

# Dealing with negative samples with multi-task learning on span-based joint entity-relation extraction

Jiamin Lu and Chenguang Xue

Key Laboratory of Water Big Data Technology of Ministry of Water Resources,  
Hohai University, Nanjing, China,  
211307050013@hhu.edu.cn

**Abstract.** Recent span-based joint extraction models have demonstrated significant advantages in both entity recognition and relation extraction. These models treat text spans as candidate entities, and span pairs as candidate relationship tuples, achieving state-of-the-art results on datasets like ADE. However, these models encounter a significant number of non-entity spans or irrelevant span pairs during the tasks, impairing model performance significantly. To address this issue, this paper introduces a span-based multitask entity-relation joint extraction model. This approach employs the multitask learning to alleviate the impact of negative samples on entity and relation classifiers. Additionally, we leverage the Intersection over Union (IoU) concept to introduce the positional information into the entity classifier, achieving a span boundary detection. Furthermore, by incorporating the entity Logits predicted by the entity classifier into the embedded representation of entity pairs, the semantic input for the relation classifier is enriched. Experimental results demonstrate that our proposed SpERT.MT model can effectively mitigate the adverse effects of excessive negative samples on the model performance. Furthermore, the model demonstrated commendable F1 scores of 73.61%, 53.72%, and 83.72% on three widely employed public datasets, namely CoNLL04, SciERC, and ADE, respectively.

**Keywords:** Natural language processing, Joint entity and relation extraction, Span-based model, Multi-task learning, Negative samples

## 1 Introduction

The task of joint entity and relation extraction involves extracting the entities and their semantic relationships from unstructured textual data. This task forms the foundational basis for numerous downstream NLP tasks such as sentiment analysis and knowledge graph construction. Due to the limitations of traditional pipeline methods that overlook interactions between Named Entity Recognition (NER) and Relation Extraction (RE) tasks and suffer from error propagation, recent research has explored approaches for joint extraction [1][8][12]. Based on the classification strategy employed for extraction, existing models can be categorized into two types: sequence labeling-based models and span-based models. The former rely on sequence labeling mechanisms [9][10] to perform label-level classification. The latter, on the other hand, adhere to the span-based paradigm [1][11][12] and engage in span-level classification.

Since the named entities contain nested entities, the traditional sequence labeling models face limitations on assigning multiple labels to a single token simultaneously. On the other hand, span-based models treat consecutive tokens as spans. For instance, in Figure 1, "Hypersensitivity," "Hypersensitivity to," and "Hypersensitivity to aspirin" represent spans of different lengths. Span-based models enumerate all possible spans and classify them into the correct categories, thereby eliminating the need for extensive feature engineering and effectively addressing the nested entity issue. However, these models encounter numerous challenges.


As the span strategy involves an enumeration process where each span (pair) needs to be classified, these models face an extremely imbalanced dataset. The dataset consists of significantly more non-entity spans (pairs) compared to entity spans (pairs). In the following text, we will refer to entity spans (pairs) as positive samples and non-entity spans (pairs) as negative samples. For example, in Figure 1, there are two entity spans, E1 and E2 (positive entity samples), and one relation R (positive relation sample). Instead, the span strategy continuously expands a sliding window to select all possible consecutive tokens, producing 64 negative entity samples. The pairwise combination of spans generates 4355 negative relation samples. Hence, it is evident that there exists a substantial disparity in the ratio between positive and negative samples. Previous span-based models did not take into account the adverse impact of excessive negative samples on model performance.

In the context of numerous negative samples, prior research in joint entity-relation extraction has employed the inclusion of a special entity class labeled as "NA" to classify negative samples. However, this approach disregards the recognition steps of Named Entity Recognition (NER) and Relation Extraction (RE) tasks, thus affecting the overall performance of the model. Nguyen et al. [19] proposed a recognition framework that breaks down RE into identification and classification stages, validating the efficacy of this approach for handling irrelevant entity pairs in RE. Our model draws inspiration from Nguyen et al.'s ideas, applying a multi-task framework to joint extraction and conducting a two-stage joint training for NER and RE.

Another challenge posed by the span-based method is the issue of hard negative samples. For instance, as shown in Figure 1, adopting the span strategy would lead to generating negative samples like "be manifested" and "a systemic anaphylactoid reaction". The former does not overlap with any entity, while the latter overlaps with the entity "systemic anaphylactoid reaction". The second type of span is termed hard negative samples. While the model can easily distinguish the first kind of negative sample, distinguishing the second kind is difficult. This challenge has been acknowledged by [15][16][13], who introduced token-level boundary detection and boundary regression to address this problem. In our approach, we imply the Intersection over Union (IoU) within the multi-task framework to analyze hard negative samples quantitatively. In Nguyen et al.'s work [19], they embedded label information from the NER task into the representation of entity pairs to enrich the information. Similarly, we found that the Logits from NER, based on the span-based strategy, also contain substantial semantic information. Consequently, in the RE stage, we embed the NER Logits information into the representation of span pairs to enhance joint extraction. In conclusion, our contributions are as follows:

- We apply a multi-task learning approach to entity-relation joint extraction. For both entity recognition and relation extraction tasks, we employ a multi-task framework to separately identify and classify candidate entities/entity pairs, aiming to mitigate the impact of negative samples on model performance.
- We introduce the Intersection over Union (IoU) computation method into the entity recognition subtask. By leveraging IoU, we incorporate positional information to address the challenge of hard negative samples.
- We merge the entity Logits generated from the entity recognition task with the embeddings of span pairs, thereby enriching the input information for relation extraction.

The rest of this paper is organized as follows. Section 2 presents a review of relevant literature on joint extraction. Section 3 provides a comprehensive description of the proposed SpERT.MT model. Subsequently, Section 4 outlines the experiments, encompassing experimental setups, results, and a detailed analysis. Finally, in Section 5, we summarize our conclusions and identify potential directions for future research.



**Fig. 1.** An example from the ADE dataset, wherein the entities E1 and E2 are delineated by shaded annotations. Employing the span-based strategy, these entities correspond to spans of length 1 and 3, respectively. The arrows symbolize the relationship R between E1 and E2.”

## 2 Related work

Joint entity and relation extraction encompass two subtasks: Named Entity Recognition (NER) and Relation Extraction (RE). Due to the error propagation and neglect of interactions between these subtasks, traditional pipeline methods, in recent years, have led to a growing interest in joint extraction [8][9][10][11][12]. Previous research can be classified into two strategies: token tagging and span classification.

The token tagging strategy treats extraction as a sequence labeling task [9][10], where each token is labeled according to the BIO (or its variant BILOU) scheme. Miwa and Sasaki [21] approached joint entity and relation extraction as a table filling problem, employing a history-based beam search to find optimal table filling solutions. Ma et al. [10] also treated joint extraction as a table filling problem. In contrast to Miwa and Sasaki [21], they incorporated the stacked model concept proposed by Miwa and Bansal [20], stacking CNN on top of pre-trained language models to capture local dependencies and predict cell labels. Li et al. [9] introduced a two-stage tagging scheme, labeling head entities and multiple tail entities separately for specific relationships. However, due to the limitation that a token cannot have multiple labels, these models cannot handle overlapping entities and relations. To address this, Wang et al. [23] and Wei et al. [13] adopted multi-turn tagging, significantly increasing the time complexity for both training and decoding processes.

The span classification strategy involves naming entity classification for spans (sub-sequences of text) and relation extraction for all possible span pairs. Models adopting the span classification strategy usually perform exhaustive search over all spans, allowing them to detect nested entities. Luan et al. [2] introduced the first span-based model and conducted beam search over the hypothesis space. Their subsequent model, DyGIE [22], employed dynamically constructed span graphs to share span representations. Wadden et al. [7] replaced the BiLSTM encoder with BERT, showcasing the advantages of using pre-trained models as encoders. Eberts and Ulges [1] encoded span representations using pre-trained models and utilized a strong negative sampling strategy to enhance the training process. Ye et al. [12] introduced the concept of pack-and-float tags, combining token tagging and span classification strategies to focus on relationships between spans/pairs. Tang et al. [11] introduced boundary regression models to learn the offset between spans and the gold entities.

However, the aforementioned studies did not focus on the data imbalance issue caused by the span classification strategy. In contrast to their research, we address the data imbalance issue through a multi-task framework. Recently, numerous experiments have employed multi-task learning methods to train NER and RE models. Zheng et al. [15], Cao et al. [16], and Tan et al. [17] used multi-task learning to capture the dependencies between entity boundaries and their classification labels. Fei et al. [18] proposed treating each unique entity type as a separate task to facilitate information exchange between biomedical NER tasks using multi-task learning. Similarly, Nguyen et al. [19] decomposed

the RE task into two subtasks and employed multi-task learning to address the relation extraction issue in imbalanced corpora. The mentioned works collectively illustrate the favorable performance of multi-task learning methods in both NER and RE tasks.

The SpERT model, proposed by Eberts et al. [1], is a span-based joint entity-relation extraction model. This model employs shallow classifiers for both NER and RE tasks and does not explicitly address the impact of data imbalance on the model. Our model is built upon the foundation of the SpERT model, and it utilizes a multi-task learning approach to individually optimize the NER and RE tasks. This is done to tackle the issue of training sample imbalance caused by the span classification strategy.

### 3 Proposed Method

The model proposed in this study, as illustrated in Figure 2, comprises four distinct modules: Vector Representation, Span Representation, Entity Recognition, and Relation Extraction. Given an input sentence, fine-tuned BERT is employed to generate token and contextual vector representations. Subsequently, the Span Representation layer enumerates all feasible spans, combining the outputs from the vector representation layer and applying width constraints to characterize these spans. Ultimately, the joint entity-relation extraction task is conducted based on the span representations. As this paper integrates a multi-task framework for both entity recognition and relation extraction, subsequent sections, apart from describing the four-layer structure of the SpERT model, will encompass the elucidation of the multi-task framework in Section 3.3, depiction of joint training loss in Section 3.6, and explication of the prediction methodology in Section 3.7.

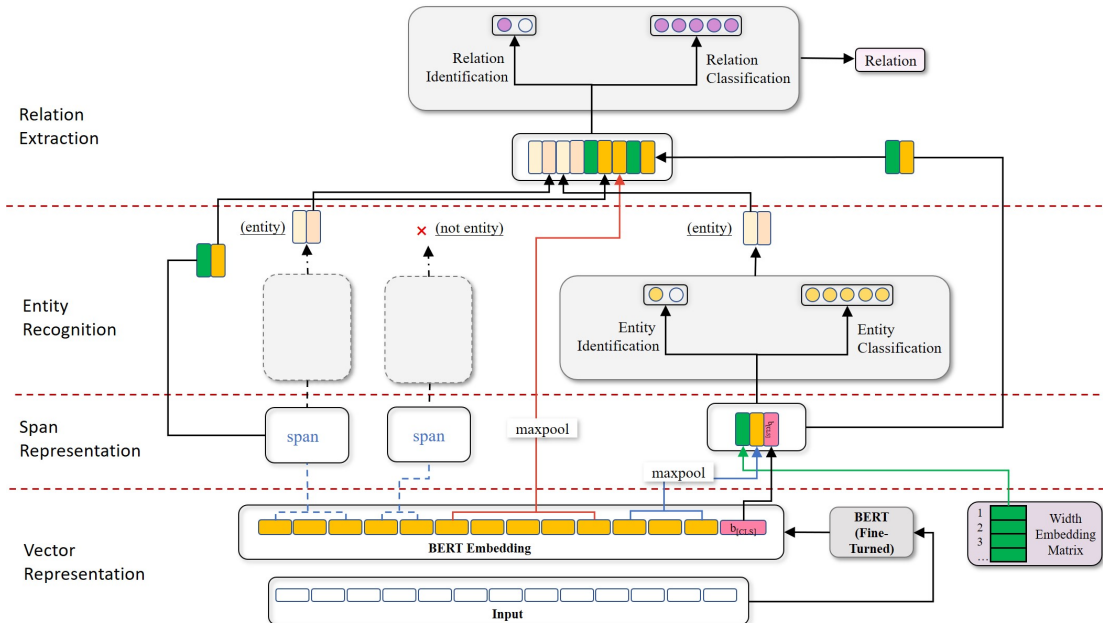


Fig. 2. Architecture of our proposed model, SpERT.MT.

### 3.1 Vector Representation

The Vector Representation layer employs fine-tuned pre-trained BERT to extract contextual information. As illustrated in the figure, the input sentence undergoes Byte Pair Encoding (BPE) and is split into a sequence of  $n$  tokens. BPE encoding represents uncommon words (e.g., "pretrain") with common subwords (e.g., "pre" and "train"), thus constraining the vocabulary and effectively handling out-of-vocabulary (OOV) and rare words. The BPE tokens are propagated through the internals of BERT, resulting in an embedded sequence  $E$  of length  $n+1$ :  $E = (e_{[CLS]}, e_1, e_2, \dots, e_n)$ , where  $e_{[CLS]}$  represents the embedded representation of the entire sentence context, and each embedding vector  $e_i \in \mathbb{R}^{d_1}$  (with  $d_1$  denoting the embedding dimension).

### 3.2 Span Representation

The Span Representation layer utilizes the embeddings  $e$  of the tokens produced by BERT to represent spans. As illustrated in Figure 2, in order to maintain a consistent representation dimension for different spans  $s := (e_i, e_{i+1}, \dots, e_j)$ , we apply a fusion process to  $s$ . We have chosen to employ max pooling for the fusion process:

$$\mathbf{v}(s) = \text{maxpool}(\mathbf{e}_i, \mathbf{e}_{i+1}, \dots, \mathbf{e}_j) \in \mathbb{R}^{d_1} \quad (1)$$

Based on prior knowledge, it is unlikely that longer spans effectively represent meaningful entities. Therefore, we set a maximum span length  $L$  for the model, primarily to reduce complexity. In addition, we incorporate width information  $w$  ( $0 \leq w \leq L$ ) into the fused span representation  $s$  using a width embedding matrix. This approach is guided by the understanding that spans of varying widths contribute distinctively to the representation:

$$\mathbf{c}(s) = \mathbf{v}(s_{ij}) \circ \mathbf{w}_k \in \mathbb{R}^{d_1+d_2} \quad (2)$$

Here,  $\circ$  denotes the concatenation of vectors, and  $w_k$  is derived from the width embedding matrix obtained through training the model.

The categories of entities within a sentence are likely to be correlated with the overall contextual information of the sentence. For instance, contextual keywords like "passing through" or "flying to" can strongly indicate the entity class "location". Therefore, we augment the embedding information of the context [CLS] to the span representation as well:

$$\mathbf{x}(s) = \mathbf{c}(s) \circ \mathbf{e}_{[CLS]} \in \mathbb{R}^{2d_1+d_2} \quad (3)$$

### 3.3 Multi-task Framework

**Identification Task** For the identification task, it is employed to determine whether the current entity or entity pair is indeed an entity or whether a relationship exists. This is a standard binary classification task, thus we opt to use binary cross-entropy as the loss function, as defined below:

$$\mathcal{H}(p, q) = -q \log(p) - (1 - q) \log(1 - p) \quad (4)$$

Here,  $q \in \{0, 1\}$  designates the true class of the span or span pair (0 represents negative sample, 1 represents positive sample), and  $p \in [0, 1]$  is the estimated probability by the model for the label  $q = 1$  class.

Therefore, the loss for the identification task is computed as follows:

$$\text{Loss}_1 = \Sigma \mathcal{H}(p, q) \quad (5)$$

**Classification Task** To enhance diversity in the loss landscape of multi-task learning and prevent the two tasks from degenerating into a single task, we opt for the pairwise ranking loss function (as proposed by Santos et al.[4]) as the loss function for classification tasks. This loss function shares the same task objective as cross-entropy loss but possesses a distinct optimization direction.

Given the span (pair) representation  $s$  and the set of entity or relation classes  $C$ , we compute the score for class label  $c \in C$  using the dot product:

$$\text{score}_c = s^T [W^{\text{classes}}]_c \quad (6)$$

Where  $W^{\text{classes}}$  is the matrix to be learned, and the number of columns hit is the number of classes.  $W_c^{\text{classes}}$  is the column vector corresponding to class  $c$ , with dimensions equal to the dimensions of  $s$ .

For a span (pair), where  $y^+$  is the correct label and  $y^-$  is not, the  $\text{scores}_{y^+}$  and  $\text{scores}_{y^-}$  are defined as the scores for  $y^+$  and  $y^-$ , respectively. Then, the rank loss can be calculated as follows:

$$L^+ = \log(1 + \exp(\gamma(m^+ - \text{score}_{y^+}))) \quad (7)$$

$$L^- = \log(1 + \exp(\gamma(m^- + \text{score}_{y^-}))) \quad (8)$$

Where  $\gamma$  is a scaling factor,  $m^+$  and  $m^-$  are margins. For non-entity classes or irrelevant entity pairs, only  $L^-$  is computed to penalize incorrect predictions. In the training step, the ranking loss selects only the highest scoring incorrect class among all error classes. Then, we optimize the pairwise ranking loss:

$$\text{Loss}_2 = \sum (L^+ + L^-) \quad (9)$$

Finally, the loss function of our multi-task framework is formulated as:

$$\text{Loss} = \alpha \text{Loss}_1 + \beta \text{Loss}_2 \quad (10)$$

### 3.4 Entity Recognition

As shown in Figure 2, the entity recognition layer directly utilizes the output from the span representation layer as input, and utilizes the multi-task framework to extract entity categories from the spans. In this section, adjustments are made based on the multi-task framework described in Section 3.3.

In the identification task of the entity recognition multi-task framework, we introduce the positional information between spans using the Intersection over Union (IoU) metric. Formally, the IoU calculation is as follows:

$$\text{IoU}(s(i_1, j_1), s(i_2, j_2)) = \begin{cases} \frac{\min(j_1, j_2) - \max(i_1, i_2) + 1}{\max(j_1, j_2) - \min(i_1, i_2) + 1}, & \min(j_1, j_2) \geq \max(i_1, i_2) \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

Where  $s(i,j)$  represents the span with left and right indices  $i$  and  $j$  in the sentence.

We compute and utilize the maximum IoU (Intersection over Union) between the span  $s$  and all entity spans in the sentence, defined as  $\text{ENIoU}(s)$ . Its calculation formula is as follows:

$$\text{ENIoU}(s) = \max(\text{IoU}(s, \text{en}), \text{en} \in E) \quad (12)$$

Where  $E$  represents the set composed of all entity spans in the sentence.

Subsequently, ENIoU is employed as a scaling factor to replace the loss function of the Identification task in the multi-task framework with dynamically scaled cross-entropy loss. The loss function is as follows:

$$\mathcal{H}(p, q) = -q(1 - \delta) \log(p) - \delta(1 - q)(1 + \text{ENIoU}) \log(1 - p) \quad (13)$$

Where  $\delta$  and ENIoU determine the scaling factor.  $\delta \in [0, 1]$  is a balancing factor to address class imbalance, and we set  $\delta < 0.5$  to emphasize the fewer positive samples. For negative samples, a higher ENIoU results in a larger loss, thereby focusing the model more on hard negative samples during training.

For the other parts of the multi-task framework, we do not make any further modifications in the entity recognition stage, and we denote the loss in the entity recognition stage as L1.

### 3.5 Relation Extraction

Unlike the direct use of span representation layer outputs as input in the entity recognition stage, the relation extraction stage makes some adjustments to the input information. For those spans that are classified as "NA" during the entity recognition multi-task stage, they will be filtered out. We only consider candidate span pairs that have been preliminarily identified as entities. As shown in Figure 2, we choose the local information between candidate span pairs as their context and obtain context representations by utilizing max-pooling to fuse their BERT embeddings.

$$\mathbf{v}(s_1, s_2) = \text{maxpool}(\mathbf{e}_i, \dots, \mathbf{e}_j) \in \mathbb{R}^{d_1+d_2} \quad (14)$$

For adjacent candidate entity pairs, we set  $\mathbf{v}(s_1, s_2) = 0$ . Additionally, we observe that the Logits from the entity recognition stage may contain useful semantic information. Therefore, we further embed the Logits from the entity recognition stage to enrich the input information. Ultimately, the multi-task input representation of the relation extraction layer is as follows:

$$\mathbf{r}(s_1, s_2) = \mathbf{c}(s_1) \circ \mathbf{v}(s_1, s_2) \circ \mathbf{c}(s_2) \circ \mathbf{p}(s_1) \circ \mathbf{p}(s_2) \in \mathbb{R}^{3d_1+2d_2+2d_3} \quad (15)$$

Where  $\mathbf{p}(s)$  represents the Logits from the entity recognition stage. Considering that the relationship between  $s_1$  and  $s_2$  is typically not symmetric, therefore,

$$\mathbf{r}(s_2, s_1) = \mathbf{c}(s_2) \circ \mathbf{v}(s_1, s_2) \circ \mathbf{c}(s_1) \circ \mathbf{p}(s_2) \circ \mathbf{p}(s_1) \in \mathbb{R}^{3d_1+2d_2+2d_3} \quad (16)$$

For the multi-task framework used in the relation extraction stage, it remains consistent with the approach described in Section 3.3. We denote the loss for the relation extraction stage as L2.

### 3.6 Joint Training Loss

We learn the parameters of width embedding  $w$  for spans (as shown in Figure 2) and the parameters of the entity/relation multi-task classifiers while fine-tuning BERT in this process. Due to our joint training approach for entity recognition and relation extraction, the loss is defined as follows:

$$\mathbf{L} = L_1 + L_2 \quad (17)$$

### 3.7 Prediction

At the prediction stage, to avoid error propagation, only the class scores obtained from the multi-task framework, denoted as  $score_c$ , are utilized, while the binary identification tasks are solely employed for optimizing network parameters.

Given a span(pair), the computation of the predicted class probability  $P$  is as follows:

$$P = \begin{cases} \operatorname{argmax}_c(\operatorname{score}_c), & \max(\operatorname{score}_c) \geq \theta \\ \text{NA}, & \max(\operatorname{score}_c) < \theta \end{cases} \quad (18)$$

Where  $\theta$  is a hyperparameter threshold. If the score for each class is below  $\theta$ , the span(pair) is predicted as non-entity or irrelevant relationship. Otherwise, the span(pair) is assigned to the class with the highest score.

## 4 Experiments

### 4.1 Datasets

- CoNLL04: The CoNLL04 dataset [5] consists of sentences with annotated named entities and relationships extracted from news articles. The dataset comprises four entity types (People, Location, Organization, Other) and five relationship types (Live-In, Work-For, Kill, Located-In, Organization-Based-In). We adopt the split strategy proposed by Gupta et al. [6] for training (1153 sentences) and testing (288 sentences) the model. Additionally, we allocate 20% of the training set for hyperparameter tuning.
- SciERC: The SciERC dataset [2] is derived from abstracts of 500 artificial intelligence research papers. It comprises six scientific entity types (Task, Method, Metric, Material, Other-Scientific-Term, Generic) and seven relationship types (Compare, Conjunction, Evaluate-For, Used-For, Feature Of, Part-Of, Hyponym-Of), totaling 2687 sentences. We adopt the same split strategy as described in [2] for training (1861 sentences), development (275 sentences), and testing (551 sentences) on the SciERC dataset.
- ADE Dataset: The ADE dataset [3] comprises 4272 sentences and 6821 relationships extracted from medical reports that describe adverse reactions caused by drug usage. The dataset encompasses two entity types (Adverse-Effect, Drug) and a single relationship type (Adverse-Effect). Due to the absence of official splits, we follow the approach of previous studies and conduct a 10-fold cross-validation for evaluation.

### 4.2 Implementation

We employed BERT as the encoder for the CoNLL04 dataset, while for ADE and SciERC, we utilized BioBERT and SciBERT with domain-specific features, respectively. As SpERT [1] served as our experimental baseline, we adopted some of its hyperparameter configurations. Specifically, we initialized the classifier weights with normally distributed random numbers ( $\mu=0$ ,  $\sigma=0.02$ ). The Adam optimizer was employed for model training, featuring linear warm-up and linear decay, with a peak learning rate set at  $5e^{-5}$ . The batch size was set to 2, and the width embedding dimension of matrix  $W$  was 25.

For the hyperparameters in the multi-task framework, we utilized the same settings for entity recognition and relation extraction in the pairwise ranking loss:  $m^+$  was set to 2.5,  $m^-$  to 0.5, and  $\gamma$  to 2. For the dynamic scaling of the cross-entropy loss in the identification phase of entity recognition,  $\sigma$  was set to 0.25 ( $\sigma \downarrow 0.5$ ). In the prediction phase, the threshold  $\theta$  was set to 0.85.



Across both entity recognition and relation extraction, we set the number of negative samples to 120 and performed training for each dataset over 20 epochs. Importantly, we employed the same set of hyperparameters for all three datasets, refraining from dataset-specific adjustments.

### 4.3 Evaluation Metrics

In the entity recognition phase, an entity is considered correct if the predicted entity type and span match the given textual span. For relation extraction, we adopted an evaluation method consistent with that proposed in [1]. Specifically, a relation is deemed correct if the relationship type between the provided span pairs and the entities represented by the spans are accurate. This evaluation criterion, in contrast to strict relation extraction, does not require verifying the correctness of the entity types associated with the spans, thus alleviating the need to focus on the results of Named Entity Recognition NER.

We evaluate the model performance on different datasets using various metrics including precision (both micro and macro), recall, and F1 score.

### 4.4 Overall Results

The table 1 presents a performance comparison between our model and several of the previous state-of-the-art entity-relation joint extraction models on the test datasets. Due to substantial variability in our model’s measurements, as also indicated in the study by Taillé et al. [14], we report the average of 10 runs for each dataset except ADE. Regarding the choice of precision evaluation, different standards are applied to different datasets: for the CoNLL04 and SciERC datasets, we use micro-averaged F1 score; for the ADE dataset, we employ macro-averaged F1 score. Additionally, for the ADE dataset, the reported results account for experiments allowing overlapping entities.

From the table 1, it is evident that our model demonstrates significant performance improvements compared to the baseline model SpERT. Moreover, our model achieves competitive results when compared to the state-of-the-art models. Particularly noteworthy is the remarkable performance achieved on the SciERC dataset, where our model attains the state-of-the-art experimental results. Specifically:

- CoNLL04: In the NER task, our model outperforms the SpERT baseline by an increase of 1.76% in F1 score. In the RE task, our model achieves a 2.14% higher F1 score compared to SpERT. When compared to state-of-the-art models, our model exhibits a notable advantage in NER due to the incorporation of IoU-enhanced span boundary detection. However, in the RE task, while we enhance input semantics, we did not specifically optimize for the RE task, leading to a noticeable gap from the state-of-the-art performance.
- SciERC: On the SciERC dataset, our model achieves state-of-the-art performance with F1 scores of 73.22% in NER and 53.72% in RE tasks. This success can be attributed to our choice of using the SciBERT model, which is specifically tailored for domain-specific encoding. Compared to several other state-of-the-art models, our experimental results surpass the state-of-the-art in all metrics except precision in the RE task, where we fall slightly behind.
- ADE: Our model also exhibits impressive performance on the ADE dataset. Compared to the SpERT baseline, we achieve improvements of 2.46% and 4.88% in NER and RE F1 scores, respectively. Under the condition of excluding the experimental results of Tang et al., our model achieves state-of-the-art performance. The data in the table

illustrates that within this cohort of state-of-the-art models, our model excels in NER, surpassing RE performance, and particularly excels in NER recall, achieving state-of-the-art results.

**Table 1.** Comparison between existing methods and our proposed SpERT.MT model on the CoNLL04, SciERC, and ADE datasets. The symbols  $\Delta$  and  $\blacktriangle$  represent evaluation using micro and macro average F1 values, respectively.

Dataset	Model	NER			RE		
		Precision	Recall	F1	Precision	Recall	F1
CoNLL04 $\Delta$	SpERT(2020)	88.25	89.64	88.94	73.04	70.00	71.47
	ECA Net(2021)	<b>91.14</b>	80.80	86.79	<b>82.42</b>	73.34	<b>77.55</b>
	TabLERT-CNN(2022)	-	-	90.50	-	-	73.20
	Tang et al.(2022)	90.80	88.80	89.80	76.60	<b>74.40</b>	75.50
	SpERT.MT(Ours)	90.13	<b>91.27</b>	<b>90.70</b>	75.36	71.94	73.61
SciERC $\Delta$	SpERT(2020)	70.87	69.79	70.33	53.40	48.54	50.84
	Tang et al.(2022)	62.40	67.10	64.70	<b>56.60</b>	48.20	52.10
	PL-Marker(2022)	-	-	69.90	-	-	53.20
	SpERT.MT(Ours)	<b>71.21</b>	<b>75.35</b>	<b>73.22</b>	53.63	<b>53.81</b>	<b>53.72</b>
ADE $\blacktriangle$	SpERT(2020)	88.99	89.59	89.28	77.77	79.96	78.84
	KECI(2021)	-	-	90.67	-	-	81.74
	ECA Net(2021)	88.52	81.11	84.63	82.28	77.17	79.62
	TableERT-CNN(2022)	-	-	89.70	-	-	80.50
	Tang et al.(2022)	<b>93.60</b>	91.50	<b>92.50</b>	<b>85.40</b>	85.00	<b>85.20</b>
	SpERT.MT(Ours)	91.35	<b>92.77</b>	92.05	81.06	<b>86.55</b>	83.72

## 4.5 Ablation Study

We conducted ablation experiments on the SciERC dataset, and the data in the table validates the effectiveness of our Multi-Task Learning (MTL) joint learning approach, the incorporation of IoU-enhanced dynamic scaling Cross-Entropy Loss (EIL), and the utilization of entity Logits. Each category of data represents the average of three experimental results. Observing the results, we note that our Multi-Task Entity-Relation Joint Extraction approach improves the F1 scores for entities and relations by 1.07% and 0.83%, respectively. This suggests that transitioning from the original single-step multi-class approach to the binary and multi-class multi-task approach contributes to enhanced model performance. Removing the EIL from our model results in a decrease of 1.01% and 0.61% in F1 scores for entities and relations, respectively. This indicates that introducing the IoU-enhanced dynamic scaling Cross-Entropy Loss during the entity recognition phase significantly benefits entity identification compared to relation extraction. In the relation extraction phase, enriching the embedding representation with entity Logits leads to F1 score improvements of 0.94% and 1.23% for entities and relations, respectively. This highlights that enhancements made during the relation extraction phase in the context of joint training for entity and relation extraction also contribute to improved entity recognition performance.

**Table 2.** Ablation study of SpERT.MT on SciERC.

Model	NER			RE		
	Precision	Recall	F1	Precision	Recall	F1
<b>SpERT.MT(SciBERT)</b>	<b>71.21</b>	<b>75.35</b>	<b>73.22</b>	<b>53.63</b>	<b>53.81</b>	<b>53.72</b>
<b>-MTL</b>	70.51	73.87	72.15	52.99	52.79	52.89
<b>-Entity Logits</b>	70.59	74.06	72.28	52.86	52.12	52.49
<b>-EIL</b>	70.86	73.62	72.21	53.13	53.09	53.11

#### 4.6 Negative Sampling

In our study, the hyperparameters of the number of negative samples and the maximum span length ( $W$ ) can influence the selection of negative samples during training. For the maximum span length ( $W$ ), we considered dataset characteristics and the hyperparameters of the SpERT model. Accordingly, we set  $W$  to 10. The focus of this experiment lies in investigating model performance (expressed through F1 scores) under scenarios with excessive negative samples in the context of a multi-task framework. We continually adjusted the hyperparameters for the number of negative samples for entities and relations to observe model performance. Taking cues from the SpERT model, we also considered the average sentence length in the dataset, setting a limit of 200 for the experiment’s maximum value. For sentences with negative sample counts exceeding the maximum limit, we randomly sampled from all samples.



**Fig. 3.** The Variation of F1 Scores for Entity and Relation Classification between Our Model and SpERT Model under Different Negative Sample Quantities.

Figure 3 illustrates the F1 scores for entity recognition and relation extraction of our model and the SpERT model on the SciERC dataset with varying numbers of negative samples. As observed in the graph, when negative samples are scarce, model performance is poor. With increasing negative sample counts, performance improves significantly, in line with the conclusion from [1] that negative samples aid model training. Performance stabilizes as negative sample counts rise, yet when the count reaches 100-120, the model’s performance generally peaks. Beyond this range, increasing negative samples can cause

performance to decline, suggesting the need for a balance between an adequate number of negative samples and a limit on their maximum value. Our model outperforms the SpERT model when a certain number of negative samples are present. The peak performance of our model occurs later than that of SpERT when a sufficient number of negative samples are available, implying better performance for our model in multi-negative sample scenarios. Additionally, our model exhibits a gentler decline in performance gradient with further increases in negative sample counts compared to the SpERT model. These results indicate that the multi-task training framework is advantageous in scenarios with excessive negative samples leading to data imbalance. Based on the experimental data, we ultimately chose a maximum negative sample limit of 120 for both entities and relations.

## 5 CONCLUSION

This paper introduces a model, SpERT.MT, based on fine-tuning pretrained BERT and utilizing a span-based strategy for joint entity-relation extraction. To address the issue of excessive negative samples resulting from span-based strategies, we employ a multi-task framework for both entity recognition and relation extraction tasks. Considering the specific characteristics of entity recognition and relation extraction, we adjusted the loss functions and input embeddings within the multi-task framework accordingly. Through experiments involving different numbers of negative samples, we demonstrate that the multi-task framework improves model performance in scenarios with an abundance of negative samples. We observed that the tasks of joint entity-relation extraction are both built upon a common span representation. In light of this, future work will involve enriching the representation information at the span level to validate our hypothesis.

## ACKNOWLEDGMENTS

The work is supported in part by the National Key R&D Program of China (Grant No. 2021YFB3900605 & 2021YFB3900601).

## References

1. M. Eberts and A. Ulges, “Span-based joint entity and relation extraction with transformer pre-training,” arXiv preprint arXiv:1909.07755, 2019.
2. Y. Luan, L. He, M. Ostendorf, and H. Hajishirzi, “Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction,” arXiv preprint arXiv:1808.09602, 2018.
3. H. Gurulingappa, A. M. Rajput, A. Roberts, J. Fluck, M. Hofmann-Apitius, and L. Toldo, “Development of a benchmark corpus to support the automatic extraction of drug-related adverse effects from medical case reports,” *Journal of biomedical informatics*, vol. 45, no. 5, pp. 885–892, 2012.
4. C. N. d. Santos, B. Xiang, and B. Zhou, “Classifying relations by ranking with convolutional neural networks,” arXiv preprint arXiv:1504.06580, 2015.
5. D. Roth and W.-t. Yih, “A linear programming formulation for global inference in natural language tasks,” in *Proceedings of the eighth conference on computational natural language learning (CoNLL-2004) at HLT-NAACL 2004*, 2004, pp. 1–8.
6. P. Gupta, H. Schütze, and B. Andrassy, “Table filling multi-task recurrent neural network for joint entity and relation extraction,” in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, 2016, pp. 2537–2547.
7. D. Wadden, U. Wennberg, Y. Luan, and H. Hajishirzi, “Entity, relation, and event extraction with contextualized span representations,” arXiv preprint arXiv:1909.03546, 2019.
8. T. Lai, H. Ji, C. Zhai, and Q. H. Tran, “Joint biomedical entity and relation extraction with knowledge-enhanced collective inference,” arXiv preprint arXiv:2105.13456, 2021.
9. R. Li, D. Li, J. Yang, F. Xiang, H. Ren, S. Jiang, and L. Zhang, “Joint extraction of entities and relations via an entity correlated attention neural model,” *Information Sciences*, vol. 581, pp. 179–193, 2021.

10. Y. Ma, T. Hiraoka, and N. Okazaki, "Joint entity and relation extraction based on table labeling using convolutional neural networks," in Proceedings of the Sixth Workshop on Structured Prediction for NLP, 2022, pp. 11–21.
11. R. Tang, Y. Chen, Y. Qin, R. Huang, B. Dong, and Q. Zheng, "Boundary assembling method for joint entity and relation extraction," Knowledge-Based Systems, vol. 250, p. 109129, 2022.
12. D. Ye, Y. Lin, P. Li, and M. Sun, "Packed levitated marker for entity and relation extraction," arXiv preprint arXiv:2109.06067, 2021.
13. Z. Wei, J. Su, Y. Wang, Y. Tian, and Y. Chang, "A novel cascade binary tagging framework for relational triple extraction," arXiv preprint arXiv:1909.03227, 2019.
14. B. Taillé, V. Guigue, G. Scoutheeten, and P. Gallinari, "Let's stop incorrect comparisons in end-to-end relation extraction!" arXiv preprint arXiv:2009.10684, 2020.
15. C. Zheng, Y. Cai, J. Xu, H. Leung, and G. Xu, "A boundary-aware neural model for nested named entity recognition," in Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). Association for Computational Linguistics, 2019.
16. J. Cao, G. Wang, C. Li, H. Ren, Y. Cai, R. C.-W. Wong, and Q. Li, "Incorporating boundary and category feature for nested named entity recognition," in Database Systems for Advanced Applications: 25th International Conference, DASFAA 2020, Jeju, South Korea, September 24–27, 2020, Proceedings, Part II 25. Springer, 2020, pp. 209–226.
17. C. Tan, W. Qiu, M. Chen, R. Wang, and F. Huang, "Boundary enhanced neural span classification for nested named entity recognition," in Proceedings of the AAAI conference on artificial intelligence, vol. 34, no. 05, 2020, pp. 9016–9023.
18. H. Fei, Y. Ren, and D. Ji, "Recognizing nested named entity in biomedical texts: a neural network model with multi-task learning," in 2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). IEEE, 2019, pp. 376–381.
19. T. H. Nguyen and R. Grishman, "Relation extraction: Perspective from convolutional neural networks," in Proceedings of the 1st workshop on vector space modeling for natural language processing, 2015, pp. 39–48.
20. M. Miwa and M. Bansal, "End-to-end relation extraction using lstms on sequences and tree structures," arXiv preprint arXiv:1601.00770, 2016.
21. M. Miwa and Y. Sasaki, "Modeling joint entity and relation extraction with table representation," in Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), 2014, pp. 1858–1869.
22. Y. Luan, D. Wadden, L. He, A. Shah, M. Ostendorf, and H. Hajishirzi, "A general framework for information extraction using dynamic span graphs," arXiv preprint arXiv:1904.03296, 2019.
23. C. Wang, A. Li, H. Tu, Y. Wang, C. Li, and X. Zhao, "An advanced bert-based decomposition method for joint extraction of entities and relations," in 2020 IEEE Fifth International Conference on Data Science in Cyberspace (DSC). IEEE, 2020, pp. 82–88.

## Authors

**Jiamin Lu** Assistant Professor at Information Department in Hohai University, China. He received his Ph.D degree in Information Science from FernUniversität in Hagen, Germany, 2014. His research interests include parallel processing on MOD (Moving Object Database), data management in Knowledge Graph construction. At present, he is mainly working on the conjunction of big data technologies and smart water applications.

**Chenguang Xue** Currently, he is pursuing his Master in Computer Science from the University of Hohai. His research interests include Natural Language Processing and Knowledge Graph Construction.