

HateMonitors: Language Agnostic Abuse Detection in Social Media

Punyajoy Saha¹[0000-0002-3952-2514], Binny Mathew¹[0000-0003-4853-0345],
Pawan Goyal¹[0000-0002-9414-8166], and Animesh
Mukherjee¹[0000-0003-4534-0044]

Indian Institute of Technology, Kharagpur, West Bengal, India - 721302

Abstract. Reducing hateful and offensive content in online social media pose a dual problem for the moderators. On the one hand, rigid censorship on social media cannot be imposed. On the other, the free flow of such content cannot be allowed. Hence, we require efficient abusive language detection system to detect such harmful content in social media. In this paper, we present our machine learning model, HateMonitor, developed for Hate Speech and Offensive Content Identification in Indo-European Languages (HASOC) [20], a shared task at FIRE 2019. We have used Gradient Boosting model, along with BERT and LASER embeddings, to make the system language agnostic. Our model came at **First position** for the German sub-task A. We have also made our model public ¹.

Keywords: Hate speech · Offensive language · Multilingual · LASER embeddings · BERT embeddings · Classification.

1 Introduction

In social media, abusive language denotes a text which contains any form of unacceptable language in a post or a comment. Abusive language can be divided into hate speech, offensive language and profanity. Hate speech is a derogatory comment that hurts an entire group in terms of ethnicity, race or gender. Offensive language is similar to derogatory comment, but it is targeted towards an individual. Profanity refers to any use of unacceptable language without a specific target. While profanity is the least threatening, hate speech has the most detrimental effect on the society.

Social media moderators are having a hard time in combating the rampant spread of hate speech² as it is closely related to the other forms of abusive language. The evolution of new slangs and multilingualism, further adding to the complexity.

¹ <https://github.com/punyajoy/HateMonitors-HASOC>

² <https://tinyurl.com/y6tgv865>

Copyright ©2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). FIRE 2019, 12-15 December 2019, Kolkata, India.

Recently, there has been a sharp rise in hate speech related incidents in India, the lynchings being the clear indication [3]. Arun et al. [3] suggests that hate speech in India is very complicated as people are not directly spreading hate but are spreading misinformation against a particular community. Hence, it has become imperative to study hate speech in Indian language.

For the first time, a shared task on abusive content detection has been released for Hindi language at HASOC 2019. This will fuel the hate speech and offensive language research for Indian languages. The inclusion of datasets for English and German language will give a performance comparison for detection of abusive content in high and low resource language.

In this paper, we focus on the detection of multilingual hate speech detection that are written in Hindi, English, and German and describe our submission (*HateMonitors*) for HASOC at FIRE 2019 competition. Our system concatenates two types of sentence embeddings to represent each tweet and use machine learning models for classification.

2 Related works

Analyzing abusive language in social media is a daunting task. Waseem et al. [33] categorizes abusive language into two sub-classes – hate speech and offensive language. In their analysis of abusive language, Classifying abusive language into these two subtypes is more challenging due to the correlation between offensive language and hate speech [10]. Nobata et al. [22] uses predefined language element and embeddings to train a regression model. With the introduction of better classification models [23, 29] and newer features [1, 10, 30], the research in hate and offensive speech detection has gained momentum.

Silva et al. [28] performed a large scale study to understand the target of such hate speech on two social media platforms: Twitter and Whisper. These target could be the Refugees and Immigrants [25], Jews [7, 14] and Muslims [4, 32]. People could become the target of hate speech based on Nationality [12], sex [5, 26], and gender [24, 16] as well. Public expressions of hate speech affects the devaluation of minority members [17], the exclusion of minorities from the society [21], and tend to diffuse through the network at a faster rate [19].

One of the key issues with the current state of the hate and offensive language research is that the majority of the research is dedicated to the English language on [15]. Few researchers have tried to solve the problem of abusive language in other languages [25, 27], but the works are mostly monolingual. Any online social media platform contains people of different ethnicity, which results in the spread of information in multiple languages. Hence, a robust classifier is needed, which can deal with abusive language in the multilingual domain. Several shared tasks like HASOC [20], HaSpeeDe [8], GermEval [34], AMI [13], HatEval [6] have focused on detection of abusive text in multiple languages recently.

3 Dataset and Task description

The dataset at HASOC 2019 were given in three languages: Hindi, English, and German. Dataset in Hindi and English had three subtasks each, while German had only two subtasks. We participated in all the tasks provided by the organisers and decided to develop a single model that would be language agnostic. We used the same model architecture for all the three languages.

3.1 Datasets

We present the statistics for HASOC dataset in Table 1. From the table, we can observe that the dataset for the German language is highly unbalanced, English and Hindi are more or less balanced for sub-task A. For sub-task B German dataset is balanced but others are unbalanced. For sub-task C both the datasets are highly unbalanced.

Table 1. This table shows the initial statistics about the training and test data

Language	English	German	Hindi
Sub-Task A	Train Test	Train Test	Train Test
HOF	2261 288	407 136	2469 605
NOT	3591 865	3142 714	2196 713
Total	5852 1153	3819 850	4665 1318
Sub-Task B	Train Test	Train Test	Train Test
HATE	1141 124	111 41	556 190
OFFN	451 71	210 77	676 197
PRFN	667 93	86 18	1237 218
Total	2261 288	407 136	2469 605
Sub-Task C	Train Test	Train Test	Train Test
TIN	2041 245	- - - -	1545 542
UNT	220 43	- - - -	924 63
Total	2261 288	- - - -	2469 605

3.2 Tasks

Sub-task A consists of building a binary classification model which can predict if a given piece of text is hateful and offensive (HOF) or not (NOT). A data point is annotated as HOF if it contains any form of non-acceptable language such as hate speech, aggression, profanity. Each of the three languages had this subtask.

Sub-task B consists of building a multi-class classification model which can predict the three different classes in the data points annotated as HOF: Hate

<https://hasoc2019.github.io/>

speech (HATE), Offensive language (OFFN), and Profane (PRFN). Again all three languages have this sub-task.

Sub-task C consists of building a binary classification model which can predict the type of offense: Targeted (TIN) and Untargeted (UNT). Sub-task C was not conducted for the German dataset.

4 System Description

In this section, we will explain the details about our system, which comprises of two sub-parts- feature generation and model selection. Figure 1 shows the architecture of our system.

4.1 Feature Generation

Preprocessing: We preprocess the tweets before performing the feature extraction. The following steps were followed:

- We remove all the URLs.
- Convert text to lowercase. This step was not applied to the Hindi language since Devanagari script does not have lowercase and uppercase characters.
- We did not normalize the mentions in the text as they could potentially reveal important information for the embeddings encoders.
- Any numerical figure was normalized to a string ‘number’.

We did not remove any punctuation and stop-words since the context of the sentence might get lost in such a process. Since we are using sentence embedding, it is essential to keep the context of the sentence intact.

Feature vectors: The preprocessed posts are then used to generate features for the classifier. For our model, we decided to generate two types of feature vector: BERT Embeddings and LASER Embeddings. For each post, we generate the BERT and LASER Embedding, which are then concatenated and fed as input to the final classifier.

Multilingual BERT embeddings: Bidirectional Encoder Representations from Transformers(BERT) [11] has played a key role in the advancement of natural language processing domain (NLP). BERT is a language model which is trained to predict the masked words in a sentence. To generate the sentence embedding for a post, we take the mean of the last 11 layers (out of 12) to get a sentence vector with length of 768.

LASER embeddings: Researchers at Facebook released a language agnostic sentence embeddings representations (LASER) [2], where the model jointly learns on 93 languages. The model takes the sentence as input and produces a vector representation of length 1024. The model is able to handle code mixing as well [31].

We use the BERT-base-multilingual-cased which has 104 languages, 12-layer, 768-hidden, 12-heads and 110M parameters

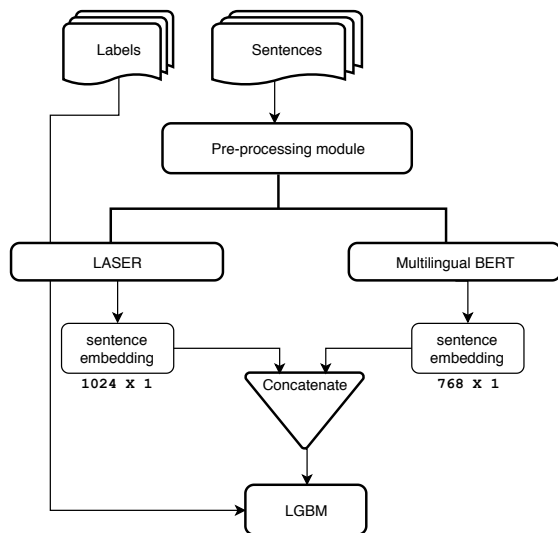


Fig. 1. Architecture of our system

We pass the preprocessed sentences through each of these embedding models and got two separate sentence representation. Further, we concatenate the embeddings into one single feature vector of length 1792, which is then passed to the final classification model.

4.2 Our Model

The amount of data in each category was insufficient to train a deep learning model. Building such deep models would lead to overfitting. So, we resorted to using simpler models such as SVM and Gradient boosted trees. Gradient boosted trees [9] are often the choice for systems where features are pre-extracted from the raw data. In the category of gradient boosted trees, Light Gradient Boosting Machine (LGBM) [18] is considered one of the most efficient in terms of memory footprint. Moreover, it has been part of winning solutions of many competition. Hence, we used LGBM as model for the downstream tasks in this competition.

5 Results

The performance of our models across different languages for sub-task A are shown in table 2. Our model got the **first position** in the German sub-task with a macro F1 score of **0.62**. The results of sub-task B and sub-task C is shown in table 3 and 4 respectively.

<https://tinyurl.com/yxmuwzla>

<https://tinyurl.com/y2g8nuuo>

Table 2. This table gives the language wise result of sub-task A by comparing the macro F1 values

Language	English	German	Hindi
HOF	0.59	0.36	0.76
NOT	0.79	0.87	0.79
Total	0.69	0.62	0.78

Table 3. This table gives the language wise result of sub-task B by comparing the macro F1 values

Language	English	German	Hindi
HATE	0.28	0.04	0.29
OFFN	0.00	0.0	0.29
PRFN	0.31	0.19	0.59
NONE	0.79	0.87	0.79
Total	0.34	0.28	0.49

Table 4. This table gives the language wise result of sub-task C by comparing the macro F1 values

Language	English	Hindi
TIN	0.51	0.63
UNT	0.11	0.17
NONE	0.79	0.79
Total	0.47	0.53

6 Discussion

In the results of subtask A, models are mainly affected by imbalance of the dataset. The training dataset of Hindi dataset was more balanced than English or German dataset. Hence, the results were around **0.78**. As the dataset in German language was highly imbalanced, the results drops to **0.62**. In subtask B, the highest F1 score reached was by the profane class for each language in table 3. The model got confused between OFFN, HATE and PRFN labels which suggests that these models are not able to capture the context in the sentence. The subtask C was again a case of imbalanced dataset as targeted(TIN) label gets the highest F1 score in table 4.

7 Conclusion

In this shared task, we experimented with zero-shot transfer learning on abusive text detection with pre-trained BERT and LASER sentence embeddings. We use an LGBM model to train the embeddings to perform downstream task. Our model for German language got the first position. The results provided a strong baseline for further research in multilingual hate speech. We have also made the models public for use by other researchers.

References

1. Alorainy, W., Burnap, P., Liu, H., Williams, M.: The enemy among us: Detecting hate speech with threats based 'othering' language embeddings. arXiv preprint arXiv:1801.07495 (2018)
2. Artetxe, M., Schwenk, H.: Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. CoRR **abs/1812.10464** (2018), <http://arxiv.org/abs/1812.10464>
3. Arun, C.: On whatsapp, rumours, and lynchings. *Economic & Political Weekly* **54**(6), 30–35 (2019)
4. Awan, I.: Islamophobia on social media: A qualitative analysis of the facebook's walls of hate. *International Journal of Cyber Criminology* **10**(1) (2016)
5. Bartlett, J., Norrie, R., Patel, S., Rumpel, R., Wibberley, S.: Misogyny on twitter. *Demos* (2014)
6. Basile, V., Bosco, C., Fersini, E., Nozza, D., Patti, V., Pardo, F.M.R., Rosso, P., Sanguinetti, M.: Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In: *Proceedings of the 13th International Workshop on Semantic Evaluation*. pp. 54–63 (2019)
7. Bilewicz, M., Winiewski, M., Kofta, M., Wójcik, A.: Harmful ideas, the structure and consequences of anti-semitic beliefs in poland. *Political Psychology* **34**(6), 821–839 (2013)
8. Bosco, C., Felice, D., Poletto, F., Sanguinetti, M., Maurizio, T.: Overview of the evalita 2018 hate speech detection task. In: *EVALITA 2018-Sixth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian*. vol. 2263, pp. 1–9. CEUR (2018)
9. Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. CoRR **abs/1603.02754** (2016), <http://arxiv.org/abs/1603.02754>
10. Davidson, T., Warmusley, D., Macy, M., Weber, I.: Automated hate speech detection and the problem of offensive language. arXiv preprint arXiv:1703.04009 (2017)
11. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. CoRR **abs/1810.04805** (2018), <http://arxiv.org/abs/1810.04805>
12. Erjavec, K., Kovačič, M.P.: “you don't understand, this is a new war!” analysis of hate speech in news web sites' comments. *Mass Communication and Society* **15**(6), 899–920 (2012)
13. Fersini, E., Nozza, D., Rosso, P.: Overview of the evalita 2018 task on automatic misogyny identification (ami). In: *EVALITA@ CLiC-it* (2018)
14. Finkelstein, J., Zannettou, S., Bradlyn, B., Blackburn, J.: A quantitative approach to understanding online antisemitism. arXiv preprint arXiv:1809.01644 (2018)
15. Fortuna, P., Nunes, S.: A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)* **51**(4), 85 (2018)
16. Gatehouse, C., Wood, M., Briggs, J., Pickles, J., Lawson, S.: Troubling vulnerability: Designing with lgbt young people's ambivalence towards hate crime reporting. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. p. 109. ACM (2018)
17. Greenberg, J., Pyszczynski, T.: The effect of an overheard ethnic slur on evaluations of the target: How to spread a social disease. *Journal of Experimental Social Psychology* **21**(1), 61–72 (1985)
18. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W.J., Ma, W., Ye, Q., Liu, T.M.: Lightgbm: A highly efficient gradient boosting decision tree. In: *NIPS* (2017)

19. Mathew, B., Dutt, R., Goyal, P., Mukherjee, A.: Spread of hate speech in online social media. In: Proceedings of the 10th ACM Conference on Web Science. pp. 173–182. ACM (2019)
20. Modha, S., Mandl, T., Majumder, P., Patel, D.: Overview of the HASOC track at FIRE 2019: Hate Speech and Offensive Content Identification in Indo-European Languages. In: Proceedings of the 11th annual meeting of the Forum for Information Retrieval Evaluation (December 2019)
21. Mullen, B., Rice, D.R.: Ethnophaulisms and exclusion: The behavioral consequences of cognitive representation of ethnic immigrant groups. *Personality and Social Psychology Bulletin* **29**(8), 1056–1067 (2003)
22. Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y., Chang, Y.: Abusive language detection in online user content. In: Proceedings of the 25th international conference on world wide web. pp. 145–153. International World Wide Web Conferences Steering Committee (2016)
23. Qian, J., ElSherief, M., Belding, E., Wang, W.Y.: Hierarchical cvae for fine-grained hate speech classification. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. pp. 3550–3559 (2018)
24. Reddy, V.: Perverts and sodomites: Homophobia as hate speech in africa. *Southern African Linguistics and Applied Language Studies* **20**(3), 163–175 (2002)
25. Ross, B., Rist, M., Carbonell, G., Cabrera, B., Kurowsky, N., Wojatzki, M.: Measuring the reliability of hate speech annotations: The case of the european refugee crisis. arXiv preprint arXiv:1701.08118 (2017)
26. Saha, P., Mathew, B., Goyal, P., Mukherjee, A.: Hateminers: Detecting hate speech against women. arXiv preprint arXiv:1812.06700 (2018)
27. Sanguinetti, M., Poletto, F., Bosco, C., Patti, V., Stranisci, M.: An italian twitter corpus of hate speech against immigrants. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018) (2018)
28. Silva, L.A., Mondal, M., Correa, D., Benevenuto, F., Weber, I.: Analyzing the targets of hate in online social media. In: ICWSM. pp. 687–690 (2016)
29. Stambach, D., Zahraei, A., Stadnikova, P., Klakow, D.: Offensive language detection with neural networks for germeval task 2018. In: 14th Conference on Natural Language Processing KONVENS 2018. p. 58 (2018)
30. Unsvåg, E.F., Gambäck, B.: The effects of user features on twitter hate speech detection. In: Proceedings of the 2nd Workshop on Abusive Language Online (ALW2). pp. 75–85 (2018)
31. Verma, S.: Code-switching: Hindi-english. *Lingua* **38**(2), 153 – 165 (1976). [https://doi.org/https://doi.org/10.1016/0024-3841\(76\)90077-2](https://doi.org/https://doi.org/10.1016/0024-3841(76)90077-2), <http://www.sciencedirect.com/science/article/pii/0024384176900772>
32. Vidgen, B., Yasseri, T.: Detecting weak and strong islamophobic hate speech on social media. arXiv preprint arXiv:1812.10400 (2018)
33. Waseem, Z., Davidson, T., Warmusley, D., Weber, I.: Understanding abuse: A typology of abusive language detection subtasks. arXiv preprint arXiv:1705.09899 (2017)
34. Wiegand, M., Siegel, M., Ruppenhofer, J.: Overview of the germeval 2018 shared task on the identification of offensive language (2018)