

Are Noisy Sentences Useless for Distant Supervised Relation Extraction?

Yu-Ming Shang,^{1,2,4} He-Yan Huang,¹ Xian-Ling Mao,¹ Xin Sun,¹ Wei Wei³

¹School of Computer Science, Beijing Institute of Technology, Beijing, China

²CETC Big Data Research Institute Co., Ltd., Guiyang, China, 550022

³Huazhong University of Science and Technology, Hu'bei, China

⁴Big Data Application on Improving Government Governance Capabilities
National Engineering Laboratory Guiyang, China, 550022
{ymshang, hhy63, maoxl, sunxin}@bit.edu.cn, Weiw@hust.edu.cn

Abstract

The noisy labeling problem has been one of the major obstacles for distant supervised relation extraction. Existing approaches usually consider that the noisy sentences are useless and will harm the model’s performance. Therefore, they mainly alleviate this problem by reducing the influence of noisy sentences, such as applying bag-level selective attention or removing noisy sentences from sentence-bags. However, the underlying cause of the noisy labeling problem is not the lack of useful information, but the missing relation labels. Intuitively, if we can allocate credible labels for noisy sentences, they will be transformed into useful training data and benefit the model’s performance. Thus, in this paper, we propose a novel method for distant supervised relation extraction, which employs unsupervised deep clustering to generate reliable labels for noisy sentences. Specifically, our model contains three modules: a sentence encoder, a noise detector and a label generator. The sentence encoder is used to obtain feature representations. The noise detector detects noisy sentences from sentence-bags, and the label generator produces high-confidence relation labels for noisy sentences. Extensive experimental results demonstrate that our model outperforms the state-of-the-art baselines on a popular benchmark dataset, and can indeed alleviate the noisy labeling problem.

Introduction

Relation Extraction, defined as the task of extracting structured relations from unstructured text, is a crucial task in natural language processing (NLP). One of the main challenges of relation extraction is the lack of large-scale manually labeled data. Thus, Mintz et al. (2009) proposed distant supervision to automatically construct training data. The assumption of distant supervision is that if two entities (e_1, e_2) have a relationship r in knowledge graph, then any sentence that mentions the two entities might express the relation r .

Obviously, this assumption is too strong and will cause the noisy labeling problem. Since it only focuses on the existence of entities in text and knowledge graph, but cannot identify the one-to-one mapping between sentences and relations. For example, as shown in Table 1, (*Barack Obama,*

	Sentence	Bag Label	Noise?	Correct Label
Bag	#1: Barack Obama was born in the United States .		Yes	born in
	#2: Barack Obama was the first African American to be elected to the president of the United States .	president of	No	president of
	#3: Barack Obama served as the 44th president of the United States from 2009 to 2017.		No	president of

Table 1: An example of sentence-bag annotated by distant supervision. “Yes” and “No” indicate whether or not each sentence is a noisy sentence. “Correct Label” means the true relationship between the entity pair expressed in each sentence.

president of, United States) is a relational triple in knowledge graph. Distant supervision will regard all sentences that contain [*Barack Obama*] _{e_1} and [*United States*] _{e_2} as the instance of relation “*president of*”. As a consequence, the first sentence which expresses the relation “*born in*” is wrongly labeled with relation “*president of*”, and becomes a noisy sentence in the sentence-bag.

Previous studies usually adopt Multi-Instance Learning (MIL) framework to address this problem (Riedel, Yao, and McCallum 2010). In this framework, the training and test process is proceeded at the sentence-bag level, where the sentence-bag contains all the sentences that mention the same triple (e_1, r, e_2). Existing MIL studies broadly fall into two categories: One is the **soft** decision methods, which tend to place soft weights on sentences to reduce the impact of noisy sentences (Lin et al. 2016; Yuan et al. 2019a; 2019b; Ye and Ling 2019). The other is the **hard** decision methods that try to remove noisy sentences from sentence-bags to eliminate their influence (Zeng et al. 2015; Feng et al. 2018; Qin, Xu, and Wang 2018).

However, previous de-noising methods ignore the essential cause of the noisy labeling problem — the lack of correct relation labels. To fill this gap, we try to solve this problem from the perspective of noisy sentences utilization, i.e., correcting their wrong labels. As shown in Table 1, “*Barack Obama was born in the United States*” is a noisy sentence in the sentence-bag. While it indeed expresses the relation “*born in*” between [*Barack Obama*] _{e_1} and [*United States*] _{e_2} .

Intuitively, if we can change its relation label from “*president of*” to “*born in*”, it will be transformed into a useful training instance. This idea brings two advantages: (1) The negative influence of noisy sentences is reduced. (2) The number of useful training data is increased.

In this paper, we propose a novel **Deep Clustering based Relation Extraction** model, named DCRE, which employs unsupervised deep clustering to generate high-confidence labels for noisy sentences. More specifically, DCRE consists of three modules: a sentence encoder, a noise detector and a label generator. The sentence encoder is adopted to derive the sentence representation and shared by the other two modules. The noise detector selects noisy sentences from sentence-bags according to the matching degree between sentences and the bag-level target relations. The sentence who scored below a certain threshold will be treated as noisy sentence. The label generator produces reliable labels for noisy sentences with the help of the deep clustering neural network. Because the results of unsupervised clustering may have errors, we further utilize clustering confidences as weights to scale the loss function. Experimental results show that our model performs better than the state-of-the-art baselines. Our contributions are as follows:

- Different from existing bag-level de-noising methods, our model tries to convert the noisy sentences as useful training data and can simultaneously decrease noisy data and increase useful data.
- To the best of our knowledge, it is the first work to apply unsupervised deep clustering to obtain more appropriate relation labels for noisy sentences.
- Extensive experiments show that our model outperforms the state-of-the-art baselines, and can effectively alleviate the noisy labeling problem.

Related Work

This paper proposes an unsupervised deep clustering based distant supervised relation extraction model. Related works to this paper mainly includes:

Distant Supervised Relation Extraction

Distant supervision (Mintz et al. 2009) is proposed to obtain large-scale training data automatically and has become the standard method for relation extraction. However, the training data generated by distant supervision often contain amounts of noisy sentences. Therefore, noise-reduction has become the mainstream of distant supervised relation extraction. According to the way of processing noisy sentences, existing de-noising methods can be divided into three categories:

The first category of methods tend to assign soft weights on sentences or sentence-bags. By conducting selective attention, it allows the model to focus more on the sentences of a higher quality and reduce the impact of noisy sentences. For example, Lin et al. (2016) employ attention mechanism by distributing different weights to each sentence to capture the bag representation. Yuan et al. (2019a) use non-independent and identically distributed relevance of sentences to obtain the weights of each sentence. Yuan et al.

(2019b) utilize cross-relation cross-bag selective attention to reduce the impact of noisy sentences. Ye and Ling (2019) consider both intra-bag and inter-bag attention in order to deal with noisy sentences at sentence-level and bag-level.

The second category of methods try to remove noisy sentences from sentence-bags through hard decision. For example, Zeng et al. (2015) select the most correct sentence from each bag and ignore other sentences. Feng et al. (2018) employ reinforcement learning to train an instance selector and remove the wrong samples from sentence-bags. Qin, Xu, and Wang (2018) also use reinforcement learning to process noisy sentences. Different from (Feng et al. 2018), they re-distribute noisy sentences into negative samples.

Different from the first two categories, the third type of approaches do not directly process noisy sentences during training stage. For example, Takamatsu, Sato, and Nakagawa (2012) use syntactic patterns to identify the latent noisy sentences and remove them during the pre-processing stage. Wang et al. (2018) avoid using noisy relation labels and employs $e_2 - e_1$ as soft label to train the model. Wu, Fan, and Zhang (2019) propose a linear layer to obtain the connection between true labels and noisy labels. Then, conduct final prediction based on only the true labels.

Similar with the first two categories, the model proposed in this paper also directly processes noisy sentences during training stage. The main difference between DCRE and other methods is: DCRE tries to convert the noisy sentences into meaningful training data. It can simultaneously reduce the number of noisy sentences and increase the number of useful sentences.

Unsupervised Deep Clustering

There are broadly two types of deep clustering algorithms:

The first category of algorithms directly take advantage of the low dimensional features learned by other neural networks, and then run conventional clustering algorithm like k -means. For example, Tian et al. (2014) utilize auto-encoder to learn feature representations, and then obtain the clustering results by conducting k -means algorithm.

The second category tries to learn the feature representation and clustering in an end-to-end mechanism. Among these methods, Deep Embedded Clustering (DEC) is a specialized clustering technique (Xie, Girshick, and Farhadi 2016). This method employs a stacked auto-encoder learning approach. After obtaining the hidden representation of the auto-encoder by pre-training, the encoder pathway is fine-tuned by a defined Kullback-Leibler divergence clustering loss. Guo et al. (2017a) consider that the defined clustering loss will corrupt the feature space and the pre-train process is too complicated, so they keep the decoder remained and add a re-construction loss. Since then, there have been increasing algorithms based on such deep clustering framework (Ghasedi Dizaji et al. 2017; Guo et al. 2017b).

The deep clustering neural architecture proposed in our work falls into the second category. It utilizes the features produced by the pre-trained sentence encoder, and then, jointly optimizes the feature representation and clustering in an end-to-end way.

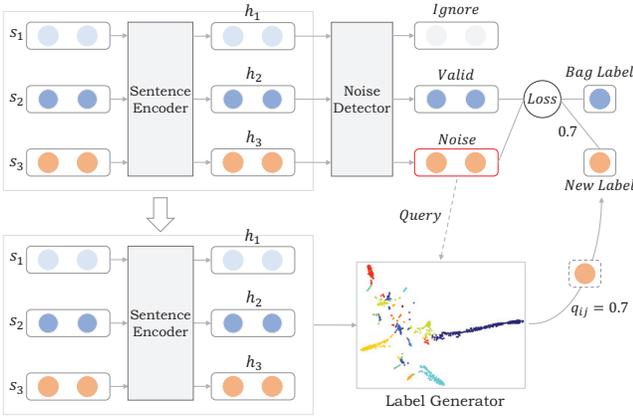


Figure 1: The architecture of DCRE, illustrating the procedure of handling one sentence-bag which contains three sentences.

Method

In this section, we present our method of distant supervised relation extraction. The architecture of the neural network is illustrated in Figure 1. It shows the procedure of handling one sentence-bag. Our model contains three main modules: a sentence encoder, a noise detector and a label generator. In the following, We first give the task definition and notations. Then, we provide detail formalization of the three modules.

Task Definition and Notation

We define the relation classes as $\mathbb{R} = \{r_1, r_2, \dots, r_k\}$, where k is the number of relations. Given a bag of sentences $\mathbb{S}_b = \{s_1, s_2, \dots, s_b\}$ consisting of b sentences and an entity pair (e_1, e_2) presenting in all sentences. The purpose of distant supervised relation extraction is to predict the relation r_i of sentence-bag \mathbb{S}_b according to the entity pair (e_1, e_2) . Therefore, the relation extraction is defined as a classification problem.

Sentence Encoder

When conducting relation extraction, sentences need to be transformed into low-dimensional vectors. We first transform a sentence into a matrix with word embeddings and position embeddings. Then, a Piece-wise Convolution (PCNN) (Zeng et al. 2015) layer is used to obtain the final sentence representation.

Word Representation In a sentence s , each word w_i is first mapped into a d_w -dimensional word embedding v_i . The position features (PFs) proposed by (Zeng et al. 2014) are adopted in this work to specify the target entity pair and make the model pay more attention to the words close to the target entities. PFs are a series of relative distances from the current word to the two entities. The position embedding $p_i^{e_1}, p_i^{e_2}$ are low dimensional vectors of PFs. The final representation x_i of each word w_i is the concatenation of the word embedding and two position embeddings $[v_i; p_i^{e_1}; p_i^{e_2}]$. Then the input sentence representation is:

$$\mathbf{X} = \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_l}, \quad (1)$$

where n_l is the length of the sentence.

PCNN We employ PCNN as our feature extractor, which mainly consists of two parts: one-dimensional convolution and piece-wise max-pooling.

One-dimensional convolution is an operation between a matrix of weights \mathbf{W} , and a matrix of inputs viewed as a sequence \mathbf{X} . \mathbf{W} is regarded as the *filter* for the convolution and x_i is a input vector associated with the i -th word in the sentence. In general, let $x_{i:j}$ refer to the concatenation of x_i to x_j , w refer to filter size. The convolution is to take the dot product of the vector \mathbf{W} with each w -gram in the sentence \mathbf{X} to obtain another sequence \mathbf{m}_i :

$$\mathbf{m}_i = \mathbf{W}^T \mathbf{x}_{i-w+1:i}. \quad (2)$$

The number of \mathbf{m}_i is $n_l - w + 1$. In our model, each sentence is padded with padding-elements such that the number of vector \mathbf{m}_i is equal to the length of the sentence n_l . The convolution result is a feature map matrix $\mathbf{M} = \{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_{n_l}\}$. The number of feature map \mathbf{M}_i is n_f , where n_f is the number of filters.

Piece-wise max-pooling is used to capture the structural information of sentences. After convolution layer, each feature map \mathbf{M}_i is divided into three parts $\{\mathbf{M}_{i1}, \mathbf{M}_{i2}, \mathbf{M}_{i3}\}$ by the position of two entities. Then, the max-pooling operation is performed on the three parts separately. The final sentence representation \mathbf{h} is the concatenation of all vectors:

$$\mathbf{h} = [p_{i1}, p_{i2}, p_{i3}], \quad (3)$$

where $p_{ik} = \max(\mathbf{M}_{ik})$, $k \in \{1, 2, 3\}$. To prevent over-fitting, the dropout strategy (Srivastava et al. 2014) is applied to sentence representation matrices.

Noise Detector

The purpose of noise detector is to select noisy sentences from sentence-bags and feed them to the label generator. Let $\mathbf{H}_b = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_b\}$, $\mathbf{H}_b \in \mathbf{R}^{b \times d_s}$ represents the representation of a sentence-bag, \mathbf{h}_i is the d_s -dimensional sentence representation produced by sentence encoder, and b is the number of sentences. Let $\mathbf{L} = \{l_1, l_2, \dots, l_k\}$, $\mathbf{L} \in \mathbf{R}^{k \times d_s}$ denotes the representations of all the relations. Firstly, a simple dot product between the sentence representation \mathbf{h}_i and the vector of bag-level relation label l_j is adopted to calculate the coupling coefficient as:

$$a_i = \mathbf{h}_i \mathbf{l}_j^T. \quad (4)$$

Then, the coupling coefficient is normalized at the bag-level through a softmax function:

$$a_i = \frac{\exp(a_i)}{\sum_b \exp(a_i)}, \quad (5)$$

where each a_i corresponds to the matching degree between each sentence and the target relation. It represents the possibility that the original relation label is correct for the

current sentence. We set a threshold ϕ to detect noisy sentences, and the sentence whose coupling coefficient is less than ϕ will be regarded as a noisy sample.

However, we can't guarantee that the sentences with a higher coefficient are not wrongly labeled. For this indetermination, our solution is to use the currently deterministic sentences. In a sentence-bag, we consider the best scored sentence as valid sample. The sentences that are neither determined as noisy sentences (scored below ϕ), nor determined as valid samples (best score) will be ignored. The reasons behind this operation are: (1) If the best scored sentence indeed expresses the target relation, it is consistent with the *expressed-at-least-once* assumption (Riedel, Yao, and McCallum 2010). This assumption believes that in a sentence-bag, *at least one sentence* might express the target relation. (2) If the best scored sentence is a noisy sample, in other words, all the sentences in the bag are noisy samples and the sentence-bag is a *noisy bag* (Ye and Ling 2019), ignoring uncertain samples is actually removing the noisy sentences. In both cases, re-labeling high-confidence noisy sentences will benefit the model's performance.

Label Generator

The label generator provides high-confidence relation labels for noisy sentences based on the deep clustering neural network. Let $\mathbf{H} = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n\}$, $\mathbf{H} \in \mathbf{R}^{n \times d_s}$ denotes the representations of all sentences produced by the pre-trained sentence encoder, and $\mathbf{L} = \{\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_k\}$, $\mathbf{L} \in \mathbf{R}^{k \times d_s}$ denotes the pre-trained relation matrix. Firstly, we project sentence representations into relation feature space:

$$\mathbf{C} = \mathbf{H}\mathbf{L}^T + \mathbf{b}, \quad (6)$$

where \mathbf{b} is a bias. This operation can be viewed as an attention with all the relations as query vectors to calculate the relation-aware sentence representations. Then, we feed \mathbf{C} into the clustering layer, which maintains cluster centers $\{\boldsymbol{\mu}_i\}_{i=1}^{n_c}$ as trainable weights, where n_c is the cluster number. We use the Student's t -distribution (Maaten and Hinton 2008) as a kernel to measure the similarity q_{ij} between the feature vector \mathbf{c}_i and the cluster center $\boldsymbol{\mu}_j$:

$$q_{ij} = \frac{(1 + \|\mathbf{c}_i - \boldsymbol{\mu}_j\|^2)^{-1}}{\sum_j (1 + \|\mathbf{c}_i - \boldsymbol{\mu}_j\|^2)^{-1}}, \quad (7)$$

where q_{ij} is the similarity between the projected sentence vector \mathbf{c}_i and the cluster center vector $\boldsymbol{\mu}_j$. It also can be interpreted as the probability of assigning the sentence s_i with relation label r_j .

The loss function of deep clustering is defined as a Kullback-Leibler divergence:

$$\mathcal{L} = KL(P\|Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}, \quad (8)$$

where P is the target distribution. The relations in NYT-10 follow a long-tail distribution, as shown in Figure 2. To alleviate this data imbalance problem, we use the same P with (Xie, Girshick, and Farhadi 2016), defined as:

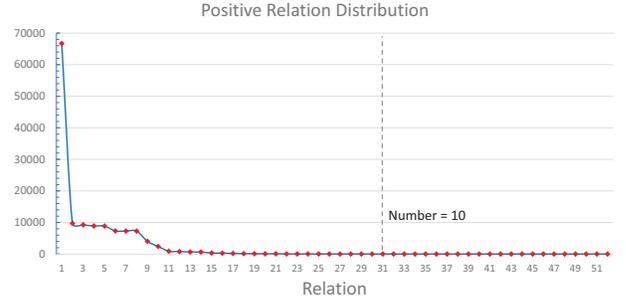


Figure 2: The distribution of 52 positive relations (exclude *NA*) in NYT-10 dataset. The horizontal axis shows different relations sorted by the number of occurrence. The vertical axis shows the number of sentences in training set. The vertical line indicates that the relation whose id is 31 appears 10 times in the training set.

$$p_{ij} = \frac{q_{ij}^2 / \sum_i q_{ij}}{\sum_j (q_{ij}^2 / \sum_i q_{ij})}. \quad (9)$$

This target distribution can normalize the loss contribution of each centroid to prevent large clusters from distorting the hidden feature space.

Noted that we only generate new labels for positive samples, i.e., the samples whose original labels are not *NA* (No relations). Because the representations of sentences which express no relations are always diverse and it's difficult to find correct labels for them. Allowing negative samples to be re-labeled will produce more noisy sentences. On the contrary, a positive sample is re-labeled as *NA* means that the noisy sentence is removed.

Scaled Loss Function

Because there is not any explicit supervision for the noisy data, it's difficult to know whether the clustering result of each sentence is correct. Thus, the new labels produced by the label generator may still be wrong.

To tackle this problem, as mentioned above, we set a threshold ϕ and select sentences with high-confidence to be noisy samples. Further more, we introduce a scaling factor q_{ij} as weight to scale the cross-entropy (Shore and Johnson 1980) loss function. The q_{ij} is obtained by equation (7), which denotes the probability of the i -th sentence is belonged to the j -th relation cluster. This scaling factor makes the new labels have different influence on the model according to their clustering confidence. Finally, the object function is defined as:

$$\mathcal{J}(\theta) = - \sum_{(x_i, y_i) \in \mathbb{V}} \log p(y_i | x_i; \Theta) - \lambda \sum_{(x_i, y_i) \in \mathbb{N}} q_{ij} \log p(y_j | x_i; \Theta), \quad (10)$$

where (x_i, y_i) is a training instance, means the target relation label of sentence x_i is y_i . y_j indicates that the new label for x_i is y_j , and $y_j \neq y_i$. λ is the coefficient that balances the two terms. \mathbb{V} is the best scored samples, \mathbb{N} is the noisy samples, Θ indicates all parameters of the model.

Experiments

Our experiments are designed to demonstrate that DCRE can alleviate the noisy labeling problem. In this section, we first introduce the dataset and evaluation metrics. Second, we show the experiment setup. Third, we compare the performance of our model with several state-of-the-art approaches. Fourth, we make parameter analysis of the threshold ϕ . Finally, we show some details of clustering results.

Data and Evaluation Metrics

We evaluate the proposed method on a widely used dataset NYT-10 (Riedel, Yao, and Mccallum 2010) which was constructed by aligning relation facts in Freebase (Bollacker et al. 2008) with the New York Times (NYT) corpus. Sentences from 2005-2006 are used for training and sentences from 2007 are used for testing. Specifically, it contains 522,611 sentences, 281,270 entity pairs, and 18,252 relational facts in the training data; and 172,448 sentences, 96,678 entity pairs and 1,950 relational facts in the test data. There are 53 unique relations including a special relation *NA* that signifies no relation between the entity pair.

Following the previous works (Yuan et al. 2019b; 2019a; Ye and Ling 2019), we evaluate our model and baselines in the held-out evaluation and present the results with precision-recall curves. In held-out evaluation, the relations extracted from testing data are automatically compared with those in Freebase. It is an approximate measure of the model without requiring costly human evaluation.

Experiment Setup

During training, we first generate the new relation labels for noisy sentences by unsupervised deep clustering, and then, train the whole model. When conducting clustering, we employed *k*-means to initialize the cluster centers for faster convergence. Both over-sampling and under-sampling strategies are applied to highlight the importance of positive samples and alleviate the data imbalance problem. For every positive sample, we obtain multiple clustering results and determine its final category by voting. Besides, we ignored 6 long-tail relations that appear less than 2 times.¹

For all the baselines, during training, we follow the settings used in their papers. Table 2 shows the main parameters used in our DCRE.

Experiment Results

Baselines The model proposed in this work directly process noisy sentences during the training stage, so we select seven related works as baselines. Among these methods, PCNN+ATT, PCNN+WN, PCNN+C2SA and PCNN+ATT_RA+BAG_ATT are soft decision methods,

¹The ignored relations are:
/business/shopping_center_owner/shopping_centers_owned,
/location/fr_region/capital,
/location/mx_state/capital,
/business/shopping_center/owner,
/location/country/languages_spoken,
/base/locations/countries/states_provinces_within.

Setting	Number
Kernel size	3
Feature maps	230
Word embedding dimension	50
Position embedding dimension	5
Pre-train learning rate	0.4
Clustering learning rate	0.004
Model learning rate	0.1
Threshold ϕ	0.1
Dropout	0.5
Coefficient λ	0.6

Table 2: Parameters Setting

PCNN, PCNN+ATT+RL1 and PCNN+ATT+RL2 are hard decision methods.

1. **PCNN+ATT**: Lin et al. (2016) propose a selective attention over sentences based on PCNN sentence encoder.
2. **PCNN+WN**: Yuan et al. (2019a) propose a non-independent and identically distributed relevance to capture the relevance of sentences in the bag.
3. **PCNN+C2SA**: Yuan et al. (2019b) propose cross-relation cross-bag selective attention to handle noisy sentences.
4. **PCNN+ATT_RA+BAG_ATT**: Ye and Ling (2019) propose intra-bag and inter-bag attention to deal with noisy sentences at both sentence-level and bag-level. It is the state-of-the-art method.
5. **PCNN**: Zeng et al. (2015) propose a method to select the most well labeled sentence from sentence-bag, and ignore the other sentences.
6. **PCNN+ATT+RL1**: Feng et al. (2018) propose a reinforcement method to remove noisy sentences from sentence-bags.
7. **PCNN+ATT+RL2**: Qin, Xu, and Wang (2018) also propose a reinforcement model. Different from (Feng et al. 2018), it redistributes noisy sentences into negative examples.

We implement PCNN+ATT, PCNN+WN, PCNN, and our DCRE. For the sake of fairness, we evaluate PCNN+C2SA², PCNN+ATT_RA+BAG_ATT³, PCNN+ATT+RL2⁴ with the codes provided by authors, and replace their training data with the one has 522,611 training sentences. For PCNN+ATT+RL1, we use the source code provided by Open-NRE⁵.

Overall Performance of DCRE The overall performance of DCRE compared with the seven baselines is shown in

²https://github.com/yuanyu255/PCNN_C2SA

³<https://github.com/ZhixiuYe/Intra-Bag-and-Inter-Bag-Attentions>

⁴<https://github.com/Panda0406/Reinforcement-Learning-Distant-Supervision-RE>

⁵<https://github.com/thunlp/OpenNRE>

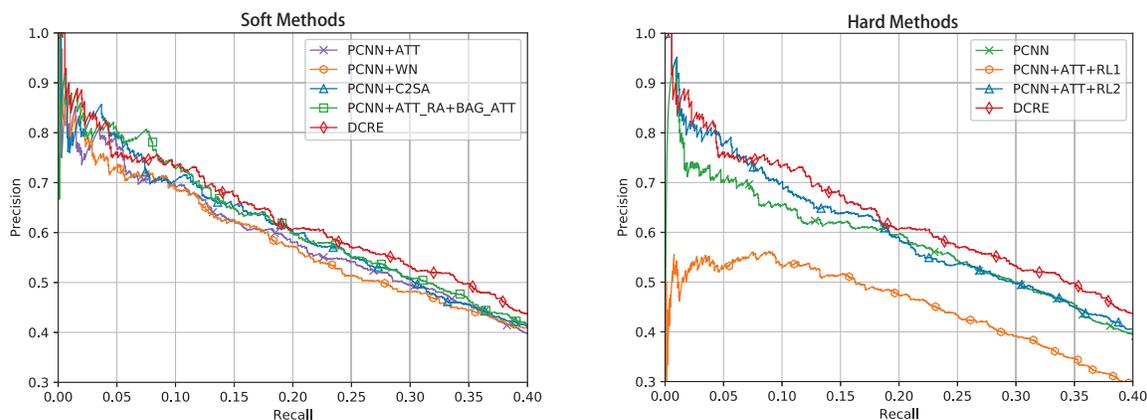


Figure 3: Comparison results with soft methods (left) and hard methods (right).

Figure 3. The left sub-figure shows the comparison results with soft decision methods. The right sub-figure exhibits the comparison results with hard decision methods. We can find that, our DCRE captures the best performance among all the baselines.

Compare DCRE with four soft decision methods, as shown in the left part of Figure 3, we have the following observations: (1) DCRE performs much better than PCNN+ATT and PCNN+WN. This illustrates that assigning low weights to noisy sentences can only reduce its negative influence, but cannot eliminate the impact of noisy sentences. (2) DCRE performs slightly better than PCNN+ATT_RA+BAG_ATT and PCNN+C2SA. Different from PCNN+ATT, PCNN+WN and DCRE, these two methods obtain the final representation according to the *superbag* (Yuan et al. 2019b), so that they utilize a wider range of information. While DCRE is still better than them. This demonstrates that the motivation of our work that finding high-confidence labels for noisy sentences is effective.

The right part of Figure 3 shows the results of comparing DCRE with three hard decision models. Ideally, the instance selector trained by reinforcement learning can remove all the noisy sentences and the hard decision methods should perform better than the soft decision methods. However, it can be observed that there is an obvious margin between DCRE and the other methods. We believe that this is mainly because: (1) As the principle of held-out evaluation, there are also noisy sentences in the test data. Therefore, it’s difficult to use the evaluating results as reward to train the instance selector. (2) Deleting noisy sentences can indeed eliminate their influence, but ignores the useful information contained in noisy sentences. This results have further verified our intuition for re-labeling the noisy sentences and convert them into useful training data.

Effect of the Threshold The most important hyperparameter of DCRE is the threshold ϕ . To analyze how the ϕ affects the performance, we conduct experiments by selecting ϕ in the set $\{0.15, 0.1, 0.05, 0\}$, and the results of different thresholds are shown in Figure 4. It can be found that the setting of ϕ is a process of reconciling contradictions

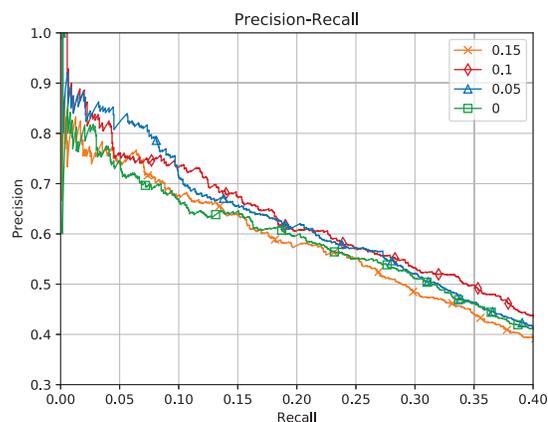


Figure 4: Effect of the threshold ϕ .

and the model performs best when $\phi = 0.1$. The reasons behind this phenomenon are: (1) A large ϕ means that some relatively high scored but not valid samples will be treated as noisy sentences. In other words, the original relation label and the new relation label of these sentences are all probably wrong. (2) A relatively small ϕ indicates that the filtered sentences are more likely to be noisy sentences. While in this situation, the recall rate is too low. (3) When $\phi = 0$, the model is equal to PCNN which only utilizes the best scored sentence.

Clustering Result We further verify the effectiveness of our deep clustering neural network by visualizing the clustering results during training. We set the number of clusters as 47, excluding 6 long-tail relations which appear less than 2 times in training data. We don’t remove more long-tail relations because the noisy sentences labeled with other relations may be clustered into these clusters.

We randomly select 1000 sentences that covering all the 47 categories and visualize them by t -SNE (Maaten and Hinton 2008), as shown in Figure 5. It can be found that there is a clear boundary between different clusters in epoch 1, but this clustering result cannot be used to re-label the

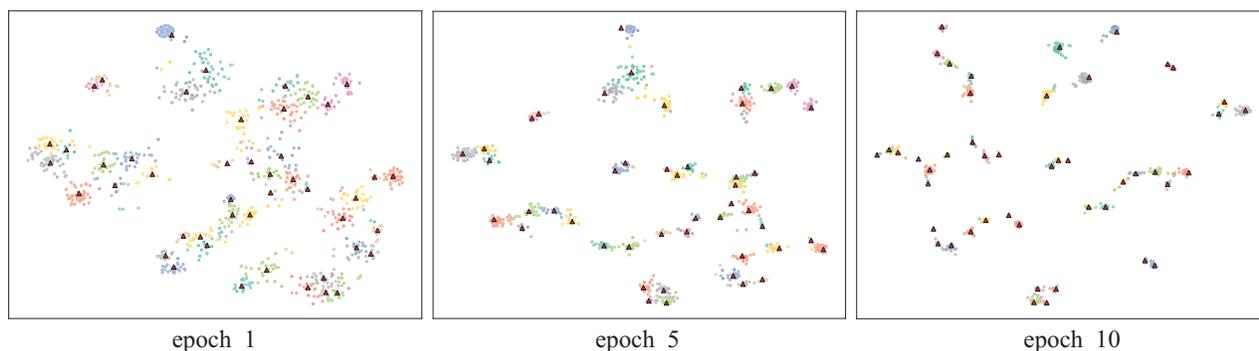


Figure 5: t -SNE visualization of clustering results on subset of NYT-10. The red triangles are cluster centers.

ID	Entity pair	Sentence	Original label	Generated label	Correct?
1	(China, Beijing)	Beijing has tried to enlist the support of Uzbekistan in fighting Islamic separatism in China 's western region of Xinjiang, while also lining up secure supplies of oil and gas.	/location/location/contains	/location/cn province /capital	No
2	(Italy, Rome)	Mr. Tomassetti's companies are named after L'Aquila, Italy , his birthplace 58 miles northeast of Rome .	/location/country/capital	/location/location/contains	Yes
3	(Saddam Hussein, Iraq)	As national journal reported in April, it was Senator Roberts who stated as the Iraq war began that the U.S. had "human intelligence that indicated the location of Saddam Hussein ."	/people/deceased person/place of death	/people/person/place lived	Yes
4	(Edith Sitwell, England)	His first book was published privately in his own country and then by a major publisher in England , where he had many supporters in the literary world, most notably Edith Sitwell and Angus Wilson.	/people/person/nationality	/people/person/place of birth	No
5	(Louisiana, New Orleans)	The book, by a New Orleans resident, John M. Barry, describes the history and politics behind a flood that killed 1,000 people and displaced 900,000 from Louisiana to Illinois.	/location/location/contains	NA	Yes

Table 3: Five sentences randomly selected in NYT-10 dataset. The text in **bold** represents the entity, the text underline represents the relation label is correct.

noisy sentences. Because the distance between the points of the same cluster is too large so that the confidence of the new labels is relatively low. It can be seen from the visualization, from left to right, the "shape" of each cluster is becoming more and more compact and the clusters are becoming increasingly well separated. Accordingly, the confidence of the new label is becoming more high.

In epoch 10, it can be found that some points of different clusters are close to each other. This phenomenon is consistent with reality, i.e., a sentence may express more than one relations. For example, *[Barack Obama]_{e1} is the 44th president of the [United States]_{e2}* is a high-quality sentence for relation "president of" and a low-quality sentence for relation "live in". Ideally, the sentence should be grouped into both clusters. In this paper, we only consider one high-confidence relation label for each noisy sentence.

Furthermore, we randomly select five re-labeled sentences during training whose new relation labels are different from original labels to show the capabilities of noisy detector and label generator. Their target entity pairs, original labels and generated labels are illustrated in Table 3. It can be found that: (1) The original labels of the five sentences are wrong. This proves the validity of the threshold ϕ . (2) The correct label for sentence 1 is */location/country/capital*. The original label is wrong because distant supervision cannot identify that the entity *[Beijing]* represents the *Chi-*

nese government. The generated label is wrong mainly due to the word *China*. (3) Exactly speaking, the correct label for sentence 4 is *people/person/place lived*. Its original label is */people/person/nationality* and generated label is */people/person/place_of_birth*. The three relations have inner connections so that it's difficult for the model to find the correct one.

Conclusion and Future Work

In this paper, we proposed an unsupervised deep clustering based distant supervised relation extraction model. Different from conventional methods which focus on reducing the influence of noisy sentences, our model tries to find new relation labels for noisy sentences and convert them into useful training data. Extensive experimental results show that the proposed method performs better than comparable approaches, and can indeed alleviate the noisy labeling problem in distant supervised relation extraction. In the future, we will explore the following directions:

- Our clustering algorithm assigns one relation label for each noisy sentence. While in reality, one sentence may express multiple relations. We will consider multi-class clustering in the future.
- The threshold ϕ is very important in DCRE. We will next develop an end-to-end method to automatically se-

lect noisy sentences and avoid manual intervention.

Acknowledgments

This work is supported by National Key R&D Plan (No.2016QY03D0602), NSFC (No. 61772076 and 61751201), NSFB (No. Z181100008918002), Major Project of Zhijiang Lab (No. 2019DH0ZX01), and Open fund of BDAIGGCNEL and CETC Big Data Research Institute Co., Ltd (No. w-2018018).

References

- Bollacker, K.; Evans, C.; Paritosh, P.; Sturge, T.; and Taylor, J. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, 1247–1250. AcM.
- Feng, J.; Huang, M.; Zhao, L.; Yang, Y.; and Zhu, X. 2018. Reinforcement learning for relation classification from noisy data. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Ghasedi Dizaji, K.; Herandi, A.; Deng, C.; Cai, W.; and Huang, H. 2017. Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization. In *Proceedings of the IEEE International Conference on Computer Vision*, 5736–5745.
- Guo, X.; Gao, L.; Liu, X.; and Yin, J. 2017a. Improved deep embedded clustering with local structure preservation. In *IJCAI*, 1753–1759.
- Guo, X.; Liu, X.; Zhu, E.; and Yin, J. 2017b. Deep clustering with convolutional autoencoders. In *International Conference on Neural Information Processing*, 373–382. Springer.
- Lin, Y.; Shen, S.; Liu, Z.; Luan, H.; and Sun, M. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 2124–2133.
- Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of machine learning research* 9(Nov):2579–2605.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 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: Volume 2-Volume 2*, 1003–1011. Association for Computational Linguistics.
- Qin, P.; Xu, W.; and Wang, W. Y. 2018. Robust distant supervision relation extraction via deep reinforcement learning. *meeting of the association for computational linguistics* 1:2137–2147.
- Riedel, S.; Yao, L.; and McCallum, A. 2010. Modeling relations and their mentions without labeled text. In *European Conference on Machine Learning Knowledge Discovery in Databases*.
- Shore, J., and Johnson, R. 1980. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on information theory* 26(1):26–37.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15(1):1929–1958.
- Takamatsu, S.; Sato, I.; and Nakagawa, H. 2012. Reducing wrong labels in distant supervision for relation extraction. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, 721–729. Association for Computational Linguistics.
- Tian, F.; Gao, B.; Cui, Q.; Chen, E.; and Liu, T.-Y. 2014. Learning deep representations for graph clustering. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Wang, G.; Zhang, W.; Wang, R.; Zhou, Y.; Chen, X.; Zhang, W.; Zhu, H.; and Chen, H. 2018. Label-free distant supervision for relation extraction via knowledge graph embedding. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2246–2255.
- Wu, S.; Fan, K.; and Zhang, Q. 2019. Improving distantly supervised relation extraction with neural noise converter and conditional optimal selector. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7273–7280.
- Xie, J.; Girshick, R.; and Farhadi, A. 2016. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, 478–487.
- Ye, Z.-X., and Ling, Z.-H. 2019. Distant supervision relation extraction with intra-bag and inter-bag attentions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2810–2819. Minneapolis, Minnesota: Association for Computational Linguistics.
- Yuan, C.; Huang, H.; Feng, C.; Liu, X.; and Wei, X. 2019a. Distant supervision for relation extraction with linear attenuation simulation and non-iid relevance embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7418–7425.
- Yuan, Y.; Liu, L.; Tang, S.; Zhang, Z.; Zhuang, Y.; Pu, S.; Wu, F.; and Ren, X. 2019b. Cross-relation cross-bag attention for distantly-supervised relation extraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 419–426.
- Zeng, D.; Liu, K.; Lai, S.; Zhou, G.; and Zhao, J. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, 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 the 2015 Conference on Empirical Methods in Natural Language Processing*, 1753–1762.