

REKnow: Enhanced Knowledge for Joint Entity and Relation Extraction

Sheng Zhang, Patrick Ng, Zhiguo Wang, Bing Xiang
AWS AI Labs

{zshe, patricng, zhiguow, bxiang} @amazon.com

Abstract

Relation extraction is an important but challenging task that aims to extract all hidden relational facts from the text. With the development of deep language models, relation extraction methods have achieved good performance on various benchmarks. However, we observe two shortcomings of previous methods: first, there is no unified framework that works well under various relation extraction settings; second, effectively utilizing external knowledge as background information is absent. In this work, we propose a knowledge-enhanced generative model to mitigate these two issues. Our generative model is a unified framework to sequentially generate relational triplets under various relation extraction settings and explicitly utilizes relevant knowledge from Knowledge Graph (KG) to resolve ambiguities. Our model achieves superior performance on multiple benchmarks and settings, including WebNLG, NYT10, and TACRED.

1 Introduction

Numerous relational facts are hidden in the emerging corpus. The extraction of relational facts is beneficial to a diverse set of downstream problems. For example, relational facts facilitate the development of dialogue systems with lower labor cost (Wen et al., 2016; Dhingra et al., 2016). Knowledge Graph (KG) construction leverages relational fact extraction of everyday news to keep its KGs up-to-date (Shi and Weninger, 2018). Even in the field of e-commerce, relation extraction of product descriptions helps build category information for products (Dong et al., 2020).

Although relation extraction (RE) has been widely studied, we observe two shortcomings of previous work. First, there is no unified framework for different settings of relation extraction tasks. Table 1 summarizes the popular relational fact extraction tasks. The formats of input and output vary

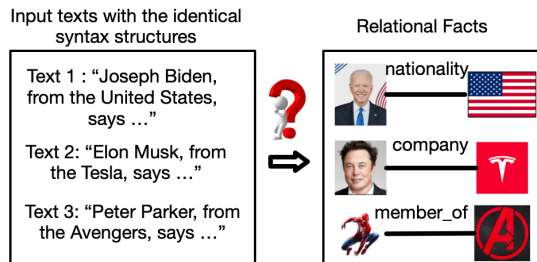


Figure 1: The three examples have identical syntax structures, but their corresponding relational facts are different. Therefore, background knowledge is very crucial to correctly extract relational facts.

by task. In the Entity Type aware Relation Classification task, models take the plain text and entity positions and entity types as input and predict relation types for all entity pairs. For the Relation Classification task, entity types are not provided. In the Joint Entity and Relation Extraction task, models only take plain texts as input and predict all possible entities and relations among them. Previous work usually designs different model frameworks to handle these tasks separately (Joshi et al., 2020; Yamada et al., 2020; Soares et al., 2019; Wei et al., 2019; Cabot and Navigli, 2021). Also, directly applying Joint Entity and Relation Extraction models to Relation Classification tasks would lead to a bad performance due to the lack of fully usage of entity information compared with task-specific models.

Second, conventional relation extraction models do not utilize external knowledge to resolve relational ambiguities. However, background knowledge of entities can be very crucial for models to figure out the right relations. For the three examples in Figure 1, they have the same syntax structure: "[entity1] from the [entity2], says". The main information for models to predict different relation types for them is the background knowledge of these involved entities. For example, the facts of "Joseph Biden is a human" and "the United States is a country" determine the "nationality" relation. The facts of "Elon Musk is

Task name	Entity Type aware Relation Classification
Text †	Born in London in 1939 the son of a Greek tycoon, Negroponte grew up in Britain, Switzerland and the United States
Entity Positions	Negroponte: (51, 61) United States: (103, 116)
Entity Types	Negroponte: PERSON United States: LOCATION
Task name	Relation Classification
Text ‡	The cutting machine contains 13 circular blades mounted on a cutting axis.
Entity Positions	machine: (12, 19) blades: (41, 47)
Task name	Joint Entity and Relation Extraction
Text §	It is Japan’s second-biggest automaker, behind Toyota and ahead of Nissan.

Table 1: Different type of relational fact extraction tasks. † example is obtained from TACRED (Zhang et al., 2017); ‡ from SemEval2010 (Hendrickx et al., 2019); § from NYT10 (Riedel et al., 2010). See Section 2 for more details.

a human" and "Tesla is an enterprise" determine the "company" relation. The facts of "Peter Parker is a fictional human" and "Avengers is a fictional organization" determine the "member_of" relation. Lack of the background knowledge, it is very difficult for models to disambiguate.

Our solution to the first issue is to utilize a generative model framework. The popular sequence-to-sequence frameworks are naturally equipped to solve multiple tasks with variations in input and output formats. Moreover, previous studies Ye et al. (2020) and Josifoski et al. (2021) have shown the effectiveness of the generative framework for each specific RE task. In this work, we aim to unify all types of RE tasks with a generative model. To address the second issue, we propose to leverage the well-organized information from the knowledge graph (KG). We employ an entity linking model to connect all entities within the plain text with their corresponding entries in KG, then explicitly feed the knowledge of these entities into the RE model for relation prediction. Previous methods (Zhang et al., 2019; Wang et al., 2020a; Liu et al., 2020a) have shown that injecting KG into models during pre-training can boost the performance in downstream tasks. However, involving KG during pre-training can be very expensive, and the model may suffer the catastrophic forgetting issue. Differently, our method directly feeds relevant knowledge into the downstream model, thus can be more efficient and effective.

In this paper, we propose a unified generative framework for **Relation Extraction with Knowledge enhancement (REKnow)**. Our generative framework is trained to sequentially generate relational facts token-by-token. Also, our framework enhances the input plain text with relevant KG information. Our experiments show that the proposed framework significantly boosts the extraction performance on multiple benchmarks and task settings. To our best knowledge, this is the first work to integrate KG into a generative framework for relational fact extraction. In summary, our contributions are threefold:

- We construct a novel knowledge-enhanced generative framework for relational fact extraction task.
- We provide KG grounding predictions for public relation extraction benchmark datasets.¹
- Our generative model achieves superior performance on multiple benchmarks and settings, including ACE2005, NYT10, and TACRED.

2 Problem Definition

All tasks in Table 1 are formally formulated as follows. We denote the plain text as $\mathcal{S} = \{w_i\}_{i=1,2,\dots,N}$, where w_i represents a word, and N is the text length. The knowledge provided in a dataset \mathcal{D} is denoted as $\mathcal{K}_{\mathcal{D}}$, such as pre-defined relation set \mathcal{R} , entity positions \mathcal{E}_{pos} , and entity types \mathcal{E}_{type} . The goal is to find all relation triples $\{\langle s_i, r_i, o_i \rangle\}_{i=1,2,\dots,T}$ in text \mathcal{S} , where s_i and o_i are subject and object entities in \mathcal{S} , r_i belonging to \mathcal{R} is the relation between subject and object, and T is the number of relation triples. The RE problem can be formulated as follows:

$$f(\mathcal{S}|\mathcal{K}_{\mathcal{D}}) = \{\langle s_i, r_i, o_i \rangle\}_{i=1,2,\dots,T}. \quad (1)$$

For different tasks, $\mathcal{K}_{\mathcal{D}}$ may contain a different set of knowledge. We consider three tasks: (1) Entity Type aware Relation Classification, where $\mathcal{K}_{\mathcal{D}}^{ETRC} = \{\mathcal{R}, \mathcal{E}_{pos}, \mathcal{E}_{type}\}$; (2) Relation Classification, where $\mathcal{K}_{\mathcal{D}}^{RC} = \{\mathcal{R}, \mathcal{E}_{pos}\}$; and (3) Joint Relation and Entity Extraction $\mathcal{K}_{\mathcal{D}}^{JREE} = \{\mathcal{R}\}$. We denote tasks (1) and (2) as entity position-aware case and task (3) as entity position-absent case for simplicity in the following sections.

¹Code, models and grounded KG will be available at [xx.xx](https://github.com/xx.xx).

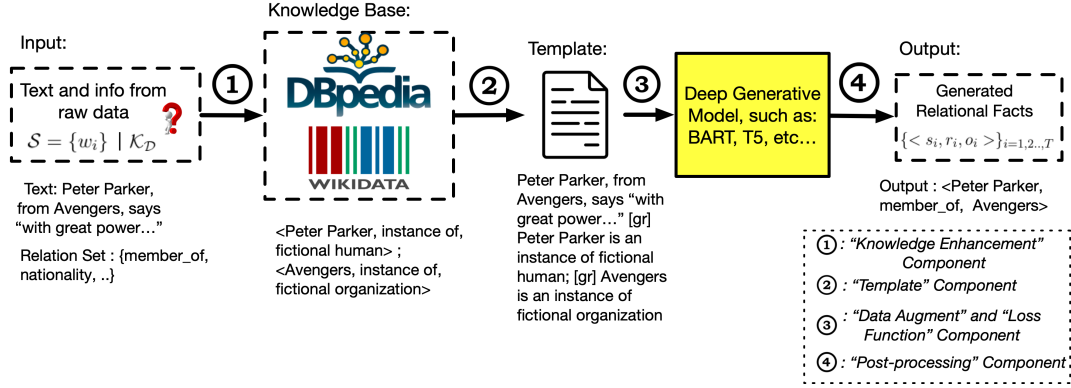


Figure 2: Overall REKnow framework. See Section 3 for more details about each component.

3 Relation Extraction with Knowledge-Enhanced Generative Model

Overview Figure 2 shows the workflow of our method. The relational fact extraction is modeled as a seq2seq generative model, where knowledge-enhanced text is passed to the encoder, and the relational facts are generated sequentially from the decoder. We use text $\mathcal{S} = \text{"Peter Parker, from Avengers, says ..."}$ and relation set $\mathcal{R} = \{member_of, nationality, \dots\}$ as a running example. First, the plain text \mathcal{S} is grounded to the public Knowledge Base (KB) Wikidata² to obtain relevant background knowledge \mathcal{K}_E , where the subscript E refers to external knowledge. In this example, we ground *"Peter Parker"* to <https://www.wikidata.org/wiki/Q23991129> and *"Avengers"* to <https://www.wikidata.org/wiki/Q322646>. We use the "instance of" property in Wikidata to retrieve the background knowledge, e.g. $\langle \text{Peter Parker, instance of, fictional human} \rangle$ and $\langle \text{Avengers, instance of, fictional organization} \rangle$. Then, we combine the plain text \mathcal{S} , knowledge from the dataset \mathcal{K}_D , and external knowledge \mathcal{K}_E via a pre-defined template \mathcal{T} . Finally, the combined text is passed into a generative language model, such as BART (Lewis et al., 2019) or T5 (Raffel et al., 2019), and the model directly generates the target relational triples token by token, e.g. *"Peter Parker member_of Avengers"*.

Knowledge-Enhancement We leverage an entity linking model to ground texts to knowledge bases. In the Relation Classification task, the span of an entity \mathcal{E}_{pos} is given. Therefore, we apply the bi-encoder entity linking method (BLINK) for

entity linking (Wu et al., 2019). BLINK uses a two-stage method that links entities to Wikipedia, based on fine-tuned BERT architectures. First, BLINK performs the retrieval step in a dense space defined by a bi-encoder that independently embeds the entity $e \in \mathcal{E}_{pos}$ and the text \mathcal{S} . In the second stage, each candidate in the KG is examined with a cross-encoder ranking model by concatenating the entity and the text. An approximate nearest neighbor search method FAISS (Johnson et al., 2019) is applied to search the large-scale KG efficiently. In the task of Joint Entity and Relation Extraction, the span of an entity is absent. We apply an efficient one-pass end-to-end Entity Linking method (ELQ) (Li et al., 2020). ELQ follows a similar architecture as BLINK, but the bi-encoder jointly performs mention detection and linking in one pass.

After we obtain the grounded Wikipedia entries for entities in the text, we use the public-source pywikibot³ to retrieve the entity knowledge from Wikidata, i.e. the knowledge provided by the "instance of" property. If multiple "instances of" properties exist, we select the one with the highest frequency over the training data. Formally, for the entity position-aware case, entity linking is formulated as:

$$\mathcal{K}_E = \{\mathcal{E}_{type}\} = EL_{BLINK}(\mathcal{S}|\{\mathcal{E}_{pos}\}). \quad (2)$$

The entity position-absent case is formulated as:

$$\mathcal{K}_E = \{\mathcal{E}_{pos}, \mathcal{E}_{type}\} = EL_{ELQ}(\mathcal{S}). \quad (3)$$

More implementation details will be discussed in Section 5.2.

Template We consider two templates to incorporate entity-related knowledge for both entity position-aware and entity position-absent cases in

²<https://www.wikidata.org/>

³<https://github.com/wikimedia/pywikibot>

our application. For text \mathcal{S} , we denote the joint knowledge from dataset and external knowledge as $\mathcal{K}_\mathcal{S} = \mathcal{K}_\mathcal{D} \cup \mathcal{K}_E$. For the entity position-aware case, we construct the template as

$$\mathcal{T}_1(\mathcal{S}|\mathcal{K}_\mathcal{S}) = \{w_1, \dots, w_{i-1}, [es], w_i^{s_t}, \dots, w_j^{s_t}, [gr], \mathbf{t}_{s_t}, w_{j+1}, \dots, w_{k-1}, [es], w_k^{o_t}, \dots, w_l^{o_t}, [gr], \mathbf{t}_{o_t}, w_{l+1}, \dots, w_N\},$$

where $[es]$ and $[gr]$ are special tokens to denote entity mention start and grounded entity information position, respectively. In the above case, the token span for subject s_t ranges from i -th to j -th token in original text, similar notation is applied to object o_t . Entity type for entity e is denoted as \mathbf{t}_e .

For the entity position-absent case, we construct the template by concatenating the entity knowledge information at the end of the original text as follows,

$$\mathcal{T}_2(\mathcal{S}|\mathcal{K}_\mathcal{S}) = \{w_1, w_2, \dots, w_N, [gr], \mathbf{e}_1 \text{ is an instance of } \mathbf{t}_{e_1}, \dots, [gr], \mathbf{e}_M \text{ is an instance of } \mathbf{t}_{e_M} \dots\},$$

where $e_i, i = 1, 2, \dots, M$ are the retrieved entities via the entity linking system, and \mathbf{t}_{e_i} are corresponding entity knowledge from public KB⁴.

Data Augmentation If multiple triples exist in a single input text, the order of target relational triples may affect the performance in the seq2seq model, which is known as *exposure bias* (Ranzato et al., 2015). In previous relation extraction work, researchers attempted to alleviate exposure bias by replacing maximum likelihood estimation (MLE) loss with reinforcement learning-based loss (Zeng et al., 2019) or unordered multi-tree structure-based loss (Zhang et al., 2020). In our work, we adapt the data augmentation strategy by shuffling the order of target relational triples and adding the shuffled triples to the original dataset during the training phase. Our ablation study in Section 5.3 shows that the data augmentation strategy can significantly improve the performance.

Loss Function We apply a generative model, such as BART (Lewis et al., 2019) and T5 (Raffel et al., 2019), to input text $\mathbf{x} = \mathcal{T}(\mathcal{S}|\mathcal{K}_\mathcal{S})$. The output of triples $\mathbf{y} =$

⁴In the T5 paper, authors added a prefix to distinguish different tasks during the pre-training stage. Therefore, if T5 pre-trained model is adopted, we add the prefix "summary:" in the template for our fine-tuning stage.

$\{s_1 \mathbf{r}_1 \mathbf{o}_1; s_2 \mathbf{r}_2 \mathbf{o}_2; \dots; s_T \mathbf{r}_T \mathbf{o}_T\}$ is generated in the text format in a single pass, where token ";" is used to separate different triples in the output. The generative model needs to model the conditional distribution $p(\mathbf{y}|\mathbf{x})$. The output tokens come from a fixed size vocabulary V . The distribution $p(\mathbf{y}|\mathbf{x})$ can be represented by the conditional probability of the next token given the previous tokens and input text:

$$p(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{|\mathbf{y}|} p(y_i|y_{<i}, \mathbf{x}),$$

where y_i denotes the i -th token in \mathbf{y} . The subject s_t and object o_t are tokens shown in the text \mathcal{S} , and \mathbf{r}_t is a relation from the relation set \mathcal{R} , i.e. $s_t, o_t \subset \mathcal{S}$, $\mathbf{r}_t \in \mathcal{R}$ for $t = 1, 2, \dots, T$. For training set $D = \{(\mathbf{x}^i, \mathbf{y}^i)_{i=1,2,\dots,|D|}\}$, we train the generative model with parameter θ using MLE by minimizing the negative log-likelihood over D :

$$\mathcal{L}(D) = - \sum_{i=1}^{|D|} \log p_\theta(\mathbf{y}^i|\mathbf{x}^i).$$

Post-processing The generated output \mathbf{y} is in the text format, which should be transformed to standard triples format $\{< s_i, \mathbf{r}_i, \mathbf{o}_i >\}_{i=1,2,\dots,T}$ for evaluation. Therefore, we apply some post-processing steps, including *split*, *delete* and *replace*, to transform the generated output. Token ";" is used to separate triples in the training stage, hence the generated output can be *split* by token ";" to obtain different triples in text format. For each generated triple, the relation should belong to the relation set \mathcal{R} . If no valid relation is found in the generated triple, we will *delete* the corresponding triple. In Relation Classification task, since the gold entities \mathcal{E} are given, we match the generated entities to the gold entities using Levenshtein similarity (Navarro, 2001). In the Relation Classification task, we would *delete* the generated triple if the generated entities show low similarity to all golden entities, i.e. $\max_{e_i \in \mathcal{E}} L_{sim}(\hat{e}, e_i) < \epsilon$ where L_{sim} denotes Levenshtein similarity, \hat{e} is the generated entity, and ϵ is a threshold⁵. Otherwise, we *replace* the generated entity with the gold entity with the highest Levenshtein similarity. In the Joint Entity and Relation Classification task, the generated entity should be shown in the original text. Therefore, we match the generated entities to the set of spans in the text, which is motivated by span-level

⁵We use $\epsilon = 0.85$ in our experiments

prediction in Name Entity Recognition (NER) research (Li et al., 2019; Yu et al., 2020; Fu et al., 2021). We *replace* a generated entity with the sub-span of text in the text with the highest Levenshtein similarity.

4 Related Work

Relational Fact Extraction We group the relational fact extraction tasks with pre-defined relation set into three categories based on different levels of information provided in the data and introduce the related methods accordingly. In the first category, for data with entity span as well as corresponding entity types, SpanBERT (Joshi et al., 2020) built the relation classification model by replacing the subject and object entities with their NER tags such as “[CLS][SUBJ-PER] was born in [OBJ-LOC], Michigan, ...”. A linear classifier was added on top of the [CLS] token to predict the relation type in their applications. LUKE (Yamada et al., 2020) constructed an additional Entity Type Embedding layer in their model construction to utilize the entity type information.

In the second category, data come with an entity span. It is the most popular and standard format in the RE task. Traditional methods (Kambhatla, 2004; Zhou et al., 2005) extracted the relation between entities based on feature engineering, which incorporated semantic, syntactic, or lexical features in the text. After the fast development of the deep language model, like BERT (Devlin et al., 2018), BERT-Entity (Soares et al., 2019) utilized the entity span information by adding a mention pooling layer on top of the BERT model.

In the third category, only text is provided in the data. In such a case, models were developed to extract both entity and relation. For pipeline-based methods, NER techniques (Chiu and Nichols, 2016; Huang et al., 2015; Fu et al., 2021) were used to detect the span of entities, then relation classification models are applied based on the detected entities. Researchers (Li and Ji, 2014; Miwa and Sasaki, 2014) also noticed that sharing parameters between NER and RE models are important in the pipeline model. For example, Pure (Zhong and Chen, 2020) built separate encoders for entities and relations, where the side product of the entity model is fed into the relation model as input. For joint extraction methods, Copy-RE (Zeng et al., 2018) built the relation extraction model based on the copy mechanism. Tplinker (Wang et al., 2020b) jointly

extracted the relation and entity in novel handshaking tagging schema, which are constructed to deal with the issue of overlapping entities. In this work, we propose a unified framework for the relation extraction task for all cases above. Compared to a unified framework for NER (Yan et al., 2021) tasks, our work focuses on relational triplet extraction.

Knowledge-Enhanced Information Extraction

With the development of KG construction, many NLP tasks benefit from the external well-organized knowledge in KG. We group knowledge-enhanced methods in the RE task into two categories based on the knowledge-ingestion stage. The first type is to use the KG in the pre-training stage. The key challenges for such knowledge-ingestion type are Heterogeneous Embedding Space (HES), which means the embedding spaces for KG and text would be heterogeneous, and Knowledge Noise (KN), which represents information in KG would divert the text from its correct meaning (Liu et al., 2020a). ERNIE (Zhang et al., 2019) solved the HES by proposing a pre-training task that requires the model to predict entities based on the given entity sequence rather than all entities in KGs. K-Adapter (Wang et al., 2020a) constructed a neural adapter to infuse different kinds of knowledge. REBEL (Cabot and Navigli, 2021) used BART-large as the base model to generate relation triplets, which is pretrained on Wikidata. However, since the method was trained on the Wikidata with 220 different relation types, it may introduce knowledge noise implicitly. MIUK (Li et al., 2021) leveraged the uncertainty characteristic in ProBase (Wu et al., 2012) across three views: mention, entity and concept view. Numerical studies about integrating knowledge into NLU tasks (Xu et al., 2021) showed that KG information can increase the performance by a significant amount.

Another type is to use the KG in downstream tasks directly. One of the popular implementations is distantly supervised relation extractors Mintz et al., 2009, which leveraged the information in Freebase (Bollacker et al., 2008) to extract relations. Liu et al., 2020b utilized KG to conduct information extraction via RL framework. Our method uses entities linking to automatically ground entities in the text to public KG. The experiment results show that the entity type in KG helps boost the performance in most cases.

Dataset	Data Size (train/val/test)	Ave. Triple Size (train/val/test)
TACRED	10,021/3,894/2,307	1.30/1.40/1.44
SemEval2010	6,507/1,493/2,717	1.00/1.00/1.00
NYT10	56,196/5,000/5,000	2.01/2.02/2.03
Webnlg	5,019/500/703	2.74/3.11/2.82
ACE2005	2,619/648/590	1.83/1.82/1.95

Table 2: Statistics of relational fact extraction benchmark datasets. Avg Triple Size represents the average size of triples in each input data.

5 Experiments

In this section, we compare our proposed method to the state-of-the-art of relational fact extraction methods — demonstrating its effectiveness and ability to extract relational facts under multiple tasks. Our codes will be made publicly available.

5.1 Data

Several public available RE datasets are evaluated in our numerical experiments. **TACRED** (Zhang et al., 2017) is a large supervised RE dataset, which is derived from the TAC-KBP relation set, with labels obtained via crowdsourcing. It contains 41 valid relation types as well as one null relation type and 12 entity types. Although alternate versions of TACRED have been published recently (Alt et al., 2020; Stoica et al., 2021), we use the original version of TACRED, with which the state of the art is mainly tested. **SemEval 2010 Task 8** (Hendrickx et al., 2019) is a benchmark dataset for evaluating classification of semantic relations between pairs of nominal entities, such as “part-whole”, “cause-effect”, etc. It defines 9 valid relation types as well as one "other" relation type and no entity type is provided. We follow the split setting as in (Soares et al., 2019) for numerical study. **NYT10** dataset (Riedel et al., 2010) was originally produced by distant supervision method from 1987-2007 New York Times news articles. Original data are adapted by (Zeng et al., 2018) for complicated relational triple extraction task, where EntityPairOverlap (EPO) and SingleEntityOverlap (SEO) widely existed in the dataset. **WebNLG** dataset (Gardent et al., 2017) was originally created for Natural Language Generation task. Zeng et al., 2018 adapted WebNLG for relational triple extraction task as well. In our experiment, we evaluate the proposed method on NYT10 and WebNLG datasets with whole entity span as in prior works (Wang et al., 2020b; Ye et al., 2020) other than datasets with the last word

Dataset	Method	Prec.	Rec.	F1
TACRED	ERNIE (Zhang et al., 2019)	80.0	66.1	68.0
	SpanBERT (Joshi et al., 2020)	70.8	70.9	70.8
	K-Adapter (Wang et al., 2020a)	68.9	75.4	72.0
	RoBERTa (Wang et al., 2020a)	70.2	72.4	71.3
	LUKE (Yamada et al., 2020)	70.4	75.1	72.7
	REKnow (ours)	75.7	73.6	74.6
	REKnow w/o KG	75.4	72.4	73.9
	REKnow using orig. entity type	75.6	73.0	74.3
	CR-CNN (Santos et al., 2015)	-	-	84.1
	BERT-Entity (Soares et al., 2019)	-	-	89.2
Semeval	BERT-MTB (Soares et al., 2019)	-	-	89.5
	REKnow (ours)	87.1	92.3	89.6
	REKnow w/o KG	86.0	91.5	88.7

Table 3: Main result for entity position-aware case. The top two performing models per dataset are marked in bold. In REKnow, the pre-trained model is T5-large and entity type is obtained by Equation (2).

of the entities (Zeng et al., 2018; Wei et al., 2019). **ACE2005** (Walker, 2006) is developed by the Linguistic Data Consortium (LDC) containing a collection of documents from a variety of domains including news and online forums. In the dataset, 6 relation types between the entities are provided. In our implementation, we focus on the RE task in ACE2005, hence we only keep the sentences with valid triples. The detailed statistics about all datasets are summarized in Table 2. For the evaluation, we report the standard micro Precision, Recall, and F1-score.

5.2 Model Implementation Details

All experiments are conducted with 8-cores NVIDIA Tesla V100 GPUs with 2.5 GHz (base) and 3.1 GHz (sustained all-core turbo) Intel Xeon 8175M processors.

Decoder For the BART model, Shakeri et al., 2020 found that using a variant of nucleus sampling (Holtzman et al., 2019) can increase the diversity of generated output compared to beam search. We observe similar findings in our preliminary experiments; therefore, we adopt the *Topk+Nucleus* decoding strategy in our decoding step for the BART model. To be specific, we pick top $k = 20$ tokens, and within top k , tokens with top 95% probability mass are picked in each generated step. For the T5 model, we simply adopt the greedy decoding strategy, which can generate triples well.

Entity Linking We use the pre-trained BLINK and ELQ model⁶ to align Wikipedia entity for en-

⁶<https://github.com/facebookresearch/BLINK> (MIT License)

Dataset	Method	Precision	Recall	F1
NYT	NovelTagging (Zheng et al., 2017)	32.8	30.6	31.7
	MultiHead (Bekoulis et al., 2018)	60.7	58.6	59.6
	ETL-Sapn (Yu et al., 2019)	85.5	71.7	78.0
	Tplinker (Wang et al., 2020b)	91.4	92.6	92.0
	CopyRE* † (Zeng et al., 2018)	61	56.6	58.7
	CopyMTL* † (Zeng et al., 2020)	75.7	68.7	72.0
	TANL (Paolini et al., 2021) †	-	-	90.8
	REBEL (Cabot and Navigli, 2021) †	-	-	92.0
	CGT † (Ye et al., 2020)	94.7	84.2	89.1
	REKnow (ours)	91.6	93.6	92.6
	REKnow w/o KG	90.5	92.7	91.6
Webnlg	NovelTagging (Zheng et al., 2017)	52.5	19.3	28.3
	MultiHead (Bekoulis et al., 2018)	57.5	54.1	55.7
	ETL-Sapn (Yu et al., 2019)	84.3	82.0	83.1
	Tplinker (Wang et al., 2020b)	88.9	84.5	86.7
	CopyRE* † (Zeng et al., 2018)	37.7	36.4	37.1
	CopyMTL* † (Zeng et al., 2020)	58.0	54.9	56.4
	CGT † (Ye et al., 2020)	92.9	75.6	83.4
	REKnow (ours)	87.5	86.6	87.0
	REKnow w/o KG	84.0	84.3	84.2
	ACE2005	Pure-Bb (Zhong and Chen, 2020)	-	-
Pure-Alb (Zhong and Chen, 2020)	-	-	69.0	
REKnow (ours)	71.0	65.7	68.3	
REKnow w/o KG	70.3	65.0	67.5	

Table 4: Main result for entity position-absent case. * denotes the methods that only generate the last token of the entity. † denotes the competitors are generative-based models. The top two performing models per dataset are marked in bold. In REKnow, the pre-trained model is T5-large and entity type is obtained by Equation (3).

tity span aware and absent cases, respectively. The entity linking models are trained on 5.9M entities from May 2019 English Wikipedia dump. We follow the same threshold score (-4.5) as in the original ELQ implementation, and the top 1 grounded Wikipedia entity is used. For dataset Semeval 2010, the entities are nominal, hence we use the property "subclass of" of the retrieved entity as entity types. For datasets TACRED, NYT10, WebNLG and ACE2005, we use property "instance of" to obtain corresponding entity types.

Dataset	REKnow	w/o KG	w/o Text	Ensemble
NYT	92.6	91.6	87.0	93.3
WebNLG	87.0	84.2	63.7	87.4
ACE2005	68.3	67.5	9.6	67.9

Table 5: Analysis of knowledge enhancement. "REKnow" is the proposed method. "REKnow w/o KG" means raw text is used. "REKnow w/o Text" denotes only found KG information is used. "Ensemble" means majority voting from the previous three methods. The F1 scores are reported.




Figure 3: F1 Score under different triple sizes for REKnow. The *triple size* is the number of triples hidden in a sentence.

Dataset	Found Ext. Info Ratio (train/val/test)	Entity Span Aware
TACRED	0.87/0.83/0.78	Yes
SemEval2010	0.47/0.45/0.45	Yes
NYT10	1.75/1.76/1.76	No
Webnlg	0.93/0.91/0.97	No
ACE2005	0.75/0.71/0.66	No

Table 6: Statistics of entity linking. "Found Ext. Info Ratio" represents the ratio between the size of entities in the text and the size of found external information in entity linking step. Low ratio represents few external information is found and vice versa.

5.3 Results

For Relation Classification tasks, we summarize the result in Table 3. The state-of-the-art methods are listed and compared. More introduction about the competing methods can be found in Section 4. "REKnow" denotes our proposed method. "REKnow w/o KG" represents no external information from KG is used. For TACRED dataset, "REKnow w/o KG" is 1.2 F1 score higher than LUKE (Yamada et al., 2020) method. It shows the effectiveness of our unified framework. Also, "REKnow" is 0.7 F1 scores higher than "REKnow w/o KG", which indicates that enhanced knowledge further boosts the performance. Since the entity types are already provided in TACRED, we also report the result using the original entity type in "REKnow using orig. entity type". It can be found that using the original entity type is 1.6 F1 scores higher than the LUKE method but slightly lower than that of the grounded entity information by our method. The reason is that original entity type in TACRED is a coarse-grained level entity type while WikiData can provide a fine-grained entity information. For SemEval2010, our method achieves performance comparable to one of the

Model	NYT	WebNLG	ACE2005	TACRED	Semeval
T5-large	92.6	87.0	68.3	74.6	89.6
BART-large	89.7	75.8	59.4	65.8	88.5

Table 7: Pre-trained model comparisons in proposed REKnow method for WebNLG. T5-large has better performance.

most popular methods BERT-MTB (Soares et al., 2019).

For Joint Entity and Relation Extraction tasks, we summarize the result in Table 4. For NYT dataset, we can find that our “REKnow” method outperforms the current SOTA generative-based model, for example, a 0.6 F1 score increase compared with the REBEL method, which is pretrained on a constructed dataset “CROCODILE” (Cabot and Navigli, 2021). And 3.5 F1 score higher than CGT (Ye et al., 2020) in WebNLG dataset. The performance of our method is comparable with SOTA extractive-based method Tplinker (Wang et al., 2020b). For ACE2005, “REKnow” achieves a good F1 score with 68.3, which is close to the SOTA method Pure-Alb (Zhong and Chen, 2020) which is a pipelined framework constructed for entity and relation extraction task.

Also, from Table 4, we notice that “REKnow” has the highest recall, which means the method can cover more potential triples. The performance of different triple sizes can be found in Figure 3. We can see that for NYT and WebNLG datasets, our method has a better performance when the triple size is greater than 1.

Analysis of Knowledge Enhancement In REKnow, both text and entity-related information from KG are used. To study the usefulness of external information, we re-train the model with different information source components. We summarize the results in Table 5. To be noticed that the founded external information in NYT dataset is very useful and achieve 87.0 F1 score without using the raw text. We also present the “Ensemble” result which ensembles the models trained with different information components by majority voting. It shows that in NYT and WebNLG, ensemble can further boost the performance. On the contrary, in the case of ACE2005, the external information may contain too much knowledge noise therefore lowering the performance.

To investigate the performance of entity linking, we also present the “found ext. info ratio” by calculating the ratio between the size of en-

Error reason	The truth is incomplete
Dataset	WebNLG (traceID : test_578)
Original Text	1634: The Bavarian Crisis is the sequel to The Grantville Gazettes .
Text with EL	1634: The Bavarian Crisis is the sequel to The Grantville Gazettes . [gr] 1634: The Bavarian Crisis is an instance of literary work [gr] The Grantville Gazette is an instance of literary work
Generated Triple	1634: The Bavarian crisis precededBy The Grantville Gazettes
Truth (s, r, o)	(“1634 : The Bavarian”, “precededBy”, “The Grantville Gazettes”)
Analysis	“1634: The Bavarian crisis” should be treated as a whole.
Error reason	Generative model has much wider searching space
Dataset	WebnlG (traceID : test_633)
Original Text	Bill Oddie ’s daughter is Kate Hardie .
Text with EL	Bill Oddie ’s daughter is Kate Hardie . [gr] Bill Oddie is an instance of human [gr] Kate Hardie is an instance of human
Generated Triple	Bill Oddie daughter Kate Hardie
Truth (s, r, o)	(“Bill Oddie”, “child”, “Kate Hardie”)
Analysis	“daughter” is not in the pre-defined relation set

Table 8: Example of Error Analysis. More examples can be found in Table A1 in the Appendix A

tities in the text and the size of found external information. The results for all RE benchmarks are summarized in Table 6. Low ratio represents few external information is found and vice versa. Based on the result, we can see that, for the SemEval2010 dataset, the founded ext. ratio is low, because the entities in SemEval2010 are nominal, and hard to find corresponding entity types from Wikidata. That explained the “REKnow w/o KG” is close to “REKnow” in Table 3. For the NYT10 dataset, the size of grounded entities is much larger than the size of actual entities. It means the entities widely exist in the plain text, and the many relations among entities are not included as triples in the dataset.

Ablation Study We study the comparison of whether to use the data augmentation component (Section 3) for entity position-absent case. In NYT dataset, removing data augmentation will decrease the F1 score 2.0. In WebNLG and ACE2005, no data augmentation will decrease the F1 score 9.8 and 6.1, respectively. It represents that exposure bias will affect the performance for generative-

based method and using data-augmentation strategy can alleviate the exposure bias significantly.

We also conduct an ablation study of the pre-trained model, BART and T5. The comparison result is summarized in Table 7. From the table, we can observe that T5-large is generally better than the BART-large model. To be specific, in datasets NYT and Semeval, the performances between T5 and BART are close. As for the rest datasets, the BART model decreases about 8 to 10 F1 scores compared with T5. It would be interesting to study the reason for the heterogeneous difference among datasets. We leave that as a direction for future investigation.

Error Analysis Based on the error analysis in Table 8, we notice that the common error reasons include incomplete truth labels. The triple provided in the dataset is (“1634: The Bavarian”, “precededBy”, “The Grantville Gazettes”), but “1634: The Bavarian crisis” is a book title — REKnow correctly predicted this as a whole entity. Another error is that generative model may create relation out of the pre-defined relation set. In the second error case in Table 8, “daughter” is not a pre-defined relation. This issue may be addressed by constraint generation, which needs future investigation.

6 Conclusion

We proposed a knowledge-enhanced unified generative model for a relational fact extraction task that generates triples sequentially with one pass. We showed that it outperforms previous state-of-the-art models on multiple relation extraction benchmarks.

References

- Christoph Alt, Aleksandra Gabryszak, and Leonhard Hennig. 2020. Tacred revisited: A thorough evaluation of the tacred relation extraction task. *arXiv preprint arXiv:2004.14855*.
- Giannis Bekoulis, Johannes Deleu, Thomas Demeester, and Chris Develder. 2018. Joint entity recognition and relation extraction as a multi-head selection problem. *Expert Systems with Applications*, 114:34–45.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250.
- 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.
- Jason PC Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional lstm-cnns. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- 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*.
- Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2016. Towards end-to-end reinforcement learning of dialogue agents for information access. *arXiv preprint arXiv:1609.00777*.
- Xin Luna Dong, Xiang He, Andrey Kan, Xian Li, Yan Liang, Jun Ma, Yifan Ethan Xu, Chenwei Zhang, Tong Zhao, Gabriel Blanco Saldana, et al. 2020. Autoknow: Self-driving knowledge collection for products of thousands of types. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2724–2734.
- Jinlan Fu, Xuanjing Huang, and Pengfei Liu. 2021. Spanner: Named entity re-/recognition as span prediction. *arXiv preprint arXiv:2106.00641*.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. Creating training corpora for nlg micro-planning. In *55th annual meeting of the Association for Computational Linguistics (ACL)*.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid O Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2019. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. *arXiv preprint arXiv:1911.10422*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.

- Martin Josifoski, Nicola De Cao, Maxime Peyrard, and Robert West. 2021. Genie: Generative information extraction. *arXiv preprint arXiv:2112.08340*.
- Nanda Kambhatla. 2004. Combining lexical, syntactic, and semantic features with maximum entropy models for information extraction. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 178–181.
- Arzoo Katiyar and Claire Cardie. 2017. Going out on a limb: Joint extraction of entity mentions and relations without dependency trees. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 917–928.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Belinda Z. Li, Sewon Min, Srinivasan Iyer, Yashar Mehdad, and Wen-tau Yih. 2020. [Efficient one-pass end-to-end entity linking for questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6433–6441, Online. Association for Computational Linguistics.
- Bo Li, Wei Ye, Canming Huang, and Shikun Zhang. 2021. Multi-view inference for relation extraction with uncertain knowledge. *arXiv preprint arXiv:2104.13579*.
- Qi Li and Heng Ji. 2014. Incremental joint extraction of entity mentions and relations. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 402–412.
- Xiaoya Li, Jingrong Feng, Yuxian Meng, Qinghong Han, Fei Wu, and Jiwei Li. 2019. A unified mrc framework for named entity recognition. *arXiv preprint arXiv:1910.11476*.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020a. K-bert: Enabling language representation with knowledge graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2901–2908.
- Ye Liu, Sheng Zhang, Rui Song, Suo Feng, and Yanghua Xiao. 2020b. Knowledge-guided open attribute value extraction with reinforcement learning. *arXiv preprint arXiv:2010.09189*.
- Yi Luan, Dave Wadden, Luheng He, Amy Shah, Mari Ostendorf, and Hannaneh Hajishirzi. 2019. A general framework for information extraction using dynamic span graphs. *arXiv preprint arXiv:1904.03296*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011.
- Makoto Miwa and Yutaka Sasaki. 2014. Modeling joint entity and relation extraction with table representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1858–1869.
- Gonzalo Navarro. 2001. A guided tour to approximate string matching. *ACM computing surveys (CSUR)*, 33(1):31–88.
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. Structured prediction as translation between augmented natural languages. *arXiv preprint arXiv:2101.05779*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Marc’ Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- 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.
- Cicero Nogueira dos Santos, Bing Xiang, and Bowen Zhou. 2015. Classifying relations by ranking with convolutional neural networks. *arXiv preprint arXiv:1504.06580*.
- Siamak Shakeri, Cicero Nogueira dos Santos, Henry Zhu, Patrick Ng, Feng Nan, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2020. End-to-end synthetic data generation for domain adaptation of question answering systems. *arXiv preprint arXiv:2010.06028*.
- Baoxu Shi and Tim Weninger. 2018. Open-world knowledge graph completion. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. *arXiv preprint arXiv:1906.03158*.
- George Stoica, Emmanouil Antonios Platanios, and Barnabás Póczos. 2021. Re-tacred: Addressing shortcomings of the tacred dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13843–13850.

- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. *arXiv preprint arXiv:1909.03546*.
- et al. Walker, Christopher. 2006. [Ace 2005 multilingual training corpus ldc2006t06](#). web download. philadelphia: Linguistic data consortium, 2006. ACE2005.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Guihong Cao, Daxin Jiang, Ming Zhou, et al. 2020a. K-adapter: Infusing knowledge into pre-trained models with adapters. *arXiv preprint arXiv:2002.01808*.
- Yucheng Wang, Bowen Yu, Yueyang Zhang, Tingwen Liu, Hongsong Zhu, and Limin Sun. 2020b. Tplinker: Single-stage joint extraction of entities and relations through token pair linking. *arXiv preprint arXiv:2010.13415*.
- Zhepei Wei, Jianlin Su, Yue Wang, Yuan Tian, and Yi Chang. 2019. A novel cascade binary tagging framework for relational triple extraction. *arXiv preprint arXiv:1909.03227*.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2019. Scalable zero-shot entity linking with dense entity retrieval. *arXiv preprint arXiv:1911.03814*.
- Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q Zhu. 2012. Probase: A probabilistic taxonomy for text understanding. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 481–492.
- Ruochen Xu, Yuwei Fang, Chenguang Zhu, and Michael Zeng. 2021. Does knowledge help general nlu? an empirical study. *arXiv preprint arXiv:2109.00563*.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. Luke: deep contextualized entity representations with entity-aware self-attention. *arXiv preprint arXiv:2010.01057*.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. A unified generative framework for various ner subtasks. *arXiv preprint arXiv:2106.01223*.
- Hongbin Ye, Ningyu Zhang, Shumin Deng, Mosha Chen, Chuanqi Tan, Fei Huang, and Huajun Chen. 2020. Contrastive triple extraction with generative transformer. *arXiv preprint arXiv:2009.06207*.
- Bowen Yu, Zhenyu Zhang, Xiaobo Shu, Yubin Wang, Tingwen Liu, Bin Wang, and Sujian Li. 2019. Joint extraction of entities and relations based on a novel decomposition strategy. *arXiv preprint arXiv:1909.04273*.
- Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. Named entity recognition as dependency parsing. *arXiv preprint arXiv:2005.07150*.
- Daojian Zeng, Haoran Zhang, and Qianying Liu. 2020. Copymtl: Copy mechanism for joint extraction of entities and relations with multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9507–9514.
- 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, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Haoran Zhang, Qianying Liu, Aysa Xuemo Fan, Heng Ji, Daojian Zeng, Fei Cheng, Daisuke Kawahara, and Sadao Kurohashi. 2020. Minimize exposure bias of seq2seq models in joint entity and relation extraction. *arXiv preprint arXiv:2009.07503*.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. [Position-aware attention and supervised data improve slot filling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*, pages 35–45.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. Ernie: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*.
- 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*.
- Zexuan Zhong and Danqi Chen. 2020. A frustratingly easy approach for entity and relation extraction. *arXiv preprint arXiv:2010.12812*.
- GuoDong Zhou, Jian Su, Jie Zhang, and Min Zhang. 2005. Exploring various knowledge in relation extraction. In *Proceedings of the 43rd annual meeting of the association for computational linguistics (acl'05)*, pages 427–434.

Appendix:

A Error Analysis

Based on the error analysis in Table A1, we find the common error reasons include incomplete truth labels, too wide generation space and missing entity linking result. In Case 2, the truth triplets should contain all possible relations between a borough “Manhattan” and some places including “Washington Heights”, “The Battery” and “Harlem River”. In Case 3, if an inversable relation, such as “org_alternate_names”, exists, the generated model may generate the triplet in a wrong order compared with golden label. In Case 4, generative model may create relation out of the pre-defined relation set. This issue may be addressed by constraint generation, which is under future investigation. In Case 5, the entity linking result may fail to provide enough background information. Hence the generated triple confuse with the relation between “sounds” and “species”.

B Hyper-parameter in Model Implementation

We adopt the following setting for model training:

- learning_rate : 5e-5
- max_source_length : 1024
- max_target_length : 128
- num_train_epochs : 3
- per_device_train_batch_size : 4 (for T5-large), 16 (for T5-base, T5-small and BART)
- nproc_per_node : 8
- optimizer : AdamW
- lr_scheduler_type : linear

Error reason :	The truth is incomplete	
Case 1	Dataset	Webnlg (traceID : test_578)
	Text with EL	1634 : The Bavarian crisis is the sequel to The Grantville Gazettes . [gr] 1634: The Bavarian Crisis is an instance of literary work [gr] The Grantville Gazette is an instance of literary work
	Generated Triple	1634 : The Bavarian crisis precededBy The Grantville Gazettes
	Truth (s, r, o)	(“1634 : The Bavarian”, “precededBy”, “The Grantville Gazettes”)
	Analysis	“1634 : The Bavarian crisis” should be treated as a whole part
Case 2	Dataset	NYT (traceID : test_582)
	Text with EL	Through a series of leisurely walks around the island – from the Battery to Washington Heights , and from Wall Street to the Harlem River – Lopate ruminates on Manhattan ’s history , architecture and inhabitants . [gr] The Battery (Manhattan) is an instance of urban park [gr] Washington Heights, Manhattan is an instance of neighborhood [gr] Wall Street is an instance of street [gr] Harlem River is an instance of strait [gr] Manhattan is an instance of borough of New York City
	Generated Triple	Washington Heights neighborhood_of Manhattan ; Harlem River neighborhood_of Manhattan ; Manhattan contains Washington Heights ; Manhattan contains Harlem River ; Manhattan contains the Battery;
	Truth (s, r, o)	(“Manhattan”, “contains”, “Washington Heights”), (“Washington Heights”, “neighborhood_of”, “Manhattan”)
	Analysis	Relations between “Manhattan” and “The Battery” or “Harlem River” should be included in the truth
Case 3	Dataset	TACRED (traceID : test_1013)
	Text with EL	The two companies were preparing to announce that AIG had agreed to sell [ms] American Life Insurance Co [gr] business , better known as [ms] Alico [gr] business , for 68 billion dollars in cash and 87 billion in MetLife equity , the report said .
	Generated Triple	American Life Insurance Co org_alternate_names Alico ;
	Truth (s, r, o)	(“Alico”, “org_alternate_names”, “American Life Insurance Co”)
	Analysis	generative model is not sensitive to inverseable triple like : org’s alternate name
Error reason :	Generative model has too wider searching space	
Case 4	Dataset	Webnlg (traceID : test_633)
	Text with EL	Bill Oddie ’s daughter is Kate Hardie . [gr] Bill Oddie is an instance of human [gr] Kate Hardie is an instance of human
	Generated Triple	Bill Oddie daughter Kate Hardie
	Truth (s, r, o)	(“Bill Oddie”, “child”, “Kate Hardie”)
	Analysis	“daughter” is not in the pre-defined relation set
Error reason :	Entity linking is missing	
Case 5	Dataset	semeval (traceID : test_876)
	Text with EL	the word “song ” is used to describe the pattern of regular and predictable [ms] sounds [gr] NA made by some [ms] species [gr] NA of whales , notably the humpback whale .
	Generated Triple	species Producer_Product sounds ; sounds Product_Producer species ;
	Truth (s, r, o)	(“sounds”, “Effect_Cause”, “species”), (“species”, “Cause_Effect”, “sounds”)
	Analysis	The model treat the relation between “sounds” and “species” as product-type relation when Entity Linking is missing in this case

Table A1: Error Analysis for REKnow.