

Improving Long-Tail Relation Extraction with Collaborating Relation-Augmented Attention

Yang Li¹, Tao Shen^{1*}, Guodong Long¹, Jing Jiang¹, Tianyi Zhou² and Chengqi Zhang¹

¹Australian AI Institute, School of Computer Science, FEIT, University of Technology Sydney

²University of Washington, Seattle

{yang.li-17, tao.shen}@student.uts.edu.au, tianyizh@uw.edu
{guodong.long, jing.jiang, chengqi.zhang}@uts.edu.au

Abstract

Wrong labeling problem and long-tail relations are two main challenges caused by distant supervision in relation extraction. Recent works alleviate the wrong labeling by selective attention via multi-instance learning, but cannot well handle long-tail relations even if hierarchies of the relations are introduced to share knowledge. In this work, we propose a novel neural network, Collaborating Relation-augmented Attention (CoRA), to handle both the wrong labeling and long-tail relations. Particularly, we first propose relation-augmented attention network as base model. It operates on sentence bag with a sentence-to-relation attention to minimize the effect of wrong labeling. Then, facilitated by the proposed base model, we introduce collaborating relation features shared among relations in the hierarchies to promote the relation-augmenting process and balance the training data for long-tail relations. Besides the main training objective to predict the relation of a sentence bag, an auxiliary objective is utilized to guide the relation-augmenting process for a more accurate bag-level representation. In the experiments on the popular benchmark dataset NYT, the proposed CoRA improves the prior state-of-the-art performance by a large margin in terms of Precision@N, AUC and Hits@K. Further analyses verify its superior capability in handling long-tail relations in contrast to the competitors.

1 Introduction

Relation extraction, as a fundamental task in natural language processing (NLP), aims to discriminate the relation between two given entities in a plain text. As recent data-driven algorithms, e.g., deep neural networks (Han et al., 2018; Du et al., 2018; Ye and Ling, 2019), have shown their capabilities in tackling NLP tasks, labor-intensive annotation and scarce training data become the major obstacles of achieving promising performance on relation extraction.

Instead of costly manual labeling, distant supervision method (Mintz et al., 2009) is proposed to auto-label the training data for relation extraction. It labels a sentence containing a pair of entities with the relation between them in a knowledge graph, e.g. Freebase (Bollacker et al., 2008). A strong assumption behind this is that a sentence containing two entities only expresses the relation existing in the knowledge graph, but this assumption will not always hold (Lin et al., 2016). Hence, it leads to two main problems regarding the training data. First, when the assumption is invalid, *wrong labeling problem* appears and degrades algorithms by introducing noisy supervision signals. This problem has been well-studied by recent works (Lin et al., 2016; Han et al., 2018; Ji et al., 2017; Li et al., 2020) operating at bag level for multi-instance learning (Hoffmann et al., 2011), where the “bag” denotes a set of sentences with the same entity pair. Most of them use selective attention to avoid wrongly-labeled sentences.

Second, the *long-tail problem* is caused by using a knowledge graph as distant supervision to auto-label a domain-specific corpus, where the knowledge graph usually suffers from long-tail relations. For example, to build NYT dataset, applying Freebase to a news corpus, New York Times, leads to $\sim 70\%$ of

*Corresponding author.

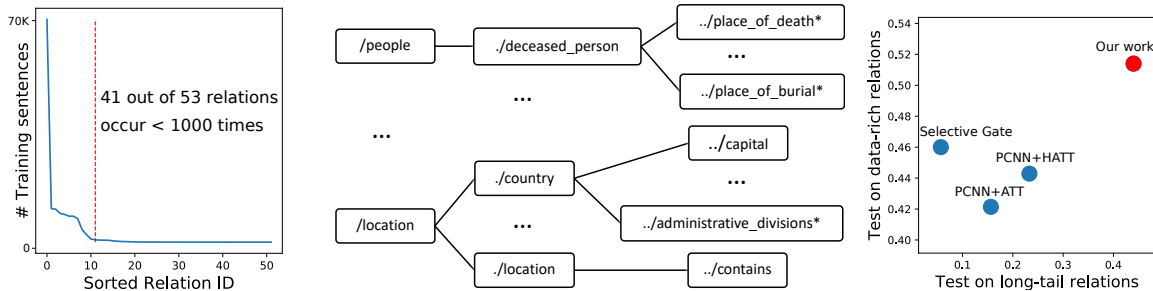


Figure 1: **Left:** Label frequency distribution of relations without *NA* in NYT dataset. Here the criterion being a long-tail relation is the number of corresponding training sentences is less than 1000. **Middle:** Relation hierarchies in Freebase knowledge graph, where relation with “*” is long-tail. Through the common high-level relation, it can be exploited that 1) multiple semantic-related low-level relations complement each other and 2) semantic knowledge is transferred from data-rich to long-tail relations. **Right:** Empirical AUC results of competitive approaches on data-rich and long-tail test subsets.

relations long-tail, as illustrated in Figure 1 (left). This problem seriously disrupts the data balance and thus becomes one of the main barriers of improvement.

To alleviate the long-tail problem, two approaches (Han et al., 2018; Zhang et al., 2019) naturally share the knowledge from data-rich relations to the long-tail ones when those relations have semantic overlap. This semantic overlap or relatedness is usually stored in the relation hierarchies of a knowledge graph, e.g., Freebase in Figure 1 (middle). Specifically, these approaches extend selective attention (Lin et al., 2016) by introducing the embeddings of high-level (i.e., coarse-grained) relations as complements to original low-level (i.e., fine-grained) relations. As such, the high-level relation embeddings are used as queries of selective attention to derive extra bag-level features. To learn the relation embeddings, Han et al. (2018) randomly initialize them followed by supervised learning in end-to-end fashion, whereas Zhang et al. (2019) combine the embeddings from both TransE (Bordes et al., 2013) pre-trained on Freebase and graph convolutional network (Defferrard et al., 2016) applied to the relation’s hierarchies.

Despite proven to improve overall and long-tail performance, they also post two issues: 1) Limited by selective attention framework, the relation embeddings are only used as the attention’s queries and thus not well-exploited to share knowledge. 2) Despite the capability in mitigating the long-tail problem, graph embeddings pre-trained on large-scale knowledge graph are time-consuming and not always off-the-shelf, hence at the cost of practicability.

In this work, we propose a novel neural network, named as **Collaborating Relation-augmented Attention (CoRA)**, to tackle distantly supervised relation extraction, where no external knowledge is introduced and the relation hierarchies are fully utilized to alleviate the long-tail problem. Specifically, as an alternative to selective attention framework, we first propose a base model, *relation-augmented attention*, operating at bag level to minimize the effect of wrong labeling, where the relation-augmenting process is fulfilled by a sentence-to-relation attention. Empowered by the base model, we then leverage the high-level relations for collaborating features in light of the relation hierarchies. Besides a further relief of wrong labeling, such features facilitate knowledge transfer among the low-level relations inheriting a common high-level relation.

Intuitively, selective attention and its hierarchical extensions learn relation label embeddings to score each sentence in a bag. In contrast, the proposed relation-augmented attention network achieves the same goal via a memory network-like structure: sentences equipped with relation features are passed into an attention-pooling (i.e., a kind of self-attention (Lin et al., 2017)) for bag-level representations. Our method is especially effective when extended to multi-granular relations – the features are enriched by cross-relation sharing, which hence benefits long-tail relations. As shown in Figure 1 (right), our proposed approach achieves consistent outstanding performance on both data-rich and long-tail relations.

We use two objectives to jointly train the CoRA. The first is predicting the relation label at bag level, which is the goal of relation extraction. As auxiliary objective, the second is guiding the model to equip each sentence with correct multi-granular relation embeddings during the augmenting process. It

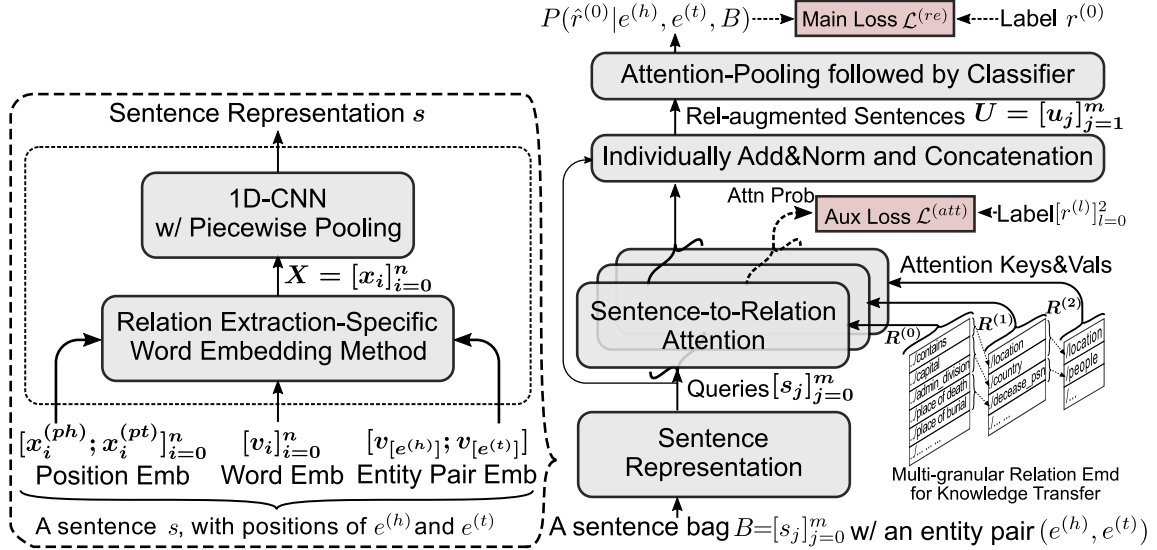


Figure 2: Our proposed **Collaborating Relation-augmented Attention (CoRA)** Network, where the *right part* is the main structure while the *left part* is a sentence embedding method for relation extraction. The illustrated relations and their hierarchies are based on NYT dataset where $M = 2$ in Eq.(14).

aims to boost downstream attention-pooling and is fulfilled by applying the multi-granular labels to the sentence-to-relation attention during training.

Our main contributions are summarized as:

- We propose a base model, named relation-augmented attention network, to handle wrong labeling problem in multi-instance learning.
- We then propose to extend the base model with the relation hierarchies, called CoRA, to further promote the performance on long-tail relations.
- We evaluate CoRA on the popular benchmark dataset NYT and set state-of-the-art results in Precision@N, AUC and Hits@K. We also verify its capability in alleviating both wrong labeling and long-tail problems via insightful analyses. The source code of this work is released at <https://github.com/YangLi1221/CoRa>.

2 Proposed Approach

This section begins with a definition of distantly supervised relation extraction with multi-granular relation labels. Then an embedding method is introduced to represent sentences. Lastly, our base model and its hierarchical extension are presented to handle wrong labeling and long-tail relations. An illustration of the model is shown in Figure 2.

2.1 Task Definition

Given a bag of sentences $B = \{s_1, \dots, s_m\}$ in which each sentence contains a pair of head $e^{(h)}$ and tail $e^{(t)}$ entities in common, the distant supervision (Mintz et al., 2009) assigns this bag with a relation label $r^{(0)}$ according to the entity pair in a knowledge graph. The goal of relation extraction is to predict the relation label $\hat{r}^{(0)}$ of an entity pair based on the corresponding sentence bag when the pair is not included in the knowledge graph. As following the hierarchical setting (Han et al., 2018; Zhang et al., 2019), labels of coarse-grained relations, $[r^{(1)}, \dots, r^{(M)}]$, can be used to share knowledge across relations.

2.2 Sentence-Level Representation

To embed each sentence s_j in a bag $B = \{s_1, \dots, s_m\}$ into latent semantic space, we derive a sentence representation from three kinds of features, including word embedding (Mikolov et al., 2013),

position embedding (Zeng et al., 2015) and entity embedding (Li et al., 2020). The integration of them has been proven crucial and effective to relation extraction by previous work (Li et al., 2020). In the following, we omit the index of a sentence, j , for a clear elaboration. Basically, a sentence s is first tokenized into a sequence of n words, $s = [w_1, \dots, w_n]$, then a word2vec method (Mikolov et al., 2013) is used to transform the discrete tokens into low-dimensional, real-valued vector embeddings, i.e., $V = [v_1, \dots, v_n] \in \mathbb{R}^{d_w \times n}$.

Word Embedding. On the one hand, *position-aware embedding* offers rich positional information for downstream modules (Zeng et al., 2014). For i -th word, the relative position is represented as the distances from the word to head $e^{(h)}$ and tail $e^{(t)}$ entities respectively. Two integer distances are then transformed into low-dimensional vectors, $\mathbf{x}_i^{(ph)}$ and $\mathbf{x}_i^{(pt)} \in \mathbb{R}^{d_p}$, by a learnable weight matrix. Consequently, a sequence of position-aware embeddings is denoted as $\mathbf{X}^{(p)} = [\mathbf{x}_1^{(p)}, \dots, \mathbf{x}_n^{(p)}] \in \mathbb{R}^{(d_w+2d_p) \times n}$ where $\mathbf{x}_i^{(p)} = [v_i, \mathbf{x}_i^{(ph)}; \mathbf{x}_i^{(pt)}]$. $[\cdot]$ denotes the operation of vector concatenation. On the other hand, *entity-aware embedding* is also crucial since the goal of relation extraction is to discriminate the relation between two entities. The embedding of head or tail entity is represented by the corresponding word embedding. Note that each entity is one entry in the vocabulary of word embedding even if it is usually composed of multiple words. Hence, a sequence of entity-aware embeddings is denoted as $\mathbf{X}^{(e)} = [\mathbf{x}_1^{(e)}, \dots, \mathbf{x}_n^{(e)}] \in \mathbb{R}^{3d_w \times n}$ where $\mathbf{x}^{(e)} = [v_i, v_{[e^{(h)}]}; v_{[e^{(t)}]}] \in \mathbb{R}^{3d_w}$. To integrate the embeddings above, a position-wise gating procedure is employed by following Li et al. (2020). That is,

$$\mathbf{A}^{(e)} = \text{sigmoid}(\lambda \cdot (\mathbf{W}^{(g1)} \mathbf{X}^{(e)} + \mathbf{b}^{(g1)})), \quad (1)$$

$$\tilde{\mathbf{X}}^{(p)} = \tanh(\mathbf{W}^{(g2)} \mathbf{X}^{(p)} + \mathbf{b}^{(g2)}), \quad (2)$$

$$\mathbf{X} = \mathbf{A}^{(e)} \circ \mathbf{X}^{(e)} + (1 - \mathbf{A}^{(e)}) \circ \tilde{\mathbf{X}}^{(p)}, \quad (3)$$

where “ \circ ” denotes element-wise product $\mathbf{W}^{(g1)} \in \mathbb{R}^{d_x \times 3d_w}$ and $\mathbf{W}^{(g2)} \in \mathbb{R}^{d_x \times (d_w+2d_p)}$ are learnable parameters, λ is a hyper-parameter to control smoothness, and $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d_x \times n}$ is the resulting sequence of word embeddings specially for relation extraction.

Piecewise Convolutional Neural Network. As a common practice in distantly supervised relation extraction, piecewise convolutional neural network (PCNN) (Zeng et al., 2015) is used to generate contextualized representations over an input sequence of word embeddings. Compared to the typical 1D-CNN with max-pooling (Zeng et al., 2014), piecewise max-pooling has the capability to capture the structure information between two entities by considering their positions. Specifically, 1D-CNN (Kim, 2014) is first invoked over the input sequence for contextualized representations. Then a piecewise max-pooling performs over the output sequence to obtain sentence-level embedding. These steps are written as

$$\mathbf{H} = [h_1, \dots, h_n] = \text{1D-CNN}(\mathbf{X}; \mathbf{W}^{(c)}, \mathbf{b}^{(c)}) \in \mathbb{R}^{d_c \times n}, \quad (4)$$

$$s = \tanh([\text{Pool}(\mathbf{H}^{(1)}); \text{Pool}(\mathbf{H}^{(2)}); \text{Pool}(\mathbf{H}^{(3)})]), \quad (5)$$

where $\mathbf{W}^{(c)} \in \mathbb{R}^{d_c \times Q \times d_x}$ is a conv kernel with window size of Q . $\mathbf{H}^{(1)}$, $\mathbf{H}^{(2)}$ and $\mathbf{H}^{(3)}$ are three consecutive parts of \mathbf{H} , obtained by dividing \mathbf{H} w.r.t. indices of head $e^{(h)}$ and tail $e^{(t)}$ entities. Consequently, $s \in \mathbb{R}^{d_h}$, where $d_h = 3d_c$, is the resulting sentence-level representation.

2.3 Relation-Augmented Attention Network

Due to the effectiveness of selective attention (Lin et al., 2016) in multi-instance learning, most recent works employ the selective attention as the baseline and then propose own approaches for improvements in wrong labeling and/or long-tail relations. However, selective attention gradually becomes a bottleneck of performance improvement. For example, Li et al. (2020) find using simple gating mechanism to replace selective attention further alleviates wrong labeling problem and significantly promotes extracting results. Intuitively, on the one hand, employing the basic PCNN and vanilla attention mechanism inevitably limits the expressive power of this framework and thus sets a barrier. On the other hand, the

relation embeddings, similar to label embeddings (Bengio et al., 2010), are crucial to distant supervision relation extraction, but are only used as attention query to score a sentence and thus not well-exploited.

In contrast, we aim to augment each sentence in a bag with the relation embeddings by a sentence-to-relation attention, and pass the relation-augmented representations of a bag’s sentences into an attention-pooling module. The attention-pooling, a kind of self-attention (Lin et al., 2017; Shen et al., 2018b; Shen et al., 2018a), is used to derive an accurate bag-level representation for relation classification. In details, we first define a relation embedding matrix $\mathbf{R}^{(0)} \in \mathbb{R}^{d_h \times N^{(0)}}$ where d_h denotes the size of hidden states and $N^{(0)}$ denotes the number of distinct relations $r^{(0)}$ in a distantly supervised relation extraction task. Then, we formulate a sentence-to-relation (sent2rel) attention as opposed to selective attention, which aims at augmenting sentence representation from §2.2 with relation information. The sentence representation \mathbf{s} is used as a query to attend the relation embedding matrix $\mathbf{R}^{(0)}$ via a dot-product compatibility function:

$$\boldsymbol{\alpha}^{(0)} = \text{softmax}(\mathbf{s}^T \mathbf{R}^{(0)}), \quad (6)$$

$$\mathbf{c}^{(0)} = \mathbf{R}^{(0)} \boldsymbol{\alpha}, \quad (7)$$

where $\text{softmax}(\cdot)$ denotes a normalization function along last dimension and $\mathbf{c}^{(0)}$ is the resulting relation-aware representation corresponding to the sentence \mathbf{s} . Then we merge the relation-aware representation $\mathbf{c}^{(0)}$ into original sentence representation \mathbf{s} by an element-wise gate mechanism with residual connection (He et al., 2016) and layer normalization (Ba et al., 2016), i.e.,

$$\boldsymbol{\beta}^{(0)} = \text{sigmoid}(\mathbf{W}^{(g)}[\mathbf{s}; \mathbf{c}^{(0)}] + \mathbf{b}^{(g)}), \quad (8)$$

$$\tilde{\mathbf{u}}^{(0)} = \boldsymbol{\beta}^{(0)} \circ \mathbf{s} + (1 - \boldsymbol{\beta}^{(0)}) \circ \mathbf{c}^{(0)}, \quad (9)$$

$$\mathbf{u}^{(0)} = \text{LayerNorm}(\mathbf{s} + \text{MLP}(\tilde{\mathbf{u}}^{(0)})), \quad (10)$$

where $\text{MLP}(\cdot)$ denotes a multi-layer perceptron to increase nonlinearity. Finally, we define relation-augmented sentence representation in our base model as

$$\mathbf{u} := \mathbf{u}^{(0)}. \quad (11)$$

Next moving to multi-instance learning, we put each sentence back to its bag $B = \{s_1, \dots, s_m\}$ so the bag of sentences with relation-augmentation is represented as $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_m] \in \mathbb{R}^{d_h \times m}$. Differing from selective attention framework, our sentence representations are augmented by the relation embeddings as elaborated above. Hence, we straightforwardly introduce an attention-pooling module to derive a bag-level representation denoising from the wrongly-labeled sentences. Specifically, the attention-pooling learns to assign each sentence with an importance score according its representation. Then it performs a weighted sum over a bag of sentence representations, where the weights are proportional to their scores. This attention is formulated as

$$\mathbf{b} = \mathbf{U} \text{softmax}(\mathbf{w}^T \mathbf{U}), \quad (12)$$

where \mathbf{w} is a learnable weight vector, and \mathbf{b} denotes the resulting bag-level representation. Lastly, an MLP is used to obtain a categorical distribution over all relations as bag-level prediction:

$$\mathbf{p} = P(\hat{r}^{(0)} | e^{(h)}, e^{(t)}, B) := \text{MLP}(\mathbf{b}) \in \mathbb{R}^{N^{(0)}}. \quad (13)$$

2.4 Collaborating Relation-Augmented Attention Network

Beyond only fine-grained relations used above, high-level relation embeddings as hierarchical knowledge can collaborate with the low-level embeddings to boost the performance by alleviating long-tail problem (Han et al., 2018; Zhang et al., 2019). Intuitively, a high-level relation, shared crossing several low-level relations, is used to represent common knowledge of low-level relations. Therefore, via the common high-level relation, 1) several low-level long-tail relations with semantic overlap mutually benefit each

other, and 2) the semantic knowledge is easily transferred from data-rich relations to long-tail ones. These common knowledge is implicitly utilized to distinguish the coarse-grained relation of a bag and thus benefits the final relation prediction. With the relation-augmented sentence representation further enriched via collaborating, we name it as **Collaborating Relations-augmented Attention (CoRA)**.

Empowered by non-trivial structure design of our base model, high-level relation embeddings can be easily integrated into the base model by re-defining Eq.(11). In particular, given the coarse-grained relation labels from low to high level, i.e., $[r^{(1)}, \dots, r^{(M)}]$, we define a list of relation embedding matrices $[\mathbf{R}^{(1)}, \dots, \mathbf{R}^{(M)}]$ in addition to $\mathbf{R}^{(0)}$ defined in last section. With these relation embedding matrices, we individually generate their corresponding relation-augmented sentence representations, i.e., $[\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(M)}]$, via the same procedure defined in Eq.(6 – 10) of §2.3. Then, we concatenate $[\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(M)}]$ in conjunction with $\mathbf{u}^{(0)}$ to re-formulate Eq.(11) as

$$\mathbf{u} := [\mathbf{u}^{(0)}; \mathbf{u}^{(1)}; \dots, \mathbf{u}^{(M)}] \in \mathbb{R}^{(1+M)d_h}. \quad (14)$$

The following procedure is identical to that in base model elaborated above, except that the learnable weight matrices are up-scaled linearly with the depth of relation hierarchies.

2.5 Training Objectives

The *main objective* for relation extraction is defined to minimize a cross-entropy loss, i.e.,

$$\mathcal{L}^{(re)} = -\frac{1}{|\mathcal{D}|} \sum_{B \in \mathcal{D}} \log P(\hat{r}^{(0)} = r^{(0)} | e^{(h)}, e^{(t)}, B), \quad (15)$$

where \mathcal{D} is the training set consisting of sentence bags. Besides, an *auxiliary objective* guides sentence-to-relation attention modules to augment each sentence with correct relation embeddings. This is critical to perform downstream attention-pooling and overcome the challenges presented by distant supervision. Given the sent2rel attention score $\alpha^{(l)}$ and relation label $r^{(l)}$ at an arbitrary l level, the loss function to achieve this objective is defined as

$$\mathcal{L}^{(att)} = -\frac{1}{|\mathcal{D}| \cdot |B| \cdot (1 + M)} \sum_{B \in \mathcal{D}} \sum_{s \in B} \sum_{l=0}^M \log \alpha_{[r^{(l)}]}^{(l)}. \quad (16)$$

where $M = 0$ for the base model in §2.3, where $M > 0$ for CoRA in §2.4. Finally, we optimize the proposed model by jointly minimizing the two loss functions above, i.e., $\mathcal{L} = \mathcal{L}^{(re)} + \mathcal{L}^{(att)}$.

3 Experiments

We evaluate our proposed network on a popular benchmark dataset and conduct several analyses for insights into our proposed model.

Dataset and Evaluation Metrics. By following previous works (Zeng et al., 2015; Lin et al., 2016; Han et al., 2018), we employ the only popular distantly supervised relation extraction dataset, New York Times (NYT) dataset (Riedel et al., 2010). It contains 53 distinct relations which includes a *NA* class denoting the relation between the entity pair is unavailable. And it consists of 570K and 172K sentences in training and test sets respectively. Two metrics, 1) area under precision-recall curve (AUC) and 2) top-n precision (P@N) are usually used to measure the effectiveness. We also use Hits@K for long-tail relations by following Zhang et al. (2019).

Setups. Following previous works, d_w, d_p, d_x, d_c, d_h and Q are 50, 5, 150, 230, 690 and 3 respectively. λ in Eq.(1) is 0.05. NYT offers two more high-level (coarse-grained) relations (i.e., $M = 2$), and the numbers of distinct relations at three levels are 53, 36 and 9. During training, we use minibatch SGD (Zeiler, 2012) with Adam (Kingma and Ba, 2015) optimizer. The learning is 0.1, batch size is 160, dropout probability is set to 0.5, weight decay of L2 regularization is 10^{-5} .

P@N (%)	One				Two				All				AUC
	100	200	300	Mean	100	200	300	Mean	100	200	300	Mean	
<i>Comparative Approaches</i>													
CNN+ATT (Lin et al., 2016)	72.0	67.0	59.5	66.2	75.5	69.0	63.3	69.3	74.3	71.5	64.5	70.1	0.35
PCNN+ATT (Lin et al., 2016)	73.3	69.2	60.8	67.8	77.2	71.6	66.1	71.6	76.2	73.1	67.4	72.2	0.39
PCNN+HATT (Han et al., 2018)	84.0	76.0	69.7	76.6	85.0	76.0	72.7	77.9	88.0	79.5	75.3	80.9	0.42
PCNN+BAG-ATT (Ye and Ling, 2019)	86.8	77.6	73.9	79.4	91.2	79.2	75.4	81.9	91.8	84.0	78.7	84.8	0.42
SeG (Li et al., 2020)	94.0	89.0	85.0	89.3	91.0	89.0	87.0	89.0	93.0	90.0	86.0	89.3	0.51
CoRA (ours)	94.0	90.5	82.0	88.8	98.0	91.0	86.3	91.8	98.0	92.5	88.3	92.9	0.53
<i>Ablations</i>													
Base* (CoRA w/o Collaborating)	90.0	89.0	85.3	88.1	93.0	90.0	85.3	89.4	93.0	90.5	87.0	90.2	0.52
Base w/o Ent Emb in §2.2	83.0	74.0	69.3	74.5	84.0	81.0	72.3	79.1	85.0	80.0	73.3	79.4	0.45
Base w/o Sent2rel Attention in §2.3	83.0	74.0	66.6	74.5	82.0	79.0	68.3	76.5	84.0	79.5	73.0	78.8	0.43
Base w/o Attention-pooling in §2.3	90.0	87.0	84.0	87.0	93.0	88.0	85.0	88.7	94.0	88.5	86.0	89.5	0.52
Base w/o Aux Obj $\mathcal{L}^{(att)}$ in Eq.(16)	80.0	70.0	65.7	71.9	83.0	74.0	68.0	75.0	85.0	80.0	70.3	78.4	0.41

Table 1: Model Evaluation and ablation study on NYT. ‘‘P@N’’ (top-n precision) denotes precision values for the entity pairs with top-100, -200 and -300 prediction confidences by randomly keeping one/two/all sentence(s) in each bag. *Base model denotes relation-augmented attention network where $M = 0$.

Comparative Approach. We compare the proposed approach with extensive previous works that are summarized as follows. A model with ‘‘*’’ denotes it is proposed for the long-tail problem.

- **PCNN+ATT** (Lin et al., 2016) proposes a selective attention to alleviate wrong labeling.
- **PCNN+HATT*** (Han et al., 2018) employs hierarchical attention to exploit the relations.
- **PCNN+BAG-ATT** (Ye and Ling, 2019) proposes intra-bag and inter-bag attentions to handle wrongly-labeled sentences at sentence level and bag level respectively.
- **PCNN+KATT*** (Zhang et al., 2019) integrates externally pre-trained graph embeddings with relation hierarchies for long-tail relations. Note, standard AUC and P@N values are not available in the paper while only Hits@K is defined and reported for long-tail settings.
- **SeG** (Li et al., 2020) focuses on one-sentence bags and proposes selective gate mechanism.

3.1 Evaluation on Benchmark

As shown in Table 1 and Figure 3 (left), we compare our CoRA with previous competitive approaches on the distantly supervised relation extraction benchmark in terms of top-n precision, AUC and PR curve. Specifically, CoRA significantly outperforms selective attention baseline, i.e., PCNN+ATT. It also surpasses selective gate framework that shows inferior performance on long-tail relations as in Figure 1 of §1. In addition, compared to PCNN+HATT also utilizing relation hierarchies, CoRA is able to achieve much better results in both P@N and AUC.

3.2 Ablation Study

To further evaluate the effectiveness of each module in the proposed framework, we conduct an extensive ablation study in the bottom of Table 1 and Figure 3 (middle). Since the performance drop is consistent in P@N and AUC, we mainly use AUC as the metric to perform following study. Compared to CoRA, base model without relation collaborating features only shows marginal precision drop when recall > 0.3 in PR-curve, but a significant drop on long-tail relations (detailed in the next section). Also, as an alternative to selective attention, our base model outperforms PCNN+ATT by a large margin. Then, removing simple entity embeddings in §2.2 leads to remarkable degeneration, verifying its importance. It is also rational to compare PCNN+ATT with ‘‘Base w/o Ent Emb’’ (+0.06 AUC) to demonstrate our relation-augmented framework is indeed better than selective attention. Then, removing ‘‘Sent2rel Attention’’, ‘‘Attention-pooling’’ and ‘‘Aux Obj’’ reduce the AUC by 0.10, 0.01 and 0.12 respectively.




Figure 3: **Left:** Precision-recall (PR) curves on NYT for model comparison. **Middle:** PR curves for ablation study. **Right:** Probability (normal) distribution of maximum attention score, $\max(\alpha^{(l)})$, in sent2rel attention, where attention accuracy is whether the the max score $\max(\alpha^{(l)})$ corresponds to $r^{(l)}$.

# Training Instance	<100			<200		
Hits@K (Macro)	10	15	20	10	15	20
PCNN+ATT (Lin et al., 2016)	<5.0	7.4	40.7	17.2	24.2	51.5
PCNN+HATT (Han et al., 2018)	29.6	51.9	61.1	41.4	60.6	68.2
PCNN+KATT (Zhang et al., 2019)	35.3	62.4	65.1	43.2	61.3	69.2
CoRA	66.6	72.0	87.0	72.7	77.3	89.4
Base	33.3	44.4	66.6	45.5	54.5	72.7
Base w/o Aux Obj	18.5	44.4	61.1	33.3	54.5	68.1
Base w/o Sent2rel Attention	5.0	33.3	61.1	22.7	45.5	68.1

Table 2: Hits@K (Macro) on the relations whose number of training instance < 100/200. “Hits@K” denotes whether a test sentence bag whose gold relation label $r^{(0)}$ falls into top- K relations ranked by their prediction confidences. “Macro” denotes macro average is applied regarding relation labels.

3.3 Evaluation on Long-Tail Relations

To prove the capability of CoRA in handling long-tail relations, we conduct an evaluation solely on long-tail relations. Our evaluation setting is identical to (Han et al., 2018; Zhang et al., 2019), where Hits@K (Macro) is used to represent statistical performance on long-tail relations. As shown in Table 2, we compare CoRA with competitors and our base models. It is observed that, CoRA improves the performance on long-tail relations by a large margin and delivers a new state-of-the-art results. Compared to previous works (PCNN+HATT/+KATT) that also leverage the relation hierarchies, our relation-augmented attention (Base) without any hierarchy even gets competitive results, not to mention pre-trained graph embeddings used in PCNN+KATT. Further comparing our base model with selective attention (PCNN+ATT), the huge performance gap demonstrates the advantages of our framework in handling both wrong labeling and long-tail relations. Finally, as shown in the table’s last row, removing the proposed sent2rel attention leads to significant decrease, which emphasizes its importance for long-tail relations.

3.4 Analysis and Case Study

Distributions of Sent2rel Attention Scores. Sent2rel attention used to incorporate multi-granular relation embeddings is an essential module in CoRA, so its normalized attention scores (i.e., attention probabilities) derived from Eq.(6) are critical to measure the knowledge transfer crossing relations. We

Example Sentence 1: <i>Muhammad.yunus</i> , who won the nobel peace prize, last year, demonstrated with <i>grameen.bank</i> , the power of microfinancing.					
Top-3 of attention score $\alpha^{(2)}$		Top-3 of attention score $\alpha^{(1)}$		Top-3 of attention score $\alpha^{(0)}$	
<i>/business:</i>	0.422	NA	0.383	NA	0.387
NA:	0.384	<i>/business/company:</i>	0.272	<i>/business/company/founders:</i>	0.197
<i>/location:</i>	0.037	<i>/business/person:</i>	0.063	<i>/business/person/company:</i>	0.063

Example Sentence 2: On sunday, though, there was a significant shift of the tectonic plates of bangladeshi politics, as <i>muhammad.yunus</i> , the founder of a microfinance empire, known as the <i>grameen.bank</i> and the winner of the 2006 nobel peace prize, announced that he would start a new party and step into the electoral fray.					
Top-3 of attention score $\alpha^{(2)}$		Top-3 of attention score $\alpha^{(1)}$		Top-3 of attention score $\alpha^{(0)}$	
<i>/business:</i>	0.755	<i>/business/company:</i>	0.679	<i>/business/company/founders:</i>	0.652
NA:	0.103	NA:	0.089	NA:	0.069
<i>/people:</i>	0.031	<i>/business/person:</i>	0.059	<i>/business/person/company:</i>	0.057

Table 3: Two example sentences with top-3 sent2rel attention scores at all relation levels. Both sentences express the same long-tail relation “*/business/company/founders*”.

show a probability distribution of maximum attention score in Figure 3 (right). Obviously, a high-level sent2rel attention tends to produce larger maximum attention score and more accurate attention target. It is easily inferred that, 1) accurate attention at high-level promotes the knowledge transfer through the relation hierarchies, and 2) attention probability distribution is more smooth at low-level to further boost embedding sharing crossing relations. To dig this out, in Table 3, we conduct a case study by showing top attention scores at all three relation levels. It is observed that attention scores and the corresponding relations are intuitively consistent with the analyses above. One exception is that *NA* class appears to be assigned with high attention score at low-level sent2rel attention, which indirectly explains 1) our base model w/o collaborating relation features only delivers inferior performance and 2) sent2rel attention for low-level relations are inaccurate.

Performance based solely on Sent2rel Module. Multi-granular relation labels are used as supervision signals for sent2rel attention modules, and the accuracy of each module is greater than 90% as in Figure 3. Therefore, it is interesting to check if the attention scores can be directly used to predict relations at bag level. We present two settings: 1) only using attention scores on fine-grained relations, i.e., $\alpha^{(0)}$, and 2) using products of attention scores at all three levels to make the best of relation hierarchies. As a result, setting 1 and 2 deliver AUC of 0.41 and 0.43 respectively, which surprisingly outperform several previous works in Table 1.

Error Analysis. To investigate the possible reasons for misclassification, we manually check several randomly-sampled error examples from the test set and find the following factors can cause wrong predictions. 1) Most of error cases demonstrate the proposed model still struggles in handling wrong labeling problem, possibly because limited expressive power of text representation is incompetent at handling noisy, imbalance data. 2) The sent2rel attention could be invalid when sibling relations have totally distinct meanings, and posts negative effects on relation extraction. For example, */people/person/children* and */people/person/profession* refer to opposite meanings. 3) Since a sentence embedding is augmented by multiple semantically-related relation embeddings, relation ambiguity problem deteriorates to post errors. For example, it is hard to distinguish */people/deceased_person/place_of_death* and */people/deceased_person/place_of_burial*.

4 Related Work

Relation Extraction. Supervised relation extraction models (Zelenko et al., 2003; GuoDong et al., 2005) require large amounts of annotated data, which is time consuming and labor intensive. To obtain a large amount of labeled data, Mintz et al. (2009) propose distant supervision method to automatically annotate data. However, it inevitably leads to the wrong labeling problem due to the strong assumption. To reduce the effect of wrong labeling problem, multi-instance learning paradigm (Riedel et al., 2010; Hoffmann et al., 2011) is proposed. To introduce the merits of deep learning into relation extraction, Zeng et

al. (2014; 2015) specifically design the position embedding and piecewise convolutional neural network to better extract the features of each sentence. To further alleviate the effect of wrong labeling problem, Lin et al. (2016) propose the selective attention framework under multi-instance learning paradigm. Recently, many works (Liu et al., 2016; Du et al., 2018; Li et al., 2020) are built upon the selective attention (Lin et al., 2016) framework to handle wrong labeling problem in distant supervision relation extraction. For example, Ye and Ling (2019) propose bag-level selective attention to share training information among the bags with the same label. Hu et al. (2019) propose a multi-layer attention-based model with joint label embedding. Li et al. (2020) propose to replace the attention with a gate mechanism especially for one-sentence bags.

Hierarchical Relation Extraction. More related to our work, to alleviate long-tail problem posted by distant supervision, it is natural to utilize relation hierarchies for knowledge transfer crossing relations. There are two existing works falling into this paradigm. Besides using the embedding of fine-grained relation as a query of selective attention, Han et al. (2018) also use embeddings of coarse-grained relations as extra queries to perform a hierarchical attention. Zhang et al. (2019) enhance embeddings of multi-granular relations by merging the embeddings from pre-trained graph model and GCN to alleviate long-tail problem.

5 Conclusion

In this paper, we propose a novel multi-instance learning framework, relation-augmented attention, as our base model for distantly supervised relation extraction. It operates at bag level to minimize the effect of the wrong labeling and leverages a sent2rel attention to alleviate the long-tail problem. By fully exploiting hierarchical knowledge of relations, we extend the base model to CoRA for boosting the performance on long-tail relations. The experiments conducted on the NYT dataset show that CoRA delivers new state-of-the-art performance in terms of both standard metrics (i.e., top-n precision, PR curve, AUC) and long-tail metrics (i.e., Hits@K). And extensive analyses provide comprehensive insights into the proposed model, and verify its capability in handling both wrong labeling and long-tail problems.

Acknowledgement

This research was funded by the Australian Government through the Australian Research Council (ARC) under the grant of LP180100654. The authors would like to appreciate anonymous reviewers for their insightful and constructive feedback.

References

- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.
- Samy Bengio, Jason Weston, and David Grangier. 2010. Label embedding trees for large multi-class tasks. In John D. Lafferty, Christopher K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and Aron Culotta, editors, *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada*, pages 163–171. Curran Associates, Inc.
- Kurt D. Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In Jason Tsong-Li Wang, editor, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, pages 1247–1250. ACM.
- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2787–2795.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon,

- and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3837–3845.
- Jinhua Du, Jingguang Han, Andy Way, and Dadong Wan. 2018. Multi-level structured self-attentions for distantly supervised relation extraction. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2216–2225. Association for Computational Linguistics.
- Zhou GuoDong, Su Jian, Zhang Jie, and Zhang Min. 2005. Exploring various knowledge in relation extraction. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 427–434. Association for Computational Linguistics.
- Xu Han, Pengfei Yu, Zhiyuan Liu, Maosong Sun, and Peng Li. 2018. Hierarchical relation extraction with coarse-to-fine grained attention. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2236–2245. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 541–550. Association for Computational Linguistics.
- Linmei Hu, Luhao Zhang, Chuan Shi, Liqiang Nie, Weili Guan, and Cheng Yang. 2019. Improving distantly-supervised relation extraction with joint label embedding. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3819–3827. Association for Computational Linguistics.
- Guoliang Ji, Kang Liu, Shizhu He, and Jun Zhao. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3060–3066. AAAI Press.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In Alessandro Moschitti, Bo Pang, and Walter Daelemans, editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751. ACL.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Yang Li, Guodong Long, Tao Shen, Tianyi Zhou, Lina Yao, Huan Huo, and Jing Jiang. 2020. Self-attention enhanced selective gate with entity-aware embedding for distantly supervised relation extraction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8269–8276. AAAI Press.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Zhouhan Lin, Minwei Feng, Cícero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Yang Liu, Chengjie Sun, Lei Lin, and Xiaolong Wang. 2016. Learning natural language inference using bidirectional LSTM model and inner-attention. *CoRR*, abs/1605.09090.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In Keh-Yih Su, Jian Su, and Janyce Wiebe, editors, *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore*, pages 1003–1011. The Association for Computer Linguistics.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In José L. Balcázar, Francesco Bonchi, Aristides Gionis, and Michèle Sebag, editors, *Machine Learning and Knowledge Discovery in Databases, European Conference, ECML PKDD 2010, Barcelona, Spain, September 20-24, 2010, Proceedings, Part III*, volume 6323 of *Lecture Notes in Computer Science*, pages 148–163. Springer.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Shirui Pan, and Chengqi Zhang. 2018a. Disan: Directional self-attention network for rnn/cnn-free language understanding. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5446–5455. AAAI Press.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. 2018b. Bi-directional block self-attention for fast and memory-efficient sequence modeling. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Zhi-Xiu Ye and Zhen-Hua Ling. 2019. Distant supervision relation extraction with intra-bag and inter-bag attentions. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2810–2819. Association for Computational Linguistics.
- Matthew D. Zeiler. 2012. ADADELTA: an adaptive learning rate method. *CoRR*, abs/1212.5701.
- Dmitry Zelenko, Chinatsu Aone, and Anthony Richardella. 2003. Kernel methods for relation extraction. *Journal of machine learning research*, 3(Feb):1083–1106.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In Jan Hajic and Junichi Tsujii, editors, *COLING 2014, 25th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, August 23-29, 2014, Dublin, Ireland*, pages 2335–2344. ACL.
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In Lluís Màrquez, Chris Callison-Burch, Jian Su, Daniele Pighin, and Yuval Marton, editors, *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1753–1762. The Association for Computational Linguistics.
- Ningyu Zhang, Shumin Deng, Zhanlin Sun, Guanying Wang, Xi Chen, Wei Zhang, and Huajun Chen. 2019. Long-tail relation extraction via knowledge graph embeddings and graph convolution networks. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 3016–3025. Association for Computational Linguistics.