

CopyMTL: Copy Mechanism for Joint Extraction of Entities and Relations with Multi-Task Learning

Daojian Zeng,^{*§} Haoran Zhang,^{*†} Qianying Liu[‡]

[§]Changsha University of Science & Technology, Changsha, 410114, China

[†]University of Illinois at Urbana-Champaign, Illinois, 61820, USA

[‡]Kyoto University, Kyoto, 606-8501, Japan

zengdj916@163.com, haoranz6@illinois.edu, ying@nlp.ist.i.kyoto-u.ac.jp

Abstract

Joint extraction of entities and relations has received significant attention due to its potential of providing higher performance for both tasks. Among existing methods, CopyRE is effective and novel, which uses a sequence-to-sequence framework and copy mechanism to directly generate the relation triplets. However, it suffers from two fatal problems. The model is extremely weak at differing the head and tail entity, resulting in inaccurate entity extraction. It also cannot predict multi-token entities (e.g. *Steven Jobs*). To address these problems, we give a detailed analysis of the reasons behind the inaccurate entity extraction problem, and then propose a simple but extremely effective model structure to solve this problem. In addition, we propose a multi-task learning framework equipped with copy mechanism, called CopyMTL, to allow the model to predict multi-token entities. Experiments reveal the problems of CopyRE and show that our model achieves significant improvement over the current state-of-the-art method by 9% in NYT and 16% in WebNLG (F1 score). Our code is available at <https://github.com/WindChimeRan/CopyMTL>

Introduction

As a key technology for automatic Knowledge Graph (KG) construction, relation extraction has received widespread attention in recent years. Relation extraction aims to automatically learn triplets (*relation, head, tail*) from the unstructured text without human intervention.

Early studies use pipeline models (Nadeau and Sekine 2007; Chan and Roth 2011), where they cast the relation extraction problem into two separate tasks, i.e. Named Entity Recognition (NER) to extract the entities and Relation Classification. They first recognize the entities and then predict the relations between entities. However, pipeline models suffer from obvious drawbacks (Roth and Yih 2007). Each component limits the performance because of the error cascading effect and there is no chance for the model to correct mistakes. In addition, such pipeline models cannot capture the explicit relation between the two subtasks (Li and Ji

Text	Steven Jobs was born on February 1955 in San Francisco, California	
CopyRE	BirthDay(<i>Jobs</i> , 1955)	BirthPlace(<i>Jobs</i> , <i>Francisco</i>)
Gold	BirthDay(<i>Steven Jobs</i> , February 1955)	BirthPlace(<i>Steven Jobs</i> , San Francisco)

Figure 1: CopyRE predicts the entity pointer refers to the word position in the source sentence. The colored tokens show the limitation of CopyRE which cannot predict multiple tokens.

2014), where joint models can benefit from such interdependencies.

Recent studies on joint models of entity and relation extraction have three major research lines: Table Filling, Tagging, and Sequence-to-Sequence (Seq2Seq). Among these approaches, the table filling method (Gupta, Schütze, and Andrassy 2016; Adel and Schütze 2017) requires the model to enumerate over all possible entity pairs, which leads to a heavy computational burden. The tagging method (Zheng et al. 2017) suffers from the overlapping relation problem that the model cannot assign different relation tags to one token. To solve this problem, the followers (Takanobu et al. 2018; Dai et al. 2019) run tagging on a sentence for multiple turns, which is akin to the table filling method together with the heavy computational burden. Relatively speaking, the Seq2Seq method is neither plagued with overlapping relations nor with excessive computations. Seq2Seq model receives the unstructured text as input and directly decodes the entity-relation triples as a sequential output. This concise approach also matches with the human annotation process, that the annotators first read the sentences, understand the meaning and then point out the entity-relation pairs sequentially.

Currently, CopyRE (Zeng et al. 2018) is the most powerful Seq2Seq based joint extraction method which expands a Seq2Seq framework with copy mechanism in the decoder. The copying mechanism allows the model to avoid the out-of-vocabulary (OOV) problem. Despite their promising result, the model still suffers from two major drawbacks.

First, the entity copying in CopyRE is unstable and it depends on an unnatural mask to differ the head (*h*) and tail (*t*) entities. Experimental results show that CopyRE nearly ran-

*Corresponding authors, equal contribution

domly predicts the head-tail order of the two entities. The model also needs an unnatural mask that masks the probability of h while predicting t . Without this mask, when predicting t , the model would choose the same token as h , and the accuracy drops to zero. After analysis, we prove that CopyRE actually uses the same distribution to model h and t , chooses the highest probability as h , and the second-highest would be chosen as t after masking the highest probability, so without this mask, it cannot differ h and t . Modeling the h and t distribution in such manner can cause various problems, the model not only is extremely weak at differing h and t , but also cannot get information about h while predicting t .

Second, CopyRE cannot extract entities that have multiple tokens. The copy-based decoder always points to the last token of any entities, which limits the applicability of the model. For example, in Fig. 1 we show that CopyRE only predicts “Jobs” rather than the whole entity “Steven Jobs” when the entity has two tokens. In real-world scene multi-token entities are common, so this can greatly drag the model performance.

To address these two problems mentioned above, we propose CopyMTL, which is a multi-task learning based model with a new architecture for entity copying. We first provide a detailed analysis of why CopyRE is unstable during copying and propose a new model architecture to improve the shortcomings. Our new model architecture merely adds one more non-linear fully connection layer so that the model predicts separate distributions for the head and tail entity, and the tail prediction receives information from the head prediction. This architecture no longer needs the unnatural mask and increases the accuracy of entity copying, resulting in the overall improvements over the state-of-the-art model.

Then we propose a multi-task learning based Seq2Seq model to predict multi-token entities. A sequence labeling layer is added at the encoding stage to assist the entity recognition process. We use multi-task learning of NER to predict the start token of each entity while the decoder points at the last token while decoding. During training, we optimize the multi-task loss function jointly.

In conclusion, the contribution of this work is as follow:

1. We analyze the reasons for the unstable performance of entity copying in CopyRE and propose a simple but effective architecture to address this problem.
2. We propose a multi-task framework to enhance the capability of handling with multi-token entities.
3. Experimental results show that our model achieves state-of-the-art results and outperforms previous approaches by a large margin.

Background

In this section, we first introduce the CopyRE model that is based on the Seq2Seq framework. Then, we give a detailed description of the two existing problems. As shown in Fig. 2, CopyRE consists two parts: an encoder and a decoder. Given a sentence $s = \{x_1, x_2 \dots x_n\}$, the encoder transforms the input s into a vector representation. The decoder predicts the relation-entity triplets (r, h, t) each three time

steps. Inspired by CopyNet (Gu et al. 2016), the first step uses Generate-Mode to predict a relation. Then, the model switches to Copy-Mode and selects the head and tail entities one by one in two different time steps.

Encoder

To model the semantics of the input sentence better, CopyRE adopts Bidirectional LSTM (BiLSTM) (Schuster and Paliwal 1997) as the encoder, which has shown great strength in many areas of NLP. Given a sentence of word embeddings $\{e_1^E, \dots, e_n^E\}$ as input, the hidden states from two directions are computed:

$$\begin{aligned} \vec{h}_i &= \overrightarrow{LSTM^E}(e_i^E, h_{i-1}) \\ \overleftarrow{h}_i &= \overleftarrow{LSTM^E}(e_i^E, h_{i+1}) \\ h_i^E &= (\vec{h}_i + \overleftarrow{h}_i) / 2 \end{aligned} \quad (1)$$

where hidden states \vec{h}_i and \overleftarrow{h}_i from two directions are averaged¹ into one vector h_i^E as output.

Decoder

The decoder uses a one direction LSTM to predict the outputs from left to right. The last hidden state of the encoder is used to initialize the decoder hidden state. The attention score is assigned to each hidden state of the encoder and then summed up to obtain an attentive sum. Then the sum is combined with the decoded hidden states at the last time step to be fed into the decoder LSTM:

$$\begin{aligned} c_t &= Attention(h_{t-1}^D, h_{1:n}^E) \\ u_t &= [e_t^D; c_t] \cdot W^u \\ h_t^D &= LSTM^D(u_t, h_{t-1}^D) \end{aligned} \quad (2)$$

where *Attention* calculates the attentive sum of all encoder hidden states $h_{1:n}^E = \{h_1^E, \dots, h_n^E\}$ according to the last decoder hidden state h_{t-1}^D . $[\cdot; \cdot]$ is the concatenation operator, e_t is the embedding of the decoder output in the last time step, $W^u \in \mathbb{R}^{(d_e+d_c) \times d_e}$ is the parameter of linear transformation. All biases are omitted for convenience.

Every three time steps form a loop in which the decoder predicts relation, last token of head and then last token of tail to form a triplet, respectively. The confidence q_t^i for each token at position i to be copied as an entity is calculated by:

$$q_t^i = [h_t^D; h_i^E] \cdot W^e \quad (3)$$

where $W^e \in \mathbb{R}^{2d_o \times 1}$.

Then, the decoder computes the logits according to the time step t (we count the time step from 1):

$$logit_t = \begin{cases} [h_t^D \cdot W^r; q^{NA}], & \text{if } t \% 3 = 1; \\ [q_t; q^{NA}], & \text{if } t \% 3 = 2; \\ [M \otimes q_t; q^{NA}], & \text{if } t \% 3 = 0. \end{cases} \quad (4)$$

¹The original paper uses concatenation, but actually they use average in the released code.

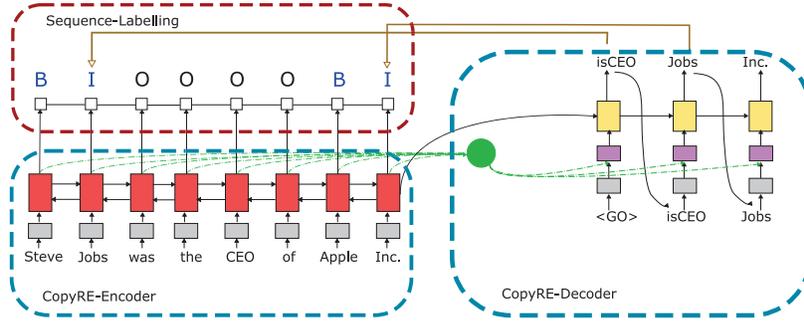


Figure 2: The overview of CopyMTL model for joint extraction of relation and entity. The CopyRE model does not contain the CopyMTL-Tagging part, i.e., the sequence-labeling part in the figure.

where $\mathbf{W}^r \in \mathbb{R}^{d_o \times rel}$, rel is the cardinality of relations, \mathbf{q}_t is the concatenation of all \mathbf{q}_i^t , \mathbf{M} is the mask which records the predicted head entity and prevents the decoder predicting it at the $t\%3 = 0$ time step. This is based on the fact that an entity cannot be both the head and the tail in the same triplet at the same time. But the mask makes no contributions for minimizing the cross entropy loss we will describe below.

Through the unnormalized logit, we can obtain the probability of output entity or relation by softmax:

$$p(y_t | y_{<t}, s) = \frac{e^{\text{logit}_t^j}}{\sum_i e^{\text{logit}_t^i}} \quad (5)$$

At time step $t\%3 = 1$ when the model should predict relation, the softmax score is calculated over all relations types; when the model should predict the entity, the softmax score is calculated over all positions in the source sentence. Then, the model can be trained via minimizing the cross entropy loss, which measures the difference between the output y_t and the label y_t^* .

$$\mathcal{L}^D = - \sum_t \log(p(y_t^* | y_{<t}, s)) \quad (6)$$

CopyRE also use padding triplets (NA, NA, NA) during training, which do not have any valid relations and entities. The confidence q^{NA} of NA-relation and NA-position of the corresponding entity is calculated through a shared parameter:

$$q^{NA} = \mathbf{h}_t^D \cdot \mathbf{W}^{NA} \quad (7)$$

where $\mathbf{W}^{NA} \in \mathbb{R}^{d_o \times 1}$.

Problems of CopyRE

As mentioned in the introduction, we found that CopyRE has two problems. First, the prediction of the entity is unstable. In detailed experiments, we observed that CopyRE cannot even fit the training set well, in which the F1 scores are approximately 0.75 and 0.40 on two datasets (see Fig. 4). In addition, if we remove the mask \mathbf{M} in Eq. (4), the F1 score will turn to zero immediately. To find out the reason behind it, we evaluate CopyRE for the predicted relations and entities in the triplets separately. The experiments show that CopyRE can gain 0.84 F1 score for relations, while the F1 score for entities dramatically drops to 0.64 (see Table 4).

In addition, when we inspected the prediction errors, we find that CopyRE is prone to mix up the order of head and tail. Thus, we can conclude that entity copying is the bottleneck of the model, which causes the performance decline.

Second, since CopyRE only predicts the last token of the entity, when the target entity contains multiple tokens, the outputs are incomplete. There are straightforward ways to solve this problem. For example, we can extend the predicted triplets to quintuple by adding the length of entities. However, such methods indirectly use or simply ignore the interactions between relation extraction and entity recognition. We propose a multi-task manner method to solve this problem and give detailed comparisons in experiments.

Our Method

As described in the last section, CopyRE is suffering the entity copying and the multi-token entity problems. We propose a model named CopyMTL (Fig. 2) to address these two problems. CopyMTL is based on a new model structure and uses a multi-task framework which adds a sequence labeling task to CopyRE encoder. In this section, we first reveal the reasons behind the entity copying problem, then propose a simple but reasonable solution. After that, we introduce an additional tagging layer of the encoder and the multi-task training procedure.

New Structure for Entity Copying

Strangely, in CopyRE entity copying highly depends on the entity mask \mathbf{M} , and the predicted distributions of head and tail entities are identical. Through our analysis, the main culprit is found in Eq. (3), who calculates the concatenation of \mathbf{h}_t^D and \mathbf{h}_i^E , then passes it to a linear transformation. Eq. (3) can be expanded and simplified to get the following form:

$$\begin{aligned} \mathbf{q}_i^t &= [\mathbf{h}_t^D; \mathbf{h}_i^E] \cdot \mathbf{W}^e \\ &= [\mathbf{h}_t^D; \mathbf{h}_i^E] \cdot [\mathbf{W}_1^e; \mathbf{W}_2^e] \\ &= \mathbf{h}_t^D \cdot \mathbf{W}_1^e + \mathbf{h}_i^E \cdot \mathbf{W}_2^e \end{aligned} \quad (8)$$

where $\mathbf{W}_1^e, \mathbf{W}_2^e \in \mathbb{R}^{d_o \times 1}$. Note that this is a summation of two scalars and the first term is independent of i . If we omit the q^{NA} , the probability of entity copying is calculated by softmax:

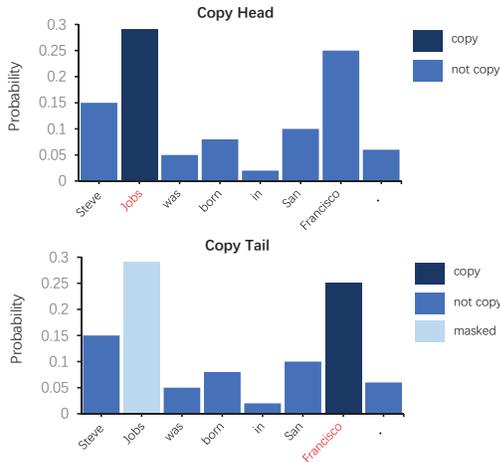


Figure 3: The problematic entity copying of CopyRE. After predicting relation *BirthPlace*, the model will copy the head entity *Jobs*, then mask the predicted head and copy the tail *Francisco*.

$$\begin{aligned}
 p(y_t | y_{<t}, s) &= \frac{e^{q_t^t}}{\sum_j e^{q_j^t}} = \frac{e^{\mathbf{h}_t^D \cdot \mathbf{W}_1^e} \cdot e^{\mathbf{h}_t^E \cdot \mathbf{W}_2^e}}{e^{\mathbf{h}_t^D \cdot \mathbf{W}_1^e} \cdot \sum_j e^{\mathbf{h}_j^E \cdot \mathbf{W}_2^e}} \\
 &= \frac{e^{\mathbf{h}_t^E \cdot \mathbf{W}_2^e}}{\sum_j e^{\mathbf{h}_j^E \cdot \mathbf{W}_2^e}}
 \end{aligned} \quad (9)$$

Abnormal dependency to the mask: In Eq. (9), we can see that $prob_t^t$ does not rely on the time step t . In other words, the output distribution of entity copying at $t \% 3 = 1$ and $t \% 3 = 2$ are identical, which causes the dependency on the mask. We visualize the output distribution of the entity copying in Fig. 3. In the figure, the model first copies the token with the highest probability, *Jobs*. Then, in the next time step, as the pointed token *Jobs* is masked, the model copies the token with the second highest probability, *Francisco*.

Unstable entity copying: Because the distributions of two time steps are the same and the mask is only used in evaluation rather than optimization, the entity copying, especially for the head entity, becomes unstable. In the training stage, CopyRE maximizes the likelihood for the head at $t \% 3 = 2$ and for the tail at $t \% 3 = 0$, while the likelihood at each time step is identical. However, as the mask is not used for optimization, there is no explicit constraint to ensure that the head has the highest probability and tail has the second highest probability. In fact, CopyRE tries to maximize both the head and the tail. Thus, which one would be the highest and be predicted at $t \% 3 = 2$ is random.

To fix the problem in Eq. (9), we simply map \mathbf{h}_t^D and \mathbf{h}_t^E to a fused feature space via one additional non-linear layer:

$$q_t^t = \sigma([\mathbf{h}_t^D; \mathbf{h}_t^E] \cdot \mathbf{W}^f) \cdot \mathbf{W}^o \quad (10)$$

where σ is the *selu*(\cdot) activation function (Klambauer et al. 2017), $\mathbf{W}^f \in \mathbb{R}^{2d_o \times d_{wf}}$ and $\mathbf{W}^o \in \mathbb{R}^{d_{wf} \times 1}$.

Due to the non-linearity of the activation function, the reduction of Eq. (9) does not hold true. Now, the entity copying depends on both i and t and there is only one target output to maximize instead of that in Fig. 3. Thus, by replacing Eq. (3) with Eq. (10), the decoder no longer needs to struggle with ranking head and tail at $t \% 3 = 2$, and the mask is no longer urgently needed². Therefore, the entity copying becomes stable with our new structure.

Sequence Labeling Layer

CopyRE only copies the last token of the entity. To predict entities with multiple tokens, we cast the problem into a sequence labeling problem and use the NER results to calibrate the entities with multiple tokens. As shown in Fig. 2, we first derive the emission potential from the encoder output. Then, an additional Conditional Random Field (CRF) layer (Lafferty, McCallum, and Pereira 2001) is employed to calculate the most probable tag for each token. We use the BIO scheme (Begin, Inside, Outside) to recognize all of the entities in the sentence. The predicted tags are used to post-process the extracted entities.

The conditional probability of target $tags^*$ given sentence s are computed by path probability:

$$p(tags^* | s) = \frac{e^{score(s, tags^*)}}{\sum_{tags'} e^{score(s, tags')}} \quad (11)$$

where the denominator is computed via dynamic programming. The unnormalized path score is defined as:

$$score(s, tags) = \sum_i \phi_{i, tag_i} + b_{tag_{i-1} \rightarrow tag_i} \quad (12)$$

where $b_{tag_{i-1} \rightarrow tag_i}$ is the transition score from tag_{i-1} to tag_i . ϕ_{i, tag_i} is the score of the tag_i for the i -th input token, which is comes from the hidden state of the Bi-LSTM at timestep i .

The loss function of the sequence labeling is:

$$\mathcal{L}^E = -\log(p(tag^* | s)) \quad (13)$$

In the inference stage, we use the NER results to post-process the decoded entities. Since we use BIO tagging scheme, there are three circumstances for the decoded last token of entities:

- 'B', a single token entity.
- 'I', an entity with multiple tokens, it will look for the token before the current token until it finds 'B'.
- 'O', a single token entity.

Training

Overall, the input sentence is fed into the encoder part. All of the hidden states of the encoder are used to label the input sequence and calculate the attention of the decoder. Initialized by the last hidden state of the encoder, the decoder generates triplets each three time steps. Thus, the loss function

²In experiments, we found that adding the mask to our method brings no enhancement.

Model	NYT			WebNLG		
	Prec	Rec	F1	Prec	Rec	F1
NovelTagging	.642	.317	.420	.525	.193	.283
CopyRE-One (ours)	.612	.530	.571	.312	.272	.291
CopyRE-Mul (ours)	.610	.566	.587	.319	.273	.294
GraphRel-1p	.629	.573	.600	.423	.392	.407
GraphRel-2p	.639	.600	.619	.447	.411	.429
CopyMTL-One	.727	.692	.709	.578	.601	.589
CopyMTL-Mul	.757	.687	.720	.580	.549	.564

Table 1: Results of the compared models on NYT and WebNLG, in which CopyRE uses less strict evaluation.

Dataset	NYT	WebNLG
Relation types	24	246
Dictionary size	90760	5928
Train sentence	56195	5019
Test sentence	5000	703

Table 2: Statistics of the datasets.

contains two parts: the encoder part introduces an additional CRF loss, and in the decoder part the cross entropy loss is used to measure the difference between the output triplets and the gold triplets.

We define the loss function as the weighted summation of encoder loss and decoder loss:

$$\mathcal{L} = \lambda \cdot \mathcal{L}^E + \mathcal{L}^D \quad (14)$$

where λ is the weight of the tagging loss.

The loss is calculated as the average over shuffled mini-batch, and the derivatives of each parameter can be computed via back-propagation.

Experiments

Datasets and Setting

We evaluated models on two datasets: New York Times (NYT) (Riedel, Yao, and McCallum 2010) and WebNLG (Gardent et al. 2017). NYT comes from the distant supervised relation extraction task (DSRE), which aims to leverage the strength of the knowledge base to generate a large-scale dataset (Mintz et al. 2009). To make joint extraction more challenging than DSRE experiment setting, Zeng et al. (2018) additionally modified the data to include more overlapping relations. WebNLG is originally used for natural language generation, in which all of the sentences are written by annotators. To avoid that the model only remembers the entity linking instead of the relation pattern, we only use the first sentence for each instance, which is the same as other baselines. The data statistics of both datasets are shown in Table 2.

Our experiments settings also followed most of the settings of CopyRE. The hidden number of LSTM was set to 1000. The max number of decoded triplets was 5. This was because the average triplet number in both dataset is about 2. We did not use “end-of-sentence” token to stop decoding,

but to decode all padding triples (*NA, NA, NA*). The embedding dimension was 100, and we used the same pretrained embeddings³. Adam (Kingma and Ba 2014) was used to optimize the neural networks and the learning rate was 0.001. The weight of L^E , λ , was set to 1.

Baselines and Evaluation Metrics

We compare CopyMTL with CopyRE (Zeng et al. 2018), NovelTagging (Zheng et al. 2017) and GraphRel (Fu, Li, and Ma 2019). NovelTagging uses sequence labeling to assign one label to each word, which contains both entity and relation information. GraphRel is the state-of-the-art model, which uses a post-editing method to revise the triplets phase by phase. For Seq2Seq model, CopyRE and our CopyMTL, we give a more detailed comparison to show the advantages of our new structure. We also evaluate the OneDecoder and MultiDecoder trick for the Seq2Seq models (denoted as -One and -Mul). The main difference between the two decoders is the parameter sharing strategy. OneDecoder uses shared parameters for predicting all triplets and is exactly what we described in the background section. MultiDecoder uses unshared decoders, each decoder predicts one triplet.

We use precision, recall, and micro-F1 score to evaluate the models. The evaluation metrics we use are stricter than that of the original CopyRE. That is to say, instead of leaving out the incomplete entity problem, the outputs of our experiments are regarded as correct only if both the relation types and all entity tokens are correct. This stricter metric meets real-world usage and the comparison is fairer to NovelTagging and GraphRel because they are not haunted by the multi-token problem. For an intuitive comparison, we also list the result of CopyRE in the table, although their evaluation is not so strict.

Comparison of Baselines

To evaluate the performance of the proposed method, we compare CopyMTL with the baseline methods. The results⁴ are shown in Table 1.

As shown, CopyMTL is the best model in WebNLG and NYT. Both the precision and the recall are significantly improved. In the NYT dataset, compared with the state-of-the-art model, GraphRel-2p, CopyMTL-One outperforms it by

³https://github.com/xiangrongzeng/copy_re

⁴As NovelTagging is significantly better than previous works, we do not add more comparisons.

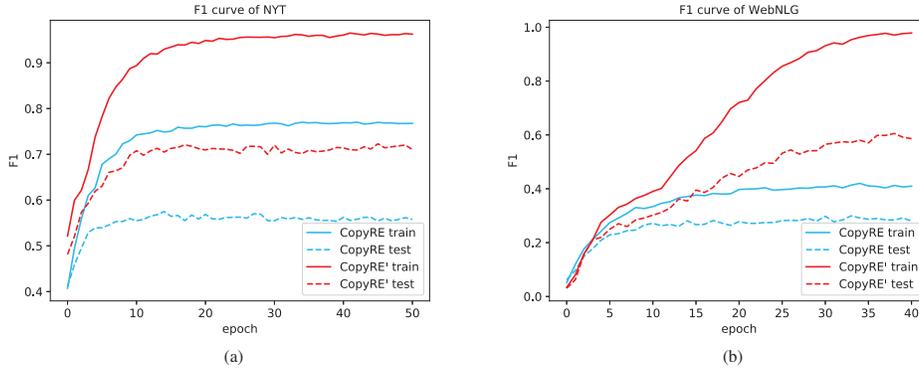


Figure 4: The training curves of CopyRE and CopyRE' on NYT and WebNLG.

Dataset	Model	Prec	Rec	F1
NYT	CopyRE	.612	.530	.571
	CopyRE'	.747	.700	.722
WebNLG	CopyRE	.312	.272	.291
	CopyRE'	.583	.629	.605

Table 3: Results of CopyRE and CopyRE' on NYT and WebNLG. These models do not consider entities with multiple tokens and use less strict evaluation that ignores entity with multiple tokens.

8.8% for precision and 9.2% for recall. In the WebNLG dataset, the effect is more significant. The improvements are 13.1% in precision and 19% in the recall. These observations verify the effectiveness of our proposed method. NovelTagging is characterized by a low recall, which is caused by its deficiency in overlapping relations. CopyRE has already solved this problem well, with 8% and 19% absolute F1 improvement in WebNLG and NYT. Our method further brings 33% and 19% F1 enhancement compared to CopyRE, which shows great potentials of Seq2Seq methods.

CopyRE argued that MultiDecoder is better than OneDecoder, which is validated by our reproduction experiment. However, with our novel CopyMTL, MultiDecoder is better than OneDecoder in NYT but worse in WebNLG. This is probably because NYT is a bigger dataset, in which MultiDecoder with more parameters works better. In practice, which decoder to use should be determined by the size of the dataset and we cannot conclude that one is better than another in every situation. For simplicity, we only discuss OneDecoder in the following sections.

Effects of the Revised Entity Copying Method

Although CopyMTL outperforms baselines by a huge margin, it is still unclear which component in CopyMTL plays the pivot role. To reveal the strength of the new model architecture, we compare CopyRE with the modified model, called CopyRE', which only substitutes Eq. (3) for Eq. (10). The comparison is in Table 3, from which we can observe that Eq. (10) is extremely effective. CopyRE' model gains

Dataset	Model	Relation	Entity
NYT	CopyRE	.846	.647
	CopyRE'	.869	.756
WebNLG	CopyRE	.767	.595
	CopyRE'	.797	.782

Table 4: F1 scores on subtasks.

13% F1 boost in NYT dataset and 31% F1 boost in WebNLG dataset.

Note that the new model architecture only considers the entity copying while the F1 score computed considers the whole triplet. In order to uncover the performance of CopyRE' in relation classification and entity recognition, we calculate the F1 scores for the two subtasks in Table 4. For the entity recognition subtask, the F1 score of CopyRE' is 10% higher in NYT and 19% higher in WebNLG. This is the main contribution of the new model architecture. For the relation classification subtask, the F1 score of CopyRE' is marginally higher (less than 3%) than that of CopyRE. This implies the better entity recognition helps relation classification learning, which confirms the argument that the interactions between two task are beneficial to each other. In the decoding stage, a more precise prediction of the entity is fed into the decoder, which aids the relation classification in the next time step.

Except for the final result, the learning processes of the two models are also different. We plot the overall F1 score varying with the training epochs in Fig. 4. The curve shows that CopyRE does not fit the training set well and the model saturates at epoch 20, where the F1 score of NYT is 75% and the F1 score of WebNLG is 40%. By contrast, CopyRE' gains 97% F1 score in the NYT training set and 97% F1 score in the WebNLG training set. In addition, the performance of CopyRE' continues increasing until epoch 40 on both datasets. The fact that the model gains lower training error which also generalizes to the test set may explain the effectiveness of CopyRE'.

Dataset	Model	Prec	Rec	F1
NYT	GraphRel-2p	.639	.600	.619
	CopyRE'5	.680	.663	.671
	CopyMTL	.727	.692	.709
WebNLG	GraphRel-2p	.447	.411	.429
	CopyRE'5	.572	.536	.553
	CopyMTL	.578	.601	.589

Table 5: Results of different multi-token models on NYT and WebNLG

Effects of Multi-Task Learning

CopyMTL aims to solve the multi-token problem. In addition to the multi-task learning used by CopyMTL, there can be other straightforward methods. For example, the decoder of CopyRE' can predict the length of the entity when it copies the entities, which forms quintuples, called CopyRE'5. This is similar to predicting both the begin and the end of the entities and should work the same.

We compare the models in Table 5, from which we can see that CopyRE'5 is worse than CopyMTL in all evaluations, but both outperform GraphRel. We conjecture that the three tasks in CopyRE'5, include relation classification, entity recognition, and entity length prediction, varying in their degree of difficulty. The entity length prediction task may interfere with the learning of other tasks, as this easier task prolongs the dependence distance of harder tasks.

In addition, we also evaluate how precisely does the encoder of CopyMTL completes the whole entities. It gains 99% F1 score in NYT and 96% F1 score in WebNLG. This evaluation is less strict than conventional NER tasks as we consider neither the types of the entities nor the entities out of relations. We can conclude that NER in joint extraction is powerful enough for triplet extraction and the main difficulty in joint extraction is to make a better prediction for both the relations and the positions of the corresponding entities.

Related Work

Extraction of entities and relations is of significance to many NLP tasks. In recent years, there have been four mainstream methods.

Pipeline methods: Previous works mainly use pipeline methods (Nadeau and Sekine 2007), a.k.a extract entities first then classify the relations. Most of the recent neural models also focus on pipeline methods, include (1) Fully-Supervised Relation Classification (Hendrickx et al. 2009) (2) Distant Supervised Relation Extraction (Mintz et al. 2009). In spite of the recent progress of neural models (Cai, Zhang, and Wang 2016; Zeng et al. 2014; Christopoulou, Miwa, and Ananiadou 2018; Qin, Xu, and Wang 2018), the pipeline methods introduce error propagation problem (Li and Ji 2014), which does harm to the overall performance.

Table filling: The joint extraction task is formalized as a table constituted by the Cartesian product of the input sentence to itself. The table blanks, except for that on the diagonal, are to be predicted as relations. The models include history-based searching (Miwa and Sasaki 2014), neural-

based prediction (Gupta, Schütze, and Andrassy 2016) and global normalization (Adel and Schütze 2017). The state-of-the-art model, GraphRel, also belongs to this genre. This model innovative takes the interaction between entities and relations into account via 2-phase GCN. The main problem of table filling is the over redundant computation for the permutation of all word pairs in a sentence. As a result, most of the blanks in the table are empty and it is the sparsity that hinders the learning of the models.

Tagging: The tagging models originally solved the tasks separately through a shared parameter: the model tags the entities first, then predicts the relations. SPTree (Miwa and Bansal 2016) used a structural neural model with the help of linguistic features. This model was promoted by an attention-based model (Katiyar and Cardie 2017). Besides, NovelTagging (Zheng et al. 2017) proposed a new tagging scheme, by which the model can predict a single tag for each word, containing both the entities and relations. However, this tagging scheme cannot handle overlapping relations because it cannot assign one token with multiple labels. To solve it, multi-pass tagging training, HRL⁵, has been proposed (Takanobu et al. 2018), based on the reinforcement learning framework and (Dai et al. 2019) leverage attention mechanism. Although these methods solve the overlapping relation problem, their nature and complexities are akin to table filling.

Seq2Seq: CopyRE (Zeng et al. 2018) is another method for solving the overlapping relation problem, which extracts triplets by a Seq2Seq framework (Sutskever, Vinyals, and Le 2014) with copy mechanism (Gu et al. 2016). But it cannot predict the entire entities. In addition, the weak performance hinders it from real-world usage. Our work resolves the problems and boosts the performance to a new level.

Conclusions and Future Work

In this paper, we revisit the CopyRE model, which jointly extracts entities and relations by a Seq2Seq model. We find that there are two problems in the model: the performance of the model is limited by the inaccurate entity copying and the generated entities are not complete. We give a theoretical analysis to reveal the reason behind the first problem, then propose a new model architecture to solve it. For the second problem, we propose a multi-task learning framework to complete the entities. Detailed experiments show the effectiveness of our method, which also outperforms the current state-of-the-art model by a huge margin.

For future work, CopyMTL still has much potential, for example, the current model can only extract a fixed number of triplets. We would also like to extend CopyMTL to extract any number of triplets. CopyMTL can build a strong baseline for future studies.

Acknowledgment

This work was supported by the National Natural Science Foundation of China(No.61602059, 61972057), “Dou-

⁵We did not use it as a compared baseline because HRL requires complicated preprocessing procedure and is not transferable to different datasets.

ble First-class” International Cooperation and Development Scientific Research Project of Changsha University of Science and Technology: 2018IC25.

References

- Adel, H., and Schütze, H. 2017. Global normalization of convolutional neural networks for joint entity and relation classification. *arXiv preprint arXiv:1707.07719*.
- Cai, R.; Zhang, X.; and Wang, H. 2016. Bidirectional recurrent convolutional neural network for relation classification. In *Proceedings of ACL*, 756–765. Berlin, Germany: Association for Computational Linguistics.
- Chan, Y. S., and Roth, D. 2011. Exploiting syntactico-semantic structures for relation extraction. In *Proceedings of ACL*, 551–560. Portland, Oregon, USA: Association for Computational Linguistics.
- Christopoulou, F.; Miwa, M.; and Ananiadou, S. 2018. A walk-based model on entity graphs for relation extraction. In *Proceedings of ACL*, 81–88. Melbourne, Australia: Association for Computational Linguistics.
- Dai, D.; Xiao, X.; Lyu, Y.; Dou, S.; She, Q.; and Wang, H. 2019. Joint extraction of entities and overlapping relations using position-attentive sequence labeling. *Proceedings of AAAI* 33(01):6300–6308.
- Fu, T.-J.; Li, P.-H.; and Ma, W.-Y. 2019. GraphRel: Modeling text as relational graphs for joint entity and relation extraction. In *Proceedings of ACL*, 1409–1418. Florence, Italy: Association for Computational Linguistics.
- Gardent, C.; Shimorina, A.; Narayan, S.; and Perez-Beltrachini, L. 2017. Creating training corpora for NLG micro-planners. In *Proceedings of ACL*, 179–188. Vancouver, Canada: Association for Computational Linguistics.
- Gu, J.; Lu, Z.; Li, H.; and Li, V. O. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of ACL*, 1631–1640. Berlin, Germany: Association for Computational Linguistics.
- Gupta, P.; Schütze, H.; and Andrassy, B. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *Proceedings of COLING 2016*, 2537–2547. Osaka, Japan: The COLING 2016 Organizing Committee.
- Hendrickx, I.; Kim, S. N.; Kozareva, Z.; Nakov, P.; Ó Séaghdha, D.; Padó, S.; Pennacchiotti, M.; Romano, L.; and Szpakowicz, S. 2009. SemEval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions (SEW-2009)*, 94–99. Boulder, Colorado: Association for Computational Linguistics.
- Katihar, A., and Cardie, C. 2017. Going out on a limb: Joint extraction of entity mentions and relations without dependency trees. In *Proceedings of ACL*, 917–928. Vancouver, Canada: Association for Computational Linguistics.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klambauer, G.; Unterthiner, T.; Mayr, A.; and Hochreiter, S. 2017. Self-normalizing neural networks. In *Advances in Neural Information Processing Systems*, 971–980.
- Lafferty, J. D.; McCallum, A.; and Pereira, F. C. N. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*, ICML ’01, 282–289. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Li, Q., and Ji, H. 2014. Incremental joint extraction of entity mentions and relations. In *Proceedings of ACL*, 402–412. Baltimore, Maryland: Association for Computational Linguistics.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of ACL-AFNLP*, 1003–1011. Suntec, Singapore: Association for Computational Linguistics.
- Miwa, M., and Bansal, M. 2016. End-to-end relation extraction using LSTMs on sequences and tree structures. In *Proceedings of ACL*, 1105–1116. Berlin, Germany: Association for Computational Linguistics.
- Miwa, M., and Sasaki, Y. 2014. Modeling joint entity and relation extraction with table representation. In *Proceedings of EMNLP*, 1858–1869. Doha, Qatar: Association for Computational Linguistics.
- Nadeau, D., and Sekine, S. 2007. A survey of named entity recognition and classification. *Linguisticae Investigationes* 30(1):3–26.
- Qin, P.; Xu, W.; and Wang, W. Y. 2018. Robust distant supervision relation extraction via deep reinforcement learning. *arXiv preprint arXiv:1805.09927*.
- Riedel, S.; Yao, L.; and McCallum, A. 2010. Modeling relations and their mentions without labeled text. In Balcázar, J. L.; Bonchi, F.; Gionis, A.; and Sebag, M., eds., *Machine Learning and Knowledge Discovery in Databases*, 148–163. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Roth, D., and Yih, W.-t. 2007. Global inference for entity and relation identification via a linear programming formulation. *Introduction to statistical relational learning* 553–580.
- Schuster, M., and Paliwal, K. K. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45(11):2673–2681.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 3104–3112.
- Takanobu, R.; Zhang, T.; Liu, J.; and Huang, M. 2018. A hierarchical framework for relation extraction with reinforcement learning. *arXiv preprint arXiv:1811.03925*.
- Zeng, D.; Liu, K.; Lai, S.; Zhou, G.; and Zhao, J. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING*, 2335–2344. Dublin, Ireland: Dublin City University and Association for Computational Linguistics.
- Zeng, X.; Zeng, D.; He, S.; Liu, K.; and Zhao, J. 2018. Extracting relational facts by an end-to-end neural model with copy mechanism. In *Proceedings of ACL*, 506–514. Melbourne, Australia: Association for Computational Linguistics.
- Zheng, S.; Wang, F.; Bao, H.; Hao, Y.; Zhou, P.; and Xu, B. 2017. Joint extraction of entities and relations based on a novel tagging scheme. In *Proceedings of ACL*, 1227–1236. Vancouver, Canada: Association for Computational Linguistics.