

problemConquero at SemEval-2020 Task 12: Transformer and Soft label-based approaches

Karishma Laud* **Jagruti Singh*** **Randeep Kumar Sahu** **Ashutosh Modi**

Indian Institute of Technology Kanpur (IITK)

{kslaud, jagriti, randeeps}@iitk.ac.in

ashutoshm@cse.iitk.ac.in

Abstract

In this paper, we present various systems submitted by our team *problemConquero* for SemEval-2020 Shared Task 12 “Multilingual Offensive Language Identification in Social Media”. We participated in all the three sub-tasks of OffensEval-2020, and our final submissions during the evaluation phase included transformer-based approaches and a soft label-based approach. BERT based fine-tuned models were submitted for each language of sub-task A (offensive tweet identification). RoBERTa based fine-tuned model for sub-task B (automatic categorization of offense types) was submitted. We submitted two models for sub-task C (offense target identification), one using soft labels and the other using BERT based fine-tuned model. Our ranks for sub-task A were Greek-19 out of 37, Turkish-22 out of 46, Danish-26 out of 39, Arabic-39 out of 53, and English-20 out of 85. We achieved a rank of 28 out of 43 for sub-task B. Our best rank for sub-task C was 20 out of 39 using BERT based fine-tuned model.

1 Introduction

There has been a rise in the use of offensive language on social media platforms. This discourages effective communication, which comes as a drawback against the goal of social media platforms. Hence a need to address the problem of offensive language identification arises. There has been an increasing awareness in communities to filter and moderate tweets. As huge number of users are online, it is difficult to manually moderate each and every content; hence the need for automation arises. As people from different regions, ethnicity interact with each other on online platforms, the detection of multilingual offensive language has become increasingly important. During the training phase of the competition, we started with various machine learning models and proceeded with the model in the evaluation phase, which gave us the highest macro F1 score on the development set. In this paper, we first describe the existing work that has been done for offensive language detection, our proposed approaches, and the results that we obtained. The implementation for our system is made available via Github¹.

2 Problem Definition

OffensEval-2020 (Zampieri et al., 2020) consisted of 3 sub-tasks.

1. Sub-task A was to classify a tweet as offensive (OFF) or not offensive (NOT). Sub-task A was carried out in five languages: English, Danish, Greek, Turkish, and Arabic.
2. Sub-task B was a classification problem that categorized offensive labeled tweet as targeted (insult or threat) or non-targeted (swear words).
3. Sub-task C was a multi-classification problem that classified the targeted offensive tweet as individual (IND), group (GRP) or others (OTH).

* Authors equally contributed to this work.

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

¹<https://github.com/karishmaslaud/OffensEval>

3 Related Work

Offensive language detection has been studied at various levels of granularity in the form of abusive language detection (Waseem et al., 2017), hate speech detection (Schmidt and Wiegand, 2017; Kshirsagar et al., 2018) and cyberbullying (Huang et al., 2018). API's like Perspective² have been developed to detect content toxicity using machine learning models. Various deep learning and ensemble methods have been proposed (Pitsilis et al., 2018) to detect abusive content online. Sentiment based approaches (Brassard-Gourdeau and Khoury, 2019) have also been used to detect toxicity.

Besides English, there have been various contributions to the detection of offensive content in other languages. For Greek, Pavlopoulos et al. (2017) describe an RNN based attention model trained on the Gazetta dataset for user content moderation. Sigurbergsson and Derczynski (2019) discuss implementation of Logistic Regression, Learned-BiLSTM, Fast-BiLSTM and AUX-Fast-BiLSTM models for Danish. For Arabic, an approach based on convolution neural network and bidirectional LSTM has been proposed (Mohaouchane et al., 2019). Abozinadah and Jones Jr (2017) use a statistical approach for detecting abusive user accounts. Alakrot et al. (2018) use an SVM classifier with n-gram features. For Turkish, Özel et al. (2017) use Naïve Bayes Multinomial, kNN classifier and SVM classifier to detect cyberbullying in which the author used both emoticons and words in the text message as features. Data collection techniques and classical methods in Greek for offensive language have been discussed in Pitenis (2019).

4 Corpus Description

For OffensEval 2020, datasets were provided by the competition organizers. For sub-task A: data was provided in the form of tweets and their corresponding labels (OFF and NOT) for Arabic (Mubarak et al., 2020), Danish (Sigurbergsson and Derczynski, 2020), Turkish (Çöltekin, 2020), and Greek (Pitenis et al., 2020) languages. For the task on the English Language (Rosenthal et al., 2020), for all the three sub-tasks A, B, and C, the data was provided in the form of tweets with their corresponding mean and standard deviation values. The mean and standard deviation scores were confidence measures obtained by semi-supervised learning methods. For sub-task A, the mean score was given to how offensive a tweet was, a higher value indicating a more offensive tweet. For sub-task B, the mean score was given to how close the tweet was to being an untargeted insult (UNT) or how far the tweet was to being a targeted insult (TIN). For sub-task C, a mean score was provided for all the labels "IND", "GRP", and "OTH" with a higher value indicating that the tweet belonged to the given category.

Multilingual datasets provided to us were unbalanced as follows: Danish consisted of approximately 13% offensive tweets and 87 percent, not offensive tweets. Similarly, Arabic (20%: "OFF"), Turkish (19%: "OFF"), and Greek (28%: "OFF") training datasets showed a similar ratio imbalance between "OFF" and "NOT" labeled tweets. We also made use of the OLID (Zampieri et al., 2019a) dataset provided in OffensEval 2019 (Zampieri et al., 2019b), which consisted of the English language tweets and their corresponding labels for sub-tasks A, B, and C conducted in OffensEval 2019. We used this as a validation dataset for sub-task A and error analysis in sub-tasks B and C. During the post-evaluation phase of the competition the organizers also released gold labels, which were actual test labels corresponding to the dataset against which we were evaluated for in the competition.

5 Sub-task A

5.1 Preprocessing

We tried various preprocessing techniques but went with those that gave us the best F1 score on the development set (validation set). Emoji replacement as a preprocessing technique, as suggested by Liu et al. (2019a) in OffensEval 2019, was used for Greek, Arabic, Turkish in sub-task A. We replaced emojis with their corresponding meaning using the emoji library³ (Liu et al., 2019a). Sentences were preprocessed before passing as input to the BERT model. For Greek, stop words and punctuation were removed by

²<https://www.perspectiveapi.com>

³<https://github.com/carpedm20/emoji>

preprocessing with Spacy⁴. For Arabic and Turkish consecutive duplicate words like “@USER @USER” were reduced to a single word “@USER”.

5.2 Pre-evaluation Phase

During the pre-evaluation phase, we implemented different classical models for Greek and Danish. For Greek and Danish language, we implemented SVM, Logistic Regression, XGBoost, and Random Forest models. We also implemented kNN for Danish.

5.3 BERT Model

BERT (Devlin et al., 2018) stands for Bidirectional Encoder Representations from Transformers released by Google AI⁵. BERT is a Masked Language Model that jointly conditions on both the left and right context in all the layers of the transformer and hence learns deep bidirectional representations. For sub-task A, English language competition mean scores were mapped to hard labels “NOT” and “OFF”. The text data having a mean score greater than 0.5 was labeled as “OFF”. We used the BERT-Base Uncased model with a linear classifier for classifying a tweet as offensive or not offensive. The model was trained using Adam optimizer (Kingma and Ba, 2015) with weight decay and learning rate of $2e-5$. We used a maximum sentence length of 64. We trained the model using an 80:20 training-validation split.

5.4 Multilingual BERT Model

BERT-Base, Multilingual Cased, is a 12 layer model trained on 104 languages. We used BERT-Base, Multilingual Cased models for Greek, Danish, Turkish, and Arabic language. For Greek, after preprocessing, we did BERT tokenization and obtained the maximum number of tokens across all the training sentences. We used this information for the truncation of sentences during training and testing. We used a maximum sentence length of 64 for Danish, Arabic and Turkish. We implemented a fine-tuned Multilingual Cased BERT-Base model with linear classifier using Adam optimizer with weight decay (L2 regularization) and a learning rate of $2e-5$.

5.5 Data Augmentation Techniques

In the post-evaluation phase, we translated OLID datasets containing 4400 “OFF” labeled tweets to Greek and Danish respectively using Google cloud translate⁶ to increase the percentage of “OFF” tweets in Greek and Danish. The translated datasets are made available via our GitHub repository.

6 Sub-task B

6.1 Preprocessing

We gave data to the BERT model with and without preprocessing, but as we were getting better results without preprocessing of text sentences, we went ahead with that. We converted the mean probability scores to hard labels “TIN” and “UNT” depending on whichever score was maximum. If the mean score was greater than or equal to 0.5, the hard label was “UNT” else “TIN” was assigned. After using the above method out of (188975) soft labeled tweets given for sub-task B, we got approximately 20% of the tweets classified as “UNT” and 0.80% as “TIN”. We then classified the processed data with the fine-tuned RoBERTa (Liu et al., 2019b) model.

6.2 RoBERTa

RoBERTa uses a robustly optimized BERT pre-training approach. It is pre-trained with dynamically masking full sentences without the next sentence prediction loss and has a larger byte-level Byte-Pair Encoding. We passed the raw tweets without preprocessing as input to the RoBERTa model. We trained RoBERTa embedding based fine-tuned model with a linear classifier using Adam optimizer (with weight decay) and a learning rate of $3e-5$. We considered the maximum sentence length of 64.

⁴<https://spacy.io>

⁵<https://github.com/google-research/bert>

⁶<https://cloud.google.com/translate/docs>

7 Sub-task C

7.1 Preprocessing

We used two approaches for sub-task C: BERT and Soft label based approach. For the BERT model, we didn't apply any preprocessing to the given text sentences as our tried combination of preprocessing techniques wasn't working out well while calculating the F1 score on the validation dataset. For the soft labels model, we replaced consecutive duplicate words with a single occurrence of the same word and removed stopwords and punctuation before giving data to the model. Data is provided in form of soft labels and the mean scores that add up to 1.

7.2 BERT

We gave the data to the BERT model without applying any preprocessing. In this approach, we converted the soft labels to hard labels. The mean score given in the form of probability for the three labels was converted to "IND", "GRP", or "OTH" depending on which label score was maximum. After using the above method out of 188,973 soft label tweets given for sub-task C, we got approximately 80.7% of the tweets classified as "IND", 13.2% as "GRP", and 0.61% as "OTH". The given data was classified using a fine-tuned BERT model. We trained the BERT model on the entire converted training dataset and used OLID as the validation dataset. The competition organizers provided an all "IND" baseline dataset during the test phase, which contained the ID of the tweet and all the labels as "IND". We measured our model's performance on both datasets.

7.3 Soft label

We used an LSTM (Hochreiter and Schmidhuber, 1997) based approach in which the loss function used the soft labels. Categorical cross-entropy loss is given by $[-\sum_{i=1}^N (\mathbf{p}_i) \log(\mathbf{q}_i)]$, where \mathbf{p}_i is the mean score of the label and \mathbf{q}_i is the softmax score calculated for each class via the LSTM, at the end of the every training iteration.

8 Experiments and Results

8.1 Pre-Evaluation Phase Results

The following experiments were conducted during the pre-evaluation phase. For multilingual models, we divided the training data in 80:20 training-validation split, and a batch size of 32 was used. BERT based models were implemented using Pytorch (Paszke et al., 2019) and Huggingface⁷ Transformers (Wolf et al., 2019) and we have fine-tuned the BERT models using McCormick and Ryan (2019). For BERT and RoBERTa models, the F1 scores calculated for both pre and post evaluation phases is the average of the F1 scores obtained across batches for a given epoch. Similarly, for BERT and RoBERTa models, the accuracy score calculated is the average of the accuracy scores across batches for a given epoch.

Greek:BERT We selected the final BERT model via early stopping, which happened around 4 epochs. We used a weight decay factor of 0.01 for non-bias and non-normalization layers while fine-tuning the BERT model. We used Adam optimizer with weight decay and a learning rate of $2e-5$. We also tried using the average sentence length across the corpus, but it did not lead to improvements in the F1 score hence we went with maximum sentence length. Results for sub-task A: Greek language for both BERT and classical methods based on the validation dataset are presented in Table 1.

Greek and Danish Classical Methods We used Scikit-learn (Pedregosa et al., 2011) toolkit for implementation. XGBoost(Chen and Guestrin, 2016) was implemented using the xgboost library⁸. For Greek and Danish, the data was slightly biased because the Greek and the Danish datasets contained 72 percent and 87 percent, not offensive tweets, respectively. To counter this unbalance in data we used class *weight = balanced*⁹ while implementing SVM and Logistic Regression methods. For classical methods (SVM, KNN, XGBoost and Logistic Regression), tf-idf¹⁰ unigram based features were used. Results for

⁷<https://github.com/huggingface/transformers/>

⁸<https://github.com/dmlc/xgboost>

⁹<https://scikit-learn.org/>

¹⁰<https://en.wikipedia.org/wiki/tf-idf/>

sub-task A: Danish language for both BERT and classical methods on the validation dataset are presented in Table 1.

Danish, Arabic, Turkish:BERT For Arabic and Turkish, we used NLTK toolkit¹¹ for tweet tokenization. We used an Adam optimizer with a learning rate of $2e-5$ and weight decay optimizer for fine-tuning the BERT model. Result for sub-task A: Turkish language based on BERT fine-tuned model is presented in Table 2 and the result for sub-task A: Arabic language based on BERT based fine-tuned model is presented in Table 2.

	Method	F1 Score	Accuracy
Greek	SVM	0.638	0.82
	Logistic Regression	0.634	0.82
	Random Forest	0.59	0.82
	XGBoost	0.612	0.81
	Bert + Linear layer	0.81	0.85
Danish	SVM	0.73	0.90
	Logistic Regression	0.73	0.90
	K-Nearest Neighbors (k=5)	0.64	0.88
	XGBoost	0.74	0.91
	Bert + Linear Layer	0.78	0.85

Table 1: Accuracy and macro F1 results on the sub-task A: Greek and Danish based on validation set.

	Method	F1 score	Accuracy
Turkish	Bert + Linear Layer	0.75	0.86
Arabic	Bert + Linear Layer	0.82	0.90
English(Task A)	Bert + Linear Layer	0.80	0.83
English(Task B)	Roberta + Linear Layer	0.85	0.91

Table 2: Accuracy and macro F1 results on the sub-task A:Turkish, Arabic, English and sub-task B based on validation set.

Sub-task A:English We trained the English data on approx 7,500,000 tweets out of the total soft label tweets given and used the OLID (Zampieri et al., 2019a) dataset to check F1 macro. As we had resource constraints, we trained 1,000,000 data tweets in one go using a batch size of 32 tweets. We continued this procedure till 7,500,000 tweets. Result is shown in Table 2.

Sub-task B: For the RoBERTa model, we trained the model for 2 epochs using 80:20 train-val split on the training data. We got the result on the validation set, as indicated in Table 2. Further training the model for 3 epochs reduced the F1 score on the validation set. Therefore, we selected the model that we obtained after 2 epochs as the final model. We also checked the F1 macro on OLID and an all “TIN” labeled testing baseline dataset, and we got the F1 macro score as mentioned in Table 3.

Dataset	F1 Score	Accuracy
All TIN Baseline	0.466	0.875
OLID	0.669	0.901

Table 3: Accuracy and macro F1 results on the sub-task B based on an all “TIN” baseline and OLID dataset.

Sub-Task C:BERT For sub-task C, using BERT based approach, we trained the model for 2 epochs. Further increasing the number of epochs decreased the F1 score on the OLID dataset; hence we went for 2 epochs. The results are shown in Table 4.

Dataset	F1 Score	Accuracy
All IND Baseline	0.337	0.765
OLID	0.302	0.523

Table 4: Accuracy and macro F1 results on the sub-task C:BERT based on an all “IND” baseline and OLID dataset.

¹¹<https://www.nltk.org>

Sub-Task C: Soft Labels For the soft label based approach for sub-task C, the model was trained for 3 epochs using 90:10 train-val split. Model with best validation loss in the 3 epochs was used. We implemented the model using Keras (Chollet and others, 2018). We used one hot encoded vectors embedding with a dimension of 128 as input for the given LSTM model built using Li (2017). The LSTM model had a 1D Spatial Dropout (Tompson et al., 2014) of 0.2 before the LSTM and applied a dropout of 0.2 for the inputs of the LSTM and a recurrent dropout (Gal and Ghahramani, 2015) of 0.2. Results are shown in Table 5.

Dataset	F1 Score	Accuracy
Development Set	0.63	0.755
All IND Baseline	0.293	0.784

Table 5: Accuracy and macro F1 results on the sub-task C:Soft Labels approach based on development set and on an all “IND” baseline.

8.2 Evaluation Phase Results

Accuracy and macro F1 results on the official test set of OffenseEval 2020 are presented in Table 6.

Sub-task	Language	F1 Score
A	Arabic	0.802
A	Danish	0.697
A	Greek	0.812
A	Turkish	0.755
A	English-Bert	0.908
B	English-Roberta	0.556
C	English-BERT	0.587
C	English-Soft Label	0.557

Table 6: Accuracy and macro F1 results on all sub-tasks based on the test dataset as provided by the competition organizers.

8.3 Post-Evaluation Phase Results

Data Augmentation Results The F1 score on the test gold labels after data augmentation is presented in Table 7 for Greek and Danish languages. We augmented Greek and Danish training datasets with OLID “OFF” labeled tweets translated to Greek and Danish, respectively.

Method	BEST F1 Score	Accuracy
BERT+Linear Greek	0.76	0.842
BERT+Linear Danish	0.724	0.878

Table 7: Accuracy and macro F1 results on sub-task A on the test gold labels:Greek and Danish language trained using augmented dataset.

9 Error Analysis

9.1 Quantitative Analysis

Data Augmentation For Greek and Danish (Section 8), the data augmentation techniques using the OLID data for Greek and Danish respectively dropped the F1 score on the test dataset for Greek and improved the F1 score for Danish. However, we are not making any claims as we are still working on improving the score on the BERT models, each trained on original and augmented data.

English Language

The confusion matrices for sub-tasks A, B, and C have been implemented using scikit-learn and the seaborn library (Waskom et al., 2017).

Sub-task A For the BERT model submitted for English sub-task A, the confusion matrix for the gold labels is as shown in Figure 1 (a). We also analyzed the tweets which were offensive but misclassified as not offensive for the gold labels dataset. Some of the tweets didn’t have an offensive tone but had

incorporated some form of offensive word¹². E.g. “D**n... it’s true. RIP Toni Morrison”. For some of the not offensive tweets that were misclassified as offensive, the tweets had a strong negative tone, but words used in it were not offensive. E.g. “One thing I hate most is a liar @elongated-nose-emoji”.

Sub-task B For the RoBERTa model implemented for sub-task B, the confusion matrix for gold labels is as shown in Figure 1 (b). We can see that the false-negative rate of “UNT” (untargeted insult) class was approximately 77% of the total “UNT” labels for testing. This served as a major error contributor for the F1 score for sub-task B. For the gold labels, some of the targeted tweets which were misclassified as “UNT” but were “TIN” included text sentences that described the target’s behavior or situation in an offensive manner. E.g. “The man eats like a f*****g animal.” or “R**e B**y what the f**k”. Also, many of the tweets misclassified as “UNT” by our model included the offensive word “s**t” used in various connotations. Some of the untargeted tweets that were misclassified as targeted included texts that were expressive of one’s feelings or opinions in an aggressive manner or self-criticizing or self-blaming texts. E.g. “I’m the reason why a lot of s**t is the way it is.” or “Looking back at old photos of me makes me physically sick like why was I allowed to be f*gly as h**l?? Who let me do that y’all fake as f**k”.




Figure 1: (a) Confusion Matrix for sub-task A:English , (b) Confusion Matrix for sub-task B on the gold labels dataset

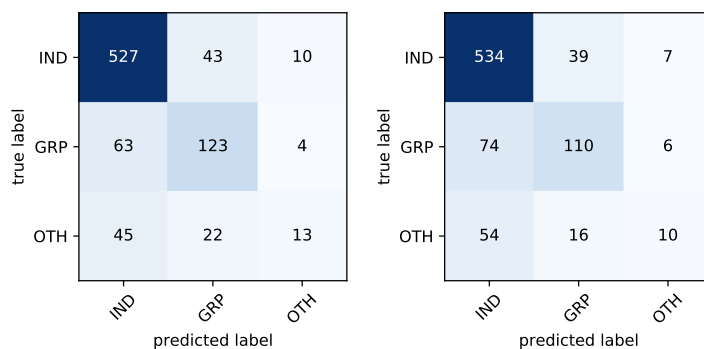


Figure 2: (a) Confusion Matrix for sub-task C: Bert , (b) Confusion Matrix for sub-task C: Soft Labels Model on the gold labels dataset

Sub-task C The confusion matrix corresponding to the Gold dataset for the BERT model is shown in Figure 2 (a), and for the soft labels model is shown in Figure 2 (b). For BERT model, 33% of the “GRP” labeled tweets got predicted as “IND” and for “OTH” hardly 16% of the true labels got classified correctly. Similarly, for the soft labels model, 38% of the “GRP” labeled tweets got predicted as “IND” and for “OTH” hardly 12.5% of the true labels got classified correctly. Here the imbalance in the labeled tweets got propagated to prediction time with “GRP” and “OTH” labels getting miss-classified. We also analyzed tweets that were actually “IND” but misclassified using the BERT model as either “GRP” or “OTH” for

¹²offensive words have been censored while analyzing texts

the gold labeled dataset. Some of the tweets seemed to have a plural entity, and an individual was either being insulted in the tweet due to these entities or the individual was being labeled as an offensive plural entity. E.g. “@USER @USER These children become miserable in your mouth, the truth is they are taken good care of by government.”. For the soft labels model, we analyzed the tweets with their true labels as “IND” but misclassified as “GRP” or “OTH”. Some of the tweets included the word “racist” or some forms of racism in the targeted insult for the individual.

10 Future Work

10.1 BERT Bidirectional LSTM

During the post-evaluation phase, for Greek, we are working on models that focus on the addition of deep learning models on top of the frozen BERT layer using Trevett and de Pablo (2017). By addition of Bidirectional GRU(Cho et al., 2014) on top of BERT layer we got an F1 score of 0.833. By the addition of Bidirectional LSTM (Bi-LSTM) on top of the BERT layer, we got an F1 score of 0.797. After concatenation of the last 4 hidden layers of BERT for both Bidirectional GRU(Bi-GRU) and Bi-LSTM models, the F1 score was 0.819 and 0.837, respectively. The results were obtained after training each model for a maximum of 10 epochs and storing the model with the largest F1 score among the 10 epochs. These results presented in Table 8 are based on the test gold labels, and the F1 and accuracy scores are obtained by averaging across batches in an epoch (Section 8). We are still in the process of tuning the hyperparameters for getting a stable output as we are experiencing a lot of deviation in our F1 scores. We have given the best F1 score obtained after the training of the model. We are also planning to take care of the deviation using an ensemble method with different dropout rates.

Method	F1 Score	Accuracy
BERT+Bi-GRU	0.833	0.906
BERT+Bi-GRU+ BERT Hidden Layers Concatenated	0.819	0.915
BERT+Bi-LSTM	0.797	0.883
BERT+Bi-LSTM + BERT Hidden Layers Concatenated	0.837	0.915

Table 8: Accuracy and macro F1 results on the sub-task A: Greek based on different techniques

10.2 Data Augmentation Techniques

We are planning on trying other data augmentation techniques in addition to the ones proposed. One of them considers using back translations (English to language X to English, Greek to language Y to Greek) as suggested by Aggarwal et al. (2019) in their paper for English sub-tasks for OffensEval 2019. We are planning to use these methods for sub-tasks A and C to improve the macro F1 score.

11 Conclusion

In this paper, we described the approaches that we used for sub-tasks A, B, and C. We used Transformer based approaches for all sub-tasks and a soft label LSTM based approach for sub-task C. We also described our future approaches based on BERT embedding with Bi-GRU and Bi-LSTM. We are planning to use the data augmentation techniques and try to improve the F1 score on the gold labels for each sub-task. We made use of google colab¹³ while developing our project, and the notebooks are made available at our GitHub page.

References

- Ehab A Abozinadah and James H Jones Jr. 2017. A statistical learning approach to detect abusive twitter accounts. In *Proceedings of the International Conference on Compute and Data Analysis*, pages 6–13.
- Piush Aggarwal, Tobias Horsmann, Michael Wojatzki, and Torsten Zesch. 2019. LTL-UDE at SemEval-2019 task 6: BERT and two-vote classification for categorizing offensiveness. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 678–682, Minneapolis, Minnesota, USA, June. Association for Computational Linguistics.

¹³<https://colab.research.google.com/>

- Azalden Alakrot, Liam Murray, and Nikola S Nikolov. 2018. Towards accurate detection of offensive language in online communication in arabic. *Procedia computer science*, 142:315–320.
- Eloi Brassard-Gourdeau and Richard Khoury. 2019. Subversive toxicity detection using sentiment information. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 1–10, Florence, Italy, August. Association for Computational Linguistics.
- 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*, KDD '16, pages 785–794, New York, NY, USA. ACM.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *CoRR*, abs/1406.1078.
- François Chollet and others. 2018. Keras: The Python Deep Learning library, June.
- 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*.
- Yarin Gal and Zoubin Ghahramani. 2015. A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. *arXiv e-prints*, page arXiv:1512.05287, December.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Qianjia Huang, Diana Inkpen, Jianhong Zhang, and David Van Bruwaene. 2018. Cyberbullying intervention based on convolutional neural networks. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 42–51, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Rohan Kshirsagar, Tyrus Cukuvac, Kathy McKeown, and Susan McGregor. 2018. Predictive embeddings for hate speech detection on twitter. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 26–32, Brussels, Belgium, October. Association for Computational Linguistics.
- Susan Li. 2017. NLP-with-Python. <https://github.com/susanli2016/NLP-with-Python>. Accessed: 21-Feb-2020.
- Ping Liu, Wen Li, and Liang Zou. 2019a. NULI at SemEval-2019 task 6: Transfer learning for offensive language detection using bidirectional transformers. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 87–91, Minneapolis, Minnesota, USA, June. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Chris McCormick and Nick Ryan. 2019. BERT Fine-Tuning Tutorial with PyTorch. (2019, July 22). Accessed: 30-Jan-2020. Retrieved from <http://www.mccormickml.com>.
- Hanane Mohaouchane, Asmaa Mourhir, and Nikola S Nikolov. 2019. Detecting offensive language on arabic social media using deep learning. In *2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pages 466–471. IEEE.
- 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*.
- Selma Ayşe Özel, Esra Saraç, Seyran Akdemir, and Hülya Aksu. 2017. Detection of cyberbullying on social media messages in turkish. In *2017 International Conference on Computer Science and Engineering (UBMK)*, pages 366–370. IEEE.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv e-prints*, page arXiv:1912.01703, December.
- John Pavlopoulos, Prodromos Malakasiotis, and Ion Androutsopoulos. 2017. Deeper attention to abusive user content moderation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1125–1135, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830.
- Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive Language Identification in Greek. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Zeses Pitenis. 2019. Detecting offensive posts in greek social media.
- Georgios K. Pitsilis, Heri Ramampiaro, and Helge Langseth. 2018. Detecting offensive language in tweets using deep learning. *CoRR*, abs/1801.04433.
- 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*. Association for Computational Linguistics, pages 1–10, Valencia, Spain.
- Gudbjartur Ingi Sigurbergsson and Leon Derczynski. 2019. Offensive language and hate speech detection for danish. *arXiv preprint arXiv:1908.04531*.
- 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.
- Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christopher Bregler. 2014. Efficient Object Localization Using Convolutional Networks. *arXiv e-prints*, page arXiv:1411.4280, November.
- Ben Trevett and César de Pablo. 2017. pytorch-sentiment-analysis. <https://github.com/bentrevett/pytorch-sentiment-analysis>. Accessed: 10-Feb-2020.
- Zeeraq Waseem, Thomas Davidson, Dana Warmusley, 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, Vancouver, BC, Canada, August. Association for Computational Linguistics.
- Michael Waskom, Olga Botvinnik, Drew O’Kane, Paul Hobson, Saulius Lukauskas, David C Gemperline, Tom Augspurger, Yaroslav Halchenko, John B. Cole, Jordi Warmenhoven, Julian de Ruyter, Cameron Pye, Stephan Hoyer, Jake Vanderplas, Santi Villalba, Gero Kunter, Eric Quintero, Pete Bachant, Marcel Martin, Kyle Meyer, Alistair Miles, Yoav Ram, Tal Yarkoni, Mike Lee Williams, Constantine Evans, Clark Fitzgerald, Brian, Chris Fonnesbeck, Antony Lee, and Adel Qalieh. 2017. mwaskom/seaborn: v0.8.1 (september 2017), September.
- 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. 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*, 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*.