

Sentence Structure and Word Relationship Modeling for Emphasis Selection *

Haoran Yang, Wai Lam

The Chinese University of Hong Kong
{hryang, wlam} @se.cuhk.edu.hk

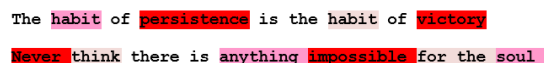
Abstract

Emphasis Selection is a newly proposed task which focuses on choosing words for emphasis in short sentences. Traditional methods only consider the sequence information of a sentence while ignoring the rich sentence structure and word relationship information. In this paper, we propose a new framework that considers sentence structure via a sentence structure graph and word relationship via a word similarity graph. The sentence structure graph is derived from the parse tree of a sentence. The word similarity graph allows nodes to share information with their neighbors since we argue that in emphasis selection, similar words are more likely to be emphasized together. Graph neural networks are employed to learn the representation of each node of these two graphs. Experimental results demonstrate that our framework can achieve superior performance.

1 Introduction

Emphasis Selection recently proposed by (Shirani et al., 2019) aims to select candidate words for emphasis in short sentences. By emphasizing words, people’s intent can be better conveyed, which is useful in a variety of applications. For example, it can be used in spoken language processing to generate more expressive sentences and be used to enable automated design assistance in authoring, i.e., labeling important parts in a paragraph or in a poster title. Although it seems that this task is highly similar to the task of keyword extraction (Gupta, 2017), these two tasks are fundamentally different. The first difference is that keyword extraction focuses on a paragraph which is composed of multiple sentences while emphasis selection aims to choose words from a short sentence. This difference implies that modeling

The work described in this paper is substantially supported by a grant from the Direct Grant of the Faculty of Engineering, CUHK (Project Code: 4055093).



The habit of persistence is the habit of victory
Never think there is anything impossible for the soul

Figure 1: Two examples of the emphasis selection. Words with darker background indicate that more people agree to emphasize.

sentence structure is more effective in emphasis selection. The second difference is that many global word statistics methods employed in keyword extraction such as TF-IDF and word co-occurrence frequency will not work in this task, because for short sentences, it is meaningless to count word frequency and whether the word should be emphasized has nothing to do with the frequency of the word. In addition, keyword extraction requires that the collected keywords are diverse, which means that if two words have similar meaning, only one should be kept. However, in emphasis selection, similar words tend to be emphasized together. Emphasis selection also shares some resemblances with entity recognition (Yadav and Bethard, 2018). But one major difference is that the parts of speech of emphasized words are more diverse and the relation of adjacent words is weaker in the emphasis selection task.

Generally speaking, emphasis selection can be modeled as a sequence classification task where the input is a sentence and the output is each word’s probability to be emphasized. Shirani et al. (Shirani et al., 2019) propose a model which is based on the Recurrent Neural Network (Mikolov et al., 2010) and KL-Divergence loss function. Despite the fact that it looks like a straightforward task, there still exist some challenges. The first challenge is about how to incorporate sentence structure information into the model. Sentence structure information includes what role (subject, predicate, object, etc.) the word plays as well as the position of the word in a sentence. Obviously, this kind of information is very useful. Existing works (Shi-

rani et al., 2019) fail to model the global structure of a sentence. The second challenge is that there is no given context except a short sentence, so it requires the model to be able to capture some common patterns or regularities of most people. More concretely, if two words are similar, they are more likely to be emphasized together. For example, in Figure 1, **persistence** and **victory** are more likely to be emphasized together. This observation can also be found in the second example: **Never** and **impossible**. Moreover, we analyze the training dataset and get a more concrete understanding of this phenomenon through the following procedures: For each training sentence, we consider the most popular emphasized word called word A. Then, we identify the most similar word called word B to the word A based on GloVe embedding (Pennington et al., 2014). We find that the word B is also emphasized with a higher probability than other words in this sentence and this phenomenon occurs in about 26% of the training dataset. Therefore, modeling this kind of relationship between words definitely can help improve the performance of models.

In this paper, we propose a sentence structure graph to handle the sentence structure issue. Specifically, the sentence structure graph is derived from the parse tree of a sentence which contains useful information for this task. For example, as illustrated in Figure 2, when the path is $S \rightarrow NP \rightarrow PRP$, the word **I** is not inclined to be emphasized since this path indicates that this word is a subject. However, when the path is $S \rightarrow VP \rightarrow S \rightarrow VP \rightarrow NP \rightarrow NN$, the word **basketball** is likely to be emphasized since this word is a noun in a verb phrase. Generally, such sentence structure graph can reveal the role of words in a sentence which is beneficial for emphasis selection. Another important information - word relationship information is captured by a word similarity graph. Through the word similarity graph, words can share information with their neighbours, resulting in similar emphasized probabilities of similar words. Next, graph neural networks (Vaswani et al., 2017; Cai and Lam, 2020; Kipf and Welling, 2017; Wu et al., 2019; Yun et al., 2019; Veličković et al., 2018) which has been demonstrated effective in modeling graph structure data are employed to learn the representation of each node of these two graphs. We conduct extensive experiments based on different word embeddings, i.e., GloVe (Pennington et al., 2014), ELMo (Peters et al., 2018), RoBERTa (Liu et al.,

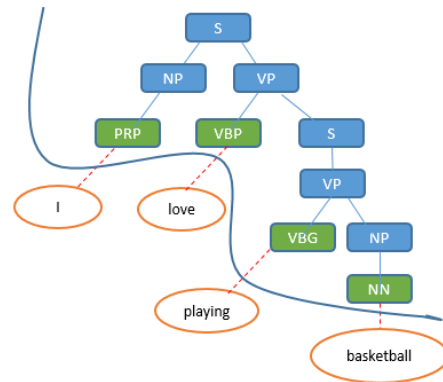


Figure 2: The part above the curve is the sentence structure graph constructed from the sentence: I love playing basketball. After the whole parse tree is encoded, the embeddings of the green nodes are used as the structure information for further classification.

2019) and the experimental results show that our model can achieve superior performance.

2 Related Work

Emphasis selection is a new task proposed by (Shirani et al., 2019) which aims to choose a subset of words to emphasize in a sentence. Shirani et al. (Shirani et al., 2019) propose a model which is based on the Recurrent Neural Network (Mikolov et al., 2010). KL-Divergence loss function is adopted to conduct the label distribution learning (LDL) (Geng and Zhao, 2014). This method achieves competitive performance over the sequence labeling model: CRF (Lafferty et al., 2001).

In Recent years, graph neural networks (Wu et al., 2019; Kipf and Welling, 2017; Yun et al., 2019; Veličković et al., 2018; Cai and Lam, 2020) have demonstrated superiority in modeling the structure of graphs. Kipf et al. (Kipf and Welling, 2017) propose a graph convolutional network which is based on the fourier theory. One drawback of this model is that the edge weight of the graph needs to be known in advance. To overcome this shortcoming, Petar et al. (Veličković et al., 2018) use a masked self-attention layer to calculate the weight of node’s neighbours dynamically and then aggregate information by conducting a weighted addition operation. Currently, graph neural networks are applied to various tasks. Feria et al. (Feria et al., 2018) construct a word graph by calculating the word embedding similarity and apply the community detection algorithm to find different communities. Through the graph, they can find named entities for a bilingual language base in an

unsupervised manner. Sun et al. (Sun et al., 2019) put forward a diverse graph pointer network for keyword extraction. They first construct a word graph based on the distance of two words and then use the graph convolutional network as an encoder to obtain each node’s representation, finally a pointer network decoder and the diverse mechanism are employed to generate diverse keywords. The graph encoder can capture document-level word salience and overcome the long-range dependency problem of RNN.

3 Methodology

We follow the same problem setting given by (Shirani et al., 2019). Suppose a sentence is composed of n words $C = (x_1, x_2, \dots, x_n)$. Our goal is to obtain a subset S of words in C as selected words for emphasis where $1 \leq |S| \leq n$.

We model this task as a prediction problem:

$$(p_1, p_2, \dots, p_n) = \text{model}(x_1, x_2, \dots, x_n) \quad (1)$$

where p_i is i -th word’s probability to be emphasized. Then S contains the top- $|S|$ words with high probability.

Figure 3 depicts the architecture of our proposed model which is composed of three parts: (i) the middle part - sequence encoder (ii) the left part - word similarity graph encoder (iii) the right part - sentence structure graph encoder. Next, we will provide a detailed description of each part.

3.1 Sequence Encoder

The sequence encoder is composed of an embedding layer and a bidirectional GRU. It is mainly used to model the sequence information, i.e., word sequence and tag sequence. Formally, given a sentence $C = (x_1, x_2, \dots, x_n)$ with n words, the embedding layer is responsible for converting each word into a d_1 -dimensional vector and converting the corresponding POS tag into a d_2 -dimensional vector:

$$(w_1, \dots, w_n) = \text{WordEmbed}(x_1, \dots, x_n) \quad (2)$$

$$(e_1, \dots, e_n) = \text{TagEmbed}(t_1, \dots, t_n) \quad (3)$$

where (t_1, \dots, t_n) is the POS tag sequence and $w_i \in R^{d_1}, e_i \in R^{d_2}$. Then the word embedding and the tag embedding are concatenated and fed into an encoder E to encode the sequence information.

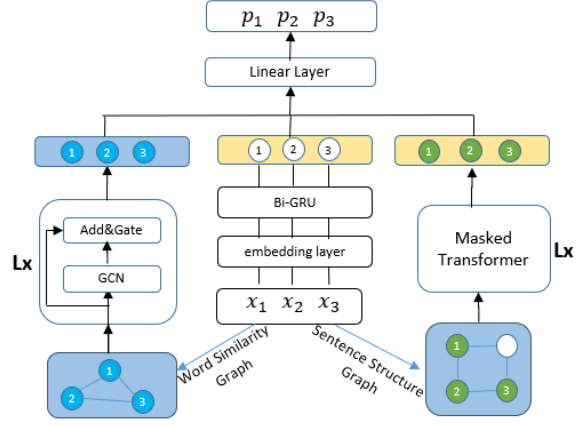


Figure 3: An overview of our model, the left part is the word similarity graph encoder, the middle part is the sequence encoder and the right part is the sentence structure graph encoder. Lx represents that there are L such blocks.

We can obtain the outputted hidden state of the encoder:

$$(h_1, \dots, h_n) = E([w_1, e_1], \dots, [w_n, e_n]) \quad (4)$$

3.2 Word Relationship Modeling

Given a sentence, we take each word as a node and the weight of the edge is calculated by the word embedding similarity. The weight matrix is denoted by $A \in R^{n \times n}$. After the graph is constructed, a L -layer graph convolutional network (GCN) (Kipf and Welling, 2017) is employed to encode the word similarity graph:

$$H^{l+1} = \text{ReLU}(D^{-\frac{1}{2}} A D^{-\frac{1}{2}} H^l W^l) \quad (5)$$

where W^l is a parameter and H^l denotes the nodes’ representation in the l -th layer. $D \in R^{n \times n}$ is a diagonal matrix and $D_{ii} = \sum_j A_{ij}$.

Recall that the WSG is a complete graph since each two words are connected by a weighted edge. There exists a serious problem: Useful information may be overwhelmed by useless information, because a majority number of words do not need to be emphasized, causing the information in words that are not emphasized dominates the words that should be emphasized. To alleviate this problem, we adopt two strategies: residual module (He et al., 2016) and gate mechanism (Gehring et al., 2017; Dauphin et al., 2017). The residual module makes the current node’s representation as the addition

between the former representation and the aggregated information from its neighbours. The gate mechanism controls the magnitude of the aggregated information. Through this way, the current node’s representation will not be significantly affected by its neighbours. Therefore Equation (5) can be rewritten as:

$$M^{l+1} = H^l W^l \quad (6)$$

$$C = s(H^l W^g) \quad (7)$$

$$H^{l+1} = M^{l+1} + D^{-\frac{1}{2}} A D^{-\frac{1}{2}} M^{l+1} \otimes C \quad (8)$$

where $s(\cdot)$ is the sigmoid function and \otimes is the point-wise multiplication.

We obtain H^0 from the word embedding matrix and obtain the L -th layer output $H^L = (w_1^L, \dots, w_n^L)$ as each node’s features of the word similarity graph.

3.3 Sentence Structure Modeling

SSG is constructed by parsing the sentence using NLTK¹ and StanfordNLP². Then, we remove the leaf nodes (which are the words) and the remaining part is the SSG. Each node of the graph is a kind of POS tag and the path from the root to a specific word can reveal what role the word plays in the sentence.

Apparently, the weight of edges is important. For example, in Figure 2, the root node S has two children nodes NP and VP. The edge (S, NP) should have a smaller weight than the edge (S, VP) since people tend not to emphasize the subject in most circumstances. Different from WSG where the weight can be calculated by the word embedding similarity explicitly, it is not appropriate to calculate the weight in the SSG by the node similarity. Hence, we integrate the idea of Transformer (Vaswani et al., 2017; Cai and Lam, 2020) and masked self-attention (Veličković et al., 2018) to the SSG modeling. Firstly, we generate three vectors: key, query, value, according to the current node’s representation:

$$k_i^{l+1} = \mathbf{W}_k^l(v_i^l) \quad (9)$$

$$q_i^{l+1} = \mathbf{W}_q^l(v_i^l) \quad (10)$$

$$v_i^{l+1} = \mathbf{W}_v^l(v_i^l) \quad (11)$$

where $\mathbf{W}_k^l, \mathbf{W}_q^l, \mathbf{W}_v^l$ are parameters. k_i^l, q_i^l, v_i^l correspond to the l -th layer key, query, value vector respectively. v_i^0 is initialized from the tag embedding matrix. Then, a masked self-attention is employed to allow nodes aggregating information only from their neighbours.

$$v_i^{l+1} = \sum_{j \in \mathcal{N}(i)} a_{ij} v_j^{l+1} \quad (12)$$

$$a_{ij} = \frac{\exp(q_i^{l+1} k_j^{l+1})}{\sum_{z \in \mathcal{N}(i)} \exp(q_i^{l+1} k_z^{l+1})} \quad (13)$$

where $\mathcal{N}(i)$ is the neighbour set of the node i . After the graph is encoded with a L -layer network, we obtain the leaf nodes (the green nodes shown in Figure 2) representation $V = (v_1^L, v_2^L, \dots, v_n^L)$.

3.4 Loss Function

After obtaining these three modules’ output, we conduct a concatenation operation and calculate the probability:

$$p_i = \text{softmax}(f([h_i, v_i^L, w_i^L])) \quad (14)$$

where $p_i \in R^3$ is i -th word’s probability distribution. f represents a fully connected neural network. We adopt negative log likelihood as the loss function:

$$L = - \sum_{C \in D_{train}} \sum_{i=1}^{|C|} \log p_{iy_i} \quad (15)$$

4 Experiment and Results

DIY	ideas	for	leafing	up	your	home
B	O	O	O	O	B	I
B	I	O	O	O	O	O
B	O	O	O	O	O	O
O	B	O	O	O	O	O
O	O	O	B	O	O	B
O	O	O	B	I	I	I
O	O	O	B	O	O	O
B	O	O	B	O	O	B
B	I	O	O	O	O	O

Table 1: An example of the labeled dataset

4.1 Dataset

We use the dataset³ provided by (Shirani et al., 2019). The dataset contains 2742 training sentences and 392 test sentences. Each sentence is

¹<https://www.nltk.org/>

²<https://stanfordnlp.github.io/CoreNLP/>

³https://github.com/RiTUAL-UH/SemEval2020_Task10_Emphasis_Selection

Methods	Match-1	Match-2	Match-3	Match-4	Average
GloVe					
CNN	0.541	0.678	0.754	0.805	0.695
RNN (Shirani et al., 2019)	0.536	0.712	0.777	0.811	0.709
Ours	0.569	0.703	0.772	0.813	0.714
Ours w/o WSG	0.563	0.710	0.778	0.810	0.715
Ours w/o SSG	0.561	0.710	0.769	0.811	0.713
ELMo					
CNN	0.574	0.729	0.795	0.832	0.733
RNN-based (Shirani et al., 2019)	0.592	0.752	0.804	0.822	0.743
Ours	0.610	0.768	0.813	0.836	0.757
Ours w/o WSG	0.604	0.742	0.804	0.827	0.744
Ours w/o SSG	0.597	0.753	0.801	0.836	0.747

Table 2: Results of our model and baselines on GloVe and ELMo. The best performance is boldfaced.

	Stay	foolish	to	stay	sane	.
Annotator	0.333(4)	0.889(1/2)	0.222(5/6)	0.444(3)	0.889(1/2)	0.222(5/6)
RNN-based	0.502(3)	0.565(2)	0.227(6)	0.460(4)	0.798(1)	0.357(5)
Ours	0.502(4)	0.784(2)	0.210(6)	0.595(3)	0.805(1)	0.288(5)

Table 3: A sample case. Numbers outside the brackets indicate the word’s probability of being emphasized. Numbers in the brackets are the ranking of the corresponding word. (a/b) means that two words have the same ranking.

labeled by nine annotators. Table 1 gives a sample record of one sentence. B, I, O represent the beginning word to be emphasized, the interior word to be emphasized, and the word not to be emphasized respectively. Since there exists different opinions about whether the word should be emphasized, the labels given by nine annotators are slightly different.

4.2 Experimental Setup

We regard each annotator’s labeling as a sample in the dataset. In other words, each sentence is associate with nine samples. In order to verify the robustness of our model, we conduct experiments on two pre-trained word embeddings: 300- d GloVe (Pennington et al., 2014) and 2048- d ELMo (Peters et al., 2018). For the above two kinds of embeddings, we adopt GRU as the encoder E . The GRU hidden state size is 512 and 1024 respectively. The word similarity graph’s node embedding size is 300 and 2048 respectively. The sentence structure graph’s node embedding size is 300 and 512 respectively. Moreover, we initialize the sentence structure graph’s node embedding by training a classifier which only uses the sentence structure graph encoder. We adopt a two-layer bidirectional GRU. The sentence structure graph and the word similarity graph are encoded by a two-

layer graph neural network. The batch size is set to 16. The negative slope of the ReLU function is set to 0.2. We use the Adam optimizer and the learning rate is 0.0001. The number of epoch is 100. We also add a dropout layer and the dropout rate is 0.5.

Since generalized pretrained language models such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) are demonstrated effective in a large bunch of downstream tasks, we also report results obtained by fine-tuning the RoBERTa on the emphasis selection dataset. There are two different experimental settings. The first setting is that only the RoBERTa model is used as the encoder E . The second setting is that a GRU layer is added on the top of the RoBERTa model, i.e., RoBERTa+GRU is the encoder E . The sentence structure model and the word relationship model remain unchanged. Adam optimizer is adopted and the learning rate is set to $1e-5$.

4.3 Evaluation Metric

We adopt **Match- m** (Shirani et al., 2019) as the evaluation metric which is defined as below:

Match- m : For a sentence C , we choose m words (denoted by $S_m^g(C)$) with the top- m probability (probability of the label B + probability of the label I) in the ground truth and m words (denoted

Methods	Match-1	Match-2	Match-3	Match-4	Average
RoBERTa					
Ours w/o both	0.635	0.756	0.803	0.832	0.757
Ours w/o WSG	0.640	0.775	0.793	0.827	0.759
Ours w/o SSG	0.633	0.760	0.804	0.839	0.759
Ours	0.633	0.779	0.803	0.833	0.762
RoBERTa+GRU					
Ours w/o both	0.607	0.755	0.795	0.822	0.745
Ours w/o WSG	0.602	0.766	0.798	0.825	0.748
Ours w/o SSG	0.607	0.758	0.801	0.837	0.747
Ours	0.600	0.761	0.806	0.838	0.751

Table 4: Results of our model and baselines based on two different architectures, RoBERTa and RoBERTa+GRU. The best performance is boldfaced.

	Thanks	for	showing	me	all	the	best	dance	moves
Annotator	0.444	0.111	0.111	0.111	0.111	0	0.888	0.555	0.444
Ours	0.412	0.057	0.332	0.092	0.063	0.024	0.406	0.599	0.387

Table 5: A failed case. Numbers are the word’s probability of being emphasized.

by $S_m^p(C)$) based on the predicted probability. The formula is defined as:

$$\text{Match-m} = \frac{\sum_{C \in D_{test}} \frac{|S_m^g(C) \cap S_m^p(C)|}{\min(|C|, m)}}{|D_{test}|} \quad (16)$$

4.4 Results and Analysis

4.4.1 Experimental Results

We compare our model with the existing model based on RNN proposed by Shirani et al. (2019) and the convolutional neural network (CNN). We report results evaluated by the metrics Match-1, Match-2, Match-3, Match-4 and the average of these four metrics. From Table 2, we can see that CNN lags behind other models on the whole.

When the word embedding is GloVe, models with at least one graph surpass RNN on almost all the metrics except Match-2. In particular, our model can achieve an improvement on Match-1 and Match-4. Our model without WSG (word similarity graph) achieves an excellent performance on Match-3 and Average. When the word embedding is ELMo, ours is superior to RNN-based on all the evaluation metrics. Compared to these two ablated models, Ours can also achieve better performance. Ours w/o WSG is better than RNN-based on all the evaluation metrics except Match-2 and Ours w/o SSG is better than RNN-based except Match-3. On the whole, models with graphs can obtain better results on most metrics compared to the baseline models, which shows the advantage of these two

components.

Experimental results based on RoBERTa are listed in Table 4. Compared with the results based on GloVe and ELMo, RoBERTa and its variants achieve higher average match score which shows that a better initialized word embedding is helpful for a better performance. For the same RoBERTa encoder, Ours can obtain the highest score on Average and Match-2. For RoBERTa+GRU encoder, Ours can obtain the highest score on Average, Match-3 and Match-4. However, one interesting finding is that RoBERTa encoder performs much better than RoBERTa+GRU encoder. Two possible reasons may interpret this phenomenon. The first reason is the overfitting problem and the second reason is that the larger network is harder to train due to some optimization issues, e.g., gradient vanishing.

4.4.2 Case Study

To gain some insights of our proposed model, we present a sample case generated by the ELMo-based model as shown in Table 3. We can see that Ours not only predicts the ranking accurately, but also obtains very close probability to the ground truth probability derived by annotators. Besides that, the probabilities of foolish and sane predicted by our model are very close than that predicted by RNN-based, which shows that the word similarity graph can impel similar words to have similar probabilities.

We also provide a failed case in Table 5. It is intrinsically harder to rank the words in this sentence even for human beings. Our model does not rank them correctly on these cases where multiple words may be emphasized.

4.4.3 Some Useful Tips

We conclude some tips on the experiment that leads to better performance. (1) We can firstly train a classifier only using the SSG, then use the pre-trained embeddings as an initialization of the sentence graph nodes embeddings. It can obtain higher score and faster convergence of the model. (2) We also consider another method to model the relationships between words using a self-attention operation proposed by Lin et al. (2017) above the hidden vectors of RNN. However, the performance is slightly degraded compared to removing this operation. So we think it is much better to model words relationships and sequence information separately.

5 Conclusions

The sentence structure graph and the word similarity graph are proposed to solve two issues found in emphasis selection. The sentence structure graph helps to model the structure information of the sentence and the word similarity graph is useful in modeling relationships between words. With the development of graph neural network, the two graphs can be properly encoded and integrated into existing models. Experimental results demonstrate that our framework can achieve superior performance.

References

- Deng Cai and Wai Lam. 2020. Graph transformer for graph-to-sequence learning. In *34th AAAI Conference on Artificial Intelligence*.
- Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language modeling with gated convolutional networks. In *34th International Conference on Machine Learning (ICML)*, page 933–941.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Miguel Ferial, Juan Paolo Balbin, and Francis Michael Bautista. 2018. Constructing a word similarity graph from vector based word representation for named entity recognition.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In *34th International Conference on Machine Learning (ICML)*, page 1243–1252.
- Xin Geng and Quan Zhao. 2014. Label distribution learning.
- Er. Tanya Gupta. 2017. Keyword extraction: A review. In *IJEAST*, pages 215–220.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations (ICLR)*.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, pages 282–289.
- Zhouhan Lin, Minwei Feng, Cicero Dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Y. Bengio. 2017. A structured self-attentive sentence embedding. In *5th International Conference on Learning Representations (ICLR)*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *INTER-SPEECH*, pages 1045–1048.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL*, page 2227–2237.
- Amirreza Shirani, Franck Dernoncourt, Paul Asente, Nedim Lipka, Seokhwan Kim, Jose Echevarria, and Tamar Solorio. 2019. Learning emphasis selection for written text in visual media from crowd-sourced label distributions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1167–1172, Florence, Italy. Association for Computational Linguistics.

- Zhiqing Sun, Jian Tang, Pan Du, Zhi-Hong Deng, and Jian-Yun Nie. 2019. Divgraphpointer: A graph pointer network for extracting diverse keyphrases. In *SIGIR*, pages 755–764.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *6th International Conference on Learning Representations (ICLR)*.
- Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2019. A comprehensive survey on graph neural networks. *CoRR*, abs/1901.00596.
- Vikas Yadav and Steven Bethard. 2018. A survey on recent advances in named entity recognition from deep learning models. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2145–2158, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J. Kim. 2019. Graph transformer networks. In *NIPS*, pages 11983–11993.