

# Deep learning for extracting protein-protein interactions from biomedical literature

Yifan Peng

Zhiyong Lu

National Center for Biotechnology Information

National Library of Medicine

National Institutes of Health

Bethesda, MD 20894

{yifan.peng, zhiyong.lu}@nih.gov

## Abstract

State-of-the-art methods for protein-protein interaction (PPI) extraction are primarily feature-based or kernel-based by leveraging lexical and syntactic information. But how to incorporate such knowledge in the recent deep learning methods remains an open question. In this paper, we propose a multichannel dependency-based convolutional neural network model (McDepCNN). It applies one channel to the embedding vector of each word in the sentence, and another channel to the embedding vector of the head of the corresponding word. Therefore, the model can use richer information obtained from different channels. Experiments on two public benchmarking datasets, AIMed and BioInfer, demonstrate that McDepCNN compares favorably to the state-of-the-art rich-feature and single-kernel based methods. In addition, McDepCNN achieves 24.4% relative improvement in F1-score over the state-of-the-art methods on cross-corpus evaluation and 12% improvement in F1-score over kernel-based methods on “difficult” instances. These results suggest that McDepCNN generalizes more easily over different corpora, and is capable of capturing long distance features in the sentences.

## 1 Introduction

With the growing amount of biomedical information available in the textual form, there has been considerable interest in applying natural language processing (NLP) techniques and machine learning (ML) methods to the biomedical litera-

ture (Huang and Lu, 2015; Leaman and Lu, 2016; Singhal et al., 2016; Peng et al., 2016). One of the most important tasks is to extract protein-protein interaction relations (Krallinger et al., 2008).

Protein-protein interaction (PPI) extraction is a task to identify interaction relations between protein entities mentioned within a document. While PPI relations can span over sentences and even cross documents, current works mostly focus on PPI in individual sentences (Pyysalo et al., 2008; Tikk et al., 2010). For example, “ARFTS” and “XIAP-BIR3” are in a PPI relation in the sentence “ARFTS<sub>PROT1</sub> specifically binds to a distinct domain in XIAP-BIR3<sub>PROT2</sub>”.

Recently, deep learning methods have achieved notable results in various NLP tasks (Manning, 2015). For PPI extraction, convolutional neural networks (CNN) have been adopted and applied effectively (Zeng et al., 2014; Quan et al., 2016; Hua and Quan, 2016). Compared with traditional supervised ML methods, the CNN model is more generalizable and does not require tedious feature engineering efforts. However, how to incorporate linguistic and semantic information into the CNN model remains an open question. Thus previous CNN-based methods have not achieved state-of-the-art performance in the PPI task (Zhao et al., 2016a).

In this paper, we propose a multichannel dependency-based convolutional neural network, McDepCNN, to provide a new way to model the syntactic sentence structure in CNN models. Compared with the widely-used one-hot CNN model (e.g., the shortest-path information is firstly transformed into a binary vector which is zero in all positions except at this shortest-path’s index, and then applied to CNN), McDepCNN utilizes a separate channel to capture the dependencies of the sentence syntactic structure.

To assess McDepCNN, we evaluated our

model on two benchmarking PPI corpora, AIMed (Bunescu et al., 2005) and BioInfer (Pyysalo et al., 2007). Our results show that McDepCNN performs better than the state-of-the-art feature- and kernel-based methods.

We further examined McDepCNN in two experimental settings: a cross-corpus evaluation and an evaluation on a subset of “difficult” PPI instances previously reported (Tikk et al., 2013). Our results suggest that McDepCNN is more generalizable and capable of capturing long distance information than kernel methods.

The rest of the manuscript is organized as follows. We first present related work. Then, we describe our model in Section 3, followed by an extensive evaluation and discussion in Section 4. We conclude in the last section.

## 2 Related work

From the ML perspective, we formulate the PPI task as a binary classification problem where discriminative classifiers are trained with a set of positive and negative relation instances. In the last decade, ML-based methods for the PPI tasks have been dominated by two main types: the feature-based vs. kernel based method. The common characteristic of these methods is to transform relation instances into a set of features or rich structural representations like trees or graphs, by leveraging linguistic analysis and knowledge resources. Then a discriminative classifier is used, such as support vector machines (Vapnik, 1995) or conditional random fields (Lafferty et al., 2001).

While these methods allow the relation extraction systems to inherit the knowledge discovered by the NLP community for the pre-processing tasks, they are highly dependent on feature engineering (Fundel et al., 2007; Van Landeghem et al., 2008; Miwa et al., 2009b; Bui et al., 2011). The difficulty with feature-based methods is that data cannot always be easily represented by explicit feature vectors.

Since natural language processing applications involve structured representations of the input data, deriving good features is difficult, time-consuming, and requires expert knowledge. Kernel-based methods attempt to solve this problem by implicitly calculating dot products for every pair of examples (Erkan et al., 2007; Airola et al., 2008; Miwa et al., 2009a; Kim et al., 2010; Chowdhury et al., 2011). Instead of extracting fea-

ture vectors from examples, they apply a similarity function between examples and use a discriminative method to label new examples (Tikk et al., 2010). However, this method also requires manual effort to design a similarity function which can not only encode linguistic and semantic information in the complex structures but also successfully discriminate between examples. Kernel-based methods are also criticized for having higher computational complexity (Collins and Duffy, 2002).

Convolutional neural networks (CNN) have recently achieved promising results in the PPI task (Zeng et al., 2014; Hua and Quan, 2016). CNNs are a type of feed-forward artificial neural network whose layers are formed by a convolution operation followed by a pooling operation (LeCun et al., 1998). Unlike feature- and kernel-based methods which have been well studied for decades, few studies investigated how to incorporate syntactic and semantic information into the CNN model. To this end, we propose a neural network model that makes use of automatically learned features (by different CNN layers) together with manually crafted ones (via domain knowledge), such as words, part-of-speech tags, chunks, named entities, and dependency graph of sentences. Such a combination in feature engineering has been shown to be effective in other NLP tasks also (e.g. (Shimaoka et al., 2017)).

Furthermore, we propose a multichannel CNN, a model that was suggested to capture different “views” of input data. In the image processing, (Krizhevsky et al., 2012) applied different RGB (red, green, blue) channels to color images. In NLP research, such models often use separate channels for different word embeddings (Yin and Schütze, 2015; Shi et al., 2016). For example, one could have separate channels for different word embeddings (Quan et al., 2016), or have one channel that is kept static throughout training and the other that is fine-tuned via backpropagation (Kim, 2014). Unlike these studies, we utilize the head of the word in a sentence as a separate channel.

## 3 CNN Model

### 3.1 Model Architecture Overview

Figure 1 illustrates the overview of our model, which takes a complete sentence with mentioned entities as input and outputs a probability vector (two elements) corresponding to whether there is a relation between the two entities. Our model

mainly consists of three layers: a multichannel embedding layer, a convolution layer, and a fully-connected layer.

## 3.2 Embedding Layer

In our model, as shown in Figure 1, each word in a sentence is represented by concatenating its word embedding, part-of-speech, chunk, named entity, dependency, and position features.

### 3.2.1 Word embedding

Word embedding is a language modeling techniques where words from the vocabulary are mapped to vectors of real numbers. It has been shown to boost the performance in NLP tasks. In this paper, we used pre-trained word embedding vectors (Pyysalo et al., 2013) learned on PubMed articles using the word2vec tool (Mikolov et al., 2013). The dimensionality of word vectors is 200.

### 3.2.2 Part-of-speech

We used the part-of-speech (POS) feature to extend the word embedding. Similar to (Zhao et al., 2016b), we divided POS into eight groups. Then each group is mapped to an eight-bit binary vector. In this way, the dimensionality of the POS feature is 8.

### 3.2.3 Chunk

We used the chunk tags obtained from Genia Tagger for each word (Tsuruoka and Tsujii, 2005). We encoded the chunk features using a one-hot scheme. The dimensionality of chunk tags is 18.

### 3.2.4 Named entity

To generalize the model, we used four types of named entity encodings for each word. The named entities were provided as input by the task data. In one PPI instance, the types of two proteins of interest are PROT1 and PROT2 respectively. The type of other proteins is PROT, and the type of other words is O. If a protein mention spans multiple words, we marked each word with the same type (we did not use a scheme such as IOB). The dimensionality of named entity is thus 4.

### 3.2.5 Dependency

To add the dependency information of each word, we used the label of “incoming” edge of that word in the dependency graph. Take the sentence from Figure 2 as an example, the dependency of “ARFTS” is “nsubj” and the dependency

of “binds” is “ROOT”. We encoded the dependency features using a one-hot scheme, and their dimensionality is 101.

### 3.2.6 Position feature

In this work, we consider the relationship of two protein mentions in a sentence. Thus, we used the position feature proposed in (Sahu et al., 2016), which consists of two relative distances,  $d1$  and  $d2$ , for representing the distances of the current word to PROT1 and PROT2 respectively. For example in Figure 2, the relative distances of the word “binds” to PROT1 (“ARFTs”) and PROT2 (“XIAP-BIR3”) are 2 and -6, respectively. Same as in Table S4 of (Zhao et al., 2016b), both  $d1$  and  $d2$  are non-linearly mapped to a ten-bit binary vector, where the first bit stands for the sign and the remaining bits for the distance.

## 3.3 Multichannel Embedding Layer

A novel aspect of McDepCNN is to add the “head” word representation of each word as the second channel of the embedding layer. For example, the second channel of the sentence in Figure 2 is “binds binds ROOT binds domain domain binds domain” as shown in Figure 1. There are several advantages of using the “head” of a word as a separate channel.

First, it intuitively incorporates the dependency graph structure into the CNN model. Compared with (Hua and Quan, 2016) which used the shortest path between two entities as the sole input for CNN, our model does not discard information outside the scope of two entities. Such information was reported to be useful (Zhou et al., 2007). Compared with (Zhao et al., 2016b) which used the shortest path as a bag-of-words sparse 0-1 vector, our model intuitively reflects the syntactic structure of the dependencies of the input sentence.

Second, together with convolution, our model can better capture longer distance dependencies than the sliding window size. As shown in Figure 2, the second channel of McDepCNN breaks the dependency graph structure into structural  $\langle \text{head word}, \text{child word} \rangle$  pairs where each word is a modifier of its previous word. In this way, it reflects the skeleton of a constituent where the second channel shadows the detailed information of all sub-constituents in the first channel. From the perspective of the sentence string, the second channel is similar to a gapped  $n$ -gram or a skipped




Figure 1: Overview of the CNN model.




Figure 2: Dependency graph.

$n$ -gram where the skipped words are based on the structure of the sentence.

### 3.4 Convolution

We applied convolution to input sentences to combine two channels and get local features (Collobert et al., 2011). Consider  $x_1, \dots, x_n$  to be the sequence of word representations in a sentence where

$$x_i = E_{word} \oplus \dots \oplus E_{poistion}, i = 1, \dots, n \quad (1)$$

Here  $\oplus$  is concatenation operation so  $x^i \in \mathbb{R}^d$  is the embedding vector for the  $i$ th word with the dimensionality  $d$ . Let  $x_{i:i+k-1}^c$  represent a window of size  $k$  in the sentence for channel  $c$ . Then the output sequence of the convolution layer is

$$con_i = f\left(\sum_c w_k^c x_{i:i+k-1}^c + b_k\right) \quad (2)$$

where  $f$  is a rectify linear unit (ReLU) function and  $b_k$  is the biased term. Both  $w_k^c$  and  $b_k$  are the learning parameters.

1-max pooling was then performed over each map, i.e., the largest number from each feature map was recorded. In this way, we obtained fixed length global features for the whole sentence. The underlying intuition is to consider only the most useful feature from the entire sentence.

$$m_k = \max_{1 \leq i \leq n-k+1} (con_i) \quad (3)$$

### 3.5 Fully Connected Layer with Softmax

To make a classifier over extracted global features, we first applied a fully connected layer to the feature vectors of multichannel obtained above.

$$O = w_o(m_3 \oplus m_5 \oplus m_7) + b_o \quad (4)$$

The final softmax then receives this vector  $O$  as input and uses it to classify the PPI; here we assume binary classification for the PPI task and hence depict two possible output states.

$$p(ppi|x, \theta) = \frac{e^{O_{ppi}}}{e^{O_{ppi}} + e^{O_{other}}} \quad (5)$$

Here,  $\theta$  is a vector of the hyper-parameters of the model, such as  $w_k^c$ ,  $b_k$ ,  $w_o$ , and  $b_o$ . Further, we used dropout technique in the output of the max pooling layer for regularization (Srivastava et al., 2014). This prevented our method from overfitting by randomly “dropping” with probability  $(1 - p)$  neurons during each forward/backward pass while training.

### 3.6 Training

To train the parameters, we used the log-likelihood of parameters on a mini-batch training with a batch size of  $m$ . We use the Adam algorithm to optimize the loss function (Kingma and Ba, 2015).

$$J(\theta) = \sum_m p(ppi^{(m)}|x^{(m)}, \theta) \quad (6)$$

### 3.7 Experimental setup

For our experiments, we used the Genia Tagger to obtain the part-of-speech, chunk tags, and named entities of each word (Tsuruoka and Tsujii, 2005). We parsed each sentence using the Bllip parser with the biomedical model (Charniak, 2000; McClosky, 2009). The universal dependencies were then obtained by applying the Stanford dependencies converter on the parse tree with the *CCProcessed* and *Universal* options (De Marneffe et al., 2014).

We implemented the model using TensorFlow (Abadi et al., 2016). All trainable variables were initialized using the Xavier algorithm (Glorot and Bengio, 2010). We set the maximum sentence length to 160. That is, longer sentences were pruned, and shorter sentences were padded with zeros. We set the learning rate to be 0.0007 and the dropping probability 0.5. During the training, we ran 250 epochs of all the training examples. For each epoch, we randomized the training examples and conducted a mini-batch training with a batch size of 128 ( $m = 128$ ).

In this paper, we experimented with three window sizes: 3, 5 and 7, each of which has 400 filters. Every filter performs convolution on the sentence matrix and generates variable-length feature maps. We got the best results using the single window of size 3 (see Section 4.2)

## 4 Results and Discussion

### 4.1 Data

We evaluated McDepCNN on two benchmarking PPI corpora, AIMed (Bunescu et al., 2005) and BioInfer (Pyysalo et al., 2007). These two corpora have different sizes (Table 1) and vary slightly in their definition of PPI (Pyysalo et al., 2008).

Tikk et al. (2010) conducted a comparison of a variety of PPI extraction systems on these two corpora<sup>1</sup>. In order to compare, we followed their

<sup>1</sup><http://mars.cs.utu.fi/PPICorpora>

Table 1: Statistics of the corpora.

Corpus	Sentences	# Positives	# Negatives
AIMed	1,955	1,000	4,834
BioInfer	1,100	2,534	7,132

experimental setup to evaluate our methods: self-interactions were excluded from the corpora and 10-fold cross-validation (CV) was performed.

### 4.2 Results and discussion

Our system performance, as measured by Precision, Recall, and F1-score, is shown in Table 2. To compare, we also include the results published in (Tikk et al., 2010; Peng et al., 2015; Van Landeghem et al., 2008; Fundel et al., 2007). Row 2 reports the results of the previous best deep learning system on these two corpora. Rows 3 and 4 report the results of two previous best single kernel-based methods, an APG kernel (Airola et al., 2008; Tikk et al., 2010) and an edit kernel (Peng et al., 2015). Rows 5-6 report the results of two rule-based systems. As can be seen, McDepCNN achieved the highest results in both precision and overall F1-score on both datasets.

Note that we did not compare our results with two recent deep-learning approaches (Hua and Quan, 2016; Quan et al., 2016). This is because unlike other previous studies, they artificially removed sentences that cannot be parsed and discarded pairs which are in a coordinate structure. Thus, our results are not directly comparable with theirs. Neither did we compare our method with (Miwa et al., 2009b) because they combined, in a rich vector, analysis from different parsers and the output of multiple kernels.

To further test the generalizability of our method, we conducted the cross-corpus experiments where we trained the model on one corpus and tested it on the other (Table 3). Here we compared our results with the shallow linguistic model which is reported as the best kernel-based method in (Tikk et al., 2013).

The cross-corpus results show that McDepCNN achieved 24.4% improvement in F-score when trained on BioInfer and tested on AIMed, and 18.2% vice versa.

To better understand the advantages of McDepCNN over kernel-based methods, we followed the lead of (Tikk et al., 2013) to compare the method performance on some known “difficult”

Table 2: Evaluation results. Performance is reported in terms of Precision, Recall, and F1-score.

Method	AIMed			BioInfer		
	P	R	F	P	R	F
1 McDepCNN	67.3	60.1	<b>63.5</b>	62.7	68.2	<b>65.3</b>
2 Deep neural network (Zhao et al., 2016a)	51.5	63.4	56.1	53.9	72.9	61.6
3 All-path graph kernel (Tikk et al., 2010)	49.2	64.6	55.3	53.3	70.1	60.0
4 Edit kernel (Peng et al., 2015)	65.3	57.3	61.1	59.9	57.6	58.7
5 Rich-feature (Van Landeghem et al., 2008)	49.0	44.0	46.0	–	–	–
6 RelEx (Fundel et al., 2007)	40.0	50.0	44.0	39.0	45.0	41.0

Table 3: Cross-corpus results. Performance is reported in terms of Precision, Recall, and F1-score.

Method	Training corpus	AIMed			BioInfer		
		P	R	F	P	R	F
McDepCNN	AIMed	–	–	–	39.5	61.4	<b>48.0</b>
	BioInfer	40.1	65.9	<b>49.9</b>	–	–	–
Shallow linguistic (Tikk et al., 2010)	AIMed	–	–	–	29.2	66.8	40.6
	BioInfer	76.8	27.2	41.5	–	–	–

Table 4: Instances that are the most difficult to classify correctly by the collection of kernels using cross-validation (Tikk et al., 2013).

Corpus	Positive difficult	Negative difficult
AIMed	61	184
BioInfer	111	295

instances in AIMed and BioInfer. This subset of difficult instances is defined as 10% of all pairs with the least number of 14 kernels being able to classify correctly (Table 4).

Table 5 shows the comparisons between McDepCNN and kernel-based methods on difficult instances. The results of McDepCNN were obtained from the difficult instances combined from AIMed and BioInfer (172 positives and 479 negatives). And the results of APG, Edit, and SL were obtained from AIMed, BioInfer, HPRD50, IEPA, and LLL (190 positives and 521 negatives) (Tikk et al., 2013). While the input datasets are different, our outcomes are remarkably higher than the prior studies. The results show that McDepCNN achieves 17.3% in F1-score on difficult instances which is more than three times better than other kernels. Since there are no examples of difficult instances that could not be classified correctly by at least one of the 14 kernel methods, below, we only list some examples that McDepCNN can classify correctly.

1. Immunoprecipitation experiments further re-

veal that the fully assembled receptor complex is composed of two **IL-6**<sub>PROT1</sub>, two **IL-6R alpha**<sub>PROT2</sub>, and two gp130 molecules.

2. The phagocyte NADPH oxidase is a complex of membrane cytochrome b558 (comprised of subunits p22-phox and gp91-phox) and three cytosol proteins (**p47-phox**<sub>PROT1</sub>, p67-phox, and p21rac) that translocate to membrane and bind to **cytochrome b558**<sub>PROT2</sub>.

Together with the conclusions in (Tikk et al., 2013), “positive pairs are more difficult to classify in longer sentences” and “most of the analyzed classifiers fail to capture the characteristics of rare positive pairs in longer sentences”, this comparison suggests that McDepCNN is probably capable of better capturing long distance features from the sentence and are more generalizable than kernel methods.

Finally, Table 6 compares the effects of different parts in McDepCNN. Here we tested McDepCNN using 10-fold of AIMed. Row 1 used a single window with the length of 3, row 2 used two windows, and row 3 used three windows. The reduced performance indicate that adding more windows did not improve the model. This is partially because the multichannel in McDepCNN has captured good context features for PPI. Second, we used the single channel and retrained the model with window size 3. The performance then dropped 1.1%. The results underscore the effectiveness of using the head word as a separate chan-

Table 5: Comparisons on the difficult instances with CV evaluation. Performance is reported in terms of Precision, Recall, and F1-score\*.

Method	P	R	F
McDepCNN	14.0	22.7	<b>17.3</b>
All-path graph kernel	4.3	7.9	5.5
Edit kernel	4.8	5.8	5.3
Shallow linguistic	3.6	7.9	4.9

\* The results of McDepCNN were obtained on the difficult instances combined from AIMed and BioInfer (172 positives and 479 negatives). The results of others (Tikk et al., 2013) were obtained from AIMed, BioInfer, HPRD50, IEPA, and LLL (190 positives and 521 negatives).

Table 6: Contributions of different parts in McDepCNN. Performance is reported in terms of Precision, Recall, and F1-score.

Method	P	R	F	$\Delta$
window = 3	67.3	60.1	63.5	
window = [3,5]	60.9	62.4	61.6	(1.9)
window = [3,5,7]	61.7	61.9	61.8	(1.7)
Single channel	62.8	62.3	62.6	(1.1)

nel in CNN.

## 5 Conclusion

In this paper, we describe a multichannel dependency-based convolutional neural network for the sentence-based PPI task. Experiments on two benchmarking corpora demonstrate that the proposed model outperformed the current deep learning model and single feature-based or kernel-based models. Further analysis suggests that our model is substantially more generalizable across different datasets. Utilizing the dependency structure of sentences as a separated channel also enables the model to capture global information more effectively.

In the future, we would like to investigate how to assemble different resources into our model, similar to what has been done to rich-feature-based methods (Miwa et al., 2009b) where the current best performance was reported (F-score of 64.0% (AIMed) and 66.7% (BioInfer)). We are also interested in extending the method to PPIs beyond the sentence boundary. Finally, we would like to test and generalize this approach to other biomedical relations such as chemical-disease relations (Wei et al., 2016).

## Acknowledgments

This work was supported by the Intramural Research Programs of the National Institutes of Health, National Library of Medicine. We are also grateful to Robert Leaman for the helpful discussion.

## References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467* <https://arxiv.org/abs/1603.04467>.
- Antti Airola, Sampo Pyysalo, Jari Björne, Tapio Pahikkala, Filip Ginter, and Tapio Salakoski. 2008. All-paths graph kernel for protein-protein interaction extraction with evaluation of cross-corpus learning. *BMC Bioinformatics* 9(Suppl 11):1–12. <https://doi.org/10.1186/1471-2105-9-s11-s2>.
- Quoc-Chinh Bui, Sophia Katrenko, and Peter M. A. Sloot. 2011. A hybrid approach to extract protein-protein interactions. *Bioinformatics* 27(2):259–265. <https://doi.org/10.1186/1471-2105-9-s11-s2>.
- Razvan C. Bunescu, Ruifang Ge, Rohit J. Kate, Edward M. Marcotte, Raymond J. Mooney, Arun K. Ramani, and Yuk Wah Wong. 2005. Comparative experiments on learning information extractors for proteins and their interactions. *Artificial Intelligence in Medicine* 33(2):139–155. <https://doi.org/10.1016/j.artmed.2004.07.016>.
- Eugene Charniak. 2000. A maximum-entropy-inspired parser. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*. pages 132–139. <http://dl.acm.org/citation.cfm?id=974323>.
- Faisal Md. Chowdhury, Alberto Lavelli, and Alessandro Moschitti. 2011. A study on dependency tree kernels for automatic extraction of protein-protein interaction. In *Proceedings of BioNLP 2011 Workshop*. pages 124–133. <http://aclweb.org/anthology/W11-0216>.
- Michael Collins and Nigel Duffy. 2002. New ranking algorithms for parsing and tagging: Kernels over discrete structures, and the voted perceptron. In *Proceedings of the 40th annual meeting on association for computational linguistics (ACL)*. Association for Computational Linguistics, pages 263–270. <http://aclweb.org/anthology/P02-1034>.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12(Aug):2493–2537.

- Marie-Catherine De Marneffe, Timothy Dozat, Natalia Silveira, Katri Haverinen, Filip Ginter, Joakim Nivre, and Christopher D Manning. 2014. Universal Stanford dependencies: A cross-linguistic typology. In *Proceedings of 9th International Conference on Language Resources and Evaluation (LREC)*. volume 14, pages 4585–4592.
- Günes Erkan, Arzucan Özgür, and Dragomir R Radev. 2007. Semi-supervised classification for extracting protein interaction sentences using dependency parsing. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*. volume 7, pages 228–237. <http://aclweb.org/anthology/D07-1024>.
- Katrin Fundel, Robert Küffner, and Ralf Zimmer. 2007. RelEx-relation extraction using dependency parse trees. *Bioinformatics* 23(3):365–371. <https://doi.org/10.1093/bioinformatics/btl616>.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*. volume 9, pages 249–256.
- Lei Hua and Chanqin Quan. 2016. A shortest dependency path based convolutional neural network for protein-protein relation extraction. *BioMed Research International* 2016. <https://doi.org/10.1155/2016/8479587>.
- Chung-Chi Huang and Zhiyong Lu. 2015. Community challenges in biomedical text mining over 10 years: success, failure and the future. *Briefings in Bioinformatics* 17(1):132–144. <https://doi.org/10.1093/bib/bbv024>.
- Seonho Kim, Juntae Yoon, Jihoon Yang, and Seog Park. 2010. Walk-weighted subsequence kernels for protein-protein interaction extraction. *BMC Bioinformatics* 11(107):1–21. <https://doi.org/10.1186/1471-2105-11-107>.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pages 1746–1751. <http://aclweb.org/anthology/D14-1181>.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*. pages 1–15. <https://arxiv.org/abs/1412.6980>.
- Martin Krallinger, Florian Leitner, Carlos Rodriguez-Penagos, Alfonso Valencia, et al. 2008. Overview of the protein-protein interaction annotation extraction task of BioCreative II. *Genome Biology* 9(Suppl 2):S4. <https://doi.org/10.1186/gb-2008-9-s2-s4>.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. pages 1097–1105. <http://dl.acm.org/citation.cfm?id=2999257>.
- John Lafferty, Andrew McCallum, and Fernando Pereira. 2001. Conditional random fields: probabilistic models for segmenting and labeling sequence data. In *Proceedings of the International Conference on Machine Learning (ICML 01)*. pages 282–289. <http://dl.acm.org/citation.cfm?id=655813>.
- Robert Leaman and Zhiyong Lu. 2016. TaggerOne: joint named entity recognition and normalization with semi-markov models. *Bioinformatics* 32(18):2839–2846. <https://doi.org/10.1093/bioinformatics/btw343>.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324. <https://doi.org/10.1109/5.726791>.
- Christopher D Manning. 2015. Computational linguistics and deep learning. *Computational Linguistics* 41(4):701–707. [https://doi.org/10.1162/coli\\_a.00239](https://doi.org/10.1162/coli_a.00239).
- David McClosky. 2009. *Any domain parsing: automatic domain adaptation for natural language parsing*. Thesis, Department of Computer Science, Brown University.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. pages 3111–3119. <http://dl.acm.org/citation.cfm?id=2999959>.
- Makoto Miwa, Rune Sætre, Yusuke Miyao, and Junichi Tsujii. 2009a. Protein-protein interaction extraction by leveraging multiple kernels and parsers. *International Journal of Medical Informatics* 78(12):e39–46. <https://doi.org/10.1016/j.ijmedinf.2009.04.010>.
- Makoto Miwa, Rune Sætre, Yusuke Miyao, and Junichi Tsujii. 2009b. A rich feature vector for protein-protein interaction extraction from multiple corpora. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*. volume 1, pages 121–130. <http://aclweb.org/anthology/D09-1013>.
- Yifan Peng, Samir Gupta, Cathy H Wu, and K Vijay-Shanker. 2015. An extended dependency graph for relation extraction in biomedical texts. In *Proceedings of BioNLP 2015 Workshop*. pages 21–30. <http://aclweb.org/anthology/W15-3803>.
- Yifan Peng, Chih-Hsuan Wei, and Zhiyong Lu. 2016. Improving chemical disease relation extraction with rich features and weakly labeled data. *Journal of Cheminformatics* 8(53):1–12. <https://doi.org/10.1186/s13321-016-0165-z>.



- Sampo Pyysalo, Antti Airola, Juho Heimonen, Jari Björne, Filip Ginter, and Tapio Salakoski. 2008. Comparative analysis of five protein-protein interaction corpora. *BMC Bioinformatics* 9(Suppl 3):1–11. <https://doi.org/10.1186/1471-2105-9-S3-S6>.
- Sampo Pyysalo, Filip Ginter, Juho Heimonen, Jari Björne, Jorma Boberg, Jouni Järvinen, and Tapio Salakoski. 2007. BioInfer: a corpus for information extraction in the biomedical domain. *BMC Bioinformatics* 8(50):1–24. <https://doi.org/10.1186/1471-2105-8-50>.
- Sampo Pyysalo, Filip Ginter, Hans Moen, Tapio Salakoski, and Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing. In *International Symposium on Languages in Biology and Medicine (LBM)*. pages 39–44.
- Chanqin Quan, Lei Hua, Xiao Sun, and Wenjun Bai. 2016. Multichannel convolutional neural network for biological relation extraction. *BioMed Research International* 2016:1–10. <https://doi.org/10.1155/2016/1850404>.
- Sunil Kumar Sahu, Ashish Anand, Krishnadev Oruganty, and Mahanandeeswar Gattu. 2016. Relation extraction from clinical texts using domain invariant convolutional neural network. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*. pages 206–215. <https://aclweb.org/anthology/W16-2928>.
- Hongjie Shi, Takashi Ushio, Mitsuru Endo, Katsuyoshi Yamagami, and Noriaki Horii. 2016. A multichannel convolutional neural network for cross-language dialog state tracking. In *2016 IEEE Workshop on Spoken Language Technology Workshop (SLT)*. IEEE, pages 559–564. <https://arxiv.org/abs/1701.06247>.
- Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui, and Sebastian Riedel. 2017. Neural architectures for fine-grained entity type classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. pages 1271–1280. <http://aclweb.org/anthology/E17-1119>.
- Ayush Singhal, Michael Simmons, and Zhiyong Lu. 2016. Text mining genotype-phenotype relationships from biomedical literature for database curation and precision medicine. *PLOS Computational Biology* 12(11):e1005017. <https://doi.org/10.1371/journal.pcbi.1005017>.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1):1929–1958.
- Domonkos Tikk, Illés Solt, Philippe Thomas, and Ulf Leser. 2013. A detailed error analysis of 13 kernel methods for protein-protein interaction extraction. *BMC Bioinformatics* 14(1):12. <https://doi.org/10.1186/1471-2105-14-12>.
- Domonkos Tikk, Philippe Thomas, Peter Palaga, Jörg Hakenberg, and Ulf Leser. 2010. A comprehensive benchmark of kernel methods to extract protein-protein interactions from literature. *PLoS Computational Biology* 6(7):1–19. <https://doi.org/10.1371/journal.pcbi.1000837>.
- Yoshimasa Tsuruoka and Jun’ichi Tsujii. 2005. Bidirectional inference with the easiest-first strategy for tagging sequence data. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT-EMNLP)*. HLT ’05, pages 467–474. <https://doi.org/10.3115/1220575.1220634>.
- Sofie Van Landeghem, Yvan Saeys, Bernard De Baets, and Yves Van de Peer. 2008. Extracting protein-protein interactions from text using rich feature vectors and feature selection. In *Proceedings of the 3rd International Symposium on Semantic Mining in Biomedicine (SMBM)*. pages 77–84.
- Vladimir N. Vapnik. 1995. *The nature of statistical learning theory*. Springer-Verlag New York, Inc.
- Chih-Hsuan Wei, Yifan Peng, Robert Leaman, Allan Peter Davis, Carolyn J. Mattingly, Jiao Li, Thomas C. Wiegers, and Zhiyong Lu. 2016. Assessing the state of the art in biomedical relation extraction: overview of the BioCreative V chemical-disease relation (CDR) task. *Database* pages 1–8. <https://doi.org/10.1093/database/baw032>.
- Wenpeng Yin and Hinrich Schütze. 2015. Multichannel variable-size convolution for sentence classification. In *Proceedings of the 19th Conference on Computational Language Learning (CoNLL)*. Association for Computational Linguistics, pages 204–214. <http://aclweb.org/anthology/K15-1021>.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, Jun Zhao, et al. 2014. Relation classification via Convolutional Deep Neural Network. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING)*. pages 2335–2344. <http://aclweb.org/anthology/C14-1220>.
- Zhehuan Zhao, Zhihao Yang, Hongfei Lin, Jian Wang, and Song Gao. 2016a. A protein-protein interaction extraction approach based on deep neural network. *International Journal of Data Mining and Bioinformatics* 15(2):145–164. <https://doi.org/10.1504/ijdbmb.2016.076534>.
- Zhehuan Zhao, Zhihao Yang, Ling Luo, Hongfei Lin, and Jian Wang. 2016b. Drug drug interaction extraction from biomedical literature using syntax convolutional neural network. *Bioinformatics* 32(22):3444–3453. <https://doi.org/10.1093/bioinformatics/btw486>.

Guodong Zhou, Min Zhang, Dong Hong, and Ji Qiaoming Zhu. 2007. Tree kernel-based relation extraction with context-sensitive structured parse tree information. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 728–736. <http://aclweb.org/anthology/D07-1076>.