between E7 event and E8 event when there are no discourse markers.

To better model correlations between events and transfer the knowledge of pre-trained models (including general knowledge and event knowledge) to the script event prediction task, we propose a novel generative approach for this task, in which a pretrained language model is fine-tuned with an event-centric pretraining objective and predicts next event within a generative paradigm. Specifically, our approach consists of two stages: event-centric pretraining stage and contrastive fine-tuning stage. (1) In the event-centric pretraining stage, We use the BART (Lewis et al. 2020) as the backbone for this task and introduce a novel event-level blank infilling strategy as the learning objective to inject event-level knowledge into the pretrained language model. We randomly mask several events in the script and then require the model to generate these masked events autoregressively. (2) In the contrastive fine-tuning stage, we design a likelihood-based contrastive loss to force the model to learn to distinguish between correct and wrong event candidates. To be specific, the model is required to reduce the generation probabilities of wrong event candidates and increase the generation probability of correct event candidates. Instead of using an additional prediction layer, we perform prediction by using sequence likelihoods generated by the generative model. Our approach models correlations between events in a soft way without any external knowledge and maintain the consistency of the contrastive fine-tuning, event-centric pretraining, and pretraining objectives. Besides, the likelihood-based prediction eliminates the need to use additional networks to make predictions and is somewhat interpretable since it scores each word in the event. The main contributions of this work are summarized as follows:

- We propose a novel generative approach for this task, in which a pretrained language model is fine-tuned with an event-centric pretraining objective and predicts the next

event within a generative paradigm.

- We introduce a novel event-level blank infilling strategy as the learning objective to inject event-level knowledge into the pretrained language model and design a likelihood-based contrastive loss to force the model to learn to distinguish between correct and wrong event candidates.
- Experimental results on the multi-choice narrative cloze (MCNC) task demonstrate that our approach achieves better results than other state-of-the-art baselines.

The Proposed Approach

As shown in Figure 1, script event prediction can be defined as predicting the most probable subsequent event given a script (an existing event chain). Formally, given a script $X = \{x_1, \dots, x_n\}$ and event candidates $Y = \{y_1, \dots, y_m\}$ where x_i and y_j represent event, this task aims to choose the correct subsequent event y_t from Y . Following Granroth-Wilding and Clark (2016), each event $e = (e_v, e_s, e_o, e_i)$ consists of a predicate e_v and three arguments: subject e_s , object e_o and indirect object e_i . If an event does not have an argument, the argument will be denoted as `null`. Our proposed approach involves modeling a conditional distribution $P(Y|X)$. Our model should learn to calculate the conditional probability distribution $P(y_i|X)$ for each event candidate $y_i (i \in 1, \dots, m)$, given the script X .

During training, our method consists of two stages: event-centric pretraining and task-specific contrastive fine-tuning. In the event-centric pretraining stage, we randomly mask some events in the script and then enable the model to generate the masked events sequentially. By doing so, our model models correlations between events and mines the event knowledge in the script. In the task-specific contrastive fine-tuning stage, we introduce wrong event candidates to make the model learn to distinguish between right and wrong event candidates. We train our model to maximize the generation probability of the right event candidate and reduce the generation probabilities of the wrong event candidates. During testing, our model calculates the generation probability of each event candidate and selects the one with the highest generation probability as the predicted subsequent event.

In recent years, pre-trained models like BART (Lewis et al. 2020) have significantly improved over various downstream tasks such as question answering, summarization, and machine translation. Therefore, we adopt BART as the underlying architecture for our model to model the conditional probability distribution $P(Y|X)$.

Event-Centric Pretraining

In the event-centric pretraining stage, we aim to train our model to model correlations between events and mine the event knowledge in the script by an event-level blank infilling strategy. For a more complete usage of the script, we concatenate the script and the right event candidate y_t to become the event sequence $S_{ori} = \{x_1, \dots, x_n, y_t\}$. To enrich the variety of mask strategies and maintain the




Figure 2: Our two-stage generative model training method. S and E correspond to the input and output of the generative model. t_i corresponds to the probability of generating the token corresponding to t_i given S and $E^{0:i-1}$. In the task-specific contrastive fine-tuning stage, the model forwards five times to get the overall scores of the five event candidates, respectively.

challenge of the event-centric pretraining stage, we randomly mask K events in S_{ori} , where K is randomly selected from 1 to 3. We denote the masked S_{ori} as S and the masked events as event sequence E where the masked events are in the order they were in S_{ori} . Since the BART model requires natural language as input and output, as shown in Figure 2, we represent each event $e = (e_v, e_s, e_o, e_i)$ as “ $e_s e_v e_o e_i$ ” and the sequence of events S as “ $\langle s \rangle x_1 \langle \text{SEP} \rangle x_2 \langle \text{SEP} \rangle \langle \text{MASK} \rangle \langle \text{SEP} \rangle \dots \langle \text{SEP} \rangle \langle \text{MASK} \rangle \langle \text{SEP} \rangle \dots \langle \text{SEP} \rangle y_t \langle s \rangle$ ” where each masking event is replaced with a token $\langle \text{MASK} \rangle$ and each event is separated by $\langle \text{SEP} \rangle$. Event sequence E is represented similarly. Note that the first event x_1 and the last event y_t may also be masked. Since the $\langle \text{SEP} \rangle$ token does not exist in the BART model, we use “.” as the separator.

With the help of the autoregressive generative model BART, the conditional probability distribution $P(E|S)$ is formulated as follows:

$$P(E|S) = \frac{1}{N_E} \sum_{n=2}^{N_E} \log P_{LM}(E^n | S, E^{1:n-1}) \quad (1)$$

where N_E is the token number of event sequence E in natural language format, E^n is the n th token, and $E^{1:n-1}$ is the first to $n-1$ th token of E .

In the event-centric pretraining stage, we aim to train our model to maximize $P(E|S)$. We test the model on the development set and use the model parameters of the best epoch as initial parameters of the model in the next stage.

Task-Specific Contrastive Fine-tuning

In the task-specific contrastive fine-tuning stage, negative event candidates are introduced to support the generative model with discriminative ability. This step further enables the model to select the correct event from the event candidates, which is crucial to the final performance of the model. We first calculate the conditional generation probability of each event candidate as their scores. We then tried three different loss functions to optimize the scores so that the score of the right event candidate is higher than that of the negative event candidate in different ways.

More specifically, we first add a token [mask] at the end of the script X , denoting the modified sequence as X_m . Then X_m and each event candidate $y_i \in Y$ are converted to a natural language format, as mentioned in Section 2.1. Since each event candidate has a different number of tokens, to avoid the model tending to choose shorter event candidates, we utilize the mean of log-probabilities of verbalizer tokens for each event as the score o_i for the event y_i .

$$o_i = \frac{1}{N_{y_i}} \sum_{n=2}^{N_{y_i}} \log P_{LM}(y_i^n | X_m, y_i^{1:n-1}) \quad (2)$$

where N_{y_i} is the token number of the event y_i . Afterwards, we use softmax function to calculate the final score s_i for each event candidate y_i :

$$s_i = \frac{\exp(o_i)}{\sum_{k=1}^M \exp(o_k)} \quad (3)$$

Finally, we define the loss function as follows:

$$\mathcal{L}_{cot} = -\log(s_t) + \frac{1}{M-1} \sum_{i=1, i \neq t}^M \left(\frac{s_i}{1-s_t} \right) \log \left(\frac{s_i}{1-s_t} \right) \quad (4)$$

where t is the subscript of the right event candidate y_t , and M is the number of event candidates. The first loss term is the conventional softmax cross entropy objective, which aims to maximize the generation probability of the right event candidate. The second loss term takes inspiration from Chen et al. (2019), which can be understood as directly maximizing the entropies of the generation probabilities of negative events conditioned on the right event candidate y_t not happening. The difference is that Chen et al. (2019) uses two different optimizers to optimize the two loss terms, and our method performs better with pilot experiments.

Further, we explore two other loss functions to explore how to optimize the generation probability of event candidates better, including traditional cross entropy loss and margin ranking loss widely used in contrastive learning.

• **Cross Entropy Loss** is calculated as follows:

$$\mathcal{L}_{cross} = -\log(s_t) \quad (5)$$

• **Margin Ranking Loss** is calculated as follows:

$$\mathcal{L}_{margin} = \sum_{1 \leq i \leq M, i \neq t} \max(m - (s_i - s_t), 0) \quad (6)$$

where m is a hyperparameter margin which controls the interval between the score of the right event candidate and the score of the negative event candidate. The loss will be zero when all s_i minus s_t is greater than the margin m .

During testing, we select the most probable event candidate with the highest score as the predicted subsequent event y_p , where $p = \arg \max_{1 \leq i \leq M} s_i$.

Experiments

In this section, we introduce the datasets, experimental setting and compared baselines. Experimental results show our method achieves state-of-the-art performance on the multi-choice narrative cloze (MCNC) task. We then perform ablation study and model training comparison to understand the effect of the model’s key components and their variants on performance. Finally, we conduct case study to show how our model predicts the subsequent event.

Datasets

In the literature on script event prediction tasks, there are two versions of benchmark datasets that are used to evaluate the methods. For the original MCNC task, Granroth-Wilding and Clark (2016) extracted event chains from the widely-used New York Times portion of the Gigaword corpus (Graff et al. 2003). We used the released codes¹ to reproduce the extraction pipeline, including pos tagging, dependency parsing, and coreference resolution, etc. After generating event chains, we manually filter the extremely short event chains and truncate the long event chain so that the length of the remaining event chain (defined as script) is nine. The last event of the script is used for the positive event candidate. Negative event candidates are randomly sampled, where the protagonist is kept the same as the protagonist of the current script, and other arguments (*object* or *indirect object*) are replaced randomly by other arguments from the same document. Finally, the dataset contains more than 1.4 million event chains, and we split the training, validation, and test sets as Bai et al. (2021). We denote it as the “original dataset,” and the statistics are shown in Table 1.

Due to the complex extraction pipeline and restricted access to Gigaword corpus, most of the existing work chooses the “public dataset” released by Li, Ding, and Liu (2018). Compared with the original dataset, the public dataset only covers the event chains that include extracted event relations between them. This filtering results in ten times smaller than the original one. We also follow the common practice of dataset split for training, validation, and testing in Table 1. For both “original” and “public” datasets, each instance has five event candidates for both datasets, of which only one choice is correct. To demonstrate the effectiveness of our proposed approach, we choose to evaluate our model on both of them.

¹https://mark.granroth-wilding.co.uk/papers/what_happens_next/

	Original Dataset	Public Dataset
Train set	1,440,295	140,331
Dev set	10,000	10,000
Test set	10,000	10,000

Table 1: The statistics of the reproduced original dataset (Granroth-Wilding and Clark 2016) and the public dataset (Li, Ding, and Liu 2018).

Experimental Settings

Our proposed two-stage method includes event-centric pre-training and task-specific contrastive fine-tuning. To compare with baselines of different sizes of datasets, we conduct experiments on BART_{base} and BART_{large}. For these two stages, the model is optimized by Adam (Kingma and Ba 2014). The learning rate and weight decay are 1e-5 and 1e-6. Our model uses an early stop strategy to select the best epoch, with patience set to 5. For BART_{base}, the batch size is 256 and 64 in the two stages, respectively. For BART_{large}, the batch size is 128 and 32 in the two stages, respectively. For the task-specific contrastive fine-tuning stage, the learning rate is chosen from {1e-5, 2e-5, 3e-5}. All the experiments are conducted on Tesla A100 GPU. We choose the model with the best result on the development set and report its result on the test set. Accuracy is adopted as the evaluation metric.

Baselines

For a comprehensive with the previous methods, we divided them into the following categories:

Event Representations: 1) **Event-Comp** (Granroth-Wilding and Clark 2016) uses the training objectives like Word2Vec (Mikolov et al. 2013) to learn event embeddings and calculates pairwise similarities between script events and event candidates. 2) **Pair-LSTM** (Wang, Zhang, and Chang 2017) uses LSTM to model the narrative order of script events. 3) **SAM-Net** (Lv et al. 2019) uses LSTM and self-attention mechanism to capture diverse event segments. 4) **MCPredictor** (Bai et al. 2021) obtains event representations from pretrained Word2Vec and enhances them with original sentence representations obtained by pretrained BERT. Moreover, multiple similar event chains are utilized to aggregate script-level information to help select the subsequent event. 5) **SCPredictor-s** is an ablation of MCPredictor, removing additional similar scripts and the original sentence information. 6) **BERT-based SCPredictor-s** replaces Word2Vec with BERT for a fair comparison under a similar amount of parameters. 7) **BART** (Lewis et al. 2020) fine-tunes the pre-trained model BART with a linear classifier.

Methods Enhanced with Structured Information: 1) **NG** (Lee, Pacheco, and Goldwasser 2020) extracts event relations from the original text of the script, then a narrative event graph is constructed based on the event relations. Finally, NG transfers this task to a link prediction task (Chang et al. 2014). 2) **RoBERTa + Rep. Fusion** (Lv, Zhu, and Hu 2020) integrates external knowledge from

Methods	Acc. (%)	ext.
<i>w/ external knowledge</i>		
SGNN + Int&Senti	56.03	Int & Senti
RoBERTa _{base} + Rep. Fusion	58.66	ASER
RoBERTa _{base} + Know. Model	59.99	ASER
<i>w/o external knowledge</i>		
Random	20.00	w/o ext.
Event-Comp	49.57	w/o ext.
PairLSTM	50.83	w/o ext.
SGNN	52.45	w/o ext.
GraphBERT	<u>60.72</u>	w/o ext.
BART _{base}	60.00	w/o ext.
Ours (BART_{base})	62.94	w/o ext.

Table 2: Base model accuracy on the test set of the public dataset. *ext.* is short for *external knowledge*.

Methods	Acc. (%)	ext.
<i>w/ external knowledge</i>		
EventBERT	63.50	BOOKCORPUS
RoBERTa _{large} + Know. Model	63.62	ASER
ClarET	64.61	BOOKCORPUS
<i>w/o external knowledge</i>		
Ours (BART_{large})	<u>64.82</u>	w/o ext.
Ours (BART_{large}) + NYT	65.88	NYT

Table 3: Large model accuracy on the test set of the public dataset. NYT denotes the New York Time portion of the Gigaword corpus.

eventuality knowledge graph, ASER (Zhang et al. 2020) and uses RoBERTa for making prediction. 3) **RoBERTa + Know. Model** (Zhou et al. 2021) learns a knowledge model from ASER to predict event relations. 4) **SGNN** (Li, Ding, and Liu 2018) constructs a narrative event evolution graph via verb con-occurrence frequency to obtain more effective event representations. 5) **SGNN+Int&Senti** incorporates external intention and sentiment knowledge from ATOMIC (Sap et al. 2019) into event representations. 6) **GraphBERT** (Du et al. 2022) builds an event graph similar to SGNN and enhances BERT with the event graph. **Event-centric Post-Pretraining:** 1) **EventBERT** (Zhou et al. 2022a) pretrains RoBERTa on BOOKCORPUS (Zhu et al. 2015) with three self-supervised contrastive learning objectives: *Correlation-based Event Ranking*, *Contradiction Event Tagging* and *Discourse Relation Ranking*. 2) **ClarET** (Zhou et al. 2021) pretrains BART on BOOKCORPUS with another three self-supervised objectives: *Whole Event Recovering*, *Contrastive Event-correlation Encoding* and *Prompt-based Event Locating*.

Results and Analyses

We first present the main experimental results of the widely-used “public” dataset in Table 2 and Table 3 using the base and large model, respectively, in order to align with the existing baselines for comparable parameters. We can draw the following observations from the results on two model settings: 1). our approach achieves the new state-of-the-art performance with a comparable amount of parameters and

Methods	Acc. (%)	ext.
<i>w/ external knowledge</i>		
NG	63.59	discourse relation
MCPredictor	<u>67.14*</u>	original text
<i>w/o external knowledge</i>		
SAM-Net	55.60	w/o ext.
SCPredictor-s	58.79*	w/o ext.
BERT-based SCPredictor-s	59.13	w/o ext.
Ours (BART_{base})	67.21	w/o ext.

Table 4: Model accuracy on the test set of the original dataset. “*” indicate our reproduced results on our preprocessed original dataset.

obtains an absolute 2.22% over the best baseline *GraphBERT* without any external knowledge. To reduce the sparsity of the constructed event graph, GraphBERT uses the verb bigrams’ frequency of two events to model the weights of event relations. By doing so, GraphBERT can only model correlations between verbs and ignore the event arguments. By contrast, our generative approach enables our model to capture correlations among all the events in one event chain. 2). Although directly fine-tuning BART achieves much stronger performance than BERT or RoBERTa models, our approach further improves and increases the accuracy up to 2.94% compared to vanilla BART_{base}. 3). Moreover, our approach even outperforms strong baselines that perform heavy event-centric post-pretraining such as ClarET (Zhou et al. 2022b) and EventBERT (Zhou et al. 2022a). Note that our approach(except Ours (BART_{large}) + NYT) requires no external knowledge (either from the corpus itself or eventuality knowledge graph like ASER (Zhang et al. 2020)). 4). To push the boundary of our proposed approach and test whether more resources can further enhance our approach, we replace the training set of the “public dataset” with the training set of the “original datasets” and then evaluate on the test set of the “public dataset”. The result (Ours (BART_{large}) + NYT in Table 3) demonstrates that more extensive training data can learn more precise event knowledge and final gains 1.27% improvement over the previous state-of-the-art model ClarET. Finally, it is worth mentioning that our event-centric pretraining stage only needs 3.2 GPU hours on 0.76M tokens, while ClarET needs 90 GPU hours at 200M tokens (external Bookcorpus). It significantly reduces the training time and requires much fewer computational resources, further demonstrating the efficiency of our approach.

Then we evaluate our approach on the original dataset, show the experimental results in Table 4 and observe the following key findings. 1). Our method still achieves comparable performance to MCPredictor, although MCPredictor improves performance by 8.35% with additional original sentence and multi-script knowledge aggregation. As we know, the original sentences where the events are extracted are critical during the training. However, the original sentences required by MCPredictor may not exist during testing. We tried to use the BM25 algorithm to retrieve the original text in the training set according to the test set’s events to replace the original text in the test set. The accuracy dropped

Methods	Acc. (%)
Ours (BART_{base})	62.94
w/o event-centric pretraining	61.08
w/o task-specific contrastive fine-tuning	40.00
replace with a linear classifier	61.77
replace with random span mask	61.44
replace with sum of log-probabilities	60.84
replace with Cross Entropy Loss	62.71
replace with Margin Ranking Loss	61.18

Table 5: Ablation study of our BART_{base} model accuracy on the test set of the public dataset.

from 67.14% to 50.27%, which proves that the original text is indispensable for MCPredictor. At the same time, MCPredictor needs to re-generate additional scripts from original text based on event candidates, which is also inconvenient for applying the model to downstream tasks. 2). To compare with MCPredictor with comparable parameters and training data, we tried changing the event encoder in SCPredictor-s from Word2Vec to BERT_{base}, and the experimental result shows similar performance. In summary, our approach can achieve comparable performance to MCPredictor without external knowledge, which benefits transferring the model to downstream tasks.

Ablation Study

Table 5 shows our ablation analysis for our approach. We first investigated the impact of the event-centric pretraining stage (row1) and found that the model’s performance dropped by 1.86% without the stage. This is because the model captures correlations between events in the stage, enabling the model to reason about unseen events based on a partial event sequence. Second, we remove our task-specific contrastive fine-tuning stage (row2) or switch to the traditional classification method as in BART (row3). We can observe that the task-specific contrastive fine-tuning stage is crucial because it enables the model to distinguish between correct and incorrect event candidates. Moreover, it is better to fine-tune the model in our generative method instead of the discriminative classification method, because our method minimizes the inconsistency between task-specific contrastive fine-tuning, event-centric pretraining and BART pretraining stages. We then conducted experiments with the random span masking strategy to verify the effectiveness of the event-level blank infilling strategy. The only difference is that the masked span is no longer a complete event (the length and number of the masked spans keep unchanged). We find that the model’s performance is deficient in the event-centric pretraining stage (35.30%), and the event knowledge in the script is not well learned, which leads to its final performance drops (row4). Another important observation is that optimizing the mean of log-probabilities (Eq. 2 in our method) is significantly better than optimizing the sum of log-probabilities (row5). The latter will tend to select shorter event candidates since they have fewer log-probability items. Finally, we explore the performance of

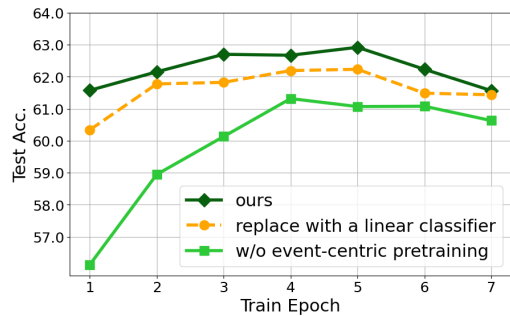


Figure 3: The training process comparison between ours and two ablations (*w/o event-centric pretraining* and *replaced with a linear classifier*). We test our model on the test set at each epoch.

the two variant losses. We find that our introduced *ComplementEntropyLoss*’s performance is slightly better than the traditional *Cross Entropy Loss* (row6) and much better than the *Margin Ranking Loss* (row7), which is widely used in contrastive learning.

Model Training Comparison

In this section, we compare the training process on the public dataset of Ours (BART_{base}) with two ablation methods (*w/o event-centric pretraining* and *replaced with a linear classifier*). From Figure 3, we can see that our method achieves higher results than the two ablation methods at every epoch. Compared with *w/o event-centric pretraining*, our method leads by a large margin in the first epoch because our method models correlations between events and learns event knowledge. Note that the model has not seen the data in the test set during the event-centric pretraining stage, so the event knowledge is generalizable. Compared with *replaced with a linear classifier*, our model leads by a more considerable margin at the initial epoch than in the later epochs, due to a more negligible difference between our task-specific contrastive fine-tuning and event-centric pretraining and pretraining objectives.

Case Study

We conduct a case study in the test dataset to show some properties of our model and illustrate how our model predicts the correct answer. First, we notice that the model can easily predict that the event candidate’s first token is “city” because “city” is the script’s protagonist. And the event candidate’s protagonist should be consistent with the script’s protagonist. In addition, the model can also perceive that “.” must be followed by $\langle /s \rangle$. Both of these show that the model can grasp some simple rules. Second, our model generates the score of each token in the event candidate in an autoregressive way. Our model is also interpretable to a certain extent. For example, the scores of “.” in E1, E2, and E3 are all high, mainly because the model thinks the event is not over after generating the front part of the event. However, for E4 and E5, the model focuses more on the plausibility of verbs and objects. Finally, this case shows how our model

Script: <s> surplus gives city . city presents budget . city devote some spending . city allocate money . city following formula . city discussed spending . city described pot . city engaged . <mask> . </s>

Event Candidates :

						sum	mean
E1	officials	named	city	.	</s>	54.2	10.8
	15.4	18.6	1.5	18.7	0		
E2	city	category	ize	links	.	51	8.5
	0	12.7	9.3	14.3	14.7	0	
E3	city	spoken	.	</s>		30	7.5
	0	15.2	14.8	0			
E4	city	thanking	officials	ceremony	.	40.9	7.2
	0	19	8	13.4	2.5	0	
E5	city	channel	ing	program	.	34.1	5.7
	0	11.4	4.5	13	5.2	0	

Figure 4: Case Study. The script corresponds to the input of our model in the contrastive stage. Below the token of the event candidate is its negative log generation probability. The higher the value, the harder it is to generate the token. **E5** is the correct event candidate.

avoids favoring shorter event candidates. In this case, E3 has the smallest sum of negative log generation probabilities, while E5 has the smallest mean of negative log generation probabilities. Compared with the ablation method (*replace with sum of log-probabilities*), our model can combine the length of events to select a more reasonable event candidate.

Related Work

Script event prediction was started by Chambers and Jurafsky (2008), who proposed the *narrative cloze* task: remove an event from a narrative event chain and adopt the model to predict the missing event. Chambers and Jurafsky (2008) define an event as a verb and its dependency, and predict the missing event by calculating the frequency of co-occurrence between events. Afterward, Granroth-Wilding and Clark (2016) expand the definition of an event as a verb and its three arguments (subject, object, indirect object) and propose the widely used *multiple choice narrative cloze* (MCNC) task: select the subsequent event from event candidates according to the historical narrative event chain, which can also be denoted as script. Compared with the narrative cloze task, MCNC is more suitable for evaluating the effectiveness of script event prediction models. Subsequent works are almost carried out based on MCNC.

Early work (Granroth-Wilding and Clark 2016) obtains the event representations through Word2Vec, then aggregates the pairwise similarities between an event candidate and script events to infer the probability that the event candidate is the subsequent event of the script. However, this method ignores the narrative order of script events, and subsequent works (Wang, Zhang, and Chang 2017; Lv et al. 2019) capture the narrative order through LSTM. Latter methods are more concerned with understanding event relations, which can be divided into two lines. The first line treats the verb co-occurrence frequency as an edge. Using this approach, Li, Ding, and Liu (2018) constructs a narra-

tive event graph, and adopts a scaled graph neural network to obtain better event representations. Following them, Ding et al. (2019) enriches event representations with external intention and sentiment knowledge. Gao et al. (2022) introduce a simultaneous weakly supervised contrastive learning and clustering method for event representation learning. Du et al. (2022) introduces the pre-trained model BERT (Devlin et al. 2019) and replaces the middle layer of the model with a graph neural network to embed the event graph information. The second line introduces event relations to help the model select the subsequent event. Lee, Pacheco, and Goldwasser (2020) mine the discourse relations from the original text by template matching method, then construct a narrative event graph, and finally transform the script event prediction task into a link prediction task on the graph. Bai et al. (2021) enriches the event representations with original text. Lv, Zhu, and Hu (2020); Zhou et al. (2021) introduce eventuality knowledge graph ASER to enhance pre-trained model RoBERTa (Liu et al. 2019). However, the second line methods introducing the discourse relation have narrow discourse relation types and cannot capture various implicit correlations. Meanwhile, to avoid the sparsity of the event graph, the first line methods model the co-occurrence frequency of verbs, ignoring other event arguments. Our method can effectively model correlations between events through the event mask strategy.

Recent event pretraining studies (Zhou et al. 2022a,b) conduct event-centric pretraining on BOOKCORPUS. Zhou et al. (2022b) and our method both adopt an event-level mask strategy. The difference is that our method masks more events, and we remove the other two pretraining objectives to focus on modeling event correlation. At the same time, our method uses the scripts in the training set for pretraining, and the learned event knowledge is closer to the MCNC task, making our model more powerful and efficient. Besides, the current methods introducing pretraining models use traditional classification fine-tuning methods when fine-tuning the models, resulting in inconsistency between pretraining and fine-tuning. Our method maximizes the consistency between pretraining and fine-tuning, enabling the model to better utilize the knowledge in the pre-trained model.

Conclusion

In this paper, we propose a novel generative approach, consisting of two stages, for the script event prediction task. For the event-centric pretraining stage, we introduce a novel event-level blank infilling strategy as the learning objective to inject event-level knowledge into the pretrained language model. In the task-specific contrastive fine-tuning stage, we design a likelihood-based contrastive loss to force the model to learn to distinguish between correct and wrong event candidates. Experimental results on the multi-choice narrative cloze (MCNC) task demonstrate that our approach achieves better results than other state-of-the-art baselines that require heavy pretraining or external knowledge. We believe our proposed approach can be a powerful ranking component for downstream tasks such as dialogue generation and story generation.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (62006062, 62176076, 62106114, 61906185), the Shenzhen Foundational Research Funding (JCYJ20210324115614039, JCYJ20200109113441941), Key Technologies Research and Development Program of Shenzhen JSGG20210802154400001, the Major Key Project of PCL2021A06, Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies 2022B1212010005k.

References

- Bai, L.; Guan, S.; Guo, J.; Li, Z.; Jin, X.; and Cheng, X. 2021. Integrating Deep Event-Level and Script-Level Information for Script Event Prediction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 9869–9878. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Chambers, N.; and Jurafsky, D. 2008. Unsupervised Learning of Narrative Event Chains. In *Proceedings of ACL-08: HLT*, 789–797. Columbus, Ohio: Association for Computational Linguistics.
- Chang, K.-W.; Yih, W.-t.; Yang, B.; and Meek, C. 2014. Typed Tensor Decomposition of Knowledge Bases for Relation Extraction. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1568–1579. Doha, Qatar: Association for Computational Linguistics.
- Chaturvedi, S.; Peng, H.; and Roth, D. 2017. Story Comprehension for Predicting What Happens Next. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1603–1614. Copenhagen, Denmark: Association for Computational Linguistics.
- Chen, H.; Wang, P.; Liu, C.; Chang, S.; Pan, J.; Chen, Y.; Wei, W.; and Juan, D. 2019. Complement Objective Training. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Danescu-Niculescu-Mizil, C.; and Lee, L. 2011. Chameleons in Imagined Conversations: A New Approach to Understanding Coordination of Linguistic Style in Dialogs. In *Proceedings of the 2nd Workshop on Cognitive Modeling and Computational Linguistics*, 76–87. Portland, Oregon, USA: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Ding, X.; Liao, K.; Liu, T.; Li, Z.; and Duan, J. 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)*, 4894–4903. Hong Kong, China: Association for Computational Linguistics.
- Du, L.; Ding, X.; Zhang, Y.; Liu, T.; and Qin, B. 2022. A Graph Enhanced BERT Model for Event Prediction. In *Findings of the Association for Computational Linguistics: ACL 2022*, 2628–2638. Dublin, Ireland: Association for Computational Linguistics.
- Gao, J.; Wang, W.; Yu, C.; Zhao, H.; Ng, W.; and Xu, R. 2022. Improving Event Representation via Simultaneous Weakly Supervised Contrastive Learning and Clustering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3036–3049.
- Graff, D.; Kong, J.; Chen, K.; and Maeda, K. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1): 34.
- Granroth-Wilding, M.; and Clark, S. 2016. What Happens Next? Event Prediction Using a Compositional Neural Network Model. In Schuurmans, D.; and Wellman, M. P., eds., *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, 2727–2733. AAAI Press.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lee, I.-T.; Pacheco, M. L.; and Goldwasser, D. 2020. Weakly-Supervised Modeling of Contextualized Event Embedding for Discourse Relations. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 4962–4972. Online: Association for Computational Linguistics.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7871–7880. Online: Association for Computational Linguistics.
- Li, Z.; Ding, X.; and Liu, T. 2018. Constructing Narrative Event Evolutionary Graph for Script Event Prediction. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 4201–4207. ijcai.org.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lv, S.; Qian, W.; Huang, L.; Han, J.; and Hu, S. 2019. SAM-Net: Integrating Event-Level and Chain-Level Attentions to Predict What Happens Next. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 6802–6809.
- Lv, S.; Zhu, F.; and Hu, S. 2020. Integrating External Event Knowledge for Script Learning. In *Proceedings of the 28th International Conference on Computational Linguistics*, 306–315. Barcelona, Spain (Online): International Committee on Computational Linguistics.

- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Sap, M.; Le Bras, R.; Allaway, E.; Bhagavatula, C.; Lourie, N.; Rashkin, H.; Roof, B.; Smith, N. A.; and Choi, Y. 2019. ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 3027–3035.
- Schank, R. C.; and Abelson, R. P. 2013. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Psychology Press.
- Wang, Z.; Zhang, Y.; and Chang, C.-Y. 2017. Integrating Order Information and Event Relation for Script Event Prediction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 57–67. Copenhagen, Denmark: Association for Computational Linguistics.
- Zhang, H.; Liu, X.; Pan, H.; Song, Y.; and Leung, C. W. 2020. ASER: A Large-scale Eventuality Knowledge Graph. In Huang, Y.; King, I.; Liu, T.; and van Steen, M., eds., *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, 201–211. ACM / IW3C2.
- Zhou, Y.; Geng, X.; Shen, T.; Long, G.; and Jiang, D. 2022a. Eventbert: A pre-trained model for event correlation reasoning. In *Proceedings of the ACM Web Conference 2022*, 850–859.
- Zhou, Y.; Geng, X.; Shen, T.; Pei, J.; Zhang, W.; and Jiang, D. 2021. Modeling Event-Pair Relations in External Knowledge Graphs for Script Reasoning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 4586–4596. Online: Association for Computational Linguistics.
- Zhou, Y.; Shen, T.; Geng, X.; Long, G.; and Jiang, D. 2022b. ClarET: Pre-training a Correlation-Aware Context-To-Event Transformer for Event-Centric Generation and Classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2559–2575. Dublin, Ireland: Association for Computational Linguistics.
- Zhu, Y.; Kiros, R.; Zemel, R. S.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, 19–27. IEEE Computer Society.