ChinaRxiv 中X

AI translation · View original & related papers at
chinarxiv.org/items/chinaxiv-202010.00067

Span-based Joint Entity and Relation Extraction Fusing Attention-based Span-specific and Contextual Semantic Representations

Authors: Ji Bin

Date: 2020-10-26T00:00:00+00:00

Abstract

Span-based joint extraction models have demonstrated their effectiveness in entity recognition and relation extraction. These models treat text spans as candidate entities and span tuples as candidate relation tuples. Span semantic representations are shared across both entity recognition and relation extraction, yet existing models fail to adequately capture the semantics of these candidate entities and relations. To address these issues, we introduce a span-based joint extraction framework with attention-based semantic representations. Specifically, attention is utilized to compute semantic representations, including span-specific and contextual representations. We further investigate the role of four attention variants in generating contextual semantic representations. Experimental results demonstrate that our model outperforms previous systems and achieves state-of-the-art results on ACE2005, CoNLL2004, and ADE.

Full Text

Preamble

Span-based Joint Entity and Relation Extraction with Attention-based Span-specific and Contextual Semantic Representations

Bin Ji, Shasha Li*, Jie Yu, Jun Ma, Qingbo Wu, Yusong Tan
College of Computer, National University of Defense Technology, Changsha, China
{jibin, shashali, yj, majun, qingbowu, yusongtan}@nudt.edu.cn

Abstract

Span-based joint extraction models have demonstrated their efficiency for entity recognition and relation extraction. These models regard text spans as

candidate entities and span tuples as candidate relation tuples. Span semantic representations are shared in both entity recognition and relation extraction, yet existing models cannot adequately capture the semantics of these candidate entities and relations. To address these limitations, we introduce a span-based joint extraction framework with attention-based semantic representations. Specifically, attentions are utilized to calculate semantic representations, including span-specific and contextual ones. We further investigate the effects of four attention variants in generating contextual semantic representations. Experiments show that our model outperforms previous systems and achieves state-of-the-art results on ACE2005, CoNLL2004, and ADE.

Introduction

This paper addresses intra-sentence joint entity and relation extraction. Joint extraction models can alleviate cascading errors and promote information utilization compared to pipelined approaches, and thus have drawn considerable attention. Typically, the joint extraction task is solved by sequence tagging based methods (Zheng et al., 2017; Chi et al., 2019).

Rather than using sequence tagging based methods, recent works attempt to solve the task with span-based joint extraction models (Dixit and Al-Onaizan, 2019; Luan et al., 2019). Typically, this approach first processes sentence text into text spans, which serve as span-based candidate entities ("spans" for short); then calculates span semantic representations and performs span classification to obtain predicted entities; next forms span-based candidate relation ("relation" for short) tuples with spans, and calculates relation semantic representations using corresponding span semantic representations; finally performs relation classification on relation semantic representations to obtain predicted relation triples. This mode further improves joint extraction performance, but suffers from three problems.

First, different tokens in a span should contribute differently to span representation, which we call span-specific features. However, existing methods treat each span token as equally important (Eberts and Ulges, 2019) or only consider span head and tail tokens (Dixit and Al-Onaizan, 2019), ignoring these significant features. Take the span "a Palestinian youth" in sentence 1 of Figure 1 [Figure 1: see original paper] as an example: the word "youth" should contribute much more to the span representation than "a" and "Palestinian" when classifying the span as "PER" . Sentence 1 illustrates this case: "The army said troops fired, and hit a boy, after [a Palestinian youth]PER threw a stone."

Second, local contextual information of relation tuples is either omitted (Luan et al., 2018; Dixit and Al-Onaizan, 2019) or simply calculated through max pooling (Eberts and Ulges, 2019) when performing relation classification, which does not sufficiently capture the information contained therein. Yet local context may contain crucial information to help predict the relations that relation tuples hold. A case study is shown in sentence 2 of Figure 1, where the word "ownership"

(in red font) can greatly help determine the relation ("PART-WHOLE") of the relation tuple. Sentence 2 demonstrates: "The weak score is primarily the result of [Starbucks]ORG current strategy of increasing equity ownership of [several foreign subsidiaries]ORG," with the relation: <several foreign subsidiaries, Starbucks, PART-WHOLE>.

Third, sentence-level contextual information is ignored in both span and relation classifications, which may provide important compensatory information for both tasks. An example is given in sentence 3 of Figure 1, where the word "state" (in red font) benefits the relation classification of the relation tuple <all of the West Bank and Gaza, Palestinians>, while "state" is neither contained in the relation tuple nor in the local context (namely "claim"), but appears elsewhere in sentence 3. Sentence 3 shows: "[Palestinians]GPE claim [all of the West Bank and Gaza]LOC for a state," with the relation: <all of the West Bank and Gaza, Palestinians, PART-WHOLE>.

To address these issues, we introduce a span-based joint extraction model with attention-based span-specific and contextual semantic representations. Specifically: 1) MLP attention is used to calculate span-specific semantic representation; 2) attention-based sentence-level contextual semantic representation for span is calculated by taking span-specific semantic representation as query and sentence token sequence semantic representations as key and value respectively; 3) local and sentence-level contextual semantic representations for relation are obtained by attention calculation with relation tuple semantic representations and corresponding token sequence semantic representations. The advantage of this approach is that we can capture the most useful information to constitute efficient span and relation semantic representations.

We adopt BERT (Devlin et al., 2019) as the default backbone network and explore three research questions. Moreover, we investigate the effects of Multi-Head attention (Vaswani et al., 2017), Dot-Product attention (Luong et al., 2015), General attention (Luong et al., 2015), and Additive attention (Bahdanau et al., 2015) in generating contextual semantic representations. Extensive experiments on three benchmark datasets show that our model consistently outperforms previous systems. In addition, Multi-Head attention firmly improves over other attention variants.

2 Related Works

Traditionally, pipelined entity and relation extraction methods are divided into two subtasks: entity detection and relation classification. Various neural networks have been widely investigated for these subtasks, such as RNNs (Huang et al., 2015; Ma and Hovy, 2016) and CNNs (Limsopatham and Collier, 2016) for entity detection, and RNNs (Zhang and Wang, 2015), CNNs (Zeng et al., 2014), and Transformer (Verga et al., 2018; Wang et al., 2019) for relation classification.

As discussed in §1, joint entity and relation extraction is typically formulated

as a sequence tagging based task. Traditionally, table-filling methods have been widely explored (Miwa and Sasaki, 2014; Gupta et al., 2016), where token labels and relation labels fill the diagonal and off-diagonal of the table respectively. Recently, many works have concentrated on leveraging deep neural networks to tackle this task, e.g., stacked bidirectional LSTM (Miwa and Bansal, 2016; Zheng et al., 2017), combination of bidirectional LSTM and CNN (Zhou et al., 2017), and combination of bidirectional LSTM and attention mechanism (Chi et al., 2019; Nguyen and Verspoor, 2019). In addition, a novel machine reading comprehension based approach (Li et al., 2019) has been proposed, which formulates the task as a Question & Answer task but still in sequence tagging based mode.

Recently, span-based joint extraction methods have been investigated to tackle problems existing in sequence tagging based methods, such as the inability to detect overlapping entities. Specifically, Dixit and Al-Onaizan (2019) implement this method by obtaining span semantic representations through a BiLSTM over concatenated ELMo, word, and character embeddings, then sharing them in both span and relation classifications. Luan et al. (2018) obtain span semantic representations generally the same as Lee et al. (2017), but reinforce them by introducing a coreference task. Following Luan et al. (2018), Luan et al. (2019) propose DyGIE, which can capture span interactions through a span graph constructed dynamically. Wadden et al. (2019) deliver further performance improvements on DyGIE by replacing the BiLSTM with BERT and introduce DyGIE++. More recently, Eberts and Ulges (2019) propose SpERT, a simple but effective span-based model that takes BERT as backbone and uses two FFNNs to classify span and relation respectively. Unlike previous works, SpERT dramatically reduces model training complexity by adopting negative sampling.

Our work follows SpERT, but differs in span-specific and contextual semantic representations. Specifically, our model obtains these semantic representations with attention mechanisms. By calculating the matching degree between target sequence semantic representations and source sequence semantic representations, the attention mechanism obtains attention scores on the source sequence, which are essentially weight scores. The more important the information, the higher weight score it holds. Classified by implementation manners of score function, attention mechanisms have multiple variants, e.g., Content-Based attention (Graves et al., 2014), Additive attention (Bahdanau et al., 2015), General attention (Luong et al., 2015), Dot-Product attention (Luong et al., 2015), and Multi-Head attention (Vaswani et al., 2017).

3 Approach

In the following, we abbreviate "semantic representation" as "representation." Figure 2 [Figure 2: see original paper] provides an overview of our model, which uses BERT as encoder following SpERT: we map word embeddings into BERT embeddings using pre-trained Transformer blocks (Vaswani et al., 2017). Based on these representations, we calculate span representations and perform span

classification & filtration (§3.1); then we organize relation tuples, calculate relation representations, and perform relation classification & filtration (§3.2); third, we investigate the effects of multiple attention variants in generating contextual representations (§3.3); finally, we introduce model training settings (§3.4).

We define a sentence and a span from the sentence to help introduce the rest, where t denotes tokens and subscripts (e.g., 1, 2, 3···) denote token indexes, as follows:

$$\text{sentence} : S = (t_1, t_2, t_3, \dots, t_n)$$

$$\text{span} : s = (t_i, t_{i+1}, t_{i+2}, \dots, t_{i+j})$$

3.1 Span Classification and Filtration

We add a NoneEntity type to the pre-defined entity types (denoted as). Spans will be classified as NoneEntity as long as they do not hold any pre-defined entity types.

As Figure 2 shows, span representation for classification comprises four parts: a) concatenation of span head and tail representations, b) span-specific representation, c) sentence-level contextual representation, and d) span width embedding. We use X_i to denote the BERT embedding of token t_i, and define the BERT embedding sequences of S and s as follows, where X_0 denotes the BERT embedding of [CLS]:

$$B_S = [X_0, X_1, X_2, X_3, \dots, X_n]$$

$$B_s = [X_i, X_{i+1}, X_{i+2}, \dots, X_{i+j}]$$

Concatenation of span head and tail representations. If a span comprises more than one token, we concatenate the BERT embeddings of span head and tail tokens. Otherwise, we duplicate the BERT embedding of the single token and concatenate them. The concatenation result for span s is:

$$H_s = [X_i; X_{i+j}]$$

Span-specific representation. Here we use MLP attention (Dixit and Al-Onaizan, 2019) to calculate span-specific representation. Taking span s as an example:

$$V_k = \text{MLP}_k(X_k)$$

$$\alpha_k = \frac{\exp(V_k)}{\sum_{m \in [i, i+j]} \exp(V_m)}$$

$$F_s = \sum_{k \in [i, i+j]} \alpha_k \cdot X_k$$

Where V_k is a scalar; α_k is the attention weight of X_k, computed by Softmax function; F_s is the span-specific representation through matrix calculation on attention weights and B_s. This way, we can evaluate the significance of each span token, with more important tokens receiving larger attention weights.

Sentence-level contextual representation. Taking F_s as query, B_S as key and value respectively, sentence-level contextual representation for span s is calculated as:

$$T_s = \text{Attention}(F_s, B_S, B_S)$$

Information beneficial for span classification will be assigned heavy weight, and the contextual representation will be incorporated into the span representation.

Span width embedding. Span width embedding allows the model to incorporate a prior over the span width. Fixed-size embeddings for each span width $1, 2, ...$ (Lee et al., 2017) are learned during model training. Thus, we can look up a width embedding W_{j+1} from the embedding matrix for s.

Span classification. The final span representation for classification is:

$$R_s = [T_s; F_s; H_s; W_{j+1}]$$

R_s first passes through a multi-layer FFNN and then is fed into a Softmax classifier, which yields a posterior for s on η (including NoneEntity), as:

$$y_s = \text{Softmax}(\text{FFNN}(R_s))$$

Span filtration. By selecting the highest-scored class, y_s estimates which entity type s holds. We only keep spans that are not classified as NoneEntity, forming a predicted entity set E. Then we perform relation classification on relation tuples derived from $\{E \otimes E\}$ to reduce search space, where $\otimes$ denotes Cartesian Product.

3.2 Relation Classification and Filtration

We add a NoneRelation type to the pre-defined relation types (denoted as γ). Let s_1, s_2 be two spans; relation tuples taken for relation classification are defined as:

$$\langle s_1, s_2 \rangle \in \{E \otimes E\}, \quad s_1 \neq s_2$$

As Figure 2 shows, relation representation for classification comprises three parts: a) concatenation of relation tuple representations, b) local contextual representation, and c) sentence-level contextual representation.

Concatenation of relation tuple representations. Before concatenating R_{s1} and R_{s2}, we first apply a multi-layer FFNN to them to reduce their dimensions. The concatenation result is:

$$H_r = [\mathrm{FFNN}(R_{s1}); \mathrm{FFNN}(R_{s2})]$$

Local contextual representation. Let B_c denote the BERT embedding sequence of local context between s_1 and s_2, as:

$$B_c = (X_m, X_{m+1}, X_{m+2}, \dots, X_{m+n})$$

The attention-based local contextual representation is calculated by taking H_r as query, B_c as key and value respectively, as:

$$F_r = \mathrm{Attention}(H_r, B_c, B_c)$$

Sentence-level contextual representation. The sentence-level contextual representation is calculated by taking H_r as query, B_S as key and value respectively, as:

$$T_r = \mathrm{Attention}(H_r, B_S, B_S)$$

Relation classification. Before F_r and T_r are incorporated into the relation representation, we first apply two different multi-layer FFNNs to them to reduce their dimensions, aiming to keep them in a proper proportion in the relation representation. The final relation representation for classification is:

$$R_r = [H_r; \mathrm{FFNN}_F(F_r); \mathrm{FFNN}_T(T_r)]$$

Akin to span classification, R_r first passes through a multi-layer FFNN and then is fed into a Softmax classifier, which yields a posterior for $\langle s_1, s_2 \rangle$ on γ (including NoneRelation), as:

$$y_r = \mathrm{Softmax}(\mathrm{FFNN}(R_r))$$

Relation filtration. By selecting the highest-scored class, y_r estimates which relation type $\langle s_1, s_2 \rangle$ holds. Only relation tuples that are not classified as NoneRelation are kept and compose predicted relation triples with predicted types.

3.3 Attention Variants

In this paper, we investigate the effects of Multi-Head attention, Additive attention, Dot-Product attention, and General attention in generating contextual representations. Their score functions are shown below:

$$\text{Multi-Head attention} : \text{score} = Q \cdot K^T$$

$$\text{Additive attention} : \text{score} = W_1 \cdot Q + W_2 \cdot K$$

$$\text{Dot-Product attention} : \text{score} = W \cdot (Q \odot K)$$

$$\text{General attention} : \text{score} = Q \cdot W \cdot K$$

Where Q, K denote query and key respectively; W denotes parameter matrix; d_K denotes dimension of K; $\cdot$ and $\odot$ denote matrix broadcast and element-wise multiplication respectively. For Multi-Head attention and Dot-Product attention, we first apply different multi-layer FFNNs to Q and K, aiming to convert them to the same dimension.

3.4 Model Training

Parameter matrices of FFNNs and attentions are learned, and BERT is fine-tuned during model training. The joint loss function of our model is defined as:

$$L = 0.4L_s + 0.6L_r$$

Where L_s denotes the cross-entropy loss of span classification and L_r denotes the binary cross-entropy loss of relation classification. Since the performance of relation classification is generally worse than span classification, we apply a larger weight to L_r, aiming to let the model focus more on relation classification.

Negative sampling (Eberts and Ulges, 2019) is adopted during model training to improve model performance and robustness. Unlike previous works, we adopt a dynamic sampling strategy, where the negative examples for both entity and relation are thirtyfold of the ground truth ones in each sentence. This strategy keeps a much more balanced data distribution in the training data.

4 Experiments

4.1 Datasets

We evaluate our model on ACE2005, CoNLL2004, and ADE, referred to as ACE05, CoNLL04, and ADE respectively in the following:

- **ACE05** (Doddington et al., 2004) English dataset comprises news articles from multiple domains, e.g., broadcast, newswire, weblog, etc. Seven coarse-grained entity types and six coarse-grained relation types are pre-defined. We follow the training/dev/test split from (Li and Ji, 2014; Li et al., 2019). It has 351 documents for training, 80 for development, and 80 for test, of which 437 contain overlapping entities.
- **CoNLL04** (Roth and Yih, 2004) comprises news articles from outlets such as WSJ and AP. We follow the training/dev/test split from (Adel and Schütze, 2017; Bekoulis et al., 2018), which consists of 910 articles for training, 243 for development, and 288 for test.
- **ADE** (Gurulingappa et al., 2012) aims to extract drug-related adverse effects from medical text, including two pre-defined entity types (namely Adverse-Effect and Drug) and a single relation type, i.e., Adverse-Effect. It consists of 4,272 sentences, of which 1,695 contain overlapping entities.

We conduct 10-fold cross-validation following (Bekoulis et al., 2018; Eberts and Ulges, 2019).

For ACE05, following (Li et al., 2019; Luan et al., 2019), an entity is considered correct if we can identify its head region and type correctly. A relation is considered correct if we can identify its argument entities and type correctly. For CoNLL04 and ADE, following (Li et al., 2019; Eberts and Ulges, 2019), we treat an entity as correct when its type and entity region match the ground truth, and treat a relation as correct when its type and argument entities match the ground truth.

4.2 Experimental Settings

We build our model upon the English cased version of BERTBASE. We set the negative sampling rate to 30, batch size for model training to 8, dropout to 0.2, and width embedding dimension to 50. For Multi-Head attention, the head number is set to 8. FFNN_F and FFNN_T contain three fully connected layers; all other FFNNs contain two layers. We set different epochs for different datasets.

For all datasets, the span width threshold is initialized to 10. In our model, we adopt weighted loss setting, as shown in equation (11), and follow Eberts and Ulges (2019) for other hyperparameter settings.

4.3 Baseline Models

We compare our model with the following:

- **DyGIE** (Luan et al., 2019) is the current span-based state-of-the-art model on ACE05. It reinforces span and relation representations by introducing a coreference task.
- **Multi-turn QA** (Li et al., 2019) is the current sequence tagging based state-of-the-art model on ACE05 and CoNLL04. It formulates joint entity and relation extraction as a multi-turn question and answer task, but still in sequence tagging based mode.
- **SpERT** (Eberts and Ulges, 2019) is the current span-based state-of-the-art model on ADE and CoNLL04. Our work follows this model but adopts attention-based span-specific and contextual representations.
- **Relation-Metric** (Tran and Kavuluru, 2019) is a sequence tagging based model in a multi-task learning scheme. It reports performances on ADE and CoNLL04, and achieves state-of-the-art on ADE.

4.4 Main Results

We compare our model with both sequence tagging based methods and span-based methods on the three benchmark datasets, and show the results in Table 1 . We denote our method as SPANMulti-Head, which means we use Multi-Head attention to calculate contextual representations. For ACE05 and CoNLL04, we report performance under micro-average metrics, and apply macro-average metrics to ADE performance, following prior works. For ACE05 and ADE, all reported performances take overlapping entities into consideration.

SPANMulti-Head consistently outperforms both sequence tagging based and span-based state-of-the-art methods on the three benchmark datasets. Compared to SpERT, SPANMulti-Head delivers performance increases in entity recognition by 1.29 (CoNLL04) and 1.31 (ADE), while showing even better improvements in relation extraction, by 2.86 (CoNLL04) and 1.89 (ADE). We attribute these performance increases to efficient span-specific representation and contextual representations. Moreover, SPANMulti-Head delivers solid performance increases compared to DyGIE by 1.19 and 2.04 in entity recognition and relation extraction on ACE05. However, it' s worth noting that DyGIE adopts a multi-task learning scheme, reinforcing span representations by introducing a coreference task, which is absent in our method.

Besides Multi-Head attention, we investigate Dot-Product attention, General attention, and Additive attention in our method. Table 2 shows performances of these attention variants on the three benchmark datasets. SPANMulti-Head consistently outperforms the other three methods. One possible reason is that in Multi-Head attention, the eight attention heads attend to different contextual information and learn features from different representation spaces. Thus, Multi-Head attention based contextual representations can better compensate span and relation representations.

5 Ablation Study

Based on SPANMulti-Head, we conduct ablations on the ACE2005 dev set to analyze the effects of different model components.

5.1 Span-specific and Sentence-level Contextual Representations for Span

Table 3 shows the effects of span-specific representation and sentence-level contextual representation for span in our model, where -SpanSpecific denotes ablating the span-specific representation by replacing $[F_s; H_s]$ in equation (4) with the max pooling of B_s; -SentenceLevel denotes ablating the sentence-level contextual representation by replacing T_s in equation (4) with the BERT embedding of [CLS] generated on B_S; base is the model performing both ablations, which represents the default span representation settings in SpERT. For ACE05, we observe that both span-specific representation and sentence-level contextual representation are helpful for both entity recognition and relation extraction, as span representations are shared in the two subtasks.

5.2 Local and Sentence-level Contextual Representations for Relation

Table 4 shows the effects of local and sentence-level contextual representations for relation in our model, where -local denotes ablating local representation by replacing F_r in equation (9) with the max pooling of B_c; -SentenceLevel denotes ablating sentence-level contextual representation by removing $\mathrm{FFNN}_T(T_r)$ from equation (9); base is the model performing both ablations, which represents the default relation representation settings in SpERT. For ACE05, we observe that both local and sentence-level contextual representations apparently benefit relation extraction, while having negligible influence on entity recognition. A convincing explanation is that these representations directly constitute relation representation, while affecting span representation only through backpropagation.

It is worth noting that local contextual representation has a greater impact on relation extraction compared to sentence-level representation. One reason is that information determining relation type mainly exists in relation tuples and the local context. Another reason is that as compensation information, sentence-level contextual representation occupies a relatively small proportion in relation representation, aiming to avoid introducing noise into relation representation.

6 Conclusion

We introduce attention-based semantic representation generating methods in span-based joint entity and relation extraction. We apply MLP attention to capture span-specific features to obtain semantically rich span representation, and calculate task-specific contextual representations with attention architecture to further reinforce span and relation representations. Our approach firmly outper-

forms both sequence tagging based and span-based state-of-the-art methods on three benchmark datasets, creating new state-of-the-art results. As future work, we would like to consider further improving relation classification performance by reducing span classification errors. We also plan to explore more advanced methods for encoding efficient span and relation representations.

Acknowledgements
This work is supported by the National Key Research and Development Program of China (2018YFB1004502) and the National Natural Science Foundation of China (61532001).

References

Heike Adel and Hinrich Schütze. 2017. Global normalization of convolutional neural networks for joint entity and relation classification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1723–1729. Association for Computational Linguistics.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Giannis Bekoulis, Johannes Deleu, Thomas Demeester, and Chris Develder. 2018. Joint entity recognition and relation extraction as a multi-head selection problem. *Expert Syst. Appl.*, 114:34–45.

Renjun Chi, Bin Wu, Linmei Hu, and Yunlei Zhang. 2019. Enhancing joint entity and relation extraction with language modeling and hierarchical attention. In *Web and Big Data - Third International Joint Conference, APWeb-WAIM 2019, Chengdu, China, August 1-3, 2019, Proceedings, Part I*, volume 11641 of *Lecture Notes in Computer Science*, pages 314–328. Springer.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Kalpit Dixit and Yaser Al-Onaizan. 2019. Span-level model for relation extraction. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28 - August 2, 2019, Volume 1: Long Papers*, pages 5308–5314. Association for Computational Linguistics.

George R. Doddington, Alexis Mitchell, Mark A. Przybocki, Lance A. Ramshaw, Stephanie M. 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 2004, May 26-28, 2004, Lisbon, Portugal. European Language Resources Association.

Markus Eberts and Adrian Ulges. 2019. Span-based joint entity and relation extraction with transformer pre-training. *arXiv preprint arXiv:1909.07755*, abs/1909.07755.

Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural turing machines. *arXiv preprint arXiv:1410.5401*.

Pankaj Gupta, Hinrich Schütze, and Bernt Andrassy. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*, pages 2537–2547. ACL.

Harsha Gurulingappa, Abdul Mateen Rajput, Angus Roberts, Juliane Fluck, Martin Hofmann-Apitius, and Luca Toldo. 2012. Development of a benchmark corpus to support the automatic extraction of drug-related adverse effects from medical case reports. *J. Biomed. Informatics*, 45(5):885–892.

Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*, abs/1508.01991.

Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 188–197. Association for Computational Linguistics.

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, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 402–412. The Association for Computer Linguistics.

Xiaoya Li, Fan Yin, Zijun Sun, Xiayu Li, Arianna Yuan, Duo Chai, Mingxin Zhou, and Jiwei Li. 2019. Entity-relation extraction as multi-turn question answering. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28 - August 2, 2019, Volume 1: Long Papers*, pages 1340–1350. Association for Computational Linguistics.

Nut Limsopatham and Nigel Collier. 2016. Bidirectional LSTM for named entity recognition in twitter messages. In *Proceedings of the 2nd Workshop on Noisy User-generated Text, NUT@COLING 2016, Osaka, Japan, December 11, 2016*, pages 145–152. The COLING 2016 Organizing Committee.

Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. Multi-task identification of entities, relations, and coreference for scientific knowledge

graph construction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3219–3232. Association for Computational Linguistics.

Yi Luan, Dave Wadden, Luheng He, Amy Shah, Mari Ostendorf, and Hannaneh Hajishirzi. 2019. A general framework for information extraction using dynamic span graphs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 3036–3046. Association for Computational Linguistics.

Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1412–1421. The Association for Computational Linguistics.

Xuezhe Ma and Eduard H. Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer 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, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.

Makoto Miwa and Yutaka Sasaki. 2014. Modeling joint entity and relation extraction with table representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1858–1869. ACL.

Dat Quoc Nguyen and Karin Verspoor. 2019. End-to-end neural relation extraction using deep biaffine attention. In *Advances in Information Retrieval - 41st European Conference on IR Research, ECIR 2019, Cologne, Germany, April 14-18, 2019, Proceedings, Part I*, volume 11437 of *Lecture Notes in Computer Science*, pages 729–738. Springer.

Dan Roth and Wen-tau Yih. 2004. A linear programming formulation for global inference in natural language tasks. In *Proceedings of the Eighth Conference on Computational Natural Language Learning, CoNLL 2004, Held in cooperation with HLT-NAACL 2004, Boston, Massachusetts, USA, May 6-7, 2004*, pages 1–8. ACL.

Tung Tran and Ramakanth Kavuluru. 2019. Neural metric learning for fast end-to-end relation extraction. *arXiv preprint arXiv:1905.07458*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.

Patrick Verga, Emma Strubell, and Andrew McCallum. 2018. Simultaneously self-attending to all mentions for full-abstract biological relation extraction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 872–884. Association for Computational Linguistics.

David Wadden, Ulme Wennberg, Yi Luan, and Hannaneh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 5783–5788. Association for Computational Linguistics.

Haoyu Wang, Ming Tan, Mo Yu, Shiyu Chang, Dakuo Wang, Kun Xu, Xiaoxiao Guo, and Saloni Potdar. 2019. Extracting multiple-relations in one-pass with pre-trained transformers. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28 - August 2, 2019, Volume 1: Long Papers*, pages 1371–1377. Association for Computational Linguistics.

Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *COLING 2014, 25th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, August 23-29, 2014, Dublin, Ireland*, pages 2335–2344. ACL.

Dongxu Zhang and Dong Wang. 2015. Relation classification via recurrent neural network. *arXiv preprint arXiv:1508.01006*, abs/1508.01006.

Suncong Zheng, Feng Wang, Hongyun Bao, Yuexing Hao, Peng Zhou, and Bo Xu. 2017. Joint extraction of entities and relations based on a novel tagging scheme. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1227–1236. Association for Computational Linguistics.

Peng Zhou, Suncong Zheng, Jiaming Xu, Zhenyu Qi, Hongyun Bao, and Bo Xu. 2017. Joint extraction of multiple relations and entities by using a hybrid neural network. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data - 16th China National Conference, CCL 2017, - and - 5th International Symposium, NLP-NABD 2017, Nanjing, China, October 13-15, 2017, Proceedings*, volume 10565 of *Lecture*

ChinaRxiv 中X

Notes in Computer Science, pages 135–146. Springer.

Note: Figure translations are in progress. See original paper for figures.

Source: ChinaXiv —Machine translation. Verify with original.