



INTERNATIONAL JOURNAL OF ADVANCE RESEARCH, IDEAS AND INNOVATIONS IN TECHNOLOGY

ISSN: 2454-132X

Impact factor: 4.295

(Volume 5, Issue 2)

Available online at: www.ijariit.com

A new mechanism on hate speech detection with hateful and offensive expressions on Twitter using various machine learning techniques

Md. Ah Hassan Rayon Hussain

fammurockzz@gmail.com

SRM Institute of Science and Technology, Chennai, Tamil Nadu

ABSTRACT

A lethal online substance has turned into a noteworthy issue in this day and age because of an exponential increment in the utilization of the web by individuals of various societies and instructive foundation. Separating hate speech and offensive language is a key test in the programmed detection of dangerous content substance. In this paper, we propose a way to deal with naturally order tweets on Twitter into three classes: hateful, offensive and clean. Utilizing Twitter dataset, in this paper, we propose a way to deal with distinguish hate expressions on Twitter. Our methodology depends on unigrams and examples that are consequently collected from the preparation set. These examples and unigrams are later utilized, among others, as highlights to prepare a machine learning calculation. Our analyses on a test set made out of 2010 tweets demonstrate that our methodology achieves an exactness equivalent to 87.4% on identifying whether a tweet is offensive or not (twofold classification), and precision equivalent to 78.4% on distinguishing whether a tweet is hateful, offensive or clean (ternary classification).

Keywords— Twitter, Hate speech, Machine learning, Sentiment analysis

1. INTRODUCTION

In the previous 10 years, we have seen exponential development in the number of individuals utilizing on the web gatherings and interpersonal organizations. Like clockwork, there are 510,000 remarks produced on Facebook [1] and around 350,000 tweets created on Twitter [2]. The individuals collaborating on these gatherings or informal organizations originate from various societies and instructive foundations. On occasion, the distinction in feelings prompts verbal strikes. In addition, unchecked the right to speak freely over the web and the veil of obscurity that the web gives instigate individuals to utilize racists slurs or defamatory terms. This can bring down the confidence of individuals, promoting psychological sickness and a negative effect on society all in all. Besides, lethal language can take different structures, for example, cyberbullying, which was one of the significant explanations for suicide [3]. This issue has appeared to be progressively vital in the most recent decade and identifying or

expelling such substance physically from the web is a monotonous undertaking. So there is a requirement for formulating a mechanized model that can identify such lethal substance on the web.

So as to handle this issue, right off the bat, we should most likely characterize lethal language. We comprehensively isolate lethal language into two classifications: hate speech and offensive language. A comparative methodology was utilized in the examinations [4] and [5]. As per Wikipedia, hate speech is characterized as "any speech that assaults an individual or gathering based on qualities, for example, race, religion, ethnic starting point, national beginning, sex, handicap, sexual introduction, or sex personality." We characterize offensive language as the content which utilizes harsh slurs or deprecatory terms.

To conquer this commotion and the non-unwavering quality of data, we propose in this work a proficient method to recognize both offensive posts and hate speeches on Twitter. Our methodology depends on composing designs, and unigrams alongside sentimental highlights to play out the detection. The rest of this paper is organized as pursues: in Section 2 we present our inspirations and portray a portion of the related work. In Section 3 we formally characterize the point of our work and depict in detail our proposed strategy for hate speech detection and how includes are extricated.

2. RELATED WORK

The analysis of abstract language on OSN has been profoundly considered and connected on various fields differing from sentiment analysis [10] [11] [12] to mockery detection [6] [7] or detection of gossipy tidbits [13] and so forth. In any case, moderately fewer works (contrasted with the previously mentioned subjects) have been routed to the hate speech detection. A portion of these works focused on sentences on the internet, for example, crafted by Warner et al. [5] and Djuric et al. [14]. The primary work achieved a precision of classification equivalent to 94% with an F1 score equivalent to 63.75% in the assignment of parallel classification, and the second achieved an exactness equivalent to 80%.

Gitari et al. [15] extricated sentences from some real "hate

destinations" in the United States. They commented on every one of the sentences into one of three classes: "emphatically hateful (SH)", "pitifully hateful (WH)", and "non-hateful (NH)". They utilized semantic highlights and syntactic examples highlights, run the classification on a test set and got an F1-score equivalent to 65.12%.

Nobata et al. [16] utilized dictionary highlights, n-gram highlights, phonetic highlights, syntactic highlights, prepared highlights, "word2vec" highlights and "comment2vec" highlights to play out the classification undertaking into two classes, and acquired a precision equivalent to 90%.

By the by, some different works focused on the detection of hateful sentences on Twitter. Kwok et al. [17] focused on the detection of hateful tweets against dark individuals. They utilized unigram highlights which gave a precision equivalent to 76% for the assignment of twofold classification. Clearly, the attention on the hate speech toward a particular sexual orientation, ethnic gathering, race or different makes the collected unigrams identified with that particular gathering. In this manner, the assembled word reference of unigrams can't be reused to recognize hate speech towards different gatherings with similar proficiency. Burnap et al. [3] utilized composed conditions (that is., the connection between words) alongside the sack of words (BoW) highlights to recognize hate speech expressions from clean speech ones.

3. PROPOSAL METHODOLOGY

In the event that you are utilizing Word, Given a lot of Tweets, the point of this work is to characterize every one of them into one of three classes which are:

- **Clean:** this class comprises of tweets which are impartial, non-offensive and present no hate speech.
- **Offensive:** this class contains tweets that are offensive, however, don't present any hate or segregative/bigot speeches
- **Hateful:** this class incorporates tweets which are offensive, and present hate, supremacist and segregative words, and expressions.

We use machine learning to play out the classification: we separate a lot of highlights from each tweet, we allude to a preparation set and play out the classification.

3.1 Data

For this work, we have collected and consolidated 3 different datasets:

A first data set freely accessible on Crowdfunder2: this data set contains in excess of 14000 tweets that have been physically characterized into one of the accompanying classes: "Hateful", "Offensive" and "Clean". Every one of the tweets on this data set has been physically clarified by three individuals.

A second data set openly accessible additionally on Crowdfunder3: which has been utilized beforehand in [19] and which has likewise been physically explained into one of the three classes: "Hateful", "Offensive" and "Not one or the other", the last alluding to the "Spotless" class referenced already.

A third data set, which has been distributed in github4 and utilized in the work [18]: Tweets on this data set are arranged into one of the accompanying three classes: "Sexism", "Prejudice" and "Not one or the other". The initial two ("Sexism", "Bigotry") alluding to explicit types of hate speech, they have been incorporated as a piece of the class "Hateful", while the tweets of the class "Not one or the other" have been disposed of in light of the fact that there is no sign whether

they are spotless or offensive (a few tweets were physically checked, and they have been distinguished as having a place with the two classes).

As expressed over, the three data sets were consolidated to make a greater data set that we split as we will depict later in this area. To play out the assignment of classification, the data set is part into three subsets as pursues:

- **A preparation set:** this set contains 21 000 tweets, circulated equally among the three classes (that is., "Clean", "Offensive" and "Hateful"): each class has 7 000 tweets. This set will be alluded to as the "preparation set" in whatever is left of this work.
- **A test set: this set contains 2010 tweets:** each class has 670 tweets. This set will be alluded to as the "test set" and will be utilized to upgrade our proposed methodology.
- **An approval set: this set contains 2010 tweets:** each class has 670 tweets. This set will be alluded to as the "approval set" and will be utilized to assess our proposed methodology.

To get a reasonable outcome, we utilize a similar number of tweets for each set. Given that the number of tweets in "Hateful" class was 8 340 and it is the least among the three classes, we set the quantity of preparing tweets for each class to 7 000 tweets, that of the test tweets to 670 tweets and that of the approval tweets to 670.

3.2 Data Pre-Processing

In this segment, we quickly depict how the tweets were preprocessed figure 1 demonstrates the diverse advances done amid this stage.

In an initial step, we tidy up the tweets. This incorporates the evacuation of URLs (which beginning either with "HTTP://" or "https ://") and labels (that is., "@user") and insignificant expressions (words written in dialects that are not bolstered by ANSI coding). This is on the grounds that these don't include any data whether the tweet may express hate or not. Specifically, for the instance of labels, if the connection between the creator of the tweet and the individual labeled is known, this data may be profitable. In any case, since no foundation is given in regards to the creator and the labeled individual, we trust that the utilization of labels isn't helpful for our work.

The second step comprises of the tokenization, Part-of-Speech (PoS) Tagging, and the lemmatization (utilizing both tokens and PoS labels) of various words. For this purpose, we utilized OpenNLP5 to play out the Natural Language Processing (NLP) errands of tokenization and lemmatization. Nonetheless, to play out the Part-of-Speech (PoS) labeling, we depend on Gate Twitter PoS Tagger [20]. This is on the grounds that OpenNLP presents poor exhibitions on PoS labeling of casual and uproarious messages, for example, tweets.

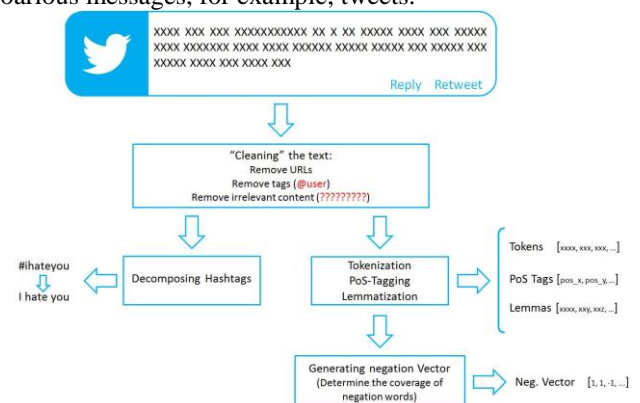


Fig. 1: Pre-processing phases of the tweets

Subsequently, we produce what we qualify as refutation vector: we identify the situation of invalidation words (e.g., "not", "never", and so forth.) and distinguish the inclusion of these words. The methodology we utilized is very straightforward and motivated from crafted by Das et al. [21]: essentially, an invalidation word covers every one of the words that tail it until the following accentuation mark or the event of a differentiation word (e.g., "yet", "notwithstanding", and so forth). Words secured by a nullification word are given an invalidation score equivalent to 1 while whatever remains of the words will be given a score equivalent to 1. This will be utilized later on the check of positive and negative words: a positive word (negative word) having a refutation score equivalent to 1 will be considered as a negative word (positive word), and it is credited the inverse of its unique score (This will be clarified in the following subsection).

On a different advance, we extricate all the hashtags, and utilize a little apparatus we formed to deteriorate it into the words that make it (e.g., the hashtag "#ihateyou" will give the articulation "I hate you") and are kept aside to be utilized when required.

3.3 Features extraction

In this subsection, we depict how includes are extricated from the tweets, and which we will utilize later to play out the classification. Be that as it may, we initially clarify the decision of our arrangements of highlights. Hate is essentially a sentiment among others, a negative sentiment to be exact. Subsequently, we trust that depending on sentiment extremity of the tweet is a critical marker of regardless of whether it tends to be a potential hateful tweet. What's more, accentuation checks and utilization of all-uppercase words can essentially change the significance of the tweet, or make unequivocal some goal covered up in content. Subsequently, such highlights should be removed alongside sentiment highlights to distinguish hate.

Nonetheless, hate shows fundamentally on the words and expressions an individual employee. Hence, the substance of the words itself is significantly more critical than the previously mentioned highlights. For this, we extricate from the preparation set, practically, a lot of words (to which we allude as unigrams) and expressions (to which we allude as examples), that are well on the way to be identified with hate and use them as additional highlights for hate detection.

As clarified at an opportune time this work (Section 2.1), in contrast to sentiment analysis, it isn't extremely helpful to depend just on the sentiment extremity of the words to recognize hate speech: not exclusively do the words' implications change as indicated by the specific situation yet in addition hate speech has distinctive appearances. Examples, in such cases, are valuable to recognize longer hateful articulation. In this way, we extricate designs alluding to words, just as a feature of speech labels, to ensure that we don't get select examples that apply to without a doubt, unmistakable circumstances, yet broad ones that reflect hate paying little mind to the substance. As it were, we ensure that an articulation extricated that demonstrates hate, is a general one that applies to various settings of hate. This will be explained later on this work, when we set a few parameters to ensure that a specific articulation happens enough occasions in a given class (that is, it isn't explicit to a solitary case or situation) and does not happen in alternate classes (that is., it's anything but a general articulation that has nothing to do with that class).

To close, basically, 4 sets of highlights are removed which we qualify as "sentiment-based highlights", "semantic highlights", "unigram highlights", and "example highlights". By consolidating these sets, we trust it is conceivable to distinguish hate speech: "sentiment highlights" enable us to extricate the extremity of the tweet, an exceptionally basic segment of hate speech (given that hateful speeches are for the most part negative ones). "Semantic highlights" enable us to locate any stressed articulation. "Unigram highlights" enable us to distinguish any express type of hate speech, while designs permit the ID of any more or certain types of hate speech. In whatever remains of this subsection, we depict how these highlights are separated.

3.3.1 Sentiment-based features: Despite the fact that the assignment of detection of hate speech varies radically from that of sentiment analysis and extremity detection, regardless it bodes well to utilize sentiment-based highlights as the most fundamental highlights that permit the detection of hate speech. This is on the grounds that hate speech is well on the way to be available in a "negative" tweet, as opposed to a "positive" one. Thusly, we first concentrate includes that would decide if a tweet is sure, negative or impartial. As referenced over, the detection of the extremity in itself isn't the motivation behind this work, yet an additional progression to encourage the principle assignment which is the detection of hate speech.

In this way, from each tweet t , we extricate the accompanying highlights: the all-out score of positive words (PW), the complete score of negative words (NW), and the proportion of passionate (positive and negative) words (t) Defined as:

$$(t) = \frac{PW}{PW+NW}; (t) \text{ is set to } 0 \text{ if the}$$

$$PW+NW$$

The tweet has no emotional words,

the number of positive slang words, the number of negative slang words, the number of positive feelings, the number of negative feelings, the number of positive hashtags, the number of negative hashtags.

The all-out score of positive words and that of negative words are extricated utilizing SentiStrength6, an apparatus that attributes sentiment scores to sentences just as the expressions of which it is formed. The scores go from - 5 to - 1 for negative words, and from 1 to 5 for positive words. Given a tweet t , we check the aggregate of the scores of individual words that have a positive extremity and credit the acquired total to the main highlights; and we do likewise for the negative words and trait the total estimation of the got whole to the second highlights (that is., the two highlights take positive qualities).

To distinguish the extremity of emoji's and slang words, we depend on two physically manufactured lexicons containing the emoji's/slang words alongside their extremity. With respect to Hashtags, we built up our very own instrument that parts a hashtag into the words that create it and utilized SentiStrength scores to settle on its extremity.

Sentiment-related highlights are great pointers of regardless of whether content is negative. As referenced over, negative content is destined to introduce hate speech. In any case, not every single negative content do. Along these lines, more highlights should be removed for detection of hate speech.

On the off chance that it has fewer words than L, it is essentially disposed of.

We extricate diverse examples as portrayed from the preparation set and spare them in three distinct records alongside their number of events. We sift through the ones that seem not exactly minpocc. Thereafter, given an example p that showed up in one of the three records (we call it C1), we measure two proportions we allude to as 12 and 13 characterized as pursues:

$$\rho_{12}(p) = \frac{N_1(p)}{N_2(p)} \tag{4}$$

$$\rho_{13}(p) = \frac{N_1(p)}{N_3(p)} \tag{5}$$

Where Ni(p) is the number of events of the example pin class I. In the event that the denominator of the proportion is 0, the esteem is set to 2.

Table 1: List of PoS tags and their corresponding simplified tags

PoS Tags	Simplified Tags
NN, NNS, NNP, NNPS	NOUN
VB, VBD, VBG, VBN, VBP, VBZ	VERB
RB, RBR, RBS	ADVERB
JJ, JJR, JJS	ADJECTIVE
CC	COORDCONJUNCTION
CD	CARDINAL
DT	DETERMINER
EX	EXISTTHERE
FW	FOREIGNWORD
IN	PREPOSITION
LS	LISTMARKER
MD	MODAL
PDT	PREDETERMINER
POS	POSSESSIVEEND
PRP, PRP\$	PRONOUN
RP	PARTICLE
SYM	SYMBOL
TO	TO
UH	INTERJECTION
WDT, WP, WP\$, WRB	WHDETERMINER
Punctuation marks	.

Just examples that fulfil the condition

$$\rho_{ij}(p) \geq Th_p \tag{6}$$

are kept, where Thp is a limit we characterize and tune. Utilizing the ideal estimations of the two parameters minpocc and Thp, 1875 examples highlights are extricated altogether. Given an example p, the relating highlight is credited to numeric esteem estimating the likeness of the tweet to that design. Accordingly, given a tweet t and an example p, we characterize the accompanying likeness work [6]:

$$res(p, t) = \begin{cases} 1, & \text{if the pattern appears in the tweets as it is,} \\ \alpha \cdot n/N, & \text{if the tweet contains } n \text{ out of the } N \text{ tags of the pattern in the correct order,} \\ 0, & \text{if the tweet doesn't contain any of the tags of the pattern.} \end{cases}$$

3.4 Parameters optimization

The proposed sets of highlights present diverse parameters that should be enhanced to get the most extreme exactness of classification. The parameters to be upgraded are the accompanying: the insignificant event of words minute the word proportions edge Th_u

- the negligible event of examples minpocc the example proportions edge Th_p
- the design length L the coefficient

To tune these parameters, each time we fix every one of the parameters with the exception of one and search for its ideal esteem. In this manner, to decide the best estimation of the parameter minuocc, we set the estimations of the diverse parameters as pursues as follows:

$$Th_u = Th_p = 1.4,$$

$$min_{pocc}^p = 3,$$

$$L=7$$

$$\alpha=.1$$

The decision of these qualities depended on a before setting of a trial in which we attempted to restrain the interims of the estimations of the parameters: we ran our investigations on every group of highlights freely utilizing the estimations of comparative parameters that we presented in past work [6]. At that point, we balanced the highlights to get the present qualities.

We attempt distinctive estimations of the parameter minuocc. The outcomes are given in figure 2. The ideal esteem was gotten for minuocc = 9.




Fig. 3: Classification accuracy (right axis) and number of words collected (left axis) for different values of the parameter minuocc

We at that point keep the estimations of the distinctive parameters as they are, set minuocc to 9, and change the parameter Th_u. Distinctive qualities from 1.1 to 2 have been checked, and figure. 6: Classification exactness (right hub) and the number of examples collected (left pivot) for various estimations of the parameter L ideal esteem was acquired for Th_u = 1:4 as appeared in figure 3 altogether, 1 373 words are collected.



Fig. 4: Classification accuracy (right axis) and number of words collected (left axis) for different values of the parameter Th_u

To decide the best length of examples (that is., L), we set the estimations of the parameters identified with unigram highlights to their ideal qualities and attempt diverse estimations of the parameter L as appeared in figure. 4.




Fig. 5: Classification accuracy (right axis) and the number of patterns collected (left axis) for different values of the parameter L

We kept alternate parameters as we set them at first. The ideal esteem was gotten for L = 5, and the complete number of examples acquired is 1875.

We continue a similar method to get the ideal estimations of minpocc and T hp. The ideal estimations of the parameters are 7 and 1:3 1:9 (in whatever is left of this work the esteem 1.4 is considered) separately as appeared in figures. 5 and 6.




Fig. 6: Classification accuracy (right axis) and number of patterns collected (left axis) for different values of the parameter minpocc




Fig. 7: Classification accuracy (right axis) and number of patterns collected (left axis) for different values of the parameter Thp

We at last set the estimations of the four parameters to their ideal and attempted distinctive estimations of. The got outcomes figure. 8: Classification exactness (right pivot) and number of examples collected (left hub) for various estimations of the parameter Thp.

Table 2: Accuracy, Precision, Recall and F1-Score of classification using different classifiers

	TP Rate	FP Rate	Prec.	Recall	F1-Score
Rand. Forest	0.592	0.204	0.605	0.592	0.589
SVM	0.57	0.276	0.643	0.57	0.599
J48graft	0.784	0.108	0.793	0.784	0.784

$$\begin{cases} \min_{occ}^u = 9, \\ Th_u = 1.4, \\ \min_{occ}^p = 7, \\ Th_p = 1.4, \\ L = 5, \\ \alpha = 0.01. \end{cases}$$

did not very much (remembering ought to have low esteem). The ideal estimation of this parameter is equivalent to 0.01. Hence, for whatever is left of this work, we considered the main case and keep the estimations of the parameters as pursues:

4. CONCLUSION

In this work, we proposed another technique to recognize hate speech on Twitter. Our proposed methodology consequently distinguishes hate speech designs and most basic unigrams and utilize these alongside sentimental and semantic highlights to order tweets into hateful, offensive and clean. Our proposed methodology achieves a precision equivalent to 87.4% for the double classification of tweets into offensive and non-offensive, and exactness equivalent to 78.4% for the ternary classification of tweets into, hateful, offensive and clean. In future work, we will endeavour to fabricate a more extravagant word reference of hate speech designs that can be utilized, alongside a unigram lexicon, to recognize hateful and offensive online writings. We will make a quantitative investigation of the nearness of hate speech among the diverse sexual orientations, age gatherings, and areas, and so forth.

5. REFERENCES

- [1] R.D. King and G.M. Sutton, "High Times for Hate Crime: Explaining the Temporal Clustering of Hate Motivated Offending", in *Criminology* pp. 871–894, 2013.
- [2] Peter J. Breckheimer, "A Haven for Hate: The Foreign and Domestic Implications of Protecting Internet Hate Speech Under the First Amendment," in *South California Law Review*, vol. 75, no. 6, Sep. 2002.
- [3] P. Burnap, and M. L. Williams, "Cyber Hate Speech on Twitter: An Application of Machine Classification and Statistical Modeling for Policy and Decision Making," in *Policy and Internet* pp. 223–242, June 2015.
- [4] A. H. Razavi, D. Inkpen, S. Uritsky, S. Matwin, "Offensive Language Detection Using Multi-level Classification," *Advances in Artificial Intelligence*, vol. 6085, pp. 16–27, Springer, Ottawa, Canada, June 2010
- [5] W. Warner and J. Hirschberg "Detecting hate speech on the World Wide Web," in *Proc. Second Workshop Language Social Media*, pp. 19–26, June 2012.
- [6] M. Bouazizi and T. Ohtsuki, "A pattern-based approach for sarcasm detection on Twitter," *IEEE Access*, Vol. 4, pp. 5477–5488, 2016.
- [7] D. Davidov, O. Tsur, and A. Rappoport, "Semi-supervised recognition of sarcastic sentences in Twitter and Amazon," In *Proc.14th Conf. on Computational Natural Language Learning*, pp. 107–116, July 2010.
- [8] M. Bouazizi and T. Ohtsuki, "Sentiment Analysis: from Binary to Multi-Class Classification - A Pattern-Based Approach for Multi-Class Sentiment Analysis in Twitter," in *Proc. IEEE ICC*, pp. 1–6, May 2016.
- [9] M. Bouazizi and T. Ohtsuki, "Sentiment analysis in Twitter: from classification to quantification of sentiments within tweets," *IEEE Globecom*, Dec. 2016, to be published.

- [10] J. M. Soler, F. Cuartero, and M. Roblizo, "Twitter as a tool for predicting elections results," in Proc. IEEE/ACM ASONAM, pp. 1194–1200, Aug. 2012.
- [11] S. Homoceanu, M. Foster, C. Lofi, and W-T. Balke, "Will I like it? Providing product overviews based on opinion excerpts," in Proc. IEEE CEC, pp. 26–33, Sept. 2011.
- [12] U. R. Hodeghatta, "Sentiment analysis of Hollywood movies on Twitter," in Proc. IEEE/ACM ASONAM, pp. 1401–1404, Aug. 2013.
- [13] Z. Zhao, P. Resnick and Q. Mei, "Enquiring Minds: Early Detection of Rumors in Social Media from Enquiry Posts," in Proc. Int. Conf. World Wide Web, pp. 1395–1405, May 2015.
- [14] N. Djuric, J. Zhou, R. Morris, M. Grbovic, V. Radosavljevic, and N. Bhamidipati, "Hate Speech Detection with Comment Embed-dings," in Proc. WWW'15 Companion, pp. 29–30, May 2015.
- [15] Njagi Dennis Gitari, Z. Zuping, Hanyurwimfura Damien, and Jun Long, "A Lexicon-based Approach for Hate Speech Detection," in, pp., Apr. 2015.
- [16] Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang, "Abusive Language Detection in Online User Content," in Proc. WWW'16, pp. 145–153, Apr. 2016.
- [17] I. Kwok, and Y. Wang, "Locate the Hate - Detecting Tweets against Blacks," in Proc. AAAI'13, pp. 1621–1622, July 2013.
- [18] Z. Waseem and D. Hovy, "Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter," in Proc. NAACL'16 Student Research Workshop, pp. 88–93, June. 2016.
- [19] T. Davidson, D. Warmley, M. Macy, and I. Weber "Automated Hate Speech Detection and the Problem of Offensive Language," in Proc. ICWSM'17, May. 2017.
- [20] L. Derczynski, A. Ritter, S. Clark, and K. Bontcheva, "Twitter part-of-speech tagging for all: Overcoming sparse and noisy data," in Proc. Int. Conf. RANLP, pp. 198–206, Sept. 2013
- [21] S. Das and M. Chen, "Yahoo! for Amazon: Extracting market sentiment from stock message boards," in Proc. 8th Asia Pacific Finance Assoc. Annu. Conf., vol. 35, pp. 43, July 2001.
- [22] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten "The WEKA data mining software: An update', SIGKDD Explor. Newsk. vol. 11, no. 1, pp. 10–18, June 2009.
- [23] G. I. Webb, "Decision Tree Grafting from the All-tests-but-one Partition," in Proc. IJCAI'99, pp. 702–707, CA, USA, Aug. 1999.