

Adaptive Ordered Information Extraction with Deep Reinforcement Learning

Wenhao Huang[◇], Jiaqing Liang^{♡†}, Zhixu Li[◇], Yanghua Xiao^{◇♠†}, Chuanjun Ji[♣]

[◇]Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University

[♡]School of Data Science, Fudan University

[♠]Fudan-Aishu Cognitive Intelligence Joint Research Center, [♣]DataGrand Inc.

whhuang21@m.fudan.edu.cn

{liangjiaqing, zhixuli, shawyh}@fudan.edu.cn

jichuanjun@datagrand.com

Abstract

Information extraction (IE) has been studied extensively. The existing methods always follow a fixed extraction order for complex IE tasks with multiple elements to be extracted in one instance such as event extraction. However, we conduct experiments on several complex IE datasets and observe that different extraction orders can significantly affect the extraction results for a great portion of instances, and the ratio of sentences that are sensitive to extraction orders increases dramatically with the complexity of the IE task. Therefore, this paper proposes a novel adaptive ordered IE paradigm to find the optimal element extraction order for different instances, so as to achieve the best extraction results. We also propose a reinforcement learning (RL) based framework to generate optimal extraction order for each instance dynamically. Additionally, we propose a co-training framework adapted to RL to mitigate the exposure bias during the extractor training phase. Extensive experiments conducted on several public datasets demonstrate that our proposed method can beat previous methods and effectively improve the performance of various IE tasks, especially for complex ones.¹

1 Introduction

Information Extraction (IE) has been studied extensively over the past few decades (Grishman, 2019). With the rapid development of pre-trained language models, simple IE tasks such as named entity recognition (Nadeau and Sekine, 2007; Li et al., 2020a) have been well solved. However, complex IE tasks with multiple elements to be extracted such as relation extraction (Pawara et al., 2017) and event extraction (Hogenboom et al., 2011) still need further exploration.

Traditional IE methods always follow a *static* extraction manner, i.e. with a pre-defined fixed

[†]Corresponding authors.

¹Resources of this paper can be found at <https://github.com/EZ-hwh/AutoExtraction>

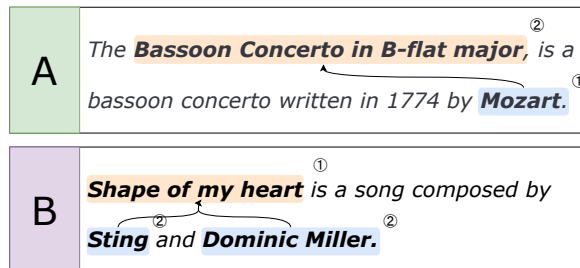


Figure 1: An example of complicated information extraction with different extraction order.

Dataset	#Ins.	#Sens ins.	Ratio
SKE21(Xie et al., 2021)	1150	100	8.70%
NYT(Riedel et al., 2010)	5000	476	9.52%
NYT10-HRL(Takanobu et al., 2019)	4006	525	13.11%
DuIE(Li et al., 2019)	15000	3618	24.12%
DuEE(Li et al., 2020b)	1492	605	40.55%
HacRED(Cheng et al., 2021)	1500	971	64.73%

Table 1: The statistic of sensitive instances on different datasets, where an instance is sensitive if different extraction orders produces different extraction results with the same model.

element order. For instance, in relation extraction task, (Xie et al., 2021) recognizes the relation and then extract the subject and object independently. (Luan et al., 2019) extracts all entity together and recognize the relation between them pair-wisely. (Wang et al., 2020) formulates joint extraction as a token pair linking problem, which follows implicit predefined order. (Lu et al., 2022) designs a unified structured extraction language to encode different IE structures, where the generation order defines the extraction order. These extraction-based IE methods assume that multiple elements of the same instance are independent of each other and thus will not affect each other’s extraction results.

However, we find that *different extraction orders highly affect the extraction results*. For example, in the A instance in Fig. 1, if we first extract the *Bassoon Concerto in B-flat major*, it will be difficult for succeeding element extraction because the

entity is long-tail. Instead, if we extract *Mozart* first, it would be much easier to extract the concerto because *Mozart* appears frequently in the corpus and he composed lots of concerto. According to our observation experiments conducted on several popular relation extraction and event extraction datasets as listed in Table 1, a significant proportion of sentences are sensitive to the extraction order, i.e., different extraction order produces different extraction results. The ratio of sensitive sentences increases dramatically with the complexity of the task on different datasets.

What’s worse, *the optimal extraction order may not be the same for different instances of the same relation or event type*. For example, for the *composer* relation, we should extract the composer *Mozart* first in A instance, but extract the song *Shape of my heart* first in B instance in Fig 1 because of its frequent appearance in corpus.

Based on the observations above, the *static* extraction paradigm following a pre-defined extraction order is insufficient to achieve reliable extraction. This motivates us to propose a *dynamic* extraction paradigm, which aims to assign an optimal element extraction order for each instance adaptively. It is nontrivial to dynamically find optimal extraction order for every instance in a dataset. On one hand, the optimal extraction order of an instance depends on not only its schema (relation or event type), but also the context of the sentence. On the other hand, multiple rounds of decisions are required to generate the optimal extraction order, where the decision of each step depends on not only the schema and sentence context, but also the extraction results of previous steps.

In this paper, we propose to adopt value-based reinforcement learning (Mnih et al., 2015) in determining the optimal extraction order for elements of an instance. Particularly, in deciding the next extraction element for an instance, every of its unextracted elements will be evaluated with a potential benefit score, which is calculated with a BERT-based model. Then, the one with the highest potential benefit score will be selected as the next extraction object. In addition, to mitigate the exposure bias that emerges during the RL-based extractor training phase, a co-training framework is adopted which can simulate the inference environment of the extraction order decision agent to help enhance its performance. It is worth mentioning that our method focuses on generating the optimal

extraction order, which is model agnostic and can be applied to various extraction paradigms.

Our main contributions are summarized below:

- First, we propose the extraction order assignment problem in complicated IE task, which can effectively affect the extractor and the extraction result.
- Second, we propose an RL based framework that dynamically generates an optimal extraction order for each sentence, which is model agnostic and can effectively guide the model towards better extraction performance.
- Third, we adopt a co-training framework for RL to alleviate the exposure bias from the extractor, which can simulate the inference environment of the extraction order decision agent to help enhance its performance.
- Fourth, the experiments conducted on several public datasets show that our method outperforms the state-of-the-art extraction models.

2 Related Work

Pipeline Information Extraction (IE) methods split the extraction process into several sub-tasks and optimize each of them. They rely on the task definition so the framework varies for different IE tasks. For relation extraction task, (Wei et al., 2020; Xie et al., 2021; Li et al., 2021) gradually extract subject, relation and object from the sentence in different order. For event extraction task, (Yang et al., 2018; Sheng et al., 2021; Yang et al., 2019) first recognize the event type and trigger and then extract the arguments with sequence tagging model or machine comprehension model.

Joint IE methods combine two or more extraction processes into one stage. Graph-based methods are the mainstream joint IE framework. They recognize the entity or text span and build a graph with co-reference, relation (Wadden et al., 2019; Luan et al., 2019), entity similarity (Xu et al., 2021) or sentence (co-occurrence) (Zhu et al., 2021). Through information propagation on the graph, they better encode the sentence and document and then decode the edge to build the final sub-graph. Generation-based IE methods are another paradigm for joint extraction, (Cabot and Navigli, 2021; Ye et al., 2021) for relation extraction task, (Zheng et al., 2019; Hsu et al., 2022; Du et al., 2021) for




Figure 2: An example of Complicated Information Extraction

event extraction, (Lu et al., 2022) for unified information extraction, they all serialize the structured extraction results into sentence or pre-defined template in a fixed order. Apart from works above, (Wang et al., 2020; Shang et al., 2022) propose one-stage joint extraction framework that decode the subject and object simultaneously.

Recently, reinforcement learning (RL) has been applied to IE task. Information extraction was augmented by using RL to acquire and incorporate external evidence in (Narasimhan et al., 2016). (Feng et al., 2018; Wang, 2020) both train a RL agent for instance selecting to denoise training data obtained via distant supervision for relation extraction. (Takanobu et al., 2019) utilizes a hierarchical reinforcement learning framework to improve the connection between entity mentions and relation types. (Zeng et al., 2019) first considers extraction order of relational facts in a sentence and then trains a sequence-to-sequence model with RL.

3 Overview

3.1 Task Definition

Fig. 2 gives an example of complicated information extraction process, which first recognizes the schema and then extracts the argument arg_i for the corresponding role role_i . Generally speaking, the task can be split into two sub-tasks: relation (event) detection and entity (argument) extraction. And we formulate the second task as multi argument extraction task. Given an instance s , the relation/event type rel and the corresponding pre-defined schema $\langle rel, \text{role}_1, \text{role}_2, \dots, \text{role}_n \rangle$, our goal is to find all the arguments in s and fill them in their corresponding roles in schema.

3.2 Solution Framework

In this work, we model the complicated information extraction as a multi-step argument extraction task. In this setting, only a role in the schema will be extracted from instance once. With the help of

extractor that can extract the arguments given the additional information and role name, we extract all arguments and fill the roles step by step.

Though there are various roles in complicated IE task, the difficulty of extracting them are completely different. For example, the role *Reference_point_time* in Fig 2 indicates the time in the context and it can be extracted without further information. Other roles like *Supplier_Consumer*, however, can not be identified with a single role name, so they should be scheduled for extraction later. We hope that the extractor can first extract the simplest role, then the next simplest one and etc. By incrementally adding the previously extracted information, the extractor can keep a good performance on extracting the difficult ones.

To achieve the goal, we need to arrange a reasonable extraction order for the extractor. However, it is hard to specify the whole extraction order once because it depends on not only the schema and context, but also the previous extracted arguments. So we regard the extraction order decision as a *Markov decision process*, where we can dynamically decide the extraction order in multi-round. Clearly, reinforcement learning is a natural way to handle this modeling. We adopt the double deep Q-network (DQN) with prioritized replay buffer as RL agent.

4 Framework

4.1 Extractor

To handle the extraction tasks with different extraction order, we have to use a powerful extractor. GlobalPointer, proposed by (Su et al., 2022), is an ideal choice as it can identify both nested and non-nested entities, and even output the scores of the entities.

We first construct the sequence consisting of the extracted elements, role name and the sentence. For an input sequence with N tokens, the BERT based encoder outputs a vector sequence $[\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N]$. Following the computation of attention matrix, we use a one-head self-attention to compute the matrix as the output of the decoder. More specifically, we first convert the vectors \mathbf{h}_i to vectors \mathbf{q}_i and \mathbf{k}_i with two linear layers.

$$\begin{aligned} \mathbf{q}_i &= \mathbf{W}^q \mathbf{h}_i + \mathbf{b}^q, \\ \mathbf{k}_i &= \mathbf{W}^k \mathbf{h}_i + \mathbf{b}^k, \end{aligned} \quad (1)$$

where \mathbf{W} and \mathbf{b} are parameters of the linear layers. Then we compute the scores of each spans with

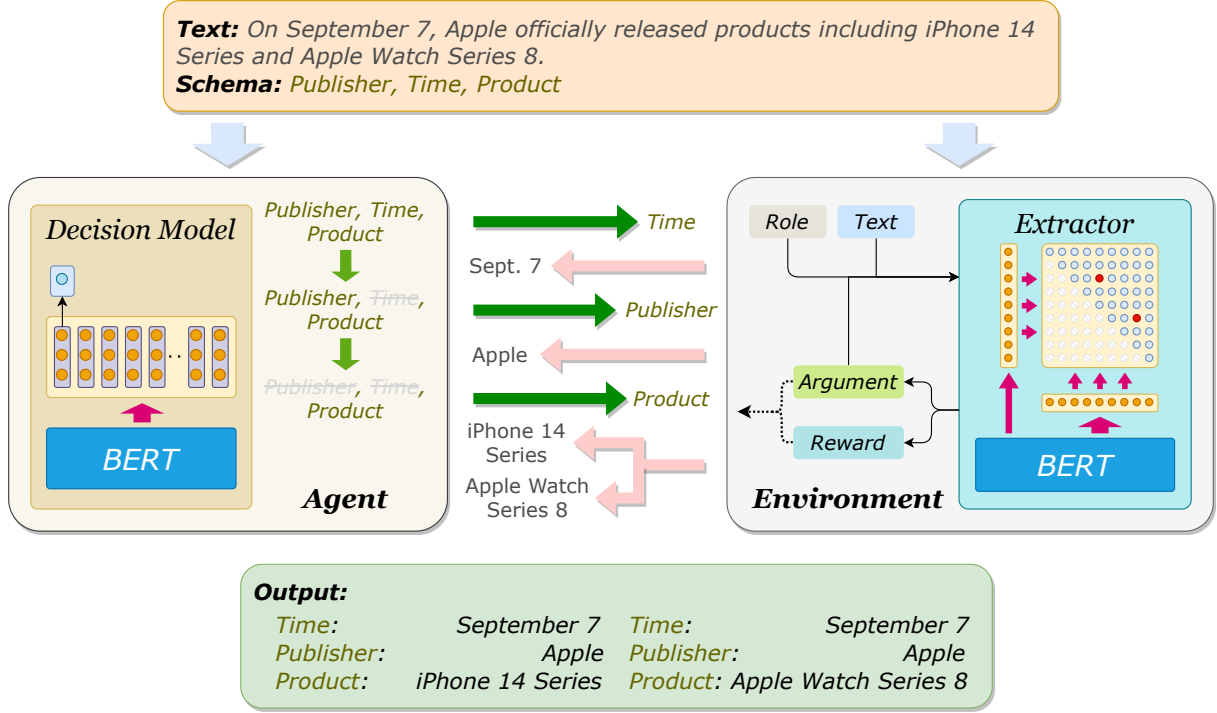


Figure 3: Reinforcement learning framework of adaptive ordered decision for IE.

the with relative position embeddings (RoPE) proposed by Reformer (Kitaev et al., 2020). The transformation matrix \mathbf{R}_i , which compose of sine and cosine function, satisfy the property that $\mathbf{R}_i^\top \mathbf{R}_j = \mathbf{R}_{j-i}$. By introducing relative positions embeddings, GlobalPointer is more sensitive to the length and span of entities, so it can better distinguish real entities.

$$\begin{aligned}
 s_\alpha(i, j) &= (\mathbf{R}_i \mathbf{q}_i)^\top (\mathbf{R}_j \mathbf{k}_j) \\
 &= \mathbf{q}_i^\top \mathbf{R}_i^\top \mathbf{R}_j \mathbf{k}_j \\
 &= \mathbf{q}_i^\top \mathbf{R}_{j-i} \mathbf{k}_j
 \end{aligned} \quad (2)$$

To better solve the label imbalance problem, we use the following loss to train our extraction model.

$$\begin{aligned}
 L &= \log \left(1 + \sum_{(i,j) \in P_\alpha} e^{-s_\alpha(i,j)} \right) + \\
 &\log \left(1 + \sum_{(i,j) \in Q_\alpha} e^{s_\alpha(i,j)} \right)
 \end{aligned} \quad (3)$$

where P_α is the head-tail set of all spans of query α , and Q_α is the head-tail set of the rest valid spans in text. In the decoding phase, all spans $t_{[i:j]}$ for which $s_\alpha(i, j) > 0$ are considered as entities that match the conditions.

For better fitting into the setting of extraction tasks that extract entities under the conditions of

the schema and role name, we enumerate all the extraction orders and match the corresponding conditions with the extraction results to construct a training instance.

4.2 MDP for extraction order

We regard the multi-role extraction order decision process as a *Markov decision process* (MDP). Fig. 3 shows the whole extraction process of an instance. In each step, the agent takes the instances and extracted arguments as input, and chooses a role unselected before as the action. The environment would take the selected role, and construct the input sequence for extractor. After collecting the extraction results to fill the role and extraction scores to assign the reward, the environment would transit to new state(s). After selecting all roles to be extracted in multiple rounds, we exchange the whole extraction history into structural output.

State We use s_t to denote the state of sentence x in extracting time step t . The state s_t consists of the extraction schema \mathcal{S} , the already extracted arguments $\hat{y}^{<t}$ and the sentence x .

$$s_t = (\mathcal{S}, \hat{y}^{<t}, x) \quad (4)$$

The state describes the extracted element in the past step. In each step, the environment would take a role selected by the agent and extracts the

corresponding arguments in the sentence with the help of extractor described in Section 4.1.

Action The action of the RL model is the next role to extract in an instance. Unlike the traditional RL environment, the action space in our model is continuously reduced at every time step. We restrict the initial action space \mathcal{A}_0 to the set of roles in the schema \mathcal{S} . After selecting role a_0 in time step 0, the extractor in environment will extract the argument and its confident score in s with the help of extractor. The action a_0 will be removed from the \mathcal{A}_0 and derive the next action space \mathcal{A}_1 . The derivation of action space can be formalized as below.

$$\mathcal{A}_t = \begin{cases} \mathcal{S}, & t = 0 \\ \mathcal{A}_{t-1} - \{a_{t-1}\}, & 0 < t < |\mathcal{S}| \end{cases} \quad (5)$$

Reward The reward is a vital important component of RL training, which is an incentive mechanism that encourages agent to better plan for the episode. For our task definition, there is a simple reward assignment. We can extract all the arguments in the sentences, and then assign a reward in terminated state to indicate whether the extracted tuple matches the golden label. But there is a transparent issue that it will majorly depends on the extractor we use. If the extractor is too strong, the results extracted following any extraction order are correct. If the extractor is too weak, the results extracted following any extraction order are incorrect. In the cases described above, different extraction order can not affect the final reward. Therefore, to better distinguish the impact of different extraction orders on the extractor, we define the reward as the score of the extraction results by the extractor. Though the extracted results recognized by the extractor for a single step is the same, the score is different given the different condition according to the Section 4.1. We regard the score as the difficulty of extracting the argument from the sentence. An extracted argument with high score indicates that it is easy to be extracted. The reward of our RL environment can be described as below.

$$\mathcal{R}(s, a) = \text{Extractor}_{score}(a|s) \quad (6)$$

where s stands for the state and a stands for the role that will be extracted chosen by the agent.

4.3 Double DQN

For traditional DQN(Mnih et al., 2013), the learning object follows the Bellman equation as below.

$$Q(s, a) = \mathcal{R}(s, a) + \gamma \cdot \max_{a' \in \mathcal{A}} Q(s', a') \quad (7)$$

Algorithm 1: The full details of our training phase for the Double DQN agent with ϵ -greedy exploration

Input : \mathcal{D} -empty replay buffer; θ -initial network parameters; θ^- -copy of θ
Input : N_b -training batch size; N^- -target network replacement frequency

```

1 for epoch = 1, ..., E do
2   | Sample instances s, S from the dataset.
3   | Nstep = #number of roles in the S
4   | for t = 1, ..., Nstep do
5     | p ← Random(0, 1)
6     | if p < 1 −  $\epsilon$  then
7       | at ← arg maxa Q(st, a;  $\theta$ )
8     | else
9       | at ← Random-Sample(At)
10    | end
11    | st+1, rt ← Transition(st, at)
12    | Store transition (st, at, rt, st+1) in D
13    | Sample random mini batch of Nb
14    | transitions (st, at, rt, st+1) from D
15    | if st+1 = done then
16      | yt = rt
17    | else
18      | yt = rt +  $\gamma$  maxa' Q(st+1, a';  $\theta^-$ )
19    | end
20    | Update parameter  $\theta$  on the loss
21    |  $\mathcal{L}(\theta) = (y_t - Q(s_t, a_t; \theta))^2$ 
22    | Replace target parameters  $\theta^- \leftarrow \theta$ 
23    | every N- steps
24  | end
25 end

```

In our task setting, the agent would choose a unextracted role and the environment would return the extracted argument with the corresponding score. Since the extractor will extract all the corresponding entities that meet the conditions at once, it is possible to extract zero to multiple answers at one time, and each extracted result will form a separated state. Due to the splitting of the state, we need to make corresponding adjustments to the original Learning object. Inspired by (Tavakoli et al., 2018), we introduce a new learning object adapted to the branching reinforcement learning environment by replacing the Q values of the next state with the

average value of the next state’s list.

$$Q(s, a) = \mathcal{R}(s, a) + \gamma \frac{1}{|S_{(s,a)}|} \sum_{s' \in S_{(s,a)}} \max_{a' \in \mathcal{A}} Q(s', a') \quad (8)$$

where $S_{(s,a)}$ is the set of the following states derived from the state s with action a , and γ is the discount factor.

To avoid suffering from the overestimation of the action values, we adopt the Double DQN (DDQN) algorithm (Van Hasselt et al., 2016) that uses the current Q-network to select the next greedy action, but evaluates it using the target. At the same time, to enable more efficient learning from the experience transitions, we adopt the prioritized experience replay buffer (Schaul et al., 2015) to replay important experience transitions.

We define the loss to be the expected value of the mean squared TD error.

$$\mathcal{L} = \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} [y - Q(s, a)]^2 \quad (9)$$

where \mathcal{D} denotes the prioritized experience replay buffer and y denotes the estimated valued by the target network.

To evaluate the value of (state, action) pair, we use Transformer based model as encoder that can take the state and action as input and encode the pair of every (state, action) pair. Specifically, we use BERT (Devlin et al., 2018) for English and RoBERTa (Liu et al., 2019) for Chinese. Formally, for an input state $\mathbf{s}_t = [t_1, t_2, \dots, t_N]$ and action $\mathbf{a}_t = [a_1, \dots, a_M]$, where the action is the candidate extraction role name, we form the sequence $x = [[\text{CLS}], \mathbf{a}_t, [\text{SEP}], \mathbf{s}_t, [\text{SEP}]]$. The BERT encoder converts these tokens into hidden vector $[\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{M+N}]$, where \mathbf{h}_i is a d -dimension vector and $d = 768$ in the Transformer based structure.

To evaluate the $Q(s, a)$ value for the corresponding state and action, we take \mathbf{h}_0 , which is the encoded vector of the first token $[\text{CLS}]$ as the representation of the state-action pair. The final output of the value evaluation module \hat{y} is define in Eq.

$$\hat{y} = \mathbf{W}\mathbf{h}_0 + \mathbf{b} \quad (10)$$

where \mathbf{W} and \mathbf{b} are trainable model parameters, representing weights and bias of the linear transformation.

4.4 Co-training framework

Beside the extraction order our agent decide would affect the final extraction order, the extractor in the

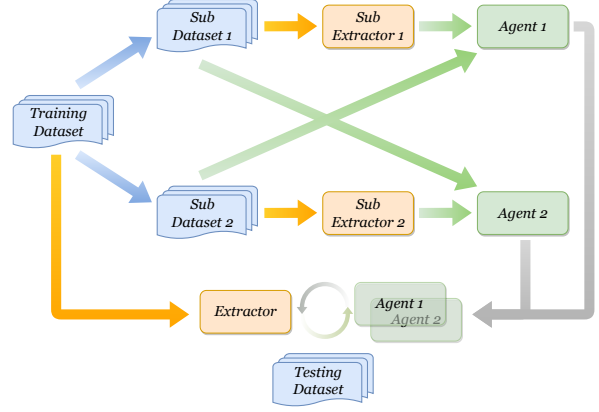


Figure 4: Framework of co-training. Blue arrow represents the training data division process. Yellow arrow represents the extractor training process. Green arrow represents the RL agent training process.

environment also matters. The argument extracted in one step not only affect the extracted tuple, but also affect the decision of the agent. However, there is a big different between the training phase and inference phase. In the training phase, the agent explores in the environment that the extractor has a good performance, which will extract the argument with high score. In the inference phase, however, the capacity of the extractor would be reduced because of the migration of the dataset from training to testing. To confirm that the agent would works in inference phase, we need to ensure similarity between the training environment and the inference environment.

We proposed a **Co-training** framework as shown in Fig 4 to simulate the environment in the testing phase. We first split the training set into two piece, which are used to train two sub extractor. Then we crossover the sub dataset and sub extractor and build two training environments. In each environment, the extractor is trained on the other piece of sub training set and extracts arguments from the sentence that it never meets. We train two agent with two environment separately. By introducing the co-training framework, we confirm the same setting on the training and inference environment. In the inference phase, we can build the environment with the test set and the extractor trained on the whole training set. Through combining the decision that two agent make, the argument in the sentences can be extracted step by step.

	NYT			NYT10-HRL			HacRED			SKE21		
	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1
NovelTagging # (Zheng et al., 2017)	62.4	31.7	42.0	59.3	38.1	46.4	30.51	2.91	5.31	-	-	-
CoType # (Ren et al., 2017)	42.3	51.1	46.3	48.6	38.6	43.0	-	-	-	-	-	-
CopyR # (Zeng et al., 2018)	61.0	56.6	58.7	56.9	45.2	50.4	13.11	9.64	11.12	-	-	-
HRL # (Takanobu et al., 2019)	-	-	-	71.4	58.6	64.4	-	-	-	-	-	-
CasRel # (Wei et al., 2020)	89.7	89.5	89.6	77.7	68.8	73.0	<u>55.24</u>	<u>43.78</u>	<u>48.85</u>	-	-	-
TPLinker # (Wang et al., 2020)	<u>91.3</u>	<u>92.5</u>	<u>91.9</u>	<u>81.19</u>	65.41	72.45	-	-	-	-	-	-
RERE# (Xie et al., 2021)	-	-	-	75.45	<u>72.50</u>	<u>73.95</u>	-	-	-	-	-	-
CasRel(Wei et al., 2020) *	87.77	90.79	89.25	76.59	68.90	72.54	57.19	44.99	50.36	87.21	75.23	80.78
TPLinker(Wang et al., 2020) *	88.61	92.29	90.41	80.37	65.11	71.94	58.96	55.78	57.33	83.86	84.77	84.32
RERE(Xie et al., 2021) *	85.68	92.45	88.93	74.43	68.46	71.32	46.42	61.37	52.86	85.65	86.37	86.01
Adaptive Order	88.92	92.83	90.84	77.21	69.81	73.32	58.36	72.43	64.64	87.99	86.87	87.42

Table 2: The main evaluation results of different models on NYT, NYT10-HRL, HacRED and SKE21. The results with only one decimal are quoted from (Wei et al., 2020). The methods with * are based on our re-implementation. The methods with # denote that the metrics are partially matched. Best exact (partial) match F1 scores are marked **bold** (underlined).

	NYT			NYT10-HRL			HacRED			DuIE			DuEE		
	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1
Sequence order	92.83	96.19	94.48	84.66	85.42	85.04	58.18	79.76	67.28	52.14	62.07	56.68	72.70	70.47	71.57
Random order	93.05	95.97	94.49	84.42	85.08	84.75	57.82	80.02	67.13	52.53	62.11	56.92	72.86	71.10	71.97
Adaptive order	93.15	96.14	94.62	84.59	85.66	85.13	57.88	81.92	67.62	53.33	62.87	57.71	73.86	72.14	72.99

Table 3: Extraction Result on different dataset with different extraction order decision.

	SKE21			HacRED			DuIE			DuEE		
	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1	Prec.	Reca.	F1
Sequence order	94.74	81.08	87.38	66.62	78.67	72.15	44.88	48.67	46.70	72.41	69.30	70.82
Random order	92.75	80.63	86.27	65.94	79.41	72.05	44.25	47.91	46.01	72.73	70.89	71.80
Adaptive order	93.37	82.43	87.56	66.31	80.56	72.74	46.60	50.76	48.59	73.88	71.69	72.76

Table 4: Extraction Result on complicated extraction case with different extraction order decision. HacRED and SKE21 are both tested on at least 5 triples of the same relation.

Table 5: Extraction Result on complicated extraction case with different extraction order decision. DuIE is restricted in at least 3 roles and DuEE is restricted in at least 5 roles.

5 Experiments

5.1 Datasets

We evaluate our methods on several public and accessible complicated information extraction datasets, including NYT, NYT10-HRL, SKE21, HacRED, DuIE, DuEE, which are challenging for many novel extraction methods. We give brief introduction to these dataset in appendix.

5.2 Comparing methods and Metrics

We compare our methods with several models on the same dataset, including NovelTagging(Zheng et al., 2017), CoType(Ren et al., 2017), HRL(Takanobu et al., 2019), and recent PLM-based model CasRel(Wei et al., 2020), TPLinker(Wang et al., 2020) and RERE(Xie et al.,

2021).

We choose the exact match that an extracted relational triple (*subject, relation, object*) is regarded as correct only if the relation and the spans of both subject and object are correct. We report the standard micro precision (Prec.), recall (Reca.) and F1 scores for the relation extraction experiments. And for the event extraction task (only DuEE in our experiment), we report the word-level metrics, which considers the correctness of the arguments in word level. We give the detail of this metric in appendix.

5.3 Main Results

Because we only consider the extraction order assignment in every instance, we add a classification module to first recognize the relation in sentence,

Instance	Extraction Process	Extraction Result
<p>Instance #1 On October 14, Redmi officially released two new entry-level smartphones, Redmi 8 and Redmi 8A.</p>	<pre> graph LR Time[Time] --> Publisher[Publisher] Publisher --> Product[Product] Time --> T1[October 14] T1 --> P1[Redmi] P1 --> P2[Redmi 8] P1 --> P3[Redmi 8A] </pre>	<p>Time: October, Publisher: Redmi, Product: Redmi 8A; Time: October, Publisher: Redmi, Product: Redmi 8A</p>
<p>Instance #2 Huami Technology releases Midong Health Watch and AMAZFIT Smart Watch 2.</p>	<pre> graph LR Publisher[Publisher] --> Product[Product] Product --> Time[Time] Publisher --> P1[Midong Health Watch] Publisher --> P2[AMAZFIT Smart Watch 2] P1 --> T1[None] P2 --> T2[None] </pre>	<p>Time: [None], Publisher: Huami Technology, Product: Midong Health Watch; Time: [None], Publisher: Midong Health Watch, Product: AMAZFIT Smart Watch 2</p>

Table 6: Instance of extracting complicated schema through dynamically assigning extraction order.

and then extract the subject and object with the extraction order RL agents assign. According to the result in Table 2, compared to the main stream relation extraction methods, ours method achieves improvement on both precision, recall and F1 score.

5.4 Extraction order

To demonstrate the effectiveness of our methods, we also conduct an experiment on different extraction order decision strategy on more challenging event argument extraction task. In this experiment, we mainly focus on the performance of extraction result in different extraction order, so ground-truth schema (relation) for every instance would be offered beforehand, we only test the correctness on the roles. We also provide the result of extracting arguments in a pre-defined order and random order, as baselines. Table 3 shows the result of different extraction order in different dataset, and our method achieve the best in every dataset. Compare to the standard relation extraction task, our method perform better on complex information extraction task (DuIE and DuEE).

5.5 Complicated Extraction settings

To further demonstrate the advantage of dynamically order extraction decision by RL, we conduct experiment with more complex extraction tasks, which contains more tuples or more arguments. For the former, we limit the minimum number of extraction tuples, and we limit the minimum number of extraction roles for the latter.

Table 4 and 5 show that compared to fixed or-

der extraction or random order extraction method, our framework has a more significant improvement over the original metrics. This is intuitive and reasonable, the extractor is more sensitive to the extraction order in more complex sentence. Besides, compare to the Table 4 and 5, we can find that our method improves the latter settings more significantly. It is because increasing the number of tuples does not increase the length of the extracted path, but only increase the difficulty of single-step extraction by the extractor. In contrast, the increase of the role number leads to an increase in the length of the extraction path, which makes the decision of extraction order more difficult. The results once again prove that the extraction order matters in complicated extraction.

5.6 Case Study

With taking RL agent into consideration, we can easily observe the extraction order in different instance. Table 6 show the instances that the extraction process in different sentence. Though two instances share the same event schema *Product release*, the RL agent assign different extraction order dynamically. The first sentence contains an obvious element of time, while the second does not, so our methods put the extraction order of time from the first to the last. The case strongly demonstrates the effectiveness of our method.

6 Conclusion

In this paper, we propose a novel adaptive order IE paradigm to find the optimal element extraction

order for different instances. We propose a RL-based framework to generate optimal extraction for each instance dynamically and a co-training framework adapted to RL to alleviate the exposure from the extractor. Extensive experiments show that our proposed method can beat previous methods and effectively improve the performance of various IE tasks, especially for complex ones.

Acknowledgements

This research is funded by National Key Research and Development Project (No. 2020AAA0109302), National Natural Science Foundation of China (No. 62102095, 62072323), Shanghai Science and Technology Innovation Action Plan (No. 22511104700, 22511105902), Shanghai Municipal Science and Technology Major Project (No.2021SHZDZX0103), and Science and Technology Commission of Shanghai Municipality Grant (No. 22511105902).

Limitations

Despite the remarkable improvement on complicated information extraction, there are still some limits of our method.

First, due to the multi-round argument extraction modeling, we discard the parallelism in element extraction. Furthermore, the MDP process interacting with the DQN further increases the computational load of the extraction process. So compare to other methods, our framework is relative slow at the inference stage.

Second, though our framework can be easily adapted to different extraction task with different schema, we still need an extra module helping identifying the relations (event types) in the instance beforehand. Because of the difference task definition and modeling (extraction task and classification task), although recognizing them potentially implies order decision making, they are beyond the scope of this paper.

Ethics statement

We hereby declare that all authors of this article are aware of and adhere to the provided ACL Code of Ethics and honor the code of conduct.

Use of Human Annotations Human annotations are only utilized in the early stages of methodological research to assess the feasibility of the proposed solution. All annotators have provided consent for

the use of their data for research purposes. We guarantee the security of all annotators throughout the annotation process, and they are justly remunerated according to local standards. Human annotations are not employed during the evaluation of our method.

Risks The datasets used in this paper have been obtained from public sources and anonymized to protect against any offensive information. Though we have taken measures to do so, we cannot guarantee that the datasets do not contain any socially harmful or toxic language.

References

- Pere-Lluís Huguet Cabot and Roberto Navigli. 2021. Rebel: Relation extraction by end-to-end language generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2370–2381.
- Qiao Cheng, Juntao Liu, Xiaoye Qu, Jin Zhao, Jiaqing Liang, Zhefeng Wang, Baoxing Huai, Nicholas Jing Yuan, and Yanghua Xiao. 2021. Hacred: A large-scale relation extraction dataset toward hard cases in practical applications. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2819–2831.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Xinya Du, Alexander M Rush, and Claire Cardie. 2021. Template filling with generative transformers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 909–914.
- Jun Feng, Minlie Huang, Li Zhao, Yang Yang, and Xiaoyan Zhu. 2018. Reinforcement learning for relation classification from noisy data. In *Proceedings of the aaai conference on artificial intelligence*, volume 32.
- Ralph Grishman. 2019. Twenty-five years of information extraction. *Natural Language Engineering*, 25(6):677–692.
- Frederik Hogenboom, Flavius Frasincar, Uzay Kaymak, and Franciska De Jong. 2011. An overview of event extraction from text. *DeRiVE@ ISWC*, pages 48–57.
- I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2022. Degree: A data-efficient generation-based event extraction model. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1890–1908.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*.
- Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. 2020a. A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering*, 34(1):50–70.
- Shuangjie Li, Wei He, Yabing Shi, Wenbin Jiang, Haijin Liang, Ye Jiang, Yang Zhang, Yajuan Lyu, and Yong Zhu. 2019. Duie: A large-scale chinese dataset for information extraction. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 791–800. Springer.
- Xianming Li, Xiaotian Luo, Chenghao Dong, Daichuan Yang, Beidi Luan, and Zhen He. 2021. Tdeer: An efficient translating decoding schema for joint extraction of entities and relations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8055–8064.
- Xinyu Li, Fayuan Li, Lu Pan, Yuguang Chen, Weihua Peng, Quan Wang, Yajuan Lyu, and Yong Zhu. 2020b. Duee: a large-scale dataset for chinese event extraction in real-world scenarios. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 534–545. Springer.
- 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. *arXiv preprint arXiv:1907.11692*.
- Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun, and Hua Wu. 2022. Unified structure generation for universal information extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5755–5772.
- Yi Luan, Dave Wadden, Luheng He, Amy Shah, Mari Ostendorf, and Hannaneh Hajishirzi. 2019. A general framework for information extraction using dynamic span graphs. 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 3036–3046.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- David Nadeau and Satoshi Sekine. 2007. A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1):3–26.
- Karthik Narasimhan, Adam Yala, and Regina Barzilay. 2016. Improving information extraction by acquiring external evidence with reinforcement learning. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2355–2365.
- Sachin Pawara, Girish K Palshikara, and Pushpak Bhat-tacharyab. 2017. Relation extraction: A survey. *arXiv preprint arXiv:1712.05191*.

- Xiang Ren, Zeqiu Wu, Wenqi He, Meng Qu, Clare R Voss, Heng Ji, Tarek F Abdelzaher, and Jiawei Han. 2017. Cotype: Joint extraction of typed entities and relations with knowledge bases. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1015–1024.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163. Springer.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Yu-Ming Shang, Heyan Huang, and Xianling Mao. 2022. Onerel: Joint entity and relation extraction with one module in one step. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11285–11293.
- Jiawei Sheng, Shu Guo, Bowen Yu, Qian Li, Yiming Hei, Lihong Wang, Tingwen Liu, and Hongbo Xu. 2021. Casee: A joint learning framework with cascade decoding for overlapping event extraction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 164–174.
- Jianlin Su, Ahmed Murtadha, Shengfeng Pan, Jing Hou, Jun Sun, Wanwei Huang, Bo Wen, and Yunfeng Liu. 2022. Global pointer: Novel efficient span-based approach for named entity recognition. *arXiv preprint arXiv:2208.03054*.
- Ryuichi Takanobu, Tianyang Zhang, Jiexi Liu, and Minlie Huang. 2019. A hierarchical framework for relation extraction with reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 7072–7079.
- Arash Tavakoli, Fabio Pardo, and Petar Kormushev. 2018. Action branching architectures for deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. 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)*, pages 5784–5789.
- Jianing Wang. 2020. Rh-net: Improving neural relation extraction via reinforcement learning and hierarchical relational searching. *arXiv preprint arXiv:2010.14255*.
- Yucheng Wang, Bowen Yu, Yueyang Zhang, Tingwen Liu, Hongsong Zhu, and Limin Sun. 2020. Tplinker: Single-stage joint extraction of entities and relations through token pair linking. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1572–1582.
- Zhepei Wei, Jianlin Su, Yue Wang, Yuan Tian, and Yi Chang. 2020. A novel cascade binary tagging framework for relational triple extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1476–1488.
- Chenhao Xie, Jiaqing Liang, Jingping Liu, Chengsong Huang, Wenhao Huang, and Yanghua Xiao. 2021. Revisiting the negative data of distantly supervised 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 1: Long Papers)*, pages 3572–3581.
- Runxin Xu, Tianyu Liu, Lei Li, and Baobao Chang. 2021. Document-level event extraction via heterogeneous graph-based interaction model with a tracker. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3533–3546.
- Hang Yang, Yubo Chen, Kang Liu, Yang Xiao, and Jun Zhao. 2018. Dcfee: A document-level chinese financial event extraction system based on automatically labeled training data. In *Proceedings of ACL 2018, System Demonstrations*, pages 50–55.
- Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. Exploring pre-trained language models for event extraction and generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5284–5294.
- Hongbin Ye, Ningyu Zhang, Shumin Deng, Mosha Chen, Chuanqi Tan, Fei Huang, and Huajun Chen. 2021. Contrastive triple extraction with generative transformer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 14257–14265.
- Xiangrong Zeng, Shizhu He, Daojian Zeng, Kang Liu, Shengping Liu, and Jun Zhao. 2019. Learning the extraction order of multiple relational facts in a sentence with reinforcement learning. 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)*, pages 367–377.
- Xiangrong Zeng, Daojian Zeng, Shizhu He, Kang Liu, and Jun Zhao. 2018. Extracting relational facts by an end-to-end neural model with copy mechanism. In *Proceedings of the 56th Annual Meeting of the*

Association for Computational Linguistics (Volume 1: Long Papers), pages 506–514.

Shun Zheng, Wei Cao, Wei Xu, and Jiang Bian. 2019. Doc2edag: An end-to-end document-level framework for chinese financial event extraction. 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)*, pages 337–346.

Suncong Zheng, Feng Wang, Hongyun Bao, Yuexing Hao, Peng Zhou, and Bo Xu. 2017. Joint extraction of entities and relations based on a novel tagging scheme. *arXiv preprint arXiv:1706.05075*.

Tong Zhu, Xiaoye Qu, Wenliang Chen, Zhefeng Wang, Baoxing Huai, Nicholas Jing Yuan, and Min Zhang. 2021. Efficient document-level event extraction via pseudo-trigger-aware pruned complete graph. *arXiv preprint arXiv:2112.06013*.

A Dataset Introduction

NYT NYT (Riedel et al., 2010) is the very early version of NYT series dataset. It is based on the articles in New York Times, contains 66, 194 text and 24 types of relation.

NYT10-HRL NYT10-HRL is an improved version of the origin NYT dataset. After preprocessed by HRL (Takanobu et al., 2019) that removed training relation not occurring in the testing test and "NA" sentence, NYT10-HRL contains 70, 339 sentences for training and 4, 006 sentences for test.

DuIE DuIE(Li et al., 2019) is a Chinese information extraction dataset. It contains more than 210,000 sentences and 48 pre-defined schema gathered from Baidu Encyclopedia, Baidu Tieba and Baidu Information Stream. Compared to the previous version, it contains 5 Multiple-O-values schema (Schema with multiple object slots and values), which greatly increase the difficulty of the task.

HacRED HacRED (Cheng et al., 2021)² is a novel challenging extraction dataset, which consists of 65,225 relational facts annotated from 9,231 wiki documents with sufficient and diverse hard cases

DuEE DuEE (Li et al., 2020b) is a Chinese document-level event extraction dataset. It contains 11,224 documents categorized into 65 event types, along with 41,520 event arguments mapped to 121 argument roles, which is the largest Chinese EE dataset.

SKE21 SKE2019 is the largest Chinese dataset available for relation extraction publish by Baidu, which contains 194,747 sentences for training. (Xie et al., 2021) manually labeled 1,150 sentences from test set with 2,765 annotated triples.

B Experiment details

Our experiments are conducted on single RTX3090 GPUs. All deep models, including the extraction model and decision model, are implemented using the PyTorch framework. We initialized the model with the bert-base-cased and chinese-roberta-wwm-ext respectively, training 10 epochs for both the extractor and classifier. As for the reinforcement learning module, we set the buffer size to 100,000 and target network update step at 200. We trained 5

²<https://github.com/qiaojiim/hacred>

Hyper-parameters	Extractor	Decision Model
Batch size	48	32
Learning rate	1e-5	1e-5
Max length	512	512

Table 7: Hyper-parameters for training extractor and agent model respectively.

epochs on the standard relation extraction task and 10 epochs for DuIE and DuEE. Additionally, the exploration parameter ϵ was initialized at 0.9 and the discount factor γ was set to 0.5 within the DQN framework. To encourage continuous exploration, the exploration rate became 0.9 times of its own in a certain steps until it reached 0.05. We calculated the necessary updating steps using the following formula.

$$\begin{aligned} \#Total_Steps &= \#Epoch \times \#Datasets \\ \#Update_steps &= \lfloor \frac{\#Total_Steps * \log 0.95}{\log 0.05 - \log 0.9} \rfloor \end{aligned} \quad (11)$$

We employed AdamW optimizer as the optimizer, with using linear scheduling with warming up proportion 10%. More details are listed in Table 7.

C Word-level Metric

We offer the word-level precision, recall and F1 score calculation formula. First we calculate the F1 score of a single arguments.

$$Set_a = Set_p \cap Set_g, \quad (12)$$

$$TextP = \frac{|Set_a|}{|Set_p|}, TextR = \frac{|Set_a|}{|Set_g|} \quad (13)$$

$$TextF1 = \frac{2 \cdot TextP \cdot TextR}{TextP + TextR} \quad (14)$$

where Set_p denotes the set of words in predict argument, and Set_g denotes the set of words in ground truth argument. We set $TextF1 = 1$ if both Set_a and Set_b are empty set.

For every predict event E_p and ground truth event E_g , we calculate their match score through calculate their mean $TextF1$ score.

$$Score(E_p, E_g) = \frac{1}{N} \sum_N TextF1(arg_p, arg_g) \quad (15)$$

where N denotes the number of the roles that event schema contains.

We separately calculate the precision for prediction event and recall for ground truth event.

$$\begin{aligned} Prec_E(E_p) &= \max_{e \in \mathbf{E}_g} Score(E_p, e) \\ Reca_E(E_g) &= \max_{e \in \mathbf{E}_g} Score(e, E_g) \end{aligned} \quad (16)$$

where \mathbf{E}_p and \mathbf{E}_g are the predict and golden event set of the same instance.

Finally, we regard the mean precision and recall of every prediction and golden annotations in the whole test set as our Prec. and Reca.

$$\begin{aligned} Prec &= \frac{1}{N_p} \sum_{E_p} Prec(E_p) \\ Reca &= \frac{1}{N_g} \sum_{E_g} Reca(E_g) \\ F1 &= \frac{2 \cdot Prec \cdot Reca}{Prec + Reca} \end{aligned} \quad (17)$$

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations section behind the main body.
- A2. Did you discuss any potential risks of your work?
Ethics statement section behind main body
- A3. Do the abstract and introduction summarize the paper's main claims?
Section 1, 2, 3, 4, 5, Limitations
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 1, 2, 4, A

- B1. Did you cite the creators of artifacts you used?
Section 1, 2, 4.1, 4.3, A
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section A
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 3
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
In this paper, all datasets used for experiments are available on website. The datasets have been anonymized protect against any offensive information.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section A
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section A

C Did you run computational experiments?

Section 5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section B

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section B

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 5, A

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section B

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.