



A Large-Scale Dataset for Hate Speech Detection on Vietnamese Social Media Texts

Son T. Luu^{1,2(✉)}, Kiet Van Nguyen^{1,2}, and Ngan Luu-Thuy Nguyen^{1,2}

¹ University of Information Technology, Ho Chi Minh City, Vietnam
{sonlt,kietnv,ngannlt}@uit.edu.vn

² Vietnam National University, Ho Chi Minh City, Vietnam

Abstract. In recent years, Vietnam witnesses the mass development of social network users on different social platforms such as Facebook, Youtube, Instagram, and Tiktok. On social media, hate speech has become a critical problem for social network users. To solve this problem, we introduce the ViHSD - a human-annotated dataset for automatically detecting hate speech on the social network. This dataset contains over 30,000 comments, each comment in the dataset has one of three labels: CLEAN, OFFENSIVE, or HATE. Besides, we introduce the data creation process for annotating and evaluating the quality of the dataset. Finally, we evaluate the dataset by deep learning and transformer models.

Keywords: Hate speech detection · Social media texts · Machine learning · Text classification

1 Introduction

There are approximately 70 million Internet users in Vietnam, and most of them are familiar with Facebook and Youtube. Nearly 70% of total Vietnamese people use Facebook, and spend an average of 2.5 h per day on it [24]. On social network, hate easily appears and spread. According to Kang and Hall [21], hate-speech does not reveal itself as hatred. It is just a mechanism to protect an individual's identity and realm from the others. Hate leads to the destruction of humanity, isolating people, and debilitating society. Within the development of social network sites, the hate appears on social media as hate-speech comments, hate-speech posts, or messages, and it spreads too fast. The existence of hate speech makes the social networking spaces toxic, threatens social network users, and bewilders the community.

The automated hate speech detection task is categorized as the supervised learning task, specifically closed to the sentiment analysis task[26]. There are several state-of-the-art approaches such as deep learning and transformer models for sentiment analysis. However, to be able to make experiment on hate speech

detection task, the datasets, especially large-scale datasets, play an important role. To handle this, we introduce ViHSD - a large-scale dataset used for automatically hate speech detection on Vietnamese social media texts to overcome the hate-speech problem on social networks. Then, we present the annotation process for our dataset and the method to ensure the quality of annotators. Finally, we evaluate our dataset on SOTA models and analyze the obtained empirical results to explore the advantages and disadvantages of the models on the dataset.

The content of the paper is structured as follows. Section 2 takes an overview on current researching for hate-speech detection. Section 3 shows statistical figures about our dataset as well as our annotation procedure. Section 4 presents the classification models applied for our dataset to solve the hate-speech detection problem. Section 5 describes our experiments on the dataset and the analytical results. Finally, Sect. 6 concludes and proposes future works.

2 Related Works

In English as well as other languages, there are many datasets constructed for hate speech detection. We divided them into two categories: flat labels and hierarchical labels. According to Cerri et al. [7], flat labels are treated as no relation between different labels. In contrast, hierarchical labels has a hierarchical structure which one or more labels can have sub-labels, or being grouped in super-labels. Besides, Hatebase [27] and Hurltex [6] are two abusive words sets used for lexicon-based approaching for the hate speech detection problem.

For flat labels, we introduce two typical and large-scale datasets in English. The first dataset is provided by Waesem and Hovy [30], which contains 17,000 tweets from Twitter and has three labels: racism, sexism, and none. The second dataset is provided by Davidson et al. [12], which contains 25,000 tweets from Twitter and also has three labels including: hate, offensive, and neither. Apart from English, there are other datasets in other languages such as: Arabic [2], and Indonesian [3].

For hierarchical labels datasets, Zampieri et al. [32] provide a multi-labelled dataset for predicting offensive posts on social media in English. This dataset serves two tasks: group-directed attacking and person-directed attacking, in which each task has binary labels. Another similar multi-labelled dataset in English are provided by Basile et al. [5] in SemEval Task 5 (2019). Other multi-labelled datasets in other non-English languages are also constructed and available such as: Portuguese [17], Spanish [16], and Indonesian [19]. Besides, multi-lingual hate speech corpora are also constructed such as hatEval with English and Spanish [5] and CONAN with English, French, and Italian [9].

The VLSP-HSD dataset provided by Vu et al. [29] is a dataset used for the VLSP 2019 shared task about Hate speech detection on Vietnamese language¹. However, the authors did not mention the annotation process and the method for evaluating the quality of the dataset. Besides, on the hate speech detection problem, many state-of-the-art models give optimistic results such as deep

¹ <https://www.aivivn.com/contests/8>.

learning models [4] and transformer language models [20]. Those models require large-scale annotated datasets, which is a challenge for low-resource languages like Vietnamese. Moreover, current researches about hate speech detection do not focus on analyzing about the sentiment aspect of Vietnamese hate speech language. Those ones are our motivation to create a new dataset called ViHSD for Vietnamese with strict annotation guidelines and evaluation process to measure inter-annotator agreement between annotators.

3 Dataset Creation

3.1 Data Preparation

We collect users' comments about entertainment, celebrities, social issues, and politics from different Vietnamese Facebook pages and YouTube videos. We select Facebook pages and YouTube channels that have a high-interactive rate, and do not restrict comments. After collected data, we remove the name entities from the comments in order to maintain the anonymity.

3.2 Annotation Guidelines

The ViHSD dataset contains three labels: HATE, OFFENSIVE, and CLEAN. Each annotator assigns one label for each comment in the dataset. In the ViHSD dataset, we have two labels denoting for hate speech comments, and one label denoting for normal comments. The detailed meanings about three labels and examples for each label are described in Table 1.

Practically, many comments in the dataset are written in informal form. Comments often contain abbreviation such as *M.n* (English: Everyone), *mik* (English: us) in **Comment 1** and *Dm* (English: f*ck) in **Comment 2**, and slangs such as *chjch* (English: f*ck), *cái lol* (English: p*ssy) in **Comment 2**. Besides, comments has the figurative meaning instead of explicit meaning. For example, the word: *lũ quan ngại* (English: dummy pessimists) in **Comment 3** is usually used by many Vietnamese Facebook users on social media platform to mention a group of people who always think pessimistically and posting negative contents.

3.3 Data Creation Process

Our annotation process contains two main phases as described in Fig. 1. The first one is the training phase, which annotators are given a detailed guidelines, and annotate for a sample of data after reading carefully. Then we compute the inter-annotator agreement by Cohen Kappa index (κ) [10]. If the inter-annotators agreement not good enough, we will re-train the annotator, and re-update the annotation guidelines if necessary. After all annotators are well-trained, we go to annotation phase. Our annotation phase is inspired from the IEEE peer review process of articles [15]. Two annotators annotate the entire dataset. If there are any different labels between two annotators, we let the third annotators annotate

Table 1. Annotation guidelines for annotating Vietnamese comments in the Hate Speech Detection task.

Label	Description	Example
CLEAN	The comments have no harassment at all.	Comment 1: M.n ơi cho mik hỏi mik theo dõi cô mà mik hk pít cô là con gái thiệt hả m.n (<i>English: Hey everyone! Is she a truly girl?</i>) (This comment is thoroughly clean, in which there are no bad words or profane language, and does not attack anyone)
OFFENSIVE	The comments contain harassment contents, even profanity words, but do not attack any specific object.	Comment 2: Đồ <u>khùng</u> (<i>English: Madness</i>) (This comment has offensive word " <u>khùng</u> ". Nevertheless, it does not contain any word that aims to a person or a group. In addition, " <u>khùng</u> " is also a slang, which means mad)
HATE	The comments have harassment and abusive contents and directly aim at an individual or a group of people based on their characteristics, religion, and nationality. Some exceptional cases happened with the HATE label: Case 1: The comments have offensive words and attack a specific object such as an individual, a community, a nation, or a religion. This case is easy to identify hate speech. Case 2: The comments have racism, harassment, and hateful meaning, however, does not contain explicit words. Case 3: The comments have racism, harassment, and hateful meaning, but showed as figurative meaning. To identify this comment, users need to have particular knowledge about social.	Comment 3: Dành cho <u>lũ quan ngại</u> (<i>English: This is for those dummy pessimists</i>) (This comment contains a phrase, which is underlined, mentions to a group of people with bad meaning) Comment 4: <u>Dm</u> Có a mới không <u>ồn</u> . Mà y rình mà <u>chịch riết</u> <u>ồn cái lol</u> (<i>English: F*ck you. I am not fine. You are fine why you're making sex ?</i>) (This comments contained many of profanities, which are underlined. Besides, it contains personal pronoun "Mày", which aims to a specific person) Comment 3: Ở <u>đấy</u> ngột ngột quá thì đưa nó qua <LOC> cho nó thoáng mát (<i>English: If this place is so stifling and not suitable for bitch like you, why don't you choose <LOC>?</i>) (This comment has the phrase <LOC> mentioned to a specific location, which has racism meaning. However, this comment does not has any bad words at all)

those labels. The fourth annotators annotate if all three annotators are disagreed. The final label are defined by Major voting. By this way, we guaranteed that each comment is annotated by one label and the objectivity for each comment. Therefore, the total time spent on annotating is less than four annotators doing with the same time.

3.4 Dataset Evaluation and Discussion

We randomly take 202 comments from the dataset and give them to four different annotators, denoted as A1, A2, A3 and A4, for annotating. Table 2 shows the inter-annotator agreement between each pair of annotators. Then, we compute

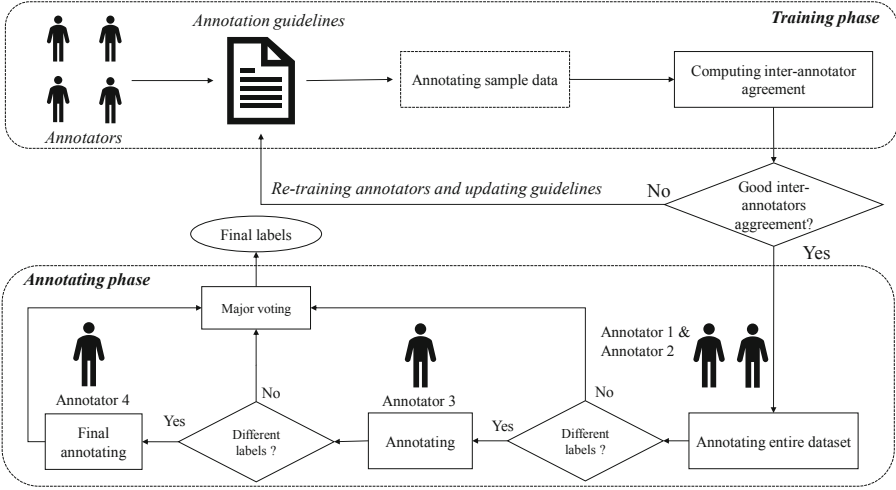


Fig. 1. Data annotation process for the ViHSD dataset. According to Eugenio [14], $\kappa > 0.5$ is acceptable.

the average inter-annotator agreement. The final inter-annotator agreement for the dataset is $\kappa = 0.52$.

Table 2. The confusion matrix between annotators in a set of 202 comments computed by Cohen Kappa index (κ).

	A1	A2	A3	A4
A1	–	0.46	0.51	0.65
A2	–	–	0.47	0.53
A3	–	–	–	0.55
A4	–	–	–	–

The ViHSD dataset was crawled from the social network so they had many abbreviations, informal words, slangs, and figurative meaning. Therefore, it confuses annotators. For example, the **Comment 1** contains the phrase: *mik*, which mean mình (English: I), and the **Comment 4** in Table 1 has the profane word *Dm* (English: m*ther f**ker). Assume that two annotators assign label for **Comment 4**, and it contains the word *Dm* written in abbreviation form. The first annotator knew about this word before, thus he/she annotates this comment as hate. The second annotator instead, annotates this comment as clean because he/she do not understand that word. The next example is the phrase *lũ quan ngãi* (English: dummy pessimists) in the **Comment 3** in Table 1. Two annotators assign label for **Comment 3**. The first annotator does not understand the real meaning of that phrase, thus he/she marks this comment as clean. In

contrast, the second annotator knows what is the real meaning of this word (See Sect. 3.2 for the meaning of that phase) so he/she knows the abusive meaning of **Comment 3** and annotates it as hate. Although the guidelines have clearly definition about the CLEAN, OFFENSIVE, and HATE labels, the annotation process is mostly impacted by the knowledge and subjective of annotators. Thus, it is necessary to re-train annotators and improve the guidelines continuously to increase the quality of annotators and the inter-annotator agreement.

3.5 Dataset Overview

Table 3. Several examples extracted from the ViHSD dataset.

#	Comments	Label
1	Nhanh thực sự (<i>English: It is really fast</i>)	clean
2	Đm chứ biết làm gì (<i>English: How f*ck damn to do that!</i>)	offensive
3	Nó học cách của thằng anh nó đó, hèn, khốn nạn (<i>English: He is coward and bastard likes his brother</i>)	hate
4	<person name> người ư? sinh vật hạ đẳng chứ người ai làm thế (<i>English: <person name> person? It is a inferior animal. Human don't do that</i>)	hate
5	Đm vcl (<i>English: God damn it!!</i>)	offensive




Fig. 2. The distributions of three labels on the train, dev, and test sets.

The ViHSD contains 33,400 comments. Each one was labelled as CLEAN (0), OFFENSIVE (1), and HATE (2). Table 3 displays some examples from the ViHSD dataset. Then we divided our dataset into training (train), development

(dev), and test sets, respectively, with proportion: 7-1-2. Figure 2 describes the distribution of data on three labels on those sets. According to Fig. 2, the distribution of data labels on the training, development, and test sets are the same, and the data are skewed to the CLEAN label.

4 Baseline Models

The problem of text classification is defined according to Aggarwal and Zhai [1] as given a set of training texts as training data $D = \{X_1, X_2, \dots, X_n\}$, in which each $X_i \in D$ has one of the label in the label set $\{1..k\}$. The training data is used to build the classification model. Then, for unlabeled data coming, the classification model predicts the label for it. In this section, we introduce two approaches for constructing prediction models on the ViHSD dataset.

4.1 Deep Neural Network Models (DNN Models)

Convolutional Neural Network (CNN) uses particular layers called the CONV layer for extracting local features in the image [23]. However, although invented for computer vision, CNN can also be applied to Natural Language Processing (NLP), in which a filter W relevant to a window of h words [22]. Besides, the pre-trained word vectors also influence the performance of the CNN model [22]. Besides, Gated Recurrent Unit (GRU) is a variant of the RNN network. It contains two recurrent networks: encoding sequences of texts into a fixed-length word vector representation, and another for decoding word vector representation back to raw input [8].

We implement the Text-CNN and the GRU models, and evaluate them on the ViHSD dataset with the *fasttext*² pre-trained word embedding of 157 different languages provided by Grave et al. [18]. This embedding transforms a word into a 300 dimension vector.

4.2 Transformer Models

The transformer model [28] is a deep neural network architecture based entirely on the attention mechanism, replaced the recurrent layers in auto encoder-decoder architectures with special called multi-head self-attention layers. Yang et al. [31] found that the transformer blocks improved the performance of the classification model. In this paper, we implement BERT[13] - the SOTA transformer model with multilingual pre-training³ such as bert-base-multilingual-uncased (m-BERT uncased) and bert-base-multilingual-cased (m-BERT cased), DistilBERT [25] - a lighter but faster variant of BERT model with multilingual cased pre-trained model⁴, and XLM-R[11] - a cross-lingual language model with xlm-roberta-base pre-trained⁵. Those multilingual pre-trained models are trained on various of languages including Vietnamese.

² <https://fasttext.cc/docs/en/crawl-vectors.html>.

³ <https://github.com/google-research/bert/blob/master/multilingual.md>.

⁴ <https://huggingface.co/distilbert-base-multilingual-cased>.

⁵ <https://huggingface.co/xlm-roberta-base>.

5 Experiment

5.1 Experiment Settings

First of all, we pre-process our dataset as follows: (1) Word-segmentation texts into words by the pyvi tool⁶, (2) Removing stopwords⁷, (3) Changing all texts into lower cases⁸, and (4) Removing special characters such as hashtags, urls, and mention tags.

Next, we run the Text-CNN model with 50 epochs, batch size equal to 256, sequence length equal to 100, and the dropout ratio is 0.5. Our model uses 2D Convolution Layer with 32 filters and size 2, 3, 5, respectively. Then, we run the GRU model with 50 epochs, sequence length equal to 100, the dropout ratio is 0.5, and the bidirectional GRU layer. We use the Adam optimizer for both Text-CNN and GRU. Finally, we implement transformer models includes BERT, XLM-R, and DistilBERT with the batch size equal to 16 for both training and evaluation, 4 epochs, sequence length equal to 100, and manual seed equal to 4.

5.2 Experiment Results

Table 4 illustrates the results of deep neural models and transformer models on the ViHSD dataset. The results are measured by Accuracy and macro-averaged F1-score. According to Table 4, Text-CNN achieves 86.69% in accuracy and 61.11% in F1-score, which is better than GRU. Transformer models such as BERT, XLM-R, and DistilBERT give better results than deep neural models in F1-score. The BERT with bert-base-multilingual-cased model (m-bert cased) obtained best result in both Accuracy and F1-score with 86.88% and 62.69% respectively on the ViHSD dataset.

Table 4. Empirical results of baseline models on the test set

	Model	Pre-trained model	Accuracy (%)	F1-macro (%)
DNN models	Text CNN	fastText	86.69	61.11
	GRU	fastText	85.41	60.47
Transformer models	BERT	bert-base-multilingual-uncased	86.60	62.38
		bert-base-multilingual-cased	86.88	62.69
	XLM-R	xlm-roberta-based	86.12	61.28
	DistilBERT	distilbert-base-multilingual-cased	86.22	62.42

Overall, the performance of transformer models are better than deep neural models, indicating the power of BERT and its variants on text classification task, especially on hate speech detection even if they were trained on various languages. Additionally, there is a large gap between the accuracy score and the F1-score, which caused by the imbalance in the dataset, as described in Sect. 3.

⁶ <https://pypi.org/project/pyvi/>.

⁷ List Vietnamese stopwords <https://github.com/stopwords/vietnamese-stopwords>.

⁸ We do not lower case texts with cased pre-trained transformer models.

5.3 Error Analysis

Figure 3 shows the confusion matrix of the m-BERT based model. Most of the offensive comments in the dataset are predicted as clean comments. Besides, Table 5 shows incorrect predictions by the m-BERT based model. The comments number 1, 2, and 3 had many special words, which are only used on the social network such as: “*dell*”, “*coin card*”, and “*éo*”. These special words make those comments had wrong predicting labels, misclassified from offensive labels to clean labels. Moreover, the comments number 4 and 5 have profane words and are written in abbreviation form and teen codes such as “*cc*”, “*lol*”. Specifically, for the fifth comments, in which “*3*” represents for “*father*” in English, combined with other bad words such as: “*Cc*” - profane word and “*m*” - represents for “*you*” in English. Generally, the fifth comment has bad meaning (see Table 5 for the English meaning) thus it is the hate speech comment. However, the classification model predicts that comment as an offensive label because it cannot identify the target objects and the profane words written in irregular form.

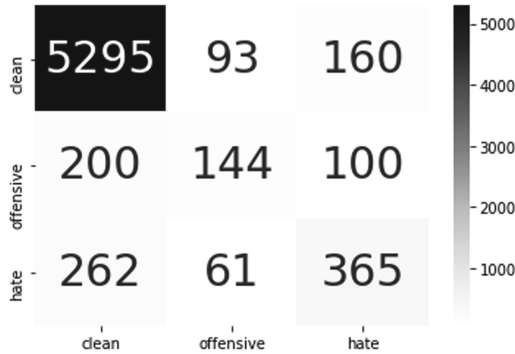


Fig. 3. Confusion matrix of the m-BERT based model on the ViHSD dataset

In general, the wrong prediction samples and most of the ViHSD dataset comments are written with irregular words, abbreviations, slangs, and teencodes. Therefore, in the pre-processing process, we need to handle those characteristic to enhance the performance of classification models. For abbreviations and slangs, we can build a Vietnamese slangs dictionary for slangs replacement and a Vietnamese abbreviation dictionary to replace the abbreviations in the comments. Besides, for words that are not found in mainstream media, we can try to normalize them to the regular words. For example, words like *giiiiii*, *ummmmm* should be normalized to *gì* (what) and *ừ* (ok). In addition, the emoji icons are also a polarity feature to define whether a comment is negative or positive, which can support for detecting the hate and offensive content in the comments.

Table 5. Wrong prediction samples in the ViHSD dataset

#	Comments	True	Predict
1	coin card :3 (<i>English: F*ck d*ck</i>)	1	0
2	<person name> dell hiểu kiểu gì :) (<i>English: I do not f*cking understand why</i>)	1	0
3	Éo thích thì biến đi thọc con mắt chóng cái tai lên xem làm gì (<i>English: Get out, damn! Prick your eyes if you don't care.</i>)	1	0
4	<person name>chửi cc (<i>English: f*ck you</i>)	2	1
5	Cc 3 m (<i>English: Like your dad's d*ck</i>)	2	1

6 Conclusion

We constructed a large-scale dataset called the ViHSD dataset for hate speech detection on Vietnamese social media texts. The dataset contains 33,400 comments annotated by humans and achieves 62.69% by the Macro F1-score with the BERT model. We also proposed an annotation process to save time for data annotating.

The current inter-annotator agreement of the ViHSD dataset is just in moderate level. Therefore, our next studies focus on improving the quality of the dataset based on the data annotation process. Besides, the best baseline model on the dataset is 62.69%, and this is a challenge for future researches to improve the performance of classification models for Vietnamese hate-speech detection task. From the error analysis, we found that it is difficult to detect the hate speech on Vietnamese social media texts due to their characteristics. Hence, we will improve the pre-processing technique for social media texts such as lexicon-based approach for teen codes and normalizing acronyms on next studies to increase the performance of this task.

References

1. Aggarwal, C.C., Zhai, C.: A survey of text classification algorithms. In: Aggarwal, C., Zhai, C. (eds.) *Mining Text Data*, pp. 163–222. Springer, Boston (2012). https://doi.org/10.1007/978-1-4614-3223-4_6
2. Albadi, N., Kurdi, M., Mishra, S.: Are they our brothers? Analysis and detection of religious hate speech in the Arabic Twittersphere. In: *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 69–76 (2018)
3. Alfina, I., Mulia, R., Fanany, M.I., Ekanata, Y.: Hate speech detection in the Indonesian language: a dataset and preliminary study. In: *2017 International Conference on Advanced Computer Science and Information Systems (ICACSIS)*, pp. 233–238 (2017)

4. Badjatiya, P., Gupta, S., Gupta, M., Varma, V.: Deep learning for hate speech detection in tweets. In: WWW 2017 Companion, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, pp. 759–760 (2017)
5. Basile, V., et al.: 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. Association for Computational Linguistics, Minneapolis (2019)
6. Bassignana, E., Basile, V., Patti, V.: Hurtlex: a multilingual lexicon of words to hurt. In: CLiC-it (2018)
7. Cerri, R., Barros, R.C., de Carvalho, A.C.: Hierarchical multi-label classification using local neural networks. *J. Comput. Syst. Sci.* **80**(1), 39–56 (2014)
8. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1724–1734. Association for Computational Linguistics, Doha (2014)
9. Chung, Y.L., Kuzmenko, E., Tekiroglu, S.S., Guerini, M.: CONAN - COUNTER NARratives through nichesourcing: a multilingual dataset of responses to fight online hate speech. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, Florence (2019)
10. Cohen, J.: A coefficient of agreement for nominal scales. *Educ. Psychol. Measur.* **20**(1), 37–46 (1960)
11. Conneau, A., et al.: Unsupervised cross-lingual representation learning at scale. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 8440–8451. Association for Computational Linguistics, Online (2020)
12. Davidson, T., Warmsley, D., Macy, M., Weber, I.: Automated hate speech detection and the problem of offensive language (2017)
13. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, vol. 1 (Long and Short Papers). Association for Computational Linguistics, Minneapolis (2019)
14. Di Eugenio, B.: On the usage of kappa to evaluate agreement on coding tasks. In: Proceedings of the Second International Conference on Language Resources and Evaluation (LREC 2000). European Language Resources Association (ELRA), Athens (2000)
15. El-Hawary, M.E.: What happens after i submit an article? [Editorial]. *IEEE Syst. Man Cybern. Mag.* **3**(2), 3–42 (2017)
16. Fersini, E., Rosso, P., Anzovino, M.: Overview of the task on automatic misogyny identification at ibereval 2018. In: IberEval@ SEPLN 2150, pp. 214–228 (2018)
17. Fortuna, P., Rocha da Silva, J., Soler-Company, J., Wanner, L., Nunes, S.: A hierarchically-labeled Portuguese hate speech dataset. In: Proceedings of the Third Workshop on Abusive Language Online, pp. 94–104. Association for Computational Linguistics, Florence (2019)
18. Grave, E., Bojanowski, P., Gupta, P., Joulin, A., Mikolov, T.: Learning word vectors for 157 languages. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018). European Language Resources Association (ELRA), Miyazaki (2018)

19. Ibrohim, M.O., Budi, I.: Multi-label hate speech and abusive language detection in Indonesian Twitter. In: Proceedings of the Third Workshop on Abusive Language Online. Association for Computational Linguistics, Florence (2019)
20. Isaksen, V., Gambäck, B.: Using transfer-based language models to detect hateful and offensive language online. In: Proceedings of the Fourth Workshop on Online Abuse and Harms. Association for Computational Linguistics, Online (2020)
21. Kang, M., Hall, P.: Hate Speech in Asia and Europe: Beyond Hate and Fear. Taylor & Francis Group, Routledge Contemporary Asia (2020)
22. Kim, Y.: Convolutional neural networks for sentence classification. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, Doha (2014)
23. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
24. Nguyen, T.N., McDonald, M., Nguyen, T.H.T., McCauley, B.: Gender relations and social media: a grounded theory inquiry of young Vietnamese women’s self-presentations on Facebook. *Gender Technol. Dev.* **24**, 1–20 (2020)
25. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter (2020)
26. Schmidt, A., Wiegand, M.: A survey on hate speech detection using natural language processing. In: Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media, pp. 1–10 (2017)
27. Tuckwood, C.: Hatebase: online database of hate speech. The Sentinel Project (2017). <https://www.hatebase.org>
28. Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)
29. Vu, X.S., Vu, T., Tran, M.V., Le-Cong, T., Nguyen, H.T.M.: HSD shared task in VLSP campaign 2019: hate speech detection for social good. In: Proceedings of VLSP 2019 (2019)
30. Waseem, Z., Hovy, D.: Hateful symbols or hateful people? Predictive features for hate speech detection on Twitter. In: Proceedings of the NAACL Student Research Workshop, pp. 88–93. Association for Computational Linguistics, San Diego (2016)
31. Yang, X., Yang, L., Bi, R., Lin, H.: A comprehensive verification of transformer in text classification. In: Sun, M., Huang, X., Ji, H., Liu, Z., Liu, Y. (eds.) CCL 2019. LNCS (LNAI), vol. 11856, pp. 207–218. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32381-3_17
32. Zampieri, M., Malmasi, S., Nakov, P., Rosenthal, S., Farra, N., Kumar, R.: Predicting the type and target of offensive posts in social media. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, vol. 1 (Long and Short Papers), Minneapolis, Minnesota (2019)