

Syntactic Dependency Representations in Neural Relation Classification

Farhad Nooralahzadeh and Lilja Øvrelid

University of Oslo

Department of Informatics

{farhadno, liljao}@ifi.uio.no

Abstract

We investigate the use of different syntactic dependency representations in a neural relation classification task and compare the CoNLL, Stanford Basic and Universal Dependencies schemes. We further compare with a syntax-agnostic approach and perform an error analysis in order to gain a better understanding of the results.

1 Introduction

The neural advances in the field of NLP challenge long held assumptions regarding system architectures. The classical NLP systems, where components of increasing complexity are combined in a pipeline architecture are being challenged by end-to-end architectures that are trained on distributed word representations to directly produce different types of analyses traditionally assigned to downstream tasks. Syntactic parsing has been viewed as a crucial component for many tasks aimed at extracting various aspects of meaning from text, but recent work challenges many of these assumptions. For the task of semantic role labeling for instance, systems that make little or no use of syntactic information, have achieved state-of-the-art results (Marcheggiani et al., 2017). For tasks where syntactic information is still viewed as useful, a variety of new methods for the incorporation of syntactic information are employed, such as recursive models over parse trees (Socher et al., 2013; Ebrahimi and Dou, 2015), tree-structured attention mechanisms (Kokkinos and Potamianos, 2017), multi-task learning (Wu et al., 2017), or the use of various types of syntactically aware input representations, such as embeddings over syntactic dependency paths (Xu et al., 2015b).

Dependency representations have by now become widely used representations for syntactic analysis, often motivated by their usefulness in

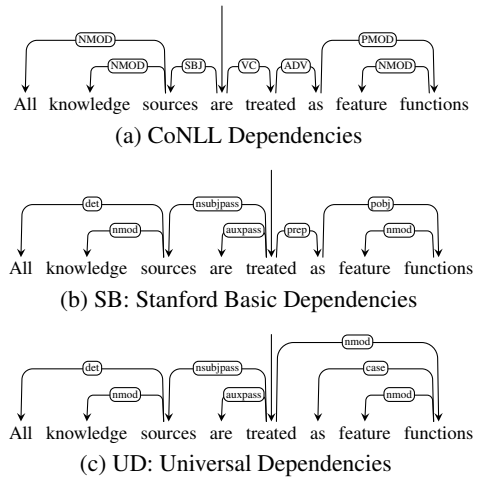


Figure 1: Dependency representations for the example sentence.

downstream application. There is currently a wide range of different types of dependency representations in use, which vary mainly in terms of choices concerning syntactic head status. Some previous studies have examined the effects of dependency representations in various downstream applications (Miyao et al., 2008; Elming et al., 2013). Most recently, the Shared Task on Extrinsic Parser Evaluation (Oepen et al., 2017) was aimed at providing better estimates of the relative utility of different types of dependency representations and syntactic parsers for downstream applications. The downstream systems in this previous work have, however, been limited to traditional (non-neural) systems and there is still a need for a better understanding of the contribution of syntactic information in neural downstream systems.

In this paper, we examine the use of syntactic representations in a neural approach to the task of relation classification. We quantify the effect of syntax by comparing to a syntax-agnostic approach and further compare different syntactic dependency representations that are used to generate embeddings over dependency paths.

2 Dependency representations

Figure 1 illustrates the three different dependency representations we compare: the so-called CoNLL-style dependencies (Johansson and Nugues, 2007) which were used for the 2007, 2008, and 2009 shared tasks of the Conference on Natural Language Learning (CoNLL), the Stanford ‘basic’ dependencies (SB) (Marneffe et al., 2006) and the Universal Dependencies (v1.3) (UD; McDonald et al., 2013; Marneffe et al., 2014; Nivre et al., 2016). We see that the analyses differ both in terms of their choices of heads vs. dependents and the inventory of dependency types. Where CoNLL analyses tend to view functional words as heads (e.g., the auxiliary verb *are*), the Stanford scheme capitalizes more on content words as heads (e.g., the main verb *treated*). UD takes the tendency to select contentful heads one step further, analyzing the prepositional complement *functions* as a head, with the preposition *as* itself as a dependent case marker. This is in contrast to the CoNLL and Stanford scheme, where the preposition is head.

For syntactic parsing we employ the parser described in Bohnet and Nivre (2012), a transition-based parser which performs joint PoS-tagging and parsing. We train the parser on the standard training sections 02-21 of the Wall Street Journal (WSJ) portion of the Penn Treebank (Marcus et al., 1993). The constituency-based treebank is converted to dependencies using two different conversion tools: (i) the `pennconverter` software¹ (Johansson and Nugues, 2007), which produces the CoNLL dependencies², and (ii) the Stanford parser using either the option to produce basic dependencies³ or its default option which is Universal Dependencies v1.3⁴. The parser achieves a labeled accuracy score of 91.23 when trained on the CoNLL08 representation, 91.31 for the Stanford basic model and 90.81 for the UD representation, when evaluated against the standard evaluation set (section 23) of the WSJ. We acknowledge that these results are not state-of-the-art parse results for English, however, the parser is straight-

¹<http://nlp.cs.lth.se/software/treebank-converter/>

²The `pennconverter` tool is run using the `rightBranching=false` flag.

³The Stanford parser is run using the `-basic` flag to produce the basic version of Stanford dependencies.

⁴Note, however, that the Stanford converter does not produce UD PoS-tags, but outputs native PTB tags.

forward to use and re-train with the different dependency representations. We also compare to another widely used parser, namely the pre-trained parsing model for English included in the Stanford CoreNLP toolkit (Manning et al., 2014), which outputs Universal Dependencies only. However, it was clearly outperformed by our version of the Bohnet and Nivre (2012) parser in the initial development experiments.

3 Relation extraction system

We evaluate the relative utility of different types of dependency representations on the task of semantic relation extraction and classification in scientific papers, SemEval Task 7 (Gábor et al., 2018). We make use of the system of Nooralahzadeh et al. (2018): a CNN classifier with dependency paths as input, which ranked 3rd (out of 28) participants in the overall evaluation of the shared task. Here, the shortest dependency path (*sdp*) connecting two target entities for each relation instance is provided by the parser and is embedded in the first layer of a CNN. We extend on their system by (i) implementing a syntax-agnostic approach, (ii) implementing hyper-parameter tuning for each dependency representation, and (iii) adding Universal Dependencies as input representation. We thus train classifiers with *sdp*s extracted from the different dependency representations discussed above and measure the effect of this information by the performance of the classifier.

3.1 Dataset and Evaluation Metrics

We use the SemEval-2018, Task 7 dataset (Gábor et al., 2018) from its *Subtask 1.1*. The training data contains abstracts of 350 papers from the ACL Anthology Corpus, annotated for concepts and semantic relations. Given an abstract of a scientific paper with pre-annotated domain concepts, the task is to perform relation classification. The classification sub-task 1.1 contains 1228 entity pairs that are annotated based on five asymmetric relations (USAGE, RESULT, MODEL-FEATURE, PART-WHOLE, TOPIC) and one symmetric relation (COMPARE). The relation instance along with its directionality are provided in both the training and the test data sets. The official evaluation metric is the macro-averaged F1-scores for the six semantic relations, therefore we will compare the impact of different dependency representations on the macro-averaged F1-scores.

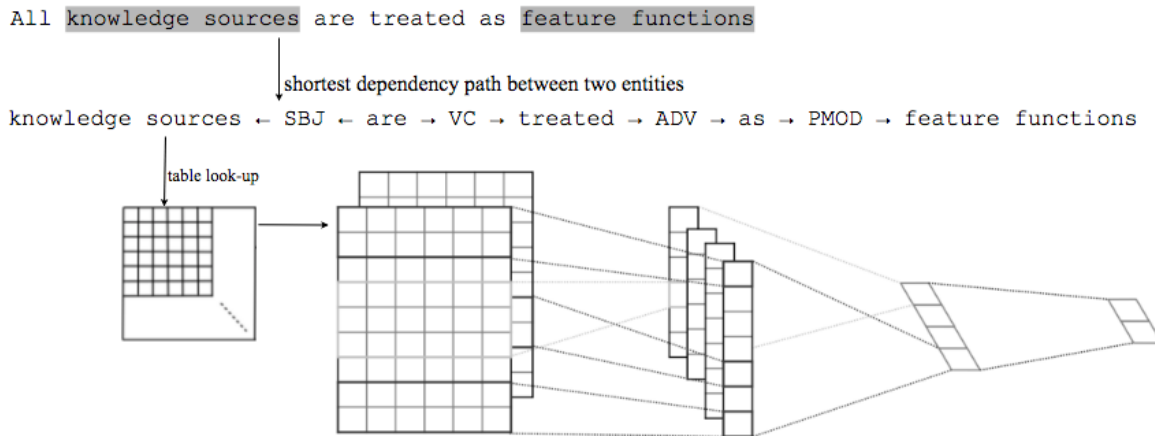


Figure 2: Model architecture with two channels for an example shortest dependency path (CNN model from Kim (2014)).

The training set for Subtask 1.1 is quite small, which is a challenge for end-to-end neural methods. To overcome this, we combined the provided datasets for Subtask 1.1 and Subtask 1.2 (relation classification on noisy data), which provides additional 350 abstracts and 1248 labeled entity pairs to train our model. This yields a positive impact (+16.00% F1) on the classification task in our initial experiments.

3.2 Pre-processing

Sentence and token boundaries are detected using the Stanford CoreNLP tool (Manning et al., 2014). Since most of the entities are multi-word units, we replace the entities with their codes in order to obtain a precise dependency path. Our example sentence *All knowledge sources are treated as feature functions*, an example of the USAGE relation between the two entities *knowledge sources* and *feature functions*, is thus transformed to All P05_1057_3 are treated as P05_1057_4.

Given an encoded sentence, we find the *sdp* connecting two target entities for each relation instance using a syntactic parser. Based on the dependency graph output by the parser, we extract the shortest dependency path connecting two entities. The path records the direction of arc traversal using left and right arrows (i.e. \leftarrow and \rightarrow) as well as the dependency relation of the traversed arcs and the predicates involved, following Xu et al. (2015a). The entity codes in the final *sdp* are replaced with the corresponding word tokens at the end of the pre-processing step.

For the sentence above, we thus extract the path: knowledge sources \leftarrow SBJ \leftarrow are \rightarrow VC \rightarrow treated \rightarrow ADV \rightarrow as \rightarrow PMOD \rightarrow feature functions

3.3 CNN model

The system is based on a CNN architecture similar to the one used for sentence classification in Kim (2014). Figure 2 provides an overview of the proposed model. It consists of 4 main layers as follows: 1) **Look-up Table and Embedding layer:** In the first step, the model takes a shortest dependency path (i.e., the words, dependency edge directions and dependency labels) between entity pairs as input and maps it into a feature vector using a look-up table operation. Each element of the dependency path (i.e. word, dependency label and arrow) is transformed into a embedding layer by looking up the embedding matrix $M \in \mathcal{R}^{d \times V}$, where d is the dimension of CNN embedding layer and V is the size of vocabulary. Each column in the embedding matrix can be initialized randomly or with pre-trained embeddings. The dependency labels and edge directions are always initialized randomly. 2) **Convolutional Layer:** The next layer performs convolutions with ReLU activation over the embeddings using multiple filter sizes and extracts feature maps. 3) **Max pooling Layer:** By applying the *max* operator, the most effective local features are generated from each feature map. 4) **Fully connected Layer:** Finally, the higher level syntactic features are fed to a fully connected *softmax* layer which outputs the probability distribution over each relation.

Representation	Hyper parameters						F1.(avg. in 5-fold)	
	Filter size	Num. Feature maps	Activation func.	L2 Reg.	Learning rate	Dropout Prob.	with default values	with optimal values
CoNLL08	4-5	1000	Softplus	1.15e+01	1.13e-03	1	73.34	74.49
SB	4-5	806	Sigmoid	8.13e-02	1.79e-03	0.87	72.83	75.05
UD v1.3	5	716	Softplus	1.66e+00	9.63E-04	1	68.93	69.57

Table 2: Hyper parameter optimization results for each model with different representation. The *max* pooling strategy consistently performs better in all model variations.

Relation	best F1 (in 5-fold)		Diff.
	without sdp	with sdp	
USAGE	60.34	80.24	+ 19.90
MODEL-FEATURE	48.89	70.00	+ 21.11
PART-WHOLE	29.51	70.27	+40.76
TOPIC	45.80	91.26	+45.46
RESULT	54.35	81.58	+27.23
COMPARE	20.00	61.82	+ 41.82
macro-averaged	50.10	76.10	+26.00

Table 1: Effect of using the shortest dependency path on each relation type.

4 Experiments

We run all the experiments with a multi-channel setting⁵ in which the first channel is initialized with pre-trained embeddings⁶ in static mode (i.e. it is not updated during training) and the second one is initialized randomly and is fine-tuned during training (non-static mode). The macro F1-score is measured by 5-fold cross validation and to deal with the effects of class imbalance, we weight the cost by the ratio of class instances, thus each observation receives a weight, depending on the class it belongs to.

4.1 Effect of syntactic information

To evaluate the effects of syntactic information in general for the relation classification task, we compare the performance of the model with and without the dependency paths. In the syntax-agnostic setup, a sentence that contains the participant entities is used as input for the CNN. We keep the value of hyper-parameters equal to the ones that are reported in the original work (Kim, 2014). To provide the *sdp* for the syntax-aware version we compare to, we use our parser with Stanford

⁵Initial experiments show that the multi-channel model works better than the single channel model

⁶We train 300-d domain-specific embeddings on the ACL Anthology corpus using the available word2vec implementation *gensim* for training.

dependencies. We find that the effect of syntactic structure varies between the different relation types. However, the *sdp* information has a clear positive impact on all the relation types (Table 1). It can be attributed to the fact that the context-based representations suffer from irrelevant subsequences or clauses when target entities occur far from each other or there are other target entities in the same sentence. The *sdp* between two entities in the dependency graph captures a condensed representation of the information required to assert a relationship between two entities (Bunescu and Mooney, 2005).

4.2 Comparison of different dependency representations

To investigate the model performance with various parser representations, we create a *sdp* for each training example using the different parse models and exploit them as input to the relation classification model. With the use of default parameters there is a chance that these favour one of the representations. In order to perform a fair comparison, we make use of Bayesian optimization (Brochu et al., 2010) in order to locate optimal hyper parameters for each of the dependency representations. We construct a Bayesian optimization procedure using a Gaussian process with 100 iterations and Expected Improvement (EI) for its acquisition functions. We set the objective function to maximize the macro F1 score over 5-fold cross validation on the training set. Here we investigate the impact of various system design choices with the following parameters:⁷ I) Filter region size: $\in \{3, 4, 5, 6, 7, 8, 9, 3-4, 4-5, 5-6, 6-7, 7-8, 8-9, 3-4-5, 4-5-6, 5-6-7, 6-7-8, 7-8-9\}$ II) Number of feature maps for each filter region size: $\in \{10 : 1000\}$ III) Activation function: $\in \{Sigmoid, ReLU, Tanh, Softplus, Idem\}$. IV) Pooling strategy: $\in \{max, avg\}$. V) L2 regular-

⁷Default values are $\{3-4-5, 128, ReLU, max, 3, 1e-3, 0.5\}$

Sentence	This indicates that there is no need to add punctuation in transcribing spoken corpora simply in order to help parsers.	class: PART_WHOLE
CoNLL08	punctuation ← obj ← add → adv → in → pmod → transcribing → obj → spoken corpora	
SB	punctuation ← dobj ← add → prep → in → pcomp → transcribing → dobj → spoken corpora	
UD v1.3	punctuation ← dobj ← add → advcl → transcribing → dobj → spoken corpora	
Sentence	In the process we also provide a formal definition of parsing motivated by an informal notion due to Lang .	class: MODEL-FEATURE
CoNLL08	formal definition → nmod → of → pmod → parsing	
SB	formal definition → prep → of → pobj → parsing	
UD v1.3	formal definition → nmod → parsing	
Sentence	This paper describes a practical "black-box" methodology for automatic evaluation of question-answering NL systems in spoken dialogue.	class: USAGE
CoNLL	"black-box" methodology → nmod → for → pmod → evaluation → nmod → of → pmod → question-answering NL systems	
SB	"black-box" methodology → prep → for → pobj → evaluation → prep → of → pobj → question-answering NL systems	
UD v1.3	"black-box" methodology → nmod → evaluation → nmod → question-answering NL systems	

Table 4: The examples for which the CoNLL/SB-based models correctly predict the relation type in 5-fold trials, whereas the UD based model has an incorrect prediction.

ization: $\in \{1e - 4 : 1e + 2\}$. VI) Learning rate: $\in \{1e - 6 : 1e - 2\}$. VII) Dropout probability ⁸: $\in \{0.1 : 1\}$. Table 2 presents the optimal values for each configuration using different dependency representations. We see that the optimized parameter settings vary for the different representations, showing the importance of tuning for these types of comparisons. The results furthermore show that the *sdp*s based on the Stanford Basic (SB) representation provide the best performance, followed by the CoNLL08 representation. We observe that the results for the UD representation are quite a bit lower than the two others.

5 Error analysis

Table 3 presents the effect of each parser representation in the classification task, broken down by relation type. We find that the UD-based model falls behind the others on the most relation types (i.e., COMPARE, MODEL-FEATURE, PART_WHOLE, TOPICS). To explore these differences in more detail, we manually inspect the instances for which the CoNLL/SB-based models correctly predict the relation type in 5-fold trials, whereas the UD-based model has an incorrect prediction. Table 4 shows some of these examples, marking the entities and the gold class of each instance and also showing the *sdp* from each representation. We observe that the UD paths are generally shorter. A striking similarity between most of the instances is the fact that one of the entities resides in a prepositional phrase. Whereas the SB and CoNLL paths explicitly represent the preposition in the path, the UD representation does not. Clearly, the difference between for

⁸The probability that each element is kept, in which 1 implies that none of the nodes are dropped out

Relation	Frq.	best F1 (in 5-fold)		
		CoNLL	SB	UD
USAGE	947	76.84	82.39	77.56
MODEL-FEATURE	498	68.27	68.54	66.36
PART_WHOLE	425	75.32	71.28	67.11
TOPIC	258	89.32	90.57	87.62
RESULT	193	82.35	81.69	82.86
COMPARE	136	66.67	66.67	54.24
macro-averaged		76.94	77.57	72.83

Table 3: Effect of using the different parser representation on each relation type.

instance the USAGE and PART_WHOLE relation may be indicated by the presence of a specific preposition (*X for Y* vs. *X of Y*). This is also interesting since this particular syntactic choice has been shown in previous work to have a negative effect on intrinsic parsing results for English (Schwartz et al., 2012).

6 Conclusion

This paper has examined the use of dependency representations for neural relation classification and has compared three widely used representations. We find that representation matters and that certain choices have clear consequences in downstream processing. Future work will extend the study to neural dependency parsers and other relation classification data sets.

References

- Bernd Bohnet and Joakim Nivre. 2012. A transition-based system for joint part-of-speech tagging and labeled non-projective dependency parsing. In *Proceedings of EMNLP*, pages 1455–1465, Jeju Island, Korea. ACL.
- Eric Brochu, Vlad M. Cora, and Nando de Freitas. 2010. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *CoRR*, abs/1012.2599.
- Razvan C. Bunescu and Raymond J. Mooney. 2005. A shortest path dependency kernel for relation extraction. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT '05*, pages 724–731, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Javid Ebrahimi and Dejing Dou. 2015. Chain based rnn for relation classification. In *HLT-NAACL*.
- Jakob Elming, Anders Johannsen, Sigrid Klerke, Emanuele Lapponi, Héctor Martínez Alonso, and Anders Søgaard. 2013. Down-stream effects of tree-to-dependency conversions. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, 2013*, pages 617–626.
- Kata Gábor, Davide Buscaldi, Anne-Kathrin Schumann, Behrang QasemiZadeh, Haifa Zargayouna, and Thierry Charnois. 2018. Semeval-2018 Task 7: Semantic relation extraction and classification in scientific papers. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA.
- Richard Johansson and Pierre Nugues. 2007. Extended constituent-to-dependency conversion for english. In *NODALIDA 2007 Proceedings*, pages 105–112. University of Tartu.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1746–1751. Association for Computational Linguistics.
- Filippos Kokkinos and Alexandros Potamianos. 2017. Structural attention neural networks for improved sentiment analysis. In *Meeting of the European Chapter of the Association for Computational Linguistics*, pages 586–591, Valencia, Spain.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Diego Marcheggiani, Anton Frolov, and Ivan Titov. 2017. A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling. In *Proceedings of the 21st Conference on Computational Natural Language Learning*, pages 411–420. Association for Computational Linguistics.
- Mitchell Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpora of English. The Penn Treebank. *Journal of Computational Linguistics*, 19:313–330.
- M. Marneffe, B. Maccartney, and C. Manning. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC-2006)*.
- 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 the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, and Oscar Täckström. 2013. Universal dependency annotation for multilingual parsing. In *Proceedings of Association for Computational Linguistics (ACL)*, pages 92–97.
- Yusuke Miyao, Rune Sætre, Kenji Sagae, Takuya Matsuzaki, and Jun’ichi Tsujii. 2008. Task-oriented evaluation of syntactic parsers and their representations. In *Meeting of the Association for Computational Linguistics*, pages 46–54, Columbus, OH, USA.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D. Manning, Ryan T. McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC*.
- Farhad Nooralahzadeh, Lilja Øvrelid, and Jan Tore Lønning. 2018. Sirius-Itg-uo at semeval-2018 task 7: Convolutional neural networks with shortest dependency paths for semantic relation extraction and classification in scientific papers. In *Proceedings of the 12th International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA.
- Stephan Oepen, Lilja Øvrelid, Jari Björne, Richard Johansson, Emanuele Lapponi, Filip Ginter, and Erik Velldal. 2017. The 2017 Shared Task on Extrinsic Parser Evaluation. Towards a reusable community infrastructure. In *Proceedings of the 2017 Shared Task on Extrinsic Parser Evaluation at the Fourth International Conference on Dependency Linguistics*

and the 15th International Conference on Parsing Technologies, pages 1–12, Pisa, Italy.

Roy Schwartz, Omri Abend, and Ari Rappoport. 2012. Learnability-based syntactic annotation design. In *Proceedings of COLING 2012*, pages 2405–2422. The COLING 2012 Organizing Committee.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Chris Manning, Andrew Ng, and Chris Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, WA, USA.

Shuangzhi Wu, Dongdong Zhang, Nan Yang, Mu Li, and Ming Zhou. 2017. Sequence-to-dependency neural machine translation. In *Meeting of the Association for Computational Linguistics*, pages 698–707, Vancouver, Canada.

Kun Xu, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2015a. Semantic relation classification via convolutional neural networks with simple negative sampling. *CoRR*, abs/1506.07650.

Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015b. Classifying relations via long short term memory networks along shortest dependency paths. In *Conference on Empirical Methods in Natural Language Processing*, pages 1785–1794, Lisbon, Portugal.