TyRaL: End-to-End Document-level Relation Extraction via Type-Constrained Rule Learning

Mierzhati Alimu[1], Chaochao Du[1] and Xiaowang Zhang[1,*]

[1] *College of Intelligence and Computing, Tianjin University, Tianjin, 300350, China*

Abstract

In recent years, Document-level Relation Extraction (DocRE) has encountered significant challenges in capturing complex entity relationships and reasoning over long-range dependencies. Existing methods primarily focus on learning implicit representations or applying chain-like logical rules, but they often overlook differences in entity types and the significance of type constraints, potentially leading to errors in relation reasoning. This poster introduces a type-constrained enhanced chain-like rule (TC rule) and proposes an end-to-end document-level relation extraction framework (TyRaL) to address this issue. By incorporating a novel rule reasoning module, TyRaL transforms the discrete rule learning problem into a parameter optimization task in continuous space, enabling both explicit and implicit learning of entity type constraint rules and thereby enhancing the model's logical consistency and interpretability. Experimental results on the standard DWIE dataset show that TyRaL significantly outperforms existing rule-enhanced methods in both F1 and Ign F1 metrics. It demonstrates superior logical modeling and semantic reasoning capabilities while offering new perspectives and solutions for research in the DocRE field.

Keywords

Document-level Relation Extraction, Logical Rules, Type Constraints

1. Introduction

DocRE aims to identify all relations between entity pairs across an entire document. It faces greater challenges in modeling long-range context and complex dependencies between entities than sentence-level extraction. Existing approaches can be broadly categorized into three groups: sequence-based models, graph-based models, and rule-constrained models. While the first two focus on learning implicit representations, they often lack logical interpretability and struggle to infer potential relations. In contrast, rule-based methods offer better transparency and reasoning capabilities. MILR[1] and BCBR[2] learn chain-like rules from annotated data, while CaDRL[3] and JMRL[4] generate such rules dynamically during training.

Although the above methods have made notable progress in improving model performance, the rules they learn are based on connecting entities through shared variables, which often overlooks differences in entity types and may lead to incorrect predictions in complex cases. For example, consider the following chain-like rule: $\texttt{hasFather}^{\text{new}}(x, y) \leftarrow \texttt{hasChild}(x, y) \wedge \texttt{Male}(y)$. In this rule, $\texttt{Male}(y)$ functions as an explicit type constraint, specifying that y is male and thereby preventing incorrect inferences—such as misclassifying a mother as a father. However, type constraints can also be inferred implicitly from relational context. For instance, if y is described as someone's brother or uncle, it is reasonable to infer that y is male, even without an explicit $\texttt{Male}(y)$ statement. These implicit type constraints are especially valuable when explicit type annotations are unavailable or incomplete.

To address this issue, we introduce Type-Constrained Enhanced Chain-like Rules (TC Rules), which extend standard logical rules with type constraints of both kinds. Based on this, we propose End-to-End Document-level Relation Extraction via Type-Constrained Rule Learning (TyRaL). This novel end-to-end framework transforms rule learning from a discrete problem into continuous parameter optimization, building a type-constrained neural reasoning module. Experimental results on the DWIE[5]

ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan

*Corresponding author.
mierzhati_0908@tju.edu.cn (M. Alimu); chuck@tju.edu.cn (C. Du); xiaowangzhang@tju.edu.cn (X. Zhang)

dataset show that TyRaL outperforms existing rule-based DocRE models in logical consistency and relation extraction.

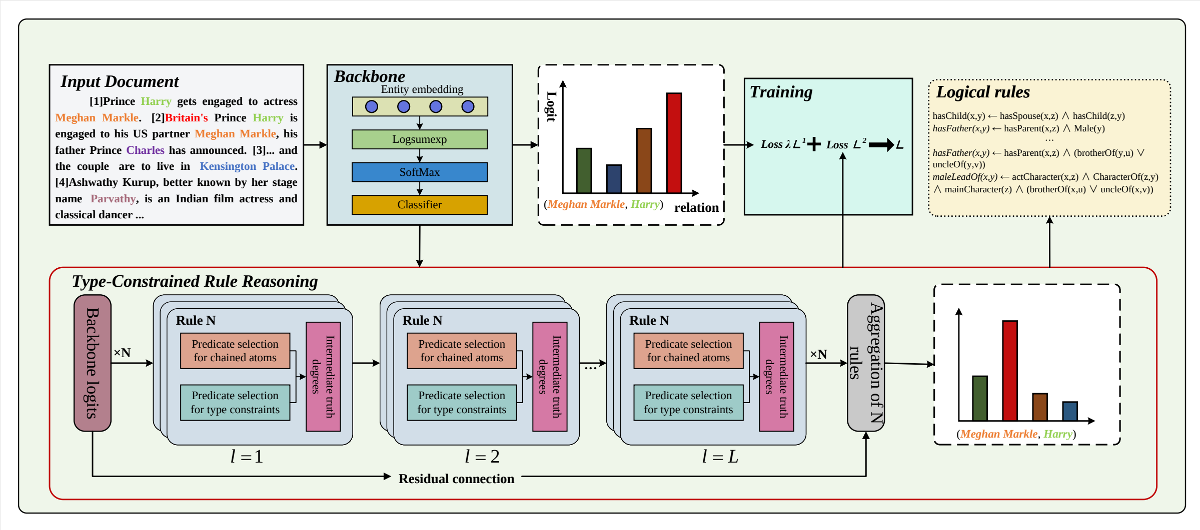

Figure 1: The overview of the TyRaL framework.

2. Approach

2.1. Problem Definition

Given a document D containing a set of named entities $E_D = \{e_i\}_{i=1}^{n}$, the goal of DocRE is to predict the semantic relation $r \in R \cup \{NA\}$ between all distinct entity pairs (e_h, e_t), where R denotes a set of predefined relation types, and NA represents no relation. An entity e_i may have multiple mentions in the document, denoted as $\left\{m_j^i\right\}_{i=1}^{N_{e_i}}$. The existence of a relationship between entities needs to be judged based on comprehensive contextual evidence between these mentions in the document.

An original DocRE model usually calculates a score vector $F(e_h, e_t, D) \in \mathbb{R}^{|R|+1}$ for each entity pair, where the k-th element represents the logit value of the k-th relation type, and the last element corresponds to "no relation" NA. During the training phase, Binary Cross-Entropy (BCE) or Adaptive Thresholding (AT) loss functions are usually used.

In the inference phase, the model uses an activation function σ (such as Softmax) to map logits to probability values, and filters them according to a threshold ϵ to predict the set of relation triples, which is in the form:

$$Y = \left\{(e_h, r, e_t) \mid [\sigma\left(F(e_h, e_t, D)\right)]_r > \epsilon\right\} \tag{1}$$

where ϵ is the set confidence threshold.

2.2. Chain-like and Type-Constrained Rules

We introduce an interpretable logical rule structure to model implicit semantic paths between entities in a document. Define a binary variable $r\,(x, y)$ to indicate whether the relation $r \in R$ holds between entities x and y . When the relation is true, $r\,(x, y) = 1$; otherwise, $r\,(x, y) = 0$.

A chain-like logical rule consists of a rule head and a rule body. The rule head represents the target relation $r_{head}\,(x, y)$, and the rule body is a conjunction of binary atoms, where each body atom shares a variable with the adjacent previous atom and another variable with the adjacent next atom, forming a chain structure. The general form of a chain-like logical rule is as follows:

$$r_{head}\,(x, y) \leftarrow r_1\,(x, z_1) \wedge r_2\,(z_1, z_2) \wedge \cdots \wedge r_l\,(z_{l-1}, y) \tag{2}$$

Building on this, we introduce TC Rules by adding unary type atoms to impose semantic restrictions on chain-like logical rules. Let C be the set of entity types. A type-constrained rule with length L is in the form:

$$r_{head}(x, y) \leftarrow c_1(x) \wedge c_2(z_1) \wedge r_1(x, z_1) \wedge r_2(z_1, z_2) \wedge \cdots \wedge r_l(z_{l-1}, y) \wedge c_{L+1}(y) \tag{3}$$

Where $c_1, c_2, \ldots, c_{L+1} \in C$ are entity types, and r_i are intermediate relation paths. This rule not only depends on the relational path structure but also requires each entity node on the path to satisfy specific type conditions, thereby improving the semantic rationality and interpretability of the rule.

2.3. Type-Constrained Rule Reasoning Module

Figure 1 shows an overview of our model framework. We propose a type-constrained rule reasoning module to enable end-to-end learning of TC Rules. This module is the core component of the TyRaL framework, and its basic idea is to transform the rule learning problem in discrete space into a parameter learning problem in continuous space, thereby simulating the reasoning process of type-constrained rules in a differentiable manner. This module is jointly trained with the DocRE to optimize the downstream relation prediction target.

Let N be the maximum number of rules to be learned, L the maximum number of atoms in each rule, and define the extended relation set as $R^* = R \cup R^- \cup \{I\}$, where $R = \{r_i\}_{1 \le i \le n}$ denotes the original relation set, $R^- = \{r_i\}_{n+1 \le i \le 2n}$ the inverse relations, and $I = r_{2n+1}$ the identity relation. We define the extended logit $F^+(x, y, d) \in \mathbb{R}^{2n+1}$, where $[F^+(x, y, d)]_i = [\sigma(F(x, y, d))]_i$ and $[F^+(x, y, d)]_{i+n} = [\sigma(F(y, x, d))]_i$ for all $1 \le i \le n$, and $[F^+(x, y, d)]_{2n+1} = 1$ if $x = y$, or 0 otherwise, with σ denoting the sigmoid function.

The goal of our rule reasoning module is, given an entity pair $(x, y) \in E_d \times E_d$ and a document d, to estimate a truth degree $s^{(N,L)}_{r,x,y,d}$ for each relation $r \in R^*$, indicating whether the relation can be inferred through at most N type-constrained rules with length L. For each original relation $r \in R$, the k-th rule $(1 \le k \le N)$, and the l-th rule atom $(1 \le l \le L)$, the intermediate truth degree $s^{(k,l)}_{r,x,y,d}$ is defined as follows:

$$s^{(k,l)}_{r,x,y,d} = \begin{cases} s^{(k,1)}_{r,x,y,d} = \varphi^{(k,1)}_r(x)\, \varphi^{(k,l+1)}_r(y) \sum\limits_{i=1}^{2n+1} \omega^{(r,k,1)}_i [F_+(x, y, d)]_i, & l = 1 \\ s^{(k,l)}_{r,x,y,d} = \varphi^{(k,l+1)}_r(y) \sum\limits_{i=1}^{2n+1} \omega^{(r,k,l)}_i \sum\limits_{(z, r_i, y) \in E_d \times R_* \times E_d} s^{(k,l-1)}_{r,x,z,d} [F_+(z, y, d)]_i, & 2 \le l \le L \end{cases} \tag{4}$$

where $\omega_{r,k,l} \in [0, 1]^{2n+1}$ is the predicate selection weight of the l-th atom in the k-th rule, normalized by Softmax to approximate one-hot, simulating the predicate selection process.

$\varphi^{(k,l)}_r(u)$ is a type constraint function representing the score that entity u satisfies specific type conditions:

$$\varphi^{(k,l)}_r(u) = \sigma_{01}\left(\alpha^{(r,k,l)} \sum_{i=1}^{m} h^{(r,k,l)}_i \mathbb{I}\left((u, Type, c_i) \in G_{type}\right) + \beta^{(r,k,l)} \sum_{i=1}^{2n} h^{(r,k,l)}_{i+m} \rho^u_{r_i}\right) \tag{5}$$

where $\sigma_{01}(x) = \max(\min(x, 1), 0)$, $h_{r,k,l} \in [0, 1]^{m+2n}$ is trainable type selection weights, and $\rho^u_{r_i} = V_u^\top B_{r_i}$ denotes the interaction between entity u and relation r_i. The parameters $\alpha^{(r,k,l)}$ and $\beta^{(r,k,l)}$ control whether explicit and implicit type constraints are applied.

The ultimate truth degree is calculated by aggregating the intermediate degrees of N rules:

$$s^{(N,L)}_{r,x,y,d} = \sum_{k=1}^{N} \alpha^{(k)}_r \cdot s^{(k,L)}_{r,x,y,d} \tag{6}$$

Where $\alpha^{(k)}_r \in [-1, 1]$ is the confidence of rule k, normalized by the Tanh activation function.

Then, we define the final logit prediction by combining output logits from the original DocRE model with the ultimate truth degrees from the type-constrained rule reasoning module:

$$\phi_r(x, y, d) = [F(x, y, d)]_r + s_{r,x,y,d}^{(N,L)} \tag{7}$$

3. Experiments

We uniformly denote the enhanced model as TyRaL-X, where X represents the name of the original DocRE model. Table 1 shows the experimental results of TyRaL on the DWIE dataset. The results indicate that TyRaL achieves stable and significant performance on all integrated DocRE backbone models, and is comprehensively superior to the original models in F1 and Ign F1 metrics, demonstrating good generality and robustness. Compared with the current state-of-the-art rule-enhanced methods, CaDRL and JMRL, TyRaL introduces key innovations in logical modeling. CaDRL relies on differentiable chain-like rule learning to improve logical consistency, while JMRL alleviates the error propagation problem through a joint training mechanism. In contrast, TyRaL proposes more refined type-constrained rules, significantly expanding the expressive power of the rules and enabling the capture of more fine-grained semantic constraints and structural relationships between entity types—rules of this kind have not been systematically modeled in existing methods. In our experiments, we adopt the F1 metric. However, some relational facts appear in both the training and the dev/test sets. As a result, a model may memorize these relations during training and achieve artificially high performance on the dev/test set, introducing evaluation bias. Such overlap is inevitable, since many common relational facts are likely to occur across different documents. Therefore, we also report the F1 scores after excluding those relational facts shared by the training and dev/test sets, which we denote as Ign F1.

Table 1
Comparison results on the DWIE dataset.

Method	PLM	Dev		Test	
		F1 (%)	Ign F1 (%)	F1 (%)	Ign F1 (%)
BiLSTM [6]	GloVe	39.66	32.14	43.54	33.88
MILR-BiLSTM [1]	GloVe	41.22 ±1.56	34.05 ±1.91	44.65 ±1.91	35.09 ±1.21
CaDRL-BiLSTM [3]	GloVe	44.02 ±4.36	38.26 ±6.16	51.43 ±7.89	42.77 ±8.89
JMRL-BiLSTM [4]	GloVe	43.68 ±4.02	37.88 ±5.74	50.70 ±7.16	42.68 ±8.80
TyRaL-BiLSTM (ours)	GloVe	**44.51** ±3.75	**39.10** ±5.30	**52.35** ±6.90	**44.23** ±8.10
GAIN [7]	$\text{BERT}_{\text{base}}$	63.81	58.89	67.45	61.36
MILR-GAIN [1]	$\text{BERT}_{\text{base}}$	65.85 ±2.04	61.22 ±2.33	69.23 ±1.78	62.77 ±1.41
CaDRL-GAIN [3]	$\text{BERT}_{\text{base}}$	66.49 ±2.68	63.51 ±4.62	70.22 ±2.77	66.63 ±5.27
JMRL-GAIN [4]	$\text{BERT}_{\text{base}}$	66.03 ±2.22	61.62 ±2.73	69.66 ±2.21	64.59 ±3.23
TyRaL-GAIN (ours)	$\text{BERT}_{\text{base}}$	**66.85** ±2.05	**62.47** ±2.40	**70.42** ±2.02	**65.35** ±3.10
ATLOP [8]	$\text{BERT}_{\text{base}}$	69.87	63.37	75.13	67.29
MILR-ATLOP [1]	$\text{BERT}_{\text{base}}$	72.05 ±2.97	67.18 ±3.81	76.51 ±1.38	69.84 ±2.55
CaDRL-ATLOP [3]	$\text{BERT}_{\text{base}}$	74.02 ±4.15	68.32 ±4.95	78.36 ±3.23	71.52 ±4.23
JMRL-ATLOP [4]	$\text{BERT}_{\text{base}}$	73.91 ±4.04	68.41 ±5.04	77.85 ±2.72	70.92 ±3.63
TyRaL-ATLOP (ours)	$\text{BERT}_{\text{base}}$	**74.76** ±3.88	**69.05** ±4.61	**78.55** ±2.51	**71.65** ±3.30

4. Limitation

While our study has made some progress, several limitations remain. First, the experiments were conducted exclusively on the DWIE dataset, which raises concerns about the generalizability of the findings to other domains and datasets. In addition, the current evaluation relies primarily on quantitative metrics and lacks case studies. We plan to address these limitations in future work.

5. Conclusion

In this poster, we propose an end-to-end learning framework, TyRaL, featuring a type-constrained rule reasoning module that simulates logical rules to enhance reasoning ability. Experiments on the DWIE dataset demonstrate its effectiveness and superiority. Future work will explore integrating logical constraints into large language models to discover more accurate and generalizable rules.

Acknowledgments

This work was supported by the Project of Science and Technology Research and Development Plan of China Railway Corporation (N2023J044).

Declaration on Generative AI

During the preparation of this work, we used ChatGPT in order to: Grammar and spelling check. After using this tool, we reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] S. Fan, S. Mo, J. Niu, Boosting document-level relation extraction by mining and injecting logical rules, in: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, 2022, pp. 10311–10323.

[2] Y. Liu, Z. Zhu, X. Zhang, Z. Feng, D. Chen, Y. Li, Document-level relationship extraction by bidirectional constraints of beta rules, in: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 2023, pp. 2256–2266.

[3] K. Zhang, P. Wu, B. Yu, K. Wu, A. Zheng, X. Huang, C. Zhu, M. Peng, H. Zan, Y. Song, Cadrl: Document-level relation extraction via context-aware differentiable rule learning, in: Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 8272–8284.

[4] K. Qi, J. Du, H. Wan, End-to-end learning of logical rules for enhancing document-level relation extraction, in: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 7247–7263.

[5] K. Zaporojets, J. Deleu, C. Develder, T. Demeester, Dwie: An entity-centric dataset for multi-task document-level information extraction, Information Processing & Management 58 (2021) 102563.

[6] Y. Yao, D. Ye, P. Li, X. Han, Y. Lin, Z. Liu, Z. Liu, L. Huang, J. Zhou, M. Sun, Docred: A large-scale document-level relation extraction dataset, arXiv preprint arXiv:1906.06127 (2019).

[7] S. Zeng, R. Xu, B. Chang, L. Li, Double graph based reasoning for document-level relation extraction, arXiv preprint arXiv:2009.13752 (2020).

[8] W. Zhou, K. Huang, T. Ma, J. Huang, Document-level relation extraction with adaptive thresholding and localized context pooling, in: Proceedings of the AAAI conference on artificial intelligence, volume 35, 2021, pp. 14612–14620.