# A Joint Extraction Framework for Cyber Threat Intelligence Triplets Based on Knowledge Enhancement and Multi-task Optimization

**Chuanyue Li**[1]

*1. School of Computer Science, Beijing University Of Technology, Beijing, 100020, China*

**Abstract:**

Triplet extraction from unstructured text is a fundamental task in the construction of cyber threat intelligence (CTI) knowledge graphs. However, traditional pipeline approaches often suffer from redundant relation prediction, entity overlap, complex contextual dependencies, and limited domain knowledge. This paper proposes a joint triplet extraction framework tailored to CTI scenarios. Our model integrates a knowledge enhancement module, potential relation prediction, relation-specific entity labeling, and a global triplet validation module. The knowledge enhancement module enriches threat texts using an external knowledge base, improving semantic understanding of security-related terms. The potential relation prediction module filters out invalid relations, while the dual BIO-based labeling mechanism addresses overlapping entities. The final validation module scores and selects the most valid triplets. We constructed a CTI-specific knowledge base from MITRE ATT&CK and other public sources, and evaluated our method on two datasets: HACKER and RE-DNRTI. Experimental results show our method outperforms strong baselines such as PRGC and OD-RTE in both precision and F1 score, particularly in noisy and complex scenarios. Ablation and sensitivity experiments demonstrate the importance of each module and the robustness of the overall framework. This research contributes a reliable and interpretable method for high-quality CTI triplet extraction.

**Keywords:** Cyber Threat Intelligence; Triplet Extraction; Knowledge Augmentation; Multi-task Learning; Relation Validation.

## 1 INTRODUCTION

In recent years, with the increasing complexity of cyberattacks, cybersecurity defenses are facing unprecedented challenges [1]. CTI as a critical tool for predicting, identifying, and defending against cyberattacks [2-3], has become an essential area of research in the field of cybersecurity. However, traditional threat intelligence extraction methods typically rely on complex manual rules or pipeline methods [4-7], which have significant limitations when handling large-scale, complex, and dynamic cybersecurity data [8]. Although traditional pipeline-based approaches to extract CTI triplets have some flexibility, they often suffer from issues such as the propagation of errors from entity recognition and low extraction efficiency due to the separation of entity recognition and relation extraction [9]. Therefore, effectively extracting structured triplets from complex CTI reports has become a pressing problem in the field of cybersecurity.

To overcome these issues, we propose a joint CTI triplet extraction framework that enhances contextual representation through knowledge insertion, predicts potential relations to reduce noise, and applies a validation mechanism to ensure the semantic coherence of triplets. The proposed model leverages a multi-task learning strategy to jointly optimize relation detection, entity labeling, and triplet validation, ensuring accuracy and robustness. Extensive experiments validate its performance and interpretability.

## 2 RELATED WORK

Triplet extraction of CTI is an important research direction in the field of cybersecurity. With the development of deep learning technology [10-12], related research has made significant progress in model architecture optimization and domain knowledge integration. However, existing methods still face many challenges when dealing with complex cybersecurity texts.

In the aspect of triplet extraction, traditional pipeline methods separate entity recognition from relation extraction, leading to error propagation and semantic information discontinuity issues. To address these problems, Wei et al. proposed the cascaded annotation framework CASREL, which for the first time modeled relations as mapping functions of entity pairs, alleviating task dependency through a two-stage decoding process. But the recall has significantly decreased in complex scenarios [13]. The PRGC framework developed by Zheng et al. reduced redundant computation through relation filtering. But it did not effectively solve the semantic ambiguity caused by domain-specific terms in the field of cybersecurity [14]. Recently, Ning et al. proposed the OD-RTE model, which using regional detection to enhance the accuracy of triplet localization. However, it has limited capability in capturing long-distance dependency relations within threat intelligence texts [15].

Knowledge-enhanced technology is a key means to improve domain adaptability. Most existing studies adopt static knowledge injection methods. Luo et al. introduced the MITRE ATT&CK knowledge base to construct distant supervision rules in the HACKER dataset, but the manually defined relationship matching patterns struggle to cover dynamically evolving threat behavior patterns [16]. Yu et al. proposed a malicious traffic detection model that utilizes transfer learning and knowledge distillation. However, the model necessitates the computation of Maximum Mean Discrepancy (MMD) and Kullback-Leibler (KL) divergence in multiple training processes, which leads to high computational complexity [17]. Sarhan et al. developed the Open-CyKG framework, which combines open information extraction (OIE) with knowledge graphs for the first time. It utilizes attention mechanisms to extract triplets from unstructured APT reports [18]. However, existing OIE methods typically rely on syntactic patterns or statistical features, which limits their generalization capabilities. Chen et al. proposed a method for multimodal named entity recognition and multimodal relation extraction, which is based on Chain-of-Thought (CoT) prompt distillation [19]. But the CoT knowledge generated by large language models may be unstable or contain inaccuracies, potentially leading to the propagation of erroneous reasoning.

Regarding the issue of nested entities and conflicting multiple relations in complex texts, Wang et al.'s TPlinker employs a matrix tagging scheme to resolve overlapping entity pairs [20]. However, its O(n²) complexity makes it challenging to handle long texts. He et al. proposed virtual prompts and designed entity-relation-aware pre-training tasks, which may require optimization with domain knowledge for specific fields, although they effectively reduce the reliance on manual prompts [21]. Xu et al. introduced a supervised multi-head self-attention to learn the correlations between words under different entity types, yet they failed to fully utilize external knowledge resources to enhance the distinction in terms of entity type differentiation [22]. Li et al. proposed a semi-supervised entity recognition method based on active learning and self-learning, but the active learning strategy is only based on uncertainty sampling which affects the generalization ability of the model [23]. These methods have made progress in general domains, but they have not fully considered the unique entity relation distribution characteristics of CTI, which limits their practical application effects.

## 3 METHODS

Figure 1 presents the overall architecture. The model consists of a knowledge enhancement module, a BERT-based encoder, and a multi-task decoder that includes:A potential relation prediction module to estimate relevant relations. A relation-specific sequence labeling module to extract subject and object entities. And an effectiveness judgment module to validate and select coherent triplets.

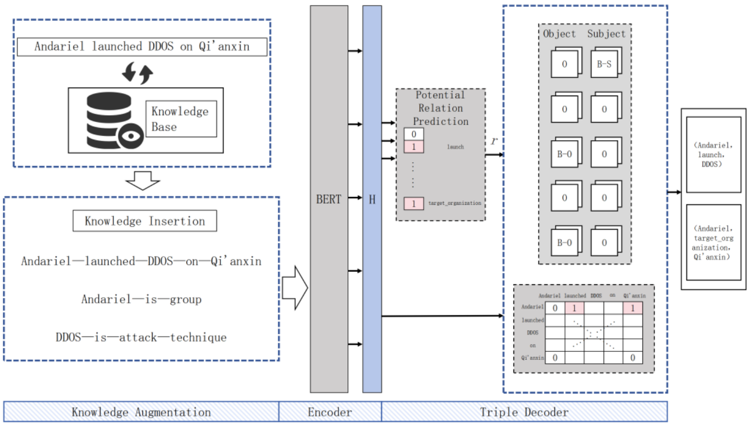

Figure 1. Framework of the joint extraction model

### 3.1 KNOWLEDGE AUGMENTATION MODULE

To effectively capture domain-specific semantics in CTI texts, our framework incorporates a dedicated knowledge enhancement module prior to encoding. This module is designed to address the inherent ambiguity and sparsity of information in raw CTI data by introducing external, structured knowledge. Specifically, we construct a comprehensive CTI knowledge base consisting of 1,887 curated entries drawn from multiple authoritative sources, including MITRE ATT&CK (covering techniques, groups, and software) and cybersecurity industry reports, as illustrated in Table 1. The entries are uniformly organized in the format of "keyword–associated knowledge", providing concise factual statements that can supplement the original text.

Table 1 Knowledge Base Sources

| Source | Quantity |
|---|---|
| https://attack.mitre.org/ | 254 |
| https://attack.mitre.org/groups/ | 298 |
| https://attack.mitre.org/software/ | 672 |
| https://cybersecurityventures.com/ | 663 |

The enhancement process begins with syntactic analysis of the input sentence, from which we extract noun phrase chunks as candidate entities or key terms. These chunks are identified based on standard dependency parsing techniques, which allow the model to isolate meaningful multi-word expressions centered around domain-specific nouns. To reduce redundancy and avoid semantic overlap, we then apply a pruning strategy that eliminates nested or repetitive phrases. For instance, if a noun phrase is a subset of a longer phrase occurring in the same sentence, only the more complete and informative term is retained. This step ensures that the subsequent knowledge matching is both efficient and semantically relevant.

Once the refined set of candidate keywords is obtained, each term is matched against entries in the external knowledge base. Upon a successful match, the corresponding factual knowledge is retrieved and appended directly to the input sentence. Rather than altering the original structure of the sentence, we concatenate the retrieved knowledge as an auxiliary segment to the end of the sentence. This strategy maintains syntactic integrity while expanding the semantic field accessible to the encoder. As illustrated in Figure 2, this process enables the model to attend not only to the explicit content of the CTI text but also to the implicit associations embedded in the external knowledge base. The enriched input thus provides a stronger foundation for accurate relation inference and entity boundary detection in subsequent stages of the model.

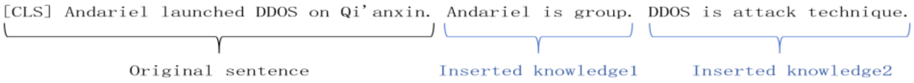
[CLS] Andariel launched DDOS on Qi'anxin. Andariel is group. DDOS is attack technique.
Original sentence
Inserted knowledge1
Inserted knowledge2

Figure 2 An Example of Knowledge Insertion

### 3.2MULTI-STAGE DECODER

A Potential Relation Prediction

Compute sentence-level relation probabilities via average pooling:

$$h^{\text{avg}} = \frac{1}{n}\sum_{i=1}^{n} h_i$$

$$P_{\text{rel}} = \sigma(W_r h^{\text{avg}} + b_r)$$

Relations with $P_{\text{rel}} \geq \lambda_1 = 0.5$ are retained, reducing redundant predictions.

B Relation-Specific Tagging

For each predicted relation $j$: Head Entity$P_{ij}^{s} = \text{softmax}(W_s[h_i + u_j] + b_s)$, Tail Entity$P_{ij}^{o} = \text{softmax}(W_o[h_i + u_j] + b_o)$. Here, $u_j$ is the relation embedding, enabling relation-specific entity detection.

C Validity Verification

Score candidate triplets $(\mathrm{h}_{\mathrm{i}}^{\mathrm{s}}, \mathrm{h}_{\mathrm{j}}^{\mathrm{o}}, u_k)$: $P_{\text{valid}} = \sigma\big(W_v[(h_i^s + u_k)\|(h_j^o + u_k)] + b_v\big)$. Triplets with $P_{\text{valid}} \geq \lambda_2 = 0.5$ are validated.

### 3.3 MULTI-TASK OPTIMIZATION

The total loss combines three objectives:

Relation Prediction ($\mathcal{L}_{\text{rel}}$): Guides the model to recognize valid relations.

$$\mathcal{L}_{\text{rel}} = -\frac{1}{R}\sum_{r=1}^{R} y_r \log \hat{p}_r + (1 - y_r)\log (1 - \hat{p}_r)$$

Entity Tagging ($\mathcal{L}_{\text{entity}}$): Ensures precise boundary detection for overlapping entities.

$$\mathcal{L}_{\text{entity}} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c} \log \hat{p}_{i,c}$$

Triplet Validity ($\mathcal{L}_{\text{triplet}}$): Penalizes inconsistent entity-relation pairs.

$$\mathcal{L}_{\text{triplet}} = -\frac{1}{|T|}\sum_{(i,j,k)\in T} y_{ijk} \log \hat{p}_{ijk} + (1 - y_{ijk})\log (1 - \hat{p}_{ijk})$$

Total Loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rel}} + \mathcal{L}_{\text{entity}} + \mathcal{L}_{\text{triplet}}$$

## 4 EXPERIMENTS

To validate the effectiveness and robustness of the proposed joint extraction framework, we conducted comprehensive experiments on two widely recognized cybersecurity datasets and performed a detailed analysis of model components and hyperparameter configurations.

We evaluated our method on two benchmark datasets: HACKER [16] and RE-DNRTI [24]. The HACKER dataset, constructed from 2,153 real-world CTI reports, includes a diverse set of attack scenarios and comprises a total of 5,747 relation triplets across 22 predefined relation types. This dataset emphasizes multi-faceted threat actor behavior and contextual entity associations in unstructured threat narratives. In contrast, the RE-DNRTI dataset is an augmented version of the DNRTI corpus, which was originally developed for CTI relationship extraction. Through distant supervision and high-quality filtering strategies, RE-DNRTI provides 13,839 annotated triplets spanning 19 relation types, with a particular focus on threat actor profiling, tool usage, and infrastructure relationships. The dataset statistics and train/dev/test splits are summarized in Table 2.

Table 2. Key Statistics and Data Splits

| Dataset | Relations | Triplets | Train/Dev/Test |
|---|---|---|---|
| HACKER | 22 | 5,747 | 2,058/258/271 |
| RE-DNRTI | 19 | 13,839 | 4,463/558/558 |

All experiments were conducted using a high-performance computing environment equipped with an NVIDIA RTX 4060 GPU and 32 GB of RAM. The implementation was based on PyTorch version 1.8.0 and Python 3.10.14. To ensure fair comparisons and reproducible results, we adopted a consistent set of hyperparameters throughout all experiments. Specifically, the batch size was set to 16, and the maximum sequence length of input tokens was fixed at 100. The model was trained for 50 epochs with an initial learning rate of 1e-4, and a dropout rate of 0.3 was applied to mitigate overfitting. We employed a warm-up strategy with a proportion of 0.1 for learning rate scheduling, and used gradient accumulation with two steps to handle longer sequences under memory constraints.

In our evaluation, we adopt standard metrics widely used in information extraction tasks, including Precision (P), Recall (R), and F1-score, to comprehensively assess the model’s ability to accurately identify and structure relation triplets. All reported results are based on the exact match criterion of triplet components.

### 4.1 COMPARATIVE EVALUATION

To assess the effectiveness of our proposed joint extraction framework, we conducted comparative experiments against two state-of-the-art models: PRGC [14] and OD-RTE [15]. PRGC is an end-to-end joint entity-relation extraction framework that utilizes a span-based tagging strategy to handle overlapping triplets. OD-RTE, on the other hand, formulates triplet extraction as a region-based object detection task, leveraging graph-based modeling of triplet regions for enhanced contextual reasoning. We evaluated all models on both HACKER and RE-DNRTI datasets under identical experimental conditions.

The results, as illustrated in Figure 3, demonstrate that our model achieves superior performance in terms of precision (P) and F1-score on both datasets. While our method yields a slightly lower recall (R) compared to OD-RTE, it consistently outperforms both baselines in F1, indicating a better balance between precision and recall. On the HACKER dataset, our model achieves an F1-score of 44.2%, surpassing OD-RTE by 5.8 percentage points. This indicates that our method is particularly effective at handling sparse relations and long-tail distribution problems commonly observed in real-world CTI scenarios. On the RE-DNRTI dataset, our framework attains an F1-score of 86.4%, significantly outperforming PRGC and OD-RTE, and showcasing its robustness in multi-relation and complex contextual settings.

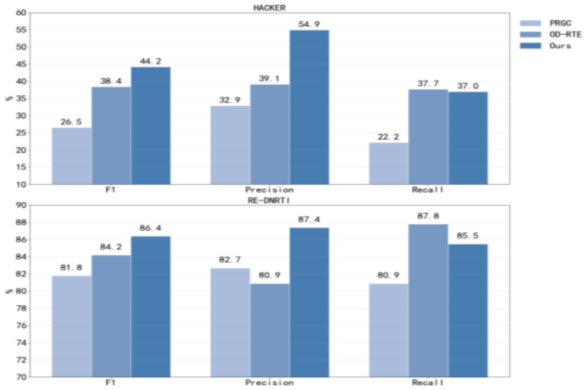

Figure 3 Comparison Experiment Results

These improvements are attributed to several innovations in our framework. The knowledge enhancement module provides strong domain priors, enabling the model to better interpret security-specific terms and relations. The dynamic attention mechanism reduces the influence of irrelevant context, while the effectiveness judgment module precisely filters valid triplet candidates. Together, these components ensure both efficiency and reliability in triplet extraction, particularly in challenging, real-world threat intelligence settings.

### 4.2 ABLATION STUDY

To further validate the contributions of individual components within our framework, we performed ablation studies by systematically removing or modifying specific modules and observing the impact on performance. The three core modules evaluated in this study include the knowledge enhancement module, the potential relation prediction module, and the triplet effectiveness judgment module. Each variant was tested on both HACKER and RE-DNRTI datasets to ensure consistency and generalizability of the findings, with results presented in Figure 4.

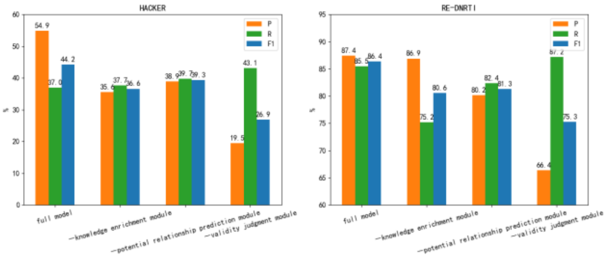

Figure 4 Comparison Experiment Results

When the knowledge enhancement module was removed, the F1-score on the HACKER dataset dropped from 44.2% to 36.6%, reflecting a significant decline in the model's understanding of complex contextual semantics. This degradation is further corroborated by noticeable drops in both precision and recall, indicating a reduced ability to detect valid relationships and entities. On RE-DNRTI, the removal of this module led to a 5.8% decline in F1, with recall falling from 85.5% to 75.2%. These results affirm the importance of external knowledge in improving both the comprehensiveness and correctness of CTI triplet extraction.

Eliminating the potential relation prediction module also negatively affected performance. Without this module, the model lost its ability to focus on likely relations and was forced to perform entity extraction for all relation types, resulting in increased noise. The F1-score dropped by 4.9% on HACKER and 5.1% on RE-DNRTI. This decline was particularly evident in HACKER due to its high complexity and label imbalance. These findings suggest that the relation prediction module is crucial for filtering out spurious relations and directing attention toward valid ones.

The most significant performance degradation occurred when the effectiveness judgment module was removed. Without this global scoring mechanism, the model could not adequately assess the coherence of entity pairs and relations, leading to a sharp decline in F1-score—down to 26.9% on HACKER and 75.3% on RE-DNRTI. Precision suffered the most, with a drop from 54.9% to 19.5% on HACKER, reflecting an inability to distinguish high-quality triplets from noise. Interestingly, recall showed a slight increase, indicating that more candidates were generated, albeit with low correctness. Overall, the ablation study confirms the indispensable role of all three modules in achieving high performance, particularly in noisy and ambiguous CTI contexts.

### 4.3 HYPERPARAMETER SENSITIVITY ANALYSIS

To explore the sensitivity of our model to threshold settings in the decoding stage, we conducted a grid-based sensitivity analysis on the two critical hyperparameters: $\lambda_1$, the threshold for potential relation prediction, and $\lambda_2$, the threshold for triplet validity scoring. The analysis was performed independently on both HACKER and RE-DNRTI datasets using the F1-score as the primary evaluation metric. The results are summarized in Figure 5.

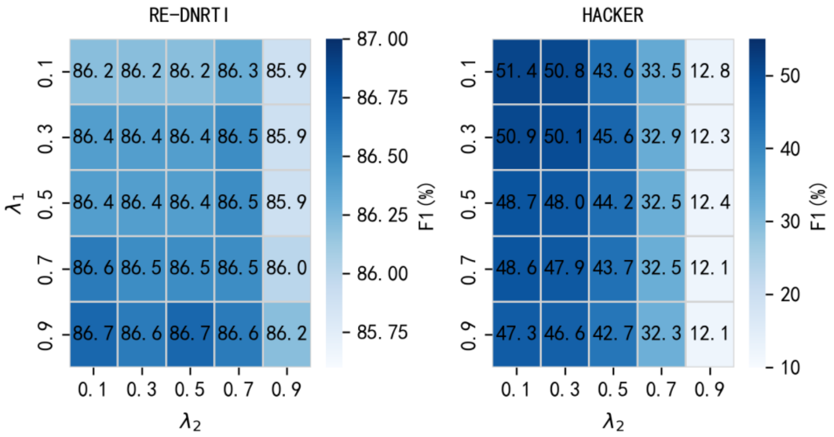

Figure 5 Hyperparameter Sensitivity Experiment Results

The RE-DNRTI dataset showed stable performance across a wide range of threshold combinations, with F1 scores consistently remaining between 85% and 87%. This robustness can be attributed to the dataset's high-quality annotations and well-defined relation boundaries, resulting from a structured distant supervision and filtering process during its construction. In contrast, the HACKER dataset exhibited high sensitivity to threshold variations. While optimal settings could yield F1 scores close to 50%, small deviations from the optimal thresholds led to substantial performance degradation, with F1 dropping below 20% in some cases.

The disparity in sensitivity between the two datasets underscores the importance of dataset quality and annotation consistency in triplet extraction tasks. In particular, the noisier HACKER dataset, which contains ambiguous or loosely defined relation types, relies more heavily on precise threshold tuning and filtering mechanisms. These findings emphasize the necessity of the effectiveness judgment module and adaptive threshold control in real-world CTI applications where data may be sparse or imprecise.

## DISCUSSION

We proposed a novel joint extraction framework tailored for CTI triplet extraction, which effectively addresses the limitations of traditional pipeline methods, including redundant relation prediction, entity overlap, and insufficient domain knowledge. Our approach introduces several key innovations: a knowledge enhancement module that leverages external CTI knowledge bases to enrich input semantics; a potential relation prediction module that filters irrelevant relations and reduces noise; a relation-specific sequence labeling strategy that handles overlapping entities; and an effectiveness judgment module that validates triplet coherence through global semantic reasoning.

Through comprehensive experiments conducted on two benchmark datasets—HACKER and RE-DNRTI—we demonstrated that our method consistently outperforms state-of-the-art baselines in terms of precision and F1-score, particularly in complex and noisy scenarios. Ablation studies confirmed the importance of each proposed module, while sensitivity analysis further highlighted the robustness of our framework under varying threshold settings. These findings collectively validate the effectiveness, interpretability, and practicality of our approach in real-world CTI applications.

Looking ahead, future work may explore the integration of large-scale pre-trained language models or prompt-based learning techniques to further improve generalization across diverse CTI domains. In addition, we aim to extend our knowledge enhancement strategies by adopting dynamic retrieval from online CTI repositories and expanding the coverage of domain-specific knowledge. Ultimately, this research provides a promising step toward more intelligent and reliable CTI knowledge graph construction, paving the way for enhanced cyber defense capabilities.

## REFRENCES

[1] Bouwman X, Le Pochat V, Foremski P, et al. Helping hands: Measuring the impact of a large threat intelligence sharing community[C]//31st USENIX Security Symposium (USENIX Security 22). 2022: 1149-1165

[2] Gao P, Shao F, Liu X, et al. A system for efficiently hunting for cyber threats in computer systems using threat intelligence[C]//2021 IEEE 37th International Conference on Data Engineering (ICDE). IEEE, 2021: 2705-2708.

[3] Gao P, Shao F, Liu X, et al. Enabling efficient cyber threat hunting with cyber threat intelligence[C]//2021 IEEE 37th International Conference on Data Engineering (ICDE). IEEE, 2021: 193-204.

[4] Satvat K, Gjomemo R, Venkatakrishnan V N. Extractor: Extracting attack behavior from threat reports[C]//2021 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, 2021: 598-615.

[5] Gao P, Liu X, Choi E, et al. Threatkg: A threat knowledge graph for automated open-source cyber threat intelligence gathering and management[J]. arXiv preprint arXiv:2212.10388, 2022

[6] Husari G, Al-Shaer E, Ahmed M, et al. Ttpdrill: Automatic and accurate extraction of threat actions from unstructured text of cti sources[C]//Proceedings of the 33rd annual computer security applications conference. 2017: 103-115.

[7] Li Z, Zeng J, Chen Y, et al. AttacKG: Constructing technique knowledge graph from cyber threat intelligence reports[C]//European Symposium on Research in Computer Security. Cham: Springer International Publishing, 2022: 589-609.

[8] Bansal M A, Sharma D R, Kathuria D M. A systematic review on data scarcity problem in deep learning: solution and applications[J]. ACM Computing Surveys (Csur), 2022, 54(10s): 1-29.

[9] Wang X, Liu J. A novel feature integration and entity boundary detection for named entity recognition in cybersecurity[J]. Knowledge-Based Systems, 2023, 260: 110114.

[10] Liu W, Zhou P, Zhao Z, et al. K-bert: Enabling language representation with knowledge graph[C]//Proceedings of the AAAI Conference on Artificial Intelligence. 2020, 34(03): 2901-2908.

[11] Shang Y M, Huang H, Mao X. Onerel: Joint entity and relation extraction with one module in one step[C]//Proceedings of the AAAI conference on artificial intelligence. 2022, 36(10): 11285-11293.

[12] Zhang Z, Zhao Y, Gao H, et al. LinkNER: linking local named entity recognition models to large language models using uncertainty[C]//Proceedings of the ACM Web Conference 2024. 2024: 4047-4058.

[13] Wei Z, Su J, Wang Y, et al. A novel cascade binary tagging framework for relational triple extraction[J]. arXiv preprint arXiv:1909.03227, 2019.

[14] Zheng H, Wen R, Chen X, et al. PRGC: Potential relation and global correspondence based joint relational triple extraction[J]. arXiv preprint arXiv:2106.09895, 2021.

[15] Ning J, Yang Z, Sun Y, et al. OD-RTE: a one-stage object detection framework for relational triple extraction[C]//Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). 2023: 11120-11135.

[16] Luo Y, Ao S, Luo N, et al. Extracting threat intelligence relations using distant supervision and neural networks[C]//Advances In Digital Forensics XVII: 17th IFIP WG 11.9 International Conference, Virtual Event, February 1−2, 2021, Revised Selected Papers 17. Springer International Publishing, 2021: 193-211.

[17] Yu K, Huang Y Y, Zhang L, et al. TL-KDNET: A Malicious Traffic Detection Model Integrating Transfer Learning and Knowledge Distillation[C]//2024 21st International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP). IEEE, 2024: 1-4.

[18] Sarhan I, Spruit M. Open-cykg: An open cyber threat intelligence knowledge graph[J]. Knowledge-Based Systems, 2021, 233: 107524.

[19] Chen F, Feng Y. Chain-of-thought prompt distillation for multimodal named entity recognition and multimodal relation extraction[J]. arXiv preprint arXiv:2306.14122, 2023.

[20] Wang Y, Yu B, Zhang Y, et al. TPLinker: Single-stage joint extraction of entities and relations through token pair linking[J]. arXiv preprint arXiv:2010.13415, 2020.

[21] He K, Huang Y, Mao R, et al. Virtual prompt pre-training for prototype-based few-shot relation extraction[J]. Expert systems with applications, 2023, 213: 118927.

[22] Xu Y, Huang H, Feng C, et al. A supervised multi-head self-attention network for nested named entity recognition[C]//Proceedings of the AAAI conference on artificial intelligence. 2021, 35(16): 14185-14193.

[23] Li T, Hu Y, Ju A, et al. Adversarial Active Learning for Named Entity Recognition in Cybersecurity[J]. Computers, Materials & Continua, 2021, 66(1).

[24] Wang X, Liu X, Ao S, et al. Dnrti: A large-scale dataset for named entity recognition in threat intelligence[C]//2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom). IEEE, 2020: 1842-1848.