# Improving Distantly-Supervised Relation Extraction through BERT-based Label & Instance Embeddings

#### Abstract

Distantly-supervised relation extraction (RE) is an effective method to scale RE to large corpora but suffers from noisy labels. Existing approaches try to alleviate noise through multi-instance learning and by providing additional information, but they treat labels as independent and seldom consider the relationship between labels and entities. Consequently, potentially valuable information goes unused. We propose REDSandT (Relation Extraction with Distant Supervision and Transformers), a novel distantly-supervised transformer-based RE method that manages to capture highly informative instance and label embeddings for RE by exploiting BERT's pre-trained model. We guide REDSandT to focus solely on relational tokens by fine-tuning BERT on a structured input, including the sub-tree connecting an entity pair and the entities' types. Using the extracted informative vectors, we shape label embeddings, which we also use as attention mechanism over instances to further reduce noise. Finally, we represent sentences by concatenating relation and instance embeddings. Experiments in the NYT-10 dataset show that REDSandT captures a broader set of relations with higher confidence, achieving state-of-theart AUC (0.424).

# 1 Introduction

Relation Extraction (RE) aims to detect semantic relationships between entity pairs in natural texts and has proven to be crucial in various natural language processing (NLP) applications, including question answering, and knowledge-base (KB) population.

Most RE methods follow a supervised approach, with the required number of labeled training data rendering the whole process time and laborintensive. To automatically construct datasets for RE, (Mintz et al., 2009) proposed to use distant supervision (DS) from a KB, assuming that if two entities exhibit a relationship in a KB, then all sentences mentioning these entities express this relation. Inevitably, this assumption generates falsepositives and leads distantly-created datasets to contain erroneous labels. To alleviate the *wrong labeling problem*, (Riedel et al., 2010) relaxed this assumption so that it does not hold for all instances and along with (Hoffmann et al., 2011; Surdeanu et al., 2012) proposed multi-instance based learning. Under this setting, classification shifts from instance-level to bag-level, with a bag consisting of all instances that contain a specific entity pair.

Current state-of-the-art RE methods try to reduce the effect of noisy instances by: i) identifying valid instances through multi-instance learning and selective attention (Lin et al., 2016), ii) reducing inner-sentence noise by capturing long-range dependencies using syntactic information from dependency parses (Mintz et al., 2009; He et al., 2018; Liu et al., 2018), specialized models like piecewise CNN (PCNN) and graph CNN (GCNN), or wordlevel attention (He et al., 2018), and iii) enhancing model effectiveness using external knowledge (i.e. KB entity types (Vashishth et al., 2018), entity descriptions (Ji, 2017; Hu et al., 2019), relation phrases (Vashishth et al., 2018)) or transfer knowledge from pre-trained models (Alt et al., 2019).

The study of the above approaches leads us to two observations. First, only (Alt et al., 2019), using a pre-trained transformer-based language model (LM), manages to recognize a broader set of relations. Second, the use of relation embeddings (He et al., 2018; Hu et al., 2019) is helpful for RE. However, these models miss information, as the former achieves lower precision in top predictions compared to standard methods (Lin et al., 2016; Vashishth et al., 2018), and the latter use label embeddings only as an attention mechanism over sentences. Consequently, we believe that there is an opportunity to: a) exploit pre-trained LMs' ability to capture deep language representations (Vaswani et al., 2017) towards better representation of both instances and labels, and b) utilize the highly informative label embeddings to boost RE.

We propose REDSandT (Relation Extraction with Distant Supervision and Transformers), a novel transformer-based RE model for distant supervision. To handle the problem of noisy instances, we guide REDSandT to focus solely on relational tokens by fine-tuning BERT on a structured input, including the sub-tree connecting an entity pair and the entities' types. The input's REspecific formation, along with BERT's knowledge from unsupervised pre-training, results in RED-SandT generating informative vectors. Using these vectors, we shape relation embeddings representing the entities' distance in vector space (Bordes et al., 2013). Relation embeddings are then used as relation-wise attention over instance representation to reduce the effect of less-informative tokens. Finally, REDSandT encodes sentences by concatenating relation and weighted-instance embeddings, with relation classification to occur at bag-level as a weighted sum over its sentences' predictions.

We chose BERT over other transformer-based models because it considers bidirectionality while training. We assume that this characteristic is important to efficiently capture entities' interactions without requiring an additional task (i.e., LM finetuning in GPT (Alt et al., 2019)).

The main contributions of this paper can be summarized as follows:

- We extend BERT to handle multi-instance learning to directly fine-tune the model in a DS setting and reduce error accumulation.
- Relation embeddings captured through BERT fine-tuned on our RE-specific input manage to recognize a wider set of relations, including relations in the long-tail.
- STP encoding effectively compresses sentences to their relational tokens and allows REDSandT to outperform (Alt et al., 2019) while preserving low complexity and trained on modest hardware.
- Experiments on the NYT-10 dataset show REDSandT to surpass state-of-the-art models (Vashishth et al., 2018; Alt et al., 2019) in AUC (1.0 & 0.2 units respectively) and performance at higher recall values, while achieving a 7-10% improvement in P@{100,200,300} over (Alt et al., 2019).

- We provide<sup>1</sup> an enhanced version of NYT-10, including STP, SDP, and entity types to facilitate experimentation.
- We provide<sup>1</sup> our source code to encourage reproducible research.

# 2 Related Work

Our work is related to distant supervision, neural relation extraction (mainly pre-trained LMs), subtree parse of input, label embedding, and entity type side information.

**Distant Supervision**: DS plays a key role in RE, as it satisfies its need for extensive training data, easily and inexpensively. The use of DS (Craven and Kumlien, 1999; Snow et al., 2005) to generate large training data for RE was proposed by (Mintz et al., 2009), who assumed that all sentences that include an entity pair, which exhibits a relationship in a KB, express the same relation. However, this assumption comes with noisy labels, especially when the KB is not directly related to the domain at hand. Multi-instance learning methods were proposed to alleviate the issue, by conducting relation classification at the bag level, with a bag including instances that mention the same entity pair (Riedel et al., 2010; Hoffmann et al., 2011).

Neural Relation Extraction: While the performance of the above approaches heavily relies on handcrafted features (POS tags, named entity tags, morphological features, etc.), the advent of neural networks in RE set the focus on model architecture. Zeng et al. (2014) propose a CNN-based method to automatically capture the semantics of sentences, while PCNN (Zeng et al., 2015) that incorporates piecewise max-pooling and multi-instance learning became the common architecture to embed sentences. PCNN is used in several approaches that handle DS noisy patterns, such as intra-bag attention (Lin et al., 2016), inter-bag attention (Ye and Ling, 2019), soft labeling (Liu et al., 2017; Wang et al., 2018) and adversarial training (Wu et al., 2018; Qin et al., 2018). Moreover, Graph CNNs (GCNNs) proved an effective way to encode syntactic information from text (Vashishth et al., 2018).

The latest development of pre-trained LMs relying on transformer architecture (Vaswani et al., 2017) has shown to capture semantic and syntactic features better (Radford and Salimans, 2018). Howard and Ruder (2018) found that they signifi-

<sup>&</sup>lt;sup>1</sup>Uploaded as .zip in submission. Freely available upon acceptance.

cantly improve text classification performance, prevent overfitting, and increase sample efficiency. Shi and Lin (2019) fine-tuned BERT (Devlin et al., 2018) on the TACRED dataset showing that simple NNs built on top of BERT improve performance. Meanwhile, Alt et al. (2019) extended GPT (Radford and Salimans, 2018) to the DS setting by incorporating a multi-instance training mechanism, proving that pre-trained LMs provide a stronger signal for DS than specific linguistic and side-information features (Vashishth et al., 2018). To this extent, we take advantage of the knowledge that these models carry to capture label embeddings and boost RE.

**Side information**: Apart from model architecture, several methods propose additional information to further reduce noise. Vashishth et al. (2018) use relation phrases and incorporate Freebase entity types achieving state-of-the-art precision at higher recall values, while (Ji, 2017; Hu et al., 2019) use entity descriptors to enhance entity and label embeddings, respectively. We incorporate more generic entities' types from spaCy's model.

**Sub-Parses of Input**: Xu et al. (2015) showed the importance of the shortest-dependency path (SDP) in reducing irrelevant to RE words. Liu et al. (2018) further reduce the noise within sentences by preserving the sub-path of the sentence that connects the two entities with their least common ancestor's parent (STP). In contrast with (Liu et al., 2018), who shape the final STP sequence by re-assigning the participating tokens into their original sequence order, we preserve the tokens' order within the STP to maintain the emerged grammar information.

Label Embedding: Label embeddings aim to embed labels in the same space with word vectors. The idea comes from computer vision, with (Wang et al., 2018) to introduce them in text classification and (Hu et al., 2019) to use them as attentionmechanism over relational tokens in distantlysupervised RE. We make use of the TransE (Bordes et al., 2013) model to shape label embeddings as the entities' distance in BERT's vector space, and we show that their use both as a feature and as attention over sentences significantly improves RE.

# **3 REDSandT**

Given a bag of sentences  $\{s_1, s_2, ..., s_n\}$  that concern a specific entity pair, REDSandT generates a probability distribution on the set of possible relations. REDSandT utilizes BERT pre-trained LM to capture the semantic and syntactic features of sentences by transferring pre-trained commonsense knowledge. We extend BERT to handle multiinstance learning, and we fine-tune the model to classify the relation linking the entity pair given the associated sentences.

During fine-tuning, we employ a structured, REspecific input to minimize architectural changes to the model (Radford and Salimans, 2018). Each sentence is adapted to a structured text, including the sentences' tokens connecting the entity pair (STP) along with the entities types. We transform the input into a (sub-)word-level distributed representation using BPE and positional embeddings from BERT fine-tuned on our corpora. Then, we form final sentence representation by concatenating relation embedding and sentence representation weighted with the relation embedding. Lastly, we use attention over the bag's sentences to shape bag representation, which is then fed to a softmax layer to get the bag 's relation distribution.

REDSandT can be summarized in three components, namely sentence encoder, bag encoder, and model training. Each component is described in detail in the following sections with the overall architecture shown in Figure 1 and 2.

#### 3.1 Sentence Encoder

Given a sentence x and an entity pair  $\langle h, t \rangle$ , RED-SandT constructs a distributed representation of the sentence by concatenating relation and instance embeddings. Overall sentence encoding is represented in Figure 1, with following sections to examine the sentence encoder parts in a bottom-up way.

#### 3.1.1 Input Representation

Relation extraction requires a structured input that can sufficiently capture the latent relation between an entity pair and its surrounding text. Our input representation encodes each sentence as a sequence of tokens, depicted in the very bottom of Figure 1.

It starts with the head entity type and token(s) followed by delimiter [H-SEP], continues with the tail entity type, and token(s) followed by delimiter [T-SEP] and ends with the token sequence of the sentence's STP path. The whole input starts and ends with special delimiters [CLS] and [SEP], respectively. In BERT, [CLS] typically acts as a pooling token representing the whole sequence for downstream tasks, such as RE.

Several other sentence encodings were attempted with ablation studies in section 5.2, revealing the importance of encoding entities' types and



Figure 1: Sentence Representation in REDSandT. The input embedding  $h_0$  to BERT is created by summing over the positional and byte pair embeddings for each token in the structured input. States  $h_t$  are obtained by selfattending over the states of the previous layer  $h_{t-1}$ . Final sentence representation is obtained by concatenating the relation embedding  $r_{ht}$ , and the final fine-tuned BERT layer  $h_L$  weighted with relation attention  $\alpha_r$ . Head and tail tokens participating in the relation embedding formation are marked with bold and dashed lines respectively.

compressing the original sentence to the belowpresented STP path.

**Sub-tree parse of input sentence**: Inspired by (Liu et al., 2018), we utilize a sub-tree parse (STP) of the input sentence to reduce the noisy words within sentence and focus on the relational tokens. STP preserves the path of the sentence that connects the two entities with their least common ancestor's parent. We select STP over SDP (Xu et al., 2015) encoding, because SDP may fail to incorporate useful relational tokens (Liu et al., 2018). Moreover, our studies in section 5.2 show SDP's weakness to recognize relations in the long tail.

In contrast with (Liu et al., 2018), who shape the final STP sequence by re-assigning the participating tokens into their original sequence order, we preserve the tokens' order within STP achieving a grammatical normalization of the original sentence.

**Entity Type special tokens:** Previous research has shown that entity type information can significantly improve RE (Liu et al., 2014; Vashishth et al., 2018). In the extent that every relation puts some constraint on the type of participating entities, we incorporate the entity type in the model's structured input (see bottom of Figure 1).

In contrast with (Vashishth et al., 2018) that use

38 Freebase-specific entity types, we incorporate 18 generic entity types, captured from recognizing NYT-10 sentence's entities with the spaCy model<sup>2</sup>. We assume these types KB-independent and easily accessible with our experiments in section 5.2 indicating their inclusion to improve performance.

#### 3.1.2 Input Embeddings

The input embedding  $h_0$  to BERT is created by summing over the positional and byte pair embeddings for each token in the structured input.

**Byte-pair tokens encoding**: To make use of subword information, we tokenize input using bytepair encoding (BPE) (Sennrich et al., 2016). We particularly use the tokenizer from the pre-trained model (30,000 tokens), which we extend with 20 task-specific tokens (e.g., [H-SEP], [T-SEP], and the 18 entity type tokens). Added tokens serve a special meaning in the input representation, thus are not split into sub-words by the tokenizer.

**Positional encoding**: Positional encoding is an essential part of BERT's attention mechanism. Precisely, BERT learns a unique position embedding to represent each of the input (sub-word) token positions within the sequence.

<sup>&</sup>lt;sup>2</sup>https://spacy.io/models/en

# 3.1.3 Sentence Representation

Input sequence is transformed into feature vectors  $(h_L)$  using BERT's pre-trained language model, fine-tuned in our task. In spite of common practice to represent the sentence by the [CLS] vector in  $h_L$  (Alt et al., 2019), we argue that not all words contribute equally to sentence representation.

By encoding the underlying relation as a function of the examining entities and by giving attention to vectors related to this underlying relation, we can further reduce sentence noise and improve precision. Core modules constitute the: relation embedding, entities-wise attention, and relation attention. We examine them below.

**Relation Embedding:** Inspired by (Bordes et al., 2013), we use TransE model to get relation embeddings. TransE model regards the embedding of the underlying relation l as the distance (difference) between h and t embeddings ( $l_i = t_i - h_i$ ), assuming that a relation r holds between an entity pair (h, t). We shape relation embedding for each sentence i by applying a linear transformation on the head and tail entities vectors, activated through a Tanh layer to capture possible nonlinearities:

$$l_i = Tanh(w_l(t_i - h_i) + b_l) \tag{1}$$

, where  $w_l$  is the underlying relation weight matrix and  $b_l \in \Re^{d_t}$  is the bias vector. We mark relation embedding as l because it represents the possible underlying relation between the two entities and not the actual relationship r. Head  $h_i$  and tail  $t_i$ embeddings reflect only the entities' related tokens, which we capture through simple entities-wise attention, shown below.

Entities-wise Attention: Head and tail embeddings participating in the relation embedding are created by summing over respective token vectors from BERT's last layer  $h_L$ . We capture these tokens through head- and tail-wise attention. Headwise attention assigns the weight  $\alpha_{it}^h$  to focus on head related tokens and tail-wise attention assigns the weight  $\alpha_{it}^t$  to focus on tail related tokens.

$$\alpha_{it}^{h} = \begin{cases} 1 & \text{if } t = \text{head in STP tokens} \\ 0 & \text{otherwise} \end{cases}$$
(2)

$$\alpha_{it}^{t} = \begin{cases} 1 & \text{if } t = \text{tail in STP tokens} \\ 0 & \text{otherwise} \end{cases}$$
(3)

Head  $h_i$  and tail  $t_i$  embeddings are then shaped as follows:

$$h_i = \sum_{t=1}^T \alpha_{it}^h \cdot h_{it} \quad (4) \quad t_i = \sum_{t=1}^T \alpha_{it}^t \cdot h_{it} \quad (5)$$

**Relation Attention**: Even though REDSandT is trained on STP that naturally preserves only relational tokens, we wanted to further reduce possible left noise on sentence-level. For this reason, we use a relation attention to emphasize on sentence tokens that are mostly related to the underlying relation  $l_i$ . We calculate relation attention  $\alpha_r$  by comparing each sentence representation against the learned representation  $l_i$  for each sentence *i*:

$$\alpha_r = \frac{exp(s_i l_i)}{\sum_{j=1}^n exp(s_j l_i)} \tag{6}$$

Then, we weight BERT' s last hidden layer  $h_L \in \Re^{d_h}$  with relation embedding:

$$h'_{L} = \sum_{t=1}^{T} \alpha_r \cdot h_{it} \tag{7}$$

Finally, sentence representation  $s_i \in \Re^{d_h*2}$ is computed as the concatenation of the relation embedding  $l_i$  and the sentence's weighted hidden representation  $h'_L$ :

$$s_i = \begin{bmatrix} l_i \; ; \; h'_L \end{bmatrix} \tag{8}$$

Several other representation techniques were tested, with the presented method to outperform.

#### 3.2 BAG Encoder

Bag encoding, i.e., aggregation of sentence representations in a bag, comes to reduce noise generated by the erroneously annotated relations accompanying DS. Assuming that not all sentences contribute equally to the bag representation, we use selective attention (Lin et al., 2016) to emphasize on sentences that better express the underlying relation.

$$B = \sum_{i} \alpha_{i} s_{i}, \tag{9}$$

As seen, selective attention represents bag as a weighted sum of the individual sentences. Attention  $\alpha_i$  is calculated by comparing each sentence representation against a learned representation r:

$$\alpha_i = \frac{exp(s_i r)}{\sum_{j=1}^n exp(s_j r)} \tag{10}$$

Finally, bag representation B is fed to a softmax classifier to obtain the probability distribution over the relations.

$$p(r) = Softmax(W_r \cdot B + b_r), \quad (11)$$

where  $W_r$  is the relation weight matrix and  $b_r \in \Re^{d_r}$  is the bias vector.

#### 3.3 Training

REDSandT utilizes a transformer model, precisely BERT, which fine-tunes on our specific setup to capture the semantic features of relational sentences. Below, we present the overall process.

#### 3.3.1 Model Pre-training

For our experiments, we use the pre-trained *bert-base-cased* language model (Devlin et al., 2018), which consists of 12 layers, 12 attention heads, and 110M parameters, with each layer being a bidirectional Transformer encoder (Vaswani et al., 2017). The model is trained on cased English text of BooksCorpus and Wikipedia with a total of 800M and 2.5K words respectively. BERT is pre-trained using two unsupervised tasks: masked LM and next sentence prediction, with masked LM being its core novelty as it allows the previously impossible bidirectional training.

# 3.3.2 Model Fine-tuning

We initialize REDSandT model' s weights with the pre-trained BERT model, and we fine-tune its 4-last layers under the multi-instance learning setting presented in Figure 2, given the specific input shown in Figure 1. We end up fine-tuning only the last four layers after experimentation.

During fine-tuning, we optimize the following



Figure 2: Transformer architecture (left) and training framework (right). Sentence representation  $s_i$  is formed as shown in Figure 1.

objective:

$$L(D) = \sum_{i=1}^{|B|} log P(l_i|B_i;\theta)$$
(12)

, where for all entity pair bags |B| in the dataset, we want to maximize the probability of correctly predicting the bag's relation given its sentences' representation and parameters.

# 4 Experimental Setup

#### 4.1 Dataset

We conduct experiments on the widely used benchmark dataset NYT-10 (Riedel et al., 2010), which was built by aligning triples in Freebase to the NYT corpus and contains 53 relations. There are 522,611 (172,448) sentences, 281,270 (96,678) entity pairs, and 18,252 (1,950) relation mentions in the train (test) set. We provide an enhanced dataset, *NYT-10enhanced*, including both STP and SDP versions of the input sentences as well as the head and tail entity types to facilitate future implementations.

# 4.2 Hyper-parameter Settings

In our experiments we utilize bert-base-cased model with hidden layer dimension  $D_h=768$ , while we fine-tune the model with *max\_seq\_length*  $D_t=64$ . Regarding model's hyper-parameters, we manually tune them on the training set, based on AUC score. We select batch\_size=32 among  $\{8, 16, 32\}, epochs=3 \text{ among } \{3, 4\}, BERT's$ fine-tuned *layers*=4 among all and last  $\{2, 4, 8\}$ , learning rate  $lr=2e^{-5}$  among  $\{2e^{-4}, 1e^{-5}, 2e^{-5}\},\$ classifier dropout p=0.4 among  $\{0.2, 0.4, 0.5\}$ , and  $weight\_decay=0.001$  among  $\{0.01, 0.001\}$ . Moreover, we fine-tune our model using the Adam optimization scheme (Kingma and Lei Ba, 2015) with  $\beta_1 = 0.9, \beta_2 = 0.999$  and a cosine learning rate decay schedule with warm-up over 0.1% of training updates. We minimize loss using cross entropy criterion weighted on dataset's classes to handle the unbalanced training set. Experiments conducted on a PC with 32.00 GB Ram, Intel i7-7800X CPU@ 3.5GHz and NVIDIA's GeForce GTX 1080 with 8GB. Training time takes  $\sim 100$  min/epoch.

#### 4.3 State-of-the-art Models

For evaluating REDSandT, we compare against following state-of-the-art models:

**Mintz** (Mintz et al., 2009): A multi-class logistic regression model under distant supervision setting.

**PCNN+ATT** (Lin et al., 2016): A CNN model with instance-level attention

**RESIDE** (Vashishth et al., 2018): A NN model that makes use of relevant side information (entity types, relational phrases) and employs Graph CNN to capture syntactic information of instances.

**DISTRE** (Alt et al., 2019): A transformer model, GPT fine-tuned for RE under the distant supervision setting.

# **5** Results

#### 5.1 Comparison with state-of-the-art Models



Figure 3 compares the precision-recall curves of REDSandT against state-of-the-art models. We observe that: (1) The NN-based approaches outperform the probabilistic method (Mintz), showing human-designed features limitation against neural networks' automatically extracted features. (2) RE-SIDE, DISTRE, and REDSandT achieve better performance than PCNN+ATT, which even exhibiting the highest precision in the beginning soon follows an abrupt decline. This reveals the importance of both side-information (i.e., entity types and relation alias), and transfer knowledge. (3) RESIDE performs the best in low recalls and generally performs well, which we attribute to the multitude of sideinformation given. (4) Although DISTRE exhibits 3.5% greater precision in medium-level recalls, it presents 2-12% lower precision in recall values <0.25 compared to RESIDE, and REDSandT. (5) Our model shows the more stable behavior, with a steady, downward trend, acting similar to RESIDE at the low and medium recalls and surpassing all baselines in the very high recall values. We believe the reason is that we use potential label information as an additional feature and as attention over the instance tokens. The learned label embeddings are of high quality since they carry common-knowledge

RE methods	AUC	P@100	P@300	P@500
Mintz	0.107	52.3	45.0	39.7
PCNN+ATT	0.341	73.0	67.3	63.6
RESIDE	0.415	81.8	74.3	69.7
DISTRE	0.422	68.0	65.3	65.0
REDSandT	0.424	78.0	73.0	67.6

Table 1: AUC and P@N evaluation results. P@N represents precision calculated for the top N rated relation instances

from the pre-trained model fine-tuned on the specific dataset and task. Moreover, the chosen pretrained model, BERT, considers bidirectionality while training, being thus able to efficiently capture head and tail interaction.

Table 1, which presents AUC and precision at various points in the P-R curve, reveals our model's precision performance to be between that of RE-SIDE and DISTRE while preserving the stateof-the-art AUC. Precisely, REDSandT' s precision does not exceed RESIDE', even though it is close enough, which suggests that additional side-information would improve our model. Meanwhile, REDSandT surpasses DISTRE' s precision, which we attribute to our selected pre-trained model that efficiently captures label embeddings. Consequently, our model is more consistent to the various points of the P-R curve.

Table 2 shows the distribution over relation types for the top 300 predictions of REDSandT and baseline models. REDSandT encompasses 10 distinct relation types, two of which (place\_founded, /geographic\_distribution) are not recognized by none of rest models. PCNN+ATT predictions are highly biased towards a set of only four relation types, while RESIDE captures three additional types. DISTRE and REDSandT manage to recognize more types than all models, emphasizing the contribution of transfer knowledge. Moreover, REDSandT correctly not recognizes /location/country/capital relation that DISTRE does, as their authors found most errors to arise from the specific predicted relation in manual evaluation. Meanwhile, we highlight REDSandT' s effectiveness in recognizing relations in the long-tail. Particularly, our model captures, founders (1.47%), neighborhood\_of (1.06%), person/children (0.47%), and sports\_team/location (0.16%) relations. Relations are listed in descending order regarding population in test set with respective percentage referenced in parentheses.

Relation	red	dis	res	pcnn
/location/contains	176	168	182	214
/person/company	38	31	26	19
/person/nationality	26	32	65	59
/admin_div/country	25	13	12	6
/neighborhood_of	22	10	3	2
/person/children	5	-	6	-
/team/location	4	2	-	-
/founders	2	2	6	-
/place_founded	1	-	-	-
/geo_distribution	1	-	-	-
/country/capital	-	17	-	-
/person/place_lived	-	22	-	-

Table 2: Relation Distribution over the top 300 predictions for PCNN+ATT(pcnn), RESIDE(res), DIS-TRE(dis) and REDSandT(red) models

Metrics	AUC	P@N(%)		
		100	200	300
REDSandT w/o $r_{ht}$	0.404	80.0	72.0	67.7
REDSandT w/o ET	0.415	78.0	74.0	71.3
REDSandT w. SDP	0.418	75.0	71.0	69.7
REDSandT w/o $a_r$	0.422	75.0	76.0	71.0
REDSandT	0.424	78.0	75.0	73.0

Table 3: Evaluation results AUC and P@N of variant models on NYT-10 dataset.

# 5.2 Ablation Studies

To assess the effectiveness of the different modules of REDSandT, we create four ablation models:

**REDSandT w/o ET**: Removes entity types from input sentence representation.

**REDSandT w/o**  $r_{ht}$ : Removes relation embedding and relation attention. We represent sentence using the [CLS] token of BERT's last hidden layer  $h_L$ .

**REDSandT w/o**  $a_r$ : Removes relation attention on instance tokens.

**REDSandT w. SDP**: Replaces STP with SDP (Xu et al., 2015) in sentence encoding.

As shown in Table 3, all modules contribute to final model' s effectiveness. Greatest impact comes from relation embeddings with their removal resulting in the highest AUC (2 units) and P@300 (5.3%) drop. Meanwhile, P@100 goes up to 80% with inspection of top 300 predictions revealing a focus on 5 relation types only, with */location/contains* to make up the 79% of these. Simple integration of entity types in input representation is the next most important feature that boosts our model.

Next, "REDSandT w. SDP", shows STP's ability to capture relational facts better, while by inspecting model' s relation distribution in top 300 predictions, we observe SDP encoding to focus primarily on */person/nationality*. Finally, removing the relation attention over instance tokens exhibits the least effect in AUC (0.002) and precision ( $\sim 2\%$ ). Meanwhile, we notice that model focuses solely on 8 relation types in the top 300 predictions.

#### 5.3 Case study: Effect of relation attention

/people/person/children	
[CLS] Le ##B ##ron James [H-SEP]PERSON_ Bryce Maximu	18
James [T-SEP] Le ##B ##ron James girlfriend delivered son Bry	ce
Maximus James [SEP]	
/location/neighborhood/neighborhood of [CLS] Sweet Auburn [H-SEP]GPE Atlanta [T-SEP] Swe	et
Auburn neighborhood was in Atlanta [SEP]	

Figure 4: Relation attention weights for *children* (top) and *neighborhood\_of* (bottom) long-tail relations.

Figure 4 shows a visualization of the relation attention weights, highlighting the different parts of the sentence that drive relation extraction, for two long-tail relations. In both cases we see that the entity type is given more weight than the entity itself. Moreover, we see which tokens affect more the relation. Tokens "girlfriend", "son", and the repetition of name "James" are predictive of the "children" relation, while tokens "neighborhood", "was", "in", along with a GPE entity type show a probable "neighborhood\_of" relation.

# 6 Conclusion

We presented a novel transformer-based relation extraction model for distant supervision. REDSandT manages to acquire high-informative instance and label embeddings and is efficient at handling the noisy labeling problem of DS by exploiting the common-sense knowledge of the pre-trained BERT model. REDSandT captures high-informative embeddings for RE by fine-tuning BERT on a REspecific structured input that focuses solely on relational arguments: the sub-tree connecting the entities along with entities' types. REDSandT utilizes these vectors to encode label embeddings, which are also used as attention mechanism over instances to reduce the effect of less-informative tokens. Finally, relation extraction occurs at baglevel by concatenating label and weighted instance embeddings. Future work includes an investigation of whether additional information, such as entity descriptors, influence REDSandT's performance and to what extent.

# References

- Christoph Alt, Marc Hübner, and Leonhard Hennig. 2019. Fine-tuning Pre-Trained Transformer Language Models to Distantly Supervised Relation Extraction. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 1388–1398. Association for Computational Linguistics.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multirelational data. In Advances in Neural Information Processing Systems, pages 2787–2795.
- Mark Craven and Johan Kumlien. 1999. Constructing biological knowledge bases by extracting information from text sources. *Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology*, pages 77–86.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova Google, and A I Language. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv* preprint arXiv:1810.04805.
- Zhengqiu He, Wenliang Chen, Zhenghua Li, Meishan Zhang, Wei Zhang, and Min Zhang. 2018. SEE: Syntax-aware entity embedding for neural relation extraction. In *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pages 5795–5802.
- Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S Weld. 2011. Knowledgebased weak supervision for information extraction of overlapping relations. In ACL-HLT 2011 - Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, volume 1, pages 541–550.
- Jeremy Howard and Sebastian Ruder. 2018. Universal Language Model Fine-tuning for Text Classification. ACL 2018 - 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers), pages 328–339.
- Linmei Hu, Luhao Zhang, Chuan Shi, Liqiang Nie, Weili Guan, and Cheng Yang. 2019. Improving Distantly-Supervised Relation Extraction with Joint Label Embedding. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), page 3821–3829.
- Liu K. He S. & Zhao J. Ji, G. 2017. Distant Supervision for Relation Extraction with Hierarchical Attention and Entity Descriptions. In *Thirty-First AAAI Conference on Artificial Intelligence*, volume 2018-July.
- Diederik P Kingma and Jimmy Lei Ba. 2015. ADAM: A METHOD FOR STOCHASTIC OPTIMIZA-TION. In *ICLR*.

- 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), pages 2124–2133.
- Tianyi Liu, Xinsong Zhang, Wanhao Zhou, and Weijia Jia. 2018. Neural Relation Extraction via Inner-Sentence Noise Reduction and Transfer Learning. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 2195–2204, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Tianyu Liu, Kexiang Wang, Baobao Chang, and Zhifang Sui. 2017. A Soft-label Method for Noisetolerant Distantly Supervised Relation Extraction. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pages 1790–1795.
- Yang Liu, Kang Liu, Liheng Xu, and Jun Zhao. 2014. Exploring fine-grained entity type constraints for distantly supervised relation extraction. In COLING 2014 - 25th International Conference on Computational Linguistics, Proceedings of COLING 2014: Technical Papers, pages 2107–2116.
- 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*, pages 1003–1011. Association for Computational Linguistics.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018. DSGAN: Generative Adversarial Training for Distant Supervision Relation Extraction. Association for Computational Linguistics, pages 496–505.
- Alec Radford and Tim Salimans. 2018. Improving Language Understanding by Generative Pre-Training. URL https://s3-us-west-2. amazonaws. com/openai-assets/research-covers/languageunsupervised/language\_ understanding\_paper. pdf., pages 1–12.
- 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, volume 6323 LNAI, pages 148–163.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 -Long Papers, volume 3, pages 1715–1725.
- Peng Shi and Jimmy Lin. 2019. Simple BERT Models for Relation Extraction and Semantic Role Labeling. *arXiv preprint arXiv:1904.05255*.

- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2005. Learning syntactic patterns for automatic hypernym discovery. In *Advances in Neural Information Processing Systems*, pages 1297–1304.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D. Manning. 2012. Multi-instance multi-label learning for relation extraction. In *EMNLP-CoNLL 2012 - 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, *Proceedings of the Conference*, 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.
- Ashish Vaswani, Google Brain, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In Advances in neural information processing systems, pages 5998–6008.
- Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. 2018. Joint embedding of words and labels for text classification. In ACL 2018 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers), volume 1, pages 2321–2331. Association for Computational Linguistics.
- Yi Wu, David Bamman, and Stuart Russell. 2018. Adversarial Training for Relation Extraction. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pages 1778–1783.
- Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. Classifying relations via long short term memory networks along shortest dependency paths. In Conference Proceedings - EMNLP 2015: Conference on Empirical Methods in Natural Language Processing, pages 1785–1794. Association for Computational Linguistics.
- Zhi-Xiu Ye and Zhen-Hua Ling. 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), pages 2810–2819. Association for Computational Linguistics (ACL).
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via Piecewise Convolutional Neural Networks. In Conference Proceedings - EMNLP 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 COLING 2014
25th International Conference on Computational Linguistics, Proceedings of COLING 2014: Technical Papers, pages 2335–2344.