

Improving Neural Relation Extraction with Positive and Unlabeled Learning

Zhengqiu He,[♣] Wenliang Chen,^{♣*} Yuyi Wang,[♡]
Wei Zhang,[♣] Guanchun Wang,[◇] Min Zhang[♣]

[♣]Institute of Artificial Intelligence, School of Computer Science and Technology, Soochow University, China
[♣]Alibaba Group, [♡]ETH Zurich, [◇]Laiye Network technology co.LTD
[♣]zqhe@stu.suda.edu.cn, {wlchen, minzhang}@suda.edu.cn
[♣]lantu.zw@alibaba-inc.com, [♡]yuwang@ethz.ch, [◇]arvid@laiye.com

Abstract

We present a novel approach to improve the performance of distant supervision relation extraction with Positive and Unlabeled (PU) Learning. This approach first applies reinforcement learning to decide whether a sentence is positive to a given relation, and then positive and unlabeled bags are constructed. In contrast to most previous studies, which mainly use selected positive instances only, we make full use of unlabeled instances and propose two new representations for positive and unlabeled bags. These two representations are then combined in an appropriate way to make bag-level prediction. Experimental results on a widely used real-world dataset demonstrate that this new approach indeed achieves significant and consistent improvements as compared to several competitive baselines.

Introduction

Relation extraction (RE) is to classify the relationship between entity pairs from plain text where the entity pairs are specified. Recently, supervised learning models have achieved lots of progress in the RE task (Riedel et al. 2013; Zeng et al. 2014). Similar to other Natural Language Processing (NLP) tasks, the supervised models for RE typically require a massive amount of training data, labeled by annotators. However, hiring annotators is costly and non-scalable, in terms of both time and money.

As an alternative solution in practice, distant supervision can automatically generate a training corpus (Mintz et al. 2009). The key idea of distant supervision is that given an entity pair $\langle e_h, e_t \rangle$ and its corresponding relation r_B from one knowledge base (KB) such as Freebase, we simply label all sentences containing the two entities by relation r_B . This solution has achieved a certain success and has been used in the state-of-the-art techniques for RE (Riedel, Yao, and McCallum 2010; Hoffmann et al. 2011; Surdeanu et al. 2012; Min et al. 2013). However, the training data generated by this solution inevitably has wrong labeling problem. For example, as illustrated in Figure 1, we have a triple including

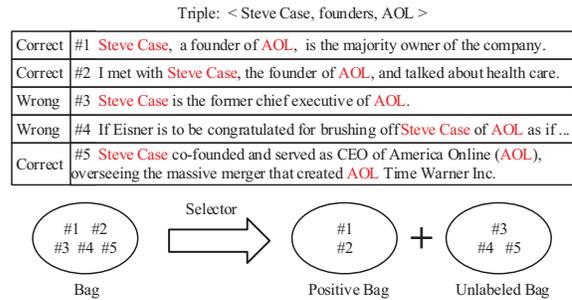


Figure 1: A sentence bag generated through distant supervision, further split into two bags after selection

entity pair $\langle \text{Steve Case, AOL} \rangle$ and its relation “founders”¹ and label five sentences by distant supervision. Though the third sentence “Steve Case is the former chief executive of AOL.” does not convey relation “founders” of “Steve Case” and “AOL”, it is wrongly labeled as a positive instance for this relation under the assumption of distant supervision. Our preliminary experimental results show that this problem largely harms the performance of relation extraction at **sentence-level**.

To address the above problem, previous studies turn to multi-instance learning (MIL) (Zeng et al. 2015; Lin et al. 2016; Ji et al. 2017), where the training set consists of bags for given entity pairs. Thus, the task of RE turns to predicting relations at **bag-level**. However, because the wrong labeling problem is still remaining, each bag in multi-instance learning may contain many sentences mentioning the same entity pair but belonging to different relations. These noisy sentences may hinder the performance of the MIL-based relation extraction models.

To alleviate the impact of noisy sentences, recent studies (Feng et al. 2018; Qin, Xu, and Wang 2018) propose to select positive sentences from noisy bags. As a selection solution, approaches based on reinforcement learning (RL) have been proposed to train RL agents to decide whether a sentence is positive. Thus the sentences are grouped into a set of positive

*The corresponding author is Wenliang Chen.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The relation “/business/company/founders” is defined in Freebase, and we shorten it to “founders”.

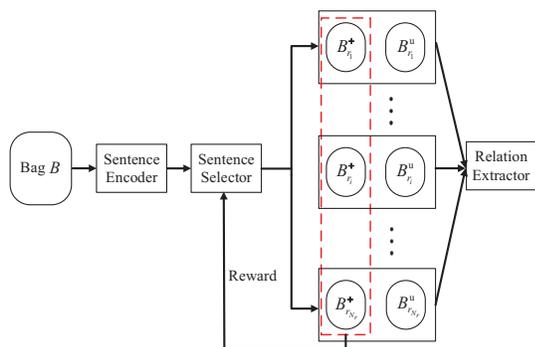


Figure 2: Infrastructure of our proposed model

instances and a set of unlabeled instances, as shown in Figure 1². There have been three solutions to utilize the sets of positive and unlabeled instances, which have achieved a certain success. One solution ignores the unlabeled sentences and only keeps the positive sentences in MIL to perform **bag-level** prediction at final relation extractor (Zeng et al. 2018). The second solution discards the unlabeled sentences and utilizes the positive sentences to perform **sentence-level** prediction at final relation extractor (Feng et al. 2018). The third solution treats the unlabeled sentences as negative instances and utilizes both positive and negative instances to perform **sentence-level** prediction at final relation extractor (Qin, Xu, and Wang 2018). However, we argue that using only the selected positive instances and ignoring the other sentences when making relation prediction at **bag-level** is suboptimal, and treating the unlabeled sentences as negative instances might not be appropriate since they might contain positive instances (for example, sentence #5 in Figure 1).

Ideally, we would wish to combine the merits of all solutions in our final relation extractor, that is using both positive and unlabeled instances in MIL for **bag-level** prediction. In this paper, we first train an RL agent to automatically select the positive sentences for each relation. As shown in Figure 1, for relation “founders”, the agent selects the first two sentences as a new positive bag, and the last three sentences are in an unlabeled bag. Therefore, our task is converted into a situation similar to a positive and unlabeled (PU) learning problem (Li and Liu 2005; Sechidis, Calvo, and Brown 2014). Classical PU learning methods, either selecting negative data points from the unlabeled collection, or simply assuming all unlabeled data points are negative (Li and Liu 2005; Sechidis, Calvo, and Brown 2014), are not suitable for our task as we mention above. Thus we propose a new method to exploit both positive and unlabeled data. That is, we first build two representations for positive and unlabeled sentences respectively, and then the two representations are combined into one as a **bag-level** representation which helps with making **bag-level** prediction.

In summary, we make the following contributions:

- In our RL-based instance selector, we design a new state representation which considers sentence embeddings, relation embeddings, and the embeddings of selected positive instances. To our best knowledge, it is the first time that the relation embedding is used in the state representation. Experimental results show that the relation embedding can significantly improve our system, as shown in Figures 3 and 4.
- In order to exploit unlabeled sentences properly, we propose a model to build representations for positive and unlabeled sentences in a bag respectively and then they are combined into one representation. Totally, we have three kinds of representations for relation extractor. The ablation study in experiments shows that all the representations make positive contributions to our system, as shown in Figures 6 and 7.

We design experiments to demonstrate the effectiveness of our system. The results on a widely used dataset indicate that our proposed model significantly outperforms the systems using RL-based instance selector in the previous studies.

Our Approach

In this section, we describe our approach in detail. Due to the distant supervision assumption, the sentences in the bag of an entity pair may not belong to the predefined relation type. The basic idea is that, we identify the relevance of a relation and every sentence in a bag, and output positive and unlabeled bags for the relation. This selection procedure can be considered as a sequential decision-making process which can be modeled by reinforcement learning. Based on both positive and unlabeled bags, we build our relation extractor.

Our framework consists of three key components as shown in Figure 2: (1) A sentence encoder, which encodes the sentences into distributed representations. (2) An instance selector based on the RL agent, which decides whether a sentence is positive for each query relation. The agent takes the distributed representation of sentence and relation as input. (3) A relation extractor, which is trained on the positive and unlabeled bags generated by the sentence selector.

Sentence Encoder

The sentence encoder transforms a sentence into its distributed representation by a (P)CNN (Zeng et al. 2014; 2015), which has shown the state-of-the-art performance in this task.

Input Representation Given a sentence q , and q_i is the i -th word in the sentence, the input is a matrix composed of n vectors $w = [w_1, \dots, w_n]$, where w_i corresponds to q_i and consists of its word embedding and position embedding. Following previous work (Zeng et al. 2015; Lin et al. 2016), we employ the skip-gram method (Mikolov et al. 2013b) to pretrain the word embedding, which will be fine-tuned afterwards. Position embedding is first successfully applied to relation extraction task by Zeng et al. (2014),

²Please note sentence #5 is positive in Figure 1, but it is grouped into the unlabeled bag by the sentence selector.

which specifies the relative distances of a word with respect to two target entities. In this way, $w_i \in \mathbb{R}^d$, and $d = d^a + 2 \times d^b$, where d^a and d^b are the dimensions of word embedding and position embedding respectively.

Convolution, Max-pooling After encoding the input words, a convolution layer and a max-pooling operation are applied to reconstruct the original input w by learning sentence features from a small window of words at a time while preserving word order information. We use K convolution filters with the same window size l . Formally, the j -th element of the output vector $c \in \mathbb{R}^K$ is

$$c[j] = \max_i (\mathbf{W}_j w_{i:i+l-1}) \quad (1)$$

where \mathbf{W}_j is the weight matrix.

Further, Zeng et al. (2015) adopt piecewise max pooling (PCNN) in relation extraction, which is a variation of CNN. Suppose the positions of the two entities are p_1 and p_2 respectively. Then, each convolution output vector is divided into three segments:

$$[0 : p_1 - 1] / [p_1 : p_2] / [p_2 + 1 : n - l] \quad (2)$$

Then the max-pooling operation is applied to each segment. Finally, we apply a non-linear transformation such as tanh on the output vector to obtain the sentence embedding. For each sentence q in bag B , we get sentence embedding x_q from the sentence encoder.

Sentence Selector

The sentence selector is an RL agent which takes a query relation and a sentence as inputs and decides whether the sentence is positive to the relation. Here we suppose that the inputs are query relation r and sentence q (with its embedding x_q) from bag B . As for bag B , after selection we obtain positive bag B_r^+ and unlabeled bag B_r^u . We first define state representation, Q-network policy function, and reward function. Then we describe how to train the RL agent.

State Representation In our state representation, we first consider the information from the sentence as Feng et al. (2018) do. To model our task as a Markov decision process (MDP) problem (Fang, Li, and Cohn 2017), we add the information from the early states into the representation of the current state. In addition, we consider the information of the relation since it is one of the key factors of selection. Thus, the state representation includes: (1) The sentence embedding x_q . (2) The relation embedding r_e of relation r , which is initialized randomly and tuned during training. (3) The embedding x_{pos} , which is the average vector of the positive sentences selected so far in the current bag. Therefore, the state representation $S_{q,r}$ is $[x_q, r_e, x_{pos}]$. Compared with the state presentation used in Feng et al. (2018), ours considers the information of relation embedding.

Q-network The agent takes an action to decide whether sentence q is positive to query relation r . The action is taken by the policy function $A_{q,r} = \operatorname{argmax} Q(S_{q,r}, A)$, where Q is the state-action table (Watkins and Dayan 1992), and A is a discrete set of actions. Following Mnih et al. (2015), we employ a two-layer neural networks as the policy function.

In our case, the action space includes two actions: positive (+1) and unlabeled (-1).

Reward We regard the selection of a bag as an episode, and a reward is calculated after all positive sentences are selected. Thus we only have a delayed reward without immediate rewards in bag B . We provide the delayed reward as feedback to the agent. We set the reward to be the conditional probability that the selected positive bag B_r^+ belongs to the query relation. For query r and Bag B , the reward function is defined as

$$R(r, B) = P(r|B_r^+) \quad (3)$$

The conditional probability $P(r|B_r^+)$ is computed as

$$P(r|B_r^+) = \frac{\exp(\mathbf{o}_r^+[r])}{\sum_{1 \leq k \leq N_r} \exp(\mathbf{o}_r^+[r_k])} \quad (4)$$

where N_r is the total number of relations, and \mathbf{o}_r^+ corresponds to the score vector associated to all relation types:

$$\mathbf{o}_r^+ = \mathbf{W} \mathbf{emb}_{B_r^+} + \mathbf{b} \quad (5)$$

where \mathbf{W} is a weight matrix, \mathbf{b} is a bias vector, and $\mathbf{emb}_{B_r^+}$ is the corresponding embedding of bag B_r^+ .

Training RL Agent We aim to maximize the expected total reward for all bags, and train the agent using Q-learning (Watkins and Dayan 1992), a standard reinforcement learning algorithm that learns policies for an agent interacting with its environment. Each episode during training corresponds to a query relation and a bag, where the agent selects the positive sentences of the relation from the bag.

Let $Q(S, A)$ be the measure of long-term cumulative reward obtained by taking action A from state S . We consider the function approximation parameterized by θ^Q , which are optimized by minimizing the loss:

$$L(\theta^Q) = \mathbb{E}_{S,A} [(V_t - Q(S, A|\theta^Q))^2] \quad (6)$$

where V_t is estimated using the Bellman equation (Sutton and Barto 1998):

$$V_t = \mathbb{E}_{S'} [R(S, A) + \gamma \max_{A'} Q(S', A'|\theta^Q)] \quad (7)$$

where $R(S, A)$ is the immediate reward, and γ is the discount factor for the value of future rewards. As mentioned above, we don't have the immediate reward. And the order of sentences in a bag should not influence the result. Thus, the immediate reward is set to 0, and γ is set to 1. As for r and B , we have

$$V_t = R(r, B) \quad (8)$$

We optimize the loss using stochastic gradient descent (SGD) and experience replay, with random mini-batch of past experience sampled for training.

Relation Extractor

After the process of the sentence selector, we build our relation extractor based on both positive and unlabeled bags. The reason we involve the unlabeled instances is to increase the robustness of the method in noisy environment of distant supervision.

We first define three types of representations from different views and then define loss functions based on the representations. Finally, we jointly train the relation extractor and the sentence selector.

Three Representations We generate three types of representations: (1) POS-based representation, which represents the selected positive bag. (2) UNL-based representation, which represents the unlabeled bag. (3) PU-combined representation, which is a combination of the above two representations. For all three representations, we project to vectors in N_r -dimension space.

POS-based Representation \mathbf{o}_r^+ : Based on the positive bag B_r^+ and relation r , we obtain the POS-based representation \mathbf{o}_r^+ by Equation (5) which is also used for the reward function of the sentence selector. We share the parameters of \mathbf{o}_r^+ between the relation extractor and the sentence selector.

UNL-based Representation \mathbf{o}_r^u : Based on the unlabeled bag B_r^u and relation r , the UNL-based representation \mathbf{o}_r^u is computed as

$$\mathbf{o}_r^u = \mathbf{W}^u \text{emb}_{B_r^u} + \mathbf{b}^u \quad (9)$$

where \mathbf{W}^u is a parameter matrix, and \mathbf{b}^u is a bias vector.

PU-combined Representation \mathbf{o}_r : This type of representation is a combination of the above two representations, which is inspired by the idea of the feature combination in traditional models. The PU-combined representation \mathbf{o}_r is computed as

$$\mathbf{o}_r = \alpha \mathbf{o}_r^+ + (1 - \alpha) \mathbf{o}_r^u \quad (10)$$

where $\alpha \in (0, 1)$ is a weight for combination.

Three Loss Functions We define three loss functions based on the representations respectively. Here we suppose that training data $\mathcal{D} = \{(B_1, r_{B_1}), \dots, (B_M, r_{B_M})\}$ consists of M sentence bags and their corresponding relation labels assigned by distant supervision.

As for POS-based representation, we use Equation (4) to calculate the probability $P(r|B_r^+)$. Then we define cross-entropy loss function as the training objective:

$$L_{pos}(\mathcal{D}) = - \sum_{i=1}^M \log P(r_{B_i} | B_{ir_{B_i}}^+) \quad (11)$$

where $B_{ir_{B_i}}^+$ is the positive bag for r_{B_i} from B_i .

As for UNL-based representation, the conditional probability $P(r|B_r^u)$ for each r is computed as

$$P(r|B_r^u) = \frac{\exp(\mathbf{o}_r^u[r])}{\sum_{1 \leq k \leq N_r} \exp(\mathbf{o}_r^u[r_k])} \quad (12)$$

The loss function is defined as

$$L_{unl}(\mathcal{D}) = - \sum_{i=1}^M \log(1 - P(r_{B_i} | B_{ir_{B_i}}^u)) \quad (13)$$

where $B_{ir_{B_i}}^u$ is the unlabeled bag for r_{B_i} from B_i .

As for PU-combined representation, we can compute the conditional probability $P(r|B)$ for bag B as

$$P(r|B) = \frac{\exp(\mathbf{o}_r[r])}{\sum_{1 \leq k \leq N_r} \exp(\mathbf{o}_r[r_k])} \quad (14)$$

Please note that for each query relation r_i , we will perform a sentence selection and calculate the probability $P(r_i|B)$. Thus the probability of one query relation is independent from the other types, and we try to optimize our model to

enlarge the probability $P(r_B|B)$ of bag label r_B for better prediction at bag-level. Therefore, we design a loss function to rank the given relation label higher than the other labels at bag-level:

$$L_{bag}(\mathcal{D}) = - \sum_{i=1}^M \left(\log P(r_{B_i} | B_i) + \log(1 - P(r_{B_i}^* | B_i)) \right) \quad (15)$$

where r_{B_i} is the given label of sentence bag B_i of focused entity pair, $r_{B_i}^*$ is the relation type that have the highest conditional probability among the other relation labels.

Finally, we combine the three loss functions for final relation extraction:

$$L(\mathcal{D}) = L_{pos}(\mathcal{D}) + L_{unl}(\mathcal{D}) + \beta L_{bag}(\mathcal{D}) \quad (16)$$

where β is a weighting factor.

Training and Testing To solve the optimization problem, we adopt mini-batch SGD to minimize the objective function. During learning, we jointly train the relation extractor and the RL agent until convergence.

In the testing phase, for each input bag, we first use the sentence selector to output the positive and unlabeled bags for each relation. Then we use Equation (14) to perform relation classification based on the positive and unlabeled bags.

Experiments

In this section, we present our experimental results and detailed analysis.

Datasets

We evaluate our model on a widely used benchmark dataset developed by Riedel, Yao, and McCallum (2010), which has also been used in many recent studies (Lin et al. 2016; Ji et al. 2017). The dataset contains 53 relation types, including a special relation ‘‘NA’’ standing for no relation between the entity pair. The training data contains 522,611 sentences, 281,270 entity pairs and 18,252 relational facts. The testing set contains 172,448 sentences, 96,678 entity pairs and 1,950 relational facts.

Evaluation Metrics

Following the practice of previous studies (Zeng et al. 2015; Lin et al. 2016), we evaluate our method in two ways, i.e., the *held-out evaluation* and *manual evaluation*.

The held-out evaluation does not consider the instances of the NA category, and only compares the entity-relation tuples produced by the system on the test data against the existing Freebase entity-relation tuples. We report the precision-recall curves of the experiments.

The manual evaluation is performed to avoid the influence of the wrong labels resulting from distant supervision and the incompleteness on Freebase data. In the manual evaluation, we manually check the newly discovered relation instances, and reports the precision of the top N discovered relational facts with the highest probabilities. We have two annotators which independently annotate the outputs. For divergent annotation results, discussions are held until agreement is reached.

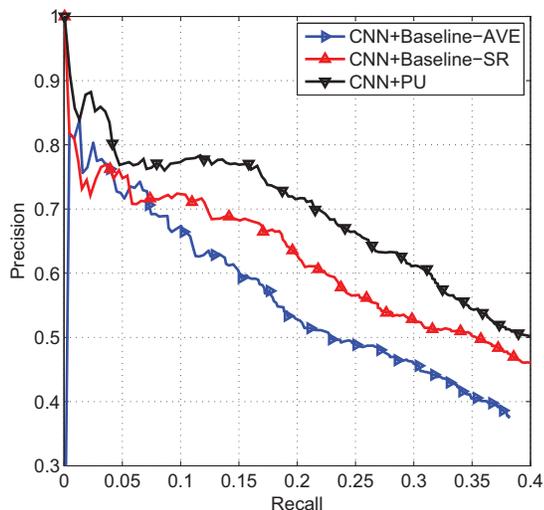


Figure 3: Comparison with baselines under CNN

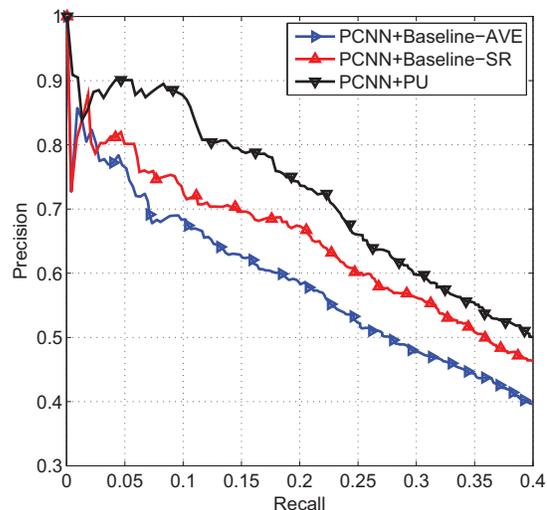


Figure 4: Comparison with baselines under PCNN

Hyperparameter Settings

In the experiments, we use the same hyperparameters of sentence encoder as previous work (Zeng et al. 2015; Lin et al. 2016), and tune other hyperparameters via three-fold cross validation on training data.

After tuning, we choose the following settings in our experiments: the dimension of word embedding d^a is set to 50, the dimension of position embedding d^b is set to 5, the number of filters K is set to 230, and the window size l of filters is 3. The batch size is fixed to 50, the dropout probability is set to 0.5. And the combination weight α is set to 0.7, β is set to 0.1. When training, we apply Adam (Kingma and Ba 2014) to optimize parameters, and the learning rate is set to 0.001.

Main Results

In this section, we compare our system (PU) with two baseline systems through held-out evaluation. We select the CNN model proposed by Zeng et al. (2014) and the PCNN model proposed by Zeng et al. (2015) as our sentence encoders. The first baseline (Baseline-AVE) is a system without RL which represents the bag as the average vector of all the sentences in bags (described in Zeng et al. (2014) and Zeng et al. (2015)). The second one (Baseline-SR) is a system where we remove the relation embedding in the state representation in the sentence selector from our final system.

Figures 3 and 4 show the results, where the higher the position of a curve is, the better the corresponding model is. From the figures, we observe that for both CNN and PCNN, our system (PU) and Baseline-SR achieve better performance than Baseline-AVE. This fact indicates that, the RL-based model can effectively select positive sentences from bags and is beneficial to relation extraction. From the results of PU and Baseline-SR, we can also see that using relation embedding as the component of state representation is helpful.

Comparison to Previous Approaches

To evaluate our proposed method, we compare our model to several advanced systems through held-out evaluation, which fall into three categories:

- Traditional discrete feature-based methods: (1) **MultiR** (Hoffmann et al. 2011) is a probabilistic graphical model with multi-instance learning under the “at-least-one” assumption. (2) **MIML** (Surdeanu et al. 2012) is a graphical model with both multi-instance and multi-label learning.
- Attention-based methods: (1) **PCNN+ATT** (Lin et al. 2016) utilizes sentence-level attention to reduce the weights of noisy instances. (2) **APCNN+D** (Ji et al. 2017) uses external background information of entities via an attention layer to help relation classification.
- Reinforcement learning-based methods: (1) **CNN+RL** (Feng et al. 2018) proposes to use RL to select sentences in the training process, and predicts the relations at sentence-level. (2) **PCNN+ATT+RL** (Qin, Xu, and Wang 2018) proposes to redistribute the distant supervision data and then trains the PCNN+ATT model on the selected data. (3) **PE+REINF** (Zeng et al. 2018) regards the relation extractor as reinforcement learning agent, and follows the expressed-at-least-once assumption to predict the bag relation.

Figure 5 shows the precision/recall curves for these methods. We can observe that: (1) Both MultiR and MIML underperform as compared to the Neural-based methods. It demonstrates that the human-designed feature cannot concisely express the semantic meaning of the sentences, and the inevitable error brought by NLP tools will hurt the performance of relation extraction. (2) Both attention-based and reinforcement learning-based methods can obtain reasonable performance. (3) Both CNN+PU and PCNN+PU consistently outperforms all other approaches by large margin.

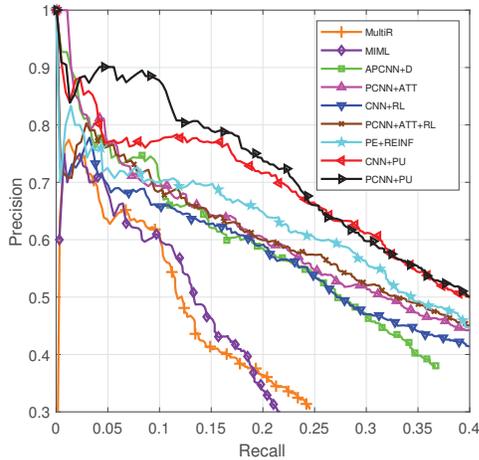


Figure 5: Comparison with previous results

Accuracy	Top 100	Top 200	Top 500
MultiR	0.83	0.74	0.59
MIML	0.85	0.75	0.61
APCNN+D	0.87	0.83	0.74
PCNN+ATT	0.86	0.83	0.73
PCNN+PU	0.92	0.89	0.80

Table 1: Manual evaluation results

Accuracy	Top 100	Top 300
CNN+PU	0.96	0.93
PCNN+PU	0.97	0.95

Table 2: Accuracy on relation extraction based on selected sentences

It indicates that the screening process of sentences and the joint training of positive and noise sentences are effective for distant supervision relation extraction. (4) PCNN+PU performs a little better as compared to CNN+PU. It means that the piecewise max-pooling strategy is helpful for encoding sentences, and the model of our model can be further improved if we have a better sentence encoder.

Manual Evaluation

Due to existence of noises resulting from distant supervision and the incomplete nature of Freebase, a majority of the tuples that were produced with high confidence are false negatives and are actually true relation instances. Thus, we can see that there is a sharp decline in the precision-recall curves in most models in Figure 5. Therefore, we perform a manual evaluation to eliminate these problems, and report the $P@100$, $P@200$, $P@500$ for PCNN+PU and four baseline models, as shown in Table 1.

We can see that PCNN+PU achieves the best performance. Moreover, the precision is higher than in the held-out evaluation. It indicates that many of the false negatives that we predicted are actually true relational facts.

Ablation Study

In this section, we perform an ablation study of our system through held-out evaluation. We want to see the contributions of the proposed representations in the model. There are three components of bag representations in our system. We compare the different settings via CNN/PCNN: (1) **PU (Ours)** is our final system based on the proposed model. (2) **PU-COMB** is a system where we remove the loss function of PU-combined representation from our final system (PU). (3) **PU-COMB-UNL** is a system where we remove the loss function of UNL-based representation from PU-COMB. PU-COMB-UNL only uses the POS-based representation with RL-based sentence selector.

Figures 6 and 7 show the results. From the figures, we find that if we remove the PU-combined representation, the system performs worse. We also find that the UNL-based representation makes the positive contribution to the systems. These facts indicate that all the three components make positive contributions to our system.

Effect of Sentence Selector

We randomly select 300 entity-relation tuples in testing set to evaluate manually, the proportion of correctly labeled is 0.76³. To demonstrate the precision of the positive sentences selected by our instance selector, we evaluate our proposed sentence selection module by manual method. We report the $P@100$ and $P@300$ entity-relation tuples generated by the sentence selection module. The results are shown in Table 2. From the table, we can see that both CNN+PU and PCNN+PU can achieve high accuracy by the sentence selection module. It indicates that the selected sentences by our selector is reliable.

Case Study

We further present a real example in Table 3 for case study. The entity-relation tuple is (*Liberia*, *capital*, *Monrovia*). There are four sentences containing the entity pair. Here, “capital” and “contains” are both relation types defined in Freebase. From the table, we can see that for relation “capital”, the RL agent selects the second sentence as positive example; for relation “contains”, the agent selects three sentences as positive examples and the first sentence as unlabeled example.

Related Work

Most existing supervised RE systems (Riedel et al. 2013; Zhou et al. 2016) require a great deal of human-annotated relation-specific training data, which is time consuming and labor intensive. To address this issue, distantly supervised RE is proposed to automatically generate training data by aligning plain text with KB (Mintz et al. 2009; Takamatsu, Sato, and Nakagawa 2012; Ritter et al. 2013). To further alleviate the wrong labeling problem arisen by distant supervision assumption, many studies formulate RE

³For a given entity-relation tuple, if a sentence in the generated bag can express the corresponding relation, then the sentence is considered to be correctly labeled.

Tuple	Generated Sentences via Distant Supervision	capital	contains
<i>capital</i> (Liberia, Monrovia)	1. Nearly two years after Charles Taylor fled [Monrovia] under pressure from advancing rebels and a force of marines on ships off [Liberia], he sits exiled in ...	unlabeled	unlabeled
	2. On the streets of [Monrovia], [Liberia] 's capital, rumors of a coup or an attack by taylor supporters, several of whom hold ...	positive	positive
	3. My sister was living in the [Monrovia] suburb of paynesville, [Liberia], with her family and a handful of orphans and other ...	unlabeled	positive
	4. They said Mr. Van Kouwenhoven, who initially ran a car import business, a hotel and a casino in [Monrovia], [Liberia], established ties with ...	unlabeled	positive

Table 3: Case study: a real example for comparison

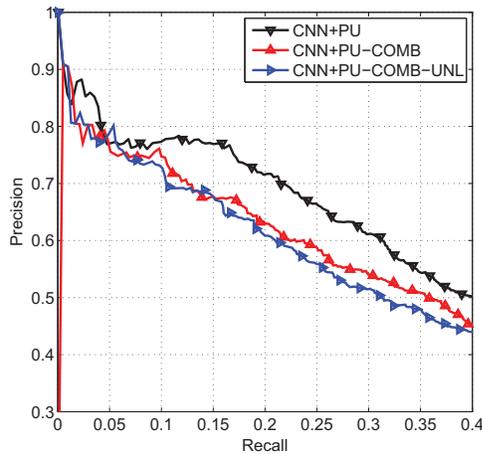


Figure 6: Ablation Study under CNN

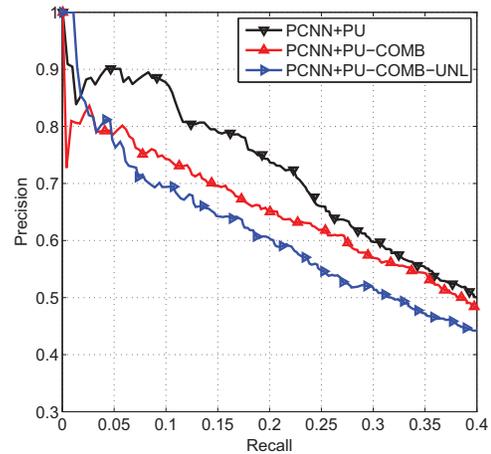


Figure 7: Ablation Study under PCNN

as a MIL problem (Riedel, Yao, and McCallum 2010; Hoffmann et al. 2011; Surdeanu et al. 2012). Zeng et al. (2015) further combine at-least-one multi-instance learning and assume that only one sentence expresses the relation for each entity pair. Lin et al. (2016) propose to use attention mechanism to dynamically reduce the weights of noisy instances, which shows promising results. Besides, Ji et al. (2017) attempt to incorporate the external information of entities for RE. In addition, He et al. (2019) utilize tree models to use syntax information to enhance the representations of entities.

At present, many researchers have begun to apply deep reinforcement learning in NLP tasks (Narasimhan, Yala, and Barzilay 2016; Li et al. 2016). As for distantly supervised relation extraction, Zeng et al. (2018) directly regard the relation extractor as an RL agent, and predict the bag relation based on the extracted sentence-level relations. Feng et al. (2018) propose to select correct sentences under the framework of reinforcement learning and then train relation extractor on the cleaned data. Qin, Xu, and Wang (2018) also use RL to determine whether a sentence is positive or negative and both positive and negative samples are used to build a classifier for rewards. Differently from the one-agent-for-

one-relation strategy of Qin, Xu, and Wang (2018), we build a single agent for all relation types. In addition, in existing approaches, they only use the selected positive instances and ignore the other sentences when making relation prediction at bag-level, or directly treat the unlabeled sentences which might contain positive instances as negative instances. In contrast, in our approach, all the unlabeled instances directly contribute to our final relation extractor.

Conclusion

In this paper, we propose a novel approach for bag-level relation classification which makes full use of the results of the sentence selector. In the first step in our approach, we apply the idea of reinforcement learning to build the sentence selector which selects positive and unlabeled sentences from noisy data. In the sentence selector, we newly add the information of relation embedding in the state representation. Then, when training the bag-level relation classifier, we consider the unlabeled instances which are often discarded by the previous studies. With proper joint training, sentence selector and relation classifier interact with each other. Finally, the experimental results show that our proposed system consistently outperforms the baseline systems. Further studies

indicate that adding relation embedding in the state representation and using unlabeled instances for training models are helpful to relation extraction.

Acknowledgments

The research work is supported by the National Key Research and Development Program of China (Grant No. 2017YFB1002104) and the National Natural Science Foundation of China (Grant Nos. 61876115, 61936010). This work is partially supported by the joint research project of Alibaba and Soochow University and a Project Funded by the Priority Academic Program Development of Jiangsu Higher Education Institutions (PAPD). Corresponding author is Wenliang Chen. We would also thank the anonymous reviewers for their detailed comments, which have helped us to improve the quality of this work.

References

- Fang, M.; Li, Y.; and Cohn, T. 2017. Learning how to active learn: A deep reinforcement learning approach. In *Proceedings of EMNLP*, 595–605.
- Feng, J.; Huang, M.; Zhao, L.; Yang, Y.; and Zhu, X. 2018. Reinforcement learning for relation classification from noisy data. In *Proceedings of AAAI*, 5779–5786.
- He, Z.; Chen, W.; Li, Z.; Zhang, W.; and Shao, Hao and Zhang, M. 2019. Syntax-aware entity representations for neural relation extraction. *Artificial Intelligence Journal* 275:602–617.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of ACL*, 541–550.
- Ji, G.; Liu, K.; He, S.; and Zhao, J. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In *Proceedings of AAAI*, 3060–3066.
- Kingma, D. P., and Ba, J. L. 2014. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Li, X.-L., and Liu, B. 2005. Learning from positive and unlabeled examples with different data distributions. In *Proceedings of ECML*, 218–229.
- Li, J.; Monroe, W.; Ritter, A.; Jurafsky, D.; Galley, M.; and Gao, J. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of EMNLP*, 1192–1202.
- Lin, Y.; Shen, S.; Liu, Z.; Luan, H.; and Sun, M. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of ACL*, 2124–2133.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013b. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS*, 3111–3119.
- Min, B.; Grishman, R.; Wan, L.; Wang, C.; and Gondek, D. 2013. Distant supervision for relation extraction with an incomplete knowledge base. In *Proceedings of NAACL*, 777–782.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of ACL*, 1003–1011.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Narasimhan, K.; Yala, A.; and Barzilay, R. 2016. Improving information extraction by acquiring external evidence with reinforcement learning. In *Proceedings of EMNLP*, 2355–2365.
- Qin, P.; Xu, W.; and Wang, W. Y. 2018. Robust distant supervision relation extraction via deep reinforcement learning. In *Proceedings of ACL*, 2137–2147.
- Riedel, S.; Yao, L.; McCallum, A.; and Marlin, B. M. 2013. Relation extraction with matrix factorization and universal schemas. In *Proceedings of NAACL*, 74–84.
- Riedel, S.; Yao, L.; and McCallum, A. 2010. Modeling relations and their mentions without labeled text. In *Proceedings of ECML PKDD*, 148–163.
- Ritter, A.; Zettlemoyer, L.; Mausam; and Etzioni, O. 2013. Modeling missing data in distant supervision for information extraction. *Transactions of the Association for Computational Linguistics* 1:367–378.
- Sechidis, K.; Calvo, B.; and Brown, G. 2014. Statistical hypothesis testing in positive unlabelled data. In *Proceedings of ECML PKDD*, 66–81.
- Surdeanu, M.; Tibshirani, J.; Nallapati, R.; and Manning, C. D. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of EMNLP*, 455–465.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Takamatsu, S.; Sato, I.; and Nakagawa, H. 2012. Reducing wrong labels in distant supervision for relation extraction. In *Proceedings of ACL*, 721–729.
- Watkins, C. J., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.
- 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.
- Zeng, D.; Liu, K.; Chen, Y.; and Zhao, J. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of EMNLP*, 1753–1762.
- Zeng, X.; He, S.; Liu, K.; and Zhao, J. 2018. Large scaled relation extraction with reinforcement learning. In *Proceedings of AAAI*, 5658–5665.
- Zhou, P.; Shi, W.; Tian, J.; Qi, Z.; Li, B.; Hao, H.; and Xu, B. 2016. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of ACL*, 207–212.