

Event Detection: Gate Diversity and Syntactic Importance Scores for Graph Convolution Neural Networks

Viet Dac Lai¹, Tuan Ngo Nguyen¹ and Thien Huu Nguyen^{1,2}

¹ Department of Computer and Information Science, University of Oregon,
Eugene, OR 97403, USA

² VinAI Research, Vietnam

{vietl, tnguyen, thien}@cs.uoregon.edu

Abstract

Recent studies on event detection (ED) have shown that the syntactic dependency graph can be employed in graph convolution neural networks (GCN) to achieve state-of-the-art performance. However, the computation of the hidden vectors in such graph-based models is agnostic to the trigger candidate words, potentially leaving irrelevant information for the trigger candidate for event prediction. In addition, the current models for ED fail to exploit the overall contextual importance scores of the words, which can be obtained via the dependency tree, to boost the performance. In this study, we propose a novel gating mechanism to filter noisy information in the hidden vectors of the GCN models for ED based on the information from the trigger candidate. We also introduce novel mechanisms to achieve the contextual diversity for the gates and the importance score consistency for the graphs and models in ED. The experiments show that the proposed model achieves state-of-the-art performance on two ED datasets.

1 Introduction

Event Detection (ED) is an important task in Information Extraction of Natural Language Processing. The main goal of this task is to identify event instances presented in text. Each event mention is associated with a word or a phrase, called an event trigger, which clearly expresses the event (Walker et al., 2006). The event detection task, precisely speaking, seeks to identify the event triggers and classify them into some types of interest. For instance, consider the following sentences:

- (1) *They'll be **fired** on at the crossing.*
- (2) *She is on her way to get **fired**.*

An ideal ED system should be able to recognize the two words “*fired*” in the sentences as the triggers of the event types “Attack” (for the first

sentence) and “End-Position” (for the second sentence).

The dominant approaches for ED involve deep neural networks to learn effective features for the input sentences, including separate models (Chen et al., 2015) and joint inference models with event argument prediction (Nguyen and Nguyen, 2019). Among those deep neural networks, graph convolutional neural networks (GCN) (Kipf and Welling, 2017) have achieved state-of-the-art performance due to the ability to exploit the syntactic dependency graph to learn effective representations for the words (Nguyen and Grishman, 2018; Liu et al., 2018; Yan et al., 2019). However, two critical issues should be addressed to further improve the performance of such models.

First, given a sentence and a trigger candidate word, the hidden vectors induced by the current GCN models are not yet customized for the trigger candidate. As such, the trigger-agnostic representations in the GCN models might retain redundant/noisy information that is not relevant to the trigger candidate. As the trigger candidate is the focused word in the sentence, that noisy information might impair the performance of the ED models. To this end, we propose to filter the noisy information from the hidden vectors of GCNs so that only the relevant information for the trigger candidate is preserved. In particular, for each GCN layer, we introduce a gate, computed from the hidden vector of the trigger candidate, serving as the irrelevant information filter for the hidden vectors. Besides, as the hidden vectors in different layers of GCNs tend to capture the contextual information at different abstract levels, we argue that the gates for the different layers should also be regulated to exhibit such abstract representation distinction. Hence, we additionally introduce a novel regularization term for the overall loss function to achieve these distinctions for the gates.

Second, the current GCN models fail to consider the overall contextual importance scores of every word in the sentence. In previous GCN models, to produce the vector representation for the trigger candidate word, the GCN models mostly focus on the closest neighbors in the dependency graphs (Nguyen and Grishman, 2018; Liu et al., 2018). However, although the non-neighboring words might not directly carry useful context information for the trigger candidate word, we argue that their overall importance scores/rankings in the sentence for event prediction can still be exploited to provide useful training signals for the hidden vectors in ED. In particular, we propose to leverage the dependency tree to induce a **graph-based** importance score for every word based on its distance to the trigger candidate. Afterward, we propose to incorporate such importance scores into the ED models by encouraging them to be consistent with another set of **model-based** importance scores that are computed from the hidden vectors of the models. Based on this consistency, we expect that graph-based scores can enhance the representation learning for ED. In our experiments, we show that our method outperforms the state-of-the-art models on the benchmark datasets for ED.

2 Related Work

Prior studies on ED involve handcrafted feature engineering for statistical models (Ahn, 2006; Ji and Grishman, 2008; Hong et al., 2011; Li et al., 2013; Mitamura et al., 2015) and deep neural networks, e.g., CNN (Chen et al., 2015, 2017; Nguyen and Grishman, 2015; Nguyen et al., 2016g), RNN (Nguyen et al., 2016; Jagannatha and Yu, 2016; Feng et al., 2016), attention mechanism (Liu et al., 2017; Chen et al., 2018), contextualized embeddings (Yang et al., 2019), and adversarial training (Wang et al., 2019). The last few years witness the success of graph convolutional neural networks for ED (Nguyen and Grishman, 2018; Liu et al., 2018; Veyseh et al., 2019; Yan et al., 2019) where the dependency trees are employed to boost the performance. However, these graph-based models have not considered representation regulation for GCNs and exploiting graph-based distances as we do in this work.

3 Model

Task Description: The goal of ED consists of identifying trigger words (**trigger identification**)

and classifying them for the event types of interest (**event classification**). Following the previous studies (Nguyen and Grishman, 2015), we combine these two tasks as a single multi-way classification task by introducing a *None* class, indicating non-event. Formally, given a sentence $X = [x_1, x_2, \dots, x_n]$ of n words, and an index t ($1 \leq t \leq n$) of the trigger candidate x_t , the goal is to predict the event type y^* for the candidate x_t . Our ED model consists of three modules: (1) Sentence Encoder, (2) GCN and Gate Diversity, and (3) Graph and Model Consistency.

Sentence Encoder: We employ the pre-trained BERT (Devlin et al., 2019) to encode the given sentence X . In particular, we create an input sequence of $[[CLS], x_1, \dots, x_n, [SEP], x_t, [SEP]]$ where $[CLS]$ and $[SEP]$ are the two special tokens in BERT. The word pieces, tokenized from the words, are fed to BERT to obtain the hidden vectors of all layers. We concatenate the vectors of the top M layers to obtain the corresponding hidden vectors for each word piece, where M is a hyper-parameter. Then, we obtain the representation of the sentence $E = \{e_1, \dots, e_n\}$ in which the vectors e_i of x_i is the average of layer-concatenated vectors of its word pieces. Finally, we feed the embedding vectors in E to a bidirectional LSTM, resulting in a sequence of hidden vectors $h^0 = \{h_1^0, \dots, h_n^0\}$.

GCN and Gate Diversity: To apply the GCN model, we first build the sentence graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ for X based on its dependency tree, where \mathcal{V}, \mathcal{E} are the sets of nodes and edges, respectively. \mathcal{V} has n nodes, corresponding to the n words X . Each edge (x_i, x_j) in \mathcal{E} amounts to a directed edge from the head x_i to the dependent x_j in the dependency tree. Following (Marcheggiani and Titov, 2017), we also include the opposite edges of the dependency edges and the self-loops in \mathcal{E} to improve the information flow in the graph.

Our GCN module contains L stacked GCN layers (Kipf and Welling, 2017), operating over the sequence of hidden vectors h^0 . The hidden vector h_i^l ($1 \leq i \leq n, 1 \leq l \leq L$) of the word x_i at the l -th layer is computed by averaging the hidden vectors of neighboring nodes of x_i at the $(l-1)$ -th layer. Formally, h_i^l is computed as follow:

$$h_i^l = \text{ReLU} \left(W^l \sum_{(x_i, x_j) \in \mathcal{E}} \frac{h_j^{l-1}}{|\{x_j\}|} \right) \quad (1)$$

where W^l is a learnable weight of the GCN layer.

The major issue of the current GCN for ED is that its hidden vectors h_i^l are induced without special awareness of the trigger candidate x_t . This might result in irrelevant information (for the trigger word candidate) in the hidden vectors of GCNs for ED, thus hindering further performance improvement. To address this problem, we propose to filter that unrelated information by introducing a gate for each GCN layer. The vector g^l for the gate at the l -th layer is computed from the embedding vector e_t of the trigger candidate: $g^l = \sigma(W_g^l e_t)$ where W_g^l are learnable parameters for the l -th layer. Then, we apply these gates over the hidden vectors of the corresponding layer via the element-wise product, resulting in the filtered vectors: $m_i^l = g^l \circ h_i^l$.

As each layer in the GCN module has access to a particular degree of neighbors, the contextual information captured in these layers is expectedly distinctive. Besides, the gates for these layers control which information is passed through, therefore, they should also demonstrate a certain degree of contextual diversity. To this end, we propose to encourage the distinction among the outcomes of these gates once they are applied to the hidden vectors in the same layers. Particularly, starting with the hidden vectors h^l of the l -layer, we apply the gates g^k (for all $(1 \leq k \leq L)$) to the vectors in h^l , which results in a sequence of filtered vectors $\bar{m}_i^{k,l} = g^k \circ h_i^l$. Afterward, we aggregate the filtered vectors obtained by the same gates using max-pooling: $\bar{m}^{k,l} = \max_pool(\bar{m}_1^{k,l}, \dots, \bar{m}_n^{k,l})$. To encourage the gate diversity, we enforce vector separation between $\bar{m}^{l,l}$ with all the other aggregated vectors from the same layer l (i.e., $\bar{m}^{k,l}$ for $k \neq l$). As such, we introduce the following cosine-based regularization term \mathcal{L}_{GD} (for Gate Diversity) into the overall loss function:

$$\mathcal{L}_{GD} = \frac{1}{L(L-1)} \sum_{l=1}^L \sum_{k=l+1}^L \text{cosine}(\bar{m}^{l,l}, \bar{m}^{l,k}) \quad (2)$$

Note that the rationale for applying the gates g^k to the hidden vectors h^l for the gate diversity is to ground the control information in the gates to the contextual information of the sentence in the hidden vectors to facilitate meaningful context-based comparison for representation learning in ED.

Graph and Model Consistency: As stated above, we seek to supervise the model using the knowledge from the dependency graph. Inspired by the contextual importance of the neighboring

words for the event prediction of the trigger candidate x_t , we compute the **graph-based importance scores** $P = p_1, \dots, p_n$ in which p_i is the negative distance from the word x_i to the trigger candidate.

In contrast, the model-based importance scores for each word x_i is computed based on the hidden vectors of the models. In particular, we first form an overall feature vector V_t that is used to predict the event type for x_t via:

$$V_t = [e_t, m_t^L, \max_pool(m_1^L, \dots, m_n^L)]$$

In this work, we argue that the hidden vector of an important word in the sentence for ED should carry more useful information to predict the event type for x_t . Therefore, we consider a word x_i as more important for the prediction of the trigger candidate x_t if its representation m_i^L is more similar to the vector V_t . We estimate the **model-based important scores** for every word x_i with respect to the candidate x_t as follow:

$$q_i = \sigma(W^v V_t) \cdot \sigma(W^m m_i^L) \quad (3)$$

where W^v and W^m are trainable parameters.

Afterward, we normalize the scores P and $Q = q_1, \dots, q_n$ using the softmax function. Finally, we minimize the KL divergence between the graph-based important scores P and the model-based importance scores Q by injecting a regularization term \mathcal{L}_{ISC} (for the graph-model Importance Score Consistency) into the overall loss function:

$$\mathcal{L}_{ISC}(P, Q) = - \sum_{i=1}^n p_i \frac{p_i}{q_i} \quad (4)$$

To predict the event type, we feed V_t into a fully connected network with softmax function in the end to estimate the probability distribution $P(\hat{y}|X, t)$. To train the model, we use the negative log-likelihood as the classification loss $\mathcal{L}_{CE} = -\log P(y^*|X, t)$. Finally, we minimize the following combined loss function to train the proposed model:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{GD} + \beta \mathcal{L}_{ISC} \quad (5)$$

where α and β are trade-off coefficients.

4 Experiments

Datasets: We evaluate our proposed model (called GatedGCN) on two ED datasets, i.e., ACE-2005 and Litbank. **ACE-2005** is a widely used

benchmark dataset for ED, which consists of 33 event types. In contrast, **Litbank** is a newly published dataset in the literature domain, annotating words with two labels *event* and *none-event* (Sims et al., 2019). Hence, on Litbank, we essentially solve trigger identification with a binary classification problem for the words.

As the sizes of the ED dataset are generally small, the pre-processing procedures (e.g., tokenization, sentence splitting, dependency parsing, and selection of negative examples) might have a significant effect on the models’ performance. For instance, the current best performance for ED on ACE-2005 is reported by (Yang et al., 2019) (i.e., 80.7% F1 score on the test set). However, once we re-implement this model and apply it to the data version pre-processed and provided by the prior work (Nguyen and Grishman, 2015, 2018), we are only able to achieve an F1 score of 76.2% on the test set. As the models share the way to split the data, we attribute such a huge performance gap to the difference in data pre-processing that highlights the need to use the same pre-processed data to measure the performance of the ED models. Consequently, in this work, we employ the exact data version that has been pre-processed and released by the early work on ED for ACE-2005 in (Nguyen and Grishman, 2015, 2018) and for Litbank in (Sims et al., 2019).

The hyper-parameters for the models in this work are tuned on the development datasets, leading to the following selected values: one layer for the BiLSTM model with 128 hidden units in the layers, $L = 2$ for the number of the GCN layers with 128 dimensions for the hidden vectors, 128 hidden units for the layers of all the feed-forward networks in this work, and $5e-5$ for the learning rate of the Adam optimizer. These values apply for both the ACE-2005 and Litbank datasets. For the trade-off coefficients α and β in the overall loss function, we use $\alpha = 0.1$ and $\beta = 0.2$ for the ACE dataset while $\alpha = 0.3$ and $\beta = 0.2$ are employed for Litbank. Finally, we use the case BERT_{base} version of BERT and freeze its parameters during training in this work. To obtain the BERT representations of the word pieces, we use $M = 12$ for ACE-2005 and $M = 4$ for Litbank (Sims et al., 2019).

Results: We compare our model with two classes of baselines on ACE-2005. The first class includes the models with non-contextualized

Model	P	R	F
CNN	71.8	66.4	69.0
NCNN	-	-	71.3
GCN-ED	77.9	68.8	73.1
DMBERT	79.1	71.3	74.9
BERT+MLP	77.8	74.6	76.2
GatedGCN (Ours)	78.8	76.3	77.6

Table 1: Performance on the ACE-2005 test set.

embedding, i.e., **CNN**: a CNN model (Nguyen and Grishman, 2015), **NCNN**: non-consecutive CNN model: (Nguyen and Grishman, 2016), and **GCN-ED**: a GCN model (Nguyen and Grishman, 2018). Note that these baselines use the same pre-processed data like ours. The second class of baselines concern the models with the contextualized embeddings, i.e., **DMBERT**: a model with dynamic pooling (Wang et al., 2019) and **BERT+MLP**: a MLP model with BERT (Yang et al., 2019). These models currently have the best-reported performance for ED on ACE-2005. Note that as these works employ different pre-processed versions of ACE-2005, we re-implement the models and tune them on our dataset version for a fair comparison. For Litbank, we use the following baselines reported in the original paper (Sims et al., 2019): **BiLSTM**: a BiLSTM model with word2vec, **BERT+BiLSTM**: a BiLSTM model with BERT, and **DMBERT** (Wang et al., 2019).

Table 1 presents the performance of the models on the ACE-2005 test set. This table shows that GatedGCN outperforms all the baselines with a significant improvement of 1.4% F1-score over the second-best model BERT+MLP. In addition, Table 2 shows the performance of the models on the Litbank test set. As can be seen, the proposed model is better than all the baseline models with 0.6% F1-score improvement over the state-of-the-art model BERT+BiLSTM. These improvements are significant on both datasets ($p < 0.05$), demonstrating the effectiveness of GatedGCN for ED.

Ablation Study: The proposed model involves three major components: (1) the **Gates** to filter irrelevant information, (2) the Gate **Diversity** to encourage contextual distinction for the gates, and (3) the **Consistency** between graph and model-based importance scores. Table 3 reports the ablation study on the ACE-2005 development set when the components are incrementally removed from the full model (note that eliminating **Gate** also removes **Diversity** at the same time). As can be seen,




Figure 1: Visualization of the model-based importance scores computed by the proposed model for several GatedGCN-successful examples. The words with bolder colors have larger importance scores in this case. Note that the golden event types “*Movement:Transport*” and “*Conflict:Attack*” are written under the trigger words in the sentences. Also, below each word in the sentences, we indicate the number of the words along the path from that word to the trigger word (i.e., the distances used in the graph-based importance scores).

Model	P	R	F
BiLSTM	70.4	60.7	65.2
+ document context	74.2	58.8	65.6
+ sentence CNN	71.6	56.4	63.1
+ subword CNN	69.2	64.8	66.9
DMBERT	65.0	76.7	70.4
BERT+BiLSTM	75.5	72.3	73.9
GatedGCN (Ours)	69.9	79.8	74.5

Table 2: Performance on the Litbank test set.

excluding any component results in significant performance reduction, clearly testifying to the benefits of the three components in the proposed model for ED.

Importance Score Visualization: In order to further demonstrate the operation of the proposed model GatedGCN for ED, we analyze the model-based importance scores for the words in test set sentences of ACE-2005 that can be correctly predicted by GatedGCN, but leads to incorrect predictions for the ablated model “-Gate-Consistency” in Table 3 (called the GatedGCN-successful examples). In particular, Figure 1 illustrates the model-based importance scores for the words in the sentences of several GatedGCN-successful examples. Among others, we find that although the trigger words are directly connected to several words (including the irrelevant ones) in these sentences, the **Gates**, **Diversity**, and **Consistency** components in GatedGCN help to better highlight the most informative words among those neighboring words by assigning them larger importance scores. This enables the representation aggregation mechanism in GCN to learn better hidden vectors, leading to improved performance for ED in this case.

5 Conclusion

We demonstrate how gating mechanisms, gate diversity, and graph structure can be used to integrating syntactic information and improve the hid-

Model	P	R	F
GatedGCN (full)	76.7	70.5	73.4
-Diversity	78.5	67.0	72.3
-Consistency	80.5	64.7	71.7
-Diversity -Consistency	79.0	63.0	70.1
-Gates	77.8	65.3	71.3
-Gates -Consistency	83.0	62.5	71.0

Table 3: Ablation study on the ACE-2005 dev set.

den vectors for ED models. The proposed model achieves state-of-the-art performance on two ED datasets. In the future, we plan to apply the proposed model for the related tasks and other settings of ED, including new type extension (Nguyen et al., 2016b; Lai and Nguyen, 2019), and few-shot learning (Lai et al., 2020a,b).

Acknowledgement

This research has been supported in part by VinGroup Innovation Foundation (VINIF) in project code VINIF.2019.DA18 and Adobe Research Gift. This research is also based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2019-19051600006 under the Better Extraction from Text Towards Enhanced Retrieval (BETTER) Program. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, the Department of Defense, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This document does not contain technology or technical data controlled under either the U.S. International Traffic in Arms Regulations or the U.S. Export Administration Regulations.

References

- David Ahn. 2006. The stages of event extraction. In *Proceedings of the Workshop on Annotating and Reasoning about Time and Events*.
- Yubo Chen, Shulin Liu, Xiang Zhang, Kang Liu, and Jun Zhao. 2017. Automatically labeled data generation for large scale event extraction. In *ACL*.
- Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. Event extraction via dynamic multi-pooling convolutional neural networks. In *ACL-IJCNLP*.
- Yubo Chen, Hang Yang, Kang Liu, Jun Zhao, and Yantao Jia. 2018. Collective event detection via a hierarchical and bias tagging networks with gated multi-level attention mechanisms. In *EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL:HLT*, pages 4171–4186.
- Xiaocheng Feng, Lifu Huang, Duyu Tang, Heng Ji, Bing Qin, and Ting Liu. 2016. A language-independent neural network for event detection. In *ACL (Volume 2: Short Papers)*, volume 2, pages 66–71.
- Yu Hong, Jianfeng Zhang, Bin Ma, Jianmin Yao, Guodong Zhou, and Qiaoming Zhu. 2011. Using cross-entity inference to improve event extraction. In *ACL*.
- Abhyuday N Jagannatha and Hong Yu. 2016. Bidirectional rnn for medical event detection in electronic health records. In *NAACL*.
- Heng Ji and Ralph Grishman. 2008. Refining event extraction through cross-document inference. In *ACL*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Viet Dac Lai, Franck Deroncourt, and Thien Huu Nguyen. 2020a. Exploiting the matching information in the support set for few shot event classification. In *Proceedings of the 24th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*.
- Viet Dac Lai, Franck Deroncourt, and Thien Huu Nguyen. 2020b. Extensively matching for few-shot learning event detection. In *Proceedings of the 1st Joint Workshop on Narrative Understanding, Storylines, and Events (NUSE) at ACL 2020*.
- Viet Dac Lai and Thien Huu Nguyen. 2019. Extending event detection to new types with learning from keywords. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019) at EMNLP 2019*.
- Qi Li, Heng Ji, and Liang Huang. 2013. Joint event extraction via structured prediction with global features. In *ACL*.
- Shulin Liu, Yubo Chen, Kang Liu, and Jun Zhao. 2017. Exploiting argument information to improve event detection via supervised attention mechanisms. In *ACL*.
- Xiao Liu, Zhunchen Luo, and He-Yan Huang. 2018. Jointly multiple events extraction via attention-based graph information aggregation. In *EMNLP*, pages 1247–1256.
- Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. In *EMNLP*.
- Teruko Mitamura, Zhengzhong Liu, and Eduard Hovy. 2015. Overview of tac kbp 2015 event nugget track. In *TAC*.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint event extraction via recurrent neural networks. In *NAACL*.
- Thien Huu Nguyen, Lisheng Fu, Kyunghyun Cho, and Ralph Grishman. 2016b. A two-stage approach for extending event detection to new types via neural networks. In *Proceedings of the 1st ACL Workshop on Representation Learning for NLP (ReplANLP)*.
- Thien Huu Nguyen and Ralph Grishman. 2015. Event detection and domain adaptation with convolutional neural networks. In *ACL-IJCNLP*.
- Thien Huu Nguyen and Ralph Grishman. 2016. Modeling skip-grams for event detection with convolutional neural networks. In *EMNLP*.
- Thien Huu Nguyen and Ralph Grishman. 2018. Graph convolutional networks with argument-aware pooling for event detection. In *AAAI*.
- Thien Huu Nguyen, Adam Meyers, and Ralph Grishman. 2016g. New york university 2016 system for kbp event nugget: A deep learning approach. In *Proceedings of Text Analysis Conference (TAC)*.
- Trung Minh Nguyen and Thien Huu Nguyen. 2019. One for all: Neural joint modeling of entities and events. In *AAAI*.
- Matthew Sims, Jong Ho Park, and David Bamman. 2019. Literary event detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634.
- Amir Poursan Ben Veyseh, Thien Huu Nguyen, and Dejing Dou. 2019. Graph based neural networks for event factuality prediction using syntactic and semantic structures. In *ACL*.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57.

Xiaozhi Wang, Xu Han, Zhiyuan Liu, Maosong Sun, and Peng Li. 2019. Adversarial training for weakly supervised event detection. In *NAACL-HLT*, pages 998–1008.

Haoran Yan, Xiaolong Jin, Xiangbin Meng, Jiafeng Guo, and Xueqi Cheng. 2019. Event detection with multi-order graph convolution and aggregated attention. In *EMNLP-IJCNLP*, pages 5770–5774.

Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. Exploring pre-trained language models for event extraction and generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5284–5294.