

Hypers at ComMA@ICON: Modelling Aggressiveness, Gender Bias and Communal Bias Identification

Sean Benhur^{1*}, Roshan Nayak^{2*}, Kanchana Sivanraju¹
Adeep Hande³, Subalalitha Chinnaudayar Navaneethakrishnan⁴,
Ruba Priyadharshini⁵, Bharathi Raja Chakravarthi⁶

¹PSG College of Arts and Science ²BMS College of Engineering

³Indian Institute of Information Technology Tiruchirappalli

⁴ SRM Institute of Science and Technology, India ⁵ULTRA Arts and Science College

⁶National University of Ireland Galway

seanbenhur@gmail.com

Abstract

Due to the exponentially increasing reach of social media, it is essential to focus on its negative aspects as it can potentially divide society and incite people into violence. In this paper, we present our system description of work on the shared task ComMA@ICON, where we have to classify how aggressive the sentence is and if the sentence is gender-biased or communal-biased. These three could be the primary reasons to cause significant problems in society. As team Hypers we have proposed an approach which utilizes different pretrained models with Attention and mean pooling methods. We were able to get Rank 3 with 0.223 Instance F1 score on Bengali, Rank 2 with 0.322 Instance F1 score on Multi-lingual set, Rank 4 with 0.129 Instance F1 score on Meitei and Rank 5 with 0.336 Instance F1 score on Hindi. The source code and the pretrained models of this work can be found here¹.

1 Introduction

The Internet is a vast network that connects devices all over the world. Due to mobile technology and affordable internet plans, users can access the Internet with ease, leading to the tremendous growth of the Internet, which is unprecedented. As of January 2021, there were 4.66 billion active internet users, 59.5% of the world's population. Users would undoubtedly want to increase their reach virtually, and hence the interaction among the people would increase. The people these days are more vocal and, at any cost, want their voices or opinions to be reached to a multitude of people. Hence, people search for a platform to share their views, and social media is an ideal place for that. This exact

mindset of people has fueled the copious amounts of social media users globally.

Social media are the technologies that allow the creation, sharing, or exchange of information, interests, ideas, and other forms of expression. Its use is an ever-increasing phenomenon of the 21st century (Livingstone and Brake, 2010). There are a plethora of social media platforms, each attracting people in unique ways. As of January 2021, there were 4.2 billion active social media users. Considering the reach of social media, they can spread people's opinions in a few minutes (Zhang and Vos, 2015). Hence it will have a significant effect on society which could be both positive as well as harmful (Harchekar, 2017). But there are instances in which the situation would go out of hand. For example, people could differ in their opinions, and people with similar views tend to form a group to denounce the group with ideas that are not the same as theirs. During the denouncement, there is a possibility that a user could show his improper behaviour, thus making offensive (Hande et al., 2021b), misogynistic (Shushkevich and Cardiff, 2019), hateful (Bhatia et al., 2021), or any kind of statements that has the potential to create controversy (Coletto et al., 2017). Such statements may be intended towards an individual or a group and are not considered to be good or acceptable in the society. As such behaviour would influence others wrongly and instigate violence or affect the mental health, leading to unpleasant situations. Hence it is necessary to flag such posts and its advisable to take them down from the social media platform and also retribute the user responsible for such posts (Hande et al., 2021a). There could be several reasons that a post by the user is considered inappropriate. Considering how important it is to regulate toxic post, in this paper we will be presenting a system to identify if the user is being aggressive on some individual or a community, or being biased regard-

Equal Contribution

¹https://github.com/seanbenhur/multilingual_aggressive_gender_bias_communal_bias_identification

Text	Language	Aggressive	Gender Bias	Communal Bias
angakpa nupini eiga unaradi fadoubi	Meitei	CAG	NGEN	NCOM
hehhh ym pkte nupi ng	Meitei	CAG	GEN	NCOM
tome hola bal chera tumi nijai jante	Bengali	OAG	GEN	NCOM
you know to whom im addressing, 'ye hindustan ke liye dimak h jo usko ander se khokla krre h' #muslimvirus	Hindi	OAG	NGEN	COM
mulle tere allah ki ma ka bhosda	Hindi	OAG	GEN	COM
gudmarani chale ke maaaaar	Multi	OAG	GEN	NCOM
jay bheem namo buddhay	Multi	OAG	NGEN	NCOM

Table 1: Samples from the dataset and their corresponding class labels for each of the tasks.

Language	NAG	CAG	OAG	NGEN	GEN	NCOM	COM	Total
Meitei	1,258	1,495	456	3,006	203	2,967	242	3,209
Bengali	1,115	494	1,782	2,120	1,271	2,975	416	3,391
Hindi	1,594	969	3,052	4,440	1,175	4,402	1,213	5,615
Multi-lang	3,966	2,956	5,289	9,564	2,647	10,342	1,869	12,211

Table 2: Samples distribution in the training set.

ing the gender (Jatmiko et al., 2020; Hande et al., 2020), or is targeting a particular religion or caste (roy, 2016).

Undoubtedly, English is the widely spoken language in the world (Crystal, 2008). But as there are no hardbound rules that users must text in English, the text found on social media could be multilingual and lack grammatical rules (Yuliah et al., 2020). Also, there could be unwanted symbols in the text (Chakravarthi et al., 2021). Considering all such challenges, in this paper, we present a model to classify the multi-lingual sentence written by the user as to how aggressive it is and if it is gendered and communal oriented text. The dataset had multilingual texts with the code-mix of English and several other languages native to India. Meitei and Bangla are native to the Indian states of Manipur and West Bengal, respectively, whereas Hindi is predominant in Northern India.

The rest of the paper is structured as follows, section 2 describes about dataset used for the shared task. The section 3 describes the models and architectures that were used for the tasks. In section 4 we discuss about the results obtained during the study, and the last section 5 concludes the work.

2 Dataset

The ComMA dataset was provided in this task (Kumar et al., 2021a; B et al., 2021). The dataset

had annotations for aggression, gender bias, and communal bias identification for multi-lingual social media sentences (Kumar et al., 2021b). The dataset comprises of code-mixed sentences has 15,000 code-mixed sentences. It is divided into 12,000 sentences for development and 3,000 sentences for the test. The data is divided into four sets, namely Meitei, Bengali, Hindi, and Multi-lingual. The Multi-lingual set comprises sentences of all three languages. The Table 1 gives an idea of how the dataset could look like. The sentences in every set are classified into one of the classes for each of the three tasks. The tasks and their classes include,

- **Aggression Classification:** The text is divided into Overtly Aggressive (OAG), Covertly Aggressive (CAG) or Non-aggressive (NAG)
- **Gender Bias Classification:** The text is divided into gendered (GEN) or non-gendered (NGEN).
- **Communal Bias Classification:** The text is divided into communal (COM) or non-communal (NCOM)

The samples count of classes is far from equal.

Language	NAG	CAG	OAG	NGEN	GEN	NCOM	COM	Total
Meitei	370	471	159	945	55	932	68	1,000
Bengali	333	157	501	624	367	879	112	991
Hindi	305	167	526	775	225	804	196	998
Multi	1,007	797	1,193	2,349	648	2,622	375	2,997

Table 3: Distribution of samples in the dev set.

Model	Arch	Language	AGG			GB			CB			Overall		
			P	R	F1	P	R	F1	P	R	F1	P	R	F1
MURIL	AP	Meitei	0.470	0.470	0.470	0.599	0.599	0.599	0.493	0.493	0.493	0.521	0.521	0.521
MURIL	MP	Meitei	0.471	0.471	0.471	0.603	0.603	0.603	0.356	0.356	0.356	0.477	0.477	0.477
csebuetnlp/banglabert	AP	Bengali	0.642	0.642	0.642	0.755	0.755	0.755	0.692	0.692	0.692	0.696	0.696	0.696
csebuetnlp/banglabert	MP	Bengali	0.635	0.635	0.635	0.762	0.762	0.762	0.612	0.612	0.612	0.670	0.670	0.670
MURIL	AP	Hindi	0.594	0.594	0.594	0.816	0.816	0.816	0.909	0.909	0.909	0.773	0.773	0.773
MURIL	MP	Hindi	0.683	0.683	0.683	0.827	0.827	0.827	0.902	0.902	0.902	0.804	0.804	0.804
MURIL	AP	Multi-Lang	0.618	0.618	0.618	0.839	0.839	0.839	0.661	0.661	0.661	0.706	0.706	0.706
MURIL	MP	Multi-Lang	0.612	0.612	0.612	0.823	0.823	0.823	0.891	0.891	0.891	0.791	0.791	0.791

Table 4: Results on the dev set. AGG: Aggressive, GB: Gender Bias, CB: Communal Bias, Arch: Architecture, AP: Attention-pooling, MP: Mean-pooling. Metrics, Micro average scores of P: Precision, R: Recall, F1: F1-score calculated on the dev set. Overall scores are the average of the aggressive, gender bias, and communal bias.

Hence the dataset is quite imbalanced. The dataset distribution is displayed in the Table 2.

3 Methodology

In this section, we describe the methodology of our systems, including data preprocessing and Model architecture. We use mean pooled, and Attention pooled pretrained models, which was shown to provide better results (Benhur and Sivanraju, 2021). We trained all the models with the batch size of 8, dropout of 0.3, linear scheduler for learning rate scheduling with $2e-5$ as an initial learning rate.

3.1 Data Preprocessing

The task dataset consists of both codemixed and native scripts; for the Bengali dataset, we converted the emojis into Bengali language using bnemo GitHub repository², we removed URLs and punctuations in the text for all the languages. Since the dataset is imbalanced, we sampled the dataset uniformly.

3.2 Pretrained Models

We finetuned pretrained transformers with custom poolers on hidden states on MURIL (Khanuja et al., 2021) for Hindi, Meitei and Multilingual datasets and BanglaBert (Bhattacharjee et al., 2021) for Bengali dataset. In this section, we describe our Pooling methods and pretrained models.

²<https://github.com/faruk-ahmad/bnemo>

3.2.1 MURIL

MURIL is a pretrained model, specifically made for Indian languages. MuRIL, the pretrained model, is trained in 16 different Indian Languages. Instead of the usual Masked Language Modelling approach, the model is trained on both Masked Language Modelling(MLM) objective and Translated Language Modelling(TLM) objective. In the TLM objective, both translated and transliterated sentence pairs are sent to the model for training. This model outperforms mBERT on all the tasks for Indian languages on the XTREME (Hu et al., 2020) benchmark.

3.2.2 BanglaBert

Banglabert is pretrained on more than 18.5 GB of a corpus in Bengali corpora. Banglabert achieves the state of the art performance on Bengali texts on five downstream tasks. It outperforms multilingual baselines with more than a 3.5 percentage score. Banglabert is pretrained using ELECTRA (Clark et al., 2020) with Replaced Token Detection(RTD) objective. In this setup, two networks, a generator network and discriminator network, are used, while training both the networks are trained jointly. The generator is trained on the Masked Language Modelling objective, where a portion of the tokens in the sentence is masked and is asked to predict the masked tokens using the rest of the input. The masked tickets are replaced by tokens sampled from the generator’s output distribution for the corresponding masks for the discrimina-

Model	Architecture	Language	Overall Micro F1	Overall Instance F1
MURIL	Attention-pooler	Meitei	0.472	0.129
MURIL	Mean-pooler	Meitei	0.436	0.080
csebuetnlp/banglabert	Attention-pooler	Bengali	0.579	0.223
csebuetnlp/banglabert	Mean-pooler	Bengali	0.572	0.201
MURIL	Attention-pooler	Hindi	0.662	0.326
MURIL	Mean-pooler	Hindi	0.683	0.336
MURIL	Attention-pooler	Multi-Lang	0.685	0.322
MURIL	Mean-pooler	Multi-Lang	0.601	0.280

Table 5: Results on the test set. Overall, the Micro F1 score is calculated by the average of aggressive, gender, and communal biases. Instance F1 score is similar to the F1 score but when all the three labels are predicted correctly.

Language	Overall Micro F1	Overall Instance F1
Meitei	0.472	0.129
Bangla	0.579	0.223
Hindi	0.683	0.336
Multi-Lang	0.685	0.322

Table 6: Results obtained when submitted to the competition.

tor input. The discriminator then has to predict whether each token is from the original sequence or not. After pretraining, the discriminator is used for finetuning.

3.3 Attention Pooler

The attention operation described in equation 1 is applied to the CLS token in last hidden state of the pretrained transformer; we hypothesize that this helps the model learn the contribution of individual tokens. Finally, the returned pooled output from the transformer is further passed to a linear layer to predict the label.

$$o = W_h^T \text{softmax}(qh_{CLS}^T)h_{CLS} \quad (1)$$

where W_h^T and q are learnable weights and h_{CLS} is the CLS representation and o is the output.

$$y = \text{softmax}(W_o^T + bo) \quad (2)$$

3.4 Mean Pooler

In the mean-pooling method, the last hidden state of the tokens are averaged on each sentence, and it is passed onto the linear layer to output the final probabilities.

4 Results

Pretrained models with different pooling methods were trained on each language set and then validated on dev sets. For the competition submissions, we submitted the model with a higher Micro F1 score on the dev set. Table 4 shows the results

of the dev set, and the Table 3 depicts the data distribution of the set used to validate the trained models. The training process was done on Tesla P100 GPU. In the test set submissions, We were able to get Rank 3 with 0.223 InstanceF1 score on Bengali, Rank 2 with 0.322 Instance F1 score on Multi-lingual set, Rank 4 with 0.129 Instance F1 score on Meitei and Rank 5 with 0.336 Instance F1 score on Hindi. The competition results are shown in Table 6. Table 5 shows the Overall Micro F1 score and Instance F1 score on the test set. The pre-trained model MURIL was not trained on Meitei, but it still achieved comparable performance on the Test set; we hypothesize that since MURIL was trained both on transliterated pairs on TLM objective and the Meitei dataset also only consisted of code-mixed texts, we get a fair results on meitei test set.

5 Conclusion

In this paper, we experimented with different pooling methods, namely Attention Pooling and Mean Pooling and pretrained models, to classify sentences, how aggressive they are, and whether gender-oriented or communal. From Table 5 its evident that attention pooling worked better in most of the cases. We have also discussed the various essential reasons why the work on this is necessary. As for future work, we will consider improving our scores, especially on multilingual and meitei datasets, and experimenting with other pretrained models.

References

2016. Social media in the domain of communal violence: a study of assam riot 2012.
- Senthil Kumar B, Aravindan Chandrabose, and Bharathi Raja Chakravarthi. 2021. [An overview of fairness in data – illuminating the bias in data pipeline](#). In *Proceedings of the First Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 34–45, Kyiv. Association for Computational Linguistics.
- Sean Benhur and Kanchana Sivanraju. 2021. [Psg@dravidian-codemix-hasoc2021: Pretrained transformers for offensive language identification in tanglish](#).
- Mehar Bhatia, Tenzin Singhay Bhotia, Akshat Agarwal, Prakash Ramesh, Shubham Gupta, Kumar Shridhar, Felix Laumann, and Ayushman Dash. 2021. [One to rule them all: Towards joint indic language hate speech detection](#).
- Abhik Bhattacharjee, Tahmid Hasan, Kazi Samin Mubasshir, Mohammad Rahman, Anindya Iqbal, and Rifat Shahriyar. 2021. [Banglabert: Combating embedding barrier for low-resource language understanding](#).
- Bharathi Raja Chakravarthi, Ruba Priyadharshini, Vigneshwaran Muralidaran, Navya Jose, Shardul Suryawanshi, Elizabeth Sherly, and John P. McCrae. 2021. [Dravidiancodemix: Sentiment analysis and offensive language identification dataset for dravidian languages in code-mixed text](#).
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#).
- Mauro Coletto, Kiran Garimella, Aristides Gionis, and Claudio Lucchese. 2017. [Automatic controversy detection in social media: A content-independent motif-based approach](#). *Online Social Networks and Media*, 3-4:22–31.
- David Crystal. 2008. [Two thousand million?](#) *English Today*, 24(1):3–6.
- Adeep Hande, Siddhanth U Hegde, Ruba Priyadharshini, Rahul Ponnusamy, Prasanna Kumar Kumaresan, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021a. [Benchmarking multi-task learning for sentiment analysis and offensive language identification in under-resourced dravidian languages](#).
- Adeep Hande, Ruba Priyadharshini, and Bharathi Raja Chakravarthi. 2020. [KanCMD: Kannada CodeMixed dataset for sentiment analysis and offensive language detection](#). In *Proceedings of the Third Workshop on Computational Modeling of People’s Opinions, Personality, and Emotion’s in Social Media*, pages 54–63, Barcelona, Spain (Online). Association for Computational Linguistics.
- Adeep Hande, Karthik Puranik, Konthala Yasaswini, Ruba Priyadharshini, Sajeetha Thavareesan, Anbukkarasi Sampath, Kogilavani Shanmugavadivel, Durairaj Thenmozhi, and Bharathi Raja Chakravarthi. 2021b. [Offensive language identification in low-resourced code-mixed dravidian languages using pseudo-labeling](#).
- Jyoti Suraj Harchekar. 2017. [Impact of social media on society](#).
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization](#).
- Mochamad Jatmiko, Muh Syukron, and Yesi Mekarsari. 2020. [Covid-19, harassment and social media: A study of gender-based violence facilitated by technology during the pandemic](#). *The Journal of Society and Media*, 4:319–347.
- 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. 2021. [Muril: Multilingual representations for indian languages](#).
- Ritesh Kumar, Bornini Lahiri, Akanksha Bansal, Enakshi Nandi, Laishram Niranjana Devi, Shyam Ratan, Siddharth Singh, Akash Bhagat, and Yogesh Dawer. 2021a. [Comma@icon: Multilingual gender biased and communal language identification task at icon-2021](#). In *Proceedings of the 18th International Conference on Natural Language Processing (ICON): COMMA@ICON 2021 Shared Task*, Silchar, India. NLP Association of India (NLP AI).
- Ritesh Kumar, Enakshi Nandi, Laishram Niranjana Devi, Shyam Ratan, Siddharth Singh, Akash Bhagat, Yogesh Dawer, and Akanksha Bansal. 2021b. [The comma dataset v0.2: Annotating aggression and bias in multilingual social media discourse](#).
- Sonia Livingstone and David R Brake. 2010. [On the rapid rise of social networking sites: New findings and policy implications](#).
- Elena Shushkevich and John Cardiff. 2019. [Automatic misogyny detection in social media: A survey](#). *Computación y Sistemas*, 23.
- Siti Yuliah, Yessy Purnamasari, and Elsa Yunita. 2020. [Grammatical errors in social media caption](#). *INTERNATIONAL JOURNAL OF LANGUAGE LITERATURE*, 8.
- Boyang Zhang and Marita Vos. 2015. [How and why some issues spread fast in social media](#). *Online Journal of Communication and Media Technologies*, 5:371–383.