

Amobee at SemEval-2019 Tasks 5 and 6: Multiple Choice CNN Over Contextual Embedding

Alon Rozental*, Dadi Biton*

Amobee Inc., Tel Aviv, Israel

{alon.rozentel,dadi.biton}@amobee.com

Abstract

This article describes Amobee’s participation in “HatEval: Multilingual detection of hate speech against immigrants and women in Twitter” (task 5) and “OffensEval: Identifying and Categorizing Offensive Language in Social Media” (task 6). The goal of task 5 was to detect hate speech targeted to women and immigrants. The goal of task 6 was to identify and categorized offensive language in social media, and identify offense target. We present a novel type of convolutional neural network called “Multiple Choice CNN” (MC-CNN) that we used over our newly developed contextual embedding, [Rozental et al. \(2019\)](#)¹. For both tasks we used this architecture and achieved 4th place out of 69 participants with an F_1 score of 0.53 in task 5, in task 6 achieved 2nd place (out of 75) in Sub-task B - automatic categorization of offense types (our model reached places 18/2/7 out of 103/75/65 for sub-tasks A, B and C respectively in task 6).

1 Introduction

Offensive language and hate speech identification are sub-fields of natural language processing that explores the automatic inference of offensive language and hate speech with its target from textual data. The motivation to explore these sub-fields is to possibly limit the hate speech and offensive language on user-generated content, particularly, on social media. One popular social media platform for researchers to study is Twitter, a social network website where people “tweet” short posts. Each post may contain URLs and/or mentions of other entities on twitter. Among these “tweets” we can find various opinions of people regarding political events, public figures, products, etc. Hence, Twitter data turned

into one of the main data sources for both academia and industry. Its unique insights are relevant for business intelligence, marketing and e-governance. This data also benefits NLP tasks such as sentiment analysis, offensive language detection, topic extraction, etc.

Both the OffensEval 2019 [task \(Zampieri et al. \(2019b\)\)](#) and HatEval 2019 [task](#) are part of the [SemEval-2019](#) workshop. OffensEval has 3 sub-tasks with over 65 groups who participate in each sub-task and HatEval has 2 sub-tasks with 69 groups.

Word embedding is one of the most popular representations of document vocabulary in low-dimensional vector. It is capable of capturing context of a word in a document, semantic and syntactic similarity, relation with other words, etc. For this work, word embedding was created with a model similar to Bidirectional Encoder Representations from Transformers (BERT), [Devlin et al. \(2018\)](#). BERT is a language representation model designed to pre-train deep bidirectional representations by jointly conditioning on both left and right context in all layers. As a result, the pre-trained BERT representations can be fine-tuned to create state-of-the-art models for a wide range of tasks, such as question answering and language inference, without substantial task-specific architecture modifications. Besides the word embedding, BERT generates a classification token, which can be used for text classification tasks.

This paper describes our system for the OffensEval 2019 and HatEval 2019 tasks, where our new contribution is the use of contextual embedding (modified BERT) together with an appropriate network architecture for such embeddings.

The paper is organized as follows: Section 2 describes the datasets we used and the pre-process

*These authors contributed equally to this work.

¹To be published.

Sub-Task A		Sub-Task B		Sub-Task C	
Label	Train	Label	Train	Label	Train
Offensive	4,400	Targeted	3,876	Group	1,074
Not offensive	8,840	Not targeted	524	Individual	2,407
				Other	395
Total	13,240	Total	4,400	Total	3,876
	(a)		(b)		(c)

Table 1: Distributions of labels in OffensEval 2019.

phase. Section 3 describes our system architecture and presents the MC-CNN. In section 4 we present the results of both tasks - the OffensEval and HatEval. Finally, in section 5 we review and conclude the system.

2 Data and Pre-Processing

We used Twitter Firehose dataset. We took a random sample of 50 million unique tweets using the Twitter Firehose service. The tweets were used to train language models and word embeddings; in the following, we will refer to this as the Tweets_50M dataset.

A modified language model, based on BERT, was trained using a large Tweets_50M dataset, containing 50 million unique tweets. We trained two models, one used to predict hate speech in posts (task 5) and the other used to predict offensive language in posts (task 6). The pre-process on the Tweets_50M dataset consists of replacing URLs and Twitter handles with special tokens and keeping only the first 80 sub-word tokens in each tweet (for our vocabulary over 99% of the tweets contain less than 80 tokens).

The language model we trained differs from Devlin et al. (2018) mainly by the addition of a latent variable that represents the topic of the tweet and the persona of the writer. The work on this model is still in progress and in this work we have used an early version of the model described in Rozental et al. (2019).

2.1 OffensEval

OffensEval 2019 is divided into three sub-tasks.

1. Sub-task A - Offensive language identification - identify whether a text contains any form of non-acceptable language (profanity) or a targeted offense.

2. Sub-task B - Automatic categorization of offense types - identify whether a text contains targeted or non-targeted profanity and swearing.
3. Sub-task C - Offense target identification - determine whether the target of the offensive text is an individual, group or other (e.g., an organization, a situation, an event, or an issue).

The official OffensEval task datasets, retrieved from social media (Twitter). Table 1 presents the label distribution for each sub-task. For further shared task description, data-sets and results regarding this task, see Zampieri et al. (2019a).

2.2 HatEval

HatEval 2019 is divided into two sub-tasks.

1. Sub-task A - Hate Speech Detection against Immigrants and Women: a two-class classification where systems have to predict whether a tweet in English with a given target (women or immigrants) is hateful or not hateful.
2. Sub-task B - Aggressive behavior and Target Classification: where systems are asked first to classify hateful tweets for English and Spanish (e.g., tweets where Hate Speech against women or immigrants has been identified) as aggressive or not aggressive, and second to identify the target harassed as individual or generic (i.e. single human or group). In this paper we will focus only on sub-task A as none of the participants overcame the baseline accuracy in sub-task B.

There were 69 groups who participated in sub-task A. Table 2 presents the label distribution in sub-

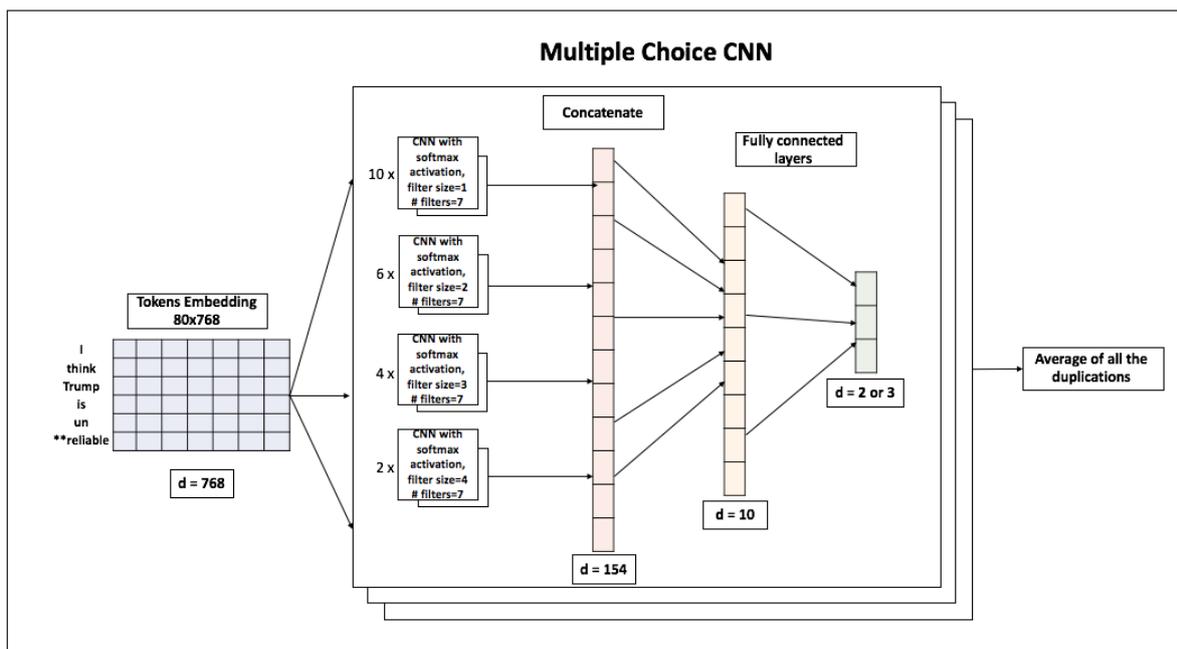


Figure 1: Architecture Of Amobee Offensive Language Detector.

Label	Train
Hate speech	4,210
Not hate speech	5,790
Total	10,000

Table 2: Distributions Of Labels In HatEval 2019.

task A. For further shared task description, datasets and results regarding this task, see Basile et al. (2019). HatEval also included Spanish task which we didn’t participate in.

3 Multiple Choice CNN

For both tasks, using our contextual word embedding, we tried several basic approaches - A feed forward network using the classification vector and an LSTM and simple CNNs Zhang and Wallace (2015) using the words vectors. These approaches overfitted very fast, even for straightforward unigram CNN with 1 filter, and their results were inferior to those obtained by similar models over a Twitter specific, Word2Vec based embedding, Mikolov et al. (2013); Rozentel et al. (2018). The fast overfitting is due to the information contained in contextual embedding which was not reflected in Word2Vec based embedding.

In order to avoid overfitting and achieve

better results we created the MC-CNN model. The motivation behind this model is to replace quantitative questions such as “how mad is the speaker?”, where the result is believed to be represented by the activation of the corresponding filter, with multiple choice questions such as “what is the speaker - happy/sad/other?”, where the number of choices denoted by the number of filters. By forcing the sum of the filter activations for each group to be equal to 1, we believe that we have achieved this effect.

The model that produced the best results is an ensemble of multiple MC-CNNs over our developed contextual embedding, described in figure 1. On top of our contextual embedding, we used four filter sizes - 1-4 sub-word token n-grams. For each filter size individual filters were divided into groups of 7 and a softmax activation applied on the output of each group. These outputs were concatenated and passed to a fully connected feed forward layer of size 10 with tanh activation before it yielded the networks’ prediction. To decrease the variance of the results, there were multiple duplications of this architecture, where the final prediction was the average of all the duplications’ output.

4 Results

We chose to use this architecture for both tasks because we believe that the BERT model output

contains most of the information about the tweet. The layers above, the MC-CNN and the fully connected layers, adapt it to the given task. We think that this model can be use for variety of NLP tasks in twitter data with the appropriate hyper-parameters tuning.

The results yielded from the architecture which was described in figure 1 for both tasks. We optimized the hyper-parameters to maximize the F₁score using categorical cross entropy loss. The tuned parameters were the activation function of the filters and the number of filters in the MC-CNNs, the size of the filter groups of the MC-CNN, and the hidden layer size. The best result were achieved with a sigmoid activation function on the filters, where the number of filters was 7 in each group. There were 10, 6, 4 and 2 filter groups for filter sizes of 1, 2, 3 and 4 respectively. The model with those hyper-parameters yielded the best results in both tasks.

At HatEval the model achieved an F₁score of 0.535. In table 3 there is the best result compared to two baselines- linear Support Vector Machine based on a TF-IDF representation (SVC), and a trivial model that assigns the most frequent label (MFC), according to the F₁score.

System	F1 (macro)	Accuracy
SVC baseline	0.451	0.492
MFC baseline	0.367	0.579
MC-CNN	0.535	0.558

Table 3: F₁Score And Accuracy Of MC-CNN Comparing To Baselines At HatEval.

At OffensEval the model achieved an F₁score of 0.787, 0.739 and 0.591 for sub-tasks A, B and C respectively. In table 4 there is the best result compared to the baseline for sub-tasks A, B and C respectively according to the F₁score.

5 Conclusion

In this paper we described the system Amobee developed for the HatEval and OffensEval tasks. It consists of our novel task specific contextual embedding and MC-CNNs with softmax activation. The use of social networks motivated us to train contextual embedding based on the Twitter dataset, and use the information learned in this language model to identify offensive language and hate speech in the text. Using MC-CNN helped overcome the overfitting caused by the

System	F1 (macro)	Accuracy
All NOT baseline	0.4189	0.7209
All OFF baseline	0.2182	0.2790
MC-CNN	0.7868	0.8384

(a) Sub-task A.

System	F1 (macro)	Accuracy
All TIN baseline	0.4702	0.8875
All UNT baseline	0.1011	0.1125
MC-CNN	0.7386	0.9042

(b) Sub-task B.

System	F1 (macro)	Accuracy
All GRP baseline	0.1787	0.3662
All IND baseline	0.2130	0.4695
All OTH baseline	0.0941	0.1643
MC-CNN	0.5909	0.7042

(c) Sub-task C.

Table 4: F₁Score And Accuracy Of MC-CNN Comparing To Baselines At OffensEval.

embedding. In order to decrease the variance of the system we used duplications of this model and averaged the results. This system reached 4th place at the HateEval task with an F₁score of 0.535, and 2nd place at sub-task B in the OffensEval task, with an F₁score of 0.739. As we mentioned, we used an early version of a Twitter specific language model to achieve the above results. We plan to release the complete, fully trained version in the near future and test it for different NLP tasks- such as topic classification, sentiment analysis, etc.

References

- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Rangel, Paolo Rosso, and Manuela Sanguinetti. 2019. Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2019)*. Association for Computational Linguistics, Minneapolis, Minnesota.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their

- compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Alon Rozental, Daniel Fleischer, and Zohar Kelrich. 2018. Amobee at iest 2018: Transfer learning from language models. *arXiv preprint arXiv:1808.08782*.
- Alon Rozental, Zohar Kelrich, and Daniel Fleischer. 2019. Latent universal task specific bert.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019a. Predicting the Type and Target of Offensive Posts in Social Media. In *Proceedings of NAACL*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019b. SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval). In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Ye Zhang and Byron Wallace. 2015. [A sensitivity analysis of \(and practitioners' guide to\) convolutional neural networks for sentence classification](#).