

UPB at SemEval-2020 Task 12: Multilingual Offensive Language Detection on Social Media by Fine-tuning a Variety of BERT-based Models

Mircea-Adrian Tanase, Dumitru-Clementin Cercel, Costin-Gabriel Chiru
University Politehnica of Bucharest, Faculty of Automatic Control and Computers
mircea.tanase@stud.acs.upb.ro
{dumitru.cercel, costin.chiru}@cs.pub.ro

Abstract

Offensive language detection is one of the most challenging problem in the natural language processing field, being imposed by the rising presence of this phenomenon in online social media. This paper describes our Transformer-based solutions for identifying offensive language on Twitter in five languages (i.e., English, Arabic, Danish, Greek, and Turkish), which was employed in Subtask A of the Offenseval 2020 shared task. Several neural architectures (i.e., BERT, mBERT, Roberta, XLM-Roberta, and ALBERT), pre-trained using both single-language and multilingual corpora, were fine-tuned and compared using multiple combinations of datasets. Finally, the highest-scoring models were used for our submissions in the competition, which ranked our team 21st of 85, 28th of 53, 19th of 39, 16th of 37, and 10th of 46 for English, Arabic, Danish, Greek, and Turkish, respectively.

1 Introduction

Social media platforms are gaining increasing popularity for both personal and political communication. Recent studies uncovered disturbing trends in communications on the Internet. For example, Pew Research Center¹ discovered that 60% of Internet users have witnessed a form of online harassment, while 41% of them have personally experienced it. The majority of the latter category says the most recent such experience occurred on a social media platform. Although most of these platforms provide ways of reporting offensive or hateful content, only 9% of the victims have considered using these tools.

Traditionally, identifying and removing offensive or hateful content on the Internet is performed by human moderators that inspect each piece of content flagged by the users and label it appropriately. This process has two major disadvantages. As previously mentioned, the first one is that the proportion of users that even considered using the tools provided by the platforms in order to report the harmful content is very small. The second one is represented by the continuously growing volume of data that needs to be analyzed.

However, the task of automated offensive language detection on social media is a very complicated problem. This is because the process of labeling an offensive language dataset has proved to be very challenging, as every individual reacts differently to the same content, and the consensus on assigning a label to a piece of content is often difficult to obtain (Waseem et al., 2017).

The SemEval-2019 shared task 6 (Zampieri et al., 2019b), *Offenseval 2019*, was the first competition oriented towards detecting offensive language in social media, specifically Twitter. The SemEval-2020 shared task 12 (Zampieri et al., 2020), *Offenseval 2020*, proposes the same problem, the novelty being a very large automatically labeled English dataset and also smaller datasets for four other languages: Arabic, Danish, Greek, and Turkish. This paper describes the Transformer-based machine learning models we used in our submissions for each language in Subtask A, where the goal is to identify the offensive language in tweets – a binary classification problem.

The remainder of the paper is structured as follows: in section 2, a brief analysis of the state-of-the-art approaches is performed. In section 3, the methods employed for automated offensive language detection

¹This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

¹<https://www.pewresearch.org/internet/2017/07/11/online-harassment-2017/>

are presented. Section 4 describes the data used in this study. Section 5 presents the evaluation process, and in section 6, conclusions are drawn.

2 Related work

Automating the task of offensive language detection becomes a necessity on the Internet of today, especially on communication platforms. The efforts in this direction have substantially grown in the research community. The first approaches on a related subject include the detection of racist texts in web pages (Greevy and Smeaton, 2004), where the authors considered part-of-speech tags as inputs for support vector machines.

This type of problems gained a lot of interest in the last decade, as advanced machine learning techniques has been developed for NLP tasks, and also computing power became increasingly accessible. Cambria et.al. (2010) proposed the detection of web trolling (i.e., posting outrageous messages that are meant to provoke an emotional response). Hate speech and offensive language detection in Twitter samples was analyzed by Davidson et al. (2017). Their study presents a framework for differentiating between profanity and hate speech and also describes the annotation process of such a dataset. Moreover, their experiments with various text preprocessing techniques are described and the logistic regression algorithm is used for hate speech and offensive language classification. Malmasi and Zampieri (2018) continued the experiments on this dataset using n-gram and skip-gram features.

More recently, neural networks gained interest in this type of problems. For example, Gambäck and Sikdar (2017) relied on a Convolutional Neural Network (CNN) (Kim, 2014) to surpass the state-of-the-art results on the previously mentioned dataset. Zhang et al. (2018) also improved these results by combining two different deep learning architectures, namely CNN and Gated Recurrent Unit (Cho et al., 2014).

There are a series of surveys and shared-tasks that took place in the last years on the subject of detecting the online offensive, abusive, hateful, or toxic content. Schmidt and Wiegand (2017) introduced a comprehensive survey on different methods to automatically recognize hate speech, focusing mostly on non-neural network approaches. Shared-tasks analyzing problems in the same areas include both editions of Abusive Language Online (Fišer et al., 2018), which focused mostly on cyberbullying, TRAC (Kumar et al., 2018), which mainly studied aggressiveness, HASOC (Mandl et al., 2019), which also addressed the problem of hate speech, the same as the SemEval-2019 Task 5 (Basile et al., 2019) competition.

3 Methodology

3.1 Baseline

As a baseline, we used a non-neural network approach, which employs the XGBoost (Chen and Guestrin, 2016) algorithm for classification and multiple text processing techniques for feature extraction:

- Firstly, the lemma of the words was extracted and the TF-IDF scores were computed for the n-grams obtained, with $n = 1, 2, 3$.
- Secondly, part-of-speech tags were extracted using the NLTK Python package (Loper and Bird, 2002) and the TF-IDF scores were computed for the tag n-grams obtained, with $n = 1, 2, 3$.
- Thirdly, TF-IDF scores were computed for character n-grams, with $n = 1, 2, 3$.
- Sentiment analysis features were obtained using the VADER tool (Hutto et al., 2015), which is based on a mapping between lexical features and sentiment scores.
- Finally, other lexical features were added, such as the number of characters, words, syllables, and the Flesch-Kincaid readability score (Kincaid et al., 1975).

3.2 BERT

Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019) is a novel deep learning architecture designed for NLP tasks by the Google team. It combines the WordPiece embeddings (Wu et al., 2016) and the Transformers (Wolf et al., 2019) which represents the last major breakthrough in NLP. The BERT models significantly outperformed the state-of-the-art approaches on various text classification or question answering benchmarks. The architecture is a multi-layer transformer encoder, with the novelty consisting of the usage of bidirectional attention instead of recurrent units.

The BERT model can be used for any NLP classification task using a technique called fine-tuning. This consists of starting with a model that has been pre-trained on a very large and comprehensive dataset and training it further for the respective classification task, in our case offensive language identification. There are several pre-trained BERT versions available, differing in terms of model size (i.e., the number of transformers) and the corpora used for pre-training. Therefore, we experimented with the following BERT-aware models:

- BERT-base, which is pre-trained on the English Wikipedia Corpus.
- BERT-base for Danish², which is pre-trained on the entire dump of Danish Wikipedia pages.
- multilingual BERT (mBERT) (Pires et al., 2019), which is pre-trained on a corpus containing the top 104 languages, considering the Wikipedia pages for each language.

3.3 Roberta and XLM-Roberta

Liu et al. (2019) analyzed the BERT model and concluded that it is under-trained, claiming that the hyperparameter choice can significantly impact the obtained results. The robust pre-training method they proposed, namely Roberta, achieved better performances in the same NLP tasks. Based on their work, XLM-Roberta was developed by Conneau et al. (2019) for multilingual NLP tasks. It is pre-trained for more than 100 languages, similarly to mBERT, which it manages to outperform. An interesting observation is that the more significant improvement in results was obtained for low-represented languages, which recommends using it on all five languages in the subtask A of the current competition. Here, we used the base architectures of Roberta for English and XLM-Roberta, both pre-trained using large amounts of specific CommonCrawl³ data.

3.4 ALBERT

ALBERT (Lan et al., 2019), A lite BERT, is a BERT variation that brings two novel parameter-reduction techniques, resulting in lower resource consumption at training time, and at the same time obtains similar performances with the original BERT model. Moreover, ALBERT uses a self-supervised loss focusing on improving modeling on inter-sentence coherence, which helps to obtain better results on NLP tasks with multi-sentence inputs. We fine-tuned the ALBERT-base model for our English language experiments.

4 Data

To analyze the influence of extending the competition-provided training datasets with other corpora constructed for related tasks, we used two additional datasets for fine-tuning the previously mentioned models. The summary of each dataset is presented in Table 1. We can observe that all the datasets have a similar structure (i.e., binary labels, unbalanced, and positive label ratio of 10%-50%).

4.1 Offenseval 2020 Dataset

The Offenseval 2020 dataset is composed of five subsets of tweets, one for each of the languages: English (Rosenthal et al., 2020), Arabic (Mubarak et al., 2020), Danish (Sigurbergsson and Derczynski, 2020), Greek (Pitenis et al., 2020), and Turkish (Çöltekin, 2020). The last four are similar in structure: for each sample, the Subtask A label is binary, revealing whether the tweet contains offensive language

²https://github.com/botxo/nordic_bert

³<https://commoncrawl.org/>

| Dataset | No. Samples | Positive Ratio (%) | Train Set ID | Validation Set ID |
|---|-------------|--------------------|------------------|-------------------|
| Offenseval 2020 English | 9,075,418 | 12.58 | Off_en_train | Off_en_val |
| Offenseval 2020 Arabic | 7,000 | 19.58 | Off_ar_train | Off_ar_val |
| Offenseval 2020 Danish | 3,000 | 12.80 | Off_da_train | Off_da_val |
| Offenseval 2020 Greek | 8,743 | 28.43 | Off_gr_train | Off_gr_val |
| Offenseval 2020 Turkish | 31,277 | 19.33 | Off_tr_train | Off_tr_val |
| Offenseval 2020 all lang. | 9,125,438 | 12.63 | Off_all_train | Off_all_val |
| Offenseval 2020 all lang., except English | 50,020 | 20.56 | Off_no_eng_train | Off_no_eng_val |
| OLID | 13,240 | 33.23 | OLID | - |
| HASOC English | 5,852 | 38.63 | HASOC_en | - |
| HASOC German | 3,819 | 10.65 | HASOC_gr | - |
| HASOC Hindi | 4,665 | 47.07 | HASOC_hi | - |
| HASOC all languages | 14,336 | 29.41 | HASOC_all | - |

Table 1: Statistics of datasets.

or not. The English subset has a more complex annotation scheme for the Subtask A: for each sample, both the average and the standard deviation of the scores assigned by a pool of semi-supervised learning models are given. After exploring the distribution of these values, the following heuristic was used for converting them to binary labels:

- If the model average score is greater than 0.6, the sample is labeled as positive.
- If the model average score is between 0.5 and 0.6, and the standard deviation is smaller than 0.1 (there is a consensus), the sample is labeled as positive.
- All other samples are labeled as negative.

Using this method, we obtained 114,223 English tweet samples labeled as positive. Finally, a 10% ratio of each language subset was put aside for validation purposes, preserving the label distribution.

4.2 OLID Dataset

Zampieri et al. (2019a) introduced the Offensive Language Identification Dataset (OLID) to the Offenseval 2019 shared task. Note that this was actually the starting point for the Offenseval 2020 dataset. Moreover, OLID contains only English language tweets, and the label for the Subtask A (i.e., offensive language identification) is binary.

4.3 HASOC Dataset

The Hate Speech and Offensive Communication dataset⁴ was proposed for the HASOC 2019 competition (Mandl et al., 2019) with the goal of identifying both hate speech and offensive content in Indo-European languages (i.e., English, German, and Hindi). Task 1 of this competition required identifying hateful or offensive tweet samples, and the labels were binary, so the dataset can be considered similar to the Offenseval 2020 subsets.

4.4 Preprocessing

The main preprocessing step is done using the BERT-specific tokenizer, which splits a sentence into tokens in a WordPiece manner. Two more Twitter-specific prior steps were performed:

- Replacing the emojis with the corresponding textual representation by using the emojiPython package⁵.

⁴<https://hasocfire.github.io/hasoc/2019/dataset.html>

⁵<https://pypi.org/project/emoji/>

- Normalizing the hashtags (e.g., "#MakeAmericaGreatAgain" is split into "Make", "America", "Great", and "Again").

For the non-English languages, we also explored the approach of translating the texts to English as a preprocessing step, in order to allow the usage of an English-only pre-trained model. We used the Yandex translation service⁶, but the translation quality proved to be poor. That is, while most of the words were correctly translated, the syntax and meaning of the sentences were lost. For instance, the Turkish tweet "yeniden doğup gelsem çocuk kalır büyümezdiim" is translated as "re-born child grow I do I remain", while a more accurate translation is "If I were born again, I would be a child and I would not grow up".

5 Experiments

All the experiments were performed using an Ubuntu machine with 64GB RAM and one NVIDIA Titan X GPU. The hardware limitations are the reason we only experimented with the base versions of the previously mentioned architectures. The Transformer Python package (Wolf et al., 2019) was used for training and evaluating the Transformer-based models, and each model was fine-tuned for four epochs. We also used the Adam algorithm (Kingma and Ba, 2014) with weight decay for optimization and a learning rate of $2e-5$.

For each language, multiple experiments were performed with the same architecture using several combinations of datasets for fine-tuning, in order to assess the impact each of them has on the performance of a certain model. This implies that, for some experiments, several datasets were simply concatenated and used as a single fine-tuning set. No additional handling of the data was required, as all of the data shared the binary label structure presented in Section 4.

The results we obtained on the validation datasets are summarized in Tables 2, 3, 4, 5, and 6 for English, Arabic, Danish, Greek, and Turkish, respectively. The reported metrics are computed for the Offenseval 2020 language-specific validation sets as described in Section 4. For each language, the highest validation set F1-score is highlighted, meaning that the corresponding model was selected and employed for predicting the language-specific competition test data in our final submission.

| Model Architecture | Pre-train Language | Preprocessing Particularities | Fine-Tuning Dataset | Acc (%) | Pr (%) | Rec (%) | F1 (%) |
|--------------------|--------------------|-------------------------------|----------------------------------|---------|--------|---------|--------------|
| Baseline | - | n-gram TFIDF POS tags | Off_en_train | 86.63 | 81.19 | 76.96 | 79.01 |
| BERT | English | - | Off_en_train | 97.90 | 91.70 | 90.69 | 91.19 |
| BERT | English | - | Off_en_train+ OLID+HASOC_en | 94.62 | 89.91 | 87.75 | 88.81 |
| mBERT | Multi | - | Off_all_train+ OLID+HASOC_all | 97.54 | 91.42 | 89.92 | 90.66 |
| Roberta | English | - | Off_en_train | 98.05 | 93.11 | 91.43 | 92.26 |
| Roberta | English | - | Off_en_train+ OLID+HASOC_en | 98.01 | 93.03 | 91.45 | 92.23 |
| ALBERT | English | - | Off_en_train+ | 97.96 | 92.10 | 91.83 | 91.96 |
| ALBERT | English | - | Off_en_train+ OLID+HASOC_en | 97.79 | 94.04 | 88.13 | 90.99 |
| XLM-Roberta | Multi | - | Off_all_train+ OLID+HASOC_all | 95.27 | 90.10 | 87.84 | 88.95 |

Table 2: Results obtained for English (on the Off_en_val validation subset).

⁶<https://tech.yandex.com/translate/>

| Model Architecture | Pre-train Language | Preprocessing Particularities | Fine-Tuning Dataset | Acc (%) | Pr (%) | Rec (%) | F1 (%) |
|--------------------|--------------------|-------------------------------|----------------------------------|---------|--------|---------|--------------|
| Baseline | - | n-gram TFIDF | Off_ar_train | 56.58 | 28.52 | 22.04 | 24.86 |
| BERT | English | Translation | Off_en_train | 78.18 | 60.52 | 57.86 | 58.96 |
| mBERT | Multi | - | Off_ar_train | 81.14 | 65.42 | 64.38 | 64.89 |
| mBERT | Multi | - | Off_no_eng_train | 86.12 | 70.48 | 68.16 | 69.30 |
| mBERT | Multi | - | Off_all_train | 89.90 | 72.63 | 71.88 | 72.25 |
| mBERT | Multi | - | Off_all_train+ OLID+HASOC_all | 89.44 | 74.81 | 71.66 | 73.20 |
| XLNet-Roberta | Multi | - | Off_ar_train | 84.42 | 68.14 | 66.78 | 67.45 |
| XLNet-Roberta | Multi | - | Off_no_eng_train | 89.28 | 72.32 | 74.19 | 73.24 |
| XLNet-Roberta | Multi | - | Off_all_train | 90.05 | 76.92 | 70.06 | 73.82 |
| XLNet-Roberta | Multi | - | Off_all_train+ OLID+HASOC_all | 90.56 | 76.82 | 74.83 | 75.81 |

Table 3: Results obtained for Arabic (on the Off_ar_val validation subset).

| Model Architecture | Pre-train Language | Preprocessing Particularities | Fine-Tuning Dataset | Acc (%) | Pr (%) | Rec (%) | F1 (%) |
|--------------------|--------------------|-------------------------------|----------------------------------|---------|--------|---------|--------------|
| Baseline | - | n-gram TFIDF | Off_da_train | 56.92 | 27.36 | 22.39 | 24.62 |
| BERT | English | Translation | Off_en_train | 82.62 | 57.12 | 42.09 | 46.41 |
| mBERT | Danish | - | Off_da_train | 93.58 | 77.12 | 65.29 | 70.71 |
| mBERT | Multi | - | Off_da_train | 89.76 | 64.93 | 50.12 | 56.57 |
| mBERT | Multi | - | Off_no_eng_train | 85.36 | 66.18 | 54.68 | 59.88 |
| mBERT | Multi | - | Off_all_train | 90.44 | 71.06 | 60.52 | 65.36 |
| mBERT | Multi | - | Off_all_train+ OLID+HASOC_all | 90.12 | 72.41 | 58.18 | 64.51 |
| XLNet-Roberta | Multi | - | Off_da_train | 85.42 | 66.19 | 52.94 | 58.82 |
| XLNet-Roberta | Multi | - | Off_no_eng_train | 88.14 | 68.41 | 56.30 | 61.76 |
| XLNet-Roberta | Multi | - | Off_all_train | 91.89 | 71.87 | 60.52 | 65.71 |
| XLNet-Roberta | Multi | - | Off_all_train+ OLID+HASOC_all | 91.89 | 66.66 | 73.68 | 70.00 |

Table 4: Results obtained for Danish (on the Off_da_val validation subset).

Results on the English Subset. Firstly, we observe that, although the baseline classifier does not obtain a negligible result, the Transformer-based models outperform it, even when pre-trained for multilingual tasks, proving the performance improvement that this type of models brings to the task of automated offensive language detection.

Secondly, ALBERT and Roberta perform better than the BERT-base architecture, thus confirming their better exploitation of the Transformer’s representative power. Furthermore, we note that even the multilingual pre-trained model performs better than the English-specific baseline, although significantly worse than the English-only pre-trained models.

Furthermore, there is no evidence that adding the OLID and HASOC English datasets to the fine-tuning data affect the results in any way, most likely because the size of these datasets is very small in comparison to the size of the Offenseval 2020 English dataset. Finally, the best performing model is Roberta fine-tuned without adding the two additional subsets.

Results on the non-English Subsets. The very low scores obtained by the baseline approach for the non-English subsets are explained by the fact that most of the pre-processing employed is English language-specific and could not be applied for other languages. The approach of automatically translating the texts and then applying an English-language pre-trained and fine-tuned model also seems to fail,

| Model Architecture | Pre-train Language | Preprocessing Particularities | Fine-Tuning Dataset | Acc (%) | P (%) | R (%) | F1 (%) |
|--------------------|--------------------|-------------------------------|----------------------------------|---------|-------|-------|--------------|
| Baseline | - | n-gram TFIDF | Off_gr_train | 58.42 | 26.37 | 22.28 | 24.15 |
| BERT | English | Translation | Off_en_train | 68.48 | 62.53 | 52.64 | 57.16 |
| mBERT | Multi | - | Off_gr_train | 80.76 | 71.62 | 68.31 | 69.92 |
| mBERT | Multi | - | Off_no_eng_train | 81.14 | 70.85 | 67.18 | 68.96 |
| mBERT | Multi | - | Off_all_train | 82.50 | 73.80 | 61.83 | 67.28 |
| mBERT | Multi | - | Off_all_train+ OLID+HASOC_all | 84.12 | 72.24 | 67.14 | 69.59 |
| XLM-Roberta | Multi | - | Off_gr_train | 79.14 | 68.23 | 65.98 | 67.08 |
| XLM-Roberta | Multi | - | Off_no_eng_train | 82.57 | 71.04 | 68.67 | 69.83 |
| XLM-Roberta | Multi | - | Off_all_train | 82.97 | 73.80 | 62.24 | 67.53 |
| XLM-Roberta | Multi | - | Off_all_train+ OLID+HASOC_all | 83.54 | 71.96 | 69.07 | 70.49 |

Table 5: Results obtained for Greek (on the Off_gr_val validation subset).

| Model Architecture | Pre-train Language | Preprocessing Particularities | Fine-Tuning Dataset | Acc (%) | P (%) | R (%) | F1 (%) |
|--------------------|--------------------|-------------------------------|----------------------------------|---------|-------|-------|--------------|
| Baseline | - | n-gram TFIDF | Off_tr_train | 61.18 | 28.83 | 19.02 | 22.91 |
| BERT | English | Translation | Off_en_train | 85.04 | 72.63 | 36.32 | 48.42 |
| mBERT | Multi | - | Off_tr_train | 83.90 | 63.12 | 54.63 | 58.56 |
| mBERT | Multi | - | Off_no_eng_train | 80.62 | 62.18 | 52.13 | 57.87 |
| mBERT | Multi | - | Off_all_train | 83.92 | 65.26 | 55.13 | 59.76 |
| mBERT | Multi | - | Off_all_train+ OLID+HASOC_all | 84.37 | 66.19 | 57.12 | 61.32 |
| XLM-Roberta | Multi | - | Off_tr_train | 81.14 | 62.44 | 52.68 | 57.14 |
| XLM-Roberta | Multi | - | Off_no_eng_train | 86.28 | 68.25 | 54.38 | 60.53 |
| XLM-Roberta | Multi | - | Off_all_train | 86.82 | 69.73 | 56.36 | 62.34 |
| XLM-Roberta | Multi | - | Off_all_train+ OLID+HASOC_all | 86.22 | 67.61 | 55.20 | 60.78 |

Table 6: Results obtained for Turkish (on the Off_tr_val validation subset).

with most of the obtained F1-scores being at least 10 points lower than the best score.

Another interesting observation is that the results are constantly improving when adding more data to the fine-tuning dataset, even if the added data is in a different language than the validation set, thus proving that the multilingual model is able to learn cross-lingual features. As expected, the XLM-Roberta model outperforms mBERT in most experimental setups. As opposed to the English language results, adding the HASOC subsets seems to improve the scores, with the sole exception of the Turkish subset. This could be partially explained by the fact that the HASOC dataset contains two other languages. For instance, the German subset from the HASOC data may have brought a performance boost to the multilingual models on the Danish validation set because Danish is more related to German, being considered a North Germanic language.

Finally, an interesting particularity can be observed for the Danish dataset. The Danish language pre-trained BERT model, fine-tuned using only the very small Danish training set, outperformed even the multilingual model, fine-tuned using all the available data. This proves that, for low-represented languages, a language-specific pre-trained model performs better than a multilingual one, even with smaller amounts of fine-tuning data.

Results on the Leaderboard. The results and the rankings obtained by our submissions can be observed in Table 7, in comparison to the best performing teams. The F1-scores obtained on the Offenseval

2020 Subtask A competition test sets are as follows: 91.05%, 82.19%, 73.80%, 81.40%, and 77.89% for English, Arabic, Danish, Greek, and Turkish, respectively.

| Language | F1-score (%) | Our Ranking | No. Participants | Leader F1-score (%) |
|----------|--------------|-------------|------------------|---------------------|
| English | 91.05 | 21 | 85 | 92.22 |
| Arabic | 82.19 | 28 | 53 | 90.17 |
| Danish | 73.80 | 19 | 39 | 81.20 |
| Greek | 81.40 | 16 | 37 | 85.20 |
| Turkish | 77.89 | 10 | 46 | 82.57 |

Table 7: The results of our submissions on the competition test sets.

6 Conclusions

This work presented our approaches to automatically detect offensive language in multilingual tweets, as part of SemEval-2020 Task 12. We proved that the deep learning solution of fine-tuning pre-trained Transformer-based models can be used successfully to classify offensive language, and we experimented with several such architectures, fine-tuned on multiple combinations of datasets. Comparing the validation set performances against the test set results, we discovered that the last of them were better for the non-English languages, which shows that our models generalize well and also that the proposed test data may be easier to classify than the development data. The smallest positive difference between validation and testing performances was obtained for the Danish subset, which may indicate that the XLM-Roberta model could have been more suitable for the Danish language too, as the small difference obtained in the validation phase could have been outweighed by the generalizing power given by the large multilingual fine-tuning dataset.

Moreover, the performances of the multilingual models increased not only with the size of the fine-tuning dataset, but also with the number of languages it contains. The results also demonstrated that the potential of the multilingual Transformer-based models in offensive language detection could be improved if larger datasets are available for non-English languages. For future work, we intend to consider a transfer learning method in order to leverage datasets that were constructed for similar tasks in the same language.

References

- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, 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*, pages 54–63.
- Erik Cambria, Praphul Chandra, Avinash Sharma, and Amir Hussain. 2010. Do not feel the trolls. *ISWC, Shanghai*.
- 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.
- Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Eleventh international aaai conference on web and social media*.
- 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.
- Darja Fišer, Ruihong Huang, Vinodkumar Prabhakaran, Rob Voigt, Zeerak Waseem, and Jacqueline Wernimont. 2018. Proceedings of the 2nd workshop on abusive language online (alw2). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*.
- Björn Gambäck and Utpal Kumar Sikdar. 2017. Using convolutional neural networks to classify hate-speech. In *Proceedings of the first workshop on abusive language online*, pages 85–90.
- Edel Greevy and Alan F Smeaton. 2004. Classifying racist texts using a support vector machine. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 468–469.
- CJ Hutto, Dennis Folds, and Darren Appling. 2015. Computationally detecting and quantifying the degree of bias in sentence-level text of news stories. In *Proceedings of Second International Conference on Human and Social Analytics*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- J Peter Kincaid, Robert P Fishburne Jr, Richard L Rogers, and Brad S Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. *Technical report, NATTC*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking aggression identification in social media. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 1–11.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. In *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*, pages 63–70.
- Shervin Malmasi and Marcos Zampieri. 2018. Challenges in discriminating profanity from hate speech. *Journal of Experimental & Theoretical Artificial Intelligence*, 30(2):187–202.
- Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019. Overview of the hasoc track at fire 2019: Hate speech and offensive content identification in indo-european languages. In *Proceedings of the 11th Forum for Information Retrieval Evaluation*, pages 14–17.
- 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*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual bert? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001.
- Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive Language Identification in Greek. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2020. A Large-Scale Semi-Supervised Dataset for Offensive Language Identification. In *arxiv*.

- 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*, pages 1–10.
- Gudbjartur Ingi Sigurbergsson and Leon Derczynski. 2020. Offensive Language and Hate Speech Detection for Danish. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Zeerak Waseem, Thomas Davidson, Dana Warmley, and Ingmar Weber. 2017. Understanding abuse: A typology of abusive language detection subtasks. In *Proceedings of the First Workshop on Abusive Language Online*, pages 78–84.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- 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 the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1415–1420.
- 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*, pages 75–86.
- 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 *European semantic web conference*, pages 745–760. Springer.