

Innovators@SMM4H'22: An Ensembles Approach for self-reporting of COVID-19 Vaccination Status Tweets

Mohammad Zohair^ϕ, Nidhir Bhavsar[&], Aakash Bhatnagar[§], Muskaan Singh[#],
Petr Motlicek[#]

^ϕJamia Millia Islamia, New Delhi, India

[§]Boston University, Boston, Massachusetts

[&]University of Potsdam, Germany

[#] Speech and Audio Processing Group, IDIAP Research Institute, Martigny, Switzerland
mohammadzohair2002@gmail.com, aakash07@bu.edu, nidbhavsar989@gmail.com,
(msingh, petr.motlicek)@idiap.ch

Abstract

With the Surge in COVID-19, the number of social media postings related to the vaccine has grown, specifically tracing the confirmed reports by the users regarding the COVID-19 vaccine dose termed as *Vaccine Surveillance*. To mitigate this research problem, we present our novel ensembled approach for self-reporting COVID-19 vaccination status tweets into two labels, namely *Vaccine Chatter* and *Self Report*. We utilize state-of-the-art models, namely BERT, RoBERTa, and XLNet. Our model provides promising results with 0.77, 0.93, and 0.66 as precision, recall, and F1-score (respectively), comparable to the corresponding median scores of 0.77, 0.9, and 0.68 (respectively). The model gave an overall accuracy of 93.43. We also present an empirical analysis of the results to present how well the tweet was able to classify and report. We release our code base here <https://github.com/Zohair0209/SMM4H-2022-Task6.git>

1 Introduction and Motivation

COVID-19 pandemic has unprecedented challenges and also advances to humanity. It provides the scientific community to advance science with wider access to openly available data (Banda et al., 2021). Vaccine surveillance became a pressing research issue with the widespread roll-out of COVID-19 vaccines. Social media platforms such as Twitter and Facebook contain abundant text data that can be utilized for research purposes (Banda et al., 2021).

In this paper, we describe our submission for the shared task ¹, for vaccine surveillance which is a very pressing issue. We aim to target people who report adverse events via their healthcare providers

¹<https://healthlanguageprocessing.org/smm4h-2022/>

to systems like Vaccine Adverse Event Reporting System (VAERS) (Chen et al., 1994), or are found documented in their electronic health record (EHR), a more robust and convenient method could be devised using self-reports from social media. In this task, we experiment with a dataset of Twitter users personally reporting vaccination status and users discussing vaccination status but not revealing their own. This task is tricky because users discuss the vaccination status of others or from news reports in similar ways than they discuss their own at a higher rate (1 to 8 on average). The dataset presents as the positive class, unambiguous tweets of users clearly stating that they have been vaccinated. All other tweets are of users discussing vaccination status. This task involves the identification of self-reported COVID-19 vaccination status in English tweets. As a two-way classification task, the two classes in the training dataset correspond to vaccination confirmation and vaccine-related chatter.

In this task, there is a class imbalance of roughly 1 to 8, implying that 1,496 tweets correspond to vaccination confirmation, and 12,197 tweets correspond to vaccine chatter.

1.1 Proposed System Architecture

This section describes the proposed methodology for classifying tweets as Self-reports or Vaccine Chatter in the context of COVID-19 vaccines. As presented in the 1, we proposed an ensembles approach for this multi-class classification problem. For this purpose, we used multiple pre-trained transformer models, which provide good results based on the dataset. Since transformer models are based on the mechanism of self-attention and differentially weight the significance of each part of the input data, they serve as an ideal option for this task. Initially, split the training dataset in the ratio of 80:20 to execute the classification task

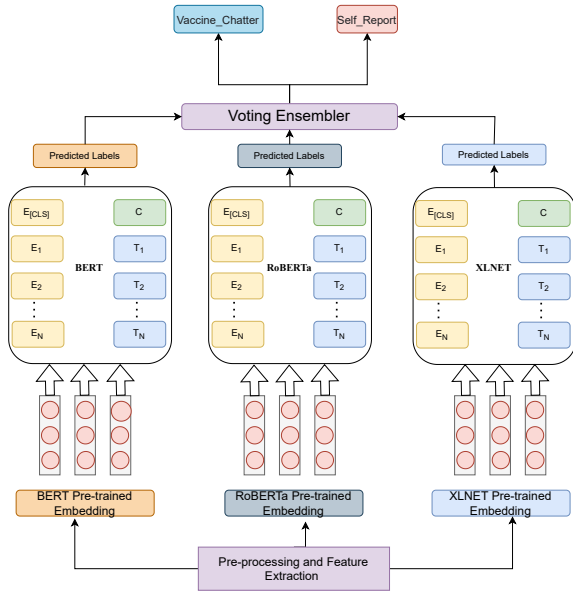


Figure 1: Proposed Architecture

and check the performance of the respective models. We had fine-tuned BERT (Sun et al., 2019), Roberta (Liu et al., 2019), and XLNet (Topal et al., 2021) models, respectively, for this dataset. As a part of the fine-tuning procedure, we tried different combinations of the optimizers, like AdamW and AdaFactor, along with other schedulers, including LinearScheduler, and CosineScheduler, PolyScheduler, etc. The loss function used in the models is Cross-Entropy loss (Ho and Wookey, 2019). Having tried various combinations, we selected the most optimum combination giving the most accurate results. Next, we ensemble these models, incorporating the voting method, which gave the combined label predictions for the respective tweets. Compared to the performance of the ensembles model with the performance of the separate individual models.

1.2 Hyperparameter

Finally, the ensemble model was considered and used for predicting the label class on the test Set as part of this task. We used the Tesla T4 GPU with 16GB RAM alongside 2560 CUDA cores for the experimental setup. We divide the entire dataset into an 80:20 training and validation split of 8 batches, with a learning rate (1e-5) and Adam optimizer (Bock and Weiß, 2019) with epsilon (1e-8). We feed a seed_val of 42. To calculate the training loss over all the batches, we use gradient descents (Ruder, 2016) clipping the norm to 1.0 to avoid exploding gradient problems.

Model	Precision	Recall	F1
Baseline	0.83	0.9	0.77
Median	0.77	0.9	0.68
Ours	0.77	0.93	0.66

Table 1: Results for our proposed ensemble approach with the median and baseline scores

Tweets	Actual	Predict
A few months ago he said masks don't work go figure. He's an idiot. I wonder how many people know we had the coronavirus last year. This one has been (I think) genetically engineered by China to be more lethal. A vaccine this year will not work next yr. Viruses mutate.	Vaccine_chatter	Vaccine_chatter
@DrAnthonyF The new coronavirus has been in the United States for seven months since March. I heard that the vaccine in the United States may not be available until March next year. I believe you and hope you can lead the Americans through the epidemic season.	Vaccine_chatter	Vaccine_chatter
I finally got my Covid vaccine today! And so far, I'm feeling fine. I'm going back on April 16th for my second dose. Covid-Vaccine COVID19Vaccination COVID19Vaccine COVID19	Self_report	Self_report
@GovMikeDeWine Got my second covid-19 vaccine yesterday! Feeling general malaise and sore arm other than that I'm doing fine. The vaccine is safe and effective. The more that are vaccinated the faster we can get back to normal.	Self_report	Self_report
Everyone in my household apart from me has had the vaccine and none of grown extra heads or got covid. Grow up and take it so we can get out of this pandemic and don't let it head into a third wave	Self_report	Vaccine_chatter
Because they haven't got their vaccine yet. Why? Because it's not their turn yet. Until when? Until when? Harini je dah ada a few under thirties without any comorbidity issues died. And yet there are elderly yang tak dapat vaccination date lagi.	Vaccine_chatter	Self_report

Table 2: Result analysis for empirical evaluation of predicted results from our proposed system

2 Results

We describe the results in Table 2. The baseline model achieved 0.83, 0.9, and 0.77 precision, recall, and F1-score (respectively). Our model gave promising results with the precision, recall, and F1-score, standing at 0.77, 0.93, and 0.66 (respectively), which was comparable to the corresponding median scores of 0.77, 0.9, and 0.68 (respectively). The proposed model achieved an overall accuracy of 93.43%.

2.1 Analysis

In this section, we present the empirical analysis of our results in Table: 2. The first and second instances are correctly labeled as Vaccine Chatter, as they do not contain any relevant phrase indicating the user's undertaking of any vaccine dose. These results suggest the perfect predictions of the model. Similarly, the third and fourth instances show that the user has taken a dose of the COVID-19 vaccine. Hence, they have been correctly classified under

the Self Report label, re-emphasizing the correct working of the model.

Contrary to these, the fifth instance was initially labeled as Self Report. However, due to the ambiguous declaration of the undertaking of COVID-19 vaccine dose, the model fails to classify it correctly and instead predicts it under the label of Vaccine Chatter. The sixth instance, on the other hand, was initially labeled as Vaccine Chatter. However, the model mislabels it due to the indecisive phrases featured in the separate tweet.

3 Conclusion

In this paper, we present our system paper submission for Innovators @ SMM4H'22. We aim to classify user tweets, indicating whether they represent Self Reports for the COVID-19 vaccines or are a part of general Vaccine Chatter. The proposed system is an ensembled voting model with fine-tuned BERT, RoBERTa, and XLNet. Given tweets in English, the submitted model classifies each vaccine-related tweet into one of the two labels: "Vaccine Chatter" and "Self Report." The system performs quite well to accomplish the desired task with an accuracy of 93.43%. In the future, we intend to work on a multitask learning framework to handle social media postings related to other medical advancements, apart from the COVID-19 vaccines. We also aim to develop models for multi-lingual postings featuring similar scenarios.

4 Acknowledgements

This work was supported by the European Union's Horizon 2020 research and innovation program under grant agreement No. 833635 (project ROX-ANNE: Real-time network, text, and speaker analytics for combating organized crime, 2019-2022).

References

- Juan M. Banda, Ramya Tekumalla, Guanyu Wang, Jingyuan Yu, Tuo Liu, Yuning Ding, Ekaterina Artemova, Elena Tutubalina, and Gerardo Chowell. 2021. [A large-scale covid-19 twitter chatter dataset for open scientific research—an international collaboration](#). *Epidemiologia*, 2(3):315–324.
- Sebastian Bock and Martin Weiß. 2019. A proof of local convergence for the adam optimizer. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Robert T Chen, Suresh C Rastogi, John R Mullen, Scott W Hayes, Stephen L Cochi, Jerome A Donlon,

and Steven G Wassilak. 1994. The vaccine adverse event reporting system (vaers). *Vaccine*, 12(6):542–550.

Yaoshiang Ho and Samuel Wookey. 2019. The real-world-weight cross-entropy loss function: Modeling the costs of mislabeling. *IEEE Access*, 8:4806–4813.

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. *arXiv preprint arXiv:1907.11692*.

Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.

Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In *China national conference on Chinese computational linguistics*, pages 194–206. Springer.

M Onat Topal, Anil Bas, and Imke van Heerden. 2021. Exploring transformers in natural language generation: Gpt, bert, and xlnet. *arXiv preprint arXiv:2102.08036*.