

# A Token-wise CNN-based Method For Sentence Compression

Weiwei Hou<sup>1,\*</sup>[0000–0002–0758–2870], Hanna Suominen<sup>1,2,3</sup>[0000–0002–4195–1641],  
Piotr Koniusz<sup>1,2</sup>[0000–0002–6340–5289], Sabrina Caldwell<sup>1</sup>[000–0003–0605–3149],  
and Tom Gedeon<sup>1</sup>[0000–0001–8356–4909]

<sup>1</sup> The Australian National University

{weiwei.hou, hanna.suominen, sabrina.caldwell, tom.gedeon}@anu.edu.au

<sup>2</sup> Data61/CSIRO, Australia

piotr.koniusz@data61.csiro.au

<sup>3</sup> University of Turku Finland

**Abstract.** Sentence compression is a *Natural Language Processing* (NLP) task aimed at shortening original sentences and preserving their key information. Its applications can benefit many fields *e.g.*, one can build tools for language education. However, current methods are largely based on *Recurrent Neural Network* (RNN) models which suffer from poor processing speed. To address this issue, in this paper, we propose a token-wise *Convolutional Neural Network*, a CNN-based model along with pre-trained *Bidirectional Encoder Representations from Transformers* (BERT) features for deletion-based sentence compression. We also compare our model with RNN-based models and fine-tuned BERT. Although one of the RNN-based models outperforms marginally other models given the same input, our CNN-based model was ten times faster than the RNN-based approach.

**Keywords:** Neural Networks · NLP · Application

## 1 Introduction

Deletion-based sentence compression refers to the task of extracting key information from a sentence by deleting some of its words. It is often used as an initial step for generating document summaries or as an intermediary step in information retrieval and machine translation. Its applications may also benefit many fields such as e-Learning and language education.

Recent studies focus on adapting methods based on neural networks to solve deletion-based sentence compression as a sequential binary classification problem (see Section 2). *Recurrent Neural Networks* (RNNs) are one of the most popular network architectures that handle sequential tasks. Models such as *Gated Recurrent Units* (GRUs), *Long-Short Term Memory* (LSTM) networks, and BiLSTM

---

\* This work is supported by the Australian National University. This paper is published at ICONIP 2020.

*Bidirectional LSTMs* (BiLSTMs) are found to be suited for this task. However, training these RNN-based models can be time consuming. Applications with poor response speed are the cause of negative user experience.

In contrast, *Convolutional Neural Networks* (CNNs) outperform RNNs in their training speed and reportedly have a similar or better performance than RNN-based models in many tasks [12, 1]. They are widely applied to tasks such as object recognition in Computer Vision. In *Natural Language Processing* (NLP), CNNs have been studied for document summarization, sentiment analysis, and sentence classification, among others. However, the majority of these methods concern the sentence or document level.

In this paper, we apply CNNs to sentence compression at the token level. However, CNNs are weaker at capturing sequential information. To circumvent this issue, we train our model with pre-trained *Bidirectional Encoder Representations from Transformers* (BERT) [4] features. In addition, we also compare our model performance against RNN and BERT fine-tuned models. We test the performance in both the correctness and efficiency.

## 2 Related Work

### 2.1 Recurrent Neural Networks

RNN-based approaches are widely applied to sequential problems such as machine translation, sentence parsing, and image captioning. Inspired by these core NLP tasks, [6] concatenated input sequence and its labels by a key word ‘GO’ as an input to a sequence-to-sequence framework. The goal was to predict sequence labels for each word succeeding ‘GO’. Their network architecture is composed of three layers of LSTMs. On top of the LSTM layers is a softmax layer that produces final outputs. Their method takes on input only word embeddings in the form of the vector representations.

Subsequently, authors of [11] and [13] discovered that adding syntactic information improves the performance. Both works included the *Part-of-Speech* (POS) tags, dependency type, and word embeddings. The results showed significant accuracy improvements. In addition, instead of concatenating embedding sequences and labels, both these studies used hidden vectors to predict labels. The difference between these two methods is that approach [11] uses a framework with three layers of BiLSTMs while approach [13] has a more complex architecture which includes a layer of *Bi-directional RNNs* (BiRNNs) and a layer of Graph Neural Networks. A year later, authors of [14] proposed an approach based on reinforcement learning which includes a BiRNN-based policy network containing two types of actions – REMOVE and RETAIN, where words marked as REMOVE are deleted to obtain a sequence of predicted compression. Such a sequence is then fed into a syntax-based evaluation step that examines the predicted compression according to two rewards. The first reward concerns the fluency of generated sentence against a syntax-based language model. The second reward is based on the comparison of the generated sentence compression

rate with an average compression rate. Their method works well on both large unlabeled and labeled datasets. However, such a model is difficult to train.

## 2.2 Bidirectional Encoder Representations from Transformers

BERT is a language representation model which takes the word position into account to capture a word with its context [4]. Unlike *Global Vectors for Word Representation* (GloVe), Word2Vec, and many other context independent word embeddings, BERT not only embeds semantic information but it also captures the structural information in a sequence. In addition, BERT enables bi-directional prediction. It uses a *Masked Language Model* mechanism that randomly masks a certain percentage of tokens in an input sequence. The objective is to use both the preceding and succeeding context to predict the masked token.

Apart from providing pre-trained language representations, BERT can be fine-tuned on related tasks. It has reportedly achieved the state-of-the-art performance in nine NLP tasks [4]. BERT is said to have also the ability to capture high-level linguistic knowledge (*e.g.*, semantics, syntax, or grammar) [3, 7, 8].

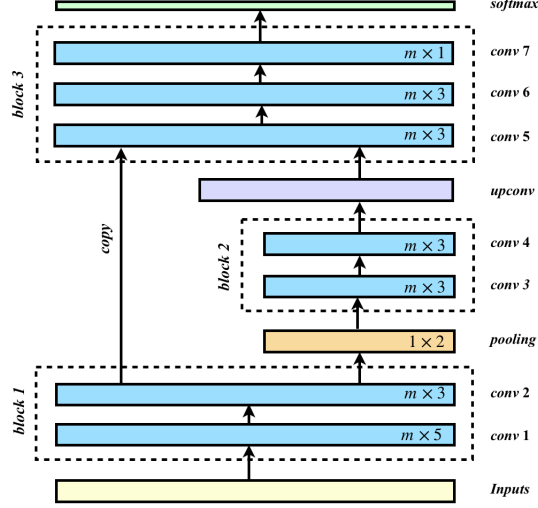
To the best of our knowledge, ours is the first work using BERT for deletion-based sentence compression tasks. We compare pre-trained BERT layers with word embeddings, POS embeddings, and dependency embeddings given the same network architecture to explore whether BERT is able to capture complex syntactic information.

## 3 Method

We define the deletion-based sentence compression task as a token level segmentation task. Specifically, we have given a sequence of  $s = \{w_1, w_2, w_3, \dots, w_i\}$  as an original sentence, where  $i$  is the number of tokens in this sequence, for  $s$ , and we have a corresponding sequence of the mask  $y = \{y_1, y_2, y_3, \dots, y_i\}$ , where  $y_i \in \{0, 1\}$  is the ground truth label of  $w_i$ . Moreover, by zero (one) we mean that a token needs to be deleted (retained) from the original sequence. The goal is to train a model to predict whether  $w_i$  in sequence  $s$  should be deleted or retained.

### 3.1 Network Architectures

Our approach is largely based on U-Net [10], with some differences. The network was originally designed for pixel-level image segmentation tasks. It is a fast and lightweight network which shows extraordinary performance in many image segmentation tasks. The reason we choose U-Net as our base network is that our task is a token level binary classification. Regular CNN cannot capture various levels of coarseness of information and “expand” it back to the token level for segmentation. In addition, we believe that the max-pooling operation can be seen as realizing “compression”.



**Fig. 1.** Graphical illustration of a CNN-based network architecture. Convolutional layers are denoted by *conv 1*, *conv 2*, *conv 3*, ..., *conv 7*.

We adapted the original U-Net network architecture to train a model for a text-based task. As Fig. 1 shows, we assume we have word  $w_i$  in sequence  $s$  where each  $w_i$  has associated with it a  $j$ -dimensional vector  $w_i = \{e_{i,1}, e_{i,2}, e_{i,3}, \dots, e_{i,j}\}$ . This setting yields a matrix of size  $i \times j$  which forms the input to the first layer of our network.

Fig. 1 shows (from the bottom to the top) that the network architecture contains three main blocks. The first block has two standard convolution layers (*conv 1* and *conv 2*). The *conv 1* layer has a  $m \times 5$  kernel;  $m$  is determined by the number of input channels (see Section 4) and *conv 2* has a  $m \times 3$ .

Moreover, between the first and second block, there is a  $1 \times 2$  pooling layer (*pool1*) with the stride size of 2 for down-sampling. The pooling output is fed to the second block which contains two  $m \times 3$  convolution layers (*conv3*, *conv4*). After the above operations, the output is halved in size.

In order to map the output into a token level, we up-sample (deconv) the output from the second block to double back the output size and then the output is concatenated with the previous output from the first block. This operation allows the network to remember “early” features. The third block includes two regular  $m \times 3$  convolutions (*conv5*, *conv6*) and one  $m \times 1$  convolutions (*conv7*). All above convolutions are followed by a *Rectified Linear Unit* (ReLU).

Lastly, we pass the convolution output  $x = \{x_1, x_2, x_3, \dots, x_i\}$  though a token-wise softmax layer to predict the mask  $y_i$  as follows:

$$P(y_i|x_i) = \text{softmax}(Wx_i), \quad (1)$$

where  $W$  is a weight matrix to be learned.

We introduce two main modifications to the original network architecture [10]. The first one is that instead of using the same kernel sizes, our network uses a mix of kernel sizes in order to capture multi-gram information. The second change is that we reduced the depth of the network due to the size of our sentences being much smaller than image sizes from the original work. Hence, empirically, a shallower network tends to perform better.

## 4 Experiments

### 4.1 Data

For our experiments, we use for the GoogleNews dataset, an automatically collected dataset according to approach [5]. The currently released data contains 200,000 sentence pairs, written in ten files. One file includes 10,000 pairs for evaluation while remaining 190,000 pairs are contained in other files.

**GoogleNewsSmall.** For parts of our experiments, we use only the 10,000 pairs based evaluation set. We call it *GoogleNewsSmall*. The reason that we choose *GoogleNewsSmall* is because one of purposes of this study is to compare the performance of CNN- and RNN-based models on deletion-based sentence compression tasks. Therefore, to ensure fairness, *GoogleNewsSmall* includes exactly the same dataset as the previous method settings [11]. Furthermore, we use the first 1,000 sentences as the testing set, the next 1,000 sentences as the validation set, and the remainder as the training set.

**GoogleNewsLarge.** We are also interested in discovering the impact of the training data size on performance. Therefore, in the second setting, we include the entire 200,000 sentence pairs in our experiments. We denote this setting as *GoogleNewsLarge*. For testing and validation, we use the setting already described above. The remaining 198,000 pairs are used as the training set.

### 4.2 Experimental Setup

**BiLSTM+Emb.** In this experiment, we use the same network settings as the one described in [11], which includes a base model architecture with three BiLSTM layers and a dropout layer between them. The input to this network is contains word embeddings only. The purpose of this experiment is to train a base-line RNN-based network.

**CNN+Emb.** In this experiment, we use our method proposed in Section 3.1. For input, we again use word embeddings as the only input to ensure the comparison with RNNs in the same testbed.

**BiLSTM+BERT.** In this experiment, we adapt the setting of **BiLSTM+Emb** by replacing the word embeddings with the last layer of pre-trained BERT outputs. We also record the training time, in seconds, from the point of starting training to full convergence.

**CNN+Multilayer-BERT.** This experiment is composed of a group of sub-experiments. We use the same network architecture as in Section 4.2. We firstly

extract the last four BERT layers (-1, -2, -3, -4) from a pre-trained model. Then, we run the experiments by feeding the network with four different input layer settings: (i) input with the layer -1, (ii) input with the layer -1, and -2, (iii) input with the layer -1, -2, and -3), and (iv) input with all last four layers. We record the training time of experiment with the layer -1 as input, from the point of start of training to the full converge.

### 4.3 Experiments on Different Network Settings

This experiment aims to test the model performance w.r.t. reducing network layers. In the first sub-experiment, we remove one convolution layer from each block (conv 2, 4, and 5 layers from the original settings) in order to test the impact of convolution layers. In the second sub-experiment, we remove pooling, conv 4, 5 of block 2, and upsampling layers. The purpose of this experiment is to compare the performance of a stack of convolution layers and our model.

### 4.4 Experiments w.r.t. the Training Size

This experiment evaluates the performance w.r.t. different sizes of training data. We vary the training size between 8,000 and 19,800. We divide the training progress into ten steps. At each step, the training size increases by 20,000. We detail this experiment in Section 4.2 by only testing on the “last four layers” setting.

The data is labeled as described in approach [5]. Word embeddings are initialised by GloVe 100-dimensional pre-trained embeddings [9]. BERT representations are extracted from the 12/768-uncased-base pre-trained model<sup>4</sup>. All experiments are conducted on a single GPU (Nvidia GeForce GTX Titan XP 12GB). For CNN-based networks, all input sequences have a fixed length of 64<sup>5</sup>. The number of input channels is and layers are the same as in BERT.

### 4.5 Quantitative Evaluations

To evaluate the performance, we use the accuracy and F1 measure. F1 scores are derived from precision and recall values [2] where the precision is defined as the percentage of predicted labels (1 = retained words) that match ground truth labels, and recall is defined as the percentage of labeled retained words in ground truth that overlap with the predictions.

Regarding the training efficiency, we evaluate it as follows. Firstly, we compare both the F1 scores and accuracy of CNN and RNN-based model as a function of each recorded time point. Secondly, we compare the training time of two models.

<sup>4</sup> <https://github.com/google-research/bert>

<sup>5</sup> We concluded that it is the best input size after calculating the mean and max length of all sentences in the entire dataset, and evaluating the efficiency of extracting BERT features.

#### 4.6 Perception-based Evaluations

In order to test the readability of outputs and their relevance to the original inputs, we asked two native English speakers to manually evaluate outputs from the **CNN-Multilayer-BERT** and **BiLSTM-Emb** (baseline). The inputs are the top 200 sentences from the test set. Evaluation methods follow approach [6].

**Table 1.** Results given different experimental settings for models trained on *Google-NewsSmall*.

Method	F1
BiLSTM+Emb	0.74
BiLSTM+BERT	0.79
BiLSTM+SynFeat[11]	0.80
LK-GNN+Emb[13]	0.80
LK-GNN+All Features[13]	<b>0.82</b>
CNN+Emb	0.72
CNN+Multilayer-BERT	0.80

**Table 2.** Results w.r.t. using different layers of BERT features (trained on *GoogleNewsSmall*).

BERT layers	F1	Accuracy
Only -1	0.78	0.80
-1, -2	0.79	0.81
-1, -2, -3	0.80	0.81
-1, -2, -3, -4	<b>0.80</b>	<b>0.82</b>

**Table 3.** Results given different network settings (trained on *GoogleNewsSmall*).

Network Setup	F1
remove conv 2, 4, and 5	0.797
remove pooling, conv 4,5 and upsampling	0.786
original setting (channel size conv4=128, conv6=256)	0.805

## 5 Results

We report the accuracy of different experimental settings in Table 1. We note that **LK-GNN+All Features** achieves the best results. **CNN+Multilayer-BERT** performs the same as **BiLSTM+SynFeat** and **LK-GNN+Emb** (0.80 = 80%). Next, we compare the performance of different models with the same input settings. When model inputs are word embeddings, the table shows that the RNN-based model outperforms the CNN-based one. The results reflect our assumption that CNNs are weaker in capturing sequential information compared to RNNs. However, the performance gap between **CNN+Emb** and the **BiLSTM+Emb** is not significant. Therefore, CNN-based models are a reasonable choice.

Looking at experiments that use the same network architecture but different input settings, we notice that models with BERT used as the input have significantly better performance than the models with Emb inputs. When comparing

BERT to add-on syntactic features, **BiLSTM+BERT** slightly under-performs **BiLSTM+SynFeat** (*ie.*, 1% lower on F1). This implies that BERT captures both syntactic and semantic information. Moreover, when we test multiple layers of BERT features, we can see that **CNN+Multilayer-BERT** performs the same as the model with add-on syntactic features. Therefore, multiple layer BERT features enhance learning ability of the model. To further investigate

**Table 4.** Perception-based Evaluations (trained on *GoogleNewsSmall*).

Method	Readability	Informativeness
BiLSTM+Emb	4.0	3.4
CNN+Multilayer-BERT	4.3	3.6

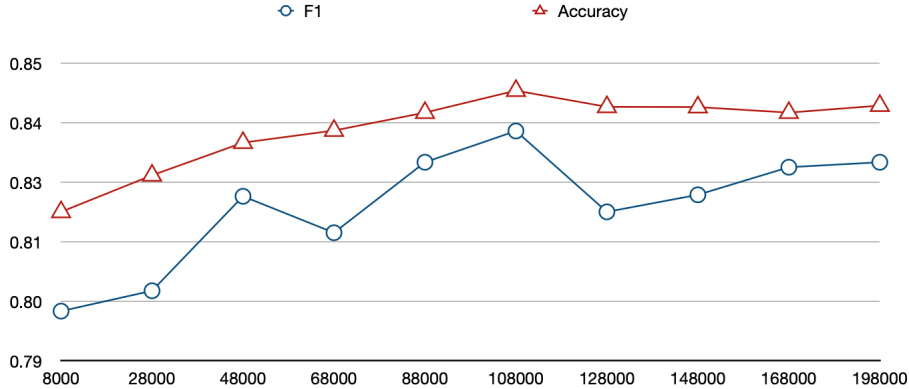
**Table 5.** Training time of CNN- *vs.* LSTM-based models. The first layer of BERT output is used as input to the model.

Run time(s)	CNN+BERT		LSTM+BERT	
	F1	Accuracy	F1	Accuracy
16	0.60	0.73	0.60	0.66
64	0.75	0.73	0.50	0.64
120	0.766	0.79	0.59	0.67
210	0.78	0.80	0.50	0.69
720	-	-	0.70	0.76
1095	-	-	0.73	0.78
1483	-	-	0.75	0.78
1863	-	-	0.75	0.79
2239	-	-	0.78	0.80
2622	-	-	0.79	0.79
3303	-	-	0.80	0.81

BERT, Table 2 shows the impact of using different layers of BERT features on the model performance. The model trained with last four BERT output layers achieves best results (F1 = 0.80, Accuracy = 0.82). However, the model that only uses the last layer of BERT under-performs by 2% (F1 score) and by 1% accuracy. We notice a trend that increasing the number of BERT layers will result in a slightly better performance. This supports our hypothesis that using multiple layers of BERT features improves the performance. Note that we tested multiple layers of BERT features on our CNN based model but not on RNN-based models because concatenating features of four BERT output layers results in an extremely long input (each layer of BERT has a hidden size equal  $4 \times 768$ ). Thus, the computation cost would be very expensive. Moreover, Table 4 shows results for the perception-based evaluations of our model compared to



the baseline. As one can see, the model trained with BERT scores better on readability and informativeness.



**Fig. 2.** Performance of CNN+Multilayer-BERT (last four layers used) in terms of F1 and accuracy. Experiments are conducted on *GoogleNewsLarge*.

The comparisons are demonstrated in Table 5. As we expected, the CNN-based model converged faster when compared to the BiLSTM-based model. It reached its best accuracy performance (F1 = 0.78, Accuracy = 0.81) after 210 seconds. However, while reaching the same F1 score, the BiLSTM-based model needs 10 times longer training time. To outperform the CNN-based models in terms of both the F1 and accuracy, the BiLSTM model took 15 times more time to train. Therefore, our model significantly outperforms BiLSTM models if F1, accuracy and time are taken into account. Table 3 shows the F1 scores for different network settings. As we can see, by removing three convolutional layers, the model slightly under-performs the original model. However, when the pooling, conv 4, conv 5 and upsampling were removed, the results are nearly 2% lower compared to the original settings. This result supports our assumption that pooling help extract compressed information.

In addition, we also investigated the impact of the training data size on the CNN-based model. Figure 2 shows the F1 scores and accuracy of our CNN-BERT model (**CNN+Multilayer-BERT**) w.r.t. different size of training data. We observed that by increasing the training size to approximately 100,000 pairs, the model performs reaches over 0.84 in both the F1 score and accuracy. For more training data, the results show no further improvement and when we increase the training size to be equal to the size of full training set, the results drop to 0.83 (F1 score). We believe that this is caused by the noise in the dataset as authors of approach [14] note that this dataset is automatically labeled based on syntactic-tree-pruning method. Noise can be introduced by syntactic tree parsing errors. Approach [14] scores 0.84 (F1 score) on their LSTM models with the same training data as in our experiments. However, under the same settings,

a model trained on 2 million training pairs reported 0.82 (F1 score) [6]. Authors of approach [14] also argued that this training data may contain errors caused by syntactic tree parsing errors during data annotation. However, we do not evaluate the quality of the data. Our results indicate that the effective training size equals 100,000.

## 6 Discussion

Section 5 showed that, the RNN-based models slightly outperform CNN-based models in their correctness; on average, their F1 scores were 2% higher. However, the CNN-based models performed over ten times faster. In addition, to improve the results of CNN-based models, we adopted BERT features as our networks inputs. The results showed that the model with four layers of BERT features achieved equal performance compared to approach with add-on syntactic features [11] given training size equal 8,000. It implies that multiple layers of BERT capture both the syntactic and semantic information. We argue that the CNN with multiple layers of BERT features was quite a reasonable setting. Since each layer of BERT features has a vector size of 678, concatenating multiple BERT layers as inputs for RNN-based models is computationally prohibitive.

In addition, we also tested the impact of the training data size on our **CNN + Multilayer-BERT** model given four distinct feature settings of BERT, and we found that the F1 scores do not improve further when the training data size reaches approximately 100,000 pairs. We believe that such a result is caused by the noise in the dataset. We believe that such a noise was introduced during the data collecting and labeling process. Similar observations were reported by authors of approaches [13, 14]. In contrast to previous works, we report what is a reasonable trade-off in terms of the dataset size.

In this paper, we did not directly compare our results with the state-of-the-art model [14] for two main reasons. Firstly, our work mainly focused on comparing the performance on different base model settings. Authors of approach [14] proposed a reinforcement learning method, implementing bidirectional RNNs as the base model of the policy network, and this setting is quite similar to approach [11]. Secondly, we tested methods that do not include any domain specific knowledge. One of reward rules in their method uses scoring the sentence compression rates. Since the data was generated by predefined rules, adding such a rule could improve the performance. Although we did not directly compare our results with theirs, we reported results given the training size of 100,000. Our method reaches 84% (F1 score) which is equal to their implemented LSTM model, and only about 1% lower than their reinforcement-based method. We believe if the base model was replaced with our CNN model, the final accuracy would be similar while enjoying faster training.

## 7 Conclusions

In this paper, we studied the performance of CNN-based models for the deletion-based sentence compression task. We first tested the correctness results against the most commonly implemented RNN-based models as well as fine-tuned BERT. Subsequently, we examined the training efficiency of both models. We also compared the results when using a pre-trained BERT language representation as an input to the models with classical word-embeddings and/or other add-on syntactic information.

Our results show that the CNN-based model requires much less training time than the RNN-based model. In addition, the pre-trained BERT language representation highlighted its ability to capture deeper information compared to classical word embedding models. BERT could also serve as a replacement of manually introduced add-on syntactic information. Finally, we observed that increasing the size of training data beyond certain point does not improve the performance further.

Our approach can potentially reduce the cost of building sentence compression applications such as language education tools. Our approach saves computational resources which promotes interactive applications while preserving their accuracy. In the future, we will use our model as a backbone in a reading assistant tool supporting university *English as a Second Language* (ESL) students in their reading activities. We will also continue to study sentence compression with the focus on approaches that can customise the output.

## References

1. Bai S., Kolter Z. J., Koltun V.: An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. Arxiv: <https://arxiv.org/abs/1803.01271> (2018)
2. Chinchor N.: MUC-4 Evaluation Metrics. In Proceedings of the 4th conference on Message understanding. pp 22–29. (1992).
3. Clark K., Khandelwal U., Levy O., Manning C. D.: What Does BERT Look At? An Analysis of BERT’s Attention. Arxiv: <https://arxiv.org/abs/1906.04341> (2019)
4. Devlin J., Chang M. W., Lee K., Toutanova K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. ArXiv: <https://arxiv.org/abs/1810.04805> (2019)
5. Filippova K., Altun Y.: Overcoming the Lack of Parallel Data in Sentence Compression. In Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing. ACL pp 1481–1491 Seattle, Washington (2013).
6. Filippova K., Alfonseca E., Colmenares C. A., Kaiser L., Vinyals O.: Sentence Compression by Deletion with LSTMs. In Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing. ACL, Lisbon, Portugal (2015).
7. Goldberg Y.: Assessing BERT’s Syntactic Abilities. <http://u.cs.biu.ac.il/yogo/bert-syntax.pdf> (2019)
8. Lie C., Andrede A. D., Osama M.: Exploring Multilingual Syntactic Sentence Representations. Arxiv: <https://arxiv.org/abs/1910.11768> (2019).

9. Pennington J., Socher R., Manning C. D.: Glove: Global vectors for word representation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing. ACL, Doha, Qatar (2014).
10. Ronneberger O., Fischer P., Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. Medical Image Computing and Computer-Assisted Intervention(MICCAI) Lecture Notes in Computer Science (9351), 234–241 (2015)
11. Wang L., Jiang J., Chieu H. L., Ong C. H., Song D., Liao L.: Can Syntax Help? Improving an LSTM-based Sentence Compression Model for New Domains. ACL, Vancouver, Canada (2017).
12. Yin W., Kann K., Yu M., Schütze H.: Comparative Study of CNN and RNN for Natural Language Processing. ArXiv: <https://arxiv.org/abs/1702.01923> (2017)
13. Zhao Y., Senuma H., Shen X., Aizawa A.: Gated Neural Network for Sentence Compression using Linguistic Knowledge. In International Conference on Applications of Natural Language to Information Systems. pp 480–491. Springer, Geidelberg (2017).
14. Zhao Y., Luo Z., Aizawa A.: A Language Model Based Evaluator for Sentence Compression. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics. pp 170–175. ACL, Melbourne, Australia (2018).