



Fine-grained Sentiment Classification using BERT

Manish Munikar, Sushil Shakya and Aakash Shrestha

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

October 24, 2019

Fine-grained Sentiment Classification using BERT

Manish Munikar*, Sushil Shakya† and Aakash Shrestha‡

Department of Electronics and Computer Engineering
Pulchowk Campus, Institute of Engineering, Tribhuvan University
Lalitpur, Nepal

*070bct520@ioe.edu.np, †070bct547@ioe.edu.np, ‡070bct501@ioe.edu.np

Abstract—Sentiment classification is an important process in understanding people’s perception towards a product, service, or topic. Many natural language processing models have been proposed to solve the sentiment classification problem. However, most of them have focused on binary sentiment classification. In this paper, we use a promising deep learning model called BERT to solve the fine-grained sentiment classification task. Experiments show that our model outperforms other popular models for this task without sophisticated architecture. We also demonstrate the effectiveness of transfer learning in natural language processing in the process.

Index Terms—sentiment classification, natural language processing, language model, pretraining

I. INTRODUCTION

Sentiment classification is a form of text classification in which a piece of text has to be classified into one of the predefined sentiment classes. It is a supervised machine learning problem. In binary sentiment classification, the possible classes are positive and negative. In fine-grained sentiment classification, there are five classes (very negative, negative, neutral, positive, and very positive). Fig 1 shows a black-box view of a fine-grained sentiment classifier model.

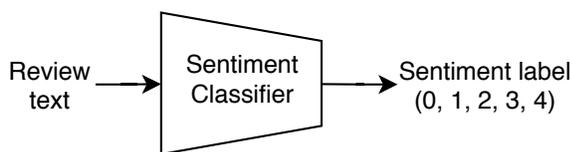


Fig. 1. High-level black-box view of a sentiment classifier showing its input and output.

Sentiment classification model, like any other machine learning model, requires its input to be a fixed-sized vector of numbers. Therefore, we need to convert a text—sequence of words represented as ASCII or Unicode—into a fixed-sized vector that encodes the meaningful information of the text. Many statistical and deep learning NLP models have been proposed just for that. Recently, there has been an explosion of developments in NLP as well as other deep learning architectures.

While transfer learning (pretraining and finetuning) has become the de-facto standard in computer vision, NLP is yet to utilize this concept fully. However, neural language models such as word vectors [1], paragraph vectors [2], and GloVe [3] have started the transfer learning revolution in NLP. Recently,

Google researchers published BERT (Bidirectional Encoder Representations from Transformers) [4], a deep bidirectional language model based on the Transformer architecture [5], and advanced the state-of-the-art in many popular NLP tasks. In this paper, we use the pretrained BERT model and fine-tune it for the fine-grained sentiment classification task on the Stanford Sentiment Treebank (SST) dataset.

The rest of the paper is organized into six sections. In Section II, we mention our motivation for this work. In Section III, we discuss related works. In Section IV, we describe the dataset we performed our experiments on. We explain our model architecture and methodology in detail in Section V. Then we present and analyze our results in Section VI. Finally, we provide our concluding remarks in Section VII.

II. MOTIVATION

We have been working on replicating the different research paper results for sentiment analysis, especially on the fine-grained Stanford Sentiment Treebank (SST) dataset. After the popularity of BERT, researchers have tried to use it on different NLP tasks, including binary sentiment classification on SST-2 (binary) dataset, and they were able to obtain state-of-the-art results as well. But we haven’t yet found any experimentation done using BERT on the SST-5 (fine-grained) dataset. Because BERT is so powerful, fast, and easy to use for downstream tasks, it is likely to give promising results in SST-5 dataset as well. This became the main motivation for pursuing this work.

III. RELATED WORK

Sentiment classification is one of the most popular tasks in NLP, and so there has been a lot of research and progress in solving this task accurately. Most of the approaches have focused on binary sentiment classification, most probably because there are large public datasets for it such as the IMDb movie review dataset [6]. In this section, we only discuss some significant deep learning NLP approaches applied to sentiment classification.

The first step in sentiment classification of a text is the embedding, where a text is converted into a fixed-size vector. Since the number of words in the vocabulary after tokenization and stemming is limited, researchers first tackled the problem of learning word embeddings. The first promising language

model was proposed by Mikolov et al. [1]. They trained continuous semantic representation of words from large unlabeled text that could be fine-tuned for downstream tasks. Pennington et al. [3] used a co-occurrence matrix and only trained on non-zero elements to efficiently learn semantic word embeddings. Bojanowski et al. [7] broke words into character n -grams for smaller vocabulary size and fast training.

The next step is to combine a variable number of word vectors into a single fixed-size document vector. The trivial way is to take the sum or the average, but they don't lose the ordering information of words and thus don't give good results. Tai et al. [8] used recursive neural networks to compute vector representation of sentences by utilizing the intrinsic tree structure of natural language sentences. Socher et al. [9] introduced a tensor-based compositionality function for better interaction between child nodes in recursive networks. They also introduced the Stanford Sentiment Treebank (SST) dataset for fine-grained sentiment classification. Tai et al. [10] applied various forms of long short-term memory (LSTM) networks and Kim [11] applied convolutional neural networks (CNN) towards sentiment classification.

All of the approaches mentioned above are context-free, i.e., they generate single word embedding for each word in the vocabulary. For instance, "bank" would have the same representation in "bank deposit" and "river bank". Recent language model research has been trying to train contextual embeddings. Peters et al. [12] extracted context-sensitive features from left-to-right and right-to-left LSTM-based language model. Devlin et al. [4] proposed BERT (Bidirectional Encoder Representations from Transformers), an attention-based Transformer architecture [5], to train deep bidirectional representations from unlabeled texts. Their architecture not only obtains state-of-the-art results on many NLP tasks but also allows a high degree of parallelism since it is not based on sequential or recurrent connections.

IV. DATASET

Stanford Sentiment Treebank (SST) [9] is one of the most popular publicly available datasets for fine-grained sentiment classification task. It contains 11,855 one-sentence movie reviews extracted from Rotten Tomatoes. Not only that, each sentence is also parsed by the Stanford constituency parser [13] into a tree structure with the whole sentence as the root node and the individual words as leaf nodes. Moreover, each node is labeled by at least three humans. In total, SST contains 215,154 unique manually labeled texts of varying lengths. Fig 2 shows a sample review from the SST dataset in a parse-tree structure with all its nodes labeled. Therefore, this dataset can be used to train models to learn the sentiment of words, phrases, and sentences together.

There are five sentiment labels in SST: 0 (very negative), 1 (negative), 2 (neutral), 3 (positive), and 4 (very positive). If we only consider positivity and negativity, we get the binary SST-2 dataset. If we consider all five labels, we get SST-5. For this research, we evaluate the performance of various models

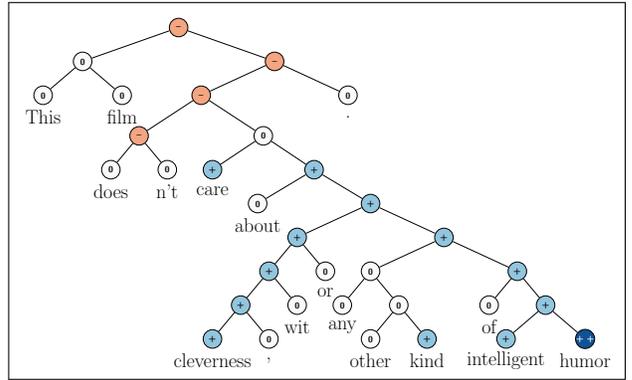


Fig. 2. A sample sentence from the SST dataset. (Source: Adapted from [9].)

on all nodes as well as on just the root nodes, and on both SST-2 and SST-5.

V. METHODOLOGY

Sentiment classification takes a natural language text as input and outputs a sentiment score $\in \{0, 1, 2, 3, 4\}$. Our method has three stages from input sentence to output score, which are described below. We use pretrained BERT model to build a sentiment classifier. Therefore, in this section, we briefly explain BERT and then describe our model architecture.

A. BERT

BERT (Bidirectional Encoder Representations from Transformers) is an embedding layer designed to train deep bidirectional representations from unlabeled texts by jointly conditioning on both left and right context in all layers. It is pretrained from a large unsupervised text corpus (such as Wikipedia dump or BookCorpus) using the following objectives:

- *Masked word prediction*: In this task, 15% of the words in the input sequence are masked out, the entire sequence is fed to a deep bidirectional Transformer [5] encoder, and then the model learns to predict the masked words.
- *Next sentence prediction*: To learn the relationship between sentences, BERT takes two sentences A and B as inputs and learns to classify whether B actually follows A or is it just a random sentence.

Unlike traditional sequential or recurrent models, the attention architecture processes the whole input sequence at once, enabling all input tokens to be processed in parallel. The layers of BERT architecture are visualized in Fig 3. Pretrained BERT model can be fine-tuned with just one additional layer to obtain state-of-the-art results in a wide range of NLP tasks [4].

There are two variants for BERT models: BERT_{BASE} and BERT_{LARGE}. The difference between them is listed in Table I.

1) *Input format*: BERT requires its input token sequence to have a certain format. First token of every sequence should be [CLS] (classification token) and there should be a [SEP] token (separation token) after every sentence. The output embedding corresponding to the [CLS] token is the

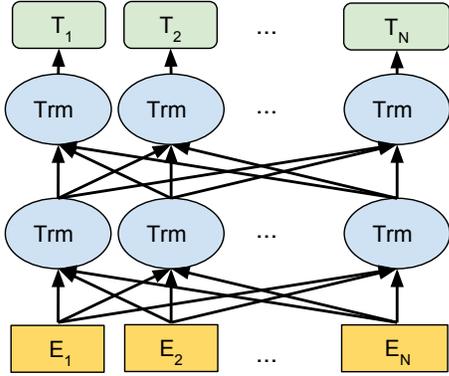


Fig. 3. BERT Architecture, where E_n is the n -th token in the input sequence, Trm is the Transformer block, and T_n is the corresponding output embedding. (Source: Adapted from [4].)

TABLE I
BERT_{BASE} VS. BERT_{LARGE}.

	BERT _{BASE}	BERT _{LARGE}
No. of layers (Transformer blocks)	12	24
No. of hidden units	768	1024
No. of self-attention heads	12	16
Total trainable parameters	110M	340M

sequence embedding that can be used for classifying the whole sequence.

B. Preprocessing

We perform the following preprocessing steps on the review text before we feed them into our model.

1) *Canonicalization*: First, we remove all the digits, punctuation symbols and accent marks, and convert everything to lowercase.

2) *Tokenization*: We then tokenize the text using the WordPiece tokenizer [14]. It breaks the words down to their prefix, root, and suffix to handle unseen words better. For example, playing \rightarrow play + ##ing.

3) *Special token addition*: Finally, we add the [CLS] and [SEP] tokens at the appropriate positions.

C. Proposed Architecture

We build a simple architecture with just a dropout regularization [15] and a softmax classifier layers on top of pretrained BERT layer to demonstrate that BERT can produce great results even without any sophisticated task-specific architecture.

Fig 4 shows the overall architecture of our model. There are four main stages. The first is the preprocessing step as described earlier. Then we compute the sequence embedding from BERT. We then apply dropout with a probability factor of 0.1 to regularize and prevent overfitting. Dropout is only applied in training phase and not in inference phase. Finally, the softmax classification layer will output the probabilities of the input text belonging to each of the class labels such

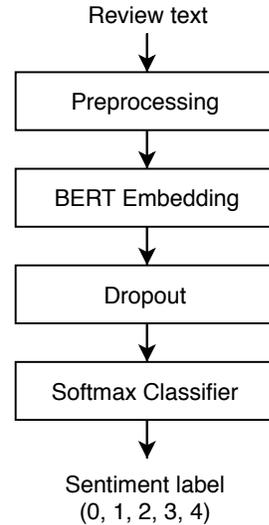


Fig. 4. Proposed architecture for fine-grained sentiment classification.

that the sum of the probabilities is 1. The softmax layer is just a fully connected neural network layer with the softmax activation function. The softmax function $\sigma : \mathbb{R}^K \rightarrow \mathbb{R}^K$ is given in (1).

$$\sigma(\mathbf{z})_i = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}} \text{ for } i = 1, \dots, K \quad (1)$$

where $\mathbf{z} = (z_1, \dots, z_K) \in \mathbb{R}^K$ is the intermediate output of the softmax layer (also called logits). The output node with the highest probability is then chosen as the predicted label for the input.

VI. EXPERIMENTS AND RESULTS

In this section, we discuss the results of our model and compare with it some of the popular models that solve the same problem, i.e., sentiment classification on the SST dataset.

A. Comparison Models

1) *Word embeddings*: In this method, the word vectors pretrained on large text corpus such as Wikipedia dump are averaged to get the document vector, which is then fed to the sentiment classifier to compute the sentiment score.

2) *Recursive networks*: Various types of recursive neural networks (RNN) have been applied on SST [9]. We compare our results with the standard RNN and the more sophisticated RNTN. Both of them were trained on SST from scratch, without pretraining.

3) *Recurrent networks*: Sophisticated recurrent networks such as left-to-right and bidirectional LSTM networks have also been applied on SST [10].

4) *Convolutional networks*: In this approach, the input sequences were passed through a 1-dimensional convolutional neural network as feature extractors [11].

TABLE II
ACCURACY (%) OF OUR MODELS ON SST DATASET COMPARED TO OTHER MODELS.¹

Model	SST-2		SST-5	
	All	Root	All	Root
Avg word vectors [9]	85.1	80.1	73.3	32.7
RNN [8]	86.1	82.4	79.0	43.2
RNTN [9]	87.6	85.4	80.7	45.7
Paragraph vectors [2]	–	87.8	–	48.7
LSTM [10]	–	84.9	–	46.4
BiLSTM [10]	–	87.5	–	49.1
CNN [11]	–	87.2	–	48.0
BERT _{BASE}	94.0	91.2	83.9	53.2
BERT _{LARGE}	94.7	93.1	84.2	55.5

¹ Some values are blank in “All” columns because the original authors of those paper did not publish their result on all phrases.

B. Evaluation Metric

Since the dataset has roughly balanced number of samples of all classes, we directly use the accuracy measure to evaluate the performance of our model and compare it with other models. The accuracy is defined simply as follows:

$$\text{accuracy} = \frac{\text{number of correct predictions}}{\text{total number of predictions}} \in [0, 1] \quad (2)$$

C. Results

The result and comparisons are shown in Table II. It shows the accuracy of various models on SST-2 and SST-5. It includes results for all phrases as well as for just the root (whole review). We can see that our model, despite being a simple architecture, performs better in terms of accuracy than many popular and sophisticated NLP models.

VII. CONCLUSION

In this paper, we used the pretrained BERT model and fine-tuned it for the fine-grained sentiment classification task on the SST dataset. Even with such a simple downstream architecture, our model was able to outperform complicated architectures like recursive, recurrent, and convolutional neural networks. Thus, we have demonstrated the transfer learning capability in NLP enabled by deep contextual language models like BERT.

ACKNOWLEDGMENT

We would like to express our gratitude towards Prof. Dr. Shashidhar Ram Joshi for his invaluable advice and guidance

on this paper. We also thank all the helpers and reviewers for their valuable input to this work.

REFERENCES

- [1] T. Mikolov, K. Chen, G. S. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” *CoRR*, vol. abs/1301.3781, 2013.
- [2] Q. Le and T. Mikolov, “Distributed representations of sentences and documents,” in *International Conference on Machine Learning*, 2014, pp. 1188–1196.
- [3] J. Pennington, R. Socher, and C. Manning, “GloVe: Global vectors for word representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *NAACL-HLT*, 2018.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [6] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, “Learning word vectors for sentiment analysis,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, June 2011, pp. 142–150.
- [7] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching word vectors with subword information,” *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.
- [8] R. Socher, J. Pennington, E. H. Huang, A. Y. Ng, and C. D. Manning, “Semi-supervised recursive autoencoders for predicting sentiment distributions,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2011, pp. 151–161.
- [9] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts, “Recursive deep models for semantic compositionality over a sentiment treebank,” in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2013, pp. 1631–1642.
- [10] K. S. Tai, R. Socher, and C. D. Manning, “Improved semantic representations from tree-structured long short-term memory networks,” in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2015, pp. 1556–1566.
- [11] Y. Kim, “Convolutional neural networks for sentence classification,” *arXiv preprint arXiv:1408.5882*, 2014.
- [12] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, “Deep contextualized word representations,” *arXiv preprint arXiv:1802.05365*, 2018.
- [13] D. Chen and C. Manning, “A fast and accurate dependency parser using neural networks,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 740–750.
- [14] M. Schuster and K. Nakajima, “Japanese and korean voice search,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 5149–5152.
- [15] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.