

OffMix-3L: A Novel Code-Mixed Dataset in Bangla-English-Hindi for Offensive Language Identification

Dhiman Goswami*, Md Nishat Raihan*, Antara Mahmud

Antonios Anastasopoulos, Marcos Zampieri

George Mason University, USA

{dgoswam, mraihan2, amahmud4, antonis, mzampier}@gmu.edu

Abstract

Code-mixing is a well-studied linguistic phenomenon when two or more languages are mixed in text or speech. Several works have been conducted on building datasets and performing downstream NLP tasks on code-mixed data. Although it is not uncommon to observe code-mixing of three or more languages, most available datasets in this domain contain code-mixed data from only two languages. In this paper, we introduce OffMix-3L, a novel offensive language identification dataset containing code-mixed data from three different languages. We experiment with several models on this dataset and observe that BanglishBERT outperforms other transformer-based models and GPT-3.5.

1 Introduction

Code-mixing and code-switching are common linguistic phenomena observed both in speech and text form. While the two terms are often used interchangeably, code-mixing is defined as the use of words or morphemes from multiple languages within a single utterance, sentence, or discourse whereas code-switching refers to the deliberate alternation between multiple languages within the same context (Thara and Poornachandran, 2018). The first case is often spontaneous while the second case is purposeful. However, both are widely observed in bilingual and multilingual communities.

As described in Anastassiou (2017), several social, linguistic, and cognitive factors are behind these two phenomena. Socially, this often serves as a sign of group identity which allows individuals to navigate multiple social and cultural affiliations. Linguistically, it is common for bilingual speakers to not be able to find a word for a specific concept in one language thus using a word from another language to help communication. Additionally, there

*These two authors contributed equally to this work.

WARNING: This paper contains examples that are offensive in nature.

are several cases even in monolingual communities, when code-mixing might be the convenient way to express a concept as in the case of English loan words such as *feedback* used in various languages.

Most commonly, code-mixing is a bilingual phenomenon. Byers-Heinlein and Lew-Williams (2013), for example, estimates that by the year 2035, over half of children enrolled in kindergarten in California will have grown up speaking a language other than English. Another study conducted by Jeffery and van Beuningen (2020) shows that bilingualism is a common practice in European countries such as Germany and Spain. However, in cosmopolitan cities and areas like New York, London, Singapore, and others, code-mixing with three or more languages is fairly common. This is also observed in countries like Luxembourg, and regions such as West Bengal, and South-East India where more than two languages are commonly used on a daily basis.

Several papers have presented code-mixed datasets for various NLP tasks (KhudaBukhsh et al., 2020; Krishnan et al., 2022). However, most of these datasets are bilingual leaving the processing of code-mixing in three or more languages largely unexplored. In this paper, we present a Bangla-Hindi-English dataset annotated for offensive language identification. To the best of our knowledge, this is one of the first datasets to contain code-mixing between more than two languages.

The main contributions of this paper are as follows:

- We introduce OffMix-3L, a novel three-language code-mixed test dataset in Bangla-Hindi-English for offensive language identification. OffMix-3L contains 1,001 instances annotated by speakers of the three languages. We made OffMix-3L freely available to the community.¹

¹<https://github.com/LanguageTechnologyLab/OffMix-3L>

- We provide a comprehensive evaluation of several monolingual, bilingual, and multilingual models on OffMix-3L.

We present OffMix-3L exclusively as a test set due to the unique and specialized nature of the task. The size of the dataset, while limiting for training purposes, offers a high-quality testing environment with gold-standard labels that will serve as a benchmark in this domain. Given the scarcity of similar datasets and the challenges associated with data collection, OffMix-3L provides an important resource for the rigorous evaluation of offensive language identification models, filling a critical gap in multi-level code-mixing research.

2 Related Work

There have been several studies describing Bangla-English (Alam, 2006; Hasan et al., 2015; Hossain and Bar, 2015; Begum and Haque, 2013; Mahbubul Alam and Quyyum, 2016), Hindi-English (Singh, 1985; Bali et al., 2014; Thara and Poornachandran, 2018) and Bangla-Hindi (Ali et al., 2019; Jose et al., 2020) code-mixing and code-switching separately. Code-mixing between these three languages has also been studied separately in NLP.

There have been few studies conducted on offensive language identification for Bangla-English code-mixed data. The work by Jahan et al. (2019) focused on detecting Bangla-English code-mixed and transliterated offensive comments on Facebook. Another Bangla-English dataset is gathered by Wadud et al. (2021), where they collected 2,200 instances.

Comparatively more work has been carried out for Hindi-English Code-mixing. Sreelakshmi et al. (2020) uses fastText (Joulin et al., 2016) to represent 10,000 instances collected from different sources. Other offensive language datasets collected from Facebook and Twitter were introduced by Bohra et al. (2018); Kumar et al. (2018); Rani et al. (2020). Mundra and Mittal (2022) proposes Fused Attention-based Network (FA-Net), which introduces a fusion of attention mechanism of collective and mutual learning between local and sequential features for Hindi-English offensive language and hate speech classification. Gupta et al. (2021) uses Character Level Embeddings, GRU, and attention layer to offensive language identification in Hindi-English code-mixed.

To the best of our knowledge, no existing work focuses specifically on Hindi-Bangla code-mixing.

However, some studies focused on multiple Indian languages code-mixing altogether including Bangla and Hindi. The work by Vasantharajan and Thayasivam (2021) focuses on offensive language identification in Dravidian languages. A few similar works include Ravikiran and Annamalai (2021); Sai and Sharma (2020); Kumar et al. (2020); Kumari and Singh (2020); Ranasinghe and Zampieri (2021).

In summary, to the best of our knowledge, there has been no work on offensive language identification for code-mixed Bangla-English-Hindi. There have also been no offensive language datasets made available for these three languages. OffMix-3L fills this gap by providing the community with a novel resource for these three languages. OffMix-3L provides the community with the opportunity to evaluate how state-of-the-art models perform on Bangla-English-Hindi.

3 The OffMix-3L Dataset

We choose a controlled data collection method, asking the volunteers to freely contribute data in Bangla, English, and Hindi. This decision stems from several challenges of extracting such specific code-mixed data from social media and other online platforms. Our approach ensures data quality and sidesteps the ethical concerns associated with using publicly available online data. Such types of datasets are often used when it is difficult to mine them from existing corpora. As examples, for fine-tuning LLMs on instructions and conversations, semi-natural datasets like Databricks (2023) and Nie (2023) have become popular.

Data Collection A group of 10 undergraduate students fluent in the three languages was asked to prepare 250 to 300 social media posts each. They were allowed to use any language including Bangla, English, and Hindi to prepare posts on several daily topics like politics, sports, education, social media rumors, etc. We also ask them to switch languages if and wherever they feel comfortable doing so. The inclusion of emojis, hashtags, and transliteration was also encouraged. The students had the flexibility to prepare the data as naturally as possible. Upon completion of this stage, we gathered 1,734 samples that contained at least one word or sub-word from each of the three languages using langdetect (Mazzocchi, 2012) an open-sourced Python tool for language identification.

Data Annotation We annotate the dataset in two steps. Firstly, we recruited three students from social science, computer science, and linguistics fluent in the three languages to serve as annotators. They annotated all 1,734 samples with one of the two labels (Non-Offensive and Offensive) with a raw agreement of 63.7%. We then take 1,106 instances, where all three annotators agree on the labels, and use them in a second step. To further ensure high-quality annotation, we recruit a second group of annotators consisting of two NLP researchers fluent in the three languages. After their annotation, we calculate a raw agreement of 91% (Kvålseth, 1989), a Cohen Kappa score of 0.82. After the two stages, we only keep the instances where both annotators agree, and we end up with a total of 1,001 instances. The label distribution is shown in Table 1.

Label	No. of Data	Percentage
Non-Offensive	522	52.15%
Offensive	479	47.85%
Total	1,001	100%

Table 1: Label distribution in OffMix-3L

Dataset Statistics A detailed description of the dataset statistics is provided in Table 2. Since the dataset was generated by people whose first language is Bangla, we observe that the majority of tokens in the dataset are in Bangla. There are several *Other* tokens in the dataset that are not from Bangla, English, or Hindi language. The *Other* tokens in the dataset primarily contain transliterated words as well as emojis and hashtags. Also, there are several misspelled words that have been classified as *Other* tokens too.

	All	Bangla	English	Hindi	Other
Tokens	87,190	31,228	6,690	14,694	34,578
Types	18,787	7,714	1,135	1,413	8,645
Avg	87.10	31.20	6.68	14.68	34.54
Std Dev	20.58	8.60	3.05	5.74	10.98

Table 2: OffMix-3L Data Card. The row Avg represents the average number of tokens with its standard deviation in row *Std Dev*.

Synthetic Train and Development Set We present OffMix-3L as a test dataset and we build a synthetic train and development set that contains Code-mixing for Bangla, English, and Hindi. We use two English training datasets annotated

with the same labels as OffMix-3L, namely OLID (Zampieri et al., 2019) and SOLID (Rosenthal et al., 2021). We randomly select 100,000 data instances randomly and we carefully choose an equal number of Non-Offensive and Offensive instances. We then use the *Random Code-mixing Algorithm* (Krishnan et al., 2022) and *r-CM* (Santy et al., 2021) to generate the synthetic Code-mixed dataset.

4 Experiments

Monolingual Models We use seven monolingual models for these experiments, five general models, and two task fine-tuned ones. The five monolingual models are DistilBERT (Sanh et al., 2019), BERT (Devlin et al., 2019), BanglaBERT (Kowsher et al., 2022), roBERTa (Liu et al., 2019), HindiBERT (Nick Doiron, 2023). BanglaBERT is trained in only Bangla and HindiBERT on only Hindi while DistilBERT, BERT, and roBERTa are trained in English only. Finally, the two English task fine-tuned models we use are HateBERT and fBERT (Caselli et al., 2021; Sarkar et al., 2021).

Bilingual Models BanglishBERT (Bhattacharjee et al., 2022) and HingBERT (Nayak and Joshi, 2022) are used as bilingual models as they are trained on both Bangla-English and Hindi-English respectively.

Multilingual Models We use mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) as multilingual models which are respectively trained on 104 and 100 languages including Bangla-English-Hindi. We also use IndicBERT (Kakwani et al., 2020) and MuRIL (Khanuja et al., 2021) which cover 12 and 17 Indian languages, respectively, including Bangla-English-Hindi. We also perform hyper-parameter tuning while using all the models to prevent overfitting.

Prompting We use prompting with GPT-3.5-turbo model (OpenAI, 2023) from OpenAI for this task. We use the API for zero-shot prompting (see Figure 1) and ask the model to label the test set.

Additionally, we run the same experiments separately on synthetic and natural datasets splitting both in a 60-20-20 way for training, evaluating, and testing purposes.

5 Results

In this experiment, synthetic data is used as a training set and natural data is used as the test set. The

Role: "You are a helpful AI assistant. You are given the task of offensive text classification."

Definition: An offensive text is something that upsets or embarrasses people because it is rude or insulting. You will be given a text to label either 'Offensive' or 'Non-Offensive'.

Task: Generate the label for this "text" in the following format: <label> Your_Predicted_Label <\label>. Thanks."

Figure 1: Sample GPT-3.5 prompt.

F1 scores of monolingual models range from 0.43 to 0.66 where BERT performs the best. mBERT is the best of all the multilingual models with an F1 score of 0.63. Besides, a zero-shot prompting technique on GPT 3.5 turbo provides a 0.57 weighted F1 score. The best task fine-tuned model is HateBERT with the F1 score of 0.60. Among all the models BanglishBERT scores 0.68 which is the best achieved F1 score. These results are available in Table 3.

Models	F1 Score
BanglishBERT	0.68
BERT	0.66
mBERT	0.63
HingBERT	0.60
MuRIL	0.60
HateBERT	0.60
fBERT	0.58
roBERTa	0.58
XLm-R	0.57
DistilBERT	0.57
GPT 3.5 Turbo	0.57
BanglaBERT	0.54
IndicBERT	0.55
HindiBERT	0.43

Table 3: Weighted F-1 score for different models: training on synthetic and tested on natural data (OffMix-3L).

We perform the same experiment using synthetic data for training and testing. We present results in Table 4. Here, mBERT and XLM-R with 0.88 F1 scores are the best-performing models.

6 Error Analysis

We observe *Other* tokens in almost 39% of the whole dataset, as shown in Table 2. These tokens occur due to transliteration which poses a challenge

Models	Weighted F1 Score
XLm-R	0.88
mBERT	0.88
BanglishBERT	0.86
BERT	0.83
HingBERT	0.82
IndicBERT	0.82
MuRIL	0.81
fBERT	0.81
HateBERT	0.81
roBERTa	0.80
DistilBERT	0.79
BanglaBERT	0.76
HindiBERT	0.73

Table 4: Weighted F-1 score for different models: training on synthetic and tested on synthetic data.

for most of the models since not all of the models are pre-trained on transliterated tokens. BanglishBERT did well since it recognizes both Bangla and English tokens. However, the total number of tokens for Hindi-English is less than Bangla-English tokens, justifying HingBERT’s inferior performance compared to BanglishBERT (see Table 3). Also, misspelled words and typos are also observed in the datasets, which are, for the most part, unknown tokens for the models, making the task even more difficult. Some examples are available in Appendix A which are classified wrongly by all the models.

7 Conclusion and Future Work

In this paper, we presented OffMix-3L, a Bangla-English-Hindi code-mixed offensive language identification dataset containing 1,001 instances. We also created 100,000 synthetic data in the same three languages for training. We evaluated various monolingual models on these two datasets. Our results show that when training on synthetic data and testing on OffMix-3L, BanglishBERT performs the best. When using synthetic data for both training and testing, multilingual models such as mBERT and XLM-R perform well. In the future, we would like to expand OffMix-3L so that it can serve as both training and testing data. Additionally, we are working on pre-training Bangla-English-Hindi trilingual code-mixing models for offensive language identification.

Acknowledgments

We thank the annotators who helped us with the annotation of OffMix-3L. We further thank the anonymous workshop reviewers for their insightful feedback. Antonios Anastasopoulos is generously supported by NSF award IIS-2125466.

Limitations

Although most datasets for the downstream tasks are scraped from social media posts in the real world, in our case these data instances are generated in a semi-natural manner, meaning that they were generated by people but not scraped from social media directly. This was done due to the complexity of extracting contents that contain a specific 3 language code-mixing in them. Also, the dataset is comparatively smaller in size, since it is costly to generate data by a specific set of people who are fluent in all 3 target languages.

Ethics Statement

The dataset introduced in this paper, which centers on the analysis of offensive language in Bangla-English-Hindi code-mixed text, adheres to the [ACL Ethics Policy](#) and seeks to make a valuable contribution to the realm of online safety. OffMix-3L serves as an important resource for the moderation of online content, which will make it easier to create safer digital environments. Moreover, the contributors and annotators of the dataset are paid respectable remuneration and also attended two sessions with psychologist before starting and after completing the work to ensure their mental health is not compromised throughout the course of this dataset preparation.

References

- Suraiya Alam. 2006. Code-mixing in bangladesh: A case study of non-government white-collar service holders and professionals. *Asian affairs*, 28(4):52–70.
- Muhammad Mooneeb Ali, Mazhar Iqbal Ranjha, and Sartaj Fakhar Jillani. 2019. Effects of code mixing in indian film songs. *Journal of Media Studies*, 31(2).
- Fotini Anastassiou. 2017. Factors associated with the code mixing and code-switching of multilingual children: An overview. *International Journal of Linguistics, Literature and Culture*, 4(3):13–26.
- Kalika Bali, Jatin Sharma, Monojit Choudhury, and Yogarshi Vyas. 2014. i am borrowing ya mixing?" an analysis of english-hindi code mixing in facebook. In *Proceedings of CodeSwitch*.
- Most Tasnim Begum and Md Mahmudul Haque. 2013. Code mixing in the ksa: A case study of expatriate bangladeshi and indian esl teachers. *Arab World English Journal*, 4(4).
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Ahmad, Kazi Samin Mubasshir, Md Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022. BanglaBERT: Language model pretraining and benchmarks for low-resource language understanding evaluation in Bangla. In *Findings of the ACL*.
- Aditya Bohra, Deepanshu Vijay, Vinay Singh, Syed Sarfaraz Akhtar, and Manish Shrivastava. 2018. A dataset of hindi-english code-mixed social media text for hate speech detection. In *Proceedings of PEO-PLES*.
- Krista Byers-Heinlein and Casey Lew-Williams. 2013. Bilingualism in the early years: What the science says. *LEARNing landscapes*, 7(1):95.
- Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. HateBERT: Retraining BERT for abusive language detection in English. In *Proceedings of WOAHA*.
- 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. In *Proceedings of ACL*.
- Databricks. 2023. [Dolly 2.0: An open source, instruction-following large language model](#). Accessed: 2023-09-10.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*.
- Vasu Gupta, Vibhu Sehra, Yashaswi Raj Vardhan, et al. 2021. Hindi-english code mixed hate speech detection using character level embeddings. In *Proceedings of ICCMC*.
- Md Kamrul Hasan, Mohd Akhand, et al. 2015. Reviewing the challenges and opportunities presented by code switching and mixing in bangla. *Journal of Education and Practice*, 6(1):103–109.
- Didar Hossain and Kapil Bar. 2015. A case study in code-mixing among jahangirnagar university students. *International Journal of English and Literature*, 6(7):123–139.
- Maliha Jahan, Istiak Ahamed, Md Rayanuzzaman Bishwas, and Swakkhar Shatabda. 2019. Abusive comments detection in bangla-english code-mixed and transliterated text. In *Proceedings of ICIET*.

- Jill V Jeffery and Catherine van Beuningen. 2020. Language education in the eu and the us: Paradoxes and parallels. *Prospects*, 48(3-4):175–191.
- Navya Jose, Bharathi Raja Chakravarthi, Shardul Suryawanshi, Elizabeth Sherly, and John P McCrae. 2020. A survey of current datasets for code-switching research. In *Proceedings of ICACCS*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hervé Jégou, and Tomáš Mikolov. 2016. Fasttext.zip: Compressing text classification models. *CoRR*, abs/1612.03651.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, NC Gokul, 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 the ACL*.
- Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, et al. 2021. Muril: Multilingual representations for indian languages. *arXiv preprint arXiv:2103.10730*.
- Ashiqur R KhudaBukhsh, Shriphani Palakodety, and Jaime G Carbonell. 2020. Harnessing code switching to transcend the linguistic barrier. In *Proceedings of IJCAI*.
- M Kowsher, Abdullah As Sami, Nusrat Jahan Protasha, Mohammad Shamsul Arefin, Pranab Kumar Dhar, and Takeshi Koshiba. 2022. Bangla-bert: transformer-based efficient model for transfer learning and language understanding. *IEEE Access*, 10:91855–91870.
- Jitin Krishnan, Antonios Anastasopoulos, Hemant Purohit, and Huzefa Rangwala. 2022. Cross-lingual text classification of transliterated hindi and malayalam. In *Proceedings of Big Data*.
- Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2020. Evaluating aggression identification in social media. In *Proceedings of TRAC*.
- Ritesh Kumar, Aishwarya N. Reganti, Akshit Bhatta, and Tushar Maheshwari. 2018. Aggression-annotated Corpus of Hindi-English Code-mixed Data. In *Proceedings of LREC*.
- Kirti Kumari and Jyoti Prakash Singh. 2020. Ai_ml_nit_patna@ trac-2: Deep learning approach for multi-lingual aggression identification. In *Proceedings of TRAC*.
- Tarald O Kvålseth. 1989. Note on cohen’s kappa. *Psychological reports*, 65(1):223–226.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Ahmad Mahbub-ul Alam and Shaima Quyyum. 2016. A sociolinguistic survey on code switching & code mixing by the native speakers of bangladesh. *Journal of Manarat International University*, 6(1):8–9.
- Daniele Mazzocchi. 2012. *langdetect: Language detection library*. Python library.
- Shikha Mundra and Namita Mittal. 2022. Fa-net: fused attention-based network for hindi english code-mixed offensive text classification. *Social Network Analysis and Mining*, 12(1):100.
- Ravindra Nayak and Raviraj Joshi. 2022. L3Cube-HingCorpus and HingBERT: A code mixed Hindi-English dataset and BERT language models. In *Proceedings of WILDRE*.
- Nick Doiron. 2023. *hindi-bert*. Accessed: 2023-09-10.
- Jianzhi Nie. 2023. *Awesome instruction datasets*. Accessed: 2023-09-10.
- OpenAI. 2023. *Gpt-3.5 turbo fine-tuning and api updates*. Accessed: 2023-08-28.
- Tharindu Ranasinghe and Marcos Zampieri. 2021. An evaluation of multilingual offensive language identification methods for the languages of india. *Information*, 12(8):306.
- Priya Rani, Shardul Suryawanshi, Koustava Goswami, Bharathi Raja Chakravarthi, Theodorus Fransen, and John Philip McCrae. 2020. A comparative study of different state-of-the-art hate speech detection methods in hindi-english code-mixed data. In *Proceedings of TRAC*.
- Manikandan Ravikiran and Subbiah Annamalai. 2021. Dosa: Dravidian code-mixed offensive span identification dataset. In *Proceedings of TRAC*.
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2021. SOLID: A large-scale semi-supervised dataset for offensive language identification. In *Findings of the ACL*.
- Siva Sai and Yashvardhan Sharma. 2020. Siva@ hasoc-dravidian-codemix-fire-2020: Multilingual offensive speech detection in code-mixed and romanized text. In *Proceedings of FIRE*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. In *Proceedings of EMC2*.
- Sebastin Santy, Anirudh Srinivasan, and Monojit Choudhury. 2021. Bertologicomix: How does code-mixing interact with multilingual bert? In *Proceedings of AdaptNLP*.
- Diptanu Sarkar, Marcos Zampieri, Tharindu Ranasinghe, and Alexander Ororbia. 2021. fBERT: A neural transformer for identifying offensive content. In *Findings of the ACL*.

Rajendra Singh. 1985. Grammatical constraints on code-mixing: Evidence from hindi-english. *Canadian Journal of Linguistics/Revue canadienne de linguistique*, 30(1):33–45.

K Sreelakshmi, B Premjith, and KP Soman. 2020. Detection of hate speech text in hindi-english code-mixed data. *Procedia Computer Science*, 171:737–744.

S Thara and Prabakaran Poornachandran. 2018. Code-mixing: A brief survey. In *Proceedings of ICACCI*.

Charangan Vasantharajan and Uthayasanker Thayasivam. 2021. Hypers@ dravidianlangtech-eacl2021: Offensive language identification in dravidian code-mixed youtube comments and posts. In *Proceedings of DravidianLangTech*.

Md Anwar Hussen Wadud, Md Abdul Hamid, Muhammad Mostafa Monowar, and Atif Alamri. 2021. L-boost: Identifying offensive texts from social media post in bengali. *Ieee Access*, 9:164681–164699.

Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. SemEval-2019 task 6: Identifying and categorizing offensive language in social media (OffensEval). In *Proceedings of SemEval*.

A Examples of Misclassified Instances

Non-Offensive: আই এম এক্সসাইটেড টু পার্টিসিপেট টু এ বুক ক্লাব টু শেয়ার মাই আইডিয়াস! এটি amazing যে কীভাবে একটি book মানুষকে একত্রিত করতে পারে এবং meaningful কথাপেকথন শুরু করতে পারে। I am looking forward to hearing ভিন্নধর্মী opinions and অনুধাবন of the book, and learning from others यह मेरी पठन सूची का विस्तार करने और नए लेखकों और शैलियों को खोजने का भी एक शानदार तरीका है। mera maanana hai ki padhana seekhane aur aatm-khoj kee kabhee na khatm hone vaalee yaatra hai #bookclub #discussiongroup

Offensive: The proliferation of young people creating adult content on social media is a reflection of the harmful and unrealistic messages we are sending to our youth about যৌনতা। हमें इन मानदंडों को चुनौती देनी चाहिए और कामुकता के अधिक सशक्त और समावेशी दृष्टिकोण को बढ़ावा देना चाहिए fantasy जगत থেকে তাদের বাস্তবে আসা উচিত har din tiktok karne ka natija #obscence #boycottAdultContent