

## **Automatic offensive language detection from Twitter data using machine learning and feature selection of metadata**

ARAUJO DE SOUZA, Gabriel and DA COSTA ABREU, Marjory  
<<http://orcid.org/0000-0001-7461-7570>>

Available from Sheffield Hallam University Research Archive (SHURA) at:  
<https://shura.shu.ac.uk/26018/>

---

This document is the Accepted Version [AM]

### **Citation:**

ARAUJO DE SOUZA, Gabriel and DA COSTA ABREU, Marjory (2020). Automatic offensive language detection from Twitter data using machine learning and feature selection of metadata. In: IEEE World Congress on Computational Intelligence (IEEE WCCI). IEEE. [Book Section]

---

### **Copyright and re-use policy**

See <http://shura.shu.ac.uk/information.html>

# Automatic offensive language detection from Twitter data using machine learning and feature selection of metadata

1<sup>st</sup> Gabriel Araujo De Souza

Federal University of Rio Grande do Norte (UFRN)  
Natal, Brazil  
gabriel\_feg@hotmail.com

2<sup>nd</sup> Márjory Da Costa-Abreu

Sheffield Hallam University  
Sheffield, UK  
m.da-costa-abreu@shu.ac.uk  
<https://orcid.org/0000-0001-7461-7570>

**Abstract**—The popularity of social networks has only increased in recent years. In theory, the use of social media was proposed so we could share our views online, keep in contact with loved ones or share good moments of life. However, the reality is not so perfect, so you have people sharing hate speech-related messages, or using it to bully specific individuals, for instance, or even creating robots where their only goal is to target specific situations or people. Identifying who wrote such text is not easy and there are several possible ways of doing it, such as using natural language processing or machine learning algorithms that can investigate and perform predictions using the metadata associated with it. In this work, we present an initial investigation of which are the best machine learning techniques to detect offensive language in tweets. After an analysis of the current trend in the literature about the recent text classification techniques, we have selected Linear SVM and Naive Bayes algorithms for our initial tests. For the preprocessing of data, we have used different techniques for attribute selection that will be justified in the literature section. After our experiments, we have obtained 92% of accuracy and 95% of recall to detect offensive language with Naive Bayes and 90% of accuracy and 92% of recall with Linear SVM. From our understanding, these results overcome our related literature and are a good indicative of the importance of the data description approach we have used.

**Index Terms**—Offensive Language Detection, Naive Bayes, Linear SVM, Attribute Selection, Twitter

## I. INTRODUCTION

In the face of popularisation of social media such as Facebook, Twitter, Instagram and Tik Tok, the communication between people has become faster and easier. In these communication mechanisms, people can express their feelings, criticism, opinions, achievements, etc. However, many times these networks are used to publicise the hate speech though offensive words. The types of the offense can be directed to most diverse aspects including Ethnicity, Economic status, Religion, Sexual Orientation, and so on. Thus, the big problem with this is relies in the fact that the offense is presented to individuals or groups and this can be potentially harmful to them [1], [2].

The challenge in automated detection of offensive language lies in the fact that, in social networks, the used language contains a specific format that belongs to the environment. In

this context, several word abbreviations are used and various forms of expression intensification and word modification, such as several letters repetitions (e.g.: looooooved, gooood) and excessive use of punctuation (e.g.: i loved!!!!, what????). Thus, the original text must be modified through a very important preprocessing stage to get a form that keeps the original sense and can somehow coincide with other similar posts [3].

For this, a software that can be able to detect this type of offenses in a social network is a very important advancement to ensure the security and mental health of users [4]. Another considered application is detecting users that commonly perform cyberbully acts and aggression. From this, measures must be taken to punish and block aggressive users [5].

Thus, this paper propose an investigation to detect offensive language in twitter data. After a review of the literature, two techniques were chosen for this task: Linear SVM and Naive Bayes. The main objective is to improve these algorithms when compared with similar experiments found in the literature. For this, some results already published are presented and a comparison of implemented techniques will discussed.

This work is organized as: Section II presents a review of the literature with the recent works about automatic detection of offensive language or similar. Section III presents the database structure and the process of data normalization. Section IV introduces the concepts of chosen machine learning algorithms to describe the formulas and steps of the algorithm. Section V shows and explains the results obtained in our experiments and Section VI reports the conclusions obtained from the work.

## II. RELATED WORKS

The investigation of hate speech in social media is a relatively new research area, but, despite that, it has created a lot of attention and have already whole events dedicated to the subject as well as an increase of related publications. This session will present and discuss some of the main findings of the most recent and relevant publications.

As an example, the International Workshop on Semantic Evaluation SemEval-2019<sup>1</sup> [6] presented tasks that focused

<sup>1</sup><http://alt.qcri.org/semeval2019>

on detection and categorising of offensive language in social media. The three main sub-task were: offensive language identification, automatic categorisation of offense types, and offense target identification. For the first sub-task, the messages were classified as offensive or not offensive. A twitter post was label as offensive if it contained any offensive or profane language. The Deep Learning BERT [7] presented better results for this task [8]. In the second sub-task, the goal was to predict the type of offense. For this, two classes were used: Insult and Untargeted. A twitter post was classified as 'Insult' when it contained an insult to an individual or group; and, a twitter post was labeled as 'Untargeted' when that the post contained non-acceptable language (swearing). The best solution found for this problem used a rule-based approach with a keyword filter, such as hashtags, signs, emoticons, and other features [9]. The last sub-task focused on the target of offenses. The used classes were 'Individual' for an offense to a unique user, 'Group' for an offense to a group of people, and 'Other' for an offense to an organisation, a situation, an event, or an issue. The team with better results also used BERT for this problem [10]. This Deep Learning also used to detect offensive language in German texts [11].

In [12], a similar application was explored. They created a new dataset using the Twitter API for twitter data classification as hate speech, offensive language or neither. In their dataset, they collected a set of 85.4 million twitter samples from about 33 thousand Twitter users. From there, they built a set of 24k labeled twitter samples. Features, such as bigram, unigram, and trigram were weighted by their TF-IDF and were used for the classification task. Other features included binary and count indicators for hashtags, mentions, retweets, and URLs. They tested a large number of classifiers: logistic regression, Naive Bayes, decision trees, random forests, and linear SVMs (Support Vector Machine). Through their experiments, it was found that Logistic Regression and Linear SVM tended to perform better results. The best model obtained an overall precision of 0.91, recall of 0.90, and F1-score of 0.9. However, the classifier did not present good results to detect hate speech, for this class, the precision and recall were 0.44 and 0.61, respectively.

A Deep Learning for classification messages of social media was proposed in [13]. The labels considered were racism, sexism or neither. For their experiments, they joined various Long Short-Term Memory (LSTM) models. In their classification, the features defined a user's tendency towards posting messages in any used classes, the set of messages posted by a user, and subsets that contained labeled messages. The method used was independent of the language. The proposed methods obtained better results to detect sexists messages (around 0.99 of precision and F1-Score). Neutral messages also presented better results (0.94 for precision), but racist messages obtained inferior results with about 0.75 and 0.70 of precision and F1-Score respectively.

The problem of detecting hate speech was expanded in [14] for vulnerable community identification. The features extracted from the messages was performed using techniques

such as convert words to vector and n-grams. For the process of hate speech detection, they used the Gated Recurrent Unit (GRU) and a variety of RNNs. These classifiers provided an accuracy of about 0.92. A Convolutional Neural Networks (CNN) was proposed in [15] for classifying offensive tweets written in English. The labels used in this work were offensive, abusive or hate-inducing. The best results obtained an accuracy of 0.83 and a precision of 0.80.

In [16], the models SVM, bidirectional Long Short-Term-Memory (BiLSTM), and CNN were used to classify messages as offensive or not offensive. In the experiments, the BiLSTM obtained a better precision to detect offensive messages (0.81). For detection of not offensive language, the precision was 0.83. SVM and CNN model obtained a precision of 0.66 and 0.78 respectively to detect offensive messages and 0.80 and 0.87 to detect not offensive messages.

Based on what was presented in the previously listed works, we have identified a lack of investigation regarding the type of features used in order to investigate offensive language. Also, a better fine tuning approach to the standard classical classification techniques was not fully explored, which can lead to poorer results. Thus, we reinstate that our main goal with this work is to investigate the quality of used features for this problem as well as the fine tuning approach of classical classification techniques.

### III. HATE SPEECH ON TWITTER POSTS: DATABASE

The used dataset was collected by [12] and contains 24783 tweets. From these, 1430 are classified as hate speech, 19190 as offensive language, and 4163 as normal language. Due to the low amount of messages classified as hate speech in comparison with the other labels and the low performance to detect hate speech with this dataset describe in [12], for this work, we have chosen to use only the labels of offensive and normal language.

The first reason for the choice of this dataset is due to the variety of examples of offensive twitters. The organization of data in this dataset was another important reason for your choice because it facilitated data processing and manipulation. The good results obtained with machine learning algorithms present in [12] for offensive language detection was the final reason for us to decide to use it.

In order to select a larger number of messages tagged as an offensive language than normal, it is necessary to balance the dataset [19]. For this, we have chosen randomly a subset of messages marked as offensive with the same size as the set of normal messages. Thus, we selected a set of 4163 normal and 4163 offensive tweets, creating a dataset with a total size of 8326 messages. We have selected randomly offensive messages to provide more diversity in the dataset. A variety of seeds can be used to test different datasets configurations. Based on related work [8], [13], [15], [16], we have decided to divide the dataset with more messages in the training dataset. Thus, we have allocated 60% of data to train and 40% to test. The choice of training and test dataset was made at random as

TABLE I  
RESUME OF RELATED WORKS

Reference	Database	Objective	Feature Extraction	Classification	Results
[8]	Offensive Language Identification Dataset (OLID)	Categorisation of offensive language in social media	Emoji substitution, HashTag segmentation and convert all the text into lowercase (Misc)	Logistic Regression (LR), LSTM and BERT	LR: 72% Acc LSTM: 76% Acc BERT: 84% Acc
[9]	Offensive Language Identification Dataset (OLID)	Automatic categorisation of offense types	Rule-based approach with a keyword filter based on a Twitter language	Modified sentence offensiveness calculation (MSOC) and RNN	RNN: 86% Acc MSOC: 92% Acc
[10]	Offensive Language Identification Dataset (OLID)	Offense target identification	Pre-processing and pre-trained word embedding based on GloVe.	BERT-Base, Multilayer Perceptron Network (MLP) and Soft Voting Classifier (SVC)	BERT-Base: 72% Acc MLP: 68% Acc SVC: 69% Acc
[11]	Dataset of German-language tweets provided in context of the GermEval Shared Task 2 (2019)	Offensive language identification for German-language texts	Replace all user mentions to a token Name	BERT	76% F1-Score
[12]	They created their own dataset described in [12]	Automatic hate-speech detection on social media is the separation of hate speech from other instances of offensive language.	Lower case, create bigram, unigram, and trigram features, TF-IDF and others	LR e SVM	Precision of 0.91 and Recall of 0.90 for offensive language detection and precision of 0.44 and recall for hate speech
[13]	A dataset of approximately 16k short messages from Twitter, that was made available by [17]	Detecting Offensive Language	Define the three features representing a user's tendency towards posting Neutral, Racist and Sexist content	An ensemble of Recurrent Neural Network (RNN) classifiers	Precision: 94% and Recall: 96%
[14]	Public hate speech datasets available in [18]	Vulnerable community identification using hate speech detection on social media	Convert words to vector and n-grams	Gated Recurrent Unit (GRU) and a variety of RNNs	RNN-GRU: 92% Acc GBT: 92% Acc RNN-LSTM: 91% Acc
[15]	Hinglish dataset HEOT and [12]	Detecting Offensive Tweets in Hindi-English	Removal of punctuations, URLs and user mentions; lower case; remove stop words and others	CNN	Accuracy of 83% and recall of 71%
[16]	Offensive Language Identification Dataset (OLID)	Predicting the Type and Target of Offensive Posts in Social Media	Pre-processing of text	SVM, bidirectional Long Short-Term-Memory (BiLSTM), and CNN	SVM: 76% precision and 78% recall BiLSTM: 82% precision and 82% recall CNN: 82% precision and 82% recall

well. Both sets contained the same number of offensive and normal messages.

#### A. Data processing

A tweet contains a diversity of elements that can confuse a text classifier, for example, user names, hashtags, URLs and emojis. This occurs because there exists a variety of different shapes in these texts which can be very complex to find patterns. Grammar errors and excessive use of repeat letters are other problems because these end up generating several different forms of the same term or word. Thus, before the

data can be analysed by any model, it is necessary to perform a preprocessing in the text to remove or reduce the mentioned problems without losing the semantic meaning of the message. It is possible to find in the literature techniques to perform normalisation [10], [15], convert to lower case [10], [15] and removal of stop words [10], [15] with the propose to generate a new text with the same sense, but in a way that provides better performance to the text classification algorithms.

1) *Data Normalisation*: This process consists of mapping divergent text that belongs to the same class in a label. In

this work, we have converted all the hashtags, user names, emojis, URLs and retweets to a tag that represents each information. Thus, every time that a hashtag is found in the text, it is replaced to tag  $\langle \text{hashtag} \rangle$ , emojis are replaced to  $\langle \text{emojis} \rangle$  and so on [10]. Other important point in the normalisation processing is the removal of all text punctuation. This is a very common process and contributed for a clean text with focus in the words [10], [15].

For an algorithm, the word "Car" is different from "car". In a tweet, the words can be written with different uses of upper case and lower case, The words can also be written completely in lower case or completely in upper case or even contain upper and lower case occurrences simultaneously. For reducing this divergence, we have converted all text to lower case.

In a text, words such as articles, pronouns, connectors, etc are considered irrelevant to the process of classification [20]. This type of text appears frequently and so can hamper the training process. Thus, these words are known as stop words. A common method in a preprocessing of text stands is to remove all stop words to create a clean text with just what is relevant. For this work, we have created a list of words that are considered stop words, and we have removed all the words in a tweet that were present in the list.

Two forms of preprocessing were used, the first, considering the tags in a process of data normalisation which we called 'Data Type A'. The second form which removed all the hashtags, emojis, URLs, Retweets, User names and stop words leaving a very clean text, which we called 'Data Type B'. All forms will be used in each proposed technique.

#### IV. METHODOLOGY

In this section, we will present the techniques that were used in our experiments. After a review of the literature, we have found many techniques to detect and classify offensive language. Some deep learning and neural networks are used for this purpose and each paper presents different results regarding performance. SVM-based solution has not appeared, nevertheless, the authors of the used dataset obtained good results [12], [16]. The Naive Bayes classifier was also not mentioned in the recent literature, but it has presented good results with similar problems [21], for this, we have decided to implement these techniques and evaluate their results. A brief description of these techniques will be presented in Sections IV-A and IV-B.

##### A. Naive Bayes Classifier

In this classifier, a table of words occurrences, also named bag of words, is created in the training process. This table contains words and a number of occurrences for this word in each class. Another information expressed in the table is the total of words for each class and the total words on the training set. The stop words are not considered [21]. An example of this table is presented in Table II.

After the training step, the mentioned table is ready to be used by the classifier. For this, the classifier receives a

TABLE II  
TABLE OF WORDS OCCURRENCES

class/word	x	y
employees	5	1
protesting	0	4
shutdown	2	3
melancholy	1	5
total	8	13

preprocessing twitter and a list of words is created, thus, for each class is calculated a score for this message that informs how pertinent this message is for this class. The class that obtained the biggest score is returned by the classifier. The score is calculated by the following equation:

$$\text{score}(y_i, W) = \log \frac{y_i^t}{t} + \sum_{i=1}^n \log \frac{w_i^{y_i} + mp}{y_i^t + m} \quad (1)$$

Where,

- $y_i$  is the  $i^{th}$  class represents in the model.
- $W$  is a set of words.
- $y_i^t$  is a total of words classified as  $y_i$  class in the training data.
- $t$  is the total of words in the training data.
- $w_i^y$  is the amount of times that the  $i^{th}$  word in the  $W$  set is classified as  $y_i$  class in the training data.
- $m$  and  $p$  are parameters:  $p = 0.5$  and  $m = 1$

##### B. Support Vector Machine

The objective of this classifier is to find a hyper-plane in a space with  $n$  dimensions, where  $n$  is a number of features that distinctly classifies the data points. In the training process, the principal goal is to find limitation points that can separate objects of distinct classes. In order to maximise the margin of the classifier, support vectors are used. These are data points that are closer to the hyperplane and influence the orientation of the hyperplane [22].

The first step to run the SVM is to prepare the data by transforming each message in a numerical feature vector. A common technique in the literature is to use TF-IDF (Term Frequency - Inverse Document Frequency) [23]. These metrics inform how interesting a word is for a type of document. However, in our experiments, we noticed a better result using the technique 'bag of words', also used for the Naive Bayes classifier and describe in IV-A. After a construction of table with examples of offensive and normal texts, each phase is associate with a score for a offensive message and other score to normal language, thus, each phase receive two features demarcating how much belongs to the set of offensive language and the set of normal language respectively. With the data, a default gradient for each class is generated and in the process of the train they are updated for each wrong classification or correct classification by the equations:

$$\omega = \omega - \alpha \cdot (2\lambda\omega) \quad (2)$$

when a correct classification is realised and

$$\omega = \omega + \alpha \cdot (y_i \cdot x_i - 2\lambda\omega) \quad (3)$$

when an incorrect classification is realised.

In the equations 2 and 3, the  $\omega$  is a gradient,  $\alpha$  is the learning rate and  $\lambda$  is a regularisation parameter defined as  $1/epochs$ . After the training process, the product is to have the weights  $w_1$  and  $w_2$  and these are used in the classification process. For this step, the same process of converting text into two numeric features is realised. By applying the weights in the values the next step, we can verify if the result is great than 1, so the text is classified as offensive, otherwise, it is normal.

## V. RESULTS

In this section, we will present the results of our experiments. The two algorithms (Linear Support Vector Machine and Naive Bayes) presented previously were fully implemented. For each algorithm, we tested the two configurations of data describes in section III-A: Data Type A and Data Type B. For the Data Type A, the Naive Bayes achieved a gain of performance of around 1.5% when compared with Data Type B. For the Data Type A, the Linear SVM demonstrated a greater difficulty in parameter setting to find satisfactory results. We tested to change the configuration of the alpha parameter with the values 0.1, 0.01 and 0.001 besides configuring the max epochs to 100 and 500, but in all possible combinations of theses parameters, the algorithm always classified any text as offensive. However, with the Data Type B, it was possible to find good results configuring the parameter values to alpha equals 0.01 and epochs equals 100.

Regarding the dataset, it was previously mentioned that the number of examples of offensive messages is greater than the number of normal messages and that to create the dataset, we have selected randomly a subset of offensive messages with the same size as the set of normal messages. For the Linear SVM, this process was generated a unique time because another process of randomisation is performed to define the order of messages evaluation in the training process. However, due to the process of classification with Naive Bayes being simpler, we tested different seeds to the process of selecting the set of offensive messages. In the LSVM classifier, we also tested different seeds for randomising the evaluation order of messages. Thus, we have tested 50 seeds ranging from 0 to 49 to each random describe process.

Table III shows the results to tests using Linear SVM Classifier. During each of 50 seeds, we have collected values of accuracy, precision, recall, and F1-Score. Three values are calculated to each mentioned metric based on all executions: The best value found to all executions that can be visualised in the second column, the average value of executions (third column), and the standard derivation of executions (fourth column). The seed with the best value of accuracy and precision simultaneously is highlighted soon after.

The seed with the best result went to the seed 6 and the results for this are 90% of accuracy, 88% of precision, 92% of recall and an F1-score of 90%.

TABLE III  
RESULTS OF LINEAR SVM CLASSIFIER WITH AVERAGE (AV) AND STANDARD DEVIATIONS (SD)

Metric	Best Value	AV	SD
Accuracy	0.900	0.523	0.136
Precision	0.883	0.272	0.260
Recall	1.0	0.300	0.442
F1-Score	0.902	0.242	0.345

TABLE IV  
RESULTS OF NAIVE BAYES CLASSIFIER WITH AVERAGE (AV) AND STANDARD DEVIATIONS (SD)

Metric	Best Value	AV	SD
Accuracy	0.922	0.912	0.004
Precision	0.899	0.882	0.006
Recall	0.964	0.950	0.005
F1-Score	0.924	0.915	0.004

Table IV presents the results to the tests using the Naive Bayes classifier. Such as the Linear SVM classifier, we have presented the best and average values of accuracy, precision, recall, and F1-Score to all executions, besides the standard derivation. The table distribution is similar to the SVM table. The seed with the best result went to the 8 and the values are 92% of accuracy, 89% of precision, 95% of recall and an F1-score of 92%.

We initially performed a statistic test (t-test) to decide if the generated values for each algorithm have had a significant statistical difference between them. For each metric, we would use all the generated values by both algorithms and apply the t-test with the one-tailed hypothesis and a significance level of 0.05. The results of the t-test can be seen in Table V. Based on the results, all calculated metrics have values with significant statistical differences between the algorithms. Therefore, it is concluded that the Naive Bayes classifier has a better performance to detect offensive language with an accuracy of 92% and a recall of 95% for the best dataset configuration.

TABLE V  
T-TEST BETWEEN NAIVE BAYES AND SVM CLASSIFIERS

Metric	T-Value	P-Value	Result
Accuracy	19.901	< 0.00001	Significant
Precision	16.383	< 0.00001	Significant
Recall	10.298	< 0.00001	Significant
F1-Score	13.652	< 0.00001	Significant

Comparing the Naive Bayes performance with the researched literature, our results are better than or equal to all works except an RNN configuration presented in [13] that obtained 94% of precision and recall of 96% to detect offensive language. Thus, we have a simpler and cheaper implementation of Linear SVM and Naive Bayes that obtained better results than a related work that used the same dataset to the same goal.

For future works, we aim to implement other text classification algorithms and evaluate the performance in comparison with our Linear SVM and Naive Bayes. An architecture to real-time tweet classification can be developed with these algorithms, mainly the Naive Bayes which obtained good results and is very fast as well as it has the simplicity of probabilistic-based analysis for classification. And finally, we are already planning to explore the use of these algorithms for classifying other types of text.

## VI. CONCLUSION

In this work, we implement a Linear SVM and Naive Bayes classifiers to detect the offensive language in tweets. During the test was noted that the Linear SVM is very sensitive to data type used in the training process. It was detected that the data normalisation with tags made it difficult for the process of parameter regulation. The tests also showed that the evaluation order of messages strongly influences the final result of the classifier, this is noted due to the high standard derivation for the tests with different seeds. This is a normal process because if big sequences of messages with the same label are given as input, for example, the weight regulation and the learning coefficient ( $\alpha$ ) make the learning arranged by the other inputs causing an imbalance of the weights. Thus, the Linear SVM needs a balanced input to obtained good results. The setting of the parameters for this algorithm proved to be a bit tricky task.

In contrast, the Naive Bayes classifier proved to be a good text classifier. One of your positive points is your simplicity and easiness of implementation that makes this algorithm very fast. The performance of this algorithm is really good, where showed be better of many techniques demonstrated in the literature. The problem found with the SVM was not present in this classifier, the standard derivation obtained in the tests is relatively low and the average value is very close to the best result.

## REFERENCES

- [1] F. Del-Vigna, A. Cimino, F. Dell-Orletta, M. Petrocchi, and M. Tesconi, "Hate me, hate me not: Hate speech detection on facebook," in *First Italian Conference on Cybersecurity*, 2017.
- [2] J. Jacobs and K. Potter, *Hate crimes: Criminal law & identity politics*. Oxford University Press on Demand, 1998.
- [3] M. Bouazizi and T. Ohtsuki, "Multi-class sentiment analysis on twitter: Classification performance and challenges," *Big Data Mining and Analytics*, vol. 2, no. 3, pp. 181–194, Sep. 2019.
- [4] G. Jalaja and C. Kavitha, *Sentiment Analysis for Text Extracted from Twitter*. Singapore: Springer Singapore, 2019, pp. 693–700.
- [5] S. Sharma and A. Jain, "Cyber social media analytics and issues: A pragmatic approach for twitter sentiment analysis," in *Advances in Computer Communication and Computational Sciences*, S. K. Bhatia, S. Tiwari, K. K. Mishra, and M. C. Trivedi, Eds. Singapore: Springer Singapore, 2019, pp. 473–484.
- [6] M. Zampieri, S. Malmasi, P. Nakov, S. Rosenthal, N. Farra, and R. Kumar, "Identifying and categorizing offensive language in social media (offenseval)," *arXiv preprint arXiv:1903.08983*, 2019.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [8] P. Liu, W. Li, and L. Zou, "Transfer learning for offensive language detection using bidirectional transformers," in *Proceedings of the 13th International Workshop on Semantic Evaluation*, 2019, pp. 87–91.
- [9] J. Han, S. Wu, and X. Liu, "Identifying and categorizing offensive language in social media," in *Proceedings of the 13th International Workshop on Semantic Evaluation*, 2019, pp. 652–656.
- [10] A. Nikolov and V. Radivchev, "Offensive tweet classification with bert and ensembles," in *Proceedings of the 13th International Workshop on Semantic Evaluation*, 2019, pp. 691–695.
- [11] J. Risch, A. Stoll, M. Ziegele, and R. Krestel, "Offensive language identification using a german bert model."
- [12] T. Davidson, D. Warmesley, M. Macy, and I. Weber, "Automated hate speech detection and the problem of offensive language," in *Proceedings of the 11th International AAAI Conference on Weblogs and Social Media*, ser. ICWSM '17, 2017.
- [13] G. Pitsilis, H. Ramampiaro, and H. Langseth, "Detecting offensive language in tweets using deep learning," *arXiv preprint arXiv:1801.04433*, 2018.
- [14] Z. Mossie and J.-H. Wang, "Vulnerable community identification using hate speech detection on social media," *Information Processing & Management*, p. 102087, 2019.
- [15] P. Mathur, R. Shah, R. Sawhney, and D. Mahata, "Detecting offensive tweets in hindi-english code-switched language," in *Proceedings of the Sixth International Workshop on Natural Language Processing for Social Media*, 2018, pp. 18–26.
- [16] M. Zampieri, S. Malmasi, P. Nakov, S. Rosenthal, N. Farra, and R. Kumar, "Predicting the type and target of offensive posts in social media," *arXiv preprint arXiv:1902.09666*, 2019.
- [17] Z. Waseem and D. Hovy, "Hateful symbols or hateful people? predictive features for hate speech detection on twitter," in *Proceedings of the NAACL student research workshop*, 2016, pp. 88–93.
- [18] H. Watanabe, M. Bouazizi, and T. Ohtsuki, "Hate speech on twitter: A pragmatic approach to collect hateful and offensive expressions and perform hate speech detection," *IEEE Access*, vol. 6, pp. 13 825–13 835, 2018.
- [19] S. Kotsiantis, D. Kanellopoulos, and P. Pintelas, "Handling imbalanced datasets: A review," *GESTS International Transactions on Computer Science and Engineering*, vol. 30, no. 1, pp. 25–36, 2006.
- [20] J. Wilbur and K. Sirotkin, "The automatic identification of stop words," *Journal of information science*, vol. 18, no. 1, pp. 45–55, 1992.
- [21] I. Rish, "An empirical study of the naive bayes classifier," in *IJCAI 2001 workshop on empirical methods in artificial intelligence*, vol. 3, no. 22, 2001, pp. 41–46.
- [22] Y.-W. Chang and C.-J. Lin, "Feature ranking using linear svm," in *Causation and Prediction Challenge*, 2008, pp. 53–64.
- [23] G. Forman, "Bns feature scaling: an improved representation over tf-idf for svm text classification," in *Proceedings of the 17th ACM conference on Information and knowledge management*. ACM, 2008, pp. 263–270.