

Partially-supervised Mention Detection

Lesly Miculicich^{†‡} James Henderson[†]

[†] Idiap Research Institute, Switzerland

[‡]École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

{`lmiculicich`, `jhenderson`}@idiap.ch

Abstract

Learning to detect entity mentions without using syntactic information can be useful for integration and joint optimization with other tasks. However, it is common to have partially annotated data for this problem. Here, we investigate two approaches to deal with partial annotation of mentions: weighted loss and soft-target classification. We also propose two neural mention detection approaches: a sequence tagging, and an exhaustive search. We evaluate our methods with coreference resolution as a downstream task, using multitask learning. The results show that the recall and F1 score improve for all methods.

1 Introduction

Mention detection is the task of identifying text spans referring to an entity: named, nominal or pronominal (Florian et al., 2004). It is a fundamental component for several downstream tasks, such as coreference resolution (Soon et al., 2001), and relation extraction (Mintz et al., 2009); and it can help to maintain coherence in large text generation (Clark et al., 2018), and contextualized machine translation (Miculicich et al., 2018). Previous studies tackled mention detection jointly with named entity recognition (Xu et al., 2017; Katiyar and Cardie, 2018; Ju et al., 2018; Wang et al., 2018). There, only certain types of entities are considered (e.g., person, location), and the goal is to recognize mention spans and their types. In this study, we are interested in discovering entity mentions, which can potentially be referred to in the text, without the use of syntactic parsing information. Our long term objective is to have a model that keeps track of entities in a document for word disambiguating language modeling and machine translation.

Data from coreference resolution is suitable for our task, but the annotation is partial in that it contains only mentions that belong to a coreference chain, not singletons. Nevertheless, the missing mentions have approximately the same distribution as the annotated ones, so we can still learn this distribution from the data. Figure 1 shows an example from Ontonotes V.5 dataset (Pradhan et al., 2012) where “the taxi driver” is annotated in sample 1 but not in 2. Thus, we approach mention detection as a partially supervised problem and investigate two simple techniques to compensate for the fact that some negative examples are true mentions: weighted loss functions and soft-target classification. By doing this, the model is encouraged to predict more false-positive samples, so it can detect potential mentions which were not annotated. We implement two neural mention detection methods: a sequence tagging approach, and an exhaustive search approach. The first method is novel, whereas the other is similar to previous work (Lee et al., 2017). We evaluate both techniques for coreference resolution by implementing a multitask learning system. We show that the proposed techniques help the model increase recall significantly with a minimal decrease in precision. In consequence, the F1 score of the mention detection and coreference resolution improves for both methods, and the exhaustive search approach yields a significant improvement over the baseline coreference resolver.

Our contributions are:

- i We investigate two techniques to deal with partially annotated data.

1. [They] informed the [taxi driver] and asked [him] to take the vehicle outside ...
2. On [our] way back , the taxi driver gave [us] an explanation ...

Figure 1: Samples from CoNLL 2012. Annotated mentions are within brackets, non-annotated ones are underlined.

- ii We propose a sequence tagging method for mention detection that can model nested mentions.
- iii We improve an exhaustive search method for mention detection.
- iv We approach mention detection and coreference resolution as multitask learning and improve both tasks' recall.

The rest of the paper is organized as follows. Sections 2 and 3 describe the two mention detection approaches we use in our experiments. Section 4 presents the proposed methods to deal with partially annotated mentions. We use coreference resolution as a proxy task for testing our methods which is described in Section 5. Section 6 contains the experimental setting and the analysis of results. Section 7 contains related work to this study. Finally, the final conclusion is drawn Section 8.

2 Sequence tagging model

Several studies have tackled mention detection and named entity recognition as a tagging problem. Some of them use one-to-one sequence tagging techniques (Lample et al., 2016; Xu et al., 2017), while others use more elaborate techniques to include nested mentions (Katiyar and Cardie, 2018; Wang et al., 2018). Here, we propose a simpler yet effective tagging approach that can manage nested mentions.

We use a sequence-to-sequence model, which allows us to tag each word with multiple labels. The words are first encoded and contextualized using a recurrent neural network, and then a sequential decoder predicts the output tag sequence. During decoding, the model keeps a pointer into the encoder, indicating the word's position, which is being tagged at each time step. The tagging is done using the following set of symbols: $\{[,], +, -\}$. The brackets “[” and “]” indicate that the tagged word is the starting or ending of a mention respectively, the symbol “+” indicates that one or more mention brackets are open, and “-” indicates that none mention bracket is open. The pointer into the encoder moves to the next word only after predicting “+” or “-”; otherwise, it remains in the same position. Figure 2 shows a tagging example indicating the alignments of words with tags.

Given a corpus of sentences $X = (x_1, \dots, x_M)$, the goal is to find the parameters Θ which maximize the log likelihood of the corresponding tag sequences $Y = (y_1, \dots, y_T)$:

$$P_{\Theta}(Y|X) = \prod_{t=1}^T P_{\Theta}(y_t|X, y_1, \dots, y_{t-1}) \quad (1)$$

The next tag probability is estimated with a softmax over the output vector of a neural network:

$$P_{\Theta}(y_t|X, y_1, \dots, y_{t-1}) = \text{softmax}(o_t) \quad (2)$$

$$o_t = \text{relu}(W_o \cdot [d_t, h_i] + b_o) \quad (3)$$

where W_o, b_o are parameters of the network, d_t is the vector representation of the tagged sequence at time-step t , modeled with a long-short term memory (LSTM) (Hochreiter and Schmidhuber, 1997), and h_i is the vector representation of the pointer's word at time t contextualized with a bidirectional LSTM (Graves and Schmidhuber, 2005).

$$(h_1, \dots, h_M) = \text{BiLSTM}(X) \quad (4)$$

$$d_t = \text{LSTM}(y_1, \dots, y_{t-1}) \quad (5)$$




Figure 2: Tagged sentence example

where the decoder is initialized with the last states of the bidirectional encoder, $d_0 = h_M$. The i -th word pointed to at time t is given by:

$$i \leftarrow \begin{cases} 0, & \text{if } t = 0 \\ i + 1, & \text{if } t > 0 \text{ and } y_{t-1} \in \{+, -\} \\ i, & \text{otherwise} \end{cases} \quad (6)$$

At decoding time, we use a beam search approach to obtain the sequence. The complexity of the model is linear with respect to the number of words. It can be parallelized at training time, given that it uses ground-truth data for the conditioned variables. However, it cannot be parallelized during decoding because of its autoregressive nature.

3 Span scoring model

Our span scoring model of mention detection is similar to the work of Lee et al. (2017) for solving coreference resolution, and to Ju et al. (2018) for nested named mention detection, as both are exhaustive search methods. The objective is to score all possible spans m_{ij} in a document, where i and j are the starting and ending word positions of the span in the document. For this purpose, we minimize the binary cross-entropy with the labels y :

$$H(y, P_{\Theta}(m)) = -\frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M (y_{m_{ij}} * \log(P_{\Theta}(m_{ij})) + (1-y_{m_{ij}}) * \log(1-P_{\Theta}(m_{ij}))) \quad (7)$$

where Θ are the parameters of the model, $y_{m_{ij}} \in [0, 1]$ is one when there is a mention from position i to j . If $y_{m_{ij}}$ is zero when there is no mention annotated, this is the same as maximizing the log-likelihood. Nevertheless, we will consider models where this is not the case.

The probability of detection is estimated as:

$$P_{\Theta}(m_{ij}) = \sigma(V \cdot \text{relu}(W_m \cdot m_{ij} + b_m)) \quad (8)$$

$$m_{ij} = \text{relu}(W_h \cdot [h_i, h_j, \tilde{x}_{ij}] + b_h) \quad (9)$$

where V, W_m, W_h are weight parameters of the model, b_m, b_h are biases, and m_{ij} is a representation of the span from position i to j . It is calculated with the contextualized representations of the starting and ending words h_i, h_j , and the average of the word embeddings \tilde{x}_{ij} :

$$(h_1, \dots, h_M) = \text{BiLSTM}(X) \quad (10)$$

$$\tilde{x}_{ij} = \frac{1}{j-i} \sum_{k=i}^j x_k \quad (11)$$

The complexity of this model is quadratic with respect to the number of words. However, it can be parallelized at training and decoding time. Lee et al. (2017) uses an attention function over the embeddings instead of an average. That approach is less memory efficient and requires the maximum length of spans as a hyperparameter. Also, they include embeddings of the span lengths which are learned during training. As shown in the experimental part, these components do not improve the performance of our model.

4 Partially annotated data

The partial annotation of coreference data for mention detection means that not labeled spans may be true mentions of entities. Thus, the approach of treating spans without mention annotations as true negative examples would be incorrect. On the other hand, the ideal solution of sampling all possible mention annotations, which are consistent with the given partial annotation, would be intractable. We want to modify the model’s loss function in such a way that, if the system predicts a false-positive, the loss is reduced. This encourages the model to favor recall over precision by predicting more mention-like spans, even when they are not labeled. We assume that it is possible to learn the true mention distribution using the annotated mention samples by extrapolating the non-annotated mentions, and we propose two ways to encourage the model to do so.

Weighted loss function: We use a weighted loss function with weight $w \in \{0, 1\}$ for negative examples only. The *sequence tagging* model makes word-wise decisions; thus, we consider words tagged as “out of mention”, $y_t = \text{“-”}$, as negative examples, while the rest are positives. Although this simplification has the potential to increase inconsistencies, e.g., having non-ending or overlapping mentions, we observe that the LSMT-based model can capture the simple grammar of the tag labels with very few mistakes. For *span scoring*, the distinction between negative and positive examples is clear, given that the decisions are made for each span.

Soft-target classification: Soft-targets allow us to have a distribution over all classes instead of having a single class annotation. Thus, we applied soft-targets to negative examples to reflect the probability that they could actually be positive ones. For *sequence tagging*, we set the target of negative examples, $y_t = \text{“-”}$, to $(\rho, \rho, \rho, 1 - 3\rho)$ corresponding to the classes $([,], +, -)$. For *span scoring*, we change the target of negative examples to $y_{neg} = \rho$. In both cases, ρ is the probability of the example being positive.

5 Coreference Resolution

We use multitask learning to train the mention detection together with coreference resolution. The weights to sum the loss functions of each task are estimated during training, as in Cipolla et al. (2018). The sentence encoder is shared, and the output of mention detection serves as input to coreference resolution. We use the coreference resolver proposed by Lee et al. (2017). It uses a pair-wise scoring function s between a mention m_k and each of its candidate antecedents m_a , defined as:

$$s(m_k, m_a) = s_c(m_k, m_a) + s_m(m_k) + s_m(m_a) \quad (12)$$

where s_c is a function that assesses whether two mentions refer to the same entity. We modified the mention detection score s_m .

For the *sequence tagging* approach, the function s_m serves as a bias value and it is calculated as:

$$s_m = v \cdot P(y_{t_i} = \text{“[”}) \cdot P(y_{t_j} = \text{“]”}) \quad (13)$$

where y_{t_i} and y_{t_j} are the labels of the first and last words of the span, and v is a scalar parameter learned during training. At test time, only mentions in the one-best output of the mention detection model are candidate mentions for the coreference resolver. During training, the set of candidate mentions includes both the spans detected by the mention detection model and the ground truth mentions. The mention decoder is run for one pass with ground-truth labels in the conditional part of the probability function (Eq. 2), to get the mention detection loss, and run for a second pass with predicted labels to provide input for the coreference task and compute the coreference loss.

For the *span scoring* approach, s_m is a function of the probability defined in Eq. 8, scaled by a parameter v learned during training.

$$s_m = v \cdot P(m_{i,j}) \quad (14)$$

Model	Rec.	Prec.	F1
Sequence tagging	73.7	77.5	75.6
Span scoring	72.7	79.2	75.8
+ span size emb.	71.6	80.1	75.6
- avg. emb. + att. emb.	72.1	78.9	75.4

Table 1: Mention detection evaluation

Instead of the end-to-end objective of Lee et al. (2017), we use a multitask objective, which adds the loss function of mention detection. We do not prune mentions with a maximum length, nor impose any maximum number of mentions per document. We use the probability of the mention detector with a threshold of τ for pruning.

6 Experiments and Results

We evaluate our model on the English OntoNotes set from the CoNLL 2012 shared-task (Pradhan et al., 2012), which has 2802 documents for training, 343 for development, and 348 for testing. The setup is the same as Lee et al. (2017) for comparison purposes, with the hyper-parameters ρ, w, τ optimized on the development set. We use the average F1 score as defined in the shared-task (Pradhan et al., 2012) for evaluation of mention detection and coreference resolution.

6.1 Mention detection

First, we evaluate our stand-alone mention detectors. For this evaluation, all unannotated mentions are treated as negative examples. Table 1 show the results on the test set with models selected using the best F1 score with $\tau=0.5$, on the development set. We can see that *sequence tagging* performs almost as well as *span scoring* in F1 score, even though the latter is an exhaustive search method. We also evaluate the *span scoring* model with different components from Lee et al. (2017). By adding the span size vector, the precision increases but the recall decreases. Replacing the average embedding \tilde{x} with attention over the embeddings requires a limited span size for memory efficiency, resulting in decreased performance.

6.2 Coreference Resolution

Table 2 shows the results obtained for our multitask systems for coreference resolution and mention detection with and without the loss modification. The *sequence tagging* method obtains lower performance compared to *span scoring*. This result can be attributed to its one-best method to select mentions, in contrast to *span scoring*, where uncertainty is fully integrated with the coreference system. The *span scoring* method performs similarly to the coreference resolution baseline, showing that the naive introduction of a loss for mention detection does not improve performance (although we find it does decrease convergence time). However, adding the modified mention loss does improve coreference performance. For *sequence tagging*, the weighted loss results in higher performance, while for the *span scoring*, soft-targets work best. In both cases, the recall increases with a small decrease in precision, which improves the F1 score of mention detection and improves coreference resolution.

6.3 Recall performance

Figure 3 shows a comparison of the mention detection methods in terms of recall. The unmodified *sequence tagging* model achieves 73.7% recall, and by introducing a weighted loss at $w=0.01$, it reaches 90.5%. The lines show the variation of recall for the *span scoring* method with respect to the detection threshold of τ . The dotted line represents the unmodified model, while the continuous line represents the model with soft-targets at $\rho=0.1$, which shows higher recall for every τ .

7 Related Work

Lee et al. (2017) proposed the first end-to-end coreference resolution that does not require heavy feature engineering for word representations. Their mention detection is done by considering all spans in a

Model	Mention			Coref.
	Rec.	Prec.	F1	Avg. F1
Lee et al. (2017)	–	–	–	67.2
Sequence tagging	73.1	84.9	78.6	59.9
+ wt. loss $w=0.01$	77.3	83.2	80.1	64.1
+ soft-target $\rho=0.1$	74.3	84.0	78.8	61.2
Span scoring	75.3	88.3	81.3	67.0
+ wt. loss $w=0.3$	76.3	88.1	81.8	67.1
+ soft-target $\rho=0.1$	78.4	87.9	82.9	67.6

Table 2: Coreference resolution evaluation (CoNLL 2012)

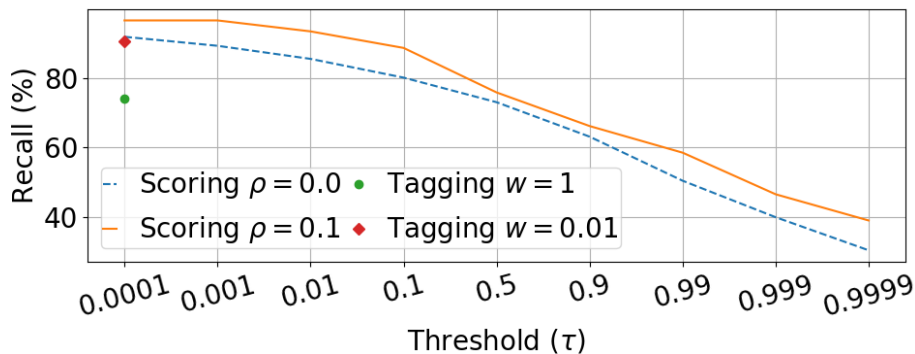


Figure 3: Recall of the *mention scoring* function with respect to the detection threshold τ . Values for the *sequence tagging* are referential

document as the candidate mentions, and the learning signal is coming indirectly from the coreference annotation. Zhang et al. (2018) used a similar approach but introducing a direct learning signal for the mention detection, which is done by adding a loss for mention detection with a scaling factor as hyperparameter. This allows a faster convergence at training time. Lee et al. (2018) proposed a high-order coreference resolution where the mention representation are inferred over several iterations of the model. However, the mention detection part is same as in (Lee et al., 2017). The following studies proposed improvements over this work (Fei et al., 2019; Joshi et al., 2019; Joshi et al., 2020) but maintaining the same method for mention detection.

Name entity recognition has been largely studied in the community. However, many of these models ignored the nested entity names. Katiyar and Cardie (2018) presents a nested named entity recognition model using a recurrent neural network that includes extra connections to handle nested mention detection. Ju et al. (2018) uses stack layers to model the nested mentions, and (Wang et al., 2018) use an stack recurrent network. Lin et al. (2019) proposed a sequence-to-nuggets architecture for nested mention detection. Li et al. (2019) uses pointer networks and adversarial learning. Shibuya and Hovy (2020) uses CRF with a iterative decoder that detect nested mentions from the outer to the inner tags. Yu et al. (2020) use a bi-affine model with a similar method as in (Lee et al., 2017).

8 Conclusion

We investigate two simple techniques to deal with partially annotated data for mention detection and propose two methods to approach it: a Weighted loss function and a soft-target classification. We evaluate them on coreference resolution and mention detection with a multitask learning approach. We show that the techniques effectively increase the recall of mentions and coreference links with a small decrease in precision, thus, improving the F1 score. In the future, we plan to use these methods to maintain coherence over long distances when reading, translating, and generating large text, by keeping track of abstract representations of entities.

Acknowledgements

We are grateful for the support of the Swiss National Science Foundation under the project LAOS, grant number “FNS-30216”

References

- Roberto Cipolla, Yarin Gal, and Alex Kendall. 2018. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7482–7491. IEEE.
- Elizabeth Clark, Yangfeng Ji, and Noah A. Smith. 2018. Neural text generation in stories using entity representations as context. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2250–2260, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Hongliang Fei, Xu Li, Dingcheng Li, and Ping Li. 2019. End-to-end deep reinforcement learning based coreference resolution. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July. Association for Computational Linguistics.
- R Florian, H Hassan, A Ittycheriah, H Jing, N Kambhatla, X Luo, N Nicolov, and S Roukos. 2004. A statistical model for multilingual entity detection and tracking. In *HLT-NAACL 2004: Main Proceedings*, pages 1–8, Boston, Massachusetts, USA, May 2 - May 7. Association for Computational Linguistics.
- Alex Graves and Jürgen Schmidhuber. 2005. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5-6):602–610.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long-short term memory. *Neural Computation*, 9(8):1735–1780.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. BERT for coreference resolution: Baselines and analysis. 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)*, pages 5803–5808, Hong Kong, China, November. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Meizhi Ju, Makoto Miwa, and Sophia Ananiadou. 2018. A neural layered model for nested named entity recognition. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1446–1459, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Arzoo Katiyar and Claire Cardie. 2018. Nested named entity recognition revisited. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 861–871, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California, June. Association for Computational Linguistics.
- 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*, pages 188–197, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-order coreference resolution with coarse-to-fine inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 687–692, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Jing Li, Deheng Ye, and Shuo Shang. 2019. Adversarial transfer for named entity boundary detection with pointer networks. In *IJCAI*, pages 5053–5059.

- Hongyu Lin, Yaojie Lu, Xianpei Han, and Le Sun. 2019. Sequence-to-nuggets: Nested entity mention detection via anchor-region networks. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5182–5192.
- Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. Document-level neural machine translation with hierarchical attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium, October-November. Association for Computational Linguistics.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011, Suntec, Singapore, August. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. Conll-2012 shared task: Modeling multilingual unrestricted coreference in ontonotes. In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea, July. Association for Computational Linguistics.
- Takashi Shibuya and Eduard Hovy. 2020. Nested named entity recognition via second-best sequence learning and decoding. *Transactions of the Association for Computational Linguistics*, 8:605–620.
- Wee Meng Soon, Hwee Tou Ng, and Daniel Chung Yong Lim. 2001. A machine learning approach to coreference resolution of noun phrases. *Computational Linguistics*, 27(4):521–544.
- Bailin Wang, Wei Lu, Yu Wang, and Hongxia Jin. 2018. A neural transition-based model for nested mention recognition. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1011–1017, Brussels, Belgium, October-November. Association for Computational Linguistics.
- Mingbin Xu, Hui Jiang, and Sedtawut Watcharawittayakul. 2017. A local detection approach for named entity recognition and mention detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1237–1247, Vancouver, Canada, July. Association for Computational Linguistics.
- Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. Named entity recognition as dependency parsing. *arXiv preprint arXiv:2005.07150*.
- Rui Zhang, Cicero Nogueira dos Santos, Michihiro Yasunaga, Bing Xiang, and Dragomir Radev. 2018. Neural coreference resolution with deep biaffine attention by joint mention detection and mention clustering. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 102–107, Melbourne, Australia, July. Association for Computational Linguistics.