

Overview of the HASOC Subtrack at FIRE 2022: Offensive Language Identification in Marathi

Tharindu Ranasinghe¹, Kai North², Damith Premasiri^{1,*} and Marcos Zampieri²

¹University of Wolverhampton, United Kingdom

²George Mason University, United States

Abstract

The widespread of offensive content online has become a reason for great concern in recent years, motivating researchers to develop robust systems capable of identifying such content automatically. With the goal of carrying out a fair evaluation of these systems, several international competitions have been organized, providing the community with important benchmark data and evaluation methods for various languages. Organized since 2019, the HASOC (Hate Speech and Offensive Content Identification) shared task is one of these initiatives. In its fourth iteration, HASOC 2022 included three subtracks for English, Hindi, and Marathi. In this paper, we report the results of the HASOC 2022 Marathi subtrack which provided participants with a dataset containing data from Twitter manually annotated using the popular OLID taxonomy. The Marathi track featured three additional subtracks, each corresponding to one level of the taxonomy: Task A - offensive content identification (offensive vs. non-offensive); Task B - categorization of offensive types (targeted vs. untargeted), and Task C - offensive target identification (individual vs. group vs. others). Overall, 59 runs were submitted by 10 teams. The best systems obtained an F1 of 0.9745 for Subtrack 3A, an F1 of 0.9207 for Subtrack 3B, and F1 of 0.9607 for Subtrack 3C. The best performing algorithms were a mixture of traditional and deep learning approaches.

Keywords

offensive language, hate speech, text categorization, Marathi

1. Introduction

Social media is the main source of information, current events, and public opinion for many individuals [1]. Users typically spend on average 2 hours per day on social media platforms, such as Twitter, Facebook, and Instagram [2]. The impact social media has on an individual's perception of the world, as well as their mental well-being, is substantial [3]. Social media often contains hateful and offensive content, including but not limited to pejorative language [4], cyber-bullying [5, 6, 7], online extremism [8], and racism [9, 10]. Such content has been shown to cause anxiety or depression [3], incite self-harm [11], or violence towards others [12].

With governments realizing the impact social media has on public health and opinion, social media companies have invested in ways to cope with hateful and offensive content [12, 13]. To achieve this, they are interested in the development of machine learning (ML) models to automatically detect such content on their social media platforms. These models have

Forum for Information Retrieval Evaluation, December 9-13, 2022, India

*Corresponding author.

✉ damith.premasiri@wlv.ac.uk (D. Premasiri)



© 2022 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)

traditionally consisted of Naive Bayes (NB) [14, 15, 9], support vector machines (SVMs) [16, 17, 18, 19], and random forests (RFs) [16, 20] while, more recently, neural networks [21, 22, 23, 24], and transformer-based models [23, 25, 26, 27] have been applied to this problem.

As evidenced by recent international competitions on this topic [28, 29, 30, 31, 32, 33, 34], a known challenge in detecting offensive and hateful content is the intrinsic subjectivity of such phenomena. In highly multilingual countries, such as India, this task becomes even more challenging. India has 22 official languages, many of which do not have large annotated datasets that can be used for the training and evaluation of hate speech and offensive language identification models. Marathi is an example of such an under-resourced language, especially when in comparison to high-resourced languages, such as Hindi or English [35].

This paper presents the results of Subtrack 3 at HASOC-2022 [36]: Offensive Language Identification in Marathi. HASOC-2022 challenged participating teams with the development of state-of-the-art hate speech and offensive language identification models. Subtrack 3 at HASOC-2022 aimed to encourage the development of such models for Marathi and thus encourage more research into the future development of models for under-resourced languages.

This paper is structured as follows. Section 2 introduces previous hate speech and offensive language shared-tasks. Section 3 introduces Subtrack 3. Section 4 describes the data used for this subtrack. Section 5 lists the approaches adopted by each participating team. Section 6 provides each teams' performance. Section 7 gives a conclusion on what is the current state-of-the-art for identifying hate speech and offensive language in Marathi and makes predictions regarding the direction of future research.

2. Related Work

Numerous shared-tasks have been organized with the goal of providing benchmarks for the automatic identification of hateful and offensive content. These shared-tasks have consisted of differing forms classification, being either binary, multi-class, or multi-label classification, and contained differing subtracks aimed at predicting the type or target of the offensive content as well as dealing with offensive content in languages other than English.

2.1. Related Tasks

TRAC 2018-2022 The first Trolling, Aggression and Cyberbullying (TRAC-2018) workshop hosted at COLING [28], challenged participating teams to create multi-class classifiers for predicting which of three levels of aggression is depicted by a Facebook comment: overt aggression, covert aggression, or nonaggression. Their dataset consisted of randomly sampled English and Hindi Facebook comments. Several issues were reported to the organizers regarding dataset quality: 1). some instances contained code-mixing between English and Hindi, and 2). some instances were considered incorrectly labelled due to their subjectivity. The second TRAC workshop (TRAC-2020) hosted at LREC [34], introduced a new language, Bengali, as well as a new domain, Youtube comments, to resolve the previous issues in data quality. TRAC-2020 also provided a new subtrack in the form of binary classification of gender: gendered (GEN) or non-gendered (NGEN) instances. The third TRAC workshop (TRAC-2022) [37] expanded this subtrack further to include the recognition of other types of offensive content, including

racial bias, religious intolerance, and casteist bias, as well as the identification of the discursive role of a given instance in relation to a previous comment. It also introduced a new domain, “COVID-19 related conversation” [37]. Performances on TRAC-2018 varied with neural networks performing on par with other approaches, such as SVMs and RFs [28]. TRAC-2020 and 2022, however, saw transformer-based models achieve the best performances [34]; teams that used an ensemble of transformers performed exceptionally well in comparison to other approaches.

HatEval 2019 Task 5 at SemEval-2019: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter (HatEval) [31], contained two subtracks. Subtrack A was a binary classification task. It saw English or Spanish tweets be classified into hatespeech or non-hatespeech instances. Subtrack B was evaluated as a multi-label classification task which required the identified instances of hatespeech to be further grouped in regards to their “aggressive attitude and the target harassed” [31]. The organizers suggested that participating teams divide this task into two binary classifications tasks, whereby each instance’s target is classified as being aggressive or non-aggressive, and its attitude is classified as being either a targeted message (individual) or a hateful message in general (generic). The first and second top performing models for Subtrack A were both SVMs having outperformed neural networks and other approaches. Performances on Subtrack B were noticeably worse than Subtrack A. The organizers contributed this to its multi-labelled nature.

OffensEval 2019-2020 Task 6 at SemEval-2019: Identifying and Categorizing Offensive Language in Social Media (OffensEval-2019) [29], presented three subtracks to participating teams ranging from binary to multi-class classification on the Offensive Language Identification Dataset (OLID) consisting of English tweets. Subtrack A saw the binary classification of tweets with non-offensive and offensive labels. Subtrack B and C were concerned with identifying the target of the offensive tweet. Subtrack B was the binary classification of instances into either targeted or untargeted insults, whereas Subtrack C was the multi-class classification of instances into three labels concerned with the target type: “individual, group, or other” [29]. Task 12 at SemEval-2020: Multilingual Offensive Language Identification in Social Media (OffensEval-2020) [33] was an extension of the previous shared-task. It included the same subtracks as SemEval-2019, but presented additional tracks and data in Arabic, Danish, English, Greek, and Turkish. Both shared-tasks saw transformer-based models achieve the best results, such as BERT, RoBERTa, and mBERT. The multilingual nature of OffensEval-2020 also allowed for cross-lingual transfer learning, which was adopted by the highest performing systems.

2.2. HASOC

HASOC 2019-2020 The Hate Speech and Offensive Content Identification in Indo-European Languages track (HASOC-2019) at FIRE [30], held three subtracks. Data was collected from Twitter and Facebook in Hindi, German, and English. Subtrack A was a binary classification task to distinguish between hateful and non-hateful instances. Subtracks B and C were more fine-grained classification tasks to identify the type (hate speech, offensive, and profanity) and target (targeted or untargeted) of each instance respectively. HASOC-2020 [32] saw the removal of Subtrack C and the creation of a new dataset with the aim of minimizing bias in the

data collection process. This was achieved through the adoption of a more randomized sampling technique that minimized the selection of types of hate-speech known to the dataset’s creators [32]. Neural networks achieved the best performances at HASOC-2019, with an LSTM model achieving the highest performance for English [21]. BERT-based models become more popular at HASOC-2020 and achieved the best performance for German. However, LSTM-based models again reported the best results for Hindi and English. Overall performances were considered to be worst for HASOC-2020 than they were for HASOC-2019 [32]. Organizers’ reported this as being due to the more realistic sampling technique used in HASOC-2020.

Marathi at HASOC-2021 HASOC-2021 [38] included the same two subtracks as HASOC-2020 for English and Hindi, but was the first to apply these subtracks to Marathi. The data used for Marathi was extracted from the *MOLD* dataset [35]. The *MOLD* dataset consists of 2,547 hateful and non-hateful tweets extracted from Twitter’s API. Its hateful tweets were obtained by searching for keywords “related to politics, entertainment, and sports along with the hashtag #Marathi” [35]. All instances were then labelled by six Marathi speakers. The best performing models for Marathi used cross-lingual transfer learning [39]. Nene et al. [39] trained an XLM-R model on pre-existing data in Hindi from HASOC 2019. They then used the returned model’s updated weights and softmax activation layer to initialize the weights of a new XLM-R model on the provided Marathi dataset. The second best performing model adopted a similar approach. Glazkova et al. [40] fine-tuned the monolingual LaBSE and multilingual XLM-R models on Marathi and Hindi data and discovered that the multilingual XLM-R model achieved the best results. The other top-10 performing models submitted to HASOC-2021’s Marathi subtrack included feature engineering approaches using word and character level n-grams [41, 42], an LSTM model [43], as well as other transformer-based models, such as mBERT and DistilBERT [44]. However, these models failed to achieve performances on par with those models fine-tuned on both Hindi and Marathi data.

3. HASOC-2022

3.1. Subtrack 3: Offensive Language Identification in Marathi

HASOC-2022 hosted three subtracks with subtracks 1 and 2 dealing with binary and multi-class classification of German and the code-mixed language, Hinglish, respectively. Subtrack 3 focused on the automatic identification of offensive language in Marathi. Subtrack 3 was further divided into three additional subtracks.

- **Subtrack 3A:** A binary classification task classifying posts into two classes: offensive (OFF), and non-offensive (NOT).
- **Subtrack 3B:** A binary classification task identifying whether a post is either targeted (TIN) or untargeted (UNT).
- **Subtrack 3C:** A more fine-grained classification task, whereby the previously identified TIN posts were further classified into individual (IND), group (GRP), or other (OTH).

4. Data

The data used for Subtrack 3 consisted of extracts taken from *MOLD 2.0* [45]. *MOLD 2.0* is an expansion of the original MOLD dataset (v1.0) [35]. It contains 3,600 tweets with an additional 1053 instances on top of the original instances provided by *MOLD* (v1.0) [35]. All instances were annotated using the popular OLID taxonomy [46] with additional information being provided in regards to the target (level B) and target type (level C) of the offensive or non-offensive tweet. The annotation was carried out by three native speakers of Marathi with an inter-annotator agreement of 0.79 Cohen’s kappa across all instances.

The *MOLD 2.0* dataset was split into an 80/20 split between train and test sets, respectively. An unequal number of offensive (1067) and non-offensive (2034) instances were provided to the participating teams resulting in the label distribution shown in Table 1.

A	B	C	Training	Test	Total
OFF	TIN	IND	503	51	554
OFF	TIN	OTH	80	56	136
OFF	TIN	GRP	157	51	208
OFF	UNT	—	327	102	429
NOT	—	—	2,034	250	2,284
All			3,101	510	3,611

Table 1

MOLD v2.0 - Distribution of labels. Table taken from [45].

5. Approaches

In total, 8 out of the 10 participating teams submitted system reports for Subtrack 3. 10 teams took part in Subtrack 3A, 7 in Subtrack 3B, and 6 in Subtrack 3C. System descriptions are provided below.

5.1. Traditional

ssncse_nlp Vani and Bharathi [47] used Scikit-learn’s CountVectorizer function [48] to generate vectors of term/token counts. They fed these vectors as features into four types of traditional classifier: RF, SVM, logistic regressor, and a KNN classifier. They discovered that their RF model achieved the greatest performance and submitted this model for all three subtracks.

satlab Bestgen [49] submitted a simplified version of their previous model used for HASOC-2021 [38]. They trained a logistic regressor taken from the LIBLinear package [50] on several character-level n-gram features. These features consisted of n-gram length, sentence position, and BM25 frequency. All n-grams with a frequency of less than two and over five characters in length were not considered. Feature scores were normalized per L2 regularization, and

parameter optimization was conducted using four-fold cross validation over several random grid searches. This model was submitted to all three subtracks.

ml_ai_iitranchi Kumari and Singh [51] experimented with a logistic regressor, SVM, and RF. They lemmatized each instance and trained their models on token frequencies obtained using Scikit-learn’s TF-IDF-Vectorizer function [48]. Their logistic regressor surpassed the performances of the other three models for Subtrack 3A, whereas their SVM achieved the best performance for subtracks 3B and 3C. They submitted their best performing models to each of the subtracks, respectively.

5.2. Deep Learning

optimize_prime Chavan et al. [52] compared the performances of various BERT-based models taken from HuggingFace that have been pre-trained on Marathi data [45]: MuRIL [53], MahaTweetBERT [54], MahaTweetBERT-Hateful [54], and MahaBERT [55]. They cleaned each instance provided by the shared-task of placeholder texts, such as “@USER”, and removed hastags, URLs, empty parentheses and new lines. Each instance was then automatically encoded and then fed into each model. The AdamW optimizer was used during training with a learning rate of 1e-5 and batch size of 32. The training was conducted over 25 epochs. The MahaTweetBERT [54] outperformed all of models when trained on a combination of instances from HASOC-2021 [38] and HASOC-2022 [36]. This model was only submitted for Subtrack 3A.

hate-busters Ripoll et al. [56] treated Subtrack 3 as a multi-class classification problem instead of a multi-label or several binary classification tasks, and attempted to predict all possible labels at once (as shown in Table 1). They applied five and ten-fold cross validation by splitting the training and test sets into n independent sets and trained an ensemble of n separate models. The final output label was derived from plurality voting. They tested two different types of model: InfoXLM [57] and L3Cube [55]. The former being trained on cross-lingual representations of 100 languages, whereas the later being trained solely on Hindi and Marathi data. The more language-specific L3Cube model achieved the greatest performance in their ensemble architecture. This ensemble was submitted for all three subtracks.

sakshi_hasoc Kalra et al. [58] fine-tuned mBERT [59], DistilmBERT [60], and MuRIL [53] on a resampled version of the provided training set to compensate for the imbalance between offensive and non-offensive instances. They oversampled the training set by including duplicates of offensive instances that made for a 0.5 ratio between the two labels. They achieved this using the RandomOverSampler function in the imblearn library [61]. MuRIL was found to achieve the best performance when trained over 4 epochs and set to a learning rate of 1e-4 and a batch size of 4. They submitted this model to Subtrack 3A.

irlab@iitbhu Chanda et al. [62] removed usernames, punctuation, URLs, and emojis. They then concatenated the tweet, its comments, and responses and fed these concatenated instances into Fast.ai [63], being an off-the-shelf deep learning package. They submitted their Fast.ai model to all three subtracks.

citk_isi Ghosh et al. [64] compared the performance of mBERT [59] and MahaBERT [55] with the latter being pre-trained on the L3 Cube-MahaCorpus. Each model was trained with a learning rate of $1e-5$ and a batch size of 5 over 5 epochs. MahaBERT was found to outperform mBERT and was submitted for all three subtracks.

6. Results

Team performances are displayed within tables 2 to 4 and are discussed throughout the following section.

Rank	Name	Number of Runs	Precision	Recall	F1
1	ssncse_nlp [47]	4	0.9758	0.9741	0.9745
2	optimize_prime [52]	2	0.9588	0.9589	0.9588
3	hate-busters [56]	5	0.9592	0.9562	0.9568
4	sakshi_hasoc [58]	3	0.9464	0.9446	0.9450
5	satlab [49]	2	0.9372	0.9373	0.9372
6	irlab@iitbhu [?]	5	0.9353	0.9355	0.9353
7	ml_ai_iiitranchi [51]	2	0.9292	0.9203	0.921
8	citk_isi [?]	1	0.9021	0.9022	0.9020
9	gunjan	3	0.8395	0.8071	0.7997
10	the shivi hunters	1	0.8053	0.7775	0.7746

Table 2

Results of Subtrack 3A. Ordered from best to least performing system.

Rank	Name	Number of Runs	Precision	Recall	F1
1	hate-busters [56]	5	0.9114	0.9342	0.9207
2	satlab [49]	2	0.9060	0.9288	0.9149
3	ssncse_nlp [47]	4	0.7587	0.7067	0.6958
4	irlab@iitbhu [?]	1	0.4837	0.6107	0.5354
5	ml_ai_iiitranchi [51]	2	0.4516	0.4371	0.4429
6	citk_isi [?]	1	0.3405	0.2869	0.3073
7	gunjan	1	0.1111	0.2073	0.0897

Table 3

Results of Subtrack 3B. Ordered from best to least performing system.

Rank	Name	Number of Runs	Precision	Recall	F1
1	satlab [49]	2	0.9365	0.9893	0.9607
2	ssncse_nlp [47]	4	0.7963	0.8663	0.7929
3	ml_ai_iiitranchi [51]	2	0.7927	0.7464	0.7423
4	hate-busters [56]	5	0.7999	0.7705	0.7255
5	irlab@iitbhu [?]	1	0.2858	0.4213	0.2891
6	citk_isi [?]	1	0.2323	0.1961	0.2064

Table 4

Results of Subtrack 3C. Ordered from best to least performing system.

Team `ssncse_nlp` [47] achieved the highest F1 of 0.974 for Subtrack 3A. Their RF model

surpassed the performance of several transformer-based models, `optimize_prime` [52] and `hate-busters` [56], that came 2nd and 3rd place respectively. Other traditional approaches also performed on par with deep learning approaches for this subtrack, with the logistic regressors submitted by `satlab` [49] and `ml_ai_iitranchi` [51] achieving similar performances to the transformer-based models provided by `irlab@iitbhu` [62] and `citk_isi` [64]. The size of the provided dataset is a possible explanation. Deep learning approaches were likely unable to infer useful feature representations that are able to differentiate between offensive and non-offensive instances, given that the provided dataset only contains 1067 offensive and 2034 non-offensive tweets. The submitted traditional approaches, on the other hand, were all engineered with features, such as token frequency or character-level n-grams, that have been proven to be useful for hate speech and offensive language identification [30, 32, 38]. Those deep learning approaches that did perform well, namely `optimize_prime` [52] and `sakshi_hasoc` [58], also increased the amount of available training data by using augmented datasets. These augmented datasets included data from HASOC-2021 [38] or oversampled instances [58] that likely increased their models' ability to infer distinguishing features.

Team `hate-busters` [56] came 1st place with an F1 of 0.921 for Subtrack 3B, closely followed by `satlab` [49] having achieved an F1 of 0.915. Team `hate-busters` was the only team that treated all three subtracks as one multi-class classification problem and that adopted an ensemble architecture. Their ensemble of L3Cude models in particular may be responsible for their high performance, as their final target output labels were generated by averaging those of multiple transformers trained on 10-folds. This is further supported by the performance of the non-ensemble-based deep learning approaches submitted by `irlab@iitbhu` [62] and `citk_isi` [64], with the latter team's submission also being pre-trained on L3Cude data. These teams scored drastically worst F1s of 0.5354 and 0.3073 respectively. However, the character-level n-gram features used by `satlab` appeared to have performed equally as well as `hate-busters`'s ensemble. The use of such features as n-gram length, sentence position, and BM25 frequency, were likely able to identify character combinations or words, including pronouns, that are useful for the identification of targeted instances [49]. In addition, it appears that `satlab` did not remove meta data, such as "`@USER`", that would likewise have been beneficial; many of the other submissions removed such data during their pre-processing [52, 62].

The 1st, 2nd, and 3rd leading teams of Subtrack 3C all used traditional approaches. The logistic regressor of `satlab` [49] came first place with F1 of 0.9607, whilst the RF and SVM of `ssncse_nlp` [47] and `ml_ai_iitranchi` [51], came 2nd and 3rd place with F1s of 0.7929 and 0.7423 respectively. The deep learning approaches submitted to this subtrack all achieved inferior performances with the ensemble submitted by `hate-busters` [56] attaining the best F1 of 0.7255 and the non-ensemble-based models submitted by `irlab@iitbhu` [62] and `citk_isi` [64] getting noticeably worst F1s of 0.2891 and 0.2064 respectively. Traditional approaches were therefore superior for Subtrack 3C. A highly likely explanation for this is once again the size of the provided dataset. Subtrack 3C was a multi-class classification task, whereby models had to differentiate between three types of target for each instance: individual (IND), group (GRP), or other (OTH). However, only 503 instances with an individual target, 157 instances with a group target, and 80 instances with a target marked as other, were provided to the participating teams (Table 1). A recommended approach for dealing with such small sample sizes is to apply cross-lingual transfer learning for transformer-based models, such as training transformers on

large datasets in Hindi and applying said models to Marathi, as shown by the winning system of HASOC-2021 [39].

7. Conclusion and Future Work

This paper has described Subtrack 3 of HASOC-2022: Offensive Language Identification in Marathi. Subtrack 3 consisted of three additional subtracks for Marathi concerned with first distinguishing between hateful and non-hateful tweets, and then identifying whether a tweet had a target as well as its target type. In total, 10 teams submitted systems to this task and those that submitted system reports have been described in Section 5. Team performances have been listed and discussed in Section 6. In general, traditional approaches achieved the best performances, especially for Subtrack 3C. A RF (ssncse_nlp) came 1st, 3rd, and 2nd place for subtracks 3A, 3B and 3C respectively [47], and a logistic regressor trained on n-gram features (satlab) came 2nd place for Subtrack 3B and 1st place for Subtrack 3C [49]. Deep learning approaches performed less well. However, those that utilized an augmented dataset (optimize_prime and sakshi_hasoc) [52, 58], or that had an ensemble architecture (hate-busters) [56] outperformed those that used a singular pre-trained deep learning model (irlab@iitbhu and citk_isi) [62, 64]. We believe this to be due to the small size of the provided dataset.

Marathi continues to be an under-resourced language for many NLP-related tasks. Subtrack 3 aimed to encourage the development of ML models that were capable of performing well, regardless of having a limited amount of available data. The results of Subtrack 3 indicate that traditional approaches can still be effective when given a small training set. Nevertheless, ensemble architectures and cross-lingual transfer learning would likely surpass the performance of traditional approaches if said training sets were supplemented with additional instances in Marathi or Hindi. We suspect these to become common approaches within future iterations of this shared-task.

Acknowledgments

We would like to thank the teams referenced in this report for participating in the competition. We further thank the organizers of the other HASOC tracks for their support.

References

- [1] R. Karlsen, T. Aalberg, Social Media and Trust in News: An Experimental Study of the Effect of Facebook on News Story Credibility, *Digital Journalism* 0 (2021) 1–17.
- [2] K. E. Riehm, K. A. Feder, K. N. Tormohlen, R. M. Crum, A. S. Young, K. M. Green, L. R. Pacek, L. N. La Flair, R. Mojtabai, Associations Between Time Spent Using Social Media and Internalizing and Externalizing Problems Among US Youth, *JAMA Psychiatry* 76 (2019) 1266–1273.
- [3] R. Bannink, S. Broeren, P. M. van de Looij-Jansen, F. G. de Waart, H. Raat, Cyber and Traditional Bullying Victimization as a Risk Factor for Mental Health Problems and Suicidal Ideation in Adolescents, *PLoS One* 9 (2014).

- [4] L. P. Dinu, I.-B. Iordache, A. S. Uban, M. Zampieri, A Computational Exploration of Pejorative Language in Social Media, in: Proceedings of EMNLP, 2021.
- [5] H. Rosa, N. Pereira, R. Ribeiro, P. C. Ferreira, J. P. Carvalho, S. Oliveira, L. Coheur, P. Paulino, A. V. Simão, I. Trancoso, Automatic Cyberbullying Detection: A Systematic Review, *Computers in Human Behavior* 93 (2019) 333–345.
- [6] M. Yao, C. Chelmiss, D.-S. Zois, Cyberbullying Ends Here: Towards Robust Detection of Cyberbullying in Social Media, in: Proceedings of WWW, 2019.
- [7] J. Shetty, K. N. Chaithali, A. M. Shetty, B. Varsha, V. Puthran, Cyber-Bullying Detection: A Comparative Analysis of Twitter Data, in: N. N. Chiplunkar, T. Fukao (Eds.), *Advances in Artificial Intelligence and Data Engineering*, Springer Singapore, Singapore, 2021, pp. 841–855.
- [8] S. Aldera, A. Emam, M. Al-Qurishi, M. Alrubaian, A. Alothaim, Online extremism detection in textual content: A systematic literature review, *IEEE Access* 9 (2021) 42384–42396.
- [9] I. Kwok, Y. Wang, Locate the Hate: Detecting Tweets against Blacks, in: Proceedings of AAAI, 2013.
- [10] S. Tulkens, L. Hilte, E. Lodewyckx, B. Verhoeven, W. Daelemans, A Dictionary-based Approach to Racism Detection in Dutch Social Media, in: Proceedings of TA-COS, 2016.
- [11] A. John, A. Glendenning, A. Marchant, P. Montgomery, A. Stewart, S. Wood, K. Lloyd, K. Hawton, Self-Harm, Suicidal Behaviours, and Cyberbullying in Children and Young People: Systematic Review, *Journal of Medical Internet Research* 20 (2018).
- [12] M. L. Williams, P. Burnap, A. Javed, H. Liu, S. Ozalp, Hate in the Machine: Anti-Black and Anti-Muslim Social Media Posts as Predictors of Offline Racially and Religiously Aggravated Crime, *The British Journal of Criminology* 60 (2019).
- [13] A. Akins, Facebook’s Oversight Board Overrules 4 Hate Speech, Misinformation Take-downs, SNL Kagan Media and Communications Report (2021).
- [14] K. Dinakar, R. Reichart, H. Lieberman, Modeling the Detection of Textual Cyberbullying, in: Proceedings of ICWSM, 2011.
- [15] Y. Chen, Y. Zhou, S. Zhu, H. Xu, Detecting Offensive Language in Social Media to Protect Adolescent Online Safety, in: Proceedings of ASE, 2012.
- [16] J. Salminen, H. Almerexhi, M. Milenkovic, S.-g. Jung, A. Jisun, H. Kwak, B. J. Jansen, Anatomy of Online Hate: Developing a Taxonomy and Machine Learning Models for Identifying and Classifying Hate in Online News Media, in: Proceedings of ICWSM, 2018.
- [17] J.-M. Xu, K.-S. Jun, X. Zhu, A. Bellmore, Learning from Bullying Traces in Social Media, in: Proceedings of NAACL, 2012.
- [18] M. Dadvar, D. Trieschnigg, R. Ordelman, F. de Jong, Improving Cyberbullying Detection with User Context, in: Proceedings of ECIR, 2013.
- [19] C. Nobata, J. Tetreault, A. Thomas, Y. Mehdad, Y. Chang, Abusive Language Detection in Online User Content, in: Proceedings of WWW, 2016.
- [20] K. Nugroho, E. Noersasongko, Purwanto, Muljono, A. Z. Fanani, Affandy, R. S. Basuki, Improving Random Forest Method to Detect Hatespeech and Offensive Word, in: Proceedings of ICOIACT, 2019.
- [21] B. Wang, Y. Ding, S. Liu, X. Zhou, YNU_Wb at HASOC 2019: Ordered Neurons LSTM with Attention for Identifying Hate Speech and Offensive Language, in: Proceedings of FIRE, 2019.

- [22] M. A. Bashar, R. Nayak, QutNocturnal at HASOC 2019: CNN for Hate Speech and Offensive Content Identification in Hindi Language, in: Proceedings of FIRE, 2019.
- [23] P. Saha, B. Mathew, P. Goyal, A. Mukherjee, HateMonitors: Language Agnostic Abuse Detection in Social Media, in: Proceedings of FIRE, 2019.
- [24] H. Hettiarachchi, T. Ranasinghe, Emoji Powered Capsule Network to Detect Type and Target of Offensive Posts in Social Media, in: Proceedings of RANLP, 2019.
- [25] R. Kumar, B. Lahiri, A. K. Ojha, A. Bansal, ComMA at HASOC 2020: Exploring Multilingual Joint Training across different Classification Tasks, in: Proceedings of FIRE, 2020.
- [26] T. Ranasinghe, M. Zampieri, H. Hettiarachchi, BRUMS at HASOC 2019: Deep Learning Models for Multilingual Hate Speech and Offensive Language Identification, in: Proceedings of FIRE, 2019.
- [27] D. Sarkar, M. Zampieri, T. Ranasinghe, A. Ororbia, fBERT: A Neural Transformer for Identifying Offensive Content, in: Proceedings of EMNLP Findings, 2021.
- [28] R. Kumar, A. K. Ojha, S. Malmasi, M. Zampieri, Benchmarking Aggression Identification in Social Media, in: Proceedings of TRAC, 2018.
- [29] M. Zampieri, S. Malmasi, P. Nakov, S. Rosenthal, N. Farra, R. Kumar, SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval), in: Proceedings of SemEval, 2019.
- [30] T. Mandl, S. Modha, P. Majumder, D. Patel, M. Dave, C. Mandlia, A. Patel, Overview of the HASOC Track at FIRE 2019: Hate Speech and Offensive Content Identification in Indo-European languages, in: Proceedings of FIRE, 2019.
- [31] 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: Proceedings of SemEval, 2019.
- [32] T. Mandl, S. Modha, A. Kumar M, B. R. Chakravarthi, Overview of the HASOC Track at FIRE 2020: Hate Speech and Offensive Language Identification in Tamil, Malayalam, Hindi, English and German, in: Proceedings of FIRE, 2020.
- [33] M. Zampieri, P. Nakov, S. Rosenthal, P. Atanasova, G. Karadzhov, H. Mubarak, L. Derczynski, Z. Pitenis, c. Çöltekin, SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media (OffensEval 2020), in: Proceedings of SemEval, 2020.
- [34] R. Kumar, A. K. Ojha, S. Malmasi, M. Zampieri, Evaluating Aggression Identification in Social Media, in: Proceedings of TRAC, 2020.
- [35] S. Gaikwad, T. Ranasinghe, M. Zampieri, C. M. Homan, Cross-lingual Offensive Language Identification for Low Resource Languages: The Case of Marathi, in: Proceedings of RANLP, 2021.
- [36] S. Satapara, P. Majumder, T. Mandl, S. Modha, H. Madhu, T. Ranasinghe, M. Zampieri, K. North, D. Premasiri, Overview of the HASOC Subtrack at FIRE 2022: Hate Speech and Offensive Content Identification in English and Indo-Aryan Languages, in: Proceedings of FIRE, 2022.
- [37] R. Kumar, A. K. Ojha, M. Zampieri, S. Malmasi, D. Kadar (Eds.), Proceedings of the Third Workshop on Threat, Aggression and Cyberbullying (TRAC 2022), Association for Computational Linguistics, Gyeongju, Republic of Korea, 2022.
- [38] S. Modha, T. Mandl, G. K. Shahi, H. Madhu, S. Satapara, T. Ranasinghe, M. Zampieri, Overview of the HASOC Subtrack at FIRE 2021: Hate Speech and Offensive Content

- Identification in English and Indo-Aryan Languages and Conversational Hate Speech, in: Proceedings of FIRE, 2021.
- [39] M. Nene, K. North, T. Ranasinghe, M. Zampieri, Transformer Models for Offensive Language Identification in Marathi, in: Proceedings of FIRE, 2021.
 - [40] A. Glazkova, M. Kadantsev, M. Glazkov, Fine-tuning of Pre-trained Transformers for Hate, Offensive, and Profane Content Detection in English and Marathi, in: Proceedings of FIRE, 2021.
 - [41] Y. Bestgem, A Simple Language-agnostic yet Very Strong Baseline System for Hate Speech and Offensive Content Identification, in: Proceedings of FIRE, 2021.
 - [42] Y. HaCohen-Kerner, M. Uzan, Detecting Offensive Language in English, Hindi, and Marathi using Classical Supervised Machine Learning Methods and Word/Char N-grams, in: Proceedings of FIRE, 2021.
 - [43] A. Velankar, H. Patil, A. Gore, S. Salunke, R. Joshi, Hate and Offensive Speech Detection in Hindi and Marathi, in: Proceedings of FIRE, 2021.
 - [44] M. Bhatia, T. S. Bhotia, A. Agarwal, P. Ramesh, S. Gupta, K. Shridhar, F. Laumann, A. Dash, One to Rule Them All: Towards Joint Indic Language Hate Speech Detection, in: Proceedings of FIRE, 2021.
 - [45] M. Zampieri, T. Ranasinghe, M. Chaudhari, S. Gaikwad, P. Krishna, M. Nene, S. Paygude, Predicting the Type and Target of Offensive Social Media Posts in Marathi, *Social Network Analysis and Mining* 12 (2022) 77.
 - [46] M. Zampieri, S. Malmasi, P. Nakov, S. Rosenthal, N. Farra, R. Kumar, Predicting the Type and Target of Offensive Posts in Social Media, in: Proceedings of NAACL, 2019.
 - [47] D. Vani, B. Bharathi, Hate Speech and Offensive Content Identification in Multiple Languages using Machine Learning Algorithms, in: Proceedings of FIRE, 2022.
 - [48] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* 12 (2011).
 - [49] Y. Bestgen, Confirming the Effectiveness of a Simple Language-Agnostic yet Very Strong System for Hate Speech and Offensive Content Identification, in: Proceedings of FIRE, 2022.
 - [50] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, C.-J. Lin, Liblinear: A library for large linear classification, *Journal of Machine Learning Research* 9 (2008) 1871–1874.
 - [51] K. Kumari, J. P. Singh, Machine Learning Approach for Hate Speech and Offensive Content Identification in English and Indo Aryan Code-Mixed Languages, in: Proceedings of FIRE, 2022.
 - [52] T. Chavan, S. Patankar, A. Kane¹, O. Gokhale, R. Joshi, A Twitter BERT Approach for Offensive Language Detection in Marathi, in: Proceedings of FIRE, 2022.
 - [53] S. Khanuja, D. Bansal, S. Mehtani, S. Khosla, A. Dey, B. Gopalan, D. K. Margam, P. Aggarwal, R. T. Nagipogu, S. Dave, et al., Muril: Multilingual Representations for Indian languages, *arXiv* (2021).
 - [54] S. Patankar, O. Gokhale, A. Kane, T. Chavan, R. Joshi, Spread Love Not Hate: Undermining the Importance of Hateful Pre-training for Hate Speech Detection, *arXiv* (2022).
 - [55] R. Joshi, L3Cube-MahaCorpus and MahaBERT: Marathi Monolingual Corpus, Marathi

- BERT Language Models, and Resources, in: Proceedings of the WILDRE-6, 2022.
- [56] M. L. Ripoll, F. Hussan, J. Attieh, G. Collell, A. Bouchekif, Multi-Lingual Contextual Hate Speech Detection Using Transformer-Based Ensembles, in: Proceedings of FIRE, 2022.
 - [57] Z. Chi, L. Dong, F. Wei, N. Yang, S. Singhal, W. Wang, X. Song, X.-L. Mao, H. Huang, M. Zhou, InfoXLM: An Information-Theoretic Framework for Cross-Lingual Language Model Pre-Training, in: Proceedings of NAACL, 2021.
 - [58] S. Kalra, K. Maheshwari, S. Goel, Y. Sharma, Hate Speech Detection for Marathi Language Using Transformer Based BERT Variants, in: Proceedings of FIRE, 2022.
 - [59] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: Proceedings of NAACL, 2019.
 - [60] V. Sanh, L. Debut, J. Chaumond, T. Wolf, DistilBERT, A Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter, arXiv (2019).
 - [61] G. Lemaître, F. Nogueira, C. K. Aridas, Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets in Machine Learning, *Journal of Machine Learning Research* 18 (2017) 1–5.
 - [62] S. Chanda, S. D. Sheth, S. Pal, Coarse and Fine-Grained Conversational Hate Speech and Offensive Content Identification in Code-Mixed Languages using Fine-Tuned Multilingual Embedding, in: Proceedings of FIRE, 2022.
 - [63] J. Howard, S. Gugger, Fastai: A Layered API for Deep Learning, *Information* 11 (2020) 108.
 - [64] K. Ghosh, A. Senapati, U. Garain, Baseline BERT models for Conversational Hate Speech Detection in Code-mixed Hinglish - German utilizing Data Augmentation and Offensive Language Identification in Marathi, in: Proceedings of FIRE, 2022.