

Detecting Hate Speech in Social Media

Shervin Malmasi

Harvard Medical School
Boston, MA, United States
smalmasi@bwh.harvard.edu

Marcos Zampieri

University of Wolverhampton
Wolverhampton, United Kingdom
marcos.zampieri@uni-koeln.de

Abstract

In this paper we examine methods to detect hate speech in social media, while distinguishing this from general profanity. We aim to establish lexical baselines for this task by applying supervised classification methods using a recently released dataset annotated for this purpose. As features, our system uses character n -grams, word n -grams and word skip-grams. We obtain results of 78% accuracy in identifying posts across three classes. Results demonstrate that the main challenge lies in discriminating profanity and hate speech from each other. A number of directions for future work are discussed.

1 Introduction

Research on safety and security in social media has grown substantially in the last decade. A particularly relevant aspect of this work is detecting and preventing the use of various forms of abusive language in blogs, micro-blogs, and social networks. A number of recent studies have been published on this issue such as the work by Xu et al. (2012) on identifying cyber-bullying, the detection of hate speech (Burnap and Williams, 2015) which was the topic of a recent survey (Schmidt and Wiegand, 2017), and the detection of racism (Tulkens et al., 2016) in user generated content.

The growing interest in this topic within the research community is evidenced by several related studies presented in Section 2 and by two recent workshops: Text Analytics for Cybersecurity and Online Safety (TA-COS)¹ held in 2016 at LREC and Abusive Language Workshop (AWL)² held in 2017 at ACL.

¹<http://www.ta-cos.org/home>

²<https://sites.google.com/site/abusive languageworkshop2017/>

In this paper we address the problem of hate speech detection using a dataset which contains English tweets annotated with three labels: (1) hate speech (HATE); (2) offensive language but no hate speech (OFFENSIVE); and (3) no offensive content (OK). Most studies on abusive language so far (Burnap and Williams, 2015; Djuric et al., 2015; Nobata et al., 2016) have been modeled as binary classification with only one positive and one negative classes (*e.g.* hate speech vs non-hate speech). As noted by Dinakar et al. (2011), systems trained on such data often rely on the frequency of offensive or non-socially acceptable words to distinguish between the two classes. Dinakar et al. (2011) stress that in some cases “the lack of profanity or negativity [can] mislead the classifier”.

Indeed, the presence of profane content does not in itself signify hate speech. General profanity is not necessarily targeted towards an individual and may be used for stylistic purposes or emphasis. On the other hand, hate speech may denigrate or threaten an individual or a group of people without the use of any profanities.

The main aim of this paper is to establish a lexical baseline for discriminating between hate speech and profanity on this standard dataset. The corpus used here provides us with an interesting opportunity to investigate how well a system can detect hate speech from other content that is generally profane. This baseline can be used to determine the difficulty of this task, and help highlight the most challenging aspects which must be addressed in future work.

The rest of this paper is organized as follows. In Section 2 we briefly outline some previous work on abusive language detection. The data is presented in Section 3, along with a description of our computational approach, features, and evaluation methodology. Results are presented in Section 4, followed by a conclusion and future perspectives in Section 5.

2 Related Work

There have been several studies on computational methods to detect abusive language published in the last few years. One example is the work by Xu et al. (2012) who apply sentiment analysis to detect bullying in tweets and use Latent Dirichlet Allocation (LDA) topic models (Blei et al., 2003) to identify relevant topics in these texts.

A number of studies have been published on hate speech detection. As previously mentioned, to the best of our knowledge all of them rely on binary classification (e.g. hate speech vs non-hate speech). Examples of such studies include the work by Kwok and Wang (2013), Djuric et al. (2015), Burnap and Williams (2015), and by Nobata et al. (2016).

Due to the availability of suitable corpora, the overwhelming majority of studies on abusive language, including ours, have used English data. However, more recently a few studies have investigated abusive language detection in other languages. Mubarak et al. (2017) addresses abusive language detection on Arabic social media and Su et al. (2017) presents a system to detect and rephrase profanity in Chinese. Hate speech and abusive language datasets have been recently annotated for German (Ross et al., 2016) and Slovene (Fišer et al., 2017) opening avenues for future work in languages other than English.

3 Methods

Next we present the Hate Speech Detection dataset used in our experiments. We applied a linear Support Vector Machine (SVM) classifier and used three groups of features extracted for these experiments: surface n -grams, word skip-grams, and Brown clusters. The classifier and features are described in more detail in Section 3.2 and Section 3.3 respectively. Finally, Section 3.4 discusses evaluation methods.

3.1 Data

In these experiments we use the aforementioned Hate Speech Detection dataset³ distributed via CrowdFlower.⁴ The dataset features 14,509 English tweets annotated by a minimum of three annotators.

Individuals in charge of the annotation of this dataset were asked to annotate each tweet and categorize them into one of three classes:

1. (HATE): contains hate speech;
2. (OFFENSIVE): contains offensive language but no hate speech;
3. (OK): no offensive content at all.

Each instance in this dataset contains the text of a tweet⁵ along with one of the three aforementioned labels. The distribution of the texts across the three classes is shown in Table 1.

Class	Texts
HATE	2,399
OFFENSIVE	4,836
OK	7,274
Total	14,509

Table 1: The distribution of classes and tweets in the Hate Speech Detection dataset.

All the texts are preprocessed to lowercase all tokens and to remove URLs and emojis.

3.2 Classifier

We use a linear SVM to perform multi-class classification in our experiments. We use the LIBLINEAR⁶ package (Fan et al., 2008) which has been shown to be very efficient for similar text classification tasks. For example, the LIBLINEAR SVM implementation has been demonstrated to be a very effective classifier for Native Language Identification (Malmasi and Dras, 2015), temporal text classification (Zampieri et al., 2016a), and language variety identification (Zampieri et al., 2016b).

3.3 Features

We use two groups of surface features in our experiments as follows:

- Surface n -grams: These are our most basic features, consisting of character n -grams (of order 2–8) and word n -grams (of order 1–3). All tokens are lowercased before extraction of n -grams; character n -grams are extracted across word boundaries.
- Word Skip-grams: Similar to the above features, we also extract 1-, 2- and 3-skip word bigrams. These features are were chosen to approximate longer distance dependencies between words, which would be hard to capture using bigrams alone.

³<https://data.world/crowdflower/hate-speech-identification>

⁴<https://www.crowdflower.com/>

⁵ Each tweet is limited to a maximum of 140 characters.

⁶<http://www.csie.ntu.edu.tw/~7Ecjlin/liblinear/>

3.4 Evaluation

To evaluate our methods we use 10-fold cross-validation. For creating the folds, we employ stratified cross-validation aiming to ensure that the proportion of classes within each partition is equal (Kohavi, 1995).

We report our results in terms of accuracy. The results obtained by our methods are compared against a majority class baseline and an oracle classifier.

The oracle takes the predictions by all the classifiers in Table 2 into account. It assigns the correct class label for an instance if at least one of the the classifiers produces the correct label for that instance. This approach establishes the *potential* or *theoretical* upper limit performance for a given dataset. Similar analysis using oracle classifiers have been previously applied to estimate the theoretical upper bound of shared tasks datasets in Native Language Identification (Malmasi et al., 2015) and similar language and language variety identification (Goutte et al., 2016).

4 Results

We start by investigating the efficacy of our features for this task. We first train a single classifier, with each of them using a type of feature. Subsequently we also train a single model combining all of our features into single space. These are compared against the majority class baseline, as well as the oracle. The results of these experiments are listed in Table 2.

Feature	Accuracy (%)
Majority Class Baseline	50.1
Oracle	91.6
Character bigrams	73.6
Character trigrams	77.2
Character 4-grams	78.0
Character 5-grams	77.9
Character 6-grams	77.2
Character 7-grams	76.5
Character 8-grams	75.8
Word unigrams	77.5
Word bigrams	73.8
Word trigrams	67.4
1-skip Word bigrams	74.0
2-skip Word bigrams	73.8
3-skip Word bigrams	73.9
All features combined	77.5

Table 2: Classification results under 10-fold cross-validation.

The majority class baseline is quite high due to the class imbalance in the data. The oracle achieves an accuracy of 91.6%, showing that none of our features are able to correctly classify a substantial portion of our samples.

We note that character n -grams perform well here, with 4-grams achieving the best performance of all features. Word unigrams also perform well, while performance degrades with bigrams, trigrams and skip-grams. However, the skip-grams may be capturing longer distance dependencies which provide complementary information to the other feature types. In tasks relying on stylistic information, it has been shown that skip-grams capture information that is very similar to syntactic dependencies (Malmasi and Cahill, 2015, §5).

Finally, the combination of all features does not achieve the performance of a character 4-grams model and causes a large dimensionality increase, with a total of 5.5 million features. It is not clear if this model is able to correctly capture the diverse information provided by the three feature types since we include more character n -gram models than word-based ones.

Next we analyze the rate of learning for these features. A learning curve for the classifier that yielded the best performance overall, character 4-grams, is shown in Figure 1.




Figure 1: Learning curve for a character 4-gram model, with standard deviation highlighted. Accuracy does not plateau with the maximal data size.

We observe that accuracy increased continuously as the amount of training instances increased, and the standard deviation of the results between the cross-validation folds decreased. This suggests that the use of more training data is likely to provide even higher accuracy. It should be noted, however, that accuracy increases at a much slower rate after 15,000 training instances.




Figure 2: Confusion matrix of the character 4-gram model for our 3 classes. The heatmap represents the proportion of correctly classified examples in each class (this is normalized as the data distribution is imbalanced). The raw numbers are also reported within each cell. We note that the HATE class is the hardest to classify and is highly confused with the OFFENSIVE class.

Finally, we also examine a confusion matrix for the character 4-gram model, as shown in Figure 2. This demonstrates that the greatest degree of confusion lies between hate speech and generally offensive material, with hate speech more frequently being confused for offensive content. A substantial amount of offensive content is also misclassified as being non-offensive. The non-offensive class achieves the best result, with the vast majority of samples being correctly classified.

5 Conclusion

In this paper we applied text classification methods to distinguish between hate speech, profanity, and other texts. We applied standard lexical features and a linear SVM classifier to establish a baseline for this task. The best result was obtained by a character 4-gram model achieving 78% accuracy. The results presented in this paper showed that distinguishing profanity from hate speech is a very challenging task.

This was to the best of our knowledge one of the first experiments to detect hate speech on social media in a scenario including non-hate speech profanity. Previous work so far (e.g. Burnap and Williams (2015) and Djuric et al. (2015)) dealt with the distinction between hate speech and socially acceptable texts in a binary classification

setting. In binary classification, Dinakar et al. (2011) note that the frequency of offensive words helps classifiers to distinguish between hate speech and socially acceptable texts.

We see a few directions in which this work could be expanded such as the use of more robust ensemble classifiers, a linguistic analysis of the most informative features, and error analysis of the misclassified instances. These aspects are presented in more detail in the next section.

5.1 Future Work

In future work we would like to investigate the performance of classifier ensembles and meta-learning for this task. Previous work has applied these techniques to a number of comparable text classification tasks, achieving success in competitive shared tasks. Examples of recent applications include automatic triage of posts in mental health forums (Malmasi et al., 2016b), detection of lexical complexity (Malmasi et al., 2016a), Native Language Identification (Malmasi and Dras, 2017), and dialect identification (Malmasi and Zampieri, 2017).

Another direction to pursue is the careful analysis of the most informative features for each class in this dataset. Our initial exploitation of the most informative words unigrams and bigrams suggests that coarse and obscene words are very informative for both HATE and OFFENSIVE words which

confuses the classifiers. For HATE we observed a prominence of words targeting ethnic and social groups. Finally, an interesting outcome that should be investigated in more detail is that many of the most informative bigrams for the OK feature grammatical words. A more detailed analysis of these features could lead to more robust feature engineering methods.

An error analysis could also help us better understand the challenges in this task. This could be used to provide insights about the classifiers' performance as well as any underlying issues with the annotation of the Hate Speech Detection dataset which, as pointed out by Ross et al. (2016), is far from trivial. Figure 2 confirms that, as expected, most confusion occurs between HATE and OFFENSIVE texts. However, we also note that a substantial amount of offensive content is misclassified as being non-offensive. The aforementioned error analysis can provide insights about this.

Acknowledgments

We would like to thank the anonymous RANLP reviewers who provided us valuable feedback to increase the quality of this paper.

We further thank the developers and the annotators who worked on the Hate Speech Dataset for making this important resource available.

References

- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent Dirichlet Allocation. *Journal of machine Learning research* 3(Jan):993–1022.
- Pete Burnap and Matthew L Williams. 2015. Cyber hate speech on twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & Internet* 7(2):223–242.
- Karthik Dinakar, Roi Reichart, and Henry Lieberman. 2011. Modeling the detection of textual cyberbullying. In *The Social Mobile Web*. pages 11–17.
- Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of the 24th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee, pages 29–30.
- Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A Library for Large Linear Classification. *Journal of Machine Learning Research* 9:1871–1874.
- Darja Fišer, Tomaž Erjavec, and Nikola Ljubešić. 2017. Legal Framework, Dataset and Annotation Schema for Socially Unacceptable On-line Discourse Practices in Slovene. In *Proceedings of the Workshop Workshop on Abusive Language Online (ALW)*. Vancouver, Canada.
- Cyril Goutte, Serge Léger, Shervin Malmasi, and Marcos Zampieri. 2016. Discriminating Similar Languages: Evaluations and Explorations. In *Proceedings of Language Resources and Evaluation (LREC)*. Portoroz, Slovenia.
- Ron Kohavi. 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *IJCAI*. volume 14, pages 1137–1145.
- Irene Kwok and Yuzhou Wang. 2013. Locate the hate: Detecting tweets against blacks. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- Shervin Malmasi and Aoife Cahill. 2015. Measuring Feature Diversity in Native Language Identification. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*. Association for Computational Linguistics, Denver, Colorado.
- Shervin Malmasi and Mark Dras. 2015. Large-scale Native Language Identification with Cross-Corpus Evaluation. In *Proceedings of NAACL-HLT 2015*. Association for Computational Linguistics, Denver, Colorado.
- Shervin Malmasi and Mark Dras. 2017. Native Language Identification using Stacked Generalization. *arXiv preprint arXiv:1703.06541*.
- Shervin Malmasi, Mark Dras, and Marcos Zampieri. 2016a. Ltg at semeval-2016 task 11: Complex word identification with classifier ensembles. In *Proceedings of SemEval*.
- Shervin Malmasi, Joel Tetreault, and Mark Dras. 2015. Oracle and Human Baselines for Native Language Identification. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*. Association for Computational Linguistics, Denver, Colorado.
- Shervin Malmasi and Marcos Zampieri. 2017. German Dialect Identification in Interview Transcriptions. In *Proceedings of the Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*.
- Shervin Malmasi, Marcos Zampieri, and Mark Dras. 2016b. Predicting Post Severity in Mental Health Forums. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology (CLPsych)*.
- Hamdy Mubarak, Darwish Kareem, and Magdy Walid. 2017. Abusive Language Detection on Arabic Social Media. In *Proceedings of the Workshop Workshop on Abusive Language Online (ALW)*. Vancouver, Canada.

- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive Language Detection in Online User Content. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, pages 145–153.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurowsky, and Michael Wotzki. 2016. Measuring the Reliability of Hate Speech Annotations: The Case of the European Refugee Crisis. In *Proceedings of the Workshop on Natural Language Processing for Computer-Mediated Communication (NLP4CMC)*. Bochum, Germany.
- 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*. Valencia, Spain, pages 1–10.
- Huei-Po Su, Chen-Jie Huang, Hao-Tsung Chang, and Chuan-Jie Lin. 2017. Rephrasing Profanity in Chinese Text. In *Proceedings of the Workshop Workshop on Abusive Language Online (ALW)*. Vancouver, Canada.
- Stéphan Tulkens, Lisa Hilte, Elise Lodewyckx, Ben Verhoeven, and Walter Daelemans. 2016. A Dictionary-based Approach to Racism Detection in Dutch Social Media. In *Proceedings of the Workshop Text Analytics for Cybersecurity and Online Safety (TA-COS)*. Portoroz, Slovenia.
- Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. Learning from bullying traces in social media. In *Proceedings of the 2012 conference of the North American chapter of the association for computational linguistics: Human language technologies*. Association for Computational Linguistics, pages 656–666.
- Marcos Zampieri, Shervin Malmasi, and Mark Dras. 2016a. Modeling language change in historical corpora: the case of Portuguese. In *Proceedings of Language Resources and Evaluation (LREC)*. Portoroz, Slovenia.
- Marcos Zampieri, Shervin Malmasi, Octavia-Maria Sulea, and Liviu P Dinu. 2016b. A Computational Approach to the Study of Portuguese Newspapers Published in Macau. In *Proceedings of Workshop on Natural Language Processing Meets Journalism (NLPMJ)*. pages 47–51.