

On the Effectiveness of the Pooling Methods for Biomedical Relation Extraction with Deep Learning

Tuan Ngo Nguyen[†], Franck Deroncourt[‡] and Thien Huu Nguyen[†]

[†] Department of Computer and Information Science, University of Oregon

[‡] Adobe Research

{tnguyen, thien}@cs.uoregon.edu, deronco@adobe.com

Abstract

Deep learning models have achieved state-of-the-art performances on many relation extraction datasets. A common element in these deep learning models involves the pooling mechanisms where a sequence of hidden vectors is aggregated to generate a single representation vector, serving as the features to perform prediction for RE. Unfortunately, the models in the literature tend to employ different strategies to perform pooling for RE, leading to the challenge to determine the best pooling mechanism for this problem, especially in the biomedical domain. In order to answer this question, in this work, we conduct a comprehensive study to evaluate the effectiveness of different pooling mechanisms for the deep learning models in biomedical RE. The experimental results suggest that dependency-based pooling is the best pooling strategy for RE in the biomedical domain, yielding the state-of-the-art performance on two benchmark datasets for this problem.

1 Introduction

In order to analyze the entities in text, it is crucial to understand how the entities are related to each other in the documents. In the literature, this problem is formalized as relation extraction (RE), an important task in information extraction. RE aims to identify the semantic relationships between two entity mentions within the same sentences in text. Due to its important applications on many areas of natural language processing (e.g., question answering, knowledge base construction), RE has been actively studied in the last decade, featuring a variety of feature-based or kernel-based models for this problem (Zelenko et al., 2002; Zhou et al., 2005; Bunescu and Mooney, 2005; Sun et al., 2011; Chan and Roth, 2010; Nguyen et al., 2009). Recently, the introduction of deep learning has produced a new generation of models for RE with

the state-of-the-art performance on many different benchmark datasets (Zeng et al., 2014; dos Santos et al., 2015; Xu et al., 2015; Liu et al., 2015; Zhou et al., 2016; Wang et al., 2016; Zhang et al., 2017, 2018b). The advantage of deep learning over the previous approaches for RE is the ability to automatically learn effective features for the sentences from data via various network architectures. The same trend has also been observed for RE in the biomedical domain where deep learning is gaining more and more attention from the research community (Mehryary et al., 2016; Björne and Salakoski, 2018; Nguyen and Verspoor, 2018; Verga et al., 2018).

The typical deep learning models for RE have involved Convolutional Neural Networks (CNN) (Zeng et al., 2014; Nguyen and Grishman, 2015b; Zeng et al., 2015; Lin et al., 2016; Zeng et al., 2017), Recurrent Neural Networks (RNN), (Miwa and Bansal, 2016; Zhang et al., 2017), Transformer (self-attention) networks (Verga et al., 2018), and Graph Convolutional Neural Networks (GCNN) (Zhang et al., 2018b). There are two major common components in such deep learning models for RE, i.e., the representation component and the pooling component. First, in the representation component, some deep learning architectures are employed to compute a sequence of vectors to represent an input sentence for RE for which each vector tends to capture the specific context information for a word in that sentence. Such word-specific representation sequence is then fed into the second pooling component (e.g., max pooling) that aggregates the representation vectors to obtain an overall vector to represent the whole input sentence for the classification problem in RE.

While there have been many work in the literature to compare different deep learning architectures for the representation component, the pos-

sible methods for the pooling component of the deep learning models have not been systematically benchmarked for RE in general and for the biomedical domain in particular. Specifically, the prior work on relation extraction with deep learning has only assumed one form of pooling in the model without considering the possible alternatives for this component. In this work, we argue that the pooling mechanisms also have significant impact on the performance of the deep learning models for RE and it is important to understand how well different pooling methods perform in this case. Consequently, in this work, we conduct a comprehensive investigation on the effectiveness of different max pooling methods for the deep learning models of RE, focusing on the biomedical domain as the case study. Our goal is to determine the best pooling methods for the deep learning models in biomedical RE. We also want to emphasize the experiments where the pooling methods are compared in a compatible manner with the same representation components and resources for the biomedical RE models in this work. Such compatible comparison is unfortunately very rare in the current literature about deep learning for RE as new models are being intensively proposed, employing a diversity of options and resources (i.e., pre-trained word embeddings, optimizers, etc.). Therefore, this is actually the first work to compare different pooling methods for deep relation extraction on the same setting.

In the experiments, we find that syntactic information (i.e., dependency parsing) can be exploited to provide the best pooling strategies for biomedical RE. In fact, our experiments also suggest that it is more beneficial to apply the syntactic information in the pooling component of the deep learning models for biomedical RE than that in the representation component. This is different from most of the prior work on relation extraction that has only employed the syntactic information in the representation component of the deep learning models (Xu et al., 2016; Miwa and Bansal, 2016). Based on the syntax-based pooling mechanism, we achieve the state-of-the-art performance on two benchmark datasets for biomedical RE.

2 Model

Relation Extraction can be seen as a multi-class classification problem that takes a sentence and two entity mentions of interest in that sentence as

the input. The goal is to predict the semantic relation between these two entity mentions according to some predefined set of relations. Formally, let $W = [w_1, w_2, \dots, w_n]$ be the input sentence where n is the number of tokens and w_i is the i -th word/token in W . As entity mentions can span multiple consecutive words/tokens, let $[s_1, e_1]$ be the span of the first entity mention M_1 where s_1 and e_1 are the indexes for the first and last token of M_1 respectively. Similarly, we define $[s_2, e_2]$ as the span for the second entity mention M_2 . For convenience, we assume that the entity mentions are not nested, i.e., $1 \leq s_1 \leq e_1 < s_2 \leq e_2 \leq n$.

2.1 Input Vector Representation

In order to encode the positions and the entity types of the two entity mentions in the input sentence, following (Zhang et al., 2018b), we first replace the tokens in the entity mentions M_1 and M_2 with the special tokens of format M_1 - $Type_1$ and M_2 - $Type_2$ respectively ($Type_1$ and $Type_2$ represent the entity types of M_1 and M_2 respectively). The purpose of this replacement is to help the models to abstract from the specific tokens/words of the entity mentions and only focus on their positions and entity types, the two most important pieces of information of the entity mentions for RE.

Given the enriched input sentence, the first step in the deep learning models for RE is to convert each word in the input sentence into a vector to facilitate the real-valued computation of the models. In this work, the vector v_i for w_i is obtained by concatenating the following two vectors:

1. The word embeddings of w_i : The embeddings for the special tokens are initialized randomly while the embeddings for the other words are retrieved from the pre-trained word embedding table provided by the *Word2Vec* toolkit with 300 dimensions (Mikolov et al., 2013).
2. The embeddings for the part-of-speech (POS) tag of w_i in W : We assign a POS tag for each word in the input sentence using the Stanford CoreNLP toolkit. The embedding for each POS tag is also randomly initialized in this case.

Note that both the word embeddings and the POS embeddings are updated during the training time of the models in this work. The word-to-vector conversion transforms the input sentence $W = [w_1, w_2, \dots, w_n]$ into a sequence of vectors $V = [v_1, v_2, \dots, v_n]$ (respectively) that would be used as the input for all the deep learning mod-

els considered in this work to ensure a compatible comparison. As mentioned in the introduction, the deep learning models for RE involves two major components, i.e., the representation component and the pooling component. We describe the options for such components in the following sections.

2.2 The Representation Component for RE

Given the input sequence of vectors $V = [v_1, v_2, \dots, v_n]$, the next step in the deep learning models for RE is to transform this vector sequence into a more abstract vector sequence $A = [a_1, a_2, \dots, a_n]$ so a_i would capture the underlying representation for the context information specific to the i -th word in the sentence. In this work, we examine the following typical architectures to obtain such an abstract sequence A for V :

1. *CNN* (Zeng et al., 2014; Nguyen and Grishman, 2015b; dos Santos et al., 2015): *CNN* is one of the early deep learning models for RE. It involves an 1D convolution layer over the input vector sequence V with multiple window sizes for the filters. *CNN* produces a sequence of vectors in which each vector capture some n -grams specific to a word in the sentence. This sequence of vectors is used as A for our purpose.

2. *BiLSTM* (Nguyen and Grishman, 2015a): In *BiLSTM*, two Long-short Term Memory Networks (LSTM) are run over the input vector sequence V in the forward and backward direction. The hidden vectors generated at the position i by the two networks are then concatenated to constitute the abstract vector a_i for this position. Due to the recurrent nature, a_i involves the context information over the whole input sentence W although a greater focus is put on the context of the current word.

3. *BiLSTM-CNN*: This models resembles the MASS model presented in (Le et al., 2018). It first applies a bidirectional LSTM layer over the input sequence V whose results are further processed by a Convolutional Neural Network (CNN) layer as in *CNN*. We also use the output of the CNN layer as the abstract vector sequence A for this model.

4. *BiLSTM-GCNN* (Zhang et al., 2018b): Similar to *BiLSTM-CNN*, *BiLSTM-GCNN* also first employs a bidirectional LSTM network to abstract the input vector sequence V . However, in the second step, different from *BiLSTM-CNN*, *BiLSTM-GCNN* introduces a Graph Convolutional Neural

Network (GCNN) layer that consumes the LSTM hidden vectors and augments the representation for a word with the representation vectors of the surrounding words in the dependency trees. The output of the GCNN layer is also a sequence of vectors to represent the contexts for the words in the sentence and functions as the abstract sequence A in our case. *BiLSTM-GCNN* (Zhang et al., 2018b) is one of the current state-of-the-art models for RE in the literature.

Note that there are many other variants of such models for RE in the literature (Xu et al., 2016; Zhang et al., 2017; Verga et al., 2018). However, as our goal in this paper is to evaluate different pooling mechanisms for RE, we focus on these standard representation learning methods to avoid the confounding effect of the complicated models, thus better revealing the effectiveness of the pooling methods.

2.3 The Pooling Component for RE

The goal of the pooling component is to aggregate the representation vectors in the abstract sequence A to constitute an overall vector F to represent the whole input sentence W and the two entity mentions of interest (i.e., $F = \text{aggregate}(A)$). The overall representation vector should be able to capture the most important features induced in A . The typical method to achieve such aggregation in the RE models is to apply the element-wise max-pooling operation over subsets of vectors in A whose results are combined to obtain the overall representation vector. While there are different methods to select the vector subsets for the max-pooling operation, the prior work for RE has only employed one particular selection method in their deep learning models (Nguyen and Grishman, 2015a; Zhang et al., 2018b; Le et al., 2018). This raises the question about the impact of the other subset selection methods for such prior RE models. Can these methods benefit from different pooling mechanisms? What are the best pooling methods for the deep learning models in RE? In order to answer these questions, besides the architectures for the representation component in the previous section, we investigate the following subset selection methods for the pooling component of the RE models in this work:

1. *ENT-ONLY*: In this pooling method, we use the subsets of the vectors corresponding to the words in the two entity mentions of interest in

A for the max-pooling operations (i.e., M_1 with the words in the range $[s_1, e_1]$ and M_2 with the words in the range $[s_2, e_2]$). This is motivated by the utmost importance of the two entity mentions of interest for RE and employed in some prior work (Nguyen and Grishman, 2015a; Zhang et al., 2018b):

$$\begin{aligned} F_{M_1} &= \text{max-pool}(a_{s_1}, a_{s_1+1}, \dots, a_{e_1}) \\ F_{M_2} &= \text{max-pool}(a_{s_2}, a_{s_2+1}, \dots, a_{e_2}) \\ F_{ENT-ONLY} &= [F_{M_1}, F_{M_2}] \end{aligned}$$

2. *ENT-SENT*: Besides the entity mentions, the other context words in the sentence might also involve important information for the relation prediction in RE. For instance, in the sentence “Acetazolamide can elevate cyclosporine levels.”, the context word “*elevate*” is crucial to determine the semantic relations between the two entity mentions of interest “*Acetazolamide* and “*cyclosporine*”. In order to capture such important contexts for pooling, the typical approach in the prior work for RE is to perform the max-pooling operation over the abstract vectors for every word in the sentence (i.e., the whole set A) (Zeng et al., 2014; dos Santos et al., 2015; Le et al., 2018). The rationale is to select the features of the abstract vectors in A with the highest values in each dimension to reveal the most important context for RE. The max-pooled vector over the whole set A is combined with the $F_{ENT-ONLY}$ vector in this method:

$$\begin{aligned} F_{SENT} &= \text{max-pool}(a_1, a_2, \dots, a_n) \\ F_{ENT-SENT} &= [F_{ENT-ONLY}, F_{SENT}] \end{aligned}$$

3. *ENT-DYM*: Similar to *ENT-SENT*, this method also seeks the important context information beyond the two entity mentions of interest. However, instead of taking the whole vector sequence A for the pooling, *ENT-DYM* divides A into three separate vector subsequences based on the start and end indexes of the first and second entity mentions (i.e., s_1 and e_2) respectively. The max-pooling operation is then applied over these three subsequences and the resulting vectors are combined to form an overall vector (i.e., dynamic

pooling) (Zeng et al., 2015):

$$\begin{aligned} F_{LEFT} &= \text{max-pool}(a_1, a_2, \dots, a_{s_1-1}) \\ F_{MIDDLE} &= \text{max-pool}(a_{s_1}, a_{s_1+1}, \dots, a_{e_2}) \\ F_{RIGHT} &= \text{max-pool}(a_{e_2+1}, a_{e_2+2}, \dots, a_n) \\ F_{ENT-DYM} &= [F_{LEFT}, F_{MIDDLE}, F_{RIGHT}, \\ &\quad F_{ENT-ONLY}] \end{aligned}$$

4. *ENT-DEP0*: The previous pooling methods have only relied on the sequential structures of the sentence where the chosen subsets of A for pooling always contain vectors for the consecutive words in the sentence. Unfortunately, such sequential pooling might introduce irrelevant words into the selected subsets of A , potentially causing noise in the pooling features and impeding the performance of the RE models. For instance, in the previous sentence example “Acetazolamide can elevate cyclosporine levels.”, the *ENT-SENT* and *ENT-DYM* methods would also include the word “*levels*” in the pooling subsets that is not very important for the relation prediction in this case. Consequently, in *ENT-DEP0*, we explore the possibility to use the dependency parse tree of the input sentence W to filter out the irrelevant words for the pooling operation. In particular, instead of considering every word in the input sentence, *ENT-DEP0* only pools over the abstract vectors in A that correspond to the words along the shortest dependency path (SDP) between the two entity mentions M_1 and M_2 in the dependency tree for W (called $SDP0(M_1, M_2)$). Note that the shortest dependency paths have been shown to be able to select the important context words for RE in many previous work (Zhou et al., 2005; Chan and Roth, 2010; Xu et al., 2016). Similar to *ENT-SENT* and *ENT-DYM*, we also include $F_{ENT-ONLY}$ in this method:

$$\begin{aligned} F_{DEP0} &= \text{max-pool}_{a_i \in SDP0(M_1, M_2)}(a_i) \\ F_{ENT-DEP0} &= [F_{DEP0}, F_{ENT-ONLY}] \end{aligned}$$

5. *ENT-DEP1*: This method is similar to *ENT-DEP0*. However, instead of directly pooling over the words in the shortest dependency path $SDP0(M_1, M_2)$, *ENT-DEP1* extends this path to also include every word that is connected to some word in $SDP0(M_1, M_2)$ via an edge in the dependency tree for W (i.e., one edge distance from $SDP0(M_1, M_2)$). We denote this extended word set by $SDP1(M_1, M_2)$ for which the corresponding abstract vectors in A would be chosen for

the max-pooling operation. The motivation for $SDP1(M_1, M_2)$ is that the representations of the words close to the shortest dependency path between M_1 and M_2 might also provide useful information to improve the performance for RE. In our experiments, we find that one edge is the optimal distance to enlarge the shortest dependency paths. Using larger distance for the pooling mechanism would hurt the performance of the deep learning models for RE:

$$F_{DEP1} = \text{max-pool}_{a_i \in SDP1(M_1, M_2)}(a_i)$$

$$F_{ENT-DEP1} = [F_{DEP1}, F_{ENT-ONLY}]$$

Once the overall representation vector F for the input sentence W and the two entity mentions of interest has been produced, we feed it into a feed-forward neural network with a softmax layer in the end to obtain the probability distribution $P(y|W, M_1, M_2) = \text{feed-forward}(F)$ over the possible relation types for our RE problem. This probability distribution would then be used for both making prediction (i.e., by taking the relation type with the highest probability) and training models (i.e., by optimizing the negative log-likelihood function).

3 Experiments

3.1 Datasets

In order to evaluate the performance of the models in this work, we employ the following biomedical datasets for RE in the experiments:

DDI-2013 (Herrero-Zazo et al., 2013): This dataset contains 730 documents from the Drugbank database, involving about 25,000 examples for the training and test sets (each example consists of a sentence and two entity mentions of interest for classification). There are 4 entity types (i.e., *drug*, *brand*, *group* and *brand_n*) and 5 relation types (i.e., *mechanism*, *advise*, *effect*, *int*, and *no_relation*) in this dataset. The *no_relation* is to indicate any example that does not belong to any relation types of interest. This dataset is severely imbalanced, containing 85% negative examples in the training dataset. In order to deal with such imbalanced data, we employ weighted sampling that equally distributes the selection probability for the positive and negative examples.

BB3 (Deléger et al., 2016). This dataset contains 95 documents; each of them involves a title and abstract from a document from the PubMed

database. There are 800 examples in this dataset divided into two separate sets (i.e., the training set and the validation set). BB3 also include a test set; however, the relation types for the examples in this test set are not provided. In order to obtain the performance of the models on the test set, the performers need to submit their system outputs to an official API that would evaluate the output and return the model performance. We train the models in this work on the training data and employ the official API to obtain their test set performance to be reported in the experiments for this dataset.

Following the prior work on these datasets (Chowdhury and Lavelli, 2013; Lever and Jones, 2016; Zhou et al., 2018; Le et al., 2018), we use the micro-averaged F1 scores as the performance measure in the experiments to ensure a compatible comparison.

3.2 Parameters and Resources

As the DDI-2013 dataset does not involve a development set, we tune the parameters for the models in this work based on the validation data of the BB3 dataset and use the selected parameters for both datasets in the experiments. The best parameters from this tuning process include the learning rate of 0.5 and momentum of 0.8 for the stochastic gradient descent (SGD) optimizer with nesterov’s momentum to optimize the models. In order to regularize the models, we apply dropout between layers with the drop rate for word embeddings set to 0.7 and other drop rates set to 0.5. We also employ the weight dropout *DropConnect* in (Wan et al., 2013) to regularize the hidden-to-hidden transition matrix within each bidirectional LSTM in the models (Merity et al., 2017). For all the models that involve bidirectional LSTMs (i.e., *BiLSTM*, *BiLSTM-CNN*, and *BiLSTM-GCNN*), two layers of bidirectional LSTMs are utilized with 300 hidden units for each LSTM network. For the models with CNN components (i.e., *CNN* and *BiLSTM-CNN*), we use one CNN layer with multiple window sizes of 2, 3, 4, and 5 for the filters (200 filters for each window size). For the *BiLSTM-GCN* model, two GCNN layers are employed with 300 hidden units in each layer. Finally, for the final feed-forward neural network to compute the probability distribution (i.e., feed-forward), we utilize two hidden layers for which 1000 hidden units are used for the first layer and the number of hidden units for the sec-

ond layer is determined by the number of relation types in the datasets.

3.3 Evaluating the Pooling Methods for RE

This section evaluates the performance of different pooling methods when they are applied to the deep learning models for RE on the two datasets DDI-2013 and BB3. In particular, we integrate each of the pooling methods in Section 2.3 (i.e., *ENT-ONLY*, *ENT-SENT*, *ENT-DYM*, *END-DEPO*, and *END-DEPI*) into each of the deep learning models in Section 2.2 (i.e., *CNN*, *BiLSTM*, *BiLSTM-CNN*, and *BiLSTM-GCN*), resulting 20 different model combinations to be investigated in this section. For each model combination, we train five versions of the model with different random seeds for parameter initialization over the training datasets. The performance of such versions over the test sets is averaged to serve as the overall model performance on the corresponding dataset. Tables 1 and 2 report the performance of the models on the DDI-2013 dataset and BB3 dataset respectively.

Model	P	R	F1
<i>CNN</i>			
+ <i>ENT-ONLY</i>	52.7	43.1	47.4
+ <i>ENT-SENT</i>	75.8	60.7	67.3
+ <i>ENT-DYM</i>	66.5	70.6	68.5
+ <i>ENT-DEPO</i>	59.8	61.5	60.6
+ <i>ENT-DEPI</i>	67.6	65.1	66.3
<i>BiLSTM</i>			
+ <i>ENT-ONLY</i>	74.0	69.4	71.6
+ <i>ENT-SENT</i>	74.8	71.7	73.1
+ <i>ENT-DYM</i>	71.5	73.4	72.4
+ <i>ENT-DEPO</i>	72.8	69.4	71.1
+ <i>ENT-DEPI</i>	71.6	76.4	73.9
<i>BiLSTM-CNN</i>			
+ <i>ENT-ONLY</i>	69.6	72.3	70.9
+ <i>ENT-SENT</i>	69.4	74.9	72.0
+ <i>ENT-DYM</i>	71.0	69.7	71.8
+ <i>ENT-DEPO</i>	72.2	69.5	70.8
+ <i>ENT-DEPI</i>	71.0	74.3	72.6
<i>BiLSTM-GCNN</i>			
+ <i>ENT-ONLY</i>	69.3	71.4	70.4
+ <i>ENT-SENT</i>	72.2	71.9	72.0
+ <i>ENT-DYM</i>	69.7	73.9	71.7
+ <i>ENT-DEPO</i>	70.1	71.1	70.6
+ <i>ENT-DEPI</i>	72.7	72.9	72.8

Table 1: Results on DDI 2013

From the tables, we have the following observations about the effectiveness of the pooling methods for RE with deep learning:

1. Comparing *ENT-SENT*, *ENT-DYM* and *ENT-ONLY*, we see that the pooling methods over the whole sentence (i.e., *ENT-SENT* and *ENT-DYM*) are significantly better than *ENT-ONLY* that only focuses on the two entity mentions of interest in

Model	P	R	F1
<i>CNN</i>			
+ <i>ENT-ONLY</i>	54.2	65.7	59.1
+ <i>ENT-SENT</i>	55.0	62.5	59.1
+ <i>ENT-DYM</i>	54.6	53.3	53.5
+ <i>ENT-DEPO</i>	55.9	65.8	60.6
+ <i>ENT-DEPI</i>	55.7	67.7	61.1
<i>BiLSTM</i>			
+ <i>ENT-ONLY</i>	58.9	59.6	59.2
+ <i>ENT-SENT</i>	60.7	59.2	59.9
+ <i>ENT-DYM</i>	50.2	66.0	56.9
+ <i>ENT-DEPO</i>	51.6	78.0	61.9
+ <i>ENT-DEPI</i>	54.7	72.6	62.4
<i>BiLSTM-CNN</i>			
+ <i>ENT-ONLY</i>	56.4	66.2	60.8
+ <i>ENT-SENT</i>	53.6	69.2	60.5
+ <i>ENT-DYM</i>	47.1	78.0	58.7
+ <i>ENT-DEPO</i>	55.9	71.4	62.5
+ <i>ENT-DEPI</i>	54.1	74.7	62.4
<i>BiLSTM-GCNN</i>			
+ <i>ENT-ONLY</i>	62.7	56.1	58.9
+ <i>ENT-SENT</i>	58.4	58.7	58.5
+ <i>ENT-DYM</i>	56.8	58.4	56.6
+ <i>ENT-DEPO</i>	55.6	67.4	60.8
+ <i>ENT-DEPI</i>	54.4	71.1	61.5

Table 2: Results on BioNLP BB3

the DDI-2013 dataset. This is true across different deep learning models in this work. However, this comparison is reversed for the BB3 dataset where *ENT-ONLY* is in general better or comparable to *ENT-SENT* and *ENT-DYM* over different deep learning models. We attribute such phenomena to the fact that the BB3 dataset often contains many entity mentions and relations within a single sentence (i.e., overlapping contexts) while the sentences in DDI-2013 tend to involve only a single relation with few entity mentions. This make *ENT-SENT* and *ENT-DYM* ineffective for BB3 as the pooling mechanisms over the whole sentence are likely to involve the contexts for the other entity mentions and relations in the sentences, causing the low quality of the resulting representations and the confusion of the model for the relation prediction. This problem is less severe in DDI-2013 as the context of the whole sentence (with a single relation) is more aligned with the important context for the relation prediction. We call the many entity mentions and relations in a single sentence of BB3 as the multiple relation effect for convenient discussion in this paper.

2. Comparing *ENT-SENT* and *ENT-DYM*, their performance are comparable in DDI-2013 (except for *CNN* where *ENT-DYM* is better); however, in the BB3 dataset, *ENT-SENT* significantly outperforms *ENT-DYM* over all the models. This suggests the amplification of the multiple relation ef-

fect in BB3 due to *ENT-DYM* where the separation of the sentence context for pooling encourages the emergence of context information for multiple relations in the final representation vector and increases the confusion of the models.

3. Comparing the syntax-based pooling methods and the non-syntax pooling methods, the pooling based on dependency paths (i.e., *ENT-DEPO*) is worse than the non-syntax pooling methods (i.e., *ENT-SENT* and *ENT-DYM*) and perform comparably with *ENT-ONLY* in the DDI-2013 dataset over all the models (except for the *CNN* model where *ENT-ONLY* is much worse). These evidences suggest that the dependency paths themselves are not able to capture effective contexts for the pooling operation beyond the entity mentions for biomedical RE in DDI-2013. However, when we switch to the BB3 dataset, it turns out that *ENT-DEPO* is significantly better than all the non-syntax pooling methods (i.e., *ENT-ONLY*, *ENT-SENT* and *ENT-DYM*) for all the comparing models. This can be explained by the multiple relation effect in BB3 for which the dependency paths help to identify the most related context words for the two given entity mentions and filter out the confusing context words for the other relations in the sentences. The models would thus become less confused with different contexts for multiple relations as those in *ENT-SENT* and *ENT-DYM* for better performance in this case.

4. Finally, among all the pooling methods, we find that *ENT-DEPI* significantly outperforms the other pooling methods across different models and datasets (except the *CNN* model on DDI-2013 and *BiLSTM* on BB3). In particular, the performance improvement is substantial over the non-syntax pooling methods in BB3 where *ENT-DEPI* is up to 2% better than *ENT-SENT*, *ENT-DYM* and *ENT-ONLY* on the absolute F1 scores. This helps to demonstrate the benefits of *ENT-DEPI* for biomedical RE to both recognize the important context words for pooling in DDI-2013 and reduce the confusion effect of the multiple relations in single sentences for the models in BB3.

3.4 Comparing the Deep Learning Models for RE

Regarding the comparison among different deep learning models, the major observations from Tables 1 and 2 include:

1. The performance of *CNN* is in general

worse than the other models with the bidirectional LSTM components (i.e., *BiLSTM*, *BiLSTM-CNN* and *BiLSTM-GCN*) over different pooling methods and datasets. This illustrates the importance of bidirectional LSTMs to capture the effective feature representations for biomedical RE.

2. Comparing *BiLSTM* and *BiLSTM-CNN*, we find that *BiLSTM* is better in DDI-2013 while *BiLSTM-CNN* achieves better performance in BB3 (over different pooling methods). In other words, the CNN layer is only helpful for the *BiLSTM* model in the BB3 dataset. This can also be attributed to the multiple relation effect in BB3 where the CNN layer helps to further abstract the representations from *BiLSTM* to better reveal the underlying structures in such confusing and complicated contexts in the sentences of BB3 for RE.

3. Graph convolutions over the dependency trees are not effective for biomedical RE as incorporating it into the *BiLSTM* model hurts the performance significantly. In particular, *BiLSTM-GCNN* is significantly worse than *BiLSTM* no matter which pooling methods are applied and which datasets are used for evaluation.

4. Interestingly, comparing the *BiLSTM* model with the *ENT-DEPI* pooling method (i.e., *BiLSTM + ENT-DEPI*) and the *BiLSTM-GCN* model with the non-syntax pooling methods (i.e., *ENT-ONLY*, *ENT-SENT* and *ENT-DYM*), we see that *BiLSTM + ENT-DEPI* is significantly better with large performance gaps over both datasets DDI-2013 and BB3. For example, *BiLSTM + ENT-DEPI* is 1.9% better than *BiLSTM-GCNN + ENT-SENT* in the DDI-2013 dataset and 3.5% better than *BiLSTM-GCNN + ENT-ONLY* in BB3 with respect to the absolute F1 scores. In fact, *BiLSTM + ENT-DEPI* also achieves the best performance among the compared models in this section for both datasets. The major difference between *BiLSTM + ENT-DEPI* and *BiLSTM-GCN* with the non-syntax pooling methods lies at the specific component of the models where the syntactic information (i.e., the dependency trees) is applied. In *BiLSTM-GCN* with the non-syntax pooling methods, the syntactic information is employed in the representation learning component while in *BiLSTM + ENT-DEPI*, the application of the syntactic information is postponed all the way to the pooling component. Our experiments thus demonstrate that it is more effective to utilize the syntactic information in the pooling component

than in the representation learning component of the deep learning models for biomedical RE. This is an interesting and unique observation given that the prior work for RE has only focused on using the syntactic information in the representation component and never explicitly investigated the effectiveness of the syntactic information for the pooling component of the deep learning models.

3.5 Comparing to the State-of-the-art Models

In order to further demonstrate the advantage of the syntactic information for the pooling component for biomedical RE, this section compares *BiLSTM + ENT-DEPI* (i.e., the best model with the *ENT-DEPI* pooling in this work) with the best reported models on the two datasets DDI-2013 and BB3. For a fair comparison between models, we select the previous single (non-ensemble) models for the comparison in this section. Tables 3 and 4 presents the model performance.

Models	P	R	F1
(Raihani and Laachfoubi, 2017)	73.6	70.1	71.8
(Zhang et al., 2018a)	74.1	71.8	72.9
(Zhou et al., 2018)	75.8	70.3	73.0
(Björne and Salakoski, 2018)	75.3	66.3	70.5
<i>BiLSTM + ENT-DEPI</i>	71.6	76.4	73.9

Table 3: Comparison with the state-of-the-art systems on the DDI-2013 test set

Models	P	R	F1
(Lever and Jones, 2016)	51.0	61.5	55.8
(Mehryary et al., 2016)	62.3	44.8	52.1
(Li et al., 2016)	56.3	58.0	57.1
(Le et al., 2018)	59.8	51.3	55.2
<i>BiLSTM + ENT-DEPI</i>	54.7	72.6	62.4

Table 4: Comparison with the state-of-the-art systems on the BB3 test set

The most important observation from the tables is that the *BiLSTM* model, once combined with the *ENT-DEPI* pooling method, significantly outperforms the previous models on DDI-2013 and BB3, establishing new state-of-the-art performance for these datasets. In particular, in the DDI-2013 dataset, *BiLSTM + ENT-DEPI* is 0.9% better than the current state-of-the-art model in (Zhou et al., 2018) while the performance improvement over the best reported model for BB3 in (Li et al., 2016) is 5.3% (over the absolute F1 scores). Such substantial improvement clearly demonstrates the ad-

vantages of the syntactic information and its delayed application in the pooling component of the deep learning models for biomedical RE.

4 Related Work

Traditional work on RE has mostly used feature engineering with syntactical information for statistical or kernel based classifiers (Zelenko et al., 2002; Zhou et al., 2005; Bunescu and Mooney, 2005; Sun et al., 2011; Chan and Roth, 2010). Recently, deep learning has been shown to advance many benchmark datasets for this RE problem due to its representation learning capacity. The typical architectures for such deep learning models involve CNN, LSTM, the attention mechanism and their variants (Zeng et al., 2014; dos Santos et al., 2015; Zhou et al., 2016; Wang et al., 2016; Nguyen and Grishman, 2015a; Miwa and Bansal, 2016; Zhang et al., 2017, 2018b). Deep learning has also been applied to biomedical RE in the last couple of years and started to demonstrate much potentials for this area (Mehryary et al., 2016; Björne and Salakoski, 2018; Nguyen and Verspoor, 2018; Verga et al., 2018).

Pooling is a common and crucial component in most of the deep learning models for RE. (Nguyen and Grishman, 2015b; dos Santos et al., 2015) apply the pooling operation over the whole sentence for RE while Zeng et al. (2015) proposes the dynamic pooling mechanism in the CNN models. However, none of these prior work systematically examines different pooling mechanisms for deep learning in RE as we do in this work.

5 Conclusion

We conduct a comprehensive study on the effectiveness of different pooling mechanisms for the deep learning models in biomedical relation extraction. Our experiments suggest that the pooling mechanisms have a significant impact on the performance of the deep learning models and a careful evaluation should be done to decide the appropriate pooling mechanism for the biomedical RE problem. From the experiments, we also find that syntactic information (i.e., dependency parsing) provides the best pooling methods for the models and biomedical RE datasets we investigate in this work (i.e., *ENT-DEPI*). We achieve the state-of-the-art performance for biomedical RE over the two datasets DDI-2013 and BB3 with such syntax-based pooling methods.

References

- Jari Björne and Tapio Salakoski. 2018. Biomedical Event Extraction Using Convolutional Neural Networks and Dependency Parsing. In *Proceedings of the BioNLP 2018 Workshop*, pages 98–108.
- Razvan Bunescu and Raymond Mooney. 2005. A Shortest Path Dependency Kernel for Relation Extraction. In *Proceedings of the EMNLP-HLT 2005*, pages 724–731.
- Yee S. Chan and Dan Roth. 2010. Exploiting background knowledge for relation extraction. In *COLING*.
- Md Faisal Mahbub Chowdhury and Alberto Lavelli. 2013. FBK-irst : A Multi-Phase Kernel Based Approach for Drug-Drug Interaction Detection and Classification that Exploits Linguistic Information. In *Proceedings of the Seventh International Workshop on Semantic Evaluation*, pages 351–355.
- Louise Deléger, Robert Bossy, Estelle Chaix, Mouhamadou Ba, Arnaud Ferré, Philippe Bessières, and Claire Nédellec. 2016. Overview of the Bacteria Biotope Task at BioNLP Shared Task 2016. In *Proceedings of the BioNLP 2016 Workshop*, pages 12–22.
- Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2015. Classifying Relations by Ranking with Convolutional Neural Networks. In *Proceedings of the IJCNLP 2015*, pages 626–634.
- María Herrero-Zazo, Isabel Segura-Bedmar, Paloma Martínez, and Thierry Declerck. 2013. [The DDI corpus: An annotated corpus with pharmacological substances and drug–drug interactions](#). *Journal of Biomedical Informatics*, 46(5):914–920.
- Hoang-Quynh Le, Duy-Cat Can, Sinh T. Vu, Thanh Hai Dang, Mohammad Taher Pilehvar, and Nigel Collier. 2018. Large-scale Exploration of Neural Relation Classification Architectures. In *Proceedings of the EMNLP 2018*, pages 2266–2277.
- Jake Lever and Steven JM Jones. 2016. VERSE: Event and Relation Extraction in the BioNLP 2016 Shared Task. In *Proceedings of the BioNLP 2016 Workshop*, pages 42–49.
- L. Li, and, and D. Huang and. 2016. [Biomedical event extraction via Long Short Term Memory networks along dynamic extended tree](#). In *Proceedings of the IEEE-BIBM 2016*, pages 739–742.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. [Neural Relation Extraction with Selective Attention over Instances](#). In *Proceedings of the ACL 2016*, pages 2124–2133.
- Yang Liu, Furu Wei, Sujian Li, Heng Ji, Ming Zhou, and Houfeng Wang. 2015. A dependency-based neural network for relation classification. *arXiv preprint arXiv:1507.04646*.
- Farrokh Mehryary, Jari Björne, Sampo Pyysalo, Tapio Salakoski, and Filip Ginter. 2016. [Deep Learning with Minimal Training Data: TurkuNLP Entry in the BioNLP Shared Task 2016](#). In *Proceedings of the BioNLP 2016 Workshop*, pages 73–81.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2017. [Regularizing and Optimizing LSTM Language Models](#). In *Proceedings of the ICLR 2018*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). *arXiv:1301.3781 [cs]*.
- Makoto Miwa and Mohit Bansal. 2016. End-to-End Relation Extraction using LSTMs on Sequences and Tree Structures. In *Proceedings of the ACL 2016*, pages 1105–1116.
- Dat Quoc Nguyen and Karin Verspoor. 2018. Convolutional neural networks for chemical-disease relation extraction are improved with character-based word embeddings. In *Proceedings of the BioNLP 2018 Workshop*, pages 129–136.
- Thien Huu Nguyen and Ralph Grishman. 2015a. Combining neural networks and log-linear models to improve relation extraction. *arXiv preprint arXiv:1511.05926*.
- Thien Huu Nguyen and Ralph Grishman. 2015b. Relation Extraction: Perspective from Convolutional Neural Networks. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 39–48.
- Truc-Vien T. Nguyen, Alessandro Moschitti, and Giuseppe Riccardi. 2009. Convolution kernels on constituent, dependency and sequential structures for relation extraction. In *EMNLP*.
- Anass Raihani and Nabil Laachfoubi. 2017. [A Rich Feature-based Kernel Approach for Drug-Drug Interaction Extraction](#). In *In Proceedings of the IJACSA 2017*, volume 8.
- Ang Sun, Ralph Grishman, and Satoshi Sekine. 2011. Semi-supervised relation extraction with large-scale word clustering. In *ACL*.
- Patrick Verga, Emma Strubell, and Andrew McCallum. 2018. [Simultaneously Self-Attending to All Mentions for Full-Abstract Biological Relation Extraction](#). In *Proceedings of the NAACL-HLT 2018*, pages 872–884.
- Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. 2013. Regularization of Neural Networks using DropConnect. In *Proceedings of the ICML 2013*, pages 1058–1066.
- Linlin Wang, Zhu Cao, Gerard de Melo, and Zhiyuan Liu. 2016. [Relation Classification via Multi-Level Attention CNNs](#). In *Proceedings of the ACL 2016*, pages 1298–1307.

- Yan Xu, Ran Jia, Lili Mou, Ge Li, Yunchuan Chen, Yangyang Lu, and Zhi Jin. 2016. Improved relation classification by deep recurrent neural networks with data augmentation. In *Proceedings of the COLING 2016*, pages 1461–1470.
- 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 *Proceedings of the EMNLP 2015*, pages 1785–1794.
- Dmitry Zelenko, Chinatsu Aone, and Anthony Richardella. 2002. [Kernel Methods for Relation Extraction](#). In *Proceedings of the EMNLP 2002*, pages 71–78.
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. [Distant Supervision for Relation Extraction via Piecewise Convolutional Neural Networks](#). In *Proceedings of the EMNLP 2015*, 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 the COLING 2014*, pages 2335–2344.
- Wenyuan Zeng, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2017. [Incorporating Relation Paths in Neural Relation Extraction](#). In *Proceedings of the EMNLP 2017*, pages 1768–1777.
- Yijia Zhang, Wei Zheng, Hongfei Lin, Jian Wang, Zhihao Yang, and Michel Dumontier. 2018a. [Drug-Drug Interaction Extraction Via Hierarchical Rnns on Sequence and Shortest Dependency Paths](#). *Bioinformatics*, 34(5):828–835.
- Yuhao Zhang, Peng Qi, and Christopher D. Manning. 2018b. [Graph Convolution over Pruned Dependency Trees Improves Relation Extraction](#). In *Proceedings of the EMNLP 2018*, pages 2205–2215.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. [Position-aware Attention and Supervised Data Improve Slot Filling](#). In *Proceedings of the EMNLP 2017*, pages 35–45.
- Deyu Zhou, Lei Miao, and Yulan He. 2018. [Position-Aware Deep Multi-Task Learning for Drug-Drug Interaction Extraction](#). *Journal of Artificial Intelligence in Medicine*, 87:1–8.
- Guodong Zhou, Jian Su, Jie Zhang, and Min Zhang. 2005. [Exploring Various Knowledge in Relation Extraction](#). In *Proceedings of the ACL 2005*, pages 427–434.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. [Attention-Based Bidirectional Long Short-Term Memory Networks for Relation Classification](#). In *Proceedings of the ACL 2016*, pages 207–212.