

KUISAIL at SemEval-2020 Task 12: BERT-CNN for Offensive Speech Identification in Social Media

Ali Safaya, Moutasem Abdullatif, Deniz Yuret
KUIS AI Lab
Koç University, İstanbul/Turkey
{asafaya19,mabdullatif18,dyuret}@ku.edu.tr

Abstract

In this paper, we describe our approach to utilize pre-trained BERT models with Convolutional Neural Networks for sub-task A of the Multilingual Offensive Language Identification shared task (OffensEval 2020), which is a part of the SemEval 2020. We show that combining CNN with BERT is better than using BERT on its own, and we emphasize the importance of utilizing pre-trained language models for downstream tasks. Our system, ranked 4th with macro averaged F1-Score of 0.897 in Arabic, 4th with score of 0.843 in Greek, and 3rd with score of 0.814 in Turkish. Additionally, we present ArabicBERT, a set of pre-trained transformer language models for Arabic that we share with the community.

1 Introduction

Recently, we have observed an increase in social media usage and a similar increase in hateful and offensive speech. Solutions to this problem vary from manual control to rule-based filtering systems; however, these methods are time-consuming or prone to errors if the full context is not taken into consideration while assessing the sentiment of the text (Saif et al., 2016).

In Subtask-A of the shared task of Multilingual Offensive Language Identification (OffensEval2020), we focus on detecting offensive language on social media platforms, more specifically, on Twitter. The organizers provided data from five different languages, which we worked on three languages of them, namely, Arabic (Mubarak et al., 2020), Greek (Pitenis et al., 2020), and Turkish (Çöltekin, 2020). More details about the annotation process have been described in task description paper (Zampieri et al., 2020).

The approach used combines the knowledge embedded in pre-trained deep bidirectional transformer BERT (Devlin et al., 2019) with Convolutional Neural Networks (CNN) for text (Kim, 2014), which is one of the most utilized approaches for text classification tasks. This combination of models has been shown to yield better results than using BERT or CNN on their own, as was shown in (Li et al., 2019), and shown in this paper. This model, and with minimum text pre-processing, ranked 4th in Arabic, 4th in Greek, and 3rd in Turkish among more than 40 participants.

In the following sections of this paper, previous work is mentioned in Section 2, next, the data is described in Section 3, then the details of the model and the other experiments are described in Section 4. Finally, the submissions and the other experiments are detailed in Section 5.

2 Background

Extensive work has been performed to solve the task of offensive speech identification, which classifies among text classification tasks. Approaches to solve this problem vary from using lexical resources, linguistic features, and meta information (Schmidt and Wiegand, 2017), to machine learning (ML) models

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

Our source code of the main model and the other experiments can be accessed through: <https://github.com/alisafaya/OffensEval2020>

(Davidson et al., 2017), and more recently, deep neural models like CNN and Long-Short Term Memory (LSTM) and their derivatives (Zhang et al., 2018).

More recent work, Zampieri et al. (2019) presented Offensive Language Identification Dataset (OLID), which is a new dataset with tweets annotated for offensive content, they experiment with various ML models, like SVM, BiLSTM and CNN.

3 Data

The data provided for this task (Zampieri et al., 2020) consists of sets of tweets which were annotated as either **Offensive** (positive) or **Non-offensive** (negative). As shown in Table 1 below, each set contains a number of positive tweets and negative tweets. In addition, the provided training data had not been split into training and development sets, so the data was split into 90% and 10% for training and development sets respectively.

	Arabic			Greek			Turkish		
	Train	Dev	Test	Train	Dev	Test	Train	Dev	Test
Negative	5,785	626	1,607	5,642	616	1,120	23,084	2,543	2,740
Positive	1,415	174	393	2,228	258	424	4,885	632	788
Total	7,200	800	2,000	7,869	874	1,545	28,581	3,175	3,528

Table 1: Tweets distribution over data sets

3.1 Data Pre-processing

Since processed texts were obtained from Twitter, a pre-processing step was needed to maximize the features that can be extracted and to obtain a clean text; Hashtags were converted into raw text by splitting the texts into words, for example: #SomeHashtagText becomes Some Hashtag Text. As an additional step only for Greek texts, all letters to converted to lowercase letters and all Greek diacritics were removed. The text was subsequently tokenized using the corresponding BERT pre-trained Wordpiece tokenizer for each language and model.

4 Model Description

The proposed model maximizes the utilization of knowledge embedded in pre-trained BERT language models by feeding the outputted contextualized embeddings of its last four hidden layers into a several filters and convolution layers of the CNN. Finally, the output of the CNN was passed to a dense layer and the predictions were obtained.

4.1 Convolutional Neural Networks

CNN for textual tasks by Kim (2014) showed superiority in text classification tasks. CNNs can be used with learned vector representations of the text (embeddings). These embeddings may either be initialized randomly and trained along with the model, or can be pre-trained vectors.

4.2 BERT

Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019) is state-of-the-art language model, which can be fine-tuned, or used directly as a feature extractor for various textual tasks. In our experiments, three pre-trained language-specific BERT models were used along with Multilingual-BERT (mBERT)¹ model. Those models are GreekBERT² model for Greek, BERTurk (Schweter, 2020) for Turkish, and ArabicBERT for Arabic.

¹Multilingual: <https://github.com/google-research/bert>

²GreekBERT: <https://github.com/nlpauieb/greek-bert>

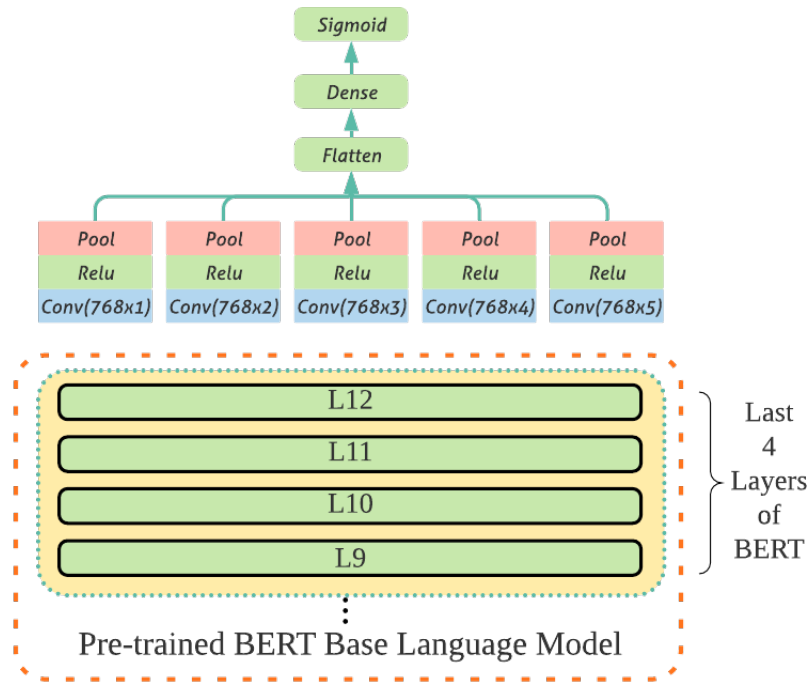


Figure 1: BERT-CNN model structure

4.3 ArabicBERT

Since there was no pre-trained BERT model for Arabic at the time of our work, four Arabic BERT language models were trained from scratch and made publicly available for use.

ArabicBERT³ is a set of BERT language models that consists of four models of different sizes trained using masked language modeling with whole word masking (Devlin et al., 2019). Models of sizes **Large**, **Base**, **Medium**, and **Mini** (Turc et al., 2019) were trained on the same data for 4M steps.

Using a corpus that consists of the unshuffled version of OSCAR data (Ortiz Suárez et al., 2020) and a recent data dump from Wikipedia, which sums up to 8.2B words, a vocabulary set of 32,000 Wordpieces was constructed. The final version of corpus contains some non-Arabic words inlines, which were not removed from sentences since that would affect some tasks like Named Entity Recognition. Although non-Arabic characters were lowered as a pre-processing step, since Arabic characters do not have upper or lower case, there is no cased and uncased version of the model. Subsequently, the corpus and the vocabulary set are not restricted to Modern Standard Arabic, they contain some dialectical (spoken) Arabic too, which boosted models performance in terms of data from social media platforms.

4.4 BERT-CNN Model Structure

As mentioned above, the main model consists of two main parts. The first part being BERT Base model, in which the text is passed through 12 layers of self-attention to obtain contextualized vector representations. The other part being CNN, which was used as a classifier.

Devlin et al. (2019) showed by comparing different combinations of layers of BERT, that the output of the last four hidden layers combined, encodes more information than the output of the top layer.

After setting the maximum sequence length of each text sample (tweet) to 64 tokens, the text was input to BERT, then the output of the last four hidden layers of base sized pre-trained BERT, was concatenated to get vector representations of size $768 \times 4 \times 64$ as shown in Figure 1. Next, these embeddings were passed in parallel into 160 convolutional filters of five different sizes (768×1 , 768×2 , 768×3 , 768×4 , 768×5), 32 filters for each size. Each kernel takes the output of the last four hidden layers of BERT as 4 different channels and applies convolution operation on it. After that, the output is passed through ReLU

³ArabicBERT: <https://github.com/alisafrica/arabic-bert>

Activation function and a Global Max-Pooling operation. Finally, the output of the pooling operation is concatenated and flattened to be later on passed through a dense layer and a Sigmoid function to get the final binary label.

This model was trained for 10 epochs with learning rate of $2e-5$, and the model with the best macro averaged F1-Score on the development set was saved.

5 Experiments and Results

Model	Arabic	Greek	Turkish	Average
SVM with TF-IDF	0.772	0.823	0.685	0.760
Multilingual BERT	0.808	0.807	0.774	0.796
Bi-LSTM	0.822	0.826	0.755	0.801
CNN-Text	0.840	0.825	0.751	0.805
BERT ⁴	0.884	0.822	0.816	0.841
BERT-CNN (Ours) ⁴	0.897	0.843	0.814	0.851

Table 2: Macro averaged F1-Scores of our submissions and the other experiments on test data

Macro-averaged F1-Score metric was used for evaluation in this shared task. The results of our submissions were shown in comparison with other experiments in Table 2.

These experiments began by building a baseline model using a classic ML approach for text classification. Additionally, the main model was compared with more recent approaches. All the models use the same train/dev/test splits.

SVM with TF-IDF⁵

The baseline model used Term Frequency-Inverse Document Frequency (TF-IDF) (Salton and Buckley, 1988) with Support Vector Machine (SVM) (Boser et al., 1992). Count Vectorizer with feature set size of 3000 was used to achieve the results demonstrated in Table 2.

CNN-Text⁶

Using CNNs with the same structure as the main model, but without pre-trained BERT as an embedder. CNN-Text model uses randomly initialized embeddings of size 300, which were trained along with the model. The difference between the results obtained using pre-trained BERT and randomly initialized embeddings was significant as shown in the Table 2 above.

BiLSTM⁶

While CNNs could be used to capture local features of the text, LSTM which have shown remarkable performance in text classification tasks, capture the temporal information. In our experiments, two layers of Bidirectional LSTM (BiLSTM) with a hidden size of 128, and randomly initialized embeddings of size 300, were used to achieve the results shown in Table 2, However, this was still outperformed by CNN-Text on average.

BERT^{6,7}

By looking at the average results of BERT model on its own, we can see the improvement that was achieved by combining BERT with CNN. Additionally we can clearly observe the advantage of using **Language-specific** pre-trained models on **Multilingual** ones.

⁴Language specific pre-trained BERT models were used for this experiment

⁵This model was built using Scikit-learn: scikit-learn.org

⁶This model was built using PyTorch: pytorch.org

⁷Transformers library was used for BERT (Wolf et al., 2019)

6 Conclusion

In this paper, the structure of BERT-CNN was described and compared with other models on the ability to identify offensive speech text in social media. It was shown that combining BERT with CNN yields better than using BERT on its own. Additionally, the pre-training process of ArabicBERT was explained. The proposed model with minimum text pre-processing was able to achieve very good results on average and our team was ranked among the highest four participating teams for all languages in the scope of the OffensEval2020.

Acknowledgements

The hardware infrastructure of this study is provided by the European Research Council (ERC) Starting Grant 714868. Also, we would like to thank Google for providing free TPUs and Credits for the pre-training process of ArabicBERT and for Huggingface.co for hosting these models on their servers.

References

- Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. 1992. A training algorithm for optimal margin classifiers. In *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory*, pages 144–152. ACM Press.
- Çağrı Çöltekin. 2020. A Corpus of Turkish Offensive Language on Social Media. In *Proceedings of the 12th International Conference on Language Resources and Evaluation*. ELRA.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated Hate Speech Detection and the Problem of Offensive Language. In *Proceedings of ICWSM*.
- 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, June. Association for Computational Linguistics.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar, October. Association for Computational Linguistics.
- W. Li, S. Gao, H. Zhou, Z. Huang, K. Zhang, and W. Li. 2019. The automatic text classification method based on bert and feature union. In *2019 IEEE 25th International Conference on Parallel and Distributed Systems (ICPADS)*, pages 774–777.
- Hamdy Mubarak, Ammar Rashed, Kareem Darwish, Younes Samih, and Ahmed Abdelali. 2020. Arabic offensive language on twitter: Analysis and experiments. *arXiv preprint arXiv:2004.02192*.
- Pedro Javier Ortiz Suárez, Laurent Romary, and Benoît Sagot. 2020. A monolingual approach to contextualized word embeddings for mid-resource languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1703–1714, Online, July. Association for Computational Linguistics.
- Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive Language Identification in Greek. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Hassan Saif, Yulan He, Miriam Fernandez, and Harith Alani. 2016. Contextual semantics for sentiment analysis of twitter. *Information Processing & Management*, 52(1):5–19, January.
- Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, 24(5):513–523, January.
- Anna Schmidt and Michael Wiegand. 2017. A Survey on Hate Speech Detection Using Natural Language Processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media. Association for Computational Linguistics*, pages 1–10, Valencia, Spain.
- Stefan Schweter. 2020. Berturk - bert models for turkish, April.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Predicting the Type and Target of Offensive Posts in Social Media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 1415–1420.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media (OffensEval 2020). In *Proceedings of SemEval*.
- Ziqi Zhang, David Robinson, and Jonathan Tepper. 2018. Detecting Hate Speech on Twitter Using a Convolution-GRU Based Deep Neural Network. In *Lecture Notes in Computer Science*. Springer Verlag.