

KEIS@JUST at SemEval-2020 Task 12: Identifying Multilingual Offensive Tweets Using Weighted Ensemble and Fine-Tuned BERT

Saja Khaled Tawalbeh, Mahmoud Hammad, Mohammad AL-Smadi

Computer Science Department

Jordan University of Science and Technology

P.O.Box: 3030 Irbid 22110, Jordan.

sajatawalbeh91@gmail.com, {m-hammad, masmadi}@just.edu.jo

Abstract

This research presents our team KEIS@JUST participation at SemEval-2020 Task 12 which represents shared task on multilingual offensive language. We participated in all the provided languages for all subtasks except sub-task-A for the English language. Two main approaches have been developed the first is performed to tackle both languages Arabic and English, a weighted ensemble consists of Bi-GRU and CNN followed by Gaussian noise and global pooling layer multiplied by weights to improve the overall performance. The second is performed for other languages, a transfer learning from BERT beside the recurrent neural networks such as Bi-LSTM and Bi-GRU followed by a global average pooling layer. Word embedding and contextual embedding have been used as features, moreover, data augmentation has been used only for the Arabic language .

1 Introduction

Natural language processing field has the researchers' attention especially with the rapid use of social media sites, for instance, Twitter, Facebook, YouTube comments, and macro blogs. Consequently, offensive, aggressive and hate-speech language identification problems that perform the automatic detection of these problems from textual data. Moreover, the main motivation to reduce the behavior of hate speech and offensive/aggressive language on user attitude and content, in particular, on social media.

The offensive detection in Arabic social media is a serious task. This refers to the Arabic language contains violent words represent both violent context and not a violent context besides Arabic is dialects language (Elfardy and Diab, 2013). For instance, the word Killing in Arabic meaning represents a violent meaning and not a violent meaning in different contexts appears in (Alhelbawy et al., 2016) tweets. The overall studies that detected on the offensive language have been applied to the English language. However, the research that regards in the Arabic language in this domain of NLP applications has been restricted according to the lack of the resources that tackle the same issue compared to the English language (Mubarak et al., 2017; Abozinadah et al., 2015). The researches that had been conducted in Arabic evaluated on small datasets collected from Twitter API such as Mubarak et al. (2017) has been evaluated their proposed approach on 1100 annotated tweets.

In this research, we describe our participation team KEIS@JUST at SemEval-2020 Task 12 which describes the shared task on offensive language as a multilingual shared task (*i.e.* Arabic, English, Danish, Greek, and Turkish). Moreover, We participated in all languages for the provided subtasks except sub-task-A for the English language. Two approaches have been implemented aim to solve the shared task, the first is performed to tackle both languages Arabic and English, a weighted ensemble consists of Bi-GRU and CNN followed by Gaussian noise and global pooling layer multiplied by weights to improve the overall performance. Consequently, the implemented approach performed to solve sub-task-A (offensive language identification), sub-task-B (automatic categorization), and sub-task-C (offense target

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

identification). The second performed for other languages, a state of art transfer learning from BERT embedding multi-cased 12A pre-trained model besides the recurrent neural networks such as Bi-LSTM and Bi-GRU followed by a global average pooling layer. Consequently, the implemented approach performed to solve sub-task-A (offensive language identification). Word embedding and contextual embedding have been used as features, moreover, data augmentation has been used only for the Arabic language and we rely on the AraVec embedding (Soliman et al., 2017) for data augmentation which aims to create more dataset that helps to train the model. To evaluate our results OffensEval 2020 (Zampieri et al., 2020) provided multilingual Dataset. The best results for the KEIS@JUST team ranked 11th place out of 56 teams with 86.55% F1-macro in the Arabic language, ranked 12th place out of 39 teams with 76.1% F1-macro in the Danish language, ranked 28th place out of 37 teams with 76.1% F1-macro in the Greek language, ranked 32th place out of 46 teams with 73.3% F1-macro in the Greek language.

2 Related Work

Offensive content on social media has recent attention (Schmidt and Wiegand, 2017; Founta et al., 2018; Malmasi and Zampieri, 2018) according to the negative effects on its users, for instance, demeaning comments or hate speech utterance. The offensive language detection on Arabic social media users considered an important step to prevent social society from these negative effects.

Several of previous researches have been presented comprehensive studies which tend to describe the main key of the proposed task Schmidt and Wiegand (2017), and (Fortuna and Nunes, 2018), moreover (Davidson et al., 2017) presents dataset for hate speech detection, (Kumar et al., 2018) presents dataset for aggressive language, and (Zampieri et al., 2019a) presents OLID dataset for the previous shared task of offensive language. Additionally, (Spertus, 1997) shows the earliest efforts in hate speech detection that performs a decision tree-based classifier. Moreover, Offensive identification for sentences have been tried for several languages behind the English such that, Arabic Mubarak et al. (2017) and (Al-Hassan and Al-Dossari, 2019), German (Ross et al., 2017; Fišer et al., 2017; Su et al., 2017).

There are lack of researches in the offensive language for Arabic research community, for instance, (Abdelfatah et al., 2017) introduced k-means for violence utterance on Twitter. MADIMARA has been used to extract morphological features as well as they used TF-IDF to represent dataset on the vector space model. (Malmasi and Zampieri, 2017) presented system to detect hate speech using lexical features and a linear SVM classifier depending on n-grams. Similarity, (Alakrot et al., 2018) introduced SVM classifier trained on word-level features. N-grams and stemming used as features. (Mulki et al., 2019) proposed L-HSAB the first dataset for hate speech and abusive language. The dataset collected from Twitter API focusing on Syrian and Lebanese tweets rich of toxic utterance. The dataset has been trained on Naive Bayes classifier.

For English language, several researchers used transformers ,for instance, (Liu et al., 2019) Proposed a fine-tuned technique for the Bidirectional Encoder Representation from Transformer (BERT) with word unigrams, word2vec, and Hatebase have been used as features. Similarity, (Zhu et al., 2019) Introduced a fine-tuned a BERT based classifier depends on linear SVM trained on character n-gram as a feature. (Pelicon et al., 2019) Proposed a fine-tuned a BERT and LSTM neural network architecture with automatically and manually crafted features were used namely: word embedding, TFIDF, POS sequences, BOW, the length of the longest punctuation sequence, and the sentiment of the tweets features. However, several researchers applied machine and deep learning, for instance, (Mahata et al., 2019) Proposed an ensemble technique consist of Convolutional Neural Network, Bidirectional LSTM with attention, and Bidirectional LSTM + Bidirectional GRU. (Han et al., 2019) Presented two approaches namely: bidirectional with GRU and probabilistic model modified sentence offensiveness calculation (MSOC) trained using word2vec embedding.

3 Methodology

Shared task on Multilingual Offensive Language Identification in Social Media (OffensEval 2020) (Zampieri et al., 2020), in this section, we will describe the shared task and the implemented system.

3.1 Task Description

OffensEval 2020 Zampieri et al. (2020) the shared task on multilingual offensive language (Sigurbergsson and Derczynski, 2020), (Mubarak et al., 2020), Pitenis et al. (2020), (Rosenthal et al., 2020b), (Rosenthal et al., 2020a), and (Çöltekin, 2020), however, the first offensive language task was organized at (Zampieri et al., 2019b), Zampieri et al. (2019a), and the recent aggressive multilingual language proposed by (Bhattacharya et al., 2020). The shared task consist of three sub-tasks the goal of each sub-task represents as the following:

3.1.1 Sub-task A

Offensive language identification, aims to identify whether a tweet contains a non-acceptable language (profanity) or an offensive content. Moreover, sub-task A is a multilingual sub-task for five languages namely: Arabic, English, Danish, Greek, and Turkish. This sub-task is a binary classification, where each tweet has a labeled offensive (OFF) or not offensive (NOT). Our team (KEIS@JUST) participate in all languages for sub-task A except sub-task A for English language.

3.1.2 Sub-task B

Automatic categorization of offense types, aims to identify whether an offense tweet contains targeted or non-targeted profanity and swearing. This sub-task is a binary classification provided only for English language, where each tweet has a labeled targeted (TIN) or untargeted (UNT).

3.1.3 Sub-task C

Offense target identification, aims to determine whether the offense target of the tweet is one of three tags namely: an individual (IND), group (GRP) or other (OTH). Other contains several tags (*i.g.* a situation, an organization, an event, or an issue). This sub-task provided only for English language.

3.2 KEIS@JUST System

3.2.1 Weighted Ensemble (KEIS-BiGRUCNN)

The main intuition of ensemble that combining the predictions comes from (Yu, 1977) that combines two regression techniques, after that (Dasarathy and Sheela, 1979) had been presented the combination of two or more models. In this research, we will perform weighted ensemble technique consist of two models namely: KEIS-BiGRU and KEIS-CNN. We have applied a weighted ensemble that aims to boost the performance of our system. The following provides more details about the implemented techniques.

- **Bidirectional-GRU (KEIS-BiGRU):**

Recurrent Neural Network (RNN) suffers from a gradient vanishing problem. Long Short Term Memory (LSTM) Hochreiter and Schmidhuber (1997) has been proposed to solve the mentioned problem. Gated Recurrent Unit (GRU) (Chung et al., 2014) as well as have been proposed to solve the gradient vanishing problem. Two gates have been used (reset and update gate) in GRU architecture.

It's started with passing a sequence of words through an embedding layer followed Bidirectional GRU layer of 128 neurons, then Gaussian Noise of (0.1). Afterword, Global Average Pooling has been used to extract the discriminative features of the input tweet to prepare that for the next layer. Dense layer of 35 neurons has been applied followed by Dropout layer of (0.2) to prevent overfitting. Finally, the output layer will be Dense of (1) neuron with sigmoid function.

- **Convolutional Neural Networks (KEIS-CNN):**

In deep learning, Convolutional Neural Networks (CNN) show the significant success in the fields

parameter	Value
number of epochs	20
batch size	128
optimizer	amsgrad
learning rate	0.01
weight for KEIS-CNN	0.4
weight for KEIS-BiGRU	0.6
kernel regularizer	L2 with (0.01)

Table 1: Proposed Models (KEIS-BiGRUCNN) Hyper-parameters

of computer vision. The first CNN architecture for text classification proposed by (Kim, 2014) which provides a remarkable enhancement on the performance of NLP tasks. Consequently, it can obtain the linguistic patterns from window of sequence words represented as embedding vectors.

It's started with passing a sequence of words through an embedding layer followed by Gaussian Noise of (0.1). As we know Conv2D have to reshape the input to be compatible to receive the previous shape. Four Conv2D layers have been used. Each individual Conv2D sharing different filter size (1, 3, 5, 7) respectively. Moreover, the number of filters is 36 for all layers which aims to obtain the local information features. Afterward, each Conv2D layer passed to Max Pool 2D layer (MaxPool2D). In the last step, each layer has been concatenated together which aims to identify better output. In order to feed the next layer, we used Dropout layer of (0.25) to reduce over-fitting followed by dense layers of 35 neurons. The output is fed into single sigmoid which can obtain the output class of the given tweet.

For training step, the implemented KEIS-BiGRUCNN ensemble approach applied Soliman et al. (2017) embedding as pre trained model with 300 dimensions for Arabic language that prepared for training step. In contrast, we applied word2vec embedding for English proposed by (Mikolov et al., 2013), the pre trained embedding available at github account ¹ with 400 dimensions. Several hyper-parameter have been used for optimization. Table 1 provides more details about the value of each parameter have been used during the training step for both approaches. It's worth mentioning that amsgrad optimizer proposed by (Tran and others, 2019) the updated version of adam optimizer with slight enhancement regards to the system performance. The final step, after we obtained the final prediction for each model, the predictions have been multiplied by the best chosen weight to enhance the over all results. The ensemble architecture is shown in Fig. 1 KEIS-BiGRUCNN has been used to solve sub-task A for Arabic language and sub-task B,C for English language.

3.2.2 BERT Fine-Tuned (BERT-Bi)

In the recent years, contextual embedding shows the significant progress in the NLP research field. Consequently, according to the reported results in several researches (*i.e.* Zhu et al. (2019)) it's it's outperform the deep learning approaches. The transformer considered as an encoder-decoder architecture applied on attention mechanisms tasks. More particularly, Google has been released BERT (Devlin et al., 2018) which stands for Bidirectional Encoder Representations from Transformers. Our intention to solve the offensive detection shared task using fine-tuned the BERT by adding Gaussian Noise layer followed by bidirectional LSTM (Hochreiter and Schmidhuber, 1997) consist of 300 neurons, GRU Chung et al. (2014) consist of 300 neurons, and global average pooling (GAP) to extract the discriminative features from the past layer and keep them to the next layer. L2 regularization and Dropout have been used to prevent overfitting. The classification layer used to find the final predictions dense layer of 1 neuron with sigmoid activation function and TruncatedNormal kernel initializers. Consequently, the implemented

¹<https://github.com/felipebravom/AffectiveTweets/releases/download/1.0.0/w2v.twitter.edinburgh10M.400d.csv.gz>

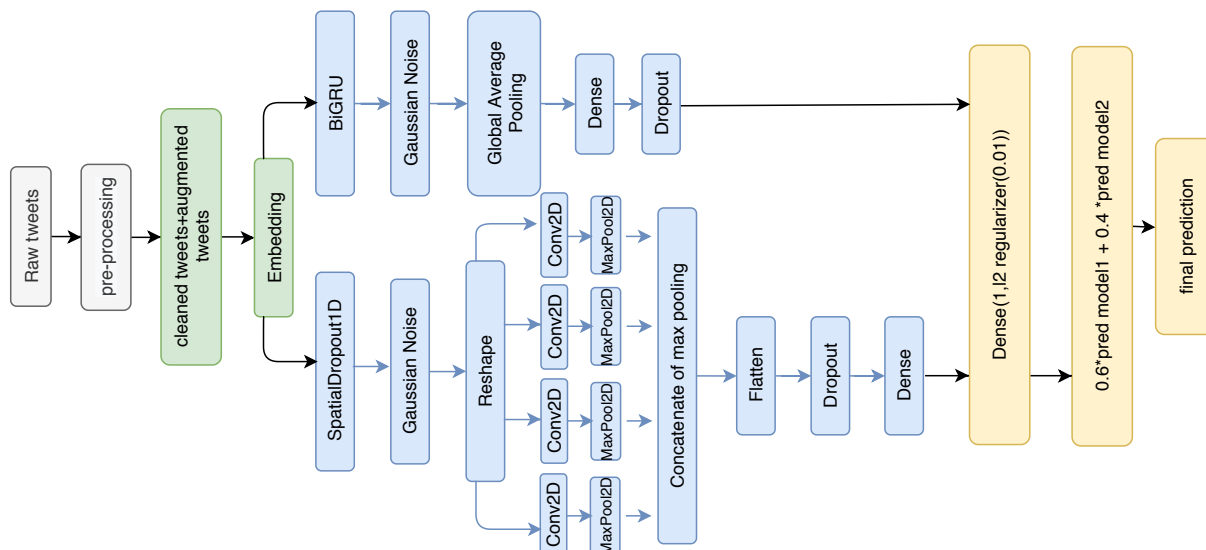


Figure 1: The implemented architecture of KEIS-BIGRU-CNN model




Figure 2: The architecture of BERT fine-tuning model (BERT-Bi)

approach called BERT-Bi.

The implemented BERT-Bi based on transfer learning architecture that has used in common specially in image classification and computer vision (Litjens et al., 2017). Moreover as we mentioned earlier in Sec 2, the applied of transformers show the promising results compared to deep learning approaches. For instance, BERT developers created several pre-trained models such as uncased, cased, and multi_cased to represent the semantic relationships among text as well as it could be applied as an independent classifier in different NLP domains (*i.e.* offensive language detection). In this research, we used multi_cased model since it trained on multi languages based on transfer learning architecture to tackle the shared task problem. The BERT-Bi architecture shown in Fig. 2 used to solve sub-task A for three languages namely: Greek, Danish, and Turkish. Whereas the special [CLS] should be added at the beginning of each tokenized tweet. the special [SEP] should be added at the end of each tokenized tweet. The Attention Mask represented as an array of 1's and 0's. In order to implement proposed model, several parameters have been used. According to the experimental results the best parameter as follows: batch size= 16, optimizer= Adam, learning rate= 2e-5, and finally BERT max length= 60.

Data File/Lang	Arabic	Danish	Greek	Turkish	Eng-B	Eng-C
Train	OFF= 1410 NOT= 5590	OFF= 398 NOT= 2568	OFF= 2539 NOT= 6260	OFF= 6132 NOT= 25630	TIN= 149550 UNT= 39424	IND= 120330 GRP= 22176 OTH= 7043
Validation	OFF= 179 NOT= 821	20% split	20% split	20% split	20% split	20% split
Test	OFF= 402 NOT=1598	OFF= 41 NOT= 288	OFF= 242 NOT= 1302	OFF= 716 NOT= 2812	UNT= 572 TIN= 850	IND= 580 GRP= 190 OTH= 80

Table 2: The Dataset distribution

id	lang	tweet	label
1689	Arabic	@USER اتعلم يا متعصب يا معدوم الذمه	OFF
1713	Danish	Haha, det er genialt!	NOT
19321	Turkish	@USER Burası da fena değil atkafalı	OFF
1159528564925984768	English	everyone talks shit in LA	TIN

Table 3: Examples from different languages that represents the dataset

4 Experimental Setup

4.1 Dataset

Multilingual dataset have been provided with five languages namely: Arabic, Danish, English, Greek, and Turkish. The annotation follows the hierarchical tagset for the previous Offensive Language Identification Dataset (OLID) Zampieri et al. (2019a). The provided dataset to tackle Task-12 at SemEval 2020 which has been obtained from Twitter using API’s. Task 12 offensEval 2020 provided three sub-tasks: (1) if the tweet offensive (OFF) or non-offensive (NOT), (2) if the tweet is targeted (TIN) or un-targeted (UNT), and (3) If the target is an individual (IND), group (GRP) or other (OTH). The provided dataset is multilingual and imbalanced refers to the distribution for each sub-task including the labels provided for the three sub-tasks with tab separated file format. Table 2 shows the distribution of the available dataset. Table 3 provide examples that represents dataset for all languages.

4.2 Data Pre-processing

The convenient process regarding social network dataset such that, Facebook and Twitter, tweets, and posts which contain such noisy data and slang language. In the raw text, it should remove the special character, punctuation marks (*, @#-(—), URLs, and user mentions. The normalization was necessary since some words written on short-cut format, the elongation was also removed (e.g congrats مبروووووك). Finally, numbers and English characters were also removed for Arabic. Moreover, the emojis have been removed.

4.3 Embeddings

Several well-known word embedding are provided to extract the vector representation of the input tweets with aims to capture the semantic features for each word and the relationship among them. Word2Vec has been provided by Mikolov et al. (2013), Glove (Pennington et al., 2014), AraVec Soliman et al. (2017) and the recent contextual embedding EIMo by (Peters et al., 2018) and BERT Devlin et al. (2018). In this research, we used AraVec, Word2Vec and the pre-trained BERT embedding to trained the performed model. It is a language representation model and becoming the state of art model for the most of NLP research.

4.4 Data Augmentation

It is a way to improve the performance of NLP models, data augmentation should appear on a deep understanding of our dataset including structure and content. The impact of using data augmentation technique will depend on that technique itself, where each one able to learn something different compare to

Model	without aug	with aug
KEIS-BiGRU	83.6%	87.6
KEIS-CNN	83.5%	87.3
KEIS-BiGRUCNN	83.8%	87.9

Table 4: The results on the validation set for Arabic language (aug refers to augmentation)

Task A Multilingual using BERT-Bi	F1	Task B/C English using BiGRUCNN	F1
Danish	78%	Task B	48.3%
Greek	78.5%	Task C	58.9%
Turkish	72%	-	-

Table 5: : The results on the validation set for other languages/all results computed using Macro-F1

others and generate a different impact as well. There are several techniques used in the data augmentation, in this research we well performed the technique depending on pre-trained AraVec Soliman et al. (2017). The first step load embedding and prepare the dataset. Afterward, making a synonym dictionary depending on the most frequent words.

4.5 Discussion

Our results extracted using SAJA CODALab user name and the team name is KEIS@JUST. The reported results on the validation set are presented in table 4 for Arabic and table 5 for other languages. Table 4 presents the results using data augmentation for Arabic language. It shows the enhancement of using the augmentation regarding the overall performance the KEIS-BiGRUCNN approach achieved F1= 87.9% on Arabic validation set. Moreover, BERT-Bi approach achieved F1= 78% on Danish validation set (see table 5). As we mentioned above in sec 3.2 presents KEIS@JUST System to present the results which consist of a)KEIS-BiGRUCNN used to solve sub-task A for Arabic language and sub-task B,C for English language. b) KEIS-BERT-Bi used to solve sub-task A for other languages. To prevent overfitting during the training step, the early stopping and checkpoints have used among the training set and the validation set and keep track of the loss value at the end of each training epoch. Moreover, the learning rate reduction has used. Fig. 3 and Fig. 4 show the model training.

5 Results and Findings

In order to evaluate the implemented approaches, F1-Macro has been used according to the shared task instruction. Table 6 presents the results of the participants models for Arabic, Danish, and Greek




Figure 3: KEIS@BiGRU loss plots for training and validation data




Figure 4: KEIS@CNN loss plots for training and validation data

Arabic Rank	F1	Danish Rank	F1	Greek Rank	F1
1	90.17%	1	81.2%	1	85.2%
2	90.15%	2	80.2%	2	85.1%
11	86.55%	12	76.1%	28	77.3%

Table 6: : The results on the test set for Arabic, Danish, and Greek languages/all results computed using Macro-F1 / Our results appears in bold

Turkish Rank	F1	Eng Rank Task B	F1	Eng Rank Task C	F1
1	81.57%	1	74.6%	1	71.45%
2	81.66%	2	73.6%	2	66.99%
32	73.3%	43	27.7%	33	48.17%

Table 7: : The results on the test set for Turkish and English language-task B and C/all results computed using Macro-F1 / Our results appears in bold

languages and macro-average results, as presented in the table, results for are competitive. Moreover, Table 7 presents the result for Turkish and English language. The KEIS@JUST team, achieved 11th place with F1= 86.55% in Arabic, achieved 12th place with F1= 76.1% in Danish, achieved 28th place with F1= 77.3% in Greek, achieved 32th place with F1= 73.3% in Turkish. In contrast, achieved F1= 27.7% in English sub-task B and achieved F1= 48.17% in English sub-task C.

6 Conclusion

In this research, presented the KEIS@JUST participation at SemEval-2020 Task 12 which represents shared task on multilingual offensive language. We have participated in all the provided languages for all subtasks except sub-task-A for the English language. Two main approaches have been developed the first one is performed to tackle both languages Arabic and English, a weighted ensemble consists of Bi-GRU and CNN followed by Gaussian noise and global pooling layer multiplied by weights to improve the overall performance. The second one performed for other languages, we investigated the main impact of developing a transfer learning approach from BERT transformer beside the recurrent neural networks such as Bi-LSTM and Bi-GRU followed by the global average pooling layer for other languages. Word embedding and contextual embedding have been used as features, moreover, we investigated how data augmentation affect the results using Arabic dataset.

Acknowledgments

This research is partially funded by Jordan University of Science and Technology, Research Grant Number: 20170107.

References

- Kareem E Abdelfatah, Gabriel Terejanu, and Ayman A Alhelbawy. 2017. Unsupervised detection of violent content in arabic social media. *Computer Science & Information Technology (CS & IT)*, 7.
- Ehab A Abozinadah, Alex V Mbaziira, and J Jones. 2015. Detection of abusive accounts with arabic tweets. *Int. J. Knowl. Eng.-IACSIT*, 1(2):113–119.
- Areej Al-Hassan and Hmood Al-Dossari. 2019. Detection of hate speech in social networks: a survey on multilingual corpus. In *6th International Conference on Computer Science and Information Technology*.
- 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.
- Ayman Alhelbawy, Poesio Massimo, and Udo Kruschwitz. 2016. Towards a corpus of violence acts in arabic social media. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1627–1631.

- Shiladitya Bhattacharya, Siddharth Singh, Ritesh Kumar, Akanksha Bansal, Akash Bhagat, Yogesh Dawer, Bornini Lahiri, and Atul Kr. Ojha. 2020. Developing a multilingual annotated corpus of misogyny and aggression.
- 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.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Belur V Dasarathy and Belur V Sheela. 1979. A composite classifier system design: concepts and methodology. *Proceedings of the IEEE*, 67(5):708–713.
- 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. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Heba Elfardy and Mona Diab. 2013. Sentence level dialect identification in arabic. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 456–461.
- Darja Fišer, Tomaž Erjavec, and Nikola Ljubešić. 2017. Legal framework, dataset and annotation schema for socially unacceptable online discourse practices in slovene. In *Proceedings of the first workshop on abusive language online*, pages 46–51.
- Paula Fortuna and Sérgio Nunes. 2018. A Survey on Automatic Detection of Hate Speech in Text. *ACM Computing Surveys (CSUR)*, 51(4):85.
- Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large Scale Crowdsourcing and Characterization of Twitter Abusive Behavior. *arXiv preprint arXiv:1802.00393*.
- Jiahui Han, Shengtan Wu, and Xinyu Liu. 2019. jhan014 at SemEval-2019 task 6: Identifying and categorizing offensive language in social media. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 652–656, Minneapolis, Minnesota, USA, June. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- 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)*, Santa Fe, USA.
- Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. 2017. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88.
- Ping Liu, Wen Li, and Liang Zou. 2019. 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.
- Debanjan Mahata, Haimin Zhang, Karan Uppal, Yaman Kumar, Rajiv Ratn Shah, Simra Shahid, Laiba Mehnaz, and Sarthak Anand. 2019. MIDAS at SemEval-2019 task 6: Identifying offensive posts and targeted offense from twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 683–690, Minneapolis, Minnesota, USA, June. Association for Computational Linguistics.
- Shervin Malmasi and Marcos Zampieri. 2017. Detecting hate speech in social media. *arXiv preprint arXiv:1712.06427*.
- Shervin Malmasi and Marcos Zampieri. 2018. Challenges in Discriminating Profanity from Hate Speech. *Journal of Experimental & Theoretical Artificial Intelligence*, 30:1–16.
- 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.

- Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. Abusive language detection on arabic social media. In *Proceedings of the First Workshop on Abusive Language Online*, pages 52–56.
- 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*.
- Hala Mulki, Hatem Haddad, Chedi Bechikh Ali, and Halima Alshabani. 2019. L-hsab: A levantine twitter dataset for hate speech and abusive language. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 111–118.
- Andraž Pelicon, Matej Martinc, and Petra Kralj Novak. 2019. Embeddia at SemEval-2019 task 6: Detecting hate with neural network and transfer learning approaches. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 604–610, Minneapolis, Minnesota, USA, June. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive language identification in greek. *arXiv preprint arXiv:2003.07459*.
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2020a. A Large-Scale Weakly Supervised Dataset for Offensive Language Identification. In *arxiv*.
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2020b. A large-scale semi-supervised dataset for offensive language identification.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurowsky, and Michael Wojatzki. 2017. Measuring the reliability of hate speech annotations: The case of the european refugee crisis. *arXiv preprint arXiv:1701.08118*.
- 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. 2020. Offensive Language and Hate Speech Detection for Danish. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Abu Bakr Soliman, Kareem Eissa, and Samhaa R El-Beltagy. 2017. Aravec: A set of arabic word embedding models for use in arabic nlp. *Procedia Computer Science*, 117:256–265.
- Ellen Spertus. 1997. Smokey: Automatic recognition of hostile messages. In *Aaai/iaai*, pages 1058–1065.
- Hui-Po Su, Zhen-Jie Huang, Hao-Tsung Chang, and Chuan-Jie Lin. 2017. Rephrasing profanity in chinese text. In *Proceedings of the First Workshop on Abusive Language Online*, pages 18–24.
- Phuong Thi Tran et al. 2019. On the convergence proof of amsgrad and a new version. *IEEE Access*, 7:61706–61716.
- Chong Ho Yu. 1977. Exploratory data analysis. *Methods*, 2:131–160.
- 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*.
- Jian Zhu, Zuoyu Tian, and Sandra Kübler. 2019. UM-IU@LING at SemEval-2019 task 6: Identifying offensive tweets using BERT and SVMs. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 788–795, Minneapolis, Minnesota, USA, June. Association for Computational Linguistics.