

Profiling Haters on Twitter using Statistical and Contextualized Embeddings

Notebook for PAN at CLEF 2021

Hamed Babaei Giglou¹, Taher Rahgooy², Jafar Razmara¹, Mostafa Rahgouy³ and Zahra Rahgooy⁴

¹*Department of Computer Science, University of Tabriz, Tabriz, Iran*

²*Department of Computer Science, University of West Florida, Florida, USA*

³*Part AI Research Center, Tehran, Iran*

⁴*Department of Computer Science, Damghan University, Semnan, Iran*

Abstract

Hate Speech (HS) in social media such as Twitter is a complex phenomenon that attracted a significant body of research in the NLP. HS Spreaders (haters) aim to spread HS via social media. In this task, we aim to identify such haters. On one hand, our proposed class-dependent LDSE representation is fed to a linear SVM classifier to identify the haters based on general commonalities. On the other hand, stylistic features of individuals are captured by using extractive summarization of the tweets in conjunction with RoBERTa embedding before classifying them using another linear SVM classifier. Experimental results expressed as accuracies 0.67 and 0.80 over English and Spanish test sets respectively show efficacy of our approach in identifying the haters across different languages.

Keywords

Hate Speech, Abusive Language, Author Profiling, Stylistic Features, Word Embedding

1. Introduction

The social medial platform enables millions to publicly share user-generated content. Regardless of different content types, a critical point of these platforms, such as Twitter, Facebook, YouTube, and Instagram, is that users can discuss content. Unfortunately, any user engaging online will always facing the risk of being targeted or harassed via abusive language, hatred expressed in the form of racism or sexism, with possible impact on his/her and the community in general. The challenge of creating effective policies to identify and appropriately respond to harassment is compounded by the difficulty of studying the phenomena at scale. Hate speech is commonly defined as any communication that disparages a person or a group on the basis of some characteristic such as race, color, ethnicity, gender, sexual orientation, nationality, religion, or other characteristics.


CLEF 2021 – Conference and Labs of the Evaluation Forum, September 21–24, 2021, Bucharest, Romania

✉ h.babaei98@ms.tabrizu.ac.ir (H. Babaei Giglou); trahgooy@students.uwf.edu (T. Rahgooy); razmara@tabrizu.ac.ir (J. Razmara); mostafa.rahgouy@partdp.ai (M. Rahgouy); za.rahgooy@gmail.com (Z. Rahgooy)

🌐 <https://hamedbabaei.github.io/> (H. Babaei Giglou); <http://www.rahgooy.com/> (T. Rahgooy)

🆔 0000-0002-2889-0522 (T. Rahgooy)

© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

To this end, in the Author Profiling task [1], we aim at identifying possible hate speech spreaders (haters) on Twitter as a first step towards preventing hate speech from being propagated among online users. Also, this task runs based on a multilingual perspective for English and Spanish languages.

The rest of the paper is organized as follows. Section 2 presents related works. Section 3 describes the proposed method. Section 4 describes the dataset, experiments and discusses the obtained results. Finally, section 5 presents our conclusions.

2. Related Works

The authors in [2] concluded that word embedding models like GloVe[3] and Word2Vec[4] are although widely used for toxic comment classification are in fact unable to handle out-of-vocabulary problem properly. However, word embeddings like FastText[5] were particularly suited for this task since it uses subword embedding. The ability to cope with unknown words is the reason why previous findings on the inferiority of word embeddings in comparison to word n-grams have become outdated. After introducing contextualized word embeddings such as BERT[6], some of the limitations of all previous word embeddings are resolved. In SemEval-2021 at the Toxic Span Detection task [7] used GloVe, GPT-2[8] and RoBERTa[9] to create empowered representation to overcome the out-of-vocabulary issue in the representation. They use the analogy that hateful words may never occur in the training time of the word embeddings like GloVe or Word2Vec, but however they are useful when they are combined with other representations since concatenations of them with LM models could create awareness signal to model.

Authors of [10] compared different deep learning and shallow approaches on a large comment dataset and propose an ensemble that outperforms all individual models. They combined different deep learning and traditional machine learning models with FastText, GloVe, word-ngrams, and char-ngrams representation. They suggest further research in representing world knowledge with embeddings to improve the distinction between paradigmatic contexts. [11] proposed character n-grams, TFIDF, BoWV(bag of word vector) that uses GloVe embedding, FastText, and random embedding representations for hate speech detection in Twitter using deep learning approaches. They investigated various combinations of SVM, LogisticRegression, CNN, LSTM, Gradient Boosted Decision Trees to handle this task. They achieved a higher F1 score with a combination of LSTM, Random Embedding, and GBDT.

At Author Profiling shared task at PAN 2020 [12], the [13] combined character and word n-grams via SVM, and they achieved best competition accuracy in Spanish.[14] also combined n-grams with stylistic features with LogisticRegression, and they achieved the best competition accuracy in English. [15] proposed a multi-feature representation that combines n-grams with word embeddings to enrich the representation for fake news spreader detection on Twitter.

3. Proposed Method

In this section, we describe the details of our proposed model. Our proposed approach aims to predict whether the user is keen to spread hate speech or not. Figure 1 depicts an overview of our

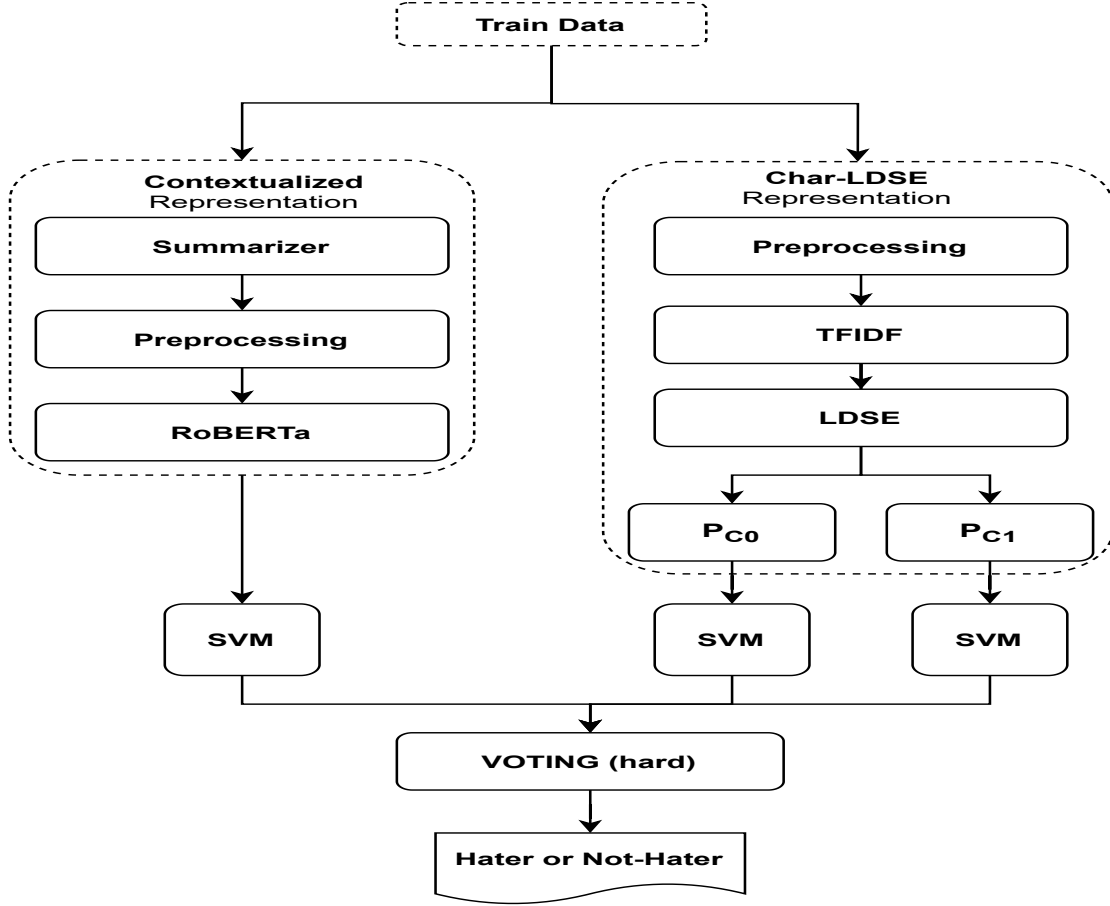


Figure 1: Architecture of Proposed Method

proposed method. We used two different representations namely, contextualized representation using RoBERTa embedding, and character-based lower dimensionality statistical embedding (Char-LDSE) [16]. In the final, representations are fed into the classification modules. In the following, we described each component in detail.

3.1. Contextualized Representation

The user content in this task consists of 200 tweets. To capture the user preferences in tweets and reduce the number of tweets to feed the word embeddings, we applied extractive text summarization. Next, we preprocessed the summary to feed RoBERTa embedding for creating a contextualized representation of user preferences. In the end, for the U users, we had a representation of $U \times R_{768}$ to feed the classification module, where R_{768} is the vectorized representation of summaries using word embedding. In the following, we describe summarization, preprocessing, and word embedding components.

3.1.1. Extractive Text Summarization:

It involves selecting phrases and sentences from the original text and including it in the final summary. We used Gensim [17], a Python library to summarize user tweets with summary ratio of **0.1** (selecting 20 tweets from user 200 tweets). The Gensim uses *TextRank* algorithm, which is based on *PageRank* algorithm for ranking search results.

1. Pre-process the given tweets (a built-in preprocessing in gensim).
2. Make a graph with tweets that are the vertices.
3. The graph has edges denoting the similarity between the two tweets at the vertices.
4. Run PageRank algorithm on this weighted graph.
5. Pick the highest-scoring vertices and append them to the summary.
6. Based on the ratio or the word count, the number of vertices to be picked is decided.

3.1.2. Preprocessing:

The preprocessing consists removal of URLs, hashtags, mentions, reserved words (RT, FAV), emojis, smileys, punctuations, special characters, and numbers from tweets. Specifically for URLs, hashtags, and mentions the following masked tags, #USER#, #URL#, #HASHTAG#.

3.1.3. RoBERTa:

It is an optimized version of BERT model. It builds on BERT's language masking strategy. RoBERTa modifies key hyperparameters in BERT, including removing BERT's next sentence pretraining objective and training with much larger mini-batches and learning rates. For the English language, we used the English version of RoBERTa base model, and for Spanish, we used SpanBERTa¹ which is of the same size as BERT-Base and is trained on 18 GB of OSCAR's Spanish corpus.

3.2. Char-LDSE Representation

Char-LDSE[16] representation is applied to capture the stylistic features of user tweets and the probability of term occurrences in hate and none-hate spreaders. First, preprocessing is applied to user tweets; next, a character n-gram matrix with TFIDF weight is created. TFIDF matrix is utilized to calculate the LDSE. Next, the weighted probability of terms per class was obtained. As a result, P_{C_0} and P_{C_1} embeddings are calculated for class 0 and class 1, respectively. In the end, using these embeddings, we calculated a matrix of $U \times F_{104}$ per class to feed classification models, Where F_{104} is the vectorized distribution of user U weighted probabilities of terms in each class (for class 0, P_{C_0} and for class 1 P_{C_1} was utilized separately)

3.2.1. Preprocessing:

The preprocessing consists removal of special characters and localization of the tweets.

¹<https://github.com/chriskhanhtran/spanish-bert>

3.2.2. TFIDF:

We apply the TFIDF weighting on the terms of the user tweets in the training set. We utilized character n-grams. Specifically, for English, we used range (2, 3), and for Spanish, we used a range of (3, 4). These ranges are obtained using a manual search. As a result, we obtain the following matrix.

$$TFIDF = \begin{bmatrix} W_{11} & W_{12} & \dots & W_{1m} & \gamma(U_{C0}) \\ W_{21} & W_{22} & \dots & W_{2m} & \gamma(U_{C0}) \\ \dots & \dots & \dots & \dots & \dots \\ W_{(n-2)1} & W_{(n-2)2} & \dots & W_{(n-2)m} & \gamma(U_{C1}) \\ W_{(n-1)1} & W_{(n-1)2} & \dots & W_{(n-1)m} & \gamma(U_{C1}) \\ W_{n1} & W_{n2} & \dots & W_{nm} & \gamma(U_{C1}) \end{bmatrix}$$

Where each row in the matrix TFIDF represents a user U_i , each column represents vocabulary term t and W_{ij} represents its TFIDF weight and γ represents the assigned class ($C0$ - class 0, $C1$ - class 1) of the user U tweets. Also, n and m represent the number of the training set (users) and vocabulary size, respectively.

3.2.3. Lower Dimensionality Statistical Embedding (LDSE):

First, using matrix TFIDF, we obtain the class-dependent term weight embedding LDSE. This embedding contains the weights of each term t for each class based on the following formulation.

$$LDSE(t, c) = \frac{\sum_{u \in \gamma(U_c)/c=\gamma(U_c)} W_{ut}}{\sum_{u \in \gamma(U_c)} W_{ut}}, \quad \forall u \in U, c \in \{C0, C1\}$$

Next, we calculated LDSE for each class:

$$P_{C0} = LDSE(t, c), \quad \forall c \in C0$$

$$P_{C1} = LDSE(t, c), \quad \forall c \in C1$$

3.2.4. Final Representation:

At the end, We employed the class-dependent LDSE, P_{C0} and P_{C1} to extract the final representation of user tweets as follows for each class separately:

$$Rep1 = F(P_{C0}), \quad Rep2 = F(P_{C1})$$

Where $F(P)$ contains the set of features showed in the followings and described in Table 1

$$F(P) = \{max, min, std, prob, Q_1, Q_2, Q_3, \dots, Q_{100}\}$$

Table 1

Set of features for each class (hater and none-hater)

avg	The average weight of $W(t, c)$ from a user content
min	The minimum weight of $W(t, c)$ from a user content
std	The standard deviation of the weight of $W(t, c)$
prob	The overall weight of $W(t, c)$ from a user content divided by total number of terms
Q_1, \dots, Q_{100}	calculating the Q-th quantile of the data.

Table 2

Number of authors in the PAN-AP-21 corpus created for profiling hate spreaders on Twitter.

Language	Traning	Testing	Total
English	200	100	300
Spanish	200	100	300

3.3. Classification Modules

There are many different types of ensembles; voting is one of them. It is one of the more general types. Voting involves training a learning algorithm to combine the predictions of several other learning algorithms. We used voting with the hard scheme, in which we trained three different classifiers with three different representations each. The Linear SVM with $C = 0.1$ was utilized as a classifier algorithm for each representation that we obtained as described in Table 3 column Representation + Model.

4. Experiments and Results

In this section, we described the Author Profiling dataset. Next, we presented experimental results on the training set. Finally, we presented the proposed model’s final results on the test set.

4.1. Dataset

Table 2 presents the statistics of the corpus that consists of 300 authors for each of the two languages, English and Spanish. For each author at least 200 Tweets collected. The corpus for each language is balanced, with 150 authors for each class (hater and none-hater spreaders). Dataset have splitted into training and test sets, following the 66/34 proportion.

4.2. Experimental Results

We conducted a few experiments with Linear SVM and different representations. We mainly focused on 5-fold cross-validation mean accuracy. According to experiments for English, ensemble modeling with different representations outperforms single representation modeling. Even in some cases (in Fold-4 and Fold-5), it gains higher accuracy than individuals. It means even contextualized representation which performs week in overall with a mean accuracy of 0.625, contributes to the ensemble model in a positive way. The same patterns are presented in

Table 3

5-Fold Cross Validation Results. In the table Rep1 is Char-LDSE (P_{C0}) and Rep2 is Char-LDSE (P_{C1})

Lan	Representation + Model	Fold-1	Fold-2	Fold-3	Fold-4	Fold-5	Mean Accuracy
en	RoBERTa + Linear SVM	0.600	0.575	0.675	0.600	0.675	0.625
	Rep1 + Linear SVM	0.700	0.625	0.650	0.650	0.750	0.675
	Rep2 + Linear SVM	0.650	0.575	0.600	0.650	0.800	0.655
	Ensemble	0.700	0.625	0.650	0.675	0.825	0.695
es	RoBERTa + Linear SVM	0.650	0.725	0.700	0.525	0.650	0.650
	Rep1 + Linear SVM	0.700	0.825	0.750	0.750	0.775	0.760
	Rep2 + Linear SVM	0.675	0.800	0.750	0.750	0.775	0.750
	Ensemble	0.675	0.825	0.750	0.750	0.775	0.755

Table 4

Submission Results

Language	Early Bird Submission	Final Submission	Best
English	0.670	0.650	0.670
Spanish	0.730	0.800	0.800
Average	0.700	0.725	0.735

Spanish regarding the fact that the second model (Rep1:LDSE(P_{C0}) + Linear SVM) performs in a similar way that Ensemble works. Consequently, we relied upon the power of the ensemble approach for the final submission.

For early bird submission, we simply used XGBoost classifier with Char n-gram with TFIDF weights for the Spanish language. However, we employed an SVM classifier with RBF kernel and LDSE representation with the tree-based feature selection model for English.

We particularly didn't rely on single representations because of the ambiguity of the user behaviors since hate spreaders may have multiple none-hate tweets in their own tweets too.

4.3. Final Evaluation

Following the previous results, for the final evaluation at TIRA platform [18], we applied statistical and contextual representations via ensemble models for hate speech spreaders detection. The obtained accuracy results for the final evaluation were as follows: in Spanish, **0.800**; in English, **0.670**; and **0.735** for both tasks. The official results are shown in Table 4 for early birds and final evaluation. We gained a better result for English at the early bird evaluation. However, for Spanish, we achieved higher accuracy at the final evaluation. In the final evaluation metrics, the best scores of the submissions between the early birds and final submissions of each participant and each language have been considered. This means that in our case, we achieved the best score for English in early bird and the best score for Spanish in the final submission, so, overall achieved accuracy is 0.735.

5. Conclusion

In this paper, we proposed a model for Profiling Hate Speech Spreaders on the Twitter task in PAN 2021. We presented statistical and contextual representations via an ensemble approach for hate speech spreaders detection. In the final, we achieved an average accuracy of 0.735. Based on our manual evaluation, our approach is very capable of distinguishing hate/none-hate speech spreaders. The proposed algorithm implemented in Python and published on GitHub² repository for research community.

References

- [1] F. Rangel, G. L. D. L. P. Sarracén, B. Chulvi, E. Fersini, P. Rosso, Profiling Hate Speech Spreaders on Twitter Task at PAN 2021, in: A. J. M. M. F. P. Guglielmo Faggioli, Nicola Ferro (Ed.), CLEF 2021 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2021.
- [2] J. Risch, R. Krestel, Toxic comment detection in online discussions, in: Deep Learning-Based Approaches for Sentiment Analysis, Springer, 2020, pp. 85–109.
- [3] J. Pennington, R. Socher, C. D. Manning, Glove: Global vectors for word representation, in: Empirical Methods in Natural Language Processing (EMNLP), 2014, pp. 1532–1543. URL: <http://www.aclweb.org/anthology/D14-1162>.
- [4] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, 2013. [arXiv:1301.3781](https://arxiv.org/abs/1301.3781).
- [5] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, [arXiv preprint arXiv:1607.04606](https://arxiv.org/abs/1607.04606) (2016).
- [6] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, 2019. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [7] H. B. Giglou, T. Rahgooy, M. Rahgouy, J. Razmara, Uot-uwf-partai at semeval-2021 task 5: Self attention based bi-gru with multi-embedding representation for toxicity highlighter, 2021. [arXiv:2104.13164](https://arxiv.org/abs/2104.13164).
- [8] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, Language models are unsupervised multitask learners (2019).
- [9] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, 2019. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692).
- [10] B. van Aken, J. Risch, R. Krestel, A. Löser, Challenges for toxic comment classification: An in-depth error analysis, in: Proceedings of the 2nd Workshop on Abusive Language Online (ALW2), Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 33–42. URL: <https://www.aclweb.org/anthology/W18-5105>. doi:10.18653/v1/W18-5105.
- [11] P. Badjatiya, S. Gupta, M. Gupta, V. Varma, Deep learning for hate speech detection in tweets, in: Proceedings of the 26th international conference on World Wide Web companion, 2017, pp. 759–760.
- [12] F. Rangel, A. Giachanou, B. Ghanem, P. Rosso, Overview of the 8th Author Profiling Task at PAN 2020: Profiling Fake News Spreaders on Twitter, in: L. Cappellato, C. Eickhoff, N. Ferro,

²<https://github.com/HamedBabaei/hate-speech>

- A. Névéal (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020. URL: <http://ceur-ws.org/Vol-2696/>.
- [13] J. Pizarro, Using N-grams to detect Fake News Spreaders on Twitter—Notebook for PAN at CLEF 2020, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéal (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020. URL: <http://ceur-ws.org/Vol-2696/>.
- [14] J. Buda, F. Bolonyai, An Ensemble Model Using N-grams and Statistical Features to Identify Fake News Spreaders on Twitter—Notebook for PAN at CLEF 2020, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéal (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020. URL: <http://ceur-ws.org/Vol-2696/>.
- [15] H. Giglou, J. Razmara, M. Rahgouy, M. Sanaei, LSACoNet: A Combination of Lexical and Conceptual Features for Analysis of Fake News Spreaders on Twitter—Notebook for PAN at CLEF 2020, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéal (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020. URL: <http://ceur-ws.org/Vol-2696/>.
- [16] F. Rangel, M. Franco-Salvador, P. Rosso, A low dimensionality representation for language variety identification, 2017. [arXiv:1705.10754](https://arxiv.org/abs/1705.10754).
- [17] R. Řehůřek, P. Sojka, Software Framework for Topic Modelling with Large Corpora, in: Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks, ELRA, Valletta, Malta, 2010, pp. 45–50.
- [18] M. Potthast, T. Gollub, M. Wiegmann, B. Stein, TIRA Integrated Research Architecture, in: N. Ferro, C. Peters (Eds.), Information Retrieval Evaluation in a Changing World - Lessons Learned from 20 Years of CLEF, Springer, 2019.