

Unified and Multilingual Author Profiling for Detecting Haters

(Notebook for PAN at CLEF 2021)

Ipek Baris Schlicht, Angel Felipe Magnoção de Paula

Universitat Politècnica de València, Spain

Abstract

This paper presents a unified user profiling framework to identify hate speech spreaders by processing their tweets regardless of the language. The framework encodes the tweets with sentence transformers and applies an attention mechanism to select important tweets for learning user profiles. Furthermore, the attention layer helps to explain why a user is a hate speech spreader by producing attention weights at both token and post level. Our proposed model outperformed the state-of-the-art multilingual transformer models.

Keywords

Hate speech detection, User profiling, Explainability, Deep Learning, Sentence Transformers, Multilingual

1. Introduction

Hate speech is a type of online harm that expresses hostility toward individuals and social groups based on race, beliefs, sexual orientation, etc. [1]. Hateful content is disseminated faster and reaches wider users than non-hateful contents through social media [2, 3]. This dissemination could trigger prejudices and violence. As a recent example of this, during the COVID-19 pandemic, people of Chinese origin suffered from discrimination and hate crimes [4, 5]. Policymakers and social media companies work hard on mitigating hate speech and the other types of abusive language [6] while keeping balance of freedom of expression. AI systems are encouraged for easing the process and understanding the rationales behind hate speech dissemination [7, 8].

In natural language processing, hate speech has been widely studied in social media (e.g [9, 10]) or as a task of news comment moderation (e.g [11, 12]). However, majority of the prior studies formulates the problem as a text classification [13, 7] that determines whether an individual post is hate speech. This year, PAN 2021 organization [14] proposed to explore the task as an author profiling problem [15]. In this case, the objective is to identify possible hate speech spreaders on Twitter as an initial effort towards preventing hate speech from being propagated among online users [15].


In a similar shared task on profiling fake news spreaders [16], many approaches rely on

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

✉ ibarsch@doctor.upv.es (I. B. Schlicht); adepau@doctor.upv.es (A. F. M. d. Paula)



© 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)

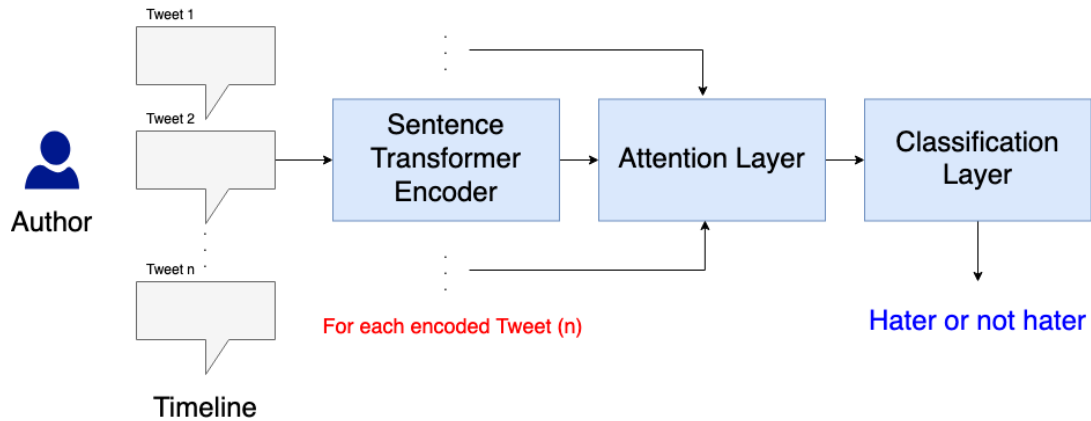


Figure 1: The Proposed Framework

appending tweets to one text for each user (e.g [17, 18, 19]) to encode the inputs. However, this approach could be problematic if not all the tweets shared by hate speech spreaders convey hatred messages, and a human moderator needs a detailed justification to ban users or delete related tweets. Furthermore, the global issues such as COVID-19 attract heated discussions from the users worldwide, thus there is a need for supporting multi-language systems to moderate those discussions. With these motivations, we propose a unified framework which is scalable to other languages and explains why a user receives a certain label based on the language used in her tweets by using token level and post level attention mechanisms [20], as shown in Figure 1. Our model outperformed multilingual DistillBERT [21] models. The source code is publicly available¹.

2. Methodology

Our proposed framework is shown in Figure 1. The input of the framework is a author profile that posts n number tweets. Each post is encoded with a Sentence Transformer, and then the encoded tweets pass through an attention layer. Finally, the output of the attention layer is fed into a classification layer which decides whether the author is a hate speech spreader or not. We give more details of each component in the subsequent sections.

2.1. Post Encodings

We encode the tweets with a Sentence-BERT (SBERT) [22], a modified BERT [23] network and consists of Siamese and Triplet network structures. SBERTs are computationally more effective than BERT models and could provide semantically more meaningful sentence representations. Like BERT models, SBERTs also have variations [24] that are publicly available. Since we have a limited resource to train our framework, and aim to use a language model that learns the

¹<https://github.com/isspek/Cross-Lingual-Cyberbullying>

usages of social language, we prefer the pre-trained SBERT that is trained on Quora corpus in 50 languages, and its knowledge is distilled [25]. The SBERT produces outputs with 768 hidden layers. We set the maximum length of the post as 32, and apply zero padding on any texts shorter than 32 tokens. The sentence embeddings are obtained by mean pooling operation on the last hidden of the outputs.

2.2. Post-Level Attention Layer

We employ an attention layer in order to learn importance scores for determining author profile vectors. First, the pooled tweets (Hp) are projected by feeding them to a linear layer which produces a hidden representation of the author profile (Hap) as shown in Equation 1. Next, a softmax layer is applied to get the similarity between the post and author profile (Hap). Lastly the similarity scores are multiplied with the author profile to obtain the attended author profile ($Hap^{attended}$), as seen in Equation 2.

$$Hap = HpWap + b^T \quad (1)$$

$$Hap^{attended} = softmax(Hp * Hap^T)Hap \quad (2)$$

2.3. Classification Layer

The classification layer consists of two linear layers. The output of the first layer is activated with the tanh function to learn the non-linearity in the features. The second layer outputs the probabilities for each class. The input of the classification layer is the attended user profile followed by a dropout layer which prevents the over-fitting. We use a cross entropy loss function for the outputs of the classification layer and an Adam optimizer with a weight decay. During training, the weights of the models are optimized by minimizing the loss, and the batches contain mixed English and Spanish samples.

3. Experiments

3.1. Dataset

PAN Profiling Hate Speech Spreader Task [15] contains a dataset in English and Spanish, whose samples were collected from Twitter. The total number of the profiles are 200 for each language, and each profile is composed of a feed of 200 tweets. The class distribution of the dataset is highly balanced. We observe a significant difference between the length of tweets by hate speech spreaders and normal profiles in the Spanish set. The statistics of the dataset are summarized in Table 1.

3.2. Preprocessing

The organizers have already cleaned the samples in the dataset. For example, certain patterns have been replaced with special tags. We extend the vocabulary of the models' tokenizers with these tags as follows:

Table 1
The statistics of the training dataset

Stats	En	Es
#Total Profiles	200	200
#Hate Speech Spreaders	100	100
#Tweets per Profile	200	200
#Mean and Std of Tweets by hate speech Spreader	67.72 ± 30.34	75.32 ± 28.91
#Mean and Std of Tweets by Normal Profiles	67.42 ± 29.05	68.47 ± 28.99

- #URL# is replaced with [URL]
- #HASHTAG# is replaced with [HASHTAG]
- #USER# is replaced with [USER]
- RT is replaced with [RT]

3.3. Baselines and Ablation Models

We compare the performance of our model with a set of baselines and an ablation model as follows:

- **DistillBERT** [21]: We use one of its version that is multilingual and cased sensitive. First each tweets of an author is joined to obtain one text. Then the joined texts for each users are fine-tuned with the DistillBERT by keeping their maximum length as 500 tokens.
- **DistillBERT***: We additionally add [POSTSTART] and [POSTEND] tags, which indicate the start and the end of the tweets, to the vocabulary of the extended DistillBERT tokenizer.
- **SBERT-Mean**: is an ablation model that replaces the attention layer with a mean pooling layer which computes the mean values of the tweets' hidden representations.

3.4. Training Settings

We train the models by applying 5-Fold Cross Validation², with the epochs of 5, learning rate as $1e-5$, batch size as 2. We use the GPU of the Google Colab³ as an environment for training the models. We use a fixed random seed of 1234 to ensure reproducible results. The official results are obtained by a TIRA machine [26].

4. Results and Discussion

We report the F1-Macro, F1-Weighted, accuracy, precision, and recall for each model. Table 2 presents the results of the 5-fold cross validation training. SBERT-Attn, the model that we propose, outperformed the other models in all metrics. When we compare SBERT-Mean and SBERT-Attn, we see that standard deviations of the SBERT-Attn are lower than the ablation

²We experiment also 10-Fold, but the models show worse performance in the test set.

³<https://colab.research.google.com/>

model. This result indicates that the attention layer enables more generalized feature representations. It also shows that the tweets by the hate speech spreader are not necessarily hatred tweets and vice versa for the non haters. For this reason, the DistillBERT models that joins the all tweets by the user to one underperformed.

Table 2

The results of the 5 Fold Cross Validation Experiment

Models	F1-Macro	F1-Weighted	Accuracy	Precision	Recall
DistillBERT	67.46 ± 5.28	67.58 ± 5.37	67.75 ± 5.15	67.04 ± 5.68	71.46 ± 1.63
DistillBERT*	61.90 ± 3.01	62.04 ± 3.22	62.25 ± 3.39	63.13 ± 4.40	59.86 ± 7.49
SBERT-Mean	69.55 ± 6.82	69.58 ± 6.71	69.75 ± 6.86	67.38 ± 3.61	77.10 ± 12.12
SBERT-Attn	73.62 ± 4.11	73.77 ± 4.12	74.0 ± 4.14	70.97 ± 5.39	81.23 ± 5.39

Table 3


Cross validation for each language and the PAN shared official result.

Mode	Language	Accuracy
Cross-Val	En	67.09 ± 7.88
	Es	80.54 ± 1.78
Official Result	En	58
	Es	77

For the submission to the PAN shared task, we leverage the 5-fold trained models to obtain the predictions on official test set. The final predictions are the majority class. Table 3 shows cross validation results for the English samples and the Spanish samples, and the official results of the PAN shared task where the accuracy is the evaluation metric. Our model obtained a result with similar range in cross-validation. The performance of the English set is worse than the Spanish one. Cultural bias or the topical difference could be reasons for the performance. We leave the detailed analysis of these issues as future work.

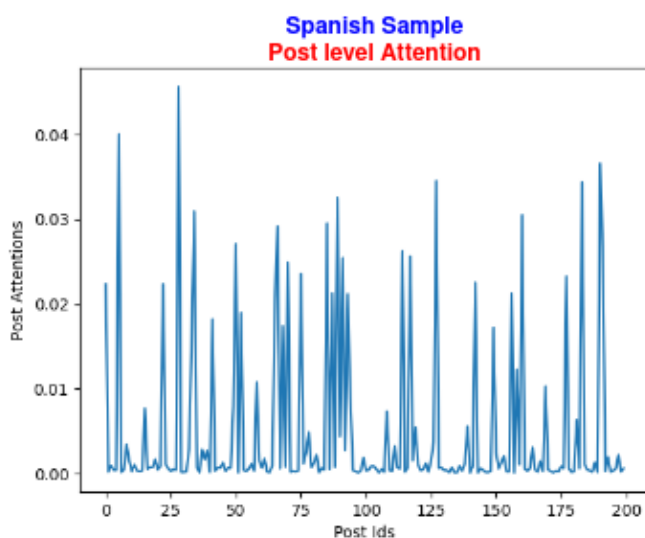
5. Visualizations

Our framework can provide explanations with tweet-level and token-level attention, as shown in Figure 2. The token-level attentions are the average of the attentions in the last layer of the SBERT and they are obtained through the self-attention mechanism. The tweet-level attentions are obtained with the attention layer, which is connected to the classification layer. The examples in the figure are the most hatred examples from the authors that are analysed. In the English example, the model pays attention to feminism. In the Spanish example, *vice presidencia* is the important entity.



Token level Attention

[CLS] [USER] [USER] Yes , you ' re a part of fem. ##ins. ##m . And that ' s because you aren ' t a man .



Token level Attention

[CLS] [USER] [USER] Le quedan grandes , como su vice. ##presl. ##den. ##cia [UNK] [SEP]

Figure 2: Attention visualizations for English and Spanish. The original sentence in English is [USER] [USER] *Yes, you're a part of feminism. And that's because you aren't a man*; and the other in Spanish is [USER] [USER] *Le quedan grandes, como su vicepresidencia* (Some emojis)

6. Conclusion

In this paper, we presented a unified framework for monitoring hate speech spreaders in multilingualism. The framework leverages multilingual SBERT representations to encode texts regardless of the language and uses an attention mechanism to determine the importance of the tweets by the author in the task. Our methods outperformed multilingual DistillBERT and SBERT that apply mean pooling on the tweets.

In the future, we plan to evaluate the method on the related user profiling tasks such as profiling fake news spreaders [16] and investigate advanced method (e.g [27]) for effectively transferring knowledge across the languages.

References

- [1] L. W. Levy, *Encyclopedia of the American Constitution*, New York: Macmillan; London: Collier Macmillan, 1986.
- [2] B. Mathew, R. Dutt, P. Goyal, A. Mukherjee, Spread of hate speech in online social media, in: *WebSci, ACM*, 2019, pp. 173–182.
- [3] C. Ziemis, B. He, S. Soni, S. Kumar, Racism is a virus: Anti-asian hate and counterhate in social media during the COVID-19 crisis, *CoRR abs/2005.12423* (2020).
- [4] S. Wang, X. Chen, Y. Li, C. Luu, R. Yan, F. Madrisotti, ‘i’m more afraid of racism than of the virus!’: racism awareness and resistance among chinese migrants and their descendants in france during the covid-19 pandemic, *European Societies* 23 (2021) S721–S742.
- [5] J. He, L. He, W. Zhou, X. Nie, M. He, Discrimination and social exclusion in the outbreak of covid-19, *International Journal of Environmental Research and Public Health* 17 (2020) 2933.
- [6] P. Nakov, V. Nayak, K. Dent, A. Bhatawdekar, S. M. Sarwar, M. Hardalov, Y. Dinkov, D. Zlatkova, G. Bouchard, I. Augenstein, Detecting abusive language on online platforms: A critical analysis, *CoRR abs/2103.00153* (2021).
- [7] A. Schmidt, M. Wiegand, A survey on hate speech detection using natural language processing, in: *Proceedings of the fifth international workshop on natural language processing for social media*, 2017, pp. 1–10.
- [8] P. Fortuna, S. Nunes, A survey on automatic detection of hate speech in text, *ACM Comput. Surv.* 51 (2018) 85:1–85:30.
- [9] V. Basile, C. Bosco, E. Fersini, D. Nozza, V. Patti, F. M. R. Pardo, P. Rosso, M. Sanguinetti, Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter, in: *SemEval@NAACL-HLT*, Association for Computational Linguistics, 2019, pp. 54–63.
- [10] F. Poletto, V. Basile, M. Sanguinetti, C. Bosco, V. Patti, Resources and benchmark corpora for hate speech detection: a systematic review, *Language Resources and Evaluation* (2020) 1–47.
- [11] D. Korencic, I. Baris, E. Fernandez, K. Leuschel, E. Salido, To block or not to block: Experiments with machine learning for news comment moderation, in: *Proceedings of the EACL Hackathon on News Media Content Analysis and Automated Report Generation*,

Association for Computational Linguistics, Online, 2021, pp. 127–133. URL: <https://www.aclweb.org/anthology/2021.hackashop-1.18>.

- [12] R. Shekhar, M. Pranjic, S. Pollak, A. Pelicon, M. Purver, Automating news comment moderation with limited resources: Benchmarking in croatian and estonian, *The Journal for Language Technology and Computational Linguistics (JLCL)*, Special Issue on Offensive Language (2020) 49.
- [13] S. MacAvaney, H.-R. Yao, E. Yang, K. Russell, N. Goharian, O. Frieder, Hate speech detection: Challenges and solutions, *PloS one* 14 (2019) e0221152.
- [14] J. Bevendorff, B. Chulvi, G. L. D. L. P. Sarracén, M. Kestemont, E. Manjavacas, I. Markov, M. Mayerl, M. Potthast, F. Rangel, P. Rosso, E. Stamatatos, B. Stein, M. Wiegmann, M. Wol-ska, , E. Zangerle, Overview of PAN 2021: Authorship Verification, Profiling Hate Speech Spreaders on Twitter, and Style Change Detection, in: *12th International Conference of the CLEF Association (CLEF 2021)*, Springer, 2021.
- [15] 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: *CLEF 2021 Labs and Workshops, Notebook Papers*, CEUR-WS.org, 2021.
- [16] F. M. R. Pardo, A. Giachanou, B. Ghanem, P. Rosso, Overview of the 8th author profiling task at PAN 2020: Profiling fake news spreaders on twitter, in: *CLEF 2020 Labs and Workshops, Notebook Papers*, volume 2696 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020.
- [17] I. Vogel, M. Meghana, Fake news spreader detection on twitter using character n-grams, in: *CLEF 2020 Labs and Workshops, Notebook Papers*, 2020.
- [18] J. Buda, F. Bolonyai, An ensemble model using n-grams and statistical features to identify fake news spreaders on twitter, in: *CLEF 2020 Labs and Workshops, Notebook Papers*, 2020.
- [19] J. Pizarro, Using n-grams to detect fake news spreaders on twitter, in: *CLEF 2020 Labs and Workshops, Notebook Papers*, volume 2696 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020.
- [20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: *NIPS*, 2017, pp. 5998–6008.
- [21] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter, *CoRR abs/1910.01108* (2019).
- [22] N. Reimers, I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, in: *EMNLP/IJCNLP (1)*, Association for Computational Linguistics, 2019, pp. 3980–3990.
- [23] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, in: *NAACL-HLT (1)*, Association for Computational Linguistics, 2019, pp. 4171–4186.
- [24] T. Wolf, J. Chaumond, L. Debut, V. Sanh, C. Delangue, A. Moi, P. Cistac, M. Funtowicz, J. Davison, S. Shleifer, et al., Transformers: State-of-the-art natural language processing, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 38–45.
- [25] N. Reimers, I. Gurevych, Making monolingual sentence embeddings multilingual using knowledge distillation, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 2020.

- [26] 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, The Information Retrieval Series*, Springer, Berlin Heidelberg New York, 2019. doi:10.1007/978-3-030-22948-1_5.
- [27] J. Pfeiffer, A. Rücklé, C. Poth, A. Kamath, I. Vulić, S. Ruder, K. Cho, I. Gurevych, Adapterhub: A framework for adapting transformers, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 46–54.