

Reward-based Input Construction for Cross-document Relation Extraction

Byeonghu Na^{1*}, Suhyeon Jo^{1*}, Yeongmin Kim¹, Il-Chul Moon^{1,2}

¹KAIST, ²summary.ai

{byeonghu.na, suhyeonjo, alsdudrla10, icmoon}@kaist.ac.kr

Abstract

Relation extraction (RE) is a fundamental task in natural language processing, aiming to identify relations between target entities in text. While many RE methods are designed for a single sentence or document, cross-document RE has emerged to address relations across multiple long documents. Given the nature of long documents in cross-document RE, extracting document embeddings is challenging due to the length constraints of pre-trained language models. Therefore, we propose REward-based Input Construction (REIC), the first learning-based sentence selector for cross-document RE. REIC extracts sentences based on relational evidence, enabling the RE module to effectively infer relations. Since supervision of evidence sentences is generally unavailable, we train REIC using reinforcement learning with RE prediction scores as rewards. Experimental results demonstrate the superiority of our method over heuristic methods for different RE structures and backbones in cross-document RE. Our code is publicly available at <https://github.com/aailabkaist/REIC>.

1 Introduction

The task of relation extraction (RE) aims to identify relations between a pair of target entities in a given text (Zeng et al., 2014; Zhang et al., 2017; Sahu et al., 2019; Xu et al., 2021). This task is fundamental in natural language processing (NLP) and provides crucial information for applications such as information retrieval (Kadry and Dietz, 2017) and question answering (Li et al., 2019; Chen et al., 2019). Most existing RE methods are limited to scenarios where the entity pair is within a single sentence (Wang et al., 2016; Zhu et al., 2019) or a single document (Nan et al., 2020; Huang et al., 2021; Xu et al., 2021). However, many scenarios require extracting relations across multiple documents, where the target entity pair may not coexist

within the same document. Therefore, recent efforts to address these challenges have led to the proposal of cross-document RE by Yao et al. (2021), and research in this area has received considerable attention (Wang et al., 2022; Wu et al., 2023; Lu et al., 2023; Son et al., 2023).

Unlike traditional RE research, cross-document RE involves extracting relational facts from large-scale long documents. For example, the DocRED dataset (Yao et al., 2019), developed for document-level RE, contains an average of 198 words per document. In contrast, the CodRED dataset (Yao et al., 2021), designed for cross-document RE, has an average of 2,416 words per document. Given the substantial length of documents in cross-document RE tasks, extracting document embeddings from pre-trained language models, a step common to all methods, poses considerable challenges. This is because pre-trained language models are limited by the maximum number of tokens they can process; for example, BERT (Devlin et al., 2019) typically has a limit of 512 tokens.

A simple approach adopted in several studies is to use text snippets around the target entities as input to pre-trained language models, as shown in the dashed boxes in the middle panel of Figure 1. Often, this text snippet is determined by either word proximity or document structure without any adjustment. However, if only a subset of sentences around the entities is extracted from long documents, there is a risk of missing important sentences that are crucial for RE. For example, as seen in the left panel of Figure 1, using snippets around the entities fails to capture the relation ‘platform’ between the entity pair (Kubuntu, x86-64).

In this paper, we introduce REward-based Input Construction (REIC), the first learning-based input sentence selection module specifically designed for cross-document RE. The right panel of Figure 1 displays the sentences selected by our model from the example. By identifying relational evidence

*Equal contribution

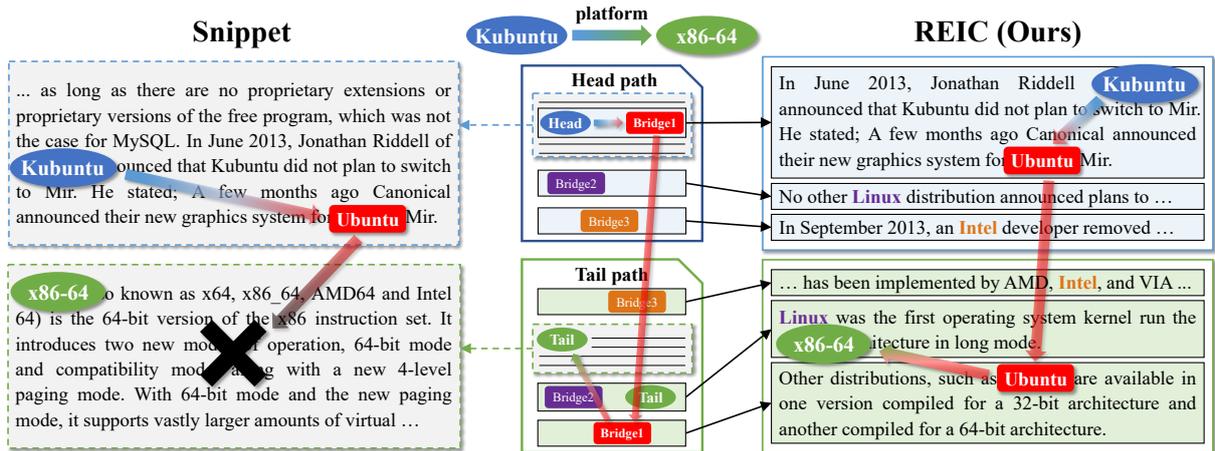


Figure 1: An illustrated comparison between Snippet and selected sentences using our REward-based Input Construction (REIC) for cross-document relation extraction. The figure depicts an example triplet (Kubuntu, x86-64, platform) with the text path (‘Mir (software)’, ‘X86-64’), including three bridge entities (Ubuntu, Linux, Intel) abbreviated as (Bridge1, Bridge2, Bridge3). Dash and solid arrows signify the selection process of Snippet and REIC, respectively, while gradient-colored arrows indicate connections between the head and tail entities. REIC selects important sentences from any position within a path to determine the relation between the head and tail entity, whereas Snippet only includes sentences located around the head or tail entity.

sentences through the selection module, the RE module infers the relation between the target entities using the reasoning chain illustrated by the gradient-colored arrows in Figure 1. Our approach is to develop a sentence selection module that computes the probability of selecting sentences based on the currently selected sentences. We show that this sentence selection process can be modeled as a Markov decision process (MDP). We specifically choose to use MDP because there is no supervision for sentence selection from corpus, so RE models have to perform iterative learning by trials. Subsequently, we train the sentence selection module using reinforcement learning (RL), where the relation prediction score obtained from the selected sentences serves as the reward. Through experimental validation, REIC outperforms other heuristic methods across various RE modules.

2 Related Work

2.1 Cross-document Relation Extraction

Traditionally, research in RE has primarily focused on the sentence- or document-level, typically dealing with documents no longer than a single paragraph, as reviewed in Appendix A. However, there has been a growing interest in cross-document RE, which aims to identify relations between entities across multiple long documents. Yao et al. (2021) laid the groundwork for cross-document RE by introducing the CodRED benchmark dataset. They

also proposed a simple baseline model that utilizes document embeddings obtained from pre-trained language models like BERT (Devlin et al., 2019). Building on this work, ECRIM (Wang et al., 2022) proposed an entity-based input filtering approach and an entity-based RE module with an attention mechanism, leading to significant performance improvements. Additionally, Wu et al. (2023) introduced a local causal estimation algorithm and a globally guided reasoning algorithm based on causality. In addition, Lu et al. (2023) presented a method for multi-hop evidence retrieval through evidence path mining and ranking, and Son et al. (2023) proposed a method for constructing explicit reasoning paths using bridging entities. In our approach, we focus on selecting essential sentences to effectively extract document information from long documents for cross-document RE. Unlike other methods, our approach involves learning-based input construction, and these surveyed models can utilize our input construction method.

2.2 Sentence Selection in NLP

Numerous studies have addressed the challenge of selecting relevant sentences from extensive and noisy textual datasets while preserving contextual coherence. These approaches include deep learning-based binary classification for each sentence (Cheng and Lapata, 2016), auto-regressive techniques to capture sentence correlations (Zhou et al., 2018; Zhong et al., 2019; Zhang et al.,

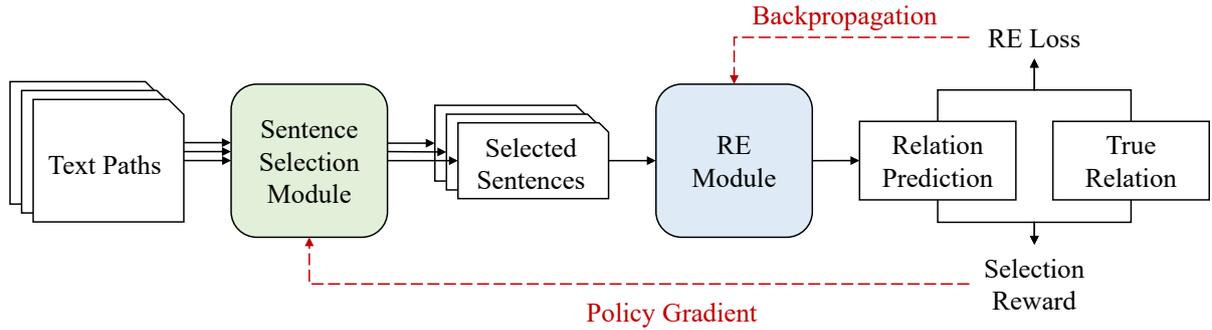


Figure 2: Overview of our cross-document RE framework. Solid black arrows represent forward propagation, while the dashed red arrows represent the loss signal. Documents from each text path are processed by the sentence selection module to construct the input for the RE module, which then obtains relation predictions.

2023b), methods that consider summary-level context (Zhong et al., 2020), and the use of diffusion models (Zhang et al., 2023a).

Other lines of research utilize RL for sentence selection (Narayan et al., 2018). This approach naturally reflects the correlation between sentences, as the selection process is optimized by considering the current state, which includes the sentences already selected. Furthermore, RL offers flexibility in reward shaping, making it suitable for specific downstream tasks. For example, Feng et al. (2018); Qin et al. (2018); Qu et al. (2019); Takanobu et al. (2019) used RL to filter noisy datasets, thereby improving the performance of sentence-level RE. Also, Xu et al. (2022) utilized RL to evidence extraction in document-level RE, and Man et al. (2022) applied RL to find input tokens for BERT embedding within documents for event-event RE. Our approach builds on previous methods for sentence selection and represents the first application of RL to cross-document RE.

3 Methods

3.1 Problem Formulation

The cross-document RE (Yao et al., 2021) aims to infer a relation $r \in \mathcal{R}$ between a pair of target entities (e^h, e^t) , where e^h is the head entity and e^t is the tail entity, from a bag of text paths $\{p_n\}_{n=1}^N$. Each text path p_n consists of two documents (d_n^h, d_n^t) that satisfy the following conditions: 1) d_n^h and d_n^t contain the head and tail entities, respectively; 2) there exists at least one entity, called a bridge entity, which is shared by both documents.

Our goal is to develop a method that effectively extracts sentences from an input document to facilitate the relation extraction between two entities. Given a document d consisting of M sentences, de-

noted by $d = \{s_m\}_{m=1}^M$, our approach is to select important sentences and obtain document embeddings using pre-trained language models such as BERT (Devlin et al., 2019). Given the maximum input length constraint of pre-trained language models, collecting significant sentences enables effective extraction of document information.

3.2 Overall Framework

We begin by describing our proposed framework as depicted in Figure 2. We divide the cross-document RE module into the input sentence selection module and the RE module. The goal of the sentence selection module is to identify important sentences for RE in each document. Next, the RE module outputs relation prediction scores based on the sentences identified by the sentence selection module.

Specifically, for each text path $p_n = (d_n^h, d_n^t)$, we utilize the sentence selection module to extract a summarized text path $\bar{p}_n = (S_n^h, S_n^t)$, containing essential sentences to recognize relations between entity pairs. Afterwards, these summarized text paths become an input to the RE module, which employs a pre-trained language model as its backbone, to produce relation prediction scores $\hat{y}(S_n^h, S_n^t)$ for each path. The RE module is trained using relation prediction loss. For example, ECRIM (Wang et al., 2022) adopts a multi-label global-threshold loss:

$$\mathcal{L}_{\text{RE}} = \log(e^\theta + \sum_{r \in \mathcal{R} \setminus \Omega^{\text{h,t}}} e^{\hat{y}_r}) + \log(e^{-\theta} + \sum_{r \in \Omega^{\text{h,t}}} e^{-\hat{y}_r}), \quad (1)$$

where θ is the threshold and $\Omega^{\text{h,t}}$ denotes the set of relations that match the relation of the target entity pair. On the other hand, we compute the selection reward to train the sentence selection module using RL, as explained in Section 3.4. It is important

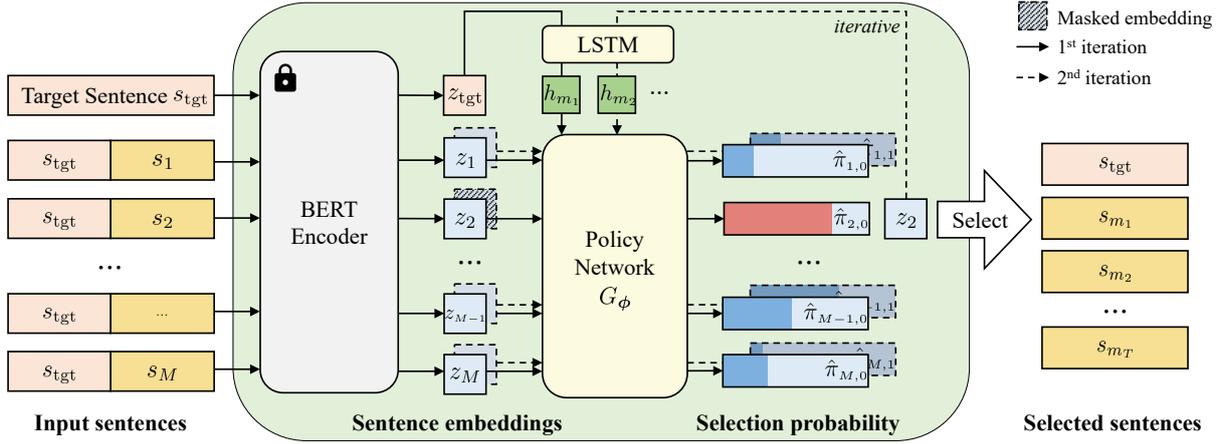


Figure 3: Illustration of the REIC module: BERT encoder, policy network, and LSTM. Each sentence s_m , combined with the target sentence s_{tgt} , is processed by the fixed BERT to obtain the embedding z_m . Then, these embeddings pass through the policy network G_ϕ to obtain the selection probability $\hat{\pi}_{m,t} := \hat{\pi}(s_m|S_t)$, and the sentence is sampled with this probability. The embedding of the selected sentence is fed to the LSTM to incorporate information from previously selected sentences into the policy network. In this way, sentences are sequentially selected, with masking applied to embeddings for subsequent selections, as indicated by the dotted line.

to note that our sentence selection module can be integrated into any existing RE modules.

3.3 REIC: Reward-based Input Construction

We present the proposed sentence selection module, called REward-based Input Construction (REIC), for cross-document RE. We build the structure of the REIC module inspired by Man et al. (2022), as shown in Figure 3. First, we obtain the representation z_m of each sentence s_m . We denote the sentence containing the target entity as s_{tgt} . We are trying to extract information relevant to the target entities in this document. Therefore, it would be helpful to include the information from the sentence that contains the target entity. To achieve this, we concatenate s_{tgt} and s_m and input them into the pre-trained BERT encoder to generate a sentence embedding z_m for the m^{th} sentence:

$$z_m = \text{BERT}([s_{\text{tgt}}, s_m]). \quad (2)$$

Note that we maintain the positions in the document when concatenating.

Next, we propose an iterative sentence selection process. We introduce the policy network G_ϕ , with trainable parameters ϕ , for generating the sentence selection probabilities. The term ‘policy’ is used because we treat the sentence selection process as a policy and apply policy gradient methods, as we will explain in Section 3.4. The policy network takes the sentence embeddings as input and outputs the probability of selecting each sentence based on the previously selected sentences.

Specifically, we denote the set of selected sentences after iteration t as S_t , and we initialize it as $S_0 = \{s_{\text{tgt}}\}$, meaning that the target sentence is selected by default. At iteration $t + 1$, we aim to select the next important sentence $s_{m_{t+1}}$ from the given set of $t + 1$ selected sentences, $S_t = \{s_{\text{tgt}}, s_{m_1}, \dots, s_{m_t}\}$. To provide information about the sentences already selected, we introduce the Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) network to the sentence embedding z_m . By feeding the embedding z_{m_t} of the selected sentence at time t into the LSTM, the hidden vector h_{m_t} of the LSTM serves as the representation of the previously selected sentences:

$$h_{m_t} = \text{LSTM}(z_{m_t}, h_{m_{t-1}}). \quad (3)$$

Then, we concatenate the embedding h_{m_t} of the selected sentences with the embedding z_m of each unselected sentence and feed them into the policy network G_ϕ to obtain the selection probabilities:

$$\hat{\pi}(s_m|S_t) = G_\phi([z_m, h_{m_t}]). \quad (4)$$

We select the sentence from the sampling with the selection probability as the next selected sentence:

$$s_{m_{t+1}} \sim \text{Categorical}(\hat{\pi}(s_m|S_t)). \quad (5)$$

This process is iterated for a maximum number of iterations T and is outlined in Algorithm 1.

Algorithm 1: Algorithm of REIC

Input: Input sentences $\{s_m\}_{i=1}^M$
Output: Selected sentences S

- 1 $\mathcal{M} \leftarrow \{1, \dots, M\}$
- 2 $z_m \leftarrow \text{BERT}([s_{\text{tgt}}, s_m])$ for $m \in \mathcal{M}$
- 3 $S \leftarrow \{s_{\text{tgt}}\}$
- 4 $z \leftarrow z_{\text{tgt}}$
- 5 $\mathcal{M} \leftarrow \mathcal{M} \setminus \{\text{tgt}\}$
- 6 $h \leftarrow \emptyset$
- 7 **for** $t \in \{1, \dots, T\}$ **do**
- 8 $h \leftarrow \text{LSTM}(z, h)$
- 9 $\hat{\pi}_m \leftarrow G_\phi([z_m, h])$ for $m \in \mathcal{M}$
- 10 $\hat{m} \sim \text{Categorical}(\hat{\pi}_m)$
- 11 $S \leftarrow S \cup \{s_{\hat{m}}\}$
- 12 $z \leftarrow z_{\hat{m}}$
- 13 $\mathcal{M} \leftarrow \mathcal{M} \setminus \{\hat{m}\}$
- 14 **end**

3.4 Training REIC with RL

We discuss the method of training the REIC module based on RL. We represent the sentence selection process as a MDP. The state space consists of combinations of selected sentences, which is represented by a set of selected sentences, S . Selecting one sentence at each iteration is considered as an action. Thus, when the action of selecting sentence s_m from the current state S is taken, it transitions to the state $S \cup \{s_m\}$. In this case, the policy, which is the distribution of actions given a state, can be represented by the selection probability from Eq. (4). Once the action is determined; the next state is also determined, so is the transition probability.

The proposed sentence selection module ultimately needs to choose sentences that perform well for the RE task. Therefore, it is necessary to set the reward function to provide signals for RE. An intuitive approach is to identify the important sentences necessary for RE and provide rewards when these sentences are selected. However, obtaining the important sentences from the large dataset requires specialized knowledge and a considerable amount of time, making it extremely challenging. Therefore, we intend to determine the reward as the relation prediction results obtained by performing RE using the selected sentences.

Specifically, for a given text path, considering the selected sentences S^h for the head document and S^t for the tail document, the reward R is,

$$R(S^h) = R(S^t) = \lambda_r \hat{y}_r(S^h, S^t), \quad (6)$$

where $\hat{y}_r(S^h, S^t)$ is the prediction score of the true relation r obtained by inputting the selected sentence pairs (S^h, S^t) into the RE module. Here, λ_r is a reward hyperparameter adjusted based on the relation type. Empirically, due to the dataset imbalance problem where there are more ‘no relation’ (N/A) cases than positive cases, we find that it is beneficial to assign a higher value to positive relations compared to ‘no relation’. In our experiments, we set $\lambda_r = 10$ for positive relations and 1 otherwise, as discussed in Section 4.4. This hyperparameter λ can be considered as a prior of Bayesian RL. Additionally, depending on the loss structure, we also consider replacing it with rewards that are proportional to the prediction scores, which is detailed in Section 4.1.

Based on the defined MDP, we update the policy network directly using one of the policy gradient methods, REINFORCE (Williams, 1992):

$$\phi \leftarrow \phi + \alpha R(S_T) \nabla \log \hat{\pi}(S_T), \quad (7)$$

$$\text{where } \log \hat{\pi}(S_T) = \sum_{t=1}^T \log \hat{\pi}(s_{m_t} | S_{t-1}). \quad (8)$$

Here, α is the learning rate. In the implementation, for the sake of time and memory efficiency, we fix the BERT encoder for sentence embedding extraction and only optimize the parameters of the policy network and the LSTM.

We train the REIC module and the RE module simultaneously, as described in Figure 2. The REIC module extracts important sentences from documents in a batch, and these selected sentences are used to predict relations through the RE module. Based on the predicted outputs, we compute the reward and the RE loss. Then, we optimize the REIC module and the RE module using policy gradient and backpropagation, respectively.

4 Experiments

4.1 Experimental Settings

We evaluate the proposed method on the cross-document RE task. We use the CodRED dataset (Yao et al., 2021), which consists of 276 relation types derived from English Wikipedia and Wikidata. Documents in CodRED contain an average of 4,939 tokens. Given our focus on evaluating cross-document reasoning, we conduct experiments in a closed setting where text paths are provided. We adopt End-to-end (Yao et al., 2021) and ECRIM (Wang et al., 2022) for the RE modules. We follow the settings provided by each model. We

RE module		Input construction module	Dev				Test	
Backbone	Structure		AUC	F1	P@500	P@1000	AUC	F1
BERT	End-to-end	Snippet*	53.83	55.17	67.80	54.80	54.47	56.81
		ECRIM-IC	51.38	52.51	66.60	52.10	48.97	52.47
		REIC (Ours)	55.80	56.51	71.20	56.70	54.71	57.38
	ECRIM	Snippet	57.14	57.85	76.00	57.70	58.88	60.79
		ECRIM-IC	62.96	61.43	78.40	61.80	62.17	60.80
		REIC (Ours)	64.50	63.77	78.60	63.80	62.59	61.27
RoBERTa	End-to-end	Snippet	56.47	57.49	72.60	57.10	56.32	56.05
		ECRIM-IC	54.20	55.55	69.60	55.70	54.35	54.05
		REIC (Ours)	58.16	58.74	74.40	58.80	57.71	58.35
	ECRIM	Snippet	63.09	63.27	80.00	63.40	64.31	64.74
		ECRIM-IC	59.36	61.81	79.60	61.90	62.12	63.06
		REIC (Ours)	66.41	63.47	80.20	63.50	65.88	65.02

* ‘Snippet + End-to-end’ shows a significant performance improvement over the results reported in Yao et al. (2021). This improvement is due to rectification of minor errors in the original implementation, as detailed in Appendix B.2.2.

Table 1: Performance comparison with input construction methods on the development (Dev) and test datasets of CodRED. Test set results are obtained from the official CodRED website on CodaLab.

consider both BERT-base and RoBERTa-large (Liu et al., 2019) models as the embedding backbone.

Regarding the REIC module, we use a fixed BERT-base (Devlin et al., 2019) model as the embedding backbone. We employ the AdamW (Loshchilov and Hutter, 2019) optimizer with a learning rate of 3e-3. Sentence embeddings are set to a dimension of 768. We use a 1-layer LSTM with a hidden vector dimension of 512. The policy network comprises a 2-layer feed-forward neural network with a hidden dimension of 512.

We use Snippet (Yao et al., 2021) and ECRIM-IC (Wang et al., 2022), which are accessible implementations, as baselines for comparing the input construction modules. Snippet uses tokens near the target entity as input, while ECRIM-IC conducts heuristic sentence filtering based on the bridge entity. In line with previous work (Yao et al., 2021), we use the AUC, F1, Precision@500 (P@500), and Precision@1000 (P@1000) as evaluation metrics. Further settings are specified in Appendix B.

Reward Function We set the reward function to be proportional to the prediction score of the given relation, but different for each RE module due to the difference in their loss structures. For the End-to-End RE module, we set the reward function by the difference between the given relation score and the highest score excluding the relation:

$$R = \lambda_r (\hat{y}_r - \max_{i \neq r} \hat{y}_i) / \hat{y}_r.$$

We set $\lambda_r = 10$ for $r \neq \text{N/A}$ and $\lambda_r = 1$ for $r = \text{N/A}$, and clip rewards below 0. For the ECRIM RE module, we set the reward function by the difference between the given relation score and threshold θ in the RE loss: $R = \lambda_r (\hat{y}_r - \theta)$, where λ_r is the same as before.

4.2 Quantitative Results

Table 1 presents the performance results of different input construction modules with various RE modules. We run experiments with all combinations of RE module backbones, their structures, and input construction module types. As shown in Table 1, REIC demonstrates the best performance across all scenarios. In particular, the model combining the REIC module with ECRIM RE module using the RoBERTa backbone exhibits state-of-the-art performance on the development set with an AUC of 66.41 and F1 of 63.47, and on the test set with an AUC of 65.88 and F1 of 65.02.

Our proposed approach outperforms previous methods: Snippet + End-to-end (Yao et al., 2021), with an increase of 1.97 in AUC and 1.34 in F1; and ECRIM-IC + ECRIM (Wang et al., 2022), with an increase of 1.54 in AUC and 2.34 in F1, on the development set when using the BERT backbone. In the case of the input construction module of ECRIM, due to its heuristic nature, its performance even decreased compared to the Snippet approach

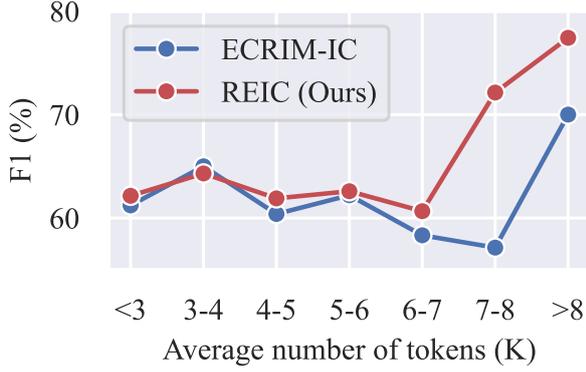


Figure 4: Experimental results on F1s based on the average number of tokens in a bag of text paths with ECRIM (BERT) RE module.

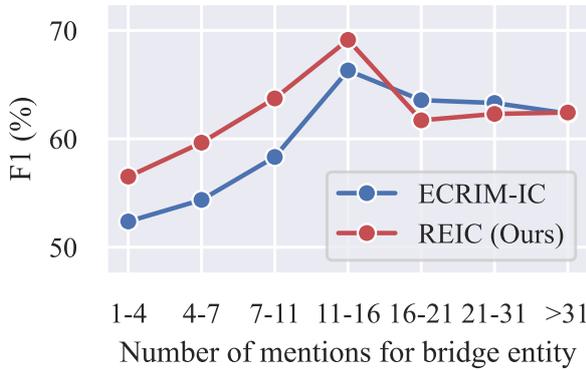


Figure 5: Experimental results on F1s based on the average number of bridge entities in a bag of text paths with ECRIM (BERT) RE module.

on the End-to-end backbone that does not leverage bridge entity information (a decrease of 2.45/2.27 in AUC for BERT/RoBERTa backbones, respectively). Conversely, our method consistently improves performance in all scenarios because the selection module is trained based on the rewards received from each RE module.

4.3 Analysis of Input Construction

Effect Analysis for Document Length To analyze the effect of document length, we partition the CodRED development set based on the average number of tokens in the text paths for each entity pair. Figure 4 presents the F1 score corresponding to the average number of tokens. We observe that our sentence selection method shows comparable performance in most cases, and especially greater effectiveness in cases with very long document lengths. This improvement is attributed to the dispersion of important sentences in longer documents, which allows our learning-based approach to select relevant sentences more effectively.

Method	N/A Bag	Positive Bag
ECRIM-IC	5.01	9.87
REIC	1.58	4.91

Table 2: Average number of bridge entity mentions in selected sentences with ECRIM (BERT) RE module.

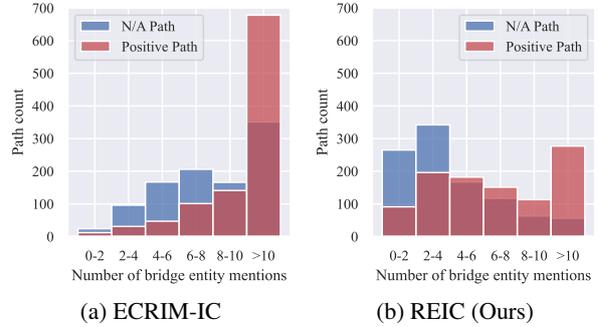


Figure 6: Histogram of the number of bridge entity mentions in sentences selected by (a) ECRIM-IC and (b) REIC for each document among entity pairs with positive relations, with ECRIM (BERT) RE module.

Effect Analysis for Number of Bridge Entities in Text Paths

To analyze the effect of the number of bridge entities in text paths, we partition the CodRED development set based on the number of mentions of bridge entities in the text paths for each entity pair. Figure 5 illustrates the variation in F1 scores according to the average number of mentions of bridge entities. Our model exhibits superior performance to ECRIM-IC, a heuristic construction method based on bridge entities, until the number of mentions reaches 16, which accounts for 64% of the entity pairs in the dataset. However, we observe that the performance of REIC is slightly lower than ECRIM-IC when the number of mentions is very high. These results suggest that while ECRIM-IC struggles to distinguish important sentences when the number of mentions is insufficient; REIC, which learns directly from sentences, is more effective. However, when the number of bridge entities is large enough, the heuristic method based on bridge entities is slightly better. Therefore, we believe that incorporating bridge entity information into the reward function of our method can provide synergistic benefits, which we identify as a topic for future work.

Analysis for Number of Bridge Entities in Selected Sentences

For the analysis of selected sentences, we examine the number of mentions of bridge entities in sentences selected by ECRIM-IC

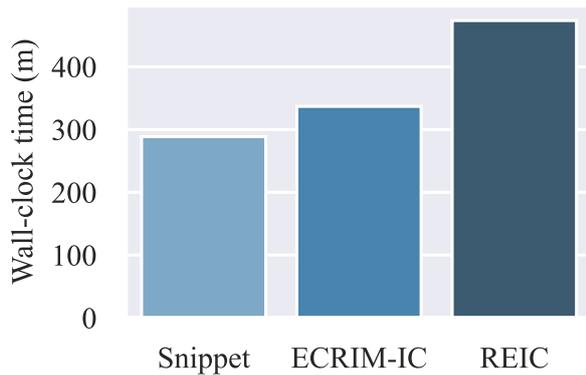


Figure 7: Comparison of training wall-clock time (minutes) per epoch with End-to-end (BERT) RE module.

Method	AUC	F1	P@500	P@1000
Snippet	53.83	55.17	67.80	54.80
ECRIM-IC	51.38	52.51	66.60	52.10
One-step	54.88	56.32	72.00	56.00
$\lambda_r = 1\forall r$	54.65	56.05	70.00	56.50
REIC (Ours)	55.80	56.51	71.20	56.70

Table 3: Ablation studies on the CodRED Dev set with End-to-end (BERT) RE module.

and REIC. Table 2 displays the average number of bridge entity mentions for each sentence selection module based on the relation type of entity pairs. Both modules extract more bridge entities from text paths of entity pairs with positive relations. ECRIM-IC extracts more bridge entities in all cases because it extracts sentences associated with bridge entities. Although our model is trained based on prediction scores without directly considering the number of bridge entities, it still extracts more bridge entities in positive relations. This observation highlights the importance of bridge entities in cross-document RE and suggests that our model effectively learns these bridge entities.

For entity pairs with positive relations, paths can either represent positive relations (positive paths) or not express the relation (N/A paths). We conduct an analysis of the number of bridge entities for each of these path types. Figure 6 plots the number of bridge entity mentions for each sentence selection module in histograms. ECRIM-IC tends to extract more bridge entities for all path types. Conversely, REIC extracts fewer bridge entities in N/A paths. We interpret this phenomenon as a preference to extract fewer unnecessary or ambiguous entities from N/A paths, thereby focusing on extracting meaningful entities to infer relations.

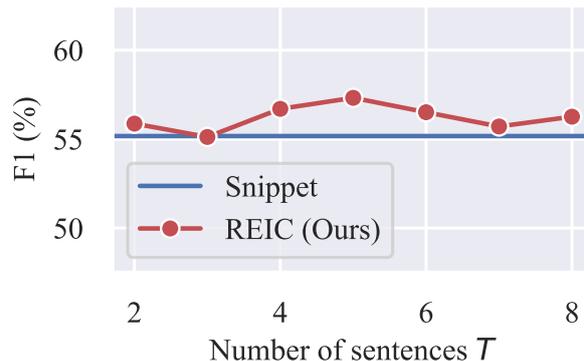


Figure 8: Ablation study for the number of selected sentences with End-to-end (BERT) RE module.

Runtime Analysis Figure 7 compares the training wall-clock time per epoch for each input construction method. Unlike other methods, REIC introduces the sentence selection module, which takes more time. REIC takes about 1.40 times as long as the ECRIM-IC training time. Additionally, we observed that REIC takes about 1.75 times as long as the ECRIM-IC inference time. From these observations, the introduction of the REIC module leads to an increased time consumption. Currently, the selection process operates at the sentence level, involving the extraction of sentence-level embeddings and the computation of selection scores. However, we believe that extending this process to larger structural units, such as paragraphs or paths, could potentially reduce the number of selections and thus the overall time required.

4.4 Ablation Studies

Multi-step vs. One-step Selection We employ a strategy of selecting one sentence at a time over T multi-steps in the sentence selection process, inspired by Man et al. (2022). We explore the option of selecting T sentences all at once in a one-step manner. In this case, we sample a sentence T times based on the selection probability obtained in the first iteration. The evaluation time of the multi-step and one-step selections is almost the same because the policy network and the LSTM are much shallower compared to other networks such as BERT. As shown in the ‘One-step’ row of Table 3, we observe an increase in performance of 0.92 in AUC with multi-step selection. We attribute this improvement to the LSTM providing information about the currently selected sentence, thus helping to select better sentences at each iteration.

(h, t, r)	Selected sentences from REIC
(Chaosium, tabletop role-playing game, product or material produced)	In 1974, “Dungeons & Dragons” brought his interest to role-playing games. He became a full-time staff member at Chaosium. A tabletop role-playing game (or pen-and-paper role-playing game) is a form of role-playing game (RPG) in which the participants describe their characters’ actions through speech. ... Most games follow the pattern established by the first published role-playing game, “Dungeons & Dragons”.

Table 4: Sentences selected by REIC for the text path (‘Sandy Petersen’, ‘Tabletop role-playing game’) containing the triplet (Chaosium, tabletop role-playing game, product or material produced). The first and second paragraphs are from the head and tail document, respectively. **Bold** font denotes the bridge entity, and typewriter font indicates the target entity. A sentence not found in both Snippet and ECRIM-IC is colored in purple.

Reward Hyperparameter λ_r . We set the reward hyperparameter to $\lambda_r = 10$ for the positive relations to ensure that paths with positive relations receive higher rewards. This adjustment aims to address the problem that N/A paths occur about 14 times as often as positive paths, as reported in Yao et al. (2021). To assess the effectiveness of this adjustment, we train a model with λ_r set to 1 for all relations. As shown in the ‘ $\lambda_r = 1 \forall r$ ’ row of Table 3, we observe a performance improvement with REIC when assigning larger rewards to positive relations compared to when rewards are equally distributed. This suggests the necessity of adjusting rewards to account for the imbalance between positive and N/A relations, which is common in RE datasets. Moreover, considering potential imbalances among positive relations, further adjustments to reward hyperparameters may be necessary, which we identify as a direction for future work.

Number of Selected Sentences T In Figure 8, we conduct an ablation study on the hyperparameter T , representing the number of selected sentences. We observe consistently better performance compared to the Snippet across different numbers of sentences. Increasing the number of sentences tends to improve performance up to a certain level, as it provides more information to the RE module. However, including more sentences than necessary may introduce noisy information, potentially affecting the ability to effectively distinguish relations.

4.5 Case Study

Table 4 presents a case study of the key sentences selected by REIC to infer the entity relation, where the sentence missed by the baselines is highlighted in color. This sentence is crucial for inferring the relation between the entity pair. We also include the complete sets of sentences extracted by each input construction method in Tables 9 to 11 of

Appendix C.2. Although ECRIM-IC extracts more bridge entities, REIC selects important sentences necessary to elucidate entity relations.

5 Conclusion

We present REward-based Input Construction (REIC), a learning-based sentence selection module tailored for cross-document RE. We tackle the limitation of document embedding extraction from long documents. REIC selects important sentences based on relational evidence, enabling the RE module to effectively infer relations. Given the unavailability of supervision for evidence sentences, we use RL to train REIC. We show the superiority of our method over heuristic methods. This highlights the potential of learning-based approaches to improve the performance of cross-document RE.

Limitations

Limitations of our method include that the introduction of a network for input construction increases the model runtime by a factor of 1.40 for training and 1.75 for inference in the ECRIM-IC runtime. To alleviate this, we precompute and store sentence embeddings from BERT to avoid repetitive BERT evaluations during training. One possible solution is to apply model compression methods, such as pruning and weight factorization, which have been used in many NLP studies (Lan et al., 2020; Gordon et al., 2020; Wang et al., 2020).

Acknowledgements

This research was supported by Research and Development on Key Supply Network Identification, Threat Analyses, Alternative Recommendation from Natural Language Data (NRF) funded by the Ministry of Education (2021R1A2C200981612).

References

- Zi-Yuan Chen, Chih-Hung Chang, Yi-Pei Chen, Jijnasa Nayak, and Lun-Wei Ku. 2019. [UHop: An unrestricted-hop relation extraction framework for knowledge-based question answering](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 345–356, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jianpeng Cheng and Mirella Lapata. 2016. [Neural summarization by extracting sentences and words](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 484–494, Berlin, Germany. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jun Feng, Minlie Huang, Li Zhao, Yang Yang, and Xiaoyan Zhu. 2018. Reinforcement learning for relation classification from noisy data. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’18/IAAI’18/EAAI’18. AAAI Press.
- Mitchell Gordon, Kevin Duh, and Nicholas Andrews. 2020. [Compressing BERT: Studying the effects of weight pruning on transfer learning](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 143–155, Online. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural computation*, 9:1735–80.
- Quzhe Huang, Shengqi Zhu, Yansong Feng, Yuan Ye, Yuxuan Lai, and Dongyan Zhao. 2021. [Three sentences are all you need: Local path enhanced document relation extraction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 998–1004, Online. Association for Computational Linguistics.
- Amina Kadry and Laura Dietz. 2017. [Open relation extraction for support passage retrieval: Merit and open issues](#). In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’17, page 1149–1152, New York, NY, USA. Association for Computing Machinery.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Xiaoya Li, Fan Yin, Zijun Sun, Xiayu Li, Arianna Yuan, Duo Chai, Mingxin Zhou, and Jiwei Li. 2019. [Entity-relation extraction as multi-turn question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1340–1350, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Keming Lu, I-Hung Hsu, Wenxuan Zhou, Mingyu Derek Ma, and Muhao Chen. 2023. [Multi-hop evidence retrieval for cross-document relation extraction](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10336–10351, Toronto, Canada. Association for Computational Linguistics.
- Youmi Ma, An Wang, and Naoaki Okazaki. 2023. [DREEAM: Guiding attention with evidence for improving document-level relation extraction](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1971–1983, Dubrovnik, Croatia. Association for Computational Linguistics.
- Hieu Man, Nghia Trung Ngo, Linh Ngo Van, and Thien Huu Nguyen. 2022. [Selecting optimal context sentences for event-event relation extraction](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):11058–11066.
- Guoshun Nan, Zhijiang Guo, Ivan Sekulic, and Wei Lu. 2020. [Reasoning with latent structure refinement for document-level relation extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1546–1557, Online. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Ranking sentences for extractive summarization with reinforcement learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1747–1759, New Orleans, Louisiana. Association for Computational Linguistics.

- Pengda Qin, Weiran Xu, and William Yang Wang. 2018. [Robust distant supervision relation extraction via deep reinforcement learning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2137–2147, Melbourne, Australia. Association for Computational Linguistics.
- Jianfeng Qu, Wen Hua, Dantong Ouyang, Xiaofang Zhou, and Ximing Li. 2019. [A fine-grained and noise-aware method for neural relation extraction](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 659–668, New York, NY, USA. Association for Computing Machinery.
- Sunil Kumar Sahu, Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. [Inter-sentence relation extraction with document-level graph convolutional neural network](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4309–4316, Florence, Italy. Association for Computational Linguistics.
- Junyoung Son, Jinsung Kim, Jungwoo Lim, Yoonna Jang, and Heuseok Lim. 2023. [Explore the way: Exploring reasoning path by bridging entities for effective cross-document relation extraction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6755–6761, Singapore. Association for Computational Linguistics.
- Ryuichi Takanobu, Tianyang Zhang, Jiexi Liu, and Minlie Huang. 2019. [A hierarchical framework for relation extraction with reinforcement learning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7072–7079.
- Fengqi Wang, Fei Li, Hao Fei, Jingye Li, Shengqiong Wu, Fangfang Su, Wenxuan Shi, Donghong Ji, and Bo Cai. 2022. [Entity-centered cross-document relation extraction](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9871–9881, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Linlin Wang, Zhu Cao, Gerard de Melo, and Zhiyuan Liu. 2016. [Relation classification via multi-level attention CNNs](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1298–1307, Berlin, Germany. Association for Computational Linguistics.
- Ziheng Wang, Jeremy Wohlwend, and Tao Lei. 2020. [Structured pruning of large language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6151–6162, Online. Association for Computational Linguistics.
- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Machine Learning*, 8(3):229–256.
- Haoran Wu, Xiuyi Chen, Zefa Hu, Jing Shi, Shuang Xu, and Bo Xu. 2023. [Local-to-global causal reasoning for cross-document relation extraction](#). *IEEE/CAA Journal of Automatica Sinica*, 10(7):1608–1621.
- Benfeng Xu, Quan Wang, Yajuan Lyu, Dai Dai, Yongdong Zhang, and Zhendong Mao. 2023. [S2ynRE: Two-stage self-training with synthetic data for low-resource relation extraction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8186–8207, Toronto, Canada. Association for Computational Linguistics.
- Benfeng Xu, Quan Wang, Yajuan Lyu, Yong Zhu, and Zhendong Mao. 2021. [Entity structure within and throughout: Modeling mention dependencies for document-level relation extraction](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14149–14157.
- Tianyu Xu, Wen Hua, Jianfeng Qu, Zhixu Li, Jiajie Xu, An Liu, and Lei Zhao. 2022. [Evidence-aware document-level relation extraction](#). In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, CIKM '22*, page 2311–2320, New York, NY, USA. Association for Computing Machinery.
- Yuan Yao, Jiaju Du, Yankai Lin, Peng Li, Zhiyuan Liu, Jie Zhou, and Maosong Sun. 2021. [CodRED: A cross-document relation extraction dataset for acquiring knowledge in the wild](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4452–4472, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. [DocRED: A large-scale document-level relation extraction dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 764–777, Florence, Italy. Association for Computational Linguistics.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. [Relation classification via convolutional deep neural network](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Shuang Zeng, Runxin Xu, Baobao Chang, and Lei Li. 2020. [Double graph based reasoning for document-level relation extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1630–1640, Online. Association for Computational Linguistics.
- Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023a. [DiffuSum: Generation enhanced extractive summarization with diffusion](#). In *Findings of the Association for Computational Linguistics: ACL 2023*,

pages 13089–13100, Toronto, Canada. Association for Computational Linguistics.

Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023b. [Extractive summarization via ChatGPT for faithful summary generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3270–3278, Singapore. Association for Computational Linguistics.

Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. [Position-aware attention and supervised data improve slot filling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 35–45, Copenhagen, Denmark. Association for Computational Linguistics.

Ming Zhong, Pengfei Liu, Yiran Chen, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2020. [Extractive summarization as text matching](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6197–6208, Online. Association for Computational Linguistics.

Ming Zhong, Pengfei Liu, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2019. [Searching for effective neural extractive summarization: What works and what's next](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1049–1058, Florence, Italy. Association for Computational Linguistics.

Qingyu Zhou, Nan Yang, Furu Wei, Shaohan Huang, Ming Zhou, and Tiejun Zhao. 2018. [Neural document summarization by jointly learning to score and select sentences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–663, Melbourne, Australia. Association for Computational Linguistics.

Hao Zhu, Yankai Lin, Zhiyuan Liu, Jie Fu, Tat-Seng Chua, and Maosong Sun. 2019. [Graph neural networks with generated parameters for relation extraction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1331–1339, Florence, Italy. Association for Computational Linguistics.

A Further Review of Related Work

Early relation extraction models were proposed to predict relationships between entities within a single sentence (Zeng et al., 2014; Wang et al., 2016; Zhang et al., 2017; Zhu et al., 2019). Some research on sentence-level relation extraction introduced a novel convolutional neural network (CNN) architecture to extract sentence-level features from word tokens (Zeng et al., 2014; Wang et al., 2016). Additionally, Zhang et al. (2017) not only presented an LSTM sequence model with entity position-aware attention, but also proposed TACRED, a large supervised dataset for sentence-level RE, serving as a benchmark dataset. GP-GNNs (Zhu et al., 2019) propose a novel graph neural network where its parameters are generated through a propagated message passing module taking sentences as input.

Yao et al. (2019) introduced DocRED, a comprehensive dataset with human-annotated document-level relation extraction, aiming to advance from sentence-level to document-level relation extraction, offering a wealth of relation facts. Several methods for document-level RE utilize document-level graphs constructed from entities or mentions, employing graph neural networks (GNN) for path reasoning to infer relations between entities (Sahu et al., 2019; Nan et al., 2020; Zeng et al., 2020). On the other hand, many document-level RE methods have emerged that do not rely on graph structures. Huang et al. (2021) presents a straightforward method to heuristically select evidence sentences. SSAN (Xu et al., 2021) integrates structural dependencies into self-attention, enhancing attention flow for relation extraction. Recent challenges for document-level RE include high memory usage and limited annotations. DREEAM (Ma et al., 2023) and S2ynRE (Xu et al., 2023) offer a memory-efficient solution and self-training to generate evidence and training data.

B Additional Experimental Settings

B.1 Datasets

The CodRED dataset, sourced from its official Github repository, comprises relation triplets, evidence, and documents, licensed under MIT. CodRED also includes processed Wikidata (CC0 license) and Wikipedia (licensed under CC BY-SA and GFDL). We processed the raw data as per repository instructions, transforming and storing documents in a Redis server for downstream RE methods. In particular, for efficiency, we stored

Split	Triplets	Text paths (mean)
Train	19,461	129,548 (6.6)
Dev	5,568	40,740 (7.3)
Test	5,535	40,524 (7.3)

Table 5: The statistics of CodRED dataset.

BERT embeddings for each sentence of every document in the server environment for operational efficiency. CodRED comprises 276 relations, and its statistics are presented in Table 5.

B.2 Implementation Details

B.2.1 Model Parameters

Table 6 displays the number of parameters for different models. Each row represents a RE module’s backbone, structure, an input construction module and the corresponding number of parameters. For BERT, the parameter count is 108,310,272, while for RoBERTa, it is 355,462,144. Our REIC model introduces a selector module, increasing the parameter count by 3,281,921 for each RE module.

Backbone	RE Module	IC Module	# of Parameters
BERT	End-to-end	Snippet	108,523,285
		REIC (Ours)	111,805,206
RoBERTa	End-to-end	Snippet	355,746,069
		REIC (Ours)	358,028,990
BERT	ECRIM	ECRIM-IC	122,225,509
		REIC (Ours)	125,844,193
RoBERTa	ECRIM	ECRIM-IC	375,111,397
		REIC (Ours)	388,063,055

Table 6: Number of parameters for each model. RE is short for relation extraction and IC is short for input construction.

B.2.2 Code Implementation

During the implementation process, we utilized the official codes of End-to-end¹ (Yao et al., 2021) and ECRIM² (Wang et al., 2022). We trained both the end-to-end and ECRIM models using a learning rate of $3e-5$. The training process involved running 2 epochs for the End-to-end and 10 epochs for the ECRIM. One change is to set the batch size to 4 and the gradient accumulation step to 4 for ECRIM.

End-to-end model underwent the modification during implementation that had a significant impact on performance. In the original code, the entity marking token ‘[UNUSED@]’, where ‘@’

¹<https://github.com/thunlp/CodRED>

²<https://github.com/MakiseKuurisu/ecrim>

Average number of tokens (K)	Number of cases
< 3	743
3-4	1071
4-5	1227
5-6	1013
6-7	655
7-8	382
> 8	477

Table 7: Number of cases in each partition by average number of tokens, which is related to Figure 4.

Number of mentions for bridge entity	Number of cases
1-4	665
4-7	889
7-11	1085
11-16	943
16-21	587
21-31	590
> 31	809

Table 8: Number of cases in each partition by average number of bridge entity, which is related to Figure 5.

is a natural number, was utilized. However, we found that this token was functionally meaningless and replaced it with ‘[unused@]’. Then, the performance of the modification, as shown in the row of ‘Snippet + End-to-end’ in Table 1, exhibited a significant improvement compared to the results reported in Yao et al. (2021).

Our REIC model implementation is also based on the official End-to-end and ECRIM codebases. Additionally, we incorporated the sentence selection module from SCS-EERE³ (Man et al., 2022) as a reference. If the selected T sentences exceed the token limit, we use only the first 512 tokens in the order they were chosen. Our REIC model was trained according to the settings specified for each RE structure, and detailed training configurations for our REIC are provided in Section 4.1. We utilized a single NVIDIA A100 80GB model for training both the baseline and our model. We performed a single run for each of the various RE backbones and structures to produce the results.

B.3 Details of Input Construction Analysis

Tables 7 and 8 provide the number of bags in each partition for each analysis.

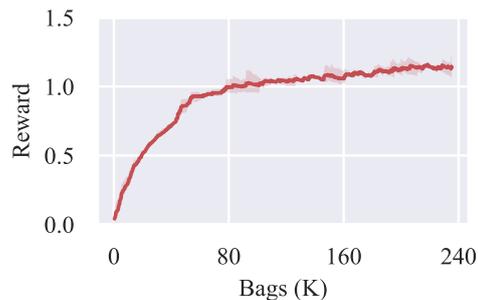


Figure 9: Reward curve during training with End-to-end (BERT) RE module. We use the exponential moving average with factor of 0.99. We ran the experiment three times. Solid line represents the mean, and the shaded region illustrates the minimum and maximum values.

C Additional Experimental Results

C.1 Reward Analysis

Figure 9 shows the selection reward curve during training. As training progresses, the reward gradually increases. After a certain initial step, the reward stabilizes at 1, which is a reward value for correct predictions of the N/A relations. This is probably influenced by the large proportion of N/A relations in the dataset and the model correctly predicts the N/A relations. Subsequently, in later stages, the reward surpasses 1, indicating that the model is increasingly focused on correctly predicting positive relations.

C.2 Case Study

Tables 9 to 11 provide the complete set of sentences selected by each input construction module for the case study analyzed in Section 4.5.

D Potential Risks

Our method may select incorrect or irrelevant sentences, resulting in inaccurate relation extraction. This risk could occur if there are limitations in the training data. In addition, there could be cases of noisy or harmful labels in the training data, which could affect the extraction of harmful sentences.

³<https://github.com/hieumdt/SCS-EERE>

Source	Selected sentences from Snippet
Head	<p>“Sandy Petersen” (born September 16, 1955) is an American game designer. Petersen was born in St. Louis, Missouri and attended University of California, Berkeley, majoring in zoology. He is a well-known fan of H. P. Lovecraft, whose work he first encountered in a World War II Armed Services Edition of “The Dunwich Horror and other Weird Tales” found in his father’s library. In 1974, “Dungeons & Dragons” brought his interest to role-playing games. He became a full-time staff member at Chaosium. His interest for role-playing games and H. P. Lovecraft were fused when he became principal author of Chaosium’s game. “Call of Cthulhu”, published 1981, and many scenarios and background pieces thereafter. While working for Chaosium he co-authored the third edition of “RuneQuest”, for which he also co-wrote the critically acclaimed “Trollpak” and a number of other Gloranthan supplements. He authored several critically acclaimed “RuneQuest” supplements for Avalon Hill and</p>
Tail	<p>A tabletop role-playing game (or pen-and-paper role-playing game) is a form of role-playing game (RPG) in which the participants describe their characters’ actions through speech. Participants determine the actions of their characters based on their characterization, and the actions succeed or fail according to a set formal system of rules and guidelines. Within the rules, players have the freedom to improvise; their choices shape the direction and outcome of the game. Unlike other types of role-playing game, tabletop RPGs are often conducted like radio drama: only the spoken component of a role is acted. This acting is not always literal, and players do not always speak exclusively in-character. Instead, players act out their role by deciding and describing what actions their characters will take within the rules of the game. In most games, a specially designated player called the game master (GM) — also known as the Dungeon Master (DM) in “Dungeons & Dragons”, Referee for all Game Designers’ Workshop games, or Storyteller for the Storytelling System — creates a setting in which each player plays the role of a single character. The GM describes the game world and</p>

Table 9: Example of the selected sentences from Snippet. The first row represents the sentences selected from the head document, and the second row represents the sentences selected from the tail document. Text in typewriter font indicates the target entity, while bold text indicates the bridge entity.

Source	Selected sentences from ECRIM-IC
Head	<p>Lovecraft, whose work he first encountered in a World War II Armed Services Edition of “The Dunwich Horror and other Weird Tales” found in his father’ s library. In 1974, “Dungeons & Dragons” brought his interest to role-playing games. He became a full-time staff member at Chaosium. His interest for role-playing games and H. P. Lovecraft were fused when he became principal author of Chaosium’ s game “Call of Cthulhu”, published 1981,</p>
Tail	<p>A tabletop role - playing game (or pen-and-paper role-playing game) is a form of role-playing game (RPG) in which the participants describe their characters’ actions through speech. Participants determine the actions of their characters based on their characterization, In most games, a specially designated player called the game master (GM) — also known as the Dungeon Master (DM) in “Dungeons & Dragons”, Referee for all Game Designers’ Workshop games, Due to the game’ s success, the term “Dungeons & Dragons” has sometimes been used as a generic term for fantasy role-playing games. TSR undertook legal action to prevent its trademark from becoming generic. The “d20 system”, based on the third edition of “Dungeons & Dragons”, was used in many modern or science fiction game settings such as “Spycraft” and the “Star Wars Roleplaying Game”. Usually a campaign setting is designed for a specific game (such as the “Forgotten Realms” setting for “Dungeons & Dragons”) or a specific genre of game (such as Medieval fantasy, World War II, or outer space / science fiction adventure). horror formed the baseline of the “World of Darkness” and “Call of Cthulhu” while “Spycraft” was based in modern-day spy thriller-oriented settings.</p>

Table 10: Example of the selected sentences from ECRIM-IC. The first row represents the sentences selected from the head document, and the second row represents the sentences selected from the tail document. Text in typewriter font indicates the target entity, while bold text indicates the bridge entity.

Source	Selected sentences from REIC
Head	Petersen was born in St. Louis, Missouri and attended University of California, Berkeley, majoring in zoology. He is a well-known fan of H. P. Lovecraft, whose work he first encountered in a World War II Armed Services Edition of “The Dunwich Horror and other Weird Tales” found in his father’s library. In 1974, “ Dungeons & Dragons ” brought his interest to role-playing games. He became a full-time staff member at Chaosium. His interest for role-playing games and H. P. Lovecraft were fused when he became principal author of Chaosium’s game “ Call of Cthulhu ”, published 1981, and many scenarios and background pieces thereafter. While working for Chaosium he co-authored the third edition of “ RuneQuest ”, for which he also co-wrote the critically acclaimed “Trollpak” and a number of other Gloranthan supplements. and is a frequent guest at conventions where he usually runs a freeform game of his own devising, and / or helps to run someone else’s game. He worked some
Tail	A tabletop role-playing game (or pen-and-paper role-playing game) is a form of role-playing game (RPG) in which the participants describe their characters’ actions through speech. Participants determine the actions of their characters based on their characterization, Most games follow the pattern established by the first published role-playing game, “ Dungeons & Dragons ”. Participants usually conduct the game as a small social gathering. One participant, called the Dungeon Master (DM) in “ Dungeons and Dragons ”, more commonly called the game master or GM, purchases or prepares a set of rules and a fictional setting in which players can act out the roles of their characters. from a single brief session (usually completed in a few hours) to a series of repeated sessions that may continue for years with an evolving cast of players and characters. Play is often episodic and mission-centric, with a series of challenges culminating in a final puzzle or enemy that must be overcome. Gygax expected to sell about 50, 000 copies total to a strictly hobbyist market. After establishing itself in boutique stores, it developed a strong, lasting fan base

Table 11: Example of the selected sentences from REIC (Ours). The first row represents the sentences selected from the head document, and the second row represents the sentences selected from the tail document. Text in typewriter font indicates the target entity, while bold text indicates the bridge entity.