

# A Walk-based Model on Entity Graphs for Relation Extraction

Fenia Christopoulou<sup>1,3</sup>, Makoto Miwa<sup>2,3</sup>, Sophia Ananiadou<sup>1,3</sup>

<sup>1</sup>National Centre for Text Mining

<sup>1</sup>School of Computer Science, The University of Manchester, United Kingdom

<sup>2</sup>Toyota Technological Institute, Nagoya, 468-8511, Japan

<sup>3</sup>Artificial Intelligence Research Center (AIRC),

National Institute of Advanced Industrial Science and Technology (AIST), Japan

{efstathia.christopoulou, sophia.ananiadou}@manchester.ac.uk

makoto-miwa@toyota-ti.ac.jp

## Abstract

We present a novel graph-based neural network model for relation extraction. Our model treats multiple pairs in a sentence simultaneously and considers interactions among them. All the entities in a sentence are placed as nodes in a fully-connected graph structure. The edges are represented with position-aware contexts around the entity pairs. In order to consider different relation paths between two entities, we construct up to  $l$ -length walks between each pair. The resulting walks are merged and iteratively used to update the edge representations into longer walks representations. We show that the model achieves performance comparable to the state-of-the-art systems on the ACE 2005 dataset without using any external tools.

## 1 Introduction

Relation extraction (RE) is a task of identifying typed relations between known entity mentions in a sentence. Most existing RE models treat each relation in a sentence individually (Miwa and Bansal, 2016; Nguyen and Grishman, 2015). However, a sentence typically contains multiple relations between entity mentions. RE models need to consider these pairs simultaneously to model the dependencies among them. The relation between a pair of interest (namely “target” pair) can be influenced by other pairs in the same sentence. The example illustrated in Figure 1 explains this phenomenon. The relation between the pair of interest *Toefting* and *capital*, can be extracted directly from the target entities or indirectly by incorporating information from other related pairs in the sentence. The person entity (PER) *Toefting* is directly related with *teammates* through the




Figure 1: Relation examples from ACE (Automatic Content Extraction) 2005 dataset (Dodington et al., 2004).

preposition *with*. Similarly, *teammates* is directly related with the geopolitical entity (GPE) *capital* through the preposition *in*. *Toefting* and *capital* can be directly related through *in* or indirectly related through *teammates*. Substantially, the path from *Toefting* to *teammates* to *capital* can additionally support the relation between *Toefting* and *capital*.

Multiple relations in a sentence between entity mentions can be represented as a graph. Neural graph-based models have shown significant improvement in modelling graphs over traditional feature-based approaches in several tasks. They are most commonly applied on knowledge graphs (KG) for knowledge graph completion (Jiang et al., 2017) and the creation of knowledge graph embeddings (Wang et al., 2017; Shi and Weninger, 2017). These models rely on paths between existing relations in order to infer new associations between entities in KGs. However, for relation extraction from a sentence, related pairs are not predefined and consequently all entity pairs need to be considered to extract relations. In addition, state-of-the-art RE models sometimes depend on external syntactic tools to build the shortest dependency path (SDP) between two entities in a sentence (Xu et al., 2015; Miwa and Bansal, 2016). This dependence on external tools leads to domain dependent models.

In this study, we propose a neural relation extraction model based on an entity graph, where entity mentions constitute the nodes and directed edges correspond to ordered pairs of entity mentions. The overview of the model is shown in Figure 2. We initialize the representation of an edge (an ordered pair of entity mentions) from the representations of the entity mentions and their context. The context representation is achieved by employing an attention mechanism on context words. We then use an iterative process to aggregate up-to  $l$ -length walk representations between two entities into a single representation, which corresponds to the final representation of the edge.

The contributions of our model can be summarized as follows:

- We propose a graph walk based neural model that considers multiple entity pairs in relation extraction from a sentence.
- We propose an iterative algorithm to form a single representation for up-to  $l$ -length walks between the entities of a pair.
- We show that our model performs comparably to the state-of-the-art without the use of external syntactic tools.

## 2 Proposed Walk-based Model

The goal of the RE task is given a sentence, entity mentions and their semantic types, to extract and classify all related entity pairs (target pairs) in the sentence. The proposed model consists of five stacked layers: embedding layer, BLSTM Layer, edge representation layer, walk aggregation layer and finally a classification layer.

As shown in Figure 2, the model receives word representations and produces simultaneously a representation for each pair in the sentence. These representations combine the target pair, its context words, their relative positions to the pair entities and walks between them. During classification they are used to predict the relation type of each pair.

### 2.1 Embedding Layer

The embedding layer involves the creation of  $n_w$ ,  $n_t$ ,  $n_p$ -dimensional vectors which are assigned to words, semantic entity types and relative positions to the target pairs. We map all words and semantic types into real-valued vectors  $\mathbf{w}$  and  $\mathbf{t}$  respectively. Relative positions to target entities are created based on the position of words in the sen-

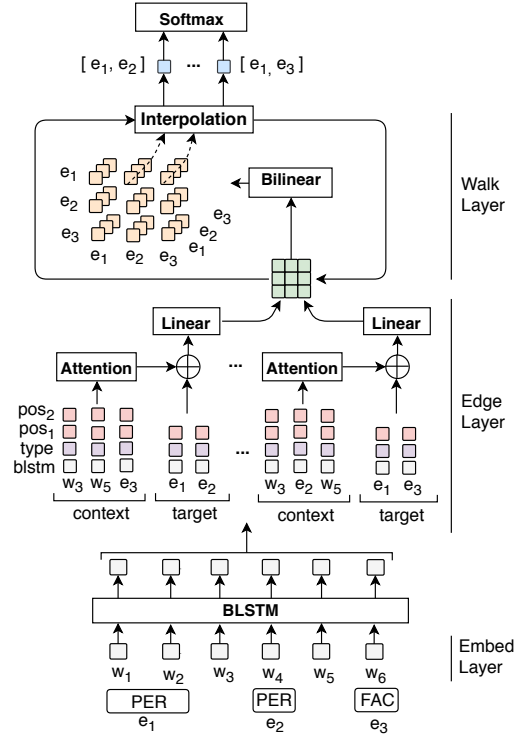


Figure 2: Overview of the walk-based model.

tence. In the example of Figure 1, the relative position of *teammates* to *capital* is  $-3$  and the relative position of *teammates* to *Toefling* is  $+16$ . We embed real-valued vectors  $\mathbf{p}$  to these positions.

### 2.2 Bidirectional LSTM Layer

The word representations of each sentence are fed into a Bidirectional Long-short Term Memory (BLSTM) layer, which encodes the context representation for every word. The BLSTM outputs new word-level representations  $\mathbf{h}$  (Hochreiter and Schmidhuber, 1997) that consider the sequence of words.

We avoid encoding target pair-dependent information in this BLSTM layer. This has two advantages: (i) the computational cost is reduced as this computation is repeated based on the number of sentences instead of the number of pairs, (ii) we can share the sequence layer among the pairs of a sentence. The second advantage is particularly important as it enables the model to indirectly learn hidden dependencies between the related pairs in the same sentence.

For each word  $t$  in the sentence, we concatenate the two representations from left-to-right and right-to-left pass of the LSTM into a  $n_e$ -dimensional vector,  $\mathbf{e}_t = [\vec{\mathbf{h}}_t; \overleftarrow{\mathbf{h}}_t]$ .

### 2.3 Edge Representation Layer

The output word representations of the BLSTM are further divided into two parts: (i) target pair representations and (ii) target pair-specific context representations. The context of a target pair can be expressed as all words in the sentence that are not part of the entity mentions. We represent a related pair as described below.

A target pair contains two entities  $e_i$  and  $e_j$ . If an entity consists of  $N$  words, we create its BLSTM representation as the average of the BLSTM representations of the corresponding words,  $\mathbf{e} = \frac{1}{|I|} \sum_{i \in I} \mathbf{e}_i$ , where  $I$  is a set with the word indices inside entity  $e$ .

We first create a representation for each pair entity and then we construct the representation for the context of the pair. The representation of an entity  $e_i$  is the concatenation of its BLSTM representation  $\mathbf{e}_i$ , the representation of its entity type  $\mathbf{t}_i$  and the representation of its relative position to entity  $e_j$ ,  $\mathbf{p}_{ij}$ . Similarly, for entity  $e_j$  we use its relative position to entity  $e_i$ ,  $\mathbf{p}_{ji}$ . Finally, the representations of the pair entities are as follows:  $\mathbf{v}_i = [\mathbf{e}_i; \mathbf{t}_i; \mathbf{p}_{ij}]$  and  $\mathbf{v}_j = [\mathbf{e}_j; \mathbf{t}_j; \mathbf{p}_{ji}]$ .

The next step involves the construction of the representation of the context for this pair. For each context word  $w_z$  of the target pair  $e_i, e_j$ , we concatenate its BLSTM representation  $\mathbf{e}_z$ , its semantic type representation  $\mathbf{t}_z$  and two relative position representations: to target entity  $e_i$ ,  $\mathbf{p}_{zi}$  and to target entity  $e_j$ ,  $\mathbf{p}_{zj}$ . The final representation for a context word  $w_z$  of a target pair is,  $\mathbf{v}_{ijz} = [\mathbf{e}_z; \mathbf{t}_z; \mathbf{p}_{zi}; \mathbf{p}_{zj}]$ . For a sentence, the context representations for all entity pairs can be expressed as a three-dimensional matrix  $\mathbf{C}$ , where rows and columns correspond to entities and the depth corresponds to the context words.

The context words representations of each target pair are then compiled into a single representation with an attention mechanism. Following the method proposed in Zhou et al. (2016), we calculate weights for the context words of the target-pair and compute their weighted average,

$$\begin{aligned} \mathbf{u} &= \mathbf{q}^\top \tanh(\mathbf{C}_{ij}), \\ \alpha &= \text{softmax}(\mathbf{u}), \\ \mathbf{c}_{ij} &= \mathbf{C}_{ij} \alpha^\top, \end{aligned} \quad (1)$$

where  $\mathbf{q} \in \mathbb{R}^{n_d}$ ,  $n_d = n_e + n_t + 2n_p$  denotes a trainable attention vector,  $\alpha$  is the attended weights vector and  $\mathbf{c}_{ij} \in \mathbb{R}^{n_d}$  is the context representation of the pair as resulted by the weighted

average. This attention mechanism is independent of the relation type. We leave relation-dependent attention as future work.

Finally, we concatenate the representations of the target entities and their context ( $\in \mathbb{R}^{n_m}$ ). We use a fully connected linear layer,  $\mathbf{W}_s \in \mathbb{R}^{n_m \times n_s}$  with  $n_s < n_m$  to reduce the dimensionality of the resulting vector. This corresponds to the representation of an edge or a one-length walk between nodes  $i$  and  $j$ :  $\mathbf{v}_{ij}^{(1)} = \mathbf{W}_s [\mathbf{v}_i; \mathbf{v}_j; \mathbf{c}_{ij}] \in \mathbb{R}^{n_s}$ .

### 2.4 Walk Aggregation Layer

Our main aim is to support the relation between an entity pair by using chains of intermediate relations between the pair entities. Thus, the goal of this layer is to generate a single representation for a finite number of different lengths walks between two target entities. To achieve this, we represent a sentence as a directed graph, where the entities constitute the graph nodes and edges correspond to the representation of the relation between the two nodes. The representation of one-length walk between a target pair  $\mathbf{v}_{ij}^{(1)}$ , serves as a building block in order to create and aggregate representations for one-to- $l$ -length walks between the pair. The walk-based algorithm can be seen as a two-step process: walk construction and walk aggregation. During the first step, two consecutive edges in the graph are combined using a modified bilinear transformation,

$$f(\mathbf{v}_{ik}^{(\lambda)}, \mathbf{v}_{kj}^{(\lambda)}) = \sigma \left( \mathbf{v}_{ik}^{(\lambda)} \odot (\mathbf{W}_b \mathbf{v}_{kj}^{(\lambda)}) \right), \quad (2)$$

where  $\mathbf{v}_{ij}^{(\lambda)} \in \mathbb{R}^{n_b}$  corresponds to walks representation of lengths one-to- $\lambda$  between entities  $e_i$  and  $e_j$ ,  $\odot$  represents element-wise multiplication,  $\sigma$  is the sigmoid non-linear function and  $\mathbf{W}_b \in \mathbb{R}^{n_b \times n_b}$  is a trainable weight matrix. This equation results in walks of lengths two-to- $2\lambda$ .

In the walk aggregation step, we linearly combine the initial walks (length one-to- $\lambda$ ) and the extended walks (length two-to- $2\lambda$ ),

$$\mathbf{v}_{ij}^{(2\lambda)} = \beta \mathbf{v}_{ij}^{(\lambda)} + (1 - \beta) \sum_{k \neq i, j} f(\mathbf{v}_{ik}^{(\lambda)}, \mathbf{v}_{kj}^{(\lambda)}), \quad (3)$$

where  $\beta$  is a weight that indicates the importance of the shorter walks. Overall, we create a representation for walks of length one-to-two using Equation (3) and  $\lambda = 1$ . We then create a representation for walks of length one-to-four by re-applying the

equation with  $\lambda = 2$ . We repeat this process until the desired maximum walk length is reached, which is equivalent to  $2\lambda = l$ .

## 2.5 Classification Layer

For the final layer of the network, we pass the resulted pair representation into a fully connected layer with a softmax function,

$$\mathbf{y} = \text{softmax}(\mathbf{W}_r \mathbf{v}_{ij}^{(l)} + \mathbf{b}_r), \quad (4)$$

where  $\mathbf{W}_r \in \mathbb{R}^{n_b \times n_r}$  is the weight matrix,  $n_r$  is the total number of relation types and  $b_r$  is the bias vector.

We use in total  $2r+1$  classes in order to consider both directions for every pair, i.e., left-to-right and right-to-left. The first argument appears first in a sentence in a left-to-right relation while the second argument appears first in a right-to-left relation. The additional class corresponds to non-related pairs, namely “no relation” class. We choose the most confident prediction for each direction and choose the positive and most confident prediction when the predictions contradict each other.

## 3 Experiments

### 3.1 Dataset

We evaluate the performance of our model on ACE 2005<sup>1</sup> for the task of relation extraction. ACE 2005 includes 7 entity types and 6 relation types between named entities. We follow the pre-processing described in [Miwa and Bansal \(2016\)](#).

### 3.2 Experimental Settings

We implemented our model using the Chainer library ([Tokui et al., 2015](#)).<sup>2</sup> The model was trained with Adam optimizer ([Kingma and Ba, 2015](#)). We initialized the word representations with existing pre-trained embeddings with dimensionality of 200.<sup>3</sup> Our model did not use any external tools except these embeddings.

The forget bias of the LSTM layer was initialized with a value equal to one following the work of [Jozefowicz et al. \(2015\)](#). We use a batchsize of 10 sentences and fix the pair representation dimensionality to 100. We use gradient clipping, dropout on the embedding and output layers and L2 regularization without regularizing the biases, to avoid

<sup>1</sup><https://catalog.ldc.upenn.edu/ldc2006t06>

<sup>2</sup><https://chainer.org/>

<sup>3</sup><https://github.com/tticoin/LSTM-ER>

Model	P	R	F1 (%)
SPTree	70.1	61.2	65.3
Baseline	72.5	53.3	61.4*
No walks $l = 1$	71.9	55.6	62.7
+ Walks $l = 2$	69.9	58.4	63.6 $\diamond$
+ Walks $l = 4$	69.7	59.5	64.2 $\diamond$
+ Walks $l = 8$	71.5	55.3	62.4

Table 1: Relation extraction performance on ACE 2005 test dataset. \* denotes significance at  $p < 0.05$  compared to SPTree,  $\diamond$  denotes significance at  $p < 0.05$  compared to the Baseline.

overfitting. We also incorporate early stopping with patience equal to five, to chose the number of training epochs and parameter averaging. We tune the model hyper-parameters on the respective development set using the RoBO Toolkit ([Klein et al., 2017](#)). Please refer to the supplementary material for the values.

We extract all possible pairs in a sentence based on the number of entities it contains. If a pair is not found in the corpus, it is assigned the “no relation” class. We report the micro precision, recall and F1 score following [Miwa and Bansal \(2016\)](#) and [Nguyen and Grishman \(2015\)](#).

## 4 Results

Table 1 illustrates the performance of our proposed model in comparison with *SPTree* system [Miwa and Bansal \(2016\)](#) on ACE 2005. We use the same data split with *SPTree* to compare with their model. We retrained their model with gold entities in order to compare the performances on the relation extraction task. The *Baseline* corresponds to a model that classifies relations by using only the representations of entities in a target pair.

As it can be observed from the table, the *Baseline* model achieves the lowest F1 score between the proposed models. By incorporating attention we can further improve the performance by 1.3 percent point (pp). The addition of 2-length walks further improves performance (0.9 pp). The best results among the proposed models are achieved for maximum 4-length walks. By using up-to 8-length walks the performance drops almost by 2 pp. We also compared our performance with [Nguyen and Grishman \(2015\) \(CNN\)](#) using their data split.<sup>4</sup> For the comparison, we applied our

<sup>4</sup>The authors kindly provided us with the data split.



# Entities	$l = 1$	$l = 2$	$l = 4$	$l = 8$
2	71.2	69.8	72.9	71.0
3	70.1	67.5	67.8	63.5*
[4, 6)	56.5	59.7	59.3	59.9
[6, 12)	59.2	64.2*	62.2	60.4
[12, 23)	54.7	59.3	62.3*	55.0

Table 2: Relation extraction performance (F1 %) on ACE 2005 development set for different number of entities. \* denotes significance at  $p < 0.05$  compared to  $l = 1$ .

best performing model ( $l = 4$ ).<sup>5</sup> The obtained performance is 65.8 / 58.4 / 61.9 in terms of P / R / F1 (%) respectively. In comparison with the performance of the *CNN* model, 71.5 / 53.9 / 61.3, we observe a large improvement in recall which results in 0.6 pp F1 increase.

We performed the Approximate Randomization test (Noreen, 1989) on the results. The best walks model has no statistically significant difference with the state-of-the-art *SPTree* model as in Table 1. This indicates that the proposed model can achieve comparable performance without any external syntactic tools.

Finally, we show the performance of the proposed model as a function of the number of entities in a sentence. Results in Table 2 reveal that for multi-pair sentences the model performs significantly better compared to the no-walks models, proving the effectiveness of the method. Additionally, it is observed that for more entity pairs, longer walks seem to be required. However, very long walks result to reduced performance ( $l = 8$ ).

## 5 Related Work

Traditionally, relation extraction approaches have incorporated a large variety of hand-crafted features to represent related entity pairs (Hermann and Blunsom, 2013; Miwa and Sasaki, 2014; Nguyen and Grishman, 2014; Gormley et al., 2015). Recent models instead employ neural network architectures and achieve state-of-the-art results without heavy feature engineering. Neural network techniques can be categorized into recurrent neural networks (RNNs) and convolutional neural networks (CNNs). The former is

<sup>5</sup>We kept the same parameters when we apply our model to the this data split. We did not remove any negative examples unlike the *CNN* model.

able to encode linguistic and syntactic properties of long word sequences, making them preferable for sequence-related tasks, e.g. natural language generation (Goyal et al., 2016), machine translation (Sutskever et al., 2014).

State-of-the-art systems have proved to achieve good performance on relation extraction using RNNs (Cai et al., 2016; Miwa and Bansal, 2016; Xu et al., 2016; Liu et al., 2015). Nevertheless, most approaches do not take into consideration the dependencies between relations in a single sentence (dos Santos et al., 2015; Nguyen and Grishman, 2015) and treat each pair separately. Current graph-based models are applied on knowledge graphs for distantly supervised relation extraction (Zeng et al., 2017). Graphs are defined on semantic types in their method, whereas we built entity-based graphs in sentences. Other approaches also treat multiple relations in a sentence (Gupta et al., 2016; Miwa and Sasaki, 2014; Li and Ji, 2014), but they fail to model long walks between entity mentions.

## 6 Conclusions

We proposed a novel neural network model for simultaneous sentence-level extraction of related pairs. Our model exploits target and context pair-specific representations and creates pair representations that encode up-to  $l$ -length walks between the entities of the pair. We compared our model with the state-of-the-art models and observed comparable performance on the ACE2005 dataset without any external syntactic tools. The characteristics of the proposed approach are summarized in three factors: the encoding of dependencies between relations, the ability to represent multiple walks in the form of vectors and the independence from external tools. Future work will aim at the construction of an end-to-end relation extraction system as well as application to different types of datasets.

## Acknowledgments

This research has been carried out with funding from AIRC/AIST, the James Elson Studentship Award, BBSRC grant BB/P025684/1 and MRC MR/N00583X/1. Results were obtained from a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO). We would finally like to thank the anonymous reviewers for their helpful comments.

## References

- Rui Cai, Xiaodong Zhang, and Houfeng Wang. 2016. Bidirectional recurrent convolutional neural network for relation classification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 756–765.
- George R Doddington, Alexis Mitchell, Mark A Przybocki, Lance A Ramshaw, Stephanie Strassel, and Ralph M Weischedel. 2004. The automatic content extraction (ace) program-tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC)*, volume 2, page 1.
- Matthew R. Gormley, Mo Yu, and Mark Dredze. 2015. Improved relation extraction with feature-rich compositional embedding models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1774–1784. Association for Computational Linguistics.
- Raghav Goyal, Marc Dymetman, and Eric Gaussier. 2016. Natural language generation through character-based rnns with finite-state prior knowledge. In *Proceedings of COLING, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1083–1092.
- Pankaj Gupta, Hinrich Schütze, and Bernt Andrassy. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *Proceedings of COLING, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2537–2547.
- Karl Moritz Hermann and Phil Blunsom. 2013. The role of syntax in vector space models of compositional semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 894–904.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Xiaotian Jiang, Quan Wang, Baoyuan Qi, Yongqin Qiu, Peng Li, and Bin Wang. 2017. Attentive path combination for knowledge graph completion. In *Asian Conference on Machine Learning*, pages 590–605.
- Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. 2015. An empirical exploration of recurrent network architectures. In *International Conference on Machine Learning*, pages 2342–2350.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization.
- A. Klein, S. Falkner, N. Mansur, and F. Hutter. 2017. Robo: A flexible and robust bayesian optimization framework in python. In *Proceedings of Workshop on Bayesian Optimization in the Conference on Neural Information Processing Systems (NIPS)*.
- Qi Li and Heng Ji. 2014. Incremental joint extraction of entity mentions and relations. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 402–412. Association for Computational Linguistics.
- Yang Liu, Furu Wei, Sujian Li, Heng Ji, Ming Zhou, and Houfeng WANG. 2015. A dependency-based neural network for relation classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 285–290. Association for Computational Linguistics.
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using lstms on sequences and tree structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1105–1116. Association for Computational Linguistics.
- Makoto Miwa and Yutaka Sasaki. 2014. Modeling joint entity and relation extraction with table representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1858–1869. Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2014. Employing word representations and regularization for domain adaptation of relation extraction. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 68–74.
- Thien Huu Nguyen and Ralph Grishman. 2015. Relation extraction: Perspective from convolutional neural networks. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 39–48. Association for Computational Linguistics.
- Eric W Noreen. 1989. *Computer-intensive methods for testing hypotheses*. Wiley New York.
- Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2015. Classifying relations by ranking with convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 626–634. Association for Computational Linguistics.
- Baoxu Shi and Tim Wenginger. 2017. Proje: Embedding projection for knowledge graph completion. In *Proceedings of AAAI Conference on Artificial Intelligence*, volume 17, pages 1236–1242.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.

Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. 2015. Chainer: a next-generation open source framework for deep learning. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in the twenty-ninth Annual Conference on Neural Information Processing Systems (NIPS)*, volume 5.

Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12):2724–2743.

Kun Xu, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2015. Semantic relation classification via convolutional neural networks with simple negative sampling. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 536–540. Association for Computational Linguistics.

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 COLING, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1461–1470.

Wenyuan Zeng, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2017. Incorporating relation paths in neural relation extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1768–1777. Association for Computational Linguistics.

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 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 207–212.

## A Hyper-parameter Settings

We tuned our proposed model using the RoBO toolkit (<https://github.com/automl/RoBO>). Table 3 provides the selected options we used for tuning the model.

Optimization Options	
Optimization method	Bohamiann
Maximizer	scipy
Acquisition function	log_ei
Number of iterations	50
Initial points	3

Table 3: Hyper-parameters optimization options.

The parameters that gave the best performance for the different models can be found in Tables 4a-4e.

Parameter	Baseline
Position dimension $n_p$	25
Type dimension $n_t$	15
LSTM dimension $n_e$	100
Input layer dropout	0.3
Output layer dropout	0.03
Learning rate	0.0018
Regularization	$3.2 \cdot 10^{-5}$
Gradient clipping	25.63

(a)

Parameter	$l = 1$
Position dimension $n_p$	25
Type dimension $n_t$	25
LSTM dimension $n_e$	100
Input layer dropout	0.13
Output layer dropout	0.38
Learning rate	0.0017
Regularization	$6.1 \cdot 10^{-5}$
Gradient clipping	30

(b)

Parameter	$l = 2$
Position dimension $n_p$	25
Type dimension $n_t$	20
LSTM dimension $n_e$	100
$\beta$	0.72
Input layer dropout	0.25
Output layer dropout	0.37
Learning rate	0.003
Regularization	0.0001
Gradient clipping	8.6

(c)

Parameter	$l = 4$
Position dimension $n_p$	25
Type dimension $n_t$	20
LSTM dimension $n_e$	100
$\beta$	0.77
Input layer dropout	0.11
Output layer dropout	0.32
Learning rate	0.002
Regularization	$5.7 \cdot 10^{-5}$
Gradient clipping	24.4

(d)

Parameter	$l = 8$
Position dimension $n_p$	25
Type dimension $n_t$	20
LSTM dimension $n_e$	100
$\beta$	0.88
Input layer dropout	0.49
Output layer dropout	0.36
Learning rate	0.001
Regularization	$1.88 \cdot 10^{-5}$
Gradient clipping	10.5

(e)

Table 4: Best hyper-parameters settings for proposed models.