Biomedical Event Causal Relation Extraction by Reasoning Optimal Entity Relation Path

Lishuang Li[1,*], Liteng Mi[1], Beibei Zhang[2], Yi Xiang[1], Yubo Feng[1],
Xueyang Qin[1], Jingyao Tang[1]
[1]School of Computer Science and Technology, Dalian University of Technology,
Dalian, Liaoning, China
[2]School of Computer Science and Technology,
Shandong Institute of Petroleum and Chemical Technology,
Dongying, Shandong, China
lils@dlut.edu.cn,{militn,lulu163163,xiangyi_job}@163.com,
argamx@126.com,qinxueyang@snnu.edu.cn,
tangjingyao@mail.dlut.edu.cn

Abstract

Biomedical Event Causal Relation Extraction (BECRE) is an important task in biomedical information extraction. Existing methods usually use pre-trained language models to learn semantic representations and then predict the event causal relation. However, these methods struggle to capture sufficient cues in biomedical texts for predicting causal relations. In this paper, we propose a Path Reasoning-based Relation-aware Network (PRRN) to explore deeper cues for causal relations using reinforcement learning. Specifically, our model reasons the relation paths between entity arguments of two events, namely entity relation path, which connects the two biomedical events through the multi-hop interactions between entities to provide richer cues for predicting event causal relations. In PRRN, we design a path reasoning module based on reinforcement learning and propose a novel reward function to encourage the model to focus on the length and contextual relevance of entity relation paths. The experimental results on two datasets suggest that PRRN brings considerable improvements over the state-of-the-art models.

1 Introduction

Biomedical Event Causal Relation Extraction (BECRE) is an important task in the field of biomedical information extraction which aims to identify the causal relation between two events in biomedical texts. As shown in Fig.1, given the context and two events, a BECRE model needs to predict whether there is a causal relation between head and tail events, i.e. "*phosphorylated* $\xrightarrow{cause}$ *degraded*". BECRE has played an important role in many downstream tasks such as knowledge question answering (Mutabazi et al., 2021), automatic decision-making, and knowledge discovery (Cañizares-Díaz et al., 2021).

BECRE is challenging because event causal relations in biomedical texts are always expressed in implicit ways. Existing methods usually utilize pre-trained language models to predict event causal relations (Akkasi and Moens, 2021). Liang et al. (2022) analyzed the impact of in-domain pre-training and distillation on the performance of BERT (Kenton and Toutanova, 2019). Akkasi and Moens (2021) used four different biomedical repositories that are publicly available to investigate the effect of the corpus used for pre-training BioBERT (Lee et al., 2020). These methods solely learned features of context relying on pre-trained language models. For further enhancing the ability of models to understand the context, Zhang et al. (2023) introduced external knowledge to build hierarchical knowledge graphs

*Corresponding author

Proceedings of the 23rd China National Conference on Computational Linguistics, pages 1087-1098, Taiyuan, China, July 25 - 28, 2024.
Volume 1: Main Conference Papers

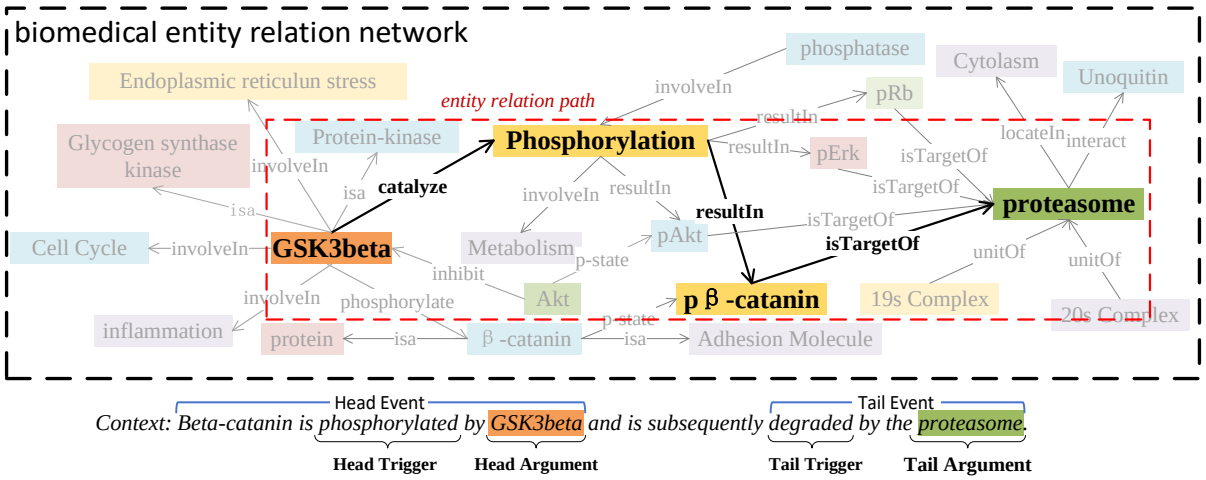

Figure 1: An illustration of expression of causal relations in biomedical texts and the entity relation path.

for the contexts. However, the above methods still have difficulty in capturing sufficient cues for event causal relations, leading to limited performance in BECRE.

To address the above issues, we pay attention to the entity relation path between biomedical events, and propose a method based on the entity relation path, named Path Reasoning-based Relation-Aware Network (PRRN). In PRRN, the entity relation path provides rich cues for predicting event causal relations. For example, in Fig.1, *"GSK3beta* $\xrightarrow{catalyze}$ *Phosphorylation* $\xrightarrow{resultIn}$ *pβ − catanin* $\xrightarrow{isTargetOf}$ *proteasome"* is a entity relation path between the event pair *(phosphorylated, degraded)*. The entity relation path describes multi-hop interactions between the head argument, *GSK3beta* and the tail argument, *proteasome*, including that *GSK3beta* can catalyze *Phosphorylation*, and *pβ-catanin* is the product of *Phosphorylation* and the target of *proteasome*. The interactions reveal an underlying cue that *β-catanin* is degraded by *proteasome* because *GSK3beta* converts *β-catanin* to *pβ-catanin* through *Phosphorylation*. So we can see that the entity relation path introduces multi-hop interactions of entities to provide rich cues for predicting event causal relation.

To obtain the optimal entity relation path, we design a path reasoning module based on reinforcement learning and propose a novel reward function. In the module, a policy-based agent network gradually approaches the optimal reasoning strategy in the way of Monte-Carlo Policy Gradient (Williams, 1992). The reward function we proposed in PRRN jointly encourages the length and contextual relevance of paths.

We conduct extensive experiments to evaluate the effectiveness of PRRN. The experimental results show that our method outperforms previous methods. In summary, the main contributions of this paper are:

- We propose a Path Reasoning-based Relation-Aware Network (PRRN) for BECRE, which explores richer cues to predict event causal relations by reasoning relation paths between entity arguments of biomedical events, i.e. entity relation paths.

- We design a path reasoning module based on reinforcement learning to reason optimal entity relation paths and propose a novel reward function to encourage the model to focus on the length and contextual relevance of entity relation paths.

- We conduct a series of experiments to verify the effectiveness of PRRN. The experimental results show that the entity relation path reasoned using reinforcement learning are beneficial for predicting event causal relation. Moreover, our method outperforms previous state-of-the-art methods on two datasets.

Proceedings of the 23rd China National Conference on Computational Linguistics, pages 1087-1098, Taiyuan, China, July 25 - 28, 2024.
Volume 1: Main Conference Papers

2 Related Work

Early studies of BECRE primarily relied on manually defined rules or shallow neural networks for event causal relation prediction (Mihăilă and Ananiadou, 2014; Hahn-Powell et al., 2016). Hahn-Powell et al. (2016) conducted rule-based and LSTM-based methods to practice BECRE. However, rule-based methods attempt to induce extremely rich and complex human language expressions with a small and limited number of rules, which is hard to extend. The modest parameter size and relatively simple model structure of LSTM-based methods are also not sufficient to support its ability to achieve a considerable level of language comprehension. In recent years, pre-trained language models have exhibited outstanding performances across many natural language processing tasks, attributed to their remarkable semantic understanding capabilities. Liang et al. (2022) and Akkasi and Moens (2021) utilized pre-trained language models to learn the semantic representation and achieved promising results. Akkasi and Moens (2021) also implemented and evaluated various techniques including a Multiview CNN (MVC), attention-based BiLSTM models, graph LSTM as well as a baseline rule-based system. Overall, compared to small-scale networks such as CNN, LSTM, and BiLSTM, pre-trained language models such as BERT (Kenton and Toutanova, 2019) and BioBERT (Lee et al., 2020) have shown better performances. However, these methods based on pre-trained language models are difficult to obtain enough semantic information solely in the context. Therefore, Zhang et al. (2023) introduced the external knowledge for BECRE and proposed the GECANet to fuse knowledge, which achieved state-of-the-art performance. In this paper, we put our emphasis on the entity relation path between biomedical events to capture richer cues for predicting event causal relations.

3 Method

Following the previous works (Zhang et al., 2023; Akkasi and Moens, 2021), we formulate BECRE as a classification problem. Given a sentence and an event pair, we predict whether there is a causal relation between the two events. Fig.2 schematically visualizes our approach, which mainly consists of three parts, including path reasoning module, representation module, and output module.

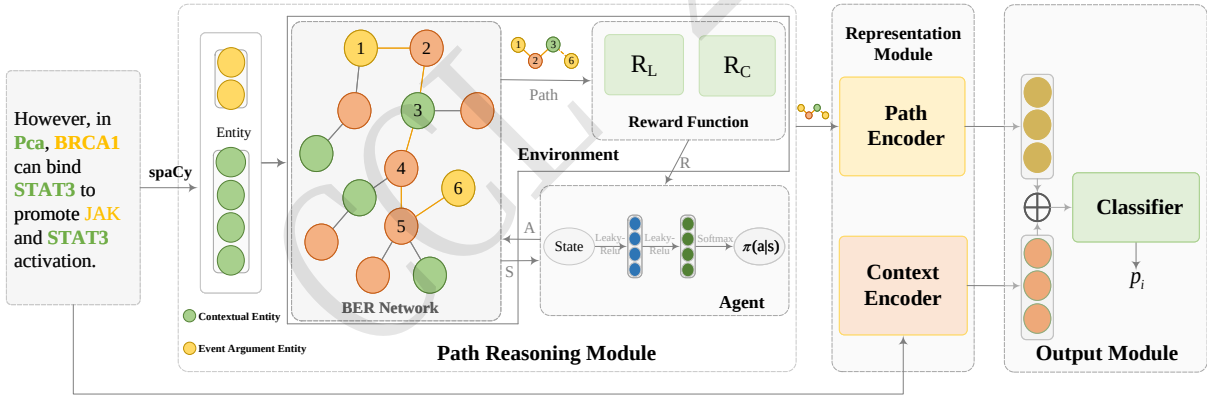

Figure 2: The illustration of our PRRN model, where "BER Network" is the abbreviation of biomedical entity relation network.

3.1 Path Reasoning Module

In this section, we introduce the process of reasoning paths based on reinforcement learning. Firstly, section 3.1.1 details the process of constructing a biomedical entity relation network. Secondly, for a given text and two events, we clarify the construction of reasoning tasks in section 3.1.2. Finally, section 3.1.3 introduces the reinforcement learning process.

3.1.1 Biomedical Entity Relation Network Construction

In this paper, we construct a biomedical entity relation network from the Unified Medical Language System (UMLS) (Bodenreider, 2004) to reason optimal entity relation paths.

UMLS is a vast knowledge base which contains rich information about biomedical concepts and mainly includes three parts: Metathesaurus, Semantic Network and SPECIALIST Lexicon & Lexical Tools. Metathesaurus contains the information of biomedical concepts and each Metathesaurus concept is assigned a concepts unique identifier (CUI). Semantic Network consists of semantic types and semantic relationships. Semantic types (STY) are broad subject categories, like *Disease* or *Syndrome* or *Clinical Drug*. Semantic relationships are relationships that exist between semantic types. Every concept in Metathesaurus is assigned at least one semantic type.

The complete UMLs contains over five million terms, or names so far and the scale makes running a reasoning model on it costly. Therefore, a sub-network is constructed from it for reasoning in this paper. Specifically, MetaMap, an open source tool that can map biomedical text to the Metathesaurus, is employed to extract CUI of biomedical entities involved in datasets. Based on the set of CUI, we query 11-hop neighbor nodes of them from the Metathesaurus and save as a sub-network. Then the Semantic Network is introduced into this sub-network to further enrich it. The STY nodes in sub-network are connected with CUI nodes by the edge "CUI2STY". At last, we add reverse of every edge to obtain the final biomedical entity relation network.

As a result, the nodes in biomedical entity relation network are divided into CUI nodes and STY nodes, and the edges can be divided into those between CUI nodes, those between STY nodes, and those between CUI and STY. The final network consists of 8311 nodes, 34075 triples, and 85 relations. After obtaining the biomedical entity relation network, TransE (Bordes et al., 2013) is used to obtain $e^d \in \mathbb{R}^d$ and $r^d \in \mathbb{R}^d$, i.e. representations of nodes and relations in the network.

3.1.2 Reasoning Task Construction

Before performing reasoning, we construct reasoning tasks to prepare inputs of reasoning module. Specifically, a python library, spaCy[1] is employed to extract all biomedical entities in sentence to obtain $E_A = \{e|e \in BiomedicalEntity,\ e \in sentence\}$. And then we utilize the syntactic analysis tool of spaCy to analyze the syntactic dependency ancestor nodes of each entities. If the k-hop ancestor node of an entity is the trigger word of the head or tail event, and $k < K_{threshold}$ ($K_{threshold}$ is a hyper-parameter), the entity will be listed into the head entities set, E_H or the tail entities set, E_T , respectively. After obtaining E_H and E_T, the inputs of reasoning module are constructed as E_{sample} in Eq.(1).

$$E_{sample} = \{(e_H, e_T), \psi_M \mid e_H \in E_H, e_T \in E_T,\ e_H \neq e_T, \psi_M = E_A - \{e_H, e_T\}\}, \tag{1}$$

where ψ_M denotes contextual entities set of size *M*, (e_H, e_T) denotes the entity pair to be reasoned. We don't take the corresponding entity pair into consideration in our method when the e_H is same as the e_T. For the convenience of locating entities in biomedical entity relation network in reasoning process, MetaMap mentioned in section 3.1.1 is employed again to map E_{sample} to $E^{CUI}_{sample} = \{(e^{CUI}_H, e^{CUI}_T), \psi^{CUI}_M\}$.

3.1.3 Reinforcement Learning Process

The Reinforcement Learning (RL) process is typically modeled as an interaction between an agent and the external environment. In this paper, the process can be understood as that a policy-based agent reasons path hop by hop for input entity pairs. At each step, it choose action based on the current state of the environment to extend the path. The external environment evaluate the performance of agent with the reward function when the reasoning of single entity pair is completed. The agent will update its parameters based on evaluations of the environment, thus gradually approaching the optimal reasoning strategy encouraged by the reward function. We detail the RL process in PRRN below and present an overview of the process in Algorithm 1.

Firstly, for the input entities in E^{CUI}_{sample}, we locate them with CUI in biomedical entity relation network to obtain their node id, $E^{id}_{sample} = \{(e^{id}_H, e^{id}_T), \psi^{id}_M\}$ and TransE representations, $E^{d}_{sample} = \{(e^{d}_H, e^{d}_T), \psi^{d}_M\}$.

[1] https://spacy.io/

Proceedings of the 23rd China National Conference on Computational Linguistics, pages 1087-1098, Taiyuan, China, July 25 - 28, 2024.
Volume 1: Main Conference Papers

Secondly, at each step *t*, the agent will choose action based on the current state of the external environment. The action space in this paper is actually the relation set in biomedical entity relation network, denoted as $\xi = \{r_1, r_2, ..., r_{85}\}$. The state s_t at step *t* is calculated as Eq.(2), where e_t denotes the node where the agent is located at step *t*, e_c denotes the pooling representation of the contextual entities and $||$ denotes the concatenation operation.

$$s_t = e_t^d \,||\, [e_T^d - e_t^d] \,||\, e_c^d \in \mathbb{R}^d,\ e_c^d = \frac{1}{M} \sum_{e \in \psi_M} e^d. \tag{2}$$

Algorithm 1 Reinforcement Learning Process

restore parameter θ from supervised learning
while i<samples_num **do**
  $\{(e_H^{CUI}, e_T^{CUI}), \psi_M^{CUI}\}_i \xrightarrow{\text{match}} \{(e_H^{id}, e_T^{id}), \psi_M^{id}\}_i$
  $\{(e_H^{id}, e_T^{id}), \psi_M^{id}\}_i \xrightarrow{\text{query}} \{(e_H^{d}, e_T^{d}), \psi_M^{d}\}_i$
  $E_{sample}^d = \{(e_H^d, e_T^d), \psi_M^d\}$
  $\{(e_H^d, e_T^d), \psi_M^d\}_i \xrightarrow{Eq.(2)} s_0$
  while $t : 0 \to t_{max}$ **do**
    choose action, $s_t \xrightarrow{\pi(a|s_t,\theta)} r_t$
    state transition, $(r_t, s_t) \xrightarrow{Eq.(2)} s_{t+1}$
    if reach e_T **then**
      break
    end if
  end while
  Path=$[(s_0, r_0), (s_1, r_1), \ldots, (s_n, *)]$
  if success **then**
    $R_A = \alpha \cdot R_L + (1 - \alpha) \cdot R_c$
  else
    $R_A = -1$
  end if
  update θ with Eq.(6)
end while

The agent in this paper is policy-based and implemented using a fully connected neural network, as the "Agent" part shown in Fig.2. The neural network is composed of two hidden layers and each layer is followed by a nonlinear activation function (i.e. LeakyRelu). The dimension of input layer and output layer is equal to the size of s_t and ξ, respectively. It maps the current state to probability distribution across the entire action space and can be represented as a policy network $\pi_\theta(s, a) = p(a|s; \theta)$, where θ denotes parameters of the network. The agent choose most promising actions based on the probability distribution on the action space at each step *t*.

The agent will randomly transfer to the next node based on the transition probability matrix (denoted as Eq.(3) in the environment after choosing the action.

$$P\left(s_{t+1} = s' \mid s_t = s, a_t = r_i\right) = \begin{cases} \frac{1}{|\varphi_{r_i}^s|} & , s' \in \varphi_{r_i}^s, \\ 0 & , others, \end{cases} \tag{3}$$

where $\varphi_{r_i}^s = \{s' \mid r_i \in r_{s,s'}\}$ denotes the neighbor nodes of s with relation r_i as the edge. In this transition strategy, if the chosen action r_i leads to an empty target state, the agent will go back to the original state s_t and make another new decision.

Repeat the above decision-transition process, the path will continuously extend until the agent reaches the tail entity or the number of reasoning steps reaches the threshold t_{max}. When the reasoning of single

entity pair is completed, the external environment will evaluate reasoning performance of the agent based on the reward function.

Given the reasoned path $p = \{e_0, r_0, e_1, r_1, ..., r_{L-1}, e_L\}$ with a length of L, the reward function in PRRN is designed for better capturing cues for predicting event causal relations and includes two components:

Length-optimized Reward. We set this reward component to encourage the agent to reach the target node with as few hops as possible. On the one hand, we observe that the shorter the path, the stronger the relation between the head and tail node it can express. Then it would be very beneficial for us to study the causal relation between events in BECRE task. On the other hand, shorter paths can also save computational resources and time and then improve reasoning efficiency. In addition, Length-optimized Reward are also aimed at mitigating the inherent challenges of sample inefficiency and high variance (Guo et al., 2021; Han et al., 2023) in policy-based RL algorithms when the path is longer. The length-optimized reward is defined as $R_L = \frac{1}{L}$.

Contextual relevance reward. This reward component encourages the agent to be closer to contextual entities during the reasoning process. The contextual entities, head entities, and tail entities share the same context and have strong correlation with each other. Therefore, using contextual entities information can quickly connect the head and tail entities, working together with length-optimized reward to improve reasoning efficiency. In addition, when the nodes on the path are closer to the contextual entity, it can better discover the cues hidden in the context about event causality, avoiding introducing noise that completely unrelated to the context. The contextual relevance reward is defined as R_C in Eq.(4).

$$R_C = \frac{1}{L}\sum_{i=0}^{L} \cos\left(e_i^d, e_c^d\right),\ e_i \in p. \tag{4}$$

Finally, the total reward R_A is obtained by weighting the above two reward components, as shown in the Eq.(5).

$$R_A = \alpha \cdot R_L + (1-\alpha) \cdot R_C, \tag{5}$$

where α is a hyper-parameter to balance two reward components. By weighting and integrating these two reward components, the agent can be encouraged to pursue both shorter path lengths and higher relevance to the context. In addition, when the agent fails to reach the target node after a certain round of interaction with the environment, we deem the reasoning as failed, and set $R_A = -1$ to give the agent a negative reward.

We update the policy network with Monte-Carlo Policy Gradient (Williams, 1992) and gradient descent, as shown in Eq.(6).

$$\begin{aligned} G_r(\theta) &= E_{\tau\sim\pi_\theta} \sum_{t=0}^{L} [-R(\tau)\nabla_\theta \log \pi(a_t|s_t;\theta)] \\ &\approx \nabla_\theta \sum_{t=0}^{L} [-\log \pi(a = r_t|s_t;\theta) R_A]. \end{aligned} \tag{6}$$

3.2 Representation Module

3.2.1 Path Encoder

Recently, recurrent neural networks have been widely used in processing sequence data such as path information. Compared to vanilla GRU, BiGRU brings more comprehensive understanding of the information in the sequence. For the reasoned entity relation path $p = \{e_0, r_0, e_1, ..., r_{L-1}, e_L\}$, the BiGRU (Chung et al., 2014) is assigned to encode it in PRRN.

We initialize the representation of the path using TransE representations of edges and nodes. These node and relation representations integrate the topological information of nodes and relations in the biomedical entity relation network, allowing us to more thoroughly utilize the background knowledge of the knowledge base. Then the representation of entity relation paths is fed into BiGRU. We concatenate

the last hidden state of BiGRU output from two directions to obtain the final representation of the entity relation path, as shown in Eq.(7).

$$V_p = BiGRU^D\left(\left[e_0^d, r_0^d, e_1^d, r_1^d, \ldots, r_{L-1}^d, e_L^d\right]\right) \in \mathbb{R}^D. \tag{7}$$

3.2.2 Context Encoder

Trigger words are important component of an event, and the semantic representation of trigger words is valuable reference for the model when predicting the relation between events. Therefore, we use BioBERT (Lee et al., 2020) to encode the text and extract the word representations corresponding to the trigger words. If the trigger words are more than one word, we extract the uniformed pooling representations of multiple trigger words, as shown in Eq.(8) and Eq.(9).

$$V_{head\_event} = \frac{1}{h_n}\sum_{i=1}^{h_n} f_c(head\_trigger_i|context) \in \mathbb{R}^D, \tag{8}$$

$$V_{tail\_event} = \frac{1}{t_n}\sum_{i=1}^{t_n} f_c(tail\_trigger_i|context) \in \mathbb{R}^D, \tag{9}$$

where h_n and t_n respectively denotes the number of trigger words for the head and tail events, and $f_c(w|context)$ denotes the contextual semantic representation of word w from BioBERT (Lee et al., 2020). Benefiting from the attention mechanism (Bordes et al., 2013) of BioBERT (Lee et al., 2020), these two word representations simultaneously contain the semantic and contextual information of trigger words, which is effective in predicting causal relations between events.

3.3 Output Module

In section 3.2, we calculate the representation of entity relation paths and the contextual semantic information of events, respectively. For these two representations, we add the representation vectors of entity relation paths to the representation vectors of head and tail events, respectively. The two fusion vectors is then concatenated to calculate the final representation of each event pair, as shown in Eq.(10).

$$V_{feature} = [V_{head\_event} + V_p] \,||\, [V_{tail\_event} + V_p]. \tag{10}$$

And then $V_{feature}$ is fed into a fully connected neural network followed by a softmax function to obtain the final prediction. As shown in Eq.(11), W_o denotes the parameters of the fully connected neural network, and p_i denotes the predicted probability of *i*-th class.

$$p_i = softmax(W_o \cdot V_{feature}),\ W_o \in \mathbb{R}^{d_{class} \times D}. \tag{11}$$

3.4 Training Process

Following the previous work (Xiong et al., 2017; Silver et al., 2016), the training of policy network in agent is divided into two stages in our method including supervised learning and retraining with reinforcement learning. Supervised learning enables the policy network to have initial reasoning ability in a vast action space, thereby accelerating the convergence speed of reinforcement learning.

In the supervised learning stage, we run the breadth-first searching algorithm on the biomedical entity relation network to search paths between the input entity pair. The searched path is employed as the supervised information to update policy network. We use the policy gradient and set the reward to 1 to update the policy network in the agent. The gradient is shown in the Eq.(12).

$$G_s(\theta) = \bigtriangledown_\theta \sum_{t=0}^{n} log\pi(a = r_t|s_t;\theta), \tag{12}$$

where s_t and r_t are the states and actions of the *t*-th hop in the path, respectively. As for the sample set used in supervised learning, we conduct the reasoning sample construction as mentioned in section 3.1.2 and extract a sample set of E^{CUI}_{sample} with size *samples_num*.

Then in the reinforcement learning stage, we use the reward function to adjust the policy network. Firstly, the agent uses the prior knowledge obtained from the supervised learning stage to reason entity relation paths hop by hop. Next, the external environment will evaluate the performance of the agent using the reward function and feedback the evaluation to the agent. Finally, the policy network calculate the policy gradient based on this evaluation to update its parameters. By repeating this process, the policy network will gradually approach the optimal reasoning policy encouraged by the reward function.

We train the context encoder, path encoder and classifier together using cross entropy loss. In this stage, the policy network is called offline and does not participate in training. The loss we applied is shown in Eq.(13).

$$L = -\frac{1}{N}\sum_{i}\sum_{c=1}^{3}\mathbb{1}_c(i)log(p_{i,c}), \tag{13}$$

where N denotes the number of samples, $\mathbb{1}_c(i)$ is an indicator function that is assigned 1 when $i = c$ and 0 when $i \neq c$ and $p_{i,c}$ denotes the predicted probability that the i-th sample belongs to c-th category.

4 Experiment

4.1 Experimental Setup

Datasets and Metrics. We evaluate our proposed method on two datasets, including Hpowell (Hahn-Powell et al., 2016) and BioCause (Mihăilă et al., 2013). Following prior works (Zhang et al., 2023; Akkasi and Moens, 2021), in Hpowell and BioCause, we only consider two types of relations, "E1 precedes E2" and "E2 precedes E1" as target relations and consider all the others as no relation, denoted as "None". Tab.1 presents the number of samples for each category in two datasets. We adopt Recall (R), Precision (P), and F1-score (F_1) as evaluation metrics.

Dataset	Class	Number
Hpowell (Hahn-Powell et al., 2016)	E1 precedes E2	163
	E2 precedes E1	27
	None	637
Biocause (Mihăilă et al., 2013)	E1 precedes E2	36
	E2 precedes E1	14
	None	3140

Table 1: Statistics of datasets used for the experiments.

Implemention Details. In the experiments, we use BioBERT-large as the context encoder. The representation dimension of token is 768. The $K_{threhold}$ mentioned in section 3.1.2 is set to 2, and t_{max} mentioned in section 3.1.3 is set to 50. We train the context encoder, path Encoder and classifier with the learning rate=5e-5 and the dropout rate=0.1. Following the previous work (Zhang et al., 2023; Akkasi and Moens, 2021), we adopt the oversampling method to alleviate the sample imbalance problem in the Hpowell and BioCause datasets. Specifically, the categories "E2 precedes E1" and "E1 precedes E2" are sampled to equal numbers, matching the quantity of samples in the "None" category. The dimension of hidden layers in the policy network of Path Reasoning Module are 512 and 1024, respectively. We train TransE model on our biomedical entity relation network for 1000 training rounds[2].

Baseline Methods. We compare the proposed PRRN with the following baseline models. Akkasi and Moens. (2021) overviewed the BECRE task and contributed some credible benchmarks, as follows: **Rule-based Model** manually summarizes some causative verbs from the corpus and designs rules to predict the event causal relation based on causative verbs. **Graph state LSTM** models the document as a graph and captures a variety of dependencies among the input words which are represented by nodes of the graph. **Multiview-CNN** integrates feature representations of context calculated from three different sizes of convolution kernels. **BiLSTM-ATTENTION** utilizes BiLSTM to calculate the word representation of tokens and obtain the representation of texts by fusing representation of each token with attention mechanism. **ELMO-LSTM** simply applies ELMO as a word representation layer, and feeds representation result into a LSTM module. **BioBERT-MVC** uses four BioBERT (denoted as BioBERT1, BioBERT2, BioBERT3 and BioBERT4 below) that pre-trained on four different public

[2]The implemention we used can be found at https://www.github.com/thunlp/Fast-TransX

biomedical repositories to obtain the representation of each token, respectively and employs Multiview-CNN as the inference layer. In addition, **GECANet** (Zhang et al., 2023) introduces knowledge triplets from an external knowledge base and calculates the edge representation, node representation, and graph representation of the triplet set separately. Then it integrates them with context representation to predict event causal relations. For the above methods, Akkasi and Moens (2021) and Zhang et al. (2023) have both detailed the experimental results utilizing both raw data and oversampled data. Within this section, our focus is solely on comparing these methods under the oversampling strategy.

4.2 Experimental Results

Source	Model	Hpowell			BioCause		
		P(%)	R(%)	F_1(%)	P(%)	R(%)	F_1(%)
Akkasi and Moens (2021)	Rule-based Model	-	-	-	33.0	21.0	25.0
	Graph state LSTM	57.0	45.0	50.0	54.0	36.0	43.0
	Multiview-CNN	60.0	44.0	52.0	77.0	20.0	32.0
	BiLSTM-ATTENTION	54.0	49.0	52.0	59.0	35.0	44.0
	ELMO-LSTM	55.0	52.0	54.0	27.0	24.0	26.0
	BioBERT1-MVC	65.0	53.0	58.0	69.0	18.0	29.0
	BioBERT2-MVC	59.0	54.0	57.0	67.0	20.0	31.0
	BioBERT3-MVC	63.0	53.0	58.0	70.0	29.0	41.0
	BioBERT4-MVC	60.0	51.0	55.0	**79.0**	31.0	44.0
Zhang et al. (2023)	GECANet	64.6	60.5	62.5	78.2	32.0	45.5
ours	PRRN	**65.7**	**61.5**	**63.6**	75.8	**44.0**	**55.7**

Table 2: Experimental results on Hpowell and BioCause, where best results are bolded, and suboptimal results are underlined.

Tab.2 shows the results on the Hpowell and BioCause datasets, respectively. Overall, our method outperforms the methods proposed in previous works on both datasets in terms of F1-score. It can be seen that pre-trained language model based method like BioBERT3-MVC significantly outperforms Multiview-CNN model by 6.0% on Hpowell dataset. This is because the quite large parameter scale and the process of pre-training on large-scale corpora offer pre-trained language model stronger semantic representation ability than the small scale network. However, the method of solely utilizing the semantic representation ability of pre-trained language models ignores the causal cues stemmed from biomedical entities interactions. Therefore, the results of BioBERT3-MVC method on two datasets are 5.6% and 11.7% lower than our method, respectively, highlighting the effectiveness of learning to capturing deep cues with entity relation paths. In addition, our method surpasses the GECANet by 1.1% and 10.2% on the two datasets, respectively. This indicates that instead of directly incorporating external knowledge, delving into the entity relation path between biomedical events to discover the causal cues is a more beneficial method for BECRE.

Furthermore, in terms of recall and precision, our approach exhibits substantial enhancements over previous methods on the Hpowell dataset. On the BioCause dataset, while our method falls slightly behind Multiview-CNN, BioBERT-MVC and GECANet in precision, 1.2%, 3.2% and 2.4%, respectively, it achieves a remarkable enhancement in recall, which ensuring the reliability of our proposed method.

4.3 Ablation Study

In this section, we aim to ablate the major components in our model and evaluate the performance of the remaining model to understand contribution of each component. We examine the following ablated models: (1) "$-PRRN_{w/o.R_L}$" and "-$PRRN_{w/o.R_C}$" exclude the reward components R_L and R_C (respectively) from the overall reward when training the agent. (2) "-$PRRN_{w/o.path}$" indicates the model predicts event causal relations without the entity relation path. (3) "-$PRRN_{w/o.RL}$" indicates the model predicts event causal relations with the entity relation path obtained from the breadth-first searching algorithm.

Proceedings of the 23rd China National Conference on Computational Linguistics, pages 1087-1098, Taiyuan, China, July 25 - 28, 2024.
Volume 1: Main Conference Papers

Tab.3 reports the performance of ablated models on the Hpowell dataset. As can be seen, removing the entity relation path or any reward components significantly hurts the overall performance of the model. The largest performance drop is due to the elimination of the entity relation path, suggesting that the entity relation path is critical to reveal the cues for predicting event causal relations. In addition, compared to the breadth-first searching algorithm, the superior performance of reinforcement learning clearly highlights the advantage to reason the entity relation path with reinforcement learning for BECRE.

Method	P(%)	R(%)	F_1(%)	Δ
-PRRN$_{w/o.path}$	60.2	57.4	58.8	-
-PRRN$_{w/o.R_L}$	65.9	59.0	62.2	+3.4
-PRRN$_{w/o.R_C}$	66.9	56.3	61.1	+2.3
-PRRN$_{w/o.RL}$	**67.3**	57.4	61.9	+3.1
PRRN(ours)	65.7	**61.6**	**63.6**	**+4.7**

Table 3: Results of ablation study on reinforcement learning. Δ denotes improvements of F_1 compared to -PRRN$_{w/o.path}$.

4.4 Entity Relation Path Encoding Strategy Analysis

Model	P(%)	R(%)	F_1(%)
LSTM	61.8	57.9	59.8
BiLSTM	64.3	60.5	62.3
GRU	**68.1**	57.4	62.6
BiGRU	65.7	**61.6**	**63.6**

Table 4: Experimental results using different entity relation path encoders on the Hpowell dataset.

Method	Hpowell	BioCause
ERP&Node	61.5(↓ 2.2)	52.1(↓ 3.6)
ERP&Edge	61.1(↓ 2.6)	53.9(↓ 1.8)
ERP&Full	**63.6**	**55.7**

Table 5: Experimental results (F_1) of retaining only edges and nodes in ERP(Entity Relation Path) separately. ↓ means the decrease compared to ERP&Full.

Effect of Entity Relation Path Encoder. To analyze the effect of different entity relation path encoders on experimental results, we utilize four different path encoders, LSTM, BiLSTM, GRU, and BiGRU. The results are shown in Tab.4. From the results, we can observe that the best performing path encoders is BiGRU. This suggest that BiGRU has the capability to utilize the sequence information in the entity relation path more effectively, enabling more accurate aggregation of information from edges and nodes. In addition, the results also indicates that adopting bidirectional encoding for the entity relation path is more effective than unidirectional encoding.

Effect of Nodes and Edges in the entity relation path. To better know the effect of different path encoding strategies, we examine another two encoding strategies, "ERP&Node" and "ERP&Edge", which keeps only nodes and edges respectively in the paths in the phase of encoding. The results are shown in Tab.5, compared to using both edges and nodes simultaneously, eliminating each of them showed a reduced effect, indicating that both edge sequences and node sequences contain causal information between biomedical events. However, the complete path including both edges and nodes is more effective in predicting event causal relations.

4.5 The effect of hyper-parameter α

In Path Reasoning Module, the hyper-parameter α is a adjustment term in Eq.(4), controlling the attention of the agent to length and contextual relevance of paths, respectively. To analyze the impact of α on model performance, we set α to different values to observe the F_1.

As shown in Fig.3, the F_1 score peaks when $\alpha = 0.5$, indicating that giving the same weight to the contextual relevance and length of the path can yield the best benefits on the Hpowell dataset. Moreover, there is a

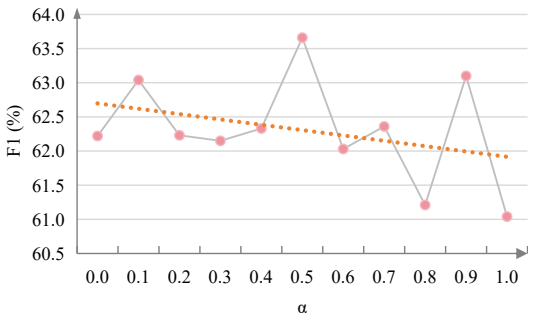

Figure 3: The effect of hyper-parameter α

discernible downward trend in the trendline, demonstrating that contextual relevance is more important than length of the path for capturing causal cues.

Conclusion

In this paper, we emphasize the importance of the entity relation path for predicting event causal relations between biomedical events and present a novel model, namely, PRRN for Biomedical Event Causal Relation Extraction. For obtaining the valuable entity relation path, we propose a path reasoning module based on reinforcement learning and design length-optimized reward and contextual relevance reward to evaluate the performance of reinforcement learning agent. Our extensive experiments show that the PRRN outperforms the previous state-of-the-art model, demonstrating the advantage of the entity relation path reasoned under the control of two reward components. In the future, we will extend our proposed method to other related tasks in events relation extraction (e.g., for event causal relation extraction in universal domain).

Acknowledgements

This work is supported by grant from the National Natural Science Foundation of China (No. 62076048).

References

Abbas Akkasi and Mari-Francine Moens. 2021. Causal relationship extraction from biomedical text using deep neural models: A comprehensive survey. *Journal of biomedical informatics*, 119:103820.

Olivier Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.

Hian Cañizares-Díaz, Alejandro Piad-Morffis, Suilan Estevez-Velarde, Yoan Gutiérrez, Yudivián Almeida Cruz, Andres Montoyo, and Rafael Muñoz. 2021. Active learning for assisted corpus construction: A case study in knowledge discovery from biomedical text. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 216–225.

Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*.

Hanlin Guo, Bowen Tan, Zhengzhong Liu, Eric P Xing, and Zhiting Hu. 2021. Text generation with efficient (soft) q-learning. *arXiv*.

Gus Hahn-Powell, Dane Bell, Marco A Valenzuela-Escárcega, and Mihai Surdeanu. 2016. This before that: Causal precedence in the biomedical domain. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*, pages 146–155.

Jiale Han, Mingxiao Feng, Wengang Zhou, and Houqiang Li. 2023. Sample efficient reinforcement learning with double importance sampling weight clipping. In *2023 IEEE Conference on Games (CoG)*, pages 1–8. IEEE.

Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.

Zhengzhong Liang, Enrique Noriega-Atala, Clayton Morrison, and Mihai Surdeanu. 2022. Low resource causal event detection from biomedical literature. In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 252–263.

Claudiu Mihăilă and Sophia Ananiadou. 2014. Semi-supervised learning of causal relations in biomedical scientific discourse. *Biomedical engineering online*, 13:1–24.

Claudiu Mihăilă, Tomoko Ohta, Sampo Pyysalo, and Sophia Ananiadou. 2013. Biocause: Annotating and analysing causality in the biomedical domain. *BMC Bioinformatics*, 14:2.

Emmanuel Mutabazi, Jianjun Ni, Guangyi Tang, and Weidong Cao. 2021. A review on medical textual question answering systems based on deep learning approaches. *Applied Sciences*, 11(12):5456.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.

Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256.

Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. Deeppath: A reinforcement learning method for knowledge graph reasoning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Beibei Zhang, Lishuang Li, Dingxin Song, and Yang Zhao. 2023. Biomedical event causal relation extraction based on a knowledge-guided hierarchical graph network. *Soft Computing*, 27(22):17369–17386.