

NLPineers@ NLU of Devanagari Script Languages 2025: Hate Speech Detection using Ensembling of BERT-based models

Anmol Guragain^{1*}, Nadika Poudel^{1*}, Rajesh Piryani², Bishesh Khanal¹

¹Nepal Applied Mathematics and Informatics Institute for Research (NAAMII), Nepal

²IRIT, Universite Toulouse III Paul Sabatier (UT3), Toulouse, France

Abstract

This paper explores hate speech detection in Devanagari-scripted languages, focusing on Hindi and Nepali, for Subtask B of the CHIP-SAL@COLING 2025 Shared Task. Using a range of transformer-based models such as XLM-RoBERTa, MURIL, and IndicBERT, we examined their effectiveness in navigating the nuanced boundary between hate speech and free expression. Our best performing model, implemented as ensemble of multilingual BERT models achieved Recall of 0.7762 (Rank 3/31 in terms of recall) and F1 score of 0.6914 (Rank 17/31). To address class imbalance, we used backtranslation for data augmentation, and cosine similarity to preserve label consistency after augmentation. This work emphasizes the need for hate speech detection in Devanagari-scripted languages and presents a foundation for further research. The code can be accessed at <https://github.com/Anmol2059/NLPineers>.

1 Introduction

Social media has become an essential part of our lives, empowering users to communicate freely and fostering a global exchange of ideas. However, it has contributed to the rapid proliferation of harmful content, including hate speech. Detecting hate speech is inherently complex due to the nuanced boundary between hate speech and legitimate free expression. What one individual may perceive as an offensive or harmful statement, another might interpret as a right to free speech, complicating the task of building an automated hate speech detection system. In languages where Devanagari script is predominantly used, such as Hindi, Marathi, and Nepali, detecting hate speech becomes even more intricate due to linguistic diversity, regional variations, and code-mixing practices.

Numerous transformer-based models have emerged to address the challenges of hate speech

detection across various high-resource languages. HateBERT (Caselli et al., 2020), a BERT model retrained on a dataset of Reddit comments from communities banned for offensive content, outperforms general BERT models in detecting abusive language in English. MC-BERT4HATE (Yang et al., 2020) presents a multi-channel architecture that integrates English, Chinese, and multilingual versions of BERT, aiming to detect hate speech across multiple languages more effectively. However, during politically charged events like elections, hate speech in Devanagari scripts intensifies on social media platforms like Twitter, producing more complex forms that require an understanding of socio-political dynamics beyond mere linguistic processing. These nuances are not well captured by general-purpose models, highlighting the need for specialized approaches.

The First Workshop on South East Asian Language Processing (Sarveswaran et al., 2025) aims to strengthen and spur NLP research and development in SEA languages. This paper aims to solve Task B: Hate Speech Detection of the Shared Task on Natural Language Understanding of Devanagari Script Languages (Thapa et al., 2025). Hate speech detection is a binary classification problem that requires determining whether a tweet is hate speech or not. The classifiers we used in this challenge include XLM-RoBERTa, MURIL, and IndicBERT.

2 Related Works

Devanagari-script languages, being low-resource, have seen relatively limited work in hate speech detection. Aggression and Misogyny Detection using BERT by (Safi Samghabadi et al., 2020) classified comments presents in English, Hindi and Bengali into one of the three aggression classes - Not Aggressive, Covertly Aggressive, Overtly Aggressive, as well as one of the two misogyny classes - Gendered and Non-Gendered scoring

*Equal contribution.

0.8579 weighted F1-measure using BERT model. In second workshop on Trolling, Aggression, and Cyberbullying (TRAC-2), (Baruah et al., 2020) work on Shared task on Misogynistic Aggression Identification achieved highest F1 score of 0.87 in Hindi language using XLMRoBERTa . Similarly, HASOC 2020: Hate Speech and Offensive Content Identification in Indo-European Languages (Mandl et al., 2020) had sub-task for Hate Speech detection in Hindi, German and English having 40 teams as participants. The best submission for Hindi used a CNN with fastText embeddings as input and the best result for English is based on a LSTM which used GloVe embeddings as input. Although there has been some work done on Hindi, it is worth noting that Nepali, which also uses the Devanagari script, has received relatively little attention in this area,(Luitel et al., 2024; Niraula et al., 2022) likely due to resource limitations. (Niraula et al., 2021) annotated 7462 records in Nepali language into four categories SEXIST, RACIST, OTHER-OFFENSIVE, and NON-OFFENSIVE using Random Forest Classifier achieving F1 scores as 0.01, 0.45,0.71 and 0.87 respectively.

Despite these few efforts, there remains a lack of performance benchmarks for multilingual BERT models on Devanagari scripts. Additionally, previous works have not explored BERT-based ensemble strategies that integrate predictions from multiple models. This gap motivated us to investigate the effectiveness of various multilingual BERT models and their ensembling approaches for hate speech detection in Devanagari scripted languages.

3 Dataset and Task

Category	Training	Evaluation
Hindi Non-Hate	7,376	1,596
Nepali Non-Hate	9,429	2,006
Hindi Hate	679	142
Nepali Hate	1,535	332
Total	19,019	4,076

Table 1: Sample distribution of data in training and evaluation sets

This experiment uses the data from the shared task, which is compiled from prior works across multiple Devanagari-script languages, including Hindi hate speech in political discourse (Jafri et al., 2024, 2023), Nepali election discourse (Thapa et al., 2023; Rauniyar et al., 2023), Bhojpuri-English sentiment modeling (Ojha, 2019), Marathi

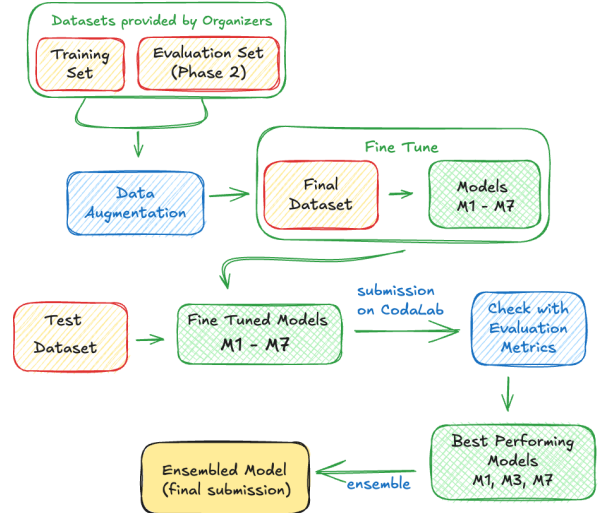


Figure 1: Experiment Workflow.

sentiment analysis (Kulkarni et al., 2021), and Sanskrit translation corpora (Aralikatte et al., 2021). In this study, the evaluation set refers to the phase 2 data provided by the challenge organizers, which is part of the development data. It is distinct from the test set, which was submitted for the challenge. Table 1 provides detailed statistics on the original dataset of the shared task and some of the example sentences are in Figure 2. The overall pipeline of this experiment is summarized on Figure 1.

4 Experimental Setup

4.1 Data Augmentation

As observed in Table 1, hate speech instances were much fewer compared to non-hate speech. To address this imbalance, data augmentation was applied to the hate speech instances using back-translation with the mBART-large-50 model (Tang et al., 2020), translating the data to English and back to the source language to introduce text variations. To ensure that the augmented data retained semantic similarity with the original data and minimize risk of unintended label changes, we calculated the cosine similarity between the embeddings of the augmented and original texts using the XLM-RoBERTa base model (Conneau et al., 2019). Only augmented data with a cosine similarity score greater than 0.9(chosen empirically) was added to the final dataset.

For the training set, data augmentation was performed on both Hindi and Nepali hate speech instances. In the evaluation set, however, augmentation was applied only to Hindi instances, as Nepali

	Model	Tokenizer / Embedding	Classifier Architecture
M1	MuRIL abusive (Das et al., 2022)	MuRIL (Hindi-Abusive) - Self	Native Head
M2	MuRIL + TabNet	MuRIL - Self	TabNet Classifier
M3	MuRIL(Khanuja et al., 2021a)	MuRIL - Self	Native Head
M4	IndicBERT(Kakwani et al., 2020)	IndicBERT - Self	Native Head
M5	IndicBERT + LSTM-CNN	IndicBERT - Self	LSTM + CNN + FC layer
M6	XLN-Roberta + Logistic Regression	XLN-Roberta - Self	Logistic Regression
M7	XLN-Roberta (Conneau et al., 2020)	XLN-Roberta - Self	Native Head
M8	FastText + LSTM	None - FastText (Hindi + Nepali)	LSTM + FC Layer

Table 2: Overview of models used. **FC** stands for Fully Connected layer, and **Native Head** refers to the model’s built-in classification head when imported from Hugging Face, indicating that these models are fine-tuned. The **Tokenizer/Embedding** column combines the tokenization method and embedding source; “Self” signifies that embeddings are generated by the model itself. Tokenizers and models are sourced from Hugging Face’s model hub, with FastText embeddings from FastText.cc.

data already had a higher representation compared to Hindi data. Additionally, to further address the class imbalance, all hate speech instances (label ‘1’) were duplicated to give the model more exposure to minority class. After augmentation, the training set grew to 13,695 instances by incorporating the original 2,214 training, 474 evaluation, and their augmented instances.

4.2 Pretrained Models

We used different pre-trained models and classifier heads, as observed in Table 2 to explore approaches that could better capture the nuances that exist in recognizing of hate speeches. The models included three BERT-based architectures—MuRIL (Khanuja et al., 2021b), XLN-RoBERTa (Conneau et al., 2020), and IndicBERT (Kakwani et al., 2020)—as well as FastText (Grave et al., 2018), an token embedding-based model. MuRIL is pre-trained on 17 Indian languages and their transliterations. XLN-RoBERTa, a large multilingual model, offers cross-lingual capabilities by training on diverse data from multiple languages. IndicBERT focuses on 12 Indian languages, including Devanagari(Hindi and Marathi), and uses a lightweight structure ideal for efficient processing. In contrast, FastText uses character-level n-grams to provide a detailed lexical representation, which is particularly beneficial for morphologically rich languages in Devanagari script. This combination allows us to leverage both deep contextual understanding and fine-grained lexical details for effective hate speech detection.

4.3 Ensemble Strategy

Our ensemble strategy leveraged the strengths of our top-performing models from Table 3. We chose M7 (XLN-Roberta) as the primary model,

M3 (MuRIL) as the secondary model, and M1 (MuRIL abusive) as the fallback model, based on each model’s unique strengths.

- **Primary Model (XLN-Roberta, Model 7):** XLN-Roberta achieved the highest recall (0.7381), making it effective at detecting hate speech and minimizing missed cases.
- **Secondary Model (MuRIL, Model 3):** When XLN-Roberta does not predict hate speech, MuRIL provides a balanced F1 score (0.6904) and accuracy (0.8744), acting as a secondary layer to catch potential cases missed by the primary model.
- **Fallback Model (MuRIL abusive, Model 1):** In cases where both primary and secondary models predict no hate speech, MuRIL abusive, with the highest precision (0.7572) and accuracy (0.8950), serves as a conservative fallback to minimize false positives.

$$\text{prediction}(x) = \begin{cases} 1 & \text{if } M_7(x) = 1 \\ 1 & \text{if } M_7(x) = 0 \text{ and } M_3(x) = 1 \\ M_1(x) & \text{otherwise} \end{cases}$$

4.4 Hyperparameters and Compute Environment

Training utilized the following hyperparameters, determined through iterative testing and practical constraints: a learning rate of 2e-5, a batch size of 16. These values were selected to balance model performance with available compute resources and processing time. We used NVIDIA GeForce RTX 3090 as compute environment.

	Model	Recall	Precision	F1 Score	Accuracy
M1	MuRIL abusive	0.6335	0.7572	0.6681	0.8950
M2	MuRIL + TabNet	0.6296	0.5874	0.5984	0.7927
M3	MuRIL	0.6877	0.6934	0.6904	0.8744
M4	IndicBERT	0.5934	0.5915	0.5924	0.8305
M5	IndicBERT + LSTM-CNN	0.6455	0.5618	0.5207	0.6271
M6	XLM-Roberta + Logistic Regression	0.6504	0.6596	0.6548	0.8619
M7	XLM-Roberta	0.7381	0.6696	0.6933	0.8472
M8	FastText + LSTM	0.5320	0.5400	0.5346	0.8270
	Ensemble (M1, M3, M7)	0.7762	0.6639	0.6914	0.8258

Table 3: Evaluation results on test set of the hate speech detection task. Dark green cells indicate the best performance in the respective metric, while dark red cells indicate the worst. Gradual shades of green represents relatively good performance.

5 Results and Discussions

The competition was hosted on the Codalab¹ platform by the organizers, where we submitted binary predictions (0 or 1) for evaluation based on recall, precision, F1 score, and accuracy. The performance of our models in the test set of challenge is shown in Table 3.

5.1 MuRIL-Based Models

As seen in Table 3, the fine-tuned MuRIL model (M1) on Devanagari script provided the highest accuracy among MuRIL-based models. The standard MuRIL model (M3) demonstrated balanced performance across all metrics. We also experimented with combining MuRIL and TabNet (M2) as suggested by (Chopra et al., 2023), but this configuration did not yield competitive results in this task.

5.2 IndicBERT

We anticipated strong performance from IndicBERT (M4) due to its training on 12 Indic languages as discussed in (Kakwani et al., 2020). However, its results were lower than expected, possibly due to the presence of Nepali text in the dataset, which IndicBERT may not be optimized for. The combination of IndicBERT with LSTM-CNN also underperformed, showing unsatisfactory results.

5.3 XLM-Roberta-Based Models

Both the plain XLM-Roberta model (M7) and XLM-Roberta with a Logistic Regression head

(M6) performed well, indicating the model’s robust generalization capabilities across different metrics. This highlights XLM-Roberta’s versatility in multilingual tasks.

5.4 FastText with LSTM

Since the evaluation set contained Devanagari script for both Nepali and Hindi, we utilized FastText embeddings of both languages and fed in LSTM based classifier (M8). However, this setup did not yield satisfactory results, likely due to the limitations of static embeddings, which struggle to capture the contextual nuances essential for accurate hate speech detection.

5.5 Ensembled Model

Our final submission was an ensemble model combining M1, M3, and M7, as described in previous sections. This ensemble achieved balanced performance, with recall, precision, F1 score, and accuracy of 0.7762, 0.6639, 0.6914, and 0.8258, respectively, effectively leveraging the strengths of multiple models.

6 Conclusion

This study highlights the potential of various BERT-based models and ensembling approach for hate speech detection in Devanagari-scripted languages, with future work planned on model robustness and scalability for real-world applications. Further research could explore additional embeddings and augmentations to enhance performance across multilingual contexts.

¹https://codalab.lisn.upsaclay.fr/competitions/20000#participate-submit_results

Limitations

This study faces several limitations, particularly due to the linguistic complexities inherent to Devanagari-scripted languages like Hindi and Nepali. Below, we outline some of the primary challenges:

- **Limitations of Data Augmentation via Backtranslation:** While backtranslation with the mBART model was used to augment hate speech samples, this approach sometimes loses the cultural nuances or tone intended in the original text. For instance, words like *tapai* and *hajar* in Nepali convey a formal or respectful tone, but during translation to English and back to Nepali, these terms are often reduced to the informal *timi*, altering the sentiment. This limitation could introduce subtle inaccuracies during model training.
- **Contextual Meaning Across Languages:** In Devanagari-scripted languages, certain words can carry vastly different meanings depending on the language context. Such linguistic ambiguities create challenges for the model, as it may misinterpret hate speech in cases where meanings differ across languages using the same script.
- **Dependency on Word Embeddings for Devanagari Script:** Devanagari script is used for multiple languages, and words in Hindi and Nepali may have similar or identical representations in embeddings, potentially leading to confusion. While BERT-based models like XLM-RoBERTa and MuRIL are designed to handle multilingual contexts, challenges persist when languages share the same script but differ in vocabulary or syntax. These issues may impact the model’s ability to differentiate nuanced expressions unique to each language.

Acknowledgements

We thank CHIPSAL@COLING organizers for providing annotated datasets and hosting the shared task. We also thank NepAI Applied Mathematics and Informatics Institute for research (NAAMII) for giving us access to training resources required for the sub-task.

References

- Rahul Aralikkatte, Miryam De Lhoneux, Anoop Kunchukuttan, and Anders Søgaard. 2021. Itihasa: A large-scale corpus for sanskrit to english translation. In *Proceedings of the 8th Workshop on Asian Translation (WAT2021)*, pages 191–197.
- Arup Baruah, Kaushik Das, Ferdous Barbhuiya, and Kuntal Dey. 2020. *Aggression identification in English, Hindi and Bangla text using BERT, RoBERTa and SVM*. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 76–82, Marseille, France. European Language Resources Association (ELRA).
- Tommaso Caselli, Valerio Basile, Jelena Mitrovic, and Michael Granitzer. 2020. *Hatebert: Retraining BERT for abusive language detection in english*. *CoRR*, abs/2010.12472.
- Abhishek Chopra, Deepak Kumar Sharma, Aashna Jha, and Uttam Ghosh. 2023. A framework for online hate speech detection on code-mixed hindi-english text and hindi text in devanagari. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 22(5):1–21.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. *Unsupervised cross-lingual representation learning at scale*. *CoRR*, abs/1911.02116.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. *Preprint*, arXiv:1911.02116.
- Mithun Das, Somnath Banerjee, and Animesh Mukherjee. 2022. Data bootstrapping approaches to improve low resource abusive language detection for indic languages. *arXiv preprint arXiv:2204.12543*.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Farhan Ahmad Jafri, Kritesh Rauniyar, Surendrabikram Thapa, Mohammad Aman Siddiqui, Matloob Khushi, and Usman Naseem. 2024. *Chunav: Analyzing hindi hate speech and targeted groups in indian election discourse*. *ACM Transactions on Asian and Low-Resource Language Information Processing*.
- Farhan Ahmad Jafri, Mohammad Aman Siddiqui, Surendrabikram Thapa, Kritesh Rauniyar, Usman Naseem, and Imran Razzak. 2023. *Uncovering political hate speech during indian election campaign: A new low-resource dataset and baselines*.

- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. IndicNLP Suite: Monolingual Corpora, Evaluation Benchmarks and Pre-trained Multilingual Language Models for Indian Languages. In *Findings of EMNLP*.
- Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, Shruti Gupta, Subhash Chandra Bose Gali, Vish Subramanian, and Partha Talukdar. 2021a. [Muril: Multilingual representations for indian languages](#). *Preprint*, arXiv:2103.10730.
- Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, Shruti Gupta, Subhash Chandra Bose Gali, Vish Subramanian, and Partha P. Talukdar. 2021b. [Muril: Multilingual representations for indian languages](#). *CoRR*, abs/2103.10730.
- Atharva Kulkarni, Meet Mandhane, Manali Likhitkar, Gayatri Kshirsagar, and Raviraj Joshi. 2021. L3cubemahasent: A marathi tweet-based sentiment analysis dataset. In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 213–220.
- Nishant Luitel, Nirajan Bekoju, Anand Kumar Sah, and Subarna Shakya. 2024. Can perplexity predict fine-tuning performance? an investigation of tokenization effects on sequential language models for nepali. *arXiv preprint arXiv:2404.18071*.
- Thomas Mandl, Sandip Modha, Marcos Zampieri, et al. 2020. Hasoc 2020: Hate speech and offensive content identification in indo-european languages. In *Proceedings of the 12th Language Resources and Evaluation Conference (LREC)*, pages 190–199.
- Nobal B. Niraula, Saurab Dulal, and Diwa Koirala. 2021. [Offensive language detection in Nepali social media](#). In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 67–75, Online. Association for Computational Linguistics.
- Nobal B Niraula, Saurab Dulal, and Diwa Koirala. 2022. Linguistic taboos and euphemisms in nepali. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 21(6):1–26.
- Atul Kr Ojha. 2019. English-bhojpuri smt system: Insights from the karaka model. *arXiv preprint arXiv:1905.02239*.
- Kritesh Rauniyar, Sweta Poudel, Shuvam Shiwakoti, Surendrabikram Thapa, Junaid Rashid, Jungeun Kim, Muhammad Imran, and Usman Naseem. 2023. Multi-aspect annotation and analysis of nepali tweets on anti-establishment election discourse. *IEEE Access*.
- Niloofer Safi Samghabadi, Parth Patwa, Srinivas PYKL, Prerana Mukherjee, Amitava Das, and Thamar Solorio. 2020. [Aggression and misogyny detection using BERT: A multi-task approach](#). In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 126–131, Marseille, France. European Language Resources Association (ELRA).
- Kengatharaiyer Sarveswaran, Bal Krishna Bal, Surendrabikram Thapa, Ashwini Vaidya, and Sana Shams. 2025. A brief overview of the first workshop on challenges in processing south asian languages (chipsal). In *Proceedings of the First Workshop on Challenges in Processing South Asian Languages (CHIPSAL)*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#). *CoRR*, abs/2008.00401.
- Surendrabikram Thapa, Kritesh Rauniyar, Farhan Ahmad Jafri, Surabhi Adhikari, Kengatharaiyer Sarveswaran, Bal Krishna Bal, Hariram Veeramani, and Usman Naseem. 2025. Natural language understanding of devanagari script languages: Language identification, hate speech and its target detection. In *Proceedings of the First Workshop on Challenges in Processing South Asian Languages (CHIPSAL)*.
- Surendrabikram Thapa, Kritesh Rauniyar, Shuvam Shiwakoti, Sweta Poudel, Usman Naseem, and Mehwish Nasim. 2023. Nehate: Large-scale annotated data shedding light on hate speech in nepali local election discourse. In *ECAI 2023*, pages 2346–2353. IOS Press.
- Ziqing Yang et al. 2020. Mc-bert4hate: Hate speech detection using multi-channel bert for different languages and translations. In *Proceedings of the 12th Language Resources and Evaluation Conference (LREC)*, pages 2159–2166. European Language Resources Association (ELRA).

A Appendix

A.1 Example Sentences from Dataset

Example Data		
Language	Label	Sentences
Training		
Nepali	Hate	पोखराका जनता सवै भन्दा भेडा लाग्यो मलाई । हामिले काठमाण्डैमा बलेनको लैरोमा सपोर्ट गरेपनि हामोर छेत्रमा स्वतन्त्रलाई सारै हेपेउँ । पार्टिगतलाई तर्साउन भएपनि केही हद सम्म स्वतन्त्रलाई भोट जानु परथ्यो ।
Nepali	Non-Hate	नया आउनै पर्छ। चाहे त्यो राष्ट्रको रुपमा होस, रा.स्व.प को रुपमा होस अथवा बिबेकशिल साझा। कहिले होला? #nonotagain #NepalVotes2022
Hindi	Hate	राहुल गांधी ने दावा करते हुए कहा कि सरकार बनेगी तो कांग्रेस पार्टी की ही बनेगी. इस दौरान राहुल ने BJP पर भी निशाना साधा है. #GoaElections2022 #RahulGandhi #AssemblyElections2022 https://t.co/crw8q959wg
Hindi	Non-Hate	अमर उजाला दैनिक समाचार में हेड लाइन में प्रकाशित "मुरदे से भी चुनाव में अशांति फैलाने का डर" के सम्बंध में खंडन। #UPPolice #AssemblyElections2022 #YourVoteMatters #AgelessDemocracy
Evaluation		
Nepali	Hate	तपाईंको मुकुण्डो उधुन धेरै दिन बाकी छैन सर ।मातैर क्रसर वालाहरुको काम फुकका भएको दिन ।तपाईंको चर्चाको राजनीति सिधिन्छ। अब उचित कानुन र नियमन वा हजुरको मुकुण्डो। No options for you. Time is running
Nepali	Non-Hate	हारने डर त्यसपछि राज्यकोषबाट रकम दुरुपयोग गर्न नपाइने भयले रवि लामिछाने माथि आक्रमणको प्रयास गरेको हुन गठबन्धनले । तोडफोड गर्दा प्रहरी र प्रशासन मुकदरशक बननु उनिहरुलाई समर्थन गर्नु हो । #NoNotAgain
Hindi	Hate	योगी ने कहा कि कांग्रेस पहले जब सत्ता में थी तो आतंकियों को प्रेरित और पोर्त्साहित करती थी और ये हिंदू संगठनो पर झूठे मुकदमें दर्ज़ करते थे। https://t.co/ispuqj8BM7 #YogiAdityanath #UttarPradesh #UttarPradeshElections2022 #BJP #Congress https://t.co/jQPWPPNPZVD
Hindi	Non-Hate	उत्तर प्रदेश विधानसभा सामान्य निर्वाचन-2022 कल 10 मार्च-2022 को मतगणना के लिए व्यवस्थाओं के सम्बन्ध में प्रेस विज्ञप्ति जारी... @ECISVEEP @SpokespersonECI #ECI #AssemblyElections2022 https://t.co/vyC5hmeD4H

Figure 2: Examples from the dataset.