

Logic Induced High-Order Reasoning Network for Event-Event Relation Extraction

Peixin Huang¹, Xiang Zhao², Minghao Hu³, Zhen Tan^{1*}, Weidong Xiao¹

¹National Key Laboratory of Information Systems Engineering, National University of Defense Technology, China

²Laboratory for Big Data and Decision, National University of Defense Technology, China

³Information Research Center of Military Science, China

{huangpeixin15, xiangzhao, tanzhen08a, wdxiao}@nudt.edu.cn, huminghao16@gmail.com

Abstract

To understand a document with multiple events, event-event relation extraction (ERE) emerges as a crucial task, aiming to discern how natural events temporally or structurally associate with each other. To achieve this goal, our work addresses the problems of temporal event relation extraction (TRE) and subevent relation extraction (SRE). The latest methods for such problems have commonly built document-level event graphs for global reasoning across sentences. However, the edges between events are usually derived from external tools heuristically, which are not always reliable and may introduce noise. Moreover, they are not capable of preserving logical constraints among event relations, e.g., coreference constraint, symmetry constraint and conjunction constraint. These constraints guarantee coherence between different relation types, enabling the generation of a unified event evolution graph. In this work, we propose a novel method named LogicERE, which performs high-order event relation reasoning through modeling logic constraints. Specifically, different from conventional event graphs, we design a logic constraint induced graph (LCG) without any external tools. LCG involves event nodes where the interactions among them can model the coreference constraint, and event pairs nodes where the interactions among them can retain the symmetry constraint and conjunction constraint. Then we perform high-order reasoning on LCG with relational graph transformer to obtain enhanced event and event pair embeddings. Finally, we further incorporate logic constraint information via a joint logic learning module. Extensive experiments demonstrate the effectiveness of the proposed method with state-of-the-art performance on benchmark datasets.

Extended version — <https://arxiv.org/abs/2412.14688>

Introduction

Interpreting news messages involves identifying how natural events temporally or structurally associate with each other from news documents, i.e., extracting event temporal relations and subevent relations. Through this process, one can induce event evolution graphs that arrange multiple-granularity events with temporal relations and subevent relations interacting among them. The event evolution graphs built through

event-event relation extraction (ERE) are important aids for future event forecasting (Chaturvedi, Peng, and Roth 2017) and hot event tracking (Zhuang, Fei, and Hu 2023b). As shown in Figure 1, ERE aims to induce such an event evolution graph, in which the event mention *storm* involves more fine-grained subevent mentions, i.e., *killed*, *died* and *anceled*. Some of those mentions follow temporal order, e.g., *died* happens BEFORE *anceled*. Generally, predicting the relations between diverse events within the same document, such that these predictions are coherent and consistent with the document, is a challenging task (Xiang and Wang 2019).

Recently, significant research efforts have been devoted to several ERE tasks, such as event temporal relation extraction (TRE) (Zhou et al. 2021; Wang, Li, and Xu 2022; Tan, Pergola, and He 2023) and subevent relation extraction (SRE) (Man et al. 2022; Hwang et al. 2022). Nonetheless, ERE is still challenging because most event relations lack explicit clue words such as *before* and *contain* in natural languages, especially when events scatter in a document. Accordingly, a few previous methods attempt to build a document-level event graph to assist in the cross-sentence inference, where the nodes are events, and the edges are designed with linguistic/discourse relations of event pairs (Zhuang, Hu, and Zhao 2023; Li and Geng 2024). Despite the success, these methods face two major issues. First, they are not capable of preserving logic constraints among relations, such as transitivity, during training time (Roth and Yih 2004). Second, the edges heuristically from external tools may introduce noise and cause exhaustive extraction (Phu and Nguyen 2021).

For the first issue, Wang et al. (2020) propose a constrained learning framework which enforces logic coherence amongst the predicted relation types through extra differentiable objectives. However, since the coherence is enforced in a soft manner, there is still room for improving coherent predictions. In this work, we show that it is possible to enforce coherence in a much stronger manner by feeding logic constraints into the event graphs. For the second issue, Chen et al. (2022) propose an event pair centered causality identification model which takes event pairs as nodes and relations of event relations as edges, refraining the external tools and enabling the causal transitivity reasoning. However, some useful prior relation constraints such as coreference are discarded. Moreover, we observe logical property information loss from the document to graph, which should be complied in global inference for

*Corresponding author.

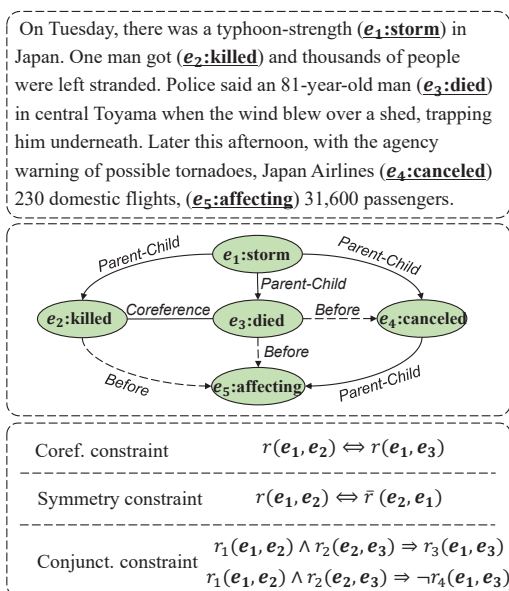


Figure 1: An example of an event evolution graph described in the document.

ERE. We summarize these properties implied in the temporal and subevent relations as three logical constraints: (1) **Coreference constraint**: Considering the example in Figure 1, given that e_2 :killed is BEFORE e_5 :affecting and e_3 :died is COREFERENCED to e_2 :killed, e_3 :died should be BEFORE e_5 :affecting. (2) **Symmetry constraint**: As shown in Figure 1, e_1 :storm is a PARENT of e_2 :killed, indicating e_2 :killed is a CHILD of e_1 :storm. (3) **Conjunction constraint**: From Figure 1, if e_1 :storm is a PARENT of e_3 :died and e_3 :died is BEFORE e_4 :canceled, the learning process should enforce e_1 :storm is a PARENT of e_4 :canceled by considering the conjunctive logic on both temporal and subevent relations. These logical constraints depict the mutuality among event relations of TRE and SRE, enabling the generation of a unified event evolution graph. While previous researches center on preserving these properties with post-learning inference or differentiable loss functions (Wang et al. 2020), there is no effective way to endow the graph model with these logical constraints for global reasoning.

In this paper, we consider the above constraints and propose a novel ERE model, logic induced high-order reasoning network (LogicERE). Our intuition is to feed the logic constraints into the event graphs for high-order event relation reasoning and prediction. Specifically, we first build a logic constraint induced graph (LCG) that models these logical constraints through interactions between events and among event pairs. Then, we encode the heterogeneous LCG with relational graph transformer. This enables the model to effectively reasoning over remote event pairs while maintaining inherent logic properties of event relations. Finally, we introduce a joint logic learning module and design two logic constraint learning objectives to further regularize the model towards consistency on logic constraints.

LogicERE models these high-order logical constraints in

two aspects. Firstly, our proposed LCG consolidates both event centered and event pair centered graphs, so that it can reason over not only coreference property among events, but also high-order symmetry and conjunction properties among event pairs. Specifically, LCG defines two types of nodes, i.e., event nodes and event pair nodes. Accordingly, there are three types of edges: (1) Event-event edge for prior event relations (e.g., coreference), which retains the coreference constraint and model the information flow among event nodes. (2) Event pair-event pair edge for two event pairs sharing at least one event, which keeps the symmetry and conjunction constraints, as well as captures the interactions among event pair nodes. (3) Event-event pair edge for an event pair and its corresponding two events, which models the information flow from event nodes to event pair nodes. As LCG preserves these inherent properties of event relations, we can get enhanced event and event pair embeddings through high-order reasoning over it. Secondly, inspired by the logic-driven framework of Li et al. (2019), we soften the logical properties through differentiable functions so as to incorporate them into multi-task learning objectives. The joint logic learning module enforces our model towards consistency with logic constraints across both TRE and SRE tasks. It is also a natural way to combine the supervision signals coming from two different tasks.

Related Work

Temporal Relation Extraction

Early studies usually employ statistical models combined with handcrafted features to extract temporal relations (Yoshikawa et al. 2009; Mani et al. 2006). These methods are of high computational complexity and are difficult to transfer to other types of relations.

Recently, with the rise of pre-trained language models, various new strategies have been applied to TRE task. To effectively model the long texts, some studies incorporate syntax information such as semantic trees or abstract meaning representation (AMR) to capture remote dependencies (Venkatchalam, Mutharaju, and Bhatia 2021; Zhang, Ning, and Huang 2022). Others construct global event graphs to enable the information flow among long-range events (Liu et al. 2021; Fei et al. 2022b). For example, Zhang, Ning, and Huang (2022) design a syntax-guided graph transformer to explore the temporal clues. Liu et al. (2021) build uncertainty-guided graphs to order the temporal events. To deepen the models’ understanding of events, some studies propose to induce external knowledge (Han, Zhou, and Peng 2020; Tan, Pergola, and He 2023). Han, Zhou, and Peng (2020) propose an end-to-end neural model that incorporates domain knowledge from the TimeBank corpus (Pustejovsky et al. 2003). Some multi-task strategies are also employed for this task (Huang et al. 2023; Knez and Zitnik 2024). These strategies can facilitate the models’ learning of complementary information from other tasks.

Subevent Relation Extraction

In earlier studies, machine learning algorithms are utilized to identify the internal structure of events (Fei et al. 2022a; Glavas et al. 2014). Afterwards, the introduction of deep

learning has led to new advances of this task. Zhou et al. (2020) adopt the multi-task learning strategy and utilize duration prediction as auxiliary task. Man et al. (2022) argue that some context sentences in documents can facilitate the recognition of subevent relations. Thus, they adopt the reinforcement learning algorithm to select informative sentences from documents to provide supplementary information for this task. Hwang et al. (2022) enforce constraint strategies into probabilistic box embedding to maintain the unique properties of subevent relations.

To sum up, existing work only focuses on either TRE or SRE, only a few studies seek to resolve both two tasks. Zhuang, Fei, and Hu (2023b); Zhuang, Hu, and Zhao (2023) adopt the dependency parser and hard pruning strategies to acquire syntactic dependency graphs that are task-aware. Zhuang, Fei, and Hu (2023a) propose to use knowledge from event ontologies as additional prompts to compensate for the lack of event-related knowledge. Li and Geng (2024) introduce the features of event argument and structure to obtain graph-enhanced event embeddings. However, the above works regard event relation extraction as a multi-class classification task, and do not guarantee any coherence between different relation types, such as symmetry and transitivity. Different from these works, our LogicERE guarantee coherence through a high-order reasoning graph which is embedded with three essential logical constraints, as well as joint logic learning.

Model

This paper focuses on the task of event-event relation extraction (ERE). The input document \mathcal{D} is represented as a sequence of n tokens $\mathcal{D} = [x_1, x_2, \dots, x_n]$. Each document contains a set of annotated event triggers (the most representative tokens for each event) $\mathcal{E}_{\mathcal{D}} = \{e_1, e_2, \dots, e_k\}$. The goal of ERE is to extract the multi-faceted event relations from the document. Particularly, we focus on two types of relations, i.e., Temporal and Subevent, corresponding to the label sets \mathcal{R}_{Temp} which contains BEFORE, AFTER, EQUAL, VAGUE, and \mathcal{R}_{Sub} which contains PARENT-CHILD, CHILD-PARENT, COREF, NOREL respectively. Note that each event pair is annotated with one relation type from either \mathcal{R}_{Temp} or \mathcal{R}_{Sub} , as the labels within two sets are mutually exclusive.

There are four major parts in our LogicERE model: (1) Sequence Encoder, which encodes event context in input document, (2) Logic Constraint Induced Graph, which build a event graph preserving logic constraints, (3) High-Order Reasoning Network on LCG, which performs high-order reasoning with relational graph transformer to obtain enhanced event and event pair embeddings, and (4) Joint Logic Learning, which further incorporates logic properties through well-designed learning objectives.

Sequence Encoder

To obtain the event representations and contextualized embeddings of the input document $\mathcal{D} = [x_t]_{t=1}^n$ (can be of any length n), we leverage pre-trained RoBERTa (Liu et al. 2019) as a base encoder. We add special tokens "[CLS]"

and "[SEP]" at the start and end of \mathcal{D} , and insert "<t>" and "</t>" at the start and end of all the events to mark event positions (Chen et al. 2022). Thus, we have:

$$h_1, h_2, \dots, h_{n'} = \text{Encoder}([x_1, x_2, \dots, x_{n'}]) \quad (1)$$

where $h_i \in \mathbb{R}^d$ is the embedding of token x_i . We employ the embeddings of token "[CLS]" and "<t>" to represent the document and the events respectively. If the document's length exceeds the limits of RoBERTa, we adopt the dynamic window mechanism to segment \mathcal{D} into several overlapping spans with specific step size, and input them into the encoder separately. Then, we average all the embeddings of "[CLS]" and "<t>" of different spans to obtain the document embedding $h_{[\text{CLS}]} \in \mathbb{R}^d$ or each event embedding $h_{e_i} \in \mathbb{R}^d$, respectively.

Logic Constraint Induced Graph

Considering the logic constraints of event relation is based on the motivation that such logic properties comprehensively define the varied interactions among those events and relations. In this section, we construct a logic constraint induced graph (LCG) which can preserve these logic properties through unifying both event centered and event pair centered graphs. Specifically, given all the events of document \mathcal{D} , LCG is formulated as $\mathcal{G} = \{\mathcal{V}, \mathcal{C}\}$, where \mathcal{V} represents the nodes and \mathcal{C} represents the edges in the graph. We highlight the following differences of \mathcal{G} from previous event graphs and event pair graphs.

First, there are two types of nodes in \mathcal{V} , i.e., the event nodes \mathcal{V}_e and the event pair nodes \mathcal{V}_{ep} . Each node in \mathcal{V}_{ep} refers to a different pair of events from \mathcal{D} . Instead of merely using events or event pairs as nodes, LCG preserves both of them, enabling high-order interactions through edges.

Second, for edges \mathcal{C} , instead of using all edges between any two nodes, we design three types of edges following the logic constraints: (1) Event-event edges \mathcal{C}_{ee} for two events that are co-referenced, which is motivated by the **coreference constraint** in Introduction. These edges are optional. \mathcal{C}_{ee} contributes to event relation reasoning as co-referenced events are expected to share the same relations with other events. Meanwhile, no additional relations exist between co-referenced events. (2) Event pair-event pair edges \mathcal{C}_{pp} for two event pairs that share at least one event, which is motivated by the **symmetry constraint** and **conjunction constraint** in Introduction. Particularly, for the TRE task, symmetry constraint exists in a pair of reciprocal relations BEFORE and AFTER, as well as two reflexive ones EQUAL and VAGUE. Similarly, the SRE task includes reciprocal relations PARENT-CHILD and CHILD-PARENT as well as reflexive ones COREF and NOREL. The conjunction constraint enables the relation transitivity in a single task, and unifies the ordered nature of TRE and the topological nature of SRE (Wang et al. 2020). (3) Event-event pair edges \mathcal{C}_{ep} for an event pair and its corresponding events. We design \mathcal{C}_{ep} to bridge the information flow between events and event pairs.

High-Order Reasoning Network on LCG

We perform high-order reasoning on LCG, which takes the relation heterogeneity into account and captures diversified

high-order interactions within events and event pairs.

Initial Node Embeddings. For global inference, we firstly initialize node embeddings. Formally, for the event node $e_i \in \mathcal{V}_e$, we take the contextualized event embeddings from the sequence encoder for initialization:

$$v_{e_i}^{(0)} = h_{e_i} \mathbf{W}_n \quad (2)$$

where 0 indicates the initial state and $\mathbf{W}_n \in \mathbb{R}^{d \times 2d}$ is a learnable weight matrix.

For the event pair node $e_{i,j} \in \mathcal{V}_{ep}$, we concatenate two corresponding event embeddings:

$$v_{e_{i,j}}^{(0)} = [h_{e_i} || h_{e_j}] \quad (3)$$

Node Embedding Update. Then, we adopt relational graph transformer (Bi et al. 2024) to enhance the node features with the relational information from neighbor nodes. Each layer l is similar to the transformer architecture. It takes a set of node embeddings $\mathbf{V}^{(l)} \in \mathbb{R}^{N \times d_{in}}$ as input, and outputs a new set of node embeddings $\mathbf{V}^{(l+1)} \in \mathbb{R}^{N \times d_{out}}$, where $N = |\mathcal{V}_e| + |\mathcal{V}_{ep}|$ is the number of nodes in LCG, d_{in} and d_{out} are the dimensions of input and output embeddings.

In each layer, to integrate information from each neighbor, we adopt a shared self-attention mechanism (Vaswani et al. 2017) to calculate the attention score:

$$\alpha_{ij} = \text{softmax}(co_{ij}) \quad (4)$$

$$co_{ij} = \frac{(v_i \mathbf{W}_q)(v_j \mathbf{W}_k)^T}{\sqrt{d_k}} \quad (5)$$

where N_i is the first order neighbor set of node i , co_{ij} measures the importance of neighbor j to i , $\mathbf{W}_q, \mathbf{W}_k \in \mathbb{R}^{d_{in} \times d_k}$ are learnable matrices, d_k is a scaling factor to assign lower attention weights to uninformative nodes.

Then we aggregate relational knowledge from the neighborhood information with weighted linear combination of the embeddings:

$$v_i^{(l+1)} = \sum_{j \in N_i} \alpha_{ij}^{(l)} (v_j^{(l)} \mathbf{W}_v^{(l)}) \quad (6)$$

where $\mathbf{W}_v^{(l)} \in \mathbb{R}^{d_{in} \times d_k}$ is a learnable matrix. We also adopt multi-head attention to attend to information from multiple attention heads. Thus, the output of the l -th layer for node i is:

$$v_i^{(l+1)} = \left(\left\| \sum_{j \in N_i} \alpha_{ij}^{(l)} (v_j^{(l)} \mathbf{W}_v^{(l)}) \right\| \right) \mathbf{W}_o^{(l)} \quad (7)$$

where C is the number of attention head and $\mathbf{W}_o^{(l)} \in \mathbb{R}^{C d_k \times d_{out}}$ is a learnable matrix.

Measure Edge Heterogeneity. It is intuitive that three types of edges in LCG contributes differently to ERE. Thus we propose to measure the edge heterogeneity and incorporate the edge features into node embeddings. Specifically, for each edge type in LCG, we learn a scalar:

$$\beta_t = r_t \mathbf{W}_r \quad (8)$$

where $1 \leq t \leq T$, T is the number of edge types, $r_t \in \mathbb{R}^{1 \times d}$ denotes the edge features specific to the edge type, $\mathbf{W}_r \in$

$\mathbb{R}^{d \times 1}$ is a learnable matrix. r_t will be randomly initialized. Then we incorporate β_t as the attention bias into the attention score to adjust the interaction strength between two adjacent nodes:

$$\tilde{\alpha}_{ij} = \text{softmax}(\beta_t + co_{ij}) \quad (9)$$

As the result, the final updated node embeddings considering the edge heterogeneity is:

$$\tilde{v}_i^{(l+1)} = \left(\left\| \sum_{c=1}^C \sum_{j \in N_i} \tilde{\alpha}_{ij}^{(l)} (v_j^{(l)} \mathbf{W}_v^{(l)}) \right\| \right) \mathbf{W}_o^{(l)} \quad (10)$$

By stacking multiple layers, the reasoning network could reach high-order interaction and maintain logic properties.

Learning and Classification. To predict whether there is the temporal or subevent relation between events e_i and e_j , we concatenate the embeddings of "[CLS]", e_i , e_j and the corresponding event pair as the logic enhanced representation. Thus, the probability distribution of the relation can be obtained through linear classification:

$$p_{e_{i,j}} = \text{softmax}([h_{[CLS]} || \tilde{v}_i || \tilde{v}_j || \tilde{v}_{e_{i,j}}] \mathbf{W}_p) \quad (11)$$

where $||$ denotes concatenation and \mathbf{W}_p is a learnable matrix.

For training, we adopt cross-entropy as the loss function:

$$\mathcal{L}_1 = - \sum_{e_i, e_j \in \mathcal{E}_D} (1 - y_{e_{i,j}}) \log(1 - p_{e_{i,j}}) + y_{e_{i,j}} \log(p_{e_{i,j}}) \quad (12)$$

where $y_{e_{i,j}}$ denotes the golden label.

Joint Logic Learning

Inspired by the logic-driven framework for consistency of Li et al. (2019), we further design two learning objectives by directly transforming the logical constraints into differentiable loss functions.

Symmetry Constraint. Symmetry constraints indicate the event pair with flipping orders will have the reversed relation, the logical formula can be written as:

$$e_i, e_j \in \mathcal{E}_D, r \in \mathcal{R}_{sym} \quad r(e_i, e_j) \leftrightarrow \bar{r}(e_j, e_i) \quad (13)$$

where \mathcal{R}_{sym} indicates the set of relations enforcing the symmetry constraint. We use the product t-norm and transformation to the negative log space and obtain the symmetry loss:

$$\mathcal{L}_{sym} = \sum_{e_i, e_j \in \mathcal{E}_D} |\log(p_{e_{i,j}}) - \log(\bar{p}_{e_{j,i}})| \quad (14)$$

Conjunction Constraint. Conjunctive constraint are applicable to any three related events e_i , e_j and e_k . It contributes to the joint learning of TRE and SRE. The conjunction constraint enforces the following logical formulas:

$$\bigwedge_{\substack{e_i, e_j, e_k \in \mathcal{E}_D \\ r_1, r_2 \in \mathcal{R}, r_3 \in D e(r_1, r_2)}} r_1(e_i, e_j) \wedge r_2(e_j, e_k) \rightarrow r_3(e_i, e_k) \quad (15)$$

$$\bigwedge_{\substack{e_i, e_j, e_k \in \mathcal{E}_D \\ r_1, r_2 \in \mathcal{R}, r_4 \notin D e(r_1, r_2)}} r_1(e_i, e_j) \wedge r_2(e_j, e_k) \rightarrow \neg r_4(e_i, e_k) \quad (16)$$

where $De(r_1, r_2)$ is a set composed of all relations from \mathcal{R} that do not conflict with r_1 and r_2 .

Similarly, the loss function specific to conjunction constraint is:

$$\mathcal{L}_{conj} = \sum_{e_i, e_j, e_k \in \mathcal{E}_D} |\mathcal{L}_{c_1}| + \sum_{e_i, e_j, e_k \in \mathcal{E}_D} |\mathcal{L}_{c_2}| \quad (17)$$

$$\mathcal{L}_{c_1} = \log(p_{e_i, j}) + \log(p_{e_j, k}) - \log(p_{e_i, k}) \quad (18)$$

$$\mathcal{L}_{c_2} = \log(p_{e_i, j}) + \log(p_{e_j, k}) - \log(1 - p_{e_i, k}) \quad (19)$$

The final loss function combines the above logic learning and event relation learning objectives, where γ are non-negative coefficients to control the influence of each loss term:

$$\mathcal{L} = \mathcal{L}_1 + \gamma_{sym} \mathcal{L}_{sym} + \gamma_{conj} \mathcal{L}_{conj} \quad (20)$$

Experiments

Datasets and Metrics

We evaluate logicERE on four widely used datasets. MATRES (Ning, Wu, and Roth 2018) and TCR (Ning et al. 2018) are used to test the performance of TRE. HiEve (Glavas et al. 2014) is used for SRE. MAVEN-ERE (Wang et al. 2022) is used to test the joint learning performance. **MATRES** is a new dataset and mainly annotates four temporal relationships, i.e., BEFORE, AFTER, EQUAL and VAGUE. **TCR** is a small-scale dataset which has the same annotation scheme as MATRES. We only apply it to the testing phase. **HiEve** is a news corpus and annotates four subevent relationships, i.e., PARENT-CHILD, CHILD-PARENT, COREF and NOREL. **MAVEN-ERE** is a unified large-scale dataset that annotates data for event coreference, temporal, causal, and subevent relations. We only experiment on the event temporal and subevent relations. For TRE, it defines six relationships. To be consistent with our framework, we only consider type BEFORE and SIMULTANEOUS, and we manually annotate reflexive relationships AFTER and VAGUE, respectively. For SRE, it defines one relationship SUBEVENT and we manually annotate corresponding reflexive relationships SUPEREVENT. Note that HiEve and MAVEN-ERE provide ground-truth event coreference annotations, but MATRES does not. We follow Chen et al. (2023) and perform pre-training on MAVEN-ERE, and then use the pre-trained model to extract coreference data for MATRES. After the preprocessing steps, we add event-event edges \mathcal{C}_{ee} to MATRES. For compatible comparison, we utilize the same data splits as in prior work for the considered datasets. We briefly summarize the data statistics for the above datasets in Table 1.

We adopt the standard micro-averaged Precision (P), Recall (R) and F1-scores (F1) as evaluation metrics. All the results are the average of five trials of different random seeds in each experiment.

Parameter Settings

Our implementation uses HuggingFace Transformers (Wolf et al. 2020) and PyTorch (Paszke et al. 2019). We employ RoBERTa-base (Liu et al. 2019) as the document encoder. As for the input of the encoder, we set the dynamic window size to 256, and divide documents into several overlapping

Dataset		Train	Dev	Test
MATRES	Document	260	21	20
	Event pairs	10,888	1,852	840
TCR	Document	-	-	25
	Event pairs	-	-	2,646
HiEve	Document	80	-	20
	Event pairs	35,001	-	7,093
MAVEN-ERE (TRE)	Document	2,913	710	857
	Event pairs	792,445	188,928	234,844
MAVEN-ERE (SRE)	Document	2,913	710	857
	Event pairs	9,193	2,826	3,822

Table 1: Data statistics for dataset MATRES, TCR, HiEve and MAVEN-ERE (TRE/SRE).

windows with a step size 32. We use AdamW (Loshchilov and Hutter 2019) optimizer and learning rate is set to $2e-5$. We adopt layer normalization (Ba, Kiros, and Hinton 2016) and dropout (Srivastava et al. 2014) between the high-order reasoning network layers. We perform early stopping and tune the hyper-parameters by grid search on the development set: heads $C \in \{1, 2, 4, 8\}$, dropout rate $\in \{0.1, 0.2, 0.3\}$ and loss coefficients $\gamma_{sym}, \gamma_{conj} \in \{0.1, 0.2, 0.4, 0.6\}$.

Baselines

We adopt the following state-of-the-art models for the MATRES, TCR, HiEve and MAVEN-ERE datasets respectively.

For the MATRES dataset, we consider the following baselines. (1) **TEMPROB+ILP** (Ning, Subramanian, and Roth 2019) incorporates temporal commonsense knowledge and integer linear programming; (2) **Hierarchical** (Adhikari et al. 2019) uses RoBERTa to encode different chunks of the document, and sets an additional BiLSTM model to aggregate representations; (3) **Joint Constrain** (Wang et al. 2020) enforces consistency with joint constrained learning objectives; (4) **Self-Training** (Ballesteros et al. 2020) relies on multitask and self-training techniques; (5) **Vanilla Classifier** (Wen and Ji 2021) is a common event relation classifier based on the RoBERTa; (6) **Relative Time** (Wen and Ji 2021) is a stack propagation framework employing relative time prediction as hints; (7) **Probabilistic Box** (Hwang et al. 2022) uses probabilistic boxes as implicit constraints to improve the consistency; (8) **TGAGCN** (Zhuang, Hu, and Zhao 2023) is based on event-specific syntactic dependency graph; (9) **OntoEnhance** (Zhuang, Fei, and Hu 2023a) fuses semantic information from event ontologies to enhance event representation; (10) **SDLG** (Zhuang, Fei, and Hu 2023b) builds syntax-based dynamic latent graph for event relation reasoning.

For TCR, the following baselines are included in our comparison. (1) **LSTM+knowledge** (Ning, Subramanian, and Roth 2019) is a variant of TEMPROB+ILP that incorporates temporal commonsense knowledge in LSTM; (2) **HGRU** (Tan, Pergola, and He 2021) maps event embeddings to hyperbolic space; (3) **Poincaré Embeddings** (Tan, Pergola, and He 2021) learns rich event representations in hyperbolic spaces; (4) We also compare with **TEMPROB+ILP**, **Vanilla**

Model	P	R	F1
TEMPROB+ILP (2019)	71.3	82.1	76.3
Hierarchical (2019)	74.2	83.1	78.4
Joint Constrain (2020)	73.4	85.0	78.3
Self-Training (2020)	-	-	81.6
Vanilla Classifier (2021)	78.1	82.5	80.2
Relative Time (2021)	78.4	85.2	81.7
Probabilistic Box (2022)	-	-	77.1
TGAGCN (2023)	81.1	83.0	82.0
OntoEnhance (2023)	79.0	86.5	82.6
SDLG (2023)	82.0	84.2	83.1
LogicERE (ours)	82.9	84.5	83.7

Table 2: Model performance on test data of MATRES for temporal relation extraction.

Classifier, TGAGCN, OntoEnhance and SDLG.

For the HiEve dataset, we choose the following baselines. (1) **StructLR** (Glavas et al. 2014) is a supervised classifier combining event, bag-of-words, location, and syntactic features; (2) **TACOLM** (Zhou et al. 2020) uses duration prediction to assist SRE; (3) Similarly, we also compare with **Joint Constrain**, **Vanilla Classifier**, **Hierarchical**, **TGAGCN**, **OntoEnhance** and **SDLG**.

For MAVEN-ERE, we select RoBERTa (Liu et al. 2019) as the main baseline, which adopts RoBERTa as the document encoder and obtains event embeddings for pair-wise classification. According to whether multiple event relations are trained simultaneously, we consider two settings, i.e., (1) **RoBERTa_{split}** and (2) **RoBERTa_{joint}** for comparisons. We also compare with (3) **GraphERE_{split}** and (4) **GraphERE_{joint}** (Li and Geng 2024), which adopt graph-enhanced event embeddings for pair-wise classification. Besides, for each task, we explore related baselines in our experiments: (5) **DocTime** (Mathur et al. 2022) constructs temporal dependency graph for TRE; (6) **MultiFeatures** (Al-dawsari and Finlayson 2019) employs discourse and narrative features for SRE.

Comparison

Table 2, 3, 4 and 5 show the performance of the models on four datasets. Here, the performance for the models in previous work is inherited from the original papers.

TRE Evaluation. From the results in Table 2 and 3 for temporal relation extraction, we can draw the following observations: 1) LogicERE significantly outperforms previous models (with $p < 0.05$) on both datasets, which demonstrates its effectiveness for TRE. 2) Compared with other methods considering logic coherence (i.e., Joint Constrain and Probabilistic Box), LogicERE gains at least 5.4% F1 scores improvements on MATRES. This indicates that LogicERE is more effective in incorporating logic constraints in graph-enhanced event embeddings, which preserves coherence of temporal relations and enriches the features in event embeddings. 3) LogicERE surpasses the graph-based model (i.e., TGAGCN, OntoEnhance and SDLG) by at least 0.6% and 1.1% F1 scores improvements on MATRES and TCR

Model	P	R	F1
LSTM+knowledge (2019)	79.3	76.9	78.1
TEMPROB+ILP (2019)	-	-	78.6
Vanilla Classifier (2021)	89.2	76.7	82.5
HGRU (2021)	88.3	79.0	83.5
Poincaré Embeddings (2021)	85.0	86.0	85.5
TGAGCN (2023)	89.2	84.3	86.7
OntoEnhance (2023)	89.6	84.3	86.8
SDLG (2023)	88.3	87.0	87.6
LogicERE (ours)	90.8	86.7	88.7

Table 3: Model performance on test set of TCR for temporal relation extraction.

Model	F1		
	PC	CP	Avg.
StructLR (2014)	52.2	63.4	57.7
Hierarchical (2019)	63.7	57.1	60.4
TACOLM (2020)	48.5	49.4	48.9
Joint Constrain (2020)	62.5	56.4	59.5
Vanilla Classifier (2021)	62.8	52.3	57.5
TGAGCN (2023)	65.2	56.0	60.6
OntoEnhance (2023)	66.7	57.3	62.0
SDLG (2023)	64.9	58.5	61.7
LogicERE (ours)	68.3	62.5	65.4

Table 4: Model performance on test set of HiEve for subevent relation extraction. We focus on the performance for PARENT-CHILD (PC), CHILD-PARENT (CP), and their micro-average (Avg.).

respectively. It shows the advantages of logic constraint induced graph for TRE. Importantly, LogicERE can accomplish state-of-the-art performance for TRE without any external knowledge. This is different from recent work that requires additional resources to secure good performance, such as ontology knowledge in OntoEnhance or dependency parsing in TGAGCN and SDLG.

SRE Evaluation. The results in Table 4 for subevent relation extraction exhibit similar observation: 1) LogicERE can effectively deal with SRE task, achieving a significant advantage over all methods (with $p < 0.05$). The reason may be that our model can benefit from high-order reasoning of subevent relations through logic induced reasoning network and can get logic-enhanced event embeddings for SRE. 2) Consistent with experimental results on MATRES and TCR, LogicERE still shows significant advantages over other logic-enhanced models (i.e., Joint Constrain) and graph-based models (i.e., TGAGCN, OntoEnhance and SDLG). This verifies the effectiveness of our model for SRE.

Joint Learning Evaluation. MAVEN-ERE contains annotations for both TRE and SRE. We experiment on it to evaluate the ability of LogicERE in jointly learning multiple event relations. From Table 5, we observe that: 1) The proposed LogicERE_{joint} surpasses all baselines in F1 scores for both event relations. Compared with the best existing

Model	TRE			SRE		
	P	R	F1	P	R	F1
DocTime (2022)	54.4	53.0	53.7	-	-	-
MultiFeatures (2019)	-	-	-	16.4	19.9	18.0
RoBERTa _{split} (2024)	53.1	53.3	53.2	27.3	20.9	23.7
RoBERTa _{joint} (2024)	50.7	56.7	53.5	19.9	24.2	21.8
GraphERE _{split} (2024)	57.1	50.0	53.3	26.1	20.8	23.1
GraphERE _{joint} (2024)	55.1	54.4	54.7	31.2	24.2	27.3
LogicERE _{split} (ours)	58.2	50.5	54.1	29.4	22.8	25.7
LogicERE _{joint} (ours)	56.3	53.7	55.0	32.8	28.2	30.3

Table 5: Model performance on test set of MAVEN-ERE for temporal relation extraction and subevent relation extraction, split/joint stands for separately or jointly training TRE and SRE tasks.

methods for each task, our model improves by 0.3% in TRE and 3.0% in SRE. 2) LogicERE has a considerable improvement in Precision. LogicERE_{joint} gains on average 1.4% Precision improvement on the two tasks (LogicERE_{split} for 1.6%). This indicates that we incorporate logic constraints into LCG which enriches the features in event embeddings through high-order reasoning, and improves the performance in Precision. 3) The models in joint setting outperform corresponding ones in split setting. Meanwhile, the improvement from LogicERE_{split} to LogicERE_{joint} are 0.9% and 4.6% on TRE and SRE respectively, indicating that learning from other types of relations boost the ERE performance. 4) SRE is a challenging task with almost all the models exhibiting less than 30% in F1 Score. However, LogicERE_{joint} has the most significant advantage of 3.0% in F1 Score over baselines. The possible reason is that the global consistency ensured by LCG and joint logic learning naturally makes up for the weak supervision signals for SRE.

In general, the experimental results show that LogicERE can effectively combine the supervision signals from two tasks, assisting in the comprehension of both temporal relation and subevent relation.

Ablation Study

We then conduct ablation study to elucidate the effectiveness of main components of our model. Likewise, we only present the results on MAVEN-ERE. In particular, we consider the following internal baselines. (1) **w/o edge heterogeneity** does not consider the edge heterogeneity, and thus the scalar β_t is removed from Eq. 9 in the relational knowledge aggregation process. (2) **w/o coreference** removes \mathcal{C}_{ee} from LCG and do not use ground-truth coreference annotations as training labels. (3) **w/o event-event pair edges** removes \mathcal{C}_{ep} from LCG. (4) **w/o symmetry objective** removes \mathcal{L}_{sym} from Eq. 20. (5) **w/o conjunction objective** removes \mathcal{L}_{conj} from Eq. 20. (6) **w/o joint logic learning** removes the joint logic learning objective and adopt \mathcal{L}_1 as the loss function.

Results are shown in Table 6. We can observe that: 1) Our full model significantly outperforms all internal baselines on MAVEN-ERE. Compared to LogicERE, w/o. edge heterogeneity drops 0.6% and 0.8% F1 scores for TRE and SRE respectively, validating the necessity of capturing the

Model	TRE	SRE
LogicERE _{joint} (full)	55.0	30.3
w/o edge heterogeneity	54.4 (-0.6)	29.5 (-0.8)
w/o coreference	54.7 (-0.3)	30.0 (-0.3)
w/o event-event pair edges	53.5 (-1.5)	28.8 (-1.5)
w/o symmetry objective	54.8 (-0.2)	30.0 (-0.3)
w/o conjunction objective	54.5 (-0.5)	29.7 (-0.6)
w/o joint logic learning	54.1 (-0.9)	28.9 (-1.4)

Table 6: Ablation results (F1 scores) on MAVEN-ERE.

semantic information of different types of edges in LCG. 2) w/o. event-event pair edges shows huge performance drop. This result demonstrates that the deep information interaction between the event and the event pair contributes to more informative event embeddings, impelling better event pair relation inference. 3) Experimental results of w/o symmetry objective, w/o conjunction objective and w/o joint logic learning demonstrate that, after removing either of the two, the F1 scores go down. It indicates that these two kinds of logic constraints both contribute to our model. Simultaneously using two kinds of knowledge further improves the overall performance.

Conclusion

We present a novel logic induced high-order reasoning network (LogicERE) to enhance the event relation reasoning with logic constraints. We first design a logic constraint induced graph (LCG) which contains interactions between events and among event pairs. Then we encode the heterogeneous LCG for high-order event relation reasoning while maintaining inherent logic properties of event relations. Finally, we further incorporate logic constraints with joint logic learning. Extensive experiments show that LogicERE can effectively utilize logic properties to enhance the event and event pair embeddings, and achieve state-of-the-art performance for both TRE and SRE. The joint learning evaluation reveals that LogicERE can effectively maintain global consistency of two types relations, assisting in the comprehension of both temporal and subevent relation.

Acknowledgments

We thank all the anonymous reviewers and meta reviewers for their valuable comments, as well as all of our team members for their support and assistance. This work was partially supported by National Key R&D Program of China (No. 2022YFB3103600), and NSFC (Nos. U23A20296, 72371245, 62476283, 62376284).

References

- Adhikari, A.; Ram, A.; Tang, R.; and Lin, J. 2019. DocBERT: BERT for Document Classification. *CoRR*, abs/1904.08398.
- Aldawsari, M.; and Finlayson, M. A. 2019. Detecting Subevents using Discourse and Narrative Features. In *Proceedings of the 57th ACL, Volume 1: Long Papers*, 4780–4790. Association for Computational Linguistics.
- Ba, L. J.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *CoRR*, abs/1607.06450.
- Ballesteros, M.; Anubhai, R.; Wang, S.; Pourdamghani, N.; Vyas, Y.; Ma, J.; Bhatia, P.; McKeown, K. R.; and Al-Onaizan, Y. 2020. Severing the Edge Between Before and After: Neural Architectures for Temporal Ordering of Events. In *Proceedings of EMNLP, Online*, 5412–5417. Association for Computational Linguistics.
- Bi, Z.; Cheng, S.; Chen, J.; Liang, X.; Xiong, F.; and Zhang, N. 2024. Relphormer: Relational Graph Transformer for Knowledge Graph Representations. *Neurocomputing*, 566: 127044.
- Chaturvedi, S.; Peng, H.; and Roth, D. 2017. Story Comprehension for Predicting What Happens Next. In *Proceedings of EMNLP*, 1603–1614. Association for Computational Linguistics.
- Chen, M.; Cao, Y.; Deng, K.; Li, M.; Wang, K.; Shao, J.; and Zhang, Y. 2022. ERGO: Event Relational Graph Transformer for Document-level Event Causality Identification. In *Proceedings of the 29th COLING*, 2118–2128. International Committee on Computational Linguistics.
- Chen, M.; Cao, Y.; Zhang, Y.; and Liu, Z. 2023. CHEER: Centrality-aware High-order Event Reasoning Network for Document-level Event Causality Identification. In *Proceedings of the 61st ACL, Volume 1: Long Papers*, 10804–10816. Association for Computational Linguistics.
- Fei, H.; Li, J.; Wu, S.; Li, C.; Ji, D.; and Li, F. 2022a. Global Inference with Explicit Syntactic and Discourse Structures for Dialogue-Level Relation Extraction. In *Proceedings of the Thirty-First IJCAI*, 4107–4113.
- Fei, H.; Wu, S.; Ren, Y.; and Zhang, M. 2022b. Matching Structure for Dual Learning. In *International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 6373–6391. PMLR.
- Glavas, G.; Snajder, J.; Moens, M.; and Kordjamshidi, P. 2014. HiEve: A Corpus for Extracting Event Hierarchies from News Stories. In *Proceedings of the Ninth LREC*, 3678–3683. European Language Resources Association (ELRA).
- Han, R.; Zhou, Y.; and Peng, N. 2020. Domain Knowledge Empowered Structured Neural Net for End-to-End Event Temporal Relation Extraction. In *Proceedings of EMNLP, Online*, 5717–5729. Association for Computational Linguistics.
- Huang, Q.; Hu, Y.; Zhu, S.; Feng, Y.; Liu, C.; and Zhao, D. 2023. More than Classification: A Unified Framework for Event Temporal Relation Extraction. In *Proceedings of the 61st ACL, Volume 1: Long Papers*, 9631–9646. Association for Computational Linguistics.
- Hwang, E.; Lee, J.; Yang, T.; Patel, D.; Zhang, D.; and McCallum, A. 2022. Event-Event Relation Extraction using Probabilistic Box Embedding. In *Proceedings of the 60th ACL, Volume 2: Short Papers*, 235–244. Association for Computational Linguistics.
- Knez, T.; and Zitnik, S. 2024. Multimodal learning for temporal relation extraction in clinical texts. *J. Am. Medical Informatics Assoc.*, 31(6): 1380–1387.
- Li, H.; and Geng, D. 2024. GraphERE: Jointly Multiple Event-Event Relation Extraction via Graph-Enhanced Event Embeddings. *CoRR*, abs/2403.12523.
- Li, T.; Gupta, V.; Mehta, M.; and Srikumar, V. 2019. A Logic-Driven Framework for Consistency of Neural Models. In *Proceedings of EMNLP-IJCNLP*, 3922–3933. Association for Computational Linguistics.
- Liu, J.; Xu, J.; Chen, Y.; and Zhang, Y. 2021. Discourse-Level Event Temporal Ordering with Uncertainty-Guided Graph Completion. In *Proceedings of the Thirtieth IJCAI, Virtual Event*, 3871–3877. 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. *CoRR*, abs/1907.11692.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations*. OpenReview.net.
- Man, H.; Ngo, N. T.; Van, L. N.; and Nguyen, T. H. 2022. Selecting Optimal Context Sentences for Event-Event Relation Extraction. In *Proceedings of Thirty-Sixth AAAI, Virtual Event*, 11058–11066. AAAI Press.
- Mani, I.; Verhagen, M.; Wellner, B.; Lee, C. M.; and Pustejovsky, J. 2006. Machine Learning of Temporal Relations. In *Proceedings of ACL*. The Association for Computer Linguistics.
- Mathur, P.; Morariu, V. I.; Kaynig-Fittkau, V.; Gu, J.; Deroncourt, F.; Tran, Q. H.; Nenkova, A.; Manocha, D.; and Jain, R. 2022. DocTime: A Document-level Temporal Dependency Graph Parser. In *Proceedings of NAACL*, 993–1009. Association for Computational Linguistics.
- Ning, Q.; Feng, Z.; Wu, H.; and Roth, D. 2018. Joint Reasoning for Temporal and Causal Relations. In *Proceedings of the 56th ACL, Volume 1: Long Papers*, 2278–2288. Association for Computational Linguistics.
- Ning, Q.; Subramanian, S.; and Roth, D. 2019. An Improved Neural Baseline for Temporal Relation Extraction. In *Proceedings of EMNLP-IJCNLP*, 6202–6208. Association for Computational Linguistics.
- Ning, Q.; Wu, H.; and Roth, D. 2018. A Multi-Axis Annotation Scheme for Event Temporal Relations. In *Proceedings*

- of the 56th ACL, Volume 1: Long Papers, 1318–1328. Association for Computational Linguistics.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E. Z.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems, NeurIPS 2019*, 8024–8035.
- Phu, M. T.; and Nguyen, T. H. 2021. Graph Convolutional Networks for Event Causality Identification with Rich Document-level Structures. In *Proceedings of NAACL-HLT, Online*, 3480–3490. Association for Computational Linguistics.
- Pustejovsky, J.; Hanks, P.; Saurí, R.; See, A.; and Lazo, M. 2003. The TimeBank corpus. *Proceedings of corpus linguistics*.
- Roth, D.; and Yih, W. 2004. A Linear Programming Formulation for Global Inference in Natural Language Tasks. In *Proceedings of the Eighth CoNLL*, 1–8. ACL.
- Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1): 1929–1958.
- Tan, X.; Pergola, G.; and He, Y. 2021. Extracting Event Temporal Relations via Hyperbolic Geometry. In *Proceedings of EMNLP, Virtual Event*, 8065–8077. Association for Computational Linguistics.
- Tan, X.; Pergola, G.; and He, Y. 2023. Event Temporal Relation Extraction with Bayesian Translational Model. In *Proceedings of the 17th EACL*, 1117–1130. Association for Computational Linguistics.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, 5998–6008.
- Venkatachalam, K.; Mutharaju, R.; and Bhatia, S. 2021. SERC: Syntactic and Semantic Sequence based Event Relation Classification. In *33rd IEEE ICTAI*, 1316–1321. IEEE.
- Wang, H.; Chen, M.; Zhang, H.; and Roth, D. 2020. Joint Constrained Learning for Event-Event Relation Extraction. In *Proceedings of EMNLP, Online*, 696–706. Association for Computational Linguistics.
- Wang, L.; Li, P.; and Xu, S. 2022. DCT-Centered Temporal Relation Extraction. In *Proceedings of the 29th COLING*, 2087–2097. International Committee on Computational Linguistics.
- Wang, X.; Chen, Y.; Ding, N.; Peng, H.; Wang, Z.; Lin, Y.; Han, X.; Hou, L.; Li, J.; Liu, Z.; Li, P.; and Zhou, J. 2022. MAVEN-ERE: A Unified Large-scale Dataset for Event Coreference, Temporal, Causal, and Subevent Relation Extraction. In *Proceedings of EMNLP*, 926–941. Association for Computational Linguistics.
- Wen, H.; and Ji, H. 2021. Utilizing Relative Event Time to Enhance Event-Event Temporal Relation Extraction. In *Proceedings of EMNLP, Virtual Event*, 10431–10437. Association for Computational Linguistics.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Scao, T. L.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. M. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of EMNLP - Demos, Online*, 38–45. Association for Computational Linguistics.
- Xiang, W.; and Wang, B. 2019. A Survey of Event Extraction From Text. *IEEE Access*, 7: 173111–173137.
- Yoshikawa, K.; Riedel, S.; Asahara, M.; and Matsumoto, Y. 2009. Jointly Identifying Temporal Relations with Markov Logic. In *Proceedings of ACL/IJCNLP*, 405–413. The Association for Computer Linguistics.
- Zhang, S.; Ning, Q.; and Huang, L. 2022. Extracting Temporal Event Relation with Syntax-guided Graph Transformer. In *Findings of NAACL*, 379–390. Association for Computational Linguistics.
- Zhou, B.; Ning, Q.; Khashabi, D.; and Roth, D. 2020. Temporal Common Sense Acquisition with Minimal Supervision. In *Proceedings of the 58th ACL, Online*, 7579–7589. Association for Computational Linguistics.
- Zhou, Y.; Yan, Y.; Han, R.; Caufield, J. H.; Chang, K.; Sun, Y.; Ping, P.; and Wang, W. 2021. Clinical Temporal Relation Extraction with Probabilistic Soft Logic Regularization and Global Inference. In *Proceedings of Thirty-Fifth AAAI, Virtual Event*, 14647–14655. AAAI Press.
- Zhuang, L.; Fei, H.; and Hu, P. 2023a. Knowledge-enhanced event relation extraction via event ontology prompt. *Inf. Fusion*, 100: 101919.
- Zhuang, L.; Fei, H.; and Hu, P. 2023b. Syntax-based dynamic latent graph for event relation extraction. *Inf. Process. Manag.*, 60(5): 103469.
- Zhuang, L.; Hu, P.; and Zhao, W. 2023. Event Relation Extraction Using Type-Guided Attentive Graph Convolutional Networks. In *Proceedings of DASFAA*, volume 13946 of *Lecture Notes in Computer Science*, 3–20. Springer.