

Leveraging 2-hop Distant Supervision from Table Entity Pairs for Relation Extraction

Xiang Deng

The Ohio State University
deng.595@buckeyemail.osu.edu

Huan Sun

The Ohio State University
sun.397@osu.edu

Abstract

Distant supervision (DS) has been widely used to automatically construct (noisy) labeled data for relation extraction (RE). Given two entities, distant supervision exploits sentences that directly mention them for predicting their semantic relation. We refer to this strategy as 1-hop DS, which unfortunately may not work well for long-tail entities with few supporting sentences. In this paper, we introduce a new strategy named 2-hop DS to enhance distantly supervised RE, based on the observation that there exist a large number of relational tables on the Web which contain entity pairs that share common relations. We refer to such entity pairs as *anchors* for each other, and collect all sentences that mention the anchor entity pairs of a given target entity pair to help relation prediction. We develop a new neural RE method REDS2 in the multi-instance learning paradigm, which adopts a hierarchical model structure to fuse information respectively from 1-hop DS and 2-hop DS. Extensive experimental results on a benchmark dataset show that REDS2 can consistently outperform various baselines across different settings by a substantial margin.¹

1 Introduction

Relation extraction (RE) aims to extract semantic relations between two entities from unstructured text and is an important task in natural language processing (NLP). Formally, given an entity pair (e_1, e_2) from a knowledge base (KB) and a sentence (instance) that mentions them, RE tries to predict if a relation r from a predefined relation set exists between e_1 and e_2 . A special relation *NA* is used if none of the predefined relations holds.

Given that it is costly to construct large-scale labeled instances for RE, distant supervision (DS)

¹Our source code and datasets are at <https://github.com/sunlab-osu/REDS2>.

Mark Aguirre - Chicago place_lived	mark_aguirre and isiah thomas , the team president , flew to chicago to attend the funeral of ray meyer.	
Mr. Basketball USA		
Mark Aguirre	Westinghouse	Chicago, IL
Rasheed Wallace	Simon Gratz	Philadelphia, PA
LeBron James	St. Vincent – St. Mary	Akron, OH
...
LeBron James - Akron	... including the akron native lebron_james , giving this quintessential college town ...	
	... another lebron_james , the high school phenomenon from akron , ohio , who ...	
Rasheed Wallace - Philly	... scored 13 points and philadelphia native , rasheed_wallace , added 10 ...	
...	...	

Figure 1: Illustration of 2-hop distant supervision. The top panel shows a target entity pair, one sentence that mentions it, and the relation under study which cannot be inferred from the sentence. The middle gives part of a table from Wikipedia page “Mr. Basketball USA”, where we can extract *anchors* for the target entity pair. The bottom shows some sentences that are associated with the anchors, which more clearly indicate the underinvestigated relation and can be utilized to extract relations between the target entity pair.

has been a popular strategy to automatically construct (noisy) training data. It assumes that if two entities hold a relation in a KB, all sentences mentioning them express the same relation. Noticing that the DS assumption does not always hold and has the wrong labeling problem, many efforts including (Riedel et al., 2010; Hoffmann et al., 2011; Surdeanu et al., 2012) have adopted the multi-instance learning paradigm to tackle the challenge, and more recently, neural models with attention mechanism have been proposed to de-emphasize the noisy instances (Lin et al., 2016; Ji et al., 2017; Han et al., 2018). Such models tend to work well when there are a large number of sentences talking about the target entity pair (Lin et al., 2016).

However, we observe that there can be a large portion of entity pairs that have very few supporting sentences (e.g., nearly 75% of entity pairs in the Riedel et al. (2010) dataset only have one single sentence mentioning them), which makes distantly supervised RE even more challenging.

The conventional distant supervision strategy only exploits instances that directly mention a target entity pair, and because of this, we refer to it as *1-hop* distant supervision. On the other hand, there are a large number of Web tables that contain relational facts about entities (Cafarella et al., 2008; Venetis et al., 2011; Wang et al., 2012). Owing to the semi-structured nature of tables, we can extract from them sets of entity pairs that share common relations, and sentences mentioning these entity pairs often have similar semantic meanings. Under this observation, we introduce a new strategy named *2-hop* distant supervision: We define entity pairs that potentially have the same relation with a given target entity pair as *anchors*, which can be found through Web tables, and aim to fully exploit the sentences that mention those anchor entity pairs to augment RE for the target entity pair. Figure 1 illustrates the 2-hop DS strategy.

The intuition behind 2-hop DS is if the target entity pair holds a certain relation, one of its anchors is likely to have that relation too and at least one sentence mentioning the anchors should express the relation. Despite being noisy, the 2-hop DS can provide extra, informative supporting sentences for the target entity pair. One straightforward approach is to merge the two bags of sentences respectively derived from 1-hop and 2-hop DS as one single set and apply existing multi-instance learning models. However, the 2-hop DS strategy also has the wrong labeling problem that already exists in 1-hop DS. Simply mixing the two sets of sentences together may mislead the prediction, especially when there is a great disparity in their size. In this paper, we propose REDS², a new neural relation extraction method in the multi-instance learning paradigm, and design a hierarchical model structure to fuse information from 1-hop and 2-hop DS. We evaluate REDS² on a widely used benchmark dataset and show that it consistently outperforms various baseline models by a large margin.

We summarize our contributions as three-fold:

- We introduce 2-hop distant supervision as an

²stands for relation extraction with 2-hop DS.

extension to the conventional distant supervision, and leverage entity pairs in Web tables as anchors to find additional supporting sentences to further improve RE.

- We propose REDS², a new neural relation extraction method based on 2-hop DS and has achieved new state-of-the-art performance in the benchmark dataset (Riedel et al., 2010).
- We release both our source code and an augmented benchmark dataset that has entity pairs aligned with those in Web tables, to facilitate future work.

2 Related Work

Distant Supervision. One main drawback of traditional supervised relation extraction models (Zelenko et al., 2003; Mooney and Bunescu, 2006) is they require adequate amounts of annotated training data, which is time consuming and labor intensive. To address this issue, Mintz et al. (2009) proposes distant supervision (DS) to automatically label data by aligning plain text with Freebase. However, DS inevitably accompanies with the wrong labeling problem. To alleviate the noise brought by DS, Riedel et al. (2010) and Hoffmann et al. (2011) introduce multi-instance learning mechanism, which is originally used to combat the problem of ambiguously-labeled training data when predicting the activity of different drugs (Dieterich et al., 1997).

Neural Relation Extraction. Early stage relation extraction (RE) methods use features extracted by NLP tools and strongly rely on the quality of features. Due to the recent success of neural models in different NLP tasks, many researchers have investigated the possibility of using neural networks to build end-to-end relation extraction models. Zeng et al. (2014) uses convolutional neural network (CNN) to encode sentences, which is further improved through piecewise-pooling (Zeng et al., 2015). Adel and Schütze (2017) and Gupta et al. (2016) use neural networks for joint entity and relation extraction. More advanced network architectures like Tree-LSTM (Miwa and Bansal, 2016) and Graph Convolution Network (Vashishth et al., 2018) are also adopted to learn better representations by using syntactic features like dependency trees. Most recent models also incorporate neural attention technology (Lin et al., 2016) as an

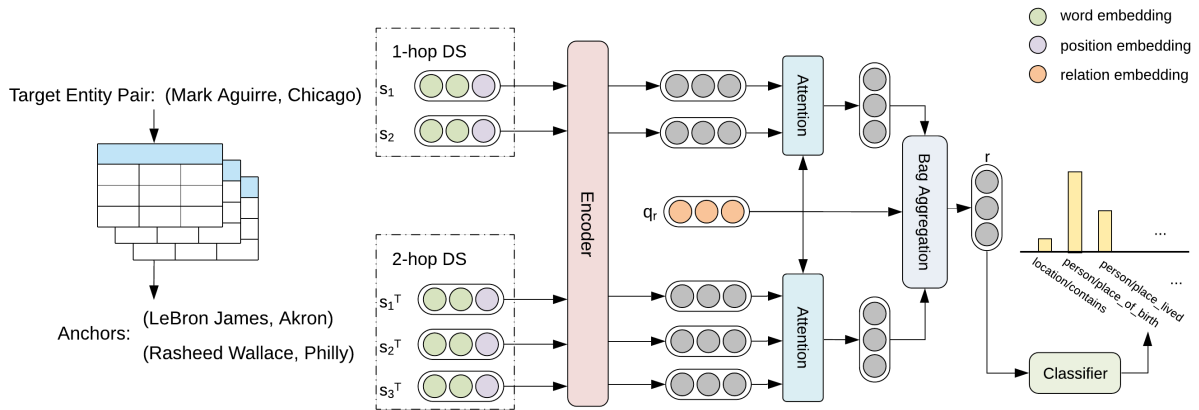


Figure 2: Overview of our method REDS2. REDS2 first obtains *anchors* of the target entity pair and constructs a 2-hop DS bag. Sentences in the 1-hop and 2-hop DS bag are individually encoded with a PCNN sentence encoder (Zeng et al., 2015). We then use selective attention and bag aggregation to get the final representation, based on which a classifier predicts scores for each candidate relation.

improvement to at-least-one multi-instance learning (Zeng et al., 2015). Han et al. (2018) further develops a hierarchical attention scheme to utilize the relation correlations and help predictions for long-tail relations.

Web Table Understanding. Aside from plain texts, there are large amounts of factual knowledge in the Web expressed in hundreds of millions of tables and other structured lists (Cafarella et al., 2008; Venetis et al., 2011), which have not been fully explored yet. Table understanding tries to match tables to KB and parse the schemas. Existing methods for table understanding mainly fall into two categories. One is based on local evidence (Venetis et al., 2011; Muñoz et al., 2014; Ritze et al., 2015). Given one table, the main idea is to first link cells to entities in KB. We can then use existing relations between linked entities to infer relations between columns and extract new facts by generalizing to all rows. However, this method requires a high overlap between table and KB, which is hampered by KB incompleteness. The other approach tries to leverage features extracted from the table header and column names (Ritze and Bizer, 2017; Cannaviccio et al., 2018). Unfortunately, a large portion of Web tables miss such metadata or contain limited information, and the second approach will fail in such cases. Although the focus of this paper is the RE task, we believe the idea of connecting Web tables and plain texts using DS can potentially benefit table understanding as well.

3 Methodology

Given a set of sentences $S = \{s_1, s_2, \dots\}$ and a target entity pair (h, t) , we will leverage the directly associated sentence bag $S_{h,t} \subseteq S$ by 1-hop distant supervision (1-hop DS bag), and the table expanded sentence bag $S_{h,t}^T \subseteq S$ by 2-hop distant supervision (2-hop DS bag), for relation extraction. $S_{h,t}$ contains all instances mentioning both h and t , while $S_{h,t}^T$ is obtained indirectly through the anchors of (h, t) found in Web tables. Following previous work (Riedel et al., 2010; Hoffmann et al., 2011), we adopt the multi-instance learning paradigm to measure the probability of (h, t) having relation r .

Figure 2 gives an overview of our framework with three major components:

- **Table-aided Instance Expansion:** Given a target entity pair (h, t) , we find its *anchor* entity pairs $\{(h_1, t_1), (h_2, t_2), \dots\}$ through Web tables. We define an anchor entity pair as two entities co-occurring with (h, t) in some table columns at least once. $S_{h,t}^T = S_{h_1,t_1} \cup S_{h_2,t_2} \cup \dots$ is then exploited to augment the directly associated bag $S_{h,t}$.
- **Sentence Encoding:** For each sentence s in bag $S_{h,t}$ or $S_{h,t}^T$, a sentence encoder is used to obtain its semantic representation s .
- **Hierarchical Bag Aggregation:** Once the embedding of each sentence is learned, we first use a sentence-level attention mechanism to get bag representation \mathbf{h} and \mathbf{h}^T , and

then aggregate them for final relation prediction.

3.1 Table-aided Instance Expansion

Now we introduce how to construct the table expanded sentence bag $S_{h,t}^T$ for a given target entity pair (h, t) by 2-hop distant supervision.

3.1.1 Web Tables

Web tables have been found to contain rich facts of entities and relations. It is estimated that out of a total of 14.1 billion tables on the Web, 154 million tables contain relational data (Cafarella et al., 2008) and Wikipedia alone is the source of nearly 1.6 million relational tables (Bhagavatula et al., 2015). Columns of a Wikipedia table can be classified into one of the following data types: ‘empty’, ‘named entity’, ‘number’, ‘date expression’, ‘long text’ and ‘other’ (Zhang, 2017). Here we only focus on named entity columns (NE-columns) and the Wikipedia page title, which can be easily linked to KB entities. These entities can be further categorized as:

A topic entity e^t that the table is centered around. We refer to the Wikipedia article where the table is found and take the entity it describes as e^t .

Subject entities $E^s = \{e_1^s, e_2^s, \dots\}$ that can act as primary keys of the table. Following previous work on Web table analysis (Venetis et al., 2011), we select the leftmost NE-column as subject column and its entities as E^s .

Body entities $E = \{e_{1,1}, e_{1,2}, \dots\}$ that compose the rest of the table. All entities in non-subject NE-columns are considered as E .

3.1.2 2-hop Distant Supervision

In the conventional distant supervision setting, each entity pair (h, t) is associated with a bag of sentences $S_{h,t}$ that directly mention h and t . The intuition behind 2-hop distant supervision is, if (h_i, t_i) and (h_j, t_j) potentially hold the same relation, we can treat them as **anchor entity pairs** for each other, and then use the 1-hop DS bag S_{h_j, t_j} to help with the prediction for (h_i, t_i) and vice versa. In this paper, we extract anchor entity pairs with the help of Web tables.

We notice that owing to the semi-structured nature of tables, (1) subject entities can usually be connected with the topic entity by the same relation. (2) Non-subject columns of a table usually have binary relationships to or are properties of the

subject column. Body entities in the same column share common relations with their corresponding subject entities. For example, in Figure 1, the topic entity is “Mr. Basketball USA”; column 1 is the subject column and contains a list of winners of “Mr. Basketball USA”; column 2 and column 3 are high school and city of the subject entity.

Formally, we consider two entity pairs (h_i, t_i) and (h_j, t_j) as anchored if there exists a Web table such that either criterion below is met:

- $h_i = h_j = e^t$ and $t_i, t_j \in E^s$.
- $h_i \in E^s$ or $t_i \in E^s$, (h_i, h_j) is in the same column (and so is (t_i, t_j)), and (h_i, t_i) is in the same row (and so is (h_j, t_j))

The 2-hop DS bag $S_{h,t}^T$ is then constructed as the union of S_{h_i, t_i} ’s, where (h_i, t_i) is an anchor entity pair of (h, t) .

3.2 Sentence Encoding

Given a sentence s consisting of n words $s = \{w_1, w_2, \dots, w_n\}$, we use a neural network with an embedding layer and an encoding layer to obtain its low-dimensional vector representation.

3.2.1 Embedding Layer

Each token is first fed into an embedding layer to embed both semantic and positional information.

Word Embedding maps words to vectors of real numbers which preserve syntactic and semantic information of words. Here we get a vector representation $\mathbf{w}_i \in \mathbb{R}^{k_w}$ for each word from a pre-trained word embedding matrix.

Position Embedding was proposed by Zeng et al. (2014). Position embedding is used to embed the positional information of each word relative to the head and tail mention. A position embedding matrix is learned in training to compute position representation $\mathbf{p}_i \in \mathbb{R}^{k_p \times 2}$.

Finally, we concatenate the word representation \mathbf{w}_i and position representation \mathbf{p}_i to build the input representation $\mathbf{x}_i \in \mathbb{R}^{k_i}$ (where $k_i = k_w + k_p \times 2$) for each word w_i .

3.2.2 Encoding Layer

A sequence of input representations $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots\}$ with a variable length is then fed through the encoding layer and converted to a fixed-sized sentence representation $\mathbf{s} \in \mathbb{R}^{k_h}$. There are many existing neural architectures that can serve as the encoding layer, such as CNN

(Zeng et al., 2014), PCNN (Zeng et al., 2015) and LSTM-RNN (Miwa and Bansal, 2016). We simply adopt PCNN here, which has been shown very powerful and efficient by a number of previous RE works.

PCNN is an extension to CNN, which first slides a convolution kernel with a window size m over the input sequence to get the hidden vectors:

$$\mathbf{h}_i = \text{CNN}(\mathbf{x}_{i-\frac{m-1}{2}:i+\frac{m-1}{2}}), \quad (1)$$

A piecewise max-pooling is then applied over the hidden vectors:

$$\begin{aligned} [\mathbf{s}^{(1)}]_j &= \max_{1 \leq i \leq i_1} \{[\mathbf{h}_i]_j\}, \\ [\mathbf{s}^{(2)}]_j &= \max_{i_i+1 \leq i \leq i_2} \{[\mathbf{h}_i]_j\}, \\ [\mathbf{s}^{(3)}]_j &= \max_{i_2+1 \leq i \leq n} \{[\mathbf{h}_i]_j\}, \end{aligned} \quad (2)$$

where i_1 and i_2 are head and tail positions. The final sentence representation \mathbf{s} is composed by concatenating these three pooling results $\mathbf{s} = [\mathbf{s}^{(1)}; \mathbf{s}^{(2)}; \mathbf{s}^{(3)}]$.

3.3 Hierarchical Bag Aggregation

After we get sentence representations $\{\mathbf{s}_1, \mathbf{s}_2, \dots\}$ and $\{\mathbf{s}_1^T, \mathbf{s}_2^T, \dots\}$ for S and S^T , to fuse key information from these two bags, we adopt a hierarchical aggregation design to obtain the final representation \mathbf{r} for prediction. We first get bag representation \mathbf{h} and \mathbf{h}^T using a sentence-level selective attention, and then employ a bag-level aggregation to compute \mathbf{r} .

3.3.1 Sentence-level Selective Attention

Since the wrong labeling problem inevitably exists in both 1-hop and 2-hop distant supervision, here we use selective attention to assign different weights to different sentences given relation r and de-emphasize the noisy sentences. The attention is calculated as follows:

$$\begin{aligned} e_i &= \mathbf{q}_r^\top \mathbf{s}_i, \\ \alpha_i &= \frac{\exp(e_i)}{\sum_{j=1}^n \exp(e_j)}, \\ \mathbf{h} &= \sum_{i=1}^n \alpha_i \mathbf{s}_i, \end{aligned} \quad (3)$$

where \mathbf{q}_r is a query vector assigned to relation r . \mathbf{h} and \mathbf{h}^T are computed respectively for the two bags S and S^T .

3.3.2 Bag-level Aggregation

Since 2-hop DS bag S^T is collected indirectly through anchor entity pairs in Web tables, despite that it brings abundant information, it also contains a massive amount of noise. Thus treating S^T equally as S may mislead the prediction, especially when their sizes are extremely imbalanced.

To automatically decide how to balance between S and S^T , we utilize information from \mathbf{h} , \mathbf{h}^T and \mathbf{q}_r to predict a weight β :

$$\beta = \sigma(\mathbf{W}[\mathbf{h}; \mathbf{h}^T; \mathbf{q}_r] + \mathbf{b}), \quad (4)$$

where vector \mathbf{W} and scalar \mathbf{b} are learnable variables and σ is the sigmoid function. Next, β is used as a weight to fuse information from 1-hop DS and 2-hop DS, determined by S and S^T of the current target entity pair and relation r . We then obtain the final representation \mathbf{r} as:

$$\mathbf{r} = \beta \mathbf{h} + (1 - \beta) \mathbf{h}^T, \quad (5)$$

Finally, we define the conditional probability $P(r|S, S^T, \theta)$ as follows,

$$P(r|S, S^T, \theta) = \frac{\exp(\mathbf{o}_r)}{\sum_{k=1}^{n_r} \exp(\mathbf{o}_k)} \quad (6)$$

where \mathbf{o} is the score vector for current target entity pair having each relation,

$$\mathbf{o} = \mathbf{M}\mathbf{r} + \mathbf{d}, \quad (7)$$

here \mathbf{M} is the representation matrix of relations, which shares weights with \mathbf{q}_r 's. \mathbf{d} is a learnable bias term.

3.4 Optimization

We adopt the cross-entropy loss as the training objective function. Given a set of target entity pairs with relations $\pi = \{(h_1, t_1, r_1), (h_2, t_2, r_2), \dots\}$, we define the loss function as follows:

$$J(\theta) = -\frac{1}{|\pi|} \sum_{i=1}^{|\pi|} \log P(r_i | S_{h_i, t_i}, S_{h_i, t_i}^T, \theta). \quad (8)$$

All models are trained with stochastic gradient descent (SGD) to minimize the objective function. The same sentence encoder is used to encode S and S^T .

	Train		Test	
	Overall	Non-NA	Overall	Non-NA
# Entity Pairs	291699	18144	96678	1761
# Entity Pairs with $ S^T > 0$	17565	6928	4832	824
% Entity Pairs with $ S^T > 0$	6.02	38.18	5.00	46.79
mean $ S $	1.69	5.24	1.62	2.78
mean $ S^T $	147.65	190.61	131.65	217.23

Table 1: Dataset statistics. We show statistics of entity pairs that hold non-NA relations separately from overall, as they are important relational facts to discover. Among non-NA entity pairs, 38.18% in training and 46.79% in testing have nonempty S^T , which respectively have 190.61 and 217.23 sentences on average.

4 Experiments

4.1 Datasets and Evaluation

We evaluate our model on the New York Times (NYT) dataset developed by Riedel et al. (2010), which is widely used in recent works. The dataset has 53 relations including a special relation NA which indicates none of the other 52 relations exists between the head and tail entity.

We use the WikiTable corpus collected by Bhagavatula et al. (2015) as our table source. It originally contains around 1.65M tables extracted from Wikipedia pages. Since the NYT dataset is already linked to Freebase, we perform entity linking on the table cells and the Wikipedia page titles using existing mapping from Wikipedia URL to Freebase MID (Machine Identifier). We then align the table corpus with NYT and construct S^T for entity pairs as detailed in section 3.1. *For both training and testing, we only use entity pairs and sentences in the original NYT training data for table-aided instance expansion.* We set the max size of S^T as 300, and randomly sample 300 sentences if $|S^T| > 300$. Statistics of our final dataset is summarized in Table 1. One can see that 38.18% and 46.79% of relational facts (i.e., entity pairs holding non-NA relations) respectively in the training and testing set can potentially benefit from leveraging 2-hop DS.

Following prior work (Mintz et al., 2009), we use the testing set for held-out evaluation, and evaluate models by comparing the predicted relational facts with those in Freebase. For evaluation, we rank the extracted relational facts based on model confidence and plot precision-recall curves. In addition, we also show the area under the curve (AUC) and precision values at specific recall rates to conduct a more comprehensive comparison.

Window size m	3
Sentence Representation Size k^h	230
Word Dimension k^w	50
Position Dimension k^p	5
Pre-train Learning Rate λ_P	0.005
Fine-tune Learning Rate λ_F	0.002
Dropout Probability p	0.5

Table 2: Parameter settings in REDS2.

4.2 Baselines

We compare REDS2 with the following baselines: **PCNN+ATT** (Lin et al., 2016). This model uses a PCNN encoder combined with selective attention over sentences. Since this is also the base block of our model, we also refer to it as BASE in this paper.

PCNN+HATT (Han et al., 2018). This is another PCNN based relation extraction model, where the authors use hierarchical attention to model the semantic correlations among relations.

RESIDE (Vashishth et al., 2018). It uses Graph Convolutional Networks (GCN) for sentence encoding, and also leverages relevant side information like relation alias and entity type.

Results of PCNN+HATT and RESIDE are directly taken from the code repositories released by the authors. For PCNN+ATT, we report results obtained by our reproduced model, which are close to those shown in (Lin et al., 2016). To simply verify the effectiveness of adding extra supporting sentences from 2-hop DS, we also compare the following vanilla method with PCNN+ATT:

BASE+MERGE. For each target entity pair (h, t) , we simply merge S and S^T as one sentence bag, and apply the trained PCNN+ATT (or, BASE) model.

4.3 Implementation Details

We preprocess the WikiTable corpus with PySpark to build index for anchor entity pairs. On a single machine with two 8-core E5 CPUs and 256 GB memory, this processing takes around 20 minutes.

We use word embeddings from (Lin et al., 2016) for initialization, which are learned by word2vec tool³ on NYT corpus. The vocabulary is composed of words that appear more than 100 times in the corpus and words in an entity mention are concatenated as a single word.

³<https://code.google.com/archive/p/word2vec/>

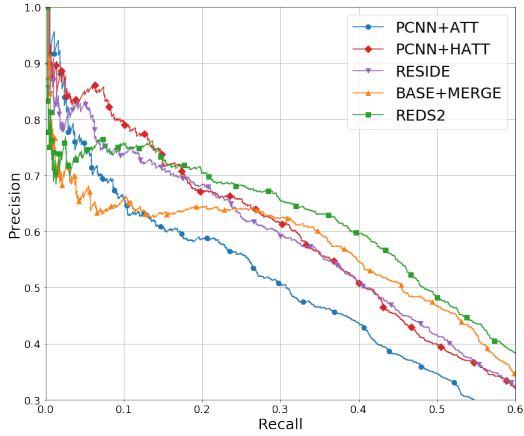


Figure 3: Precision-recall curves for the proposed model and various baselines.

Method	P@0.1	P@0.2	P@0.3	AUC
PCNN+ATT	65.9	58.7	50.4	35.8
PCNN+HATT	79.3	67.2	61.6	42.0
RESIDE	73.6	68.4	59.5	41.5
BASE+MERGE	65.9	64.5	62.4	41.2
REDS2	75.9	70.4	65.5	44.7

Table 3: Comparison on Precision@recall and AUC.

To see the effect of 2-hop DS more directly, we set most parameters in REDS2 following Lin et al. (2016). Since the original NYT dataset only contains training and testing set, we randomly sample 20% training data for development. We first pre-train a PCNN+ATT model with only S and sentence-level selective attention. This BASE model converges in around 100 epochs. We then fine-tune the entire model with S^T and bag-level aggregation added, which can finish within 50 epochs. Some key parameter settings in REDS2 are summarized in Table 2.

In testing phase, inference using 2-hop DS is slower, because the average size of S^T is about 100 times that of S . With single 2080ti GPU, one full pass of testing data takes around 37s using REDS2, compared with 12s using BASE model.

4.4 Results

4.4.1 Overall Evaluation Results

Evaluation results on all target entity pairs in testing set are shown in Figure 3 and Table 3, from which we make the following observations:

(1) Figure 3 shows all models obtain a reasonable precision when recall is smaller than 0.05. With the recall gradually increasing, the performance of models with 2-hop DS drops slower




Figure 4: Precision-recall curves on the subset of test entity pairs whose S^T is not empty, to better show the effect of hierarchical bag aggregation design.

than those existing methods without. From Figure 3, we can see simply merging S^T with S in BASE+MERGE can boost the performance of basic PCNN+ATT model, and even achieves higher precision than state-of-the-art models like PCNN+HATT when recall is greater than 0.3. This demonstrates that models utilizing 2-hop DS are more robust and remain a reasonable precision when including more lower-ranked relational facts which tend to be more challenging to predict because of insufficient evidence.

(2) As shown in both Figure 3 and Table 3, REDS2 achieves the best results among all the models. Even when compared with PCNN+HATT and RESIDE which adopt extra relation hierarchy and side information from KB, our model still enjoys a significant performance gain. This is because our method can take advantage of the rich entity pair correlations in Web tables and leverage the extra information brought by 2-hop DS. We anticipate our REDS2 model can be further improved by using more advanced sentence encoders and extra mechanisms like reinforcement learning (Feng et al., 2018) and adversarial training (Wu et al., 2017), which we leave for future work.

4.4.2 Effect of Hierarchical Bag Aggregation

To further show the effect of our hierarchical bag aggregation design, here we also plot precision-recall curves in Figure 4 on a subset of entity pairs in the test set (i.e., 4832 in total according to Table 1) whose table expanded sentence bag S^T is not empty.

One main challenge of using 2-hop DS is it brings more noise. As shown in Table 1, for en-

Test Mode	SINGLE				MULTIPLE											
					ONE				TWO				ALL			
Metric	P@0.1	P@0.2	P@0.3	AUC	P@0.1	P@0.2	P@0.3	AUC	P@0.1	P@0.2	P@0.3	AUC	P@0.1	P@0.2	P@0.3	AUC
PCNN+ATT	57.3	53.0	41.3	30.0	72.9	66.4	57.5	39.0	75.3	69.7	63.8	45.9	80.5	68.6	64.2	48.6
RESIDE	66.9	60.5	51.8	36.2	79.7	70.4	57.4	42.1	77.7	76.1	67.6	46.5	83.2	80.3	73.2	51.9
PCNN+HATT	70.1	59.1	50.8	34.8	74.5	67.4	57.2	40.1	78.4	68.9	65.9	44.5	86.4	75.5	70.0	49.7
BASE+MERGE	62.2	58.5	55.2	35.6	71.4	74.9	73.2	51.3	70.0	71.8	73.9	52.0	70.7	71.1	73.2	52.0
REDS2	69.1	61.1	57.4	37.5	82.4	81.9	78.1	56.3	81.4	79.6	76.9	57.2	82.4	79.6	76.6	57.6

Table 4: Comparison on Precision@recall and AUC under different testing settings, detailed in Section 4.4.3.

max $ S^T $	P@0.1	P@0.2	P@0.3	AUC
10	57.9	55.8	51.5	36.2
50	69.4	65.7	60.8	42.2
100	70.4	66.7	62.3	43.2
200	72.8	68.4	63.4	44.0
300	75.9	70.4	65.5	44.7

Table 5: Effect of the table expanded sentence bag size $|S^T|$ on Precision@recall and AUC.

tity pair with nonempty S^T , the size of S^T is usually tens of times the size of S . From Figure 4 we can see BASE+MERGE performs much worse compared with PCNN+ATT when recall is smaller than 0.2. This is because 2-hop DS bag tends to be much larger than 1-hop DS bag, and the model has a larger chance to attend to the noisy sentences obtained from 2-hop DS. While ignoring the information in its directly associated sentences. We alleviate this problem by introducing hierarchical structure to first aggregate the two sets separately and then weight and sum them together. The proposed REDS2 model has a comparable precision with PCNN+ATT in the beginning and gradually outperform it.

4.4.3 Effect of Sentence Number

Number of sentences from 1-hop DS. In the originally testing set, there are 79176 entity pairs that are associated with only one sentence, out of which 1149 actually have relations. We hope our model can improve performance on these long-tail entities. Following Lin et al. (2016), we design the following test settings to evaluate the effect of sentence number: the ‘‘SINGLE’’ test setting contains all entity pairs that correspond to only one sentence; the ‘‘MULTIPLE’’ test setting contains the rest of entity pairs that have at least two sentences associated. We further construct the ‘‘ONE’’ testing setting where we randomly select one sentence for each entity pair; the ‘‘TWO’’ setting where we randomly select two sentences for each entity pair and the ‘‘ALL’’ setting where

Relation: country.capital	
1-hop	... the golden gate bridge and the petronas towers in kuala lumpur, malaysia , was experienced ...
2-hop	a friend from cardiff , the capital city of wales , lives for complex ...

Table 6: An example for case study, where the sentence with the highest attention weight is selected respectively from 1-hop and 2-hop sentence bag.

we use all the associated sentences from MULTIPLE. We use all sentences in S^T for each entity pair if it is nonempty. Results are shown in Table 4, from which we can see that REDS2 and BASE+MERGE have 25.0% and 18.7% improvements under AUC compared with PCNN+ATT in the SINGLE setting. Although the performance of all models generally improves as the sentence number increases in MULTIPLE setting, models leveraging 2-hop DS are more stable and have smaller changes. These observations indicate that 2-hop DS is helpful when information obtained by 1-hop DS is insufficient.

Number of sentences from 2-hop DS. We also evaluate how the number of sentences obtained by 2-hop DS will affect the performance of our proposed model. In Table 5, we show the performance of REDS2 with different numbers of sentences sampled from S^T . We observe that: (1) Performance of REDS2 improves as the number of sentences sampled increases. This shows that the selective attention over S^T can effectively take advantage of the extra information from 2-hop DS while filtering out noisy sentences. (2) Even with 50 randomly sampled sentences, our model REDS2 still has a higher AUC than all baselines in Table 3. This indicates information obtained by 2-hop DS is redundant, even a small portion can be beneficial to relation extraction. How to sample a representative set effectively is worth further exploring in future work.

4.5 RE for Entity Pairs with Empty 1-hop Sentence Bag

We observe that there are large amounts of entity pairs in the table corpus that have no associated sentences but have anchor entity pairs mentioned in the text corpus. By leveraging 2-hop distant supervision, we can do relation extraction for this set of entity pairs.

We extract a total number of 251917 entity pairs from the WikiTable dataset which do not exist in the NYT dataset but have at least one anchor entity pair that appear in the original NYT training data. We randomly sample 10000 examples and evaluate our trained model on them. Surprisingly, the relation extraction result is even better than the result on the NYT test data in Table 3, with an overall AUC of 54.7 and a P@0.3 of 71.1. This can be explained partly by two observations: (1) The table corpus generates higher-quality entity pairs, 18% of extracted entity pairs have non-NA relations, compared with only 1.8% in NYT test data. (2) The newly extracted entity pairs have 14 useful anchor entity pairs and 175 2-hop DS sentences on average, which give ample information for prediction. This study shows that for two entities that have no directly associated sentences, it is possible to utilize the 2-hop DS to predict their relations accurately.

4.6 Case Study and Error Analysis

In addition to the motivating example from the training set shown in Figure 1, we also demonstrate how 2-hop DS helped relation extraction using an example from the testing set in Table 6. As we can see, the sentence with the highest attention weight in 1-hop DS bag does not express the desired relation between the target entity pair whereas that in 2-hop DS bag clearly indicates the `country.capital` relation.

We also conduct an error analysis by analyzing examples where REDS2 gives worse predictions than BASE (e.g., assigns a lower score to a correct relation or a higher score to a wrong relation), and 50 examples with most disparity in the two methods' scores are selected. We find that 29 examples have wrong labels caused by KB incompleteness and our model in fact makes the right prediction. 11 examples are due to errors in column processing (e.g., errors in NE/subject column selection and entity linking), 9 are caused by anchor entity pairs with different relations (e.g.,

(Greece, Atlanta) and (Mexico, Xalapa) are in the same table “*National Records in High Jump*” under columns (Nation, Place), but only the latter has relation `location.contains`), and 1 is because of wrong information in the original table.

5 Conclusion and Future Work

This paper introduces 2-hop distant supervision for relation extraction, based on the intuition that entity pairs in relational Web tables often share common relations. Given a target entity pair, we define and find its anchor entity pairs via Web tables and collect all sentences that mention the anchor entity pairs to help relation prediction. We develop a new neural RE method REDS2 in the multi-instance learning paradigm which fuses information from 1-hop DS and 2-hop DS using a hierarchical model structure, and substantially outperforms existing RE methods on a benchmark dataset. Interesting future work includes: (1) Given that information from 2-hop DS is redundant and noisy, we can explore smarter sampling and/or better bag-level aggregation methods to capture the most representative information. (2) Metadata in Web tables like headers and column names also contain rich information, which can be incorporated to further improve RE performance.

6 Acknowledgments

This research was sponsored in part by the Army Research Office under cooperative agreements NSF Grant IIS1815674, W911NF-17-1-0412, Fujitsu gift grant, and Ohio Supercomputer Center [8]. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notice herein.

References

- Heike Adel and Hinrich Schütze. 2017. Global normalization of convolutional neural networks for joint entity and relation classification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1723–1729.
- Chandra Sekhar Bhagavatula, Thanapon Noraset, and Doug Downey. 2015. Tabel: entity linking in web

- tables. In *International Semantic Web Conference*, pages 425–441.
- Michael J Cafarella, Alon Halevy, Daisy Zhe Wang, Eugene Wu, and Yang Zhang. 2008. Webtables: exploring the power of tables on the web. *Proceedings of the VLDB Endowment*, 1(1):538–549.
- Matteo Cannavicchio, Lorenzo Ariemma, Denilson Barbosa, and Paolo Merialdo. 2018. Leveraging wikipedia table schemas for knowledge graph augmentation. In *Proceedings of the 21st International Workshop on the Web and Databases*, page 5. ACM.
- Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. 1997. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71.
- Jun Feng, Minlie Huang, Li Zhao, Yang Yang, and Xiaoyan Zhu. 2018. Reinforcement learning for relation classification from noisy data. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Pankaj Gupta, Hinrich Schütze, and Bernt Andrassy. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2537–2547.
- Xu Han, Pengfei Yu, Zhiyuan Liu, Maosong Sun, and Peng Li. 2018. Hierarchical relation extraction with coarse-to-fine grained attention. In *Proceedings of EMNLP*.
- 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.
- Guoliang Ji, Kang Liu, Shizhu He, and Jun Zhao. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- 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 (Volume 1: Long Papers)*, volume 1, pages 2124–2133.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. Association for Computational Linguistics.
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using lstms on sequences and tree structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1105–1116.
- Raymond J Mooney and Razvan C Bunescu. 2006. Subsequence kernels for relation extraction. In *Advances in neural information processing systems*, pages 171–178.
- Emir Muñoz, Aidan Hogan, and Alessandra Mileo. 2014. Using linked data to mine rdf from wikipedia’s tables. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 533–542. ACM.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163. Springer.
- Dominique Ritze and Christian Bizer. 2017. Matching web tables to dbpedia - A feature utility study. In *Proceedings of the 20th International Conference on Extending Database Technology, EDBT 2017, Venice, Italy, March 21-24, 2017.*, pages 210–221.
- Dominique Ritze, Oliver Lehmberg, and Christian Bizer. 2015. Matching html tables to dbpedia. In *Proceedings of the 5th International Conference on Web Intelligence, Mining and Semantics*, page 10. ACM.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 455–465.
- Shikhar Vashishth, Rishabh Joshi, Sai Suman Prayaga, Chiranjib Bhattacharyya, and Partha Talukdar. 2018. Reside: Improving distantly-supervised neural relation extraction using side information. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1257–1266.
- Petros Venetis, Alon Halevy, Jayant Madhavan, Marius Pasca, Warren Shen, Fei Wu, Gengxin Miao, and Chung Wu. 2011. Recovering semantics of tables on the web. *Proceedings of the VLDB Endowment*, 4(9):528–538.
- Jingjing Wang, Haixun Wang, Zhongyuan Wang, and Kenny Q Zhu. 2012. Understanding tables on the web. In *International Conference on Conceptual Modeling*, pages 141–155. Springer.
- Yi Wu, David Bamman, and Stuart Russell. 2017. Adversarial training for relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1778–1783.

- 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, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1753–1762.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344.
- Ziqi Zhang. 2017. Effective and efficient semantic table interpretation using tableminer+. *Semantic Web*, 8(6):921–957.