

MasonPerplexity at ClimateActivism 2024: Integrating Advanced Ensemble Techniques and Data Augmentation for Climate Activism Stance and Hate Event Identification

Al Nahian Bin Emran*, Amrita Ganguly*, Sadiya Sayara Chowdhury Puspoo
Dhiman Goswami, Md Nishat Raihan

George Mason University, USA
{abinemra, agangul}@gmu.edu

Abstract

The task of identifying public opinions on social media, particularly regarding climate activism and the detection of hate events, has emerged as a critical area of research in our rapidly changing world. With a growing number of people voicing either to support or oppose to climate-related issues - understanding these diverse viewpoints has become increasingly vital. Our team, *MasonPerplexity*, participates in a significant research initiative focused on this subject. We extensively test various models and methods, discovering that our most effective results are achieved through ensemble modeling, enhanced by data augmentation techniques like back-translation. In the specific components of this research task, our team achieved notable positions, ranking 5th, 1st, and 6th in the respective sub-tasks, thereby illustrating the effectiveness of our approach in this important field of study.

1 Introduction

In the ever-evolving landscape of climate change activism, encouraging meaningful conversations and comprehending how things change throughout events depends critically on the ability to recognize hate speech and the understanding of attitude during these events. This paper presents our effort in the Shared Task on Hate Speech and Stance Detection during Climate Activism (Thapa et al., 2024), where our goal is to develop effective models for hate speech detection, target identification, and stance detection.

This task consists of three subtasks that work together to support an integrated approach to event identification. The goal of the first subtask is to identify whether the given text contains hate speech or not. The second subtask focuses on identifying if people, groups, or communities are targets of hate speech. Lastly, Stance Detection provides insight into the dynamics of climate activism protests

by assessing the support, opposition, or neutrality indicated within texts.

Our paper serves as a comprehensive system description, outlining the approaches and models used to address these subtasks within the framework of activist events related to climate change. We present our ensemble method for identifying hate speech, which combines robust models like XLM-roBERTa-large (Conneau et al., 2019), BERTweet-large (Ushio and Camacho-Collados, 2021a), and fBERT (Sarkar et al., 2021). Notably, for Target Detection, the best-performing model is BERTweet-large (Ushio and Camacho-Collados, 2021b) while BERTweet-base (Nguyen et al., 2020) excels in Stance Detection.

We also discuss our fine-tuning strategies and dataset augmentation techniques, demonstrating our commitment to refining model performance. Our approach's effectiveness is demonstrated by our remarkable F1 scores of 0.8885, 0.7858, and 0.7373. Furthermore, our team named *MasonPerplexity* has secured 5th, 1st, and 6th ranks in the respective subtasks, underscoring the competence of our models in comparison to peers.

Through this paper, we aim to contribute to the advancement of hate speech and stance detection in the context of climate activism, fostering a safer and more informed space for dialogue and understanding during crucial events. We employ ensemble methods to better classify the texts - our approach increases the accuracy metrics for the first sub-task where we encounter a comparatively larger amount of data. We also use data augmentation methods, which further improve our results.

2 Related Works

The paper (Parihar et al., 2021) explores the rising concern of hate speech on the internet and its potential impact. It emphasizes machine learning and deep learning models in automatically identifying hate speech. In (Malmasi and Zampieri, 2017),

* denotes Equal Contribution

English tweets are subjected to supervised classification using n-gram features and a linear SVM classifier. Even at 78% accuracy, it is still difficult to discern offensive language from hate speech. By combining recurrent neural networks and user features, (Pitsilis et al., 2018) outperforms current systems and achieves a remarkable F-score of 0.9320 on a Twitter dataset. Additionally, Warner and Hirschberg (Warner and Hirschberg, 2012) define hate speech and despite limitations in capturing larger language patterns, SVM classification can detect anti-Semitic speech with 94% accuracy.

The literature on target detection in hate speech unveils valuable insights through various studies in the field. (Lemmens et al., 2021) focuses on Dutch Facebook comments, exploring hateful metaphors to enhance hate speech type and target detection. The study incorporates manual metaphor annotations as features for SVM and BERT models, observing improvements in F1 scores. Conversely, (Zampieri et al., 2019) proposes a hierarchical annotation scheme for offensive language in English tweets, creating the OLID dataset. The study employs SVM, CNN, and BiLSTM models, achieving notable results and providing a valuable resource for offensive language research.

Stance detection, a crucial aspect of NLP, involves determining a person’s position towards a concept. (Küçük and Can, 2021) outlines the significance and challenges in this domain, emphasizing its relation to sentiment analysis, emotion detection, and other tasks. It highlights the evolution facilitated by shared tasks, varied approaches, including traditional SVMs and newer LSTM models, and the necessity of annotated datasets. Additionally, (Upadhyaya et al., 2023) introduces a multitasking approach, enhancing performance on multiple datasets, and showcasing the potential of incorporating auxiliary tasks. Furthermore, (Küçük and Can, 2018) contributes a valuable stance-annotated Turkish Twitter dataset, showcasing the diversity of research efforts in stance detection.

3 Datasets

From the tables for Subtask 1, Subtask 2, and Subtask 3, it is evident that the dataset (Shiwakoti et al., 2024) is imbalanced across different labels.

3.1 Hate Speech Detection

In subtask A, the distribution between NON-HATE and HATE is heavily skewed towards NON-HATE,

with approximately 87.66% in the training set, 87.83% in the evaluation set, and 87.96% in the test set. This indicates a significant class imbalance, which may pose challenges for model training and evaluation.

3.2 Target Detection

In subtask B, there is an imbalance among the labels INDIVIDUAL, ORGANIZATION, and COMMUNITY. The majority of instances belong to the individual category, with around 80.54% in the training set, 80.00% in the evaluation set, and 80.67% in the test set.

3.3 Stance Detection

Subtask C exhibits an imbalance, between the SUPPORT, OPPOSE, and NEUTRAL labels. SUPPORT dominates the dataset, comprising 59.42% in the training set, 57.46% in the evaluation set, and 58.96% in the test set, where OPPOSE has respective percentages 9.61%, 9.80%, and 9.03%.

In summary, the dataset for all three subtasks is not well-balanced, and addressing this imbalance may be crucial for developing models that generalize well across different classes.

Subtask A			
Label	Train	Eval	Test
NON-HATE	87.66	87.83	87.96
HATE	12.34	12.17	12.04

Table 1: label wise data percentage of subtask A

Subtask B			
Label	Train	Eval	Test
INDIVIDUAL	80.54	80.00	80.67
ORGANIZATION	15.02	15.33	15.33
COMMUNITY	4.44	4.67	4.00

Table 2: label wise data percentage of subtask B

Subtask C			
Label	Train	Eval	Test
SUPPORT	59.42	57.46	58.96
OPPOSE	9.61	9.80	9.03
NEUTRAL	30.97	32.74	32.01

Table 3: label wise data percentage of subtask C

4 Experiments

In subtask A, we initially employ GPT3.5 (OpenAI, 2023) zero shot and few shot prompting with Test F1 score 0.66 and 0.73. The prompt provided to GPT3.5 is available in Figure 1.

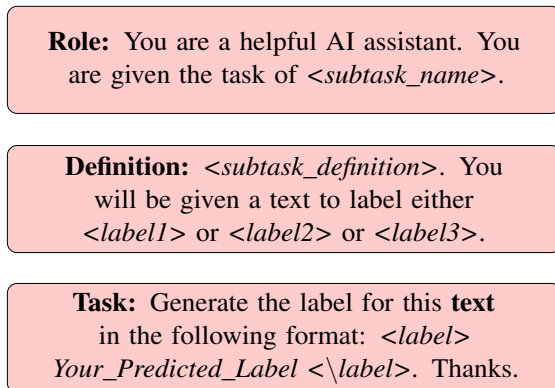


Figure 1: Sample GPT-3.5 prompt.

Then we use BERTweet large (Ushio and Camacho-Collados, 2021a) (Ushio et al., 2022), XLM-R (Conneau et al., 2019), HATE-BERT (Caselli et al., 2021) and fBERT (Sarkar et al., 2021). Following this, we adopt a weighted ensemble approach (Ensemble 1) for the best three models (BERTweet, XLM-R, fBERT). Similarly, we perform another weighted ensemble approach (Ensemble 2) with the same models only replacing BERTweet with HATE-BERT, as these two models show the same F1 score on test data with the same setting. However, the former ensemble strategy yields the highest F1 score for this task.

To address class imbalance in subtask A, we implement back translation by converting the training data of those specific labels that have a smaller ratio with respect to the whole training set through various languages, including Xosha to Twi to English, Lao to Pashto to Yoruba to English, Yoruba to Somali to Kinyarwanda to English, and Zulu to Oromo to Shona to Tsonga to English. This approach significantly contributed to improving the overall F1 score of Ensemble 1 from 0.85 to 0.88 and Ensemble 2 from 0.86 to 0.89.

We follow the approach of back translation of (Raihan et al., 2023). For this, we select languages that demonstrate limited or no cultural overlap with the original language featured in the dataset. Xosha, Twi, Lao, Pashto, Yoruba, Somali, Kinyarwanda, Zulu, Oromo, Shona, and Tsonga are languages that are very diverse culturally and geographically. This diversity underscores the significance of con-

sidering a wide range of cultural and geographical influences when working with these languages. By intentionally selecting these languages without cultural overlap, we introduce a purposeful aspect of diversity, mitigating potential biases, and enhancing the dataset with a broader spectrum of linguistic expressions. Moreover, the Ensemble method with majority voting is also proven helpful in this type of case where a single model may not label the data correctly due to class imbalance (Goswami et al., 2023). For instance, when two out of three models predict a sentence as a hate event, the sentence is subsequently labeled as a hate event through the application of majority voting.

In subtask B, we utilize BERTweet-large (Ushio and Camacho-Collados, 2021a), BERT base (Devlin et al., 2018), and XLM-R (Conneau et al., 2019). Additionally, like subtask 1, we implement back translation using the same language sequences mentioned earlier to address class imbalance. Notably, BERTweet large (Ushio and Camacho-Collados, 2021a) demonstrates the highest F1- score among these models. We also use GPT3.5 zero shot and few shot prompting with 0.63 and 0.64 test F1 scores.

BERTweet-large (Ushio and Camacho-Collados, 2021a), BERT base (Devlin et al., 2018), and BERTweet base (Nguyen et al., 2020) models are applied in subtask C for stance detection. Among these models, the BERTweet base achieves the highest F1 score. F1 score for GPT3.5 zero shots and few shot prompting are 0.63 and 0.67.

Hyperparameters of all the models used excluding GPT3.5 in the experiments are available in Figure 4.

Parameter	Value
Learning Rate	$1e - 5$
Train Batch Size	8
Test Batch Size	8
Epochs	5
Dropout	0.2

Table 4: Training Configuration Parameters

5 Results

The results in Tables 5, 6, and 7 provide a comprehensive evaluation of various NLP models across the three subtasks of the shared task.

In subtask A, our ensemble approach (Ensemble 2 with HATE-BERT, XLM-R and fBERT models)

secures the fifth rank. For subtask B, BERTweet large secures the top position (Rank 1), while in subtask C, we achieve the sixth rank utilizing BERT-Base.

Model	Eval F1	Test F1
GPT3.5-(ZERO SHOT)	–	0.66
GPT3.5-(FEW SHOT)	–	0.73
HATE-BERT	0.88	0.83
BERTWEET-LARGE	0.89	0.84
XLM-R	0.89	0.85
F-BERT	0.90	0.85
*ENSEMBLE 1	0.90	0.85
**ENSEMBLE 2	0.91	0.86
HATE-BERT (AUG.)	0.91	0.87
BERTWEET-LARGE (AUG.)	0.92	0.87
XLM-R (AUG.)	0.91	0.88
F-BERT (AUG.)	0.93	0.88
*ENSEMBLE 1 (AUG.)	0.93	0.88
**ENSEMBLE 2 (AUG.)	0.94	0.89

Table 5: Results of subtask A (before and after data augmentation). *Ensemble 1 (BERTweet-large, XLM-R, fBERT), **Ensemble 2 (HATE-BERT, XLM-R, fBERT)

Model	Eval F1	Test F1
GPT3.5-(ZERO SHOT)	–	0.63
GPT3.5-(FEW SHOT)	–	0.64
XLM-R	0.75	0.60
BERT-BASE	0.86	0.69
BERTWEET-LARGE	0.97	0.79

Table 6: Results of subtask B.

Model	Eval F1	TEST F1
GPT3.5-(ZERO SHOT)	–	0.63
GPT3.5-(FEW SHOT)	–	0.67
BERT-BASE	0.71	0.69
BERTWEET-LARGE	0.71	0.70
BERTWEET-BASE	0.80	0.74

Table 7: Results of subtask C.

6 Error Analysis

Upon evaluating our models’ performance across the three subtasks, we identify several key sources of errors that contributed to limiting our scores.

In subtask A on hate speech detection, our ensemble model struggles with longer text segments that express hate in subtle or nuanced ways. The models are not always able to pick up on the underlying mocking or criticism woven into complex

rhetorical devices. Additionally, sarcasm and irony continue to pose challenges, as models interpret literally what is meant to convey the opposite meaning.

For subtask B on target identification, errors frequently occur in distinguishing between organizations and communities as categories. Our models have difficulty consistently applying the definitions and criteria that delineate these two groups as targets of hate speech. There are also inconsistencies in labeling individual people who are associated with or represent a broader community.

Regarding subtask C on stance detection, our models struggle to some extent with longer text segments, having more trouble identifying stances from among nuanced discussions. Shorter, more direct statements of opposition or support were simpler for the models to categorize accurately.

To visualize label-wise models’ performance we can see the Figures 2, 3, and 4 of confusion matrices for all the subtasks.

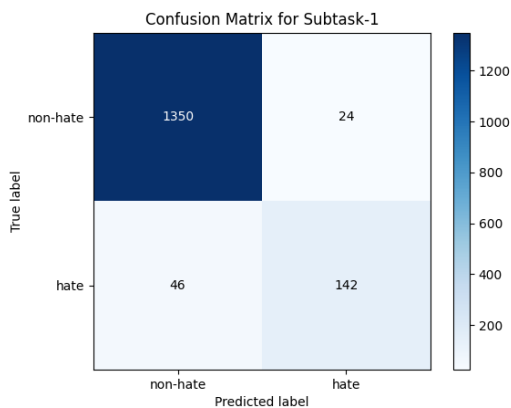


Figure 2: Confusion Matrix for Hate Speech Detection




Figure 3: Confusion Matrix for Target Detection




Figure 4: Confusion Matrix for Stance Detection

7 Conclusion

In conclusion, the *MasonPerplexity* team has made significant strides in the domain of detecting climate activism stances and hate events on social media. Through a comprehensive evaluation of various models, our research underscores the efficacy of ensemble modeling coupled with data augmentation techniques like back-translation. Our achievements in the shared task, marked by rankings of 5th, 1st, and 6th in the respective subtasks, reflect the potential of our methodologies in addressing the complexities of sentiment analysis in the context of climate activism.

There are several avenues for future research. Firstly, addressing the challenge of label imbalance in our dataset could enhance the accuracy and reliability of our models. Exploring advanced techniques in data sampling or synthetic data generation may provide viable solutions. Secondly, the refinement of label quality through more rigorous annotation processes or leveraging semi-supervised learning techniques could further improve model performance. Finally, the integration of Large Language Model (LLM) fine-tuning presents a promising direction. Fine-tuning pre-trained models specifically for the nuances of climate activism discourse and hate speech detection could yield more nuanced and contextually aware results. Additionally, expanding our research to include multilingual datasets would enhance the applicability and relevance of our findings in a global context, fostering a more comprehensive understanding of public sentiment on climate issues worldwide.

Limitations

This study encounters some limitations that affect its outcomes. The first issue is with the balance of labels in our dataset. We have more examples of some types of data than others, a problem known as label imbalance. This imbalance can lead our model to be better at recognizing the more common types and not as good with the rare ones, creating a bias in our results. The second limitation is the quality of the labels themselves. In our dataset, some labels are incorrect or not consistent with each other. This poor quality can confuse the model, making it harder for it to learn correctly and possibly leading to inaccurate results. Lastly, we did not fine-tune Large Language Models (LLMs) for our specific task. Fine-tuning is a process where a pre-trained model, like an LLM, is further trained on a specific type of data. Not doing this fine-tuning means we may not be taking full advantage of the LLM’s capabilities, which can improve our model’s understanding of complex patterns in climate activism and hate event data. However, due to a lack of computing resources, we are not fine-tuning.

Acknowledgment

We would like to thank the shared task organizers for providing us with the dataset used in our study. Moreover, we also want to express our gratitude to Dr. Marcos Zampieri for his effective guidelines throughout the span of the competition.

Ethics Statement

The present study, which centers on the identification of Climate Activism Stance and Hate Event, rigorously adheres to the [ACL Ethics Policy](#) and seeks to make a valuable contribution to the realm of online safety. The dataset is supplied to us by the organizers and has undergone anonymization to secure the privacy of the users. The technology in question possesses the potential to serve as a beneficial instrument for the moderation of online content, thereby facilitating the creation of safer digital environments. However, it is imperative to exercise caution and implement stringent regulations to prevent its potential misuse for purposes such as monitoring or censorship.

References

- Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. [HateBERT: Retraining BERT for abusive language detection in English](#). In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 17–25, Online. Association for Computational Linguistics.
- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Dhiman Goswami, Md Nishat Raihan, Sadiya Sayara Chowdhury Puspo, and Marcos Zampieri. 2023. [nlpBDpatriots at BLP-2023 task 2: A transfer learning approach towards Bangla sentiment analysis](#). In *Proceedings of the First Workshop on Bangla Language Processing (BLP-2023)*, pages 286–292, Singapore. Association for Computational Linguistics.
- Dilek Küçük and Fazli Can. 2021. [Stance detection: Concepts, approaches, resources, and outstanding issues](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 2673–2676, New York, NY, USA. Association for Computing Machinery.
- Dilek Küçük and Fazli Can. 2018. [Stance detection on tweets: An svm-based approach](#). *CoRR*, abs/1803.08910.
- Jens Lemmens, Ilija Markov, and Walter Daelemans. 2021. Improving hate speech type and target detection with hateful metaphor features. In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*.
- Shervin Malmasi and Marcos Zampieri. 2017. Detecting hate speech in social media. In *Proceedings of the International Conference Recent Advances in Natural Language Processing (RANLP)*, pages 467–472.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. BERTweet: A pre-trained language model for English Tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14.
- OpenAI. 2023. [Gpt-3.5 turbo fine-tuning and api updates](#). Accessed: 2023-08-28.
- Anil Singh Parihar, Surendrabikram Thapa, and Sushruti Mishra. 2021. [Hate speech detection using natural language processing: Applications and challenges](#). In *2021 5th International Conference on Trends in Electronics and Informatics (ICOEI)*, pages 1302–1308.
- Georgios K. Pitsilis, Heri Ramampiaro, and Helge Langseth. 2018. [Effective hate-speech detection in twitter data using recurrent neural networks](#). *Applied Intelligence*, 48:4730 – 4742.
- Md Nishat Raihan, Dhiman Goswami, Sadiya Sayara Chowdhury Puspo, and Marcos Zampieri. 2023. [nlpBDpatriots at BLP-2023 task 1: Two-step classification for violence inciting text detection in Bangla - leveraging back-translation and multilinguality](#). In *Proceedings of the First Workshop on Bangla Language Processing (BLP-2023)*, pages 179–184, Singapore. Association for Computational Linguistics.
- Diptanu Sarkar, Marcos Zampieri, Tharindu Ranasinghe, and Alexander Ororbia. 2021. [fbert: A neural transformer for identifying offensive content](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1792–1798.
- Shuvam Shiwakoti, Surendrabikram Thapa, Kritesh Rauniyar, Akshyat Shah, Aashish Bhandari, and Usman Naseem. 2024. Analyzing the dynamics of climate change discourse on twitter: A new annotated corpus and multi-aspect classification. *Preprint*.
- Surendrabikram Thapa, Kritesh Rauniyar, Farhan Ahmad Jafri, Shuvam Shiwakoti, Hariram Veeramani, Raghav Jain, Guneet Singh Kohli, Ali Hürriyetoğlu, and Usman Naseem. 2024. Stance and hate event detection in tweets related to climate activism - shared task at case 2024. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*.
- Apoorva Upadhyaya, Marco Fisichella, and Wolfgang Nejdl. 2023. [Toxicity, morality, and speech act guided stance detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4464–4478, Singapore. Association for Computational Linguistics.
- Asahi Ushio and Jose Camacho-Collados. 2021a. [T-NER: An all-round python library for transformer-based named entity recognition](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 53–62, Online. Association for Computational Linguistics.
- Asahi Ushio and Jose Camacho-Collados. 2021b. [T-NER: An all-round python library for transformer-based named entity recognition](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 53–62, Online. Association for Computational Linguistics.
- Asahi Ushio, Leonardo Neves, Vitor Silva, Francesco Barbieri, and Jose Camacho-Collados. 2022. Named Entity Recognition in Twitter: A Dataset and Analysis on Short-Term Temporal Shifts. In *The 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th*

International Joint Conference on Natural Language Processing, Online. Association for Computational Linguistics.

William Warner and Julia Hirschberg. 2012. Detecting hate speech on the world wide web. pages 19–26.

Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Predicting the type and target of offensive posts in social media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.