

Emotion Analysis of Writers and Readers of Japanese Tweets on Vaccinations

Patrick Ramos

Ateneo de Manila University, Philippines
patrick.john.ramos@obf.ateneo.edu

Kiki Ferawati and Kongmeng Liew and Eiji Aramaki and Shoko Wakamiya

Nara Institute of Science and Technology, Japan
{kiki.ferawati.kb6, liew.kongmeng, aramaki, wakamiya}@is.naist.jp

Abstract

Public opinion in social media is increasingly becoming a critical factor in pandemic control. Understanding the emotions of a population towards vaccinations and COVID-19 may be valuable in convincing members to become vaccinated. We investigated the emotions of Japanese Twitter users towards Tweets related to COVID-19 vaccination. Using the WRIME dataset, which provides emotion ratings for Japanese Tweets sourced from writers (Tweet posters) and readers, we fine-tuned a BERT model to predict levels of emotional intensity. This model achieved a training accuracy of $MSE = 0.356$. A separate dataset of 20,254 Japanese Tweets containing COVID-19 vaccine-related keywords was also collected, on which the fine-tuned BERT was used to perform emotion analysis. Afterwards, a correlation analysis between the extracted emotions and a set of vaccination measures in Japan was conducted. The results revealed that surprise and fear were the most intense emotions predicted by the model for writers and readers, respectively, on the vaccine-related Tweet dataset. The correlation analysis also showed that vaccinations were weakly positively correlated with predicted levels of writer joy, writer/reader anticipation, and writer/reader trust. Code will be made available at <https://github.com/PatrickJohnRamos/BERT-Japan-vaccination>.

1 Introduction

Vaccination against COVID-19 has been demonstrated to reduce the spread of the virus (Jones et al., 2021; Hall et al., 2021; Voysey et al., 2021). However, vaccine hesitancy can prevent vaccine uptake, increasing risk of infection. Searching for and understanding causes of vaccine hesitancy can lead to more effective methodologies in convincing community members to become vaccinated. One possible area of understanding vaccine hesitancy is the emotions felt towards vaccines and the COVID-

19 pandemic. For example, fear felt towards vaccination might discourage one from receiving the vaccine. Meanwhile, fear felt towards contracting COVID-19 could encourage one to become vaccinated against it. Leveraging emotions has also been proposed in communication to reduce COVID-19 vaccine hesitancy (Chou and Budenz, 2020).

There have been several attempts at extracting emotions regarding COVID-19 vaccines, particularly from social media. The wealth of information available on social networking services such as Twitter already makes them a popular source for mining sentiment and emotions in other areas such as politics (Bose et al., 2019) and consumerism (Rathan et al., 2018). Social media information extraction has seen continued research during the COVID-19 pandemic, with multiple works specifically seeking to mine sentiments and emotions surrounding the pandemic and vaccination (Boon-Itt and Skunkan, 2020; Sakti et al., 2021; Aygün et al., 2021; Niu et al., 2022).

Existing emotion analysis studies on COVID-19-related Japanese Twitter corpora only focus on the emotions of writers, or those who post Tweets (Lee et al., 2020; Bashar, 2021). However, the emotions of readers, or those who read Tweets, are not necessarily the same as those of writers. For instance, if a writer expresses disgust towards vaccination, a reader might express anger out of disagreement in response. These reader emotions may also contain useful information in understanding vaccine hesitancy. We contribute to this research area by using a BERT to extract emotions of both writers and readers towards Tweets related to COVID-19 vaccines and comparing the predicted emotions to vaccination uptake. We do this by fine-tuning a BERT (Devlin et al., 2019) to predict intensity scores for Plutchik’s eight emotions (Plutchik, 1980) for writers and readers from Tweets containing keywords related to COVID-19 vaccination, and performing a correlation analysis between the mined emotions

and a set of vaccination measures in Japan. Note that one limitation of our study is that the dataset of COVID-19 vaccine-related Tweets does not guarantee that COVID-19 or vaccination are the topics of the texts, or are necessarily the object of any inferred emotion.

We find that the emotion most prominently predicted by the model for writers on the vaccine-related Tweet dataset is surprise, with fear being the most intensely predicted emotion for readers. Additionally, writer joy, writer/reader anticipation, and writer/reader trust are weakly positively correlated with vaccinations.

2 Related Work

Prior to the adoption of deep learning, emotion analysis of Tweets was performed with a combination of feature engineering, lexicon-based approach, and traditional off-the-shelf classifiers. Balabantaray et al. (2012) and Wang et al. (2012) engineered overlapping sets of features from Tweets, with shared features including n-grams, POS, adjectives, and lexicon-based sentiment polarity scores. Balabantaray et al. (2012) fed the features to an SVM while Wang et al. (2012) used linear and Naive Bayes classifiers. EmpaTweet (2012) used a similar set of features and also an SVM, but exchanged sentiment polarity scores for synonym rings, hypernyms, and LDA topic scores.

However, these classical methods have been outperformed by more contemporary and dedicated sequence modelling techniques such as RNNs and LSTMs. Vateekul and Koomsubha (2016) demonstrated the superiority of LSTMs in emotion analysis over SVMs and Naive Bayes on Thai Twitter text, while Colnerič and Demšar (2018) showed the effectiveness of character-based RNNs.

The introduction of Transformers (Vaswani et al., 2017) as the new state-of-the-art sequence modelling architecture and their increasing ubiquity has also led to their application in social media emotion analysis for a variety of languages. BERT models can outperform CNNs and BiLSTMs on English Twitter (Harb et al., 2020); Naive Bayes, logistic regression, and SVMs on Romanian Twitter (Ciobotaru and Dinu, 2021); and TF-IDF and word2vec on Arabic Twitter (Al-Twairesh, 2021). During the COVID-19 pandemic, a RoBERTa (Liu et al., 2019) was fine-tuned on emotion analysis to classify the emotions of Tweets containing hashtags related to the pandemic (Choudrie et al., 2021).

Meanwhile, emotion analyses of Japanese Tweets related to COVID-19 have been conducted using traditional techniques such as lexicon-based methods (Bashar, 2021) and frequency analysis (Lee et al., 2020).

The work most similar to ours is that of Niu et al. (2022), who perform sentiment analysis of COVID-19-related Japanese Tweets; investigate the correlation of the mined sentiments with infections, deaths, and vaccinations; and conduct additional analyses comparing multiple vaccine brands. Our study differs from this by extracting emotions rather than sentiments from both writers and readers, and comparing these mined features to vaccinations, vaccinated people, and fully vaccinated people.

3 Method

3.1 Dataset

To gauge the emotions of the Japanese public towards COVID-19 vaccines and related topics, a dataset of 20,254 vaccine-related Tweets from December 2021 containing any of the keywords “ワクチン” (“vaccine”), “モデルナ” (“Moderna”), “ファイザー” (“Pfizer”), or “オミクロン” (“Omicron”) was constructed. “Moderna” and “Pfizer” were specifically selected as keywords as these are brands of COVID-19 vaccines commonly administered in Japan. The dataset was created by sampling 15 random minutes from each day of December 2021 for each keyword, and scraping all Tweets containing the assigned keyword for each sampled minute. A distribution of the dataset according to keyword is shown in Table 1.

Table 2 compares our constructed dataset against WRIME (Kajiwara et al., 2021), the fine-tuning dataset we discuss in Section 3.2. Our dataset contains longer texts on average. Postpositional particles, nouns, punctuation marks, and verbs are among the most common parts of speech in both datasets. However, auxiliary verbs are more common in the vaccine-related Tweet dataset while non-punctuation symbols are more common in WRIME.

3.2 Fine-tuning BERT for Emotion Analysis

We used BERT to perform emotion analysis by fine-tuning it to extract writer and reader emotional intensity scores from text, with higher scores indicating higher intensities. Writer emotions refer to emotions felt by the writers of a Tweet as they

Keyword	Number of Tweets
vaccine	13,664
Moderna	349
Pfizer	1,352
Omicron	7,761

Table 1: Number of Tweets per keyword in the vaccine Tweet dataset. Keywords are translated from Japanese. Some Tweets contain multiple keywords.

	Vaccine-related Tweets	WRIME
Average # of tokens per text	51	23
Five most common POS (descending)	ADP, AUX, NOUN, PUNCT, VERB	ADP, NOUN, PUNCT, SYM, VERB

Table 2: Average number of tokens and the five most common parts of speech using UniDic.¹

write the post, while reader emotions refer to the emotions felt by the readers of a Tweet as they read it. We followed Plutchik’s (1980) framework and fine-tuned BERT to extract intensity scores for eight emotions: joy, sadness, anticipation, surprise, anger, fear, disgust, and trust.

Using BERT for emotion analysis was straightforward and required only a simple modification to the head of the BERT model. First, input texts were tokenized using MeCab (Kudo et al., 2004) and WordPiece (Wu et al., 2016). The tokenized inputs were then each prepended with [CLS] classification tokens and fed through the BERT model. After the last layer, each [CLS] token was linearly projected into 16 class scores, representing the 8 emotions of writers and readers.

A PyTorch (Paszke et al., 2019) implementation of BERT_{BASE} (110M parameters) pre-trained on Japanese Wikipedia from HuggingFace (Wolf et al., 2020)² was fine-tuned on WRIME (Kajiwara et al., 2021). WRIME is a dataset for emotional intensity estimation comprised of 43,200 Japanese Tweets annotated with Plutchik’s 8 emotions by the posts’ writers and 3 “reader” annotators hired by the dataset authors to read the posts.³ Each emotion is annotated across 4 levels (0 to 3) of increasing intensity, with 0 referring to no presence and 3 re-

²We use the `cl-tohoku/bert-base-japanese-v2` Japanese BERT checkpoint available at <https://huggingface.co/cl-tohoku/bert-base-japanese-v2>.

³As of February 2022, there are two versions of this dataset. We use Version 1. WRIME is available at <https://github.com/ids-cv/wrime>.

ferring to strong intensity. To create only a single set of reader emotion scores per data point, we averaged the scores across the three readers. BERT was then trained with mean squared error loss to directly predict the emotional intensity scores of writers and readers for given Tweets. A sample data point can be seen in Table 3.

Our fine-tuning was inspired by common BERT fine-tuning procedures (Devlin et al., 2019). We fine-tuned BERT on WRIME for 3 epochs using the AdamW optimizer (Loshchilov and Hutter, 2017) with a learning rate of $2e-5$, $\beta_1=0.9$, $\beta_2=0.999$, weight decay of 0.01, linear decay, a warmup ratio of 0.01, and a batch size of 32. Training was conducted with an NVIDIA Tesla K80 and finished in 3 hours. The fine-tuned BERT was dubbed “emotion analysis BERT.”

We evaluated emotion analysis BERT by comparing its mean squared errors to those of two baselines based on Kajiwara et al. (2021). The first was a bag-of-words and linear regression model (BoW+LinReg). Each text was tokenized using MeCab, vectorized into a bag-of-words using the 2000 most common words in the vocabulary, and then fed into a linear regression model. While Kajiwara et al. (2021) used logistic regression, we used linear regression for a fairer comparison with emotion analysis BERT, which directly predicts intensity scores. The second model we compared to uses fastText (Bojanowski et al., 2017) and an SVM (fastText+SVM). Each word of was embedded using fastText, after which the average of all the embeddings in the sequence were used as input into an SVM that regresses the emotional intensity scores. We used a linear kernel and $C = 100$.

3.3 Inferring Emotions from Vaccine-related Tweets

Emotions from the vaccine-related Tweet dataset were extracted using emotion analysis BERT following the procedure described in Section 3.2. Tweets were tokenized, prepended with a [CLS] token, and processed through emotion analysis BERT, with the [CLS] token being projected into writer and reader emotion scores after the last layer. Note that this is purely inference and no training is done using the vaccine Tweet data.

Tweet	車のタイヤがパンクしてた。。いたずらの可能性が高いんだって。。(The tire of my car was flat. I heard that it might be mischief.)								
Annotator	Joy	Sadness	Anticipation	Surprise	Anger	Fear	Disgust	Trust	
Writer	0	3	0	1	3	0	0	0	
Reader	0	2.33	0	2.33	0.33	1	0.67	0	

Table 3: Sample training data point. While WRIME has annotations for three separate readers, we create only one set by averaging the reader scores per emotion.

Model	Joy	Sadness	Anticipation	Surprise	Anger	Fear	Disgust	Trust	Overall
Writer Emotions									
BoW+LinReg	0.889	0.830	0.849	0.617	0.605	0.536	0.822	0.399	0.692
fastText+SVM	1.141	0.973	1.066	0.801	0.754	0.628	1.032	0.394	0.849
BERT (ours)	0.658	0.688	0.746	0.542	0.486	0.462	0.664	0.400	0.581
Reader Emotions									
BoW+LinReg	0.351	0.270	0.344	0.172	0.049	0.201	0.175	0.037	0.200
fastText+SVM	0.374	0.297	0.422	0.177	0.047	0.269	0.221	0.090	0.237
BERT (ours)	0.192	0.178	0.211	0.139	0.032	0.147	0.123	0.029	0.131
Both Emotions									
BoW+LinReg	0.615	0.550	0.597	0.394	0.327	0.368	0.499	0.218	0.446
fastText+SVM	0.758	0.635	0.744	0.489	0.400	0.449	0.626	0.242	0.543
BERT (ours)	0.425	0.433	0.479	0.341	0.259	0.304	0.394	0.214	0.356

Table 4: Mean squared errors for each emotion in the test split of WRIME.

4 Experimental Results and Discussion

4.1 Fine-tuning results

Results for emotional intensity prediction on WRIME are shown in Table 4. The only emotion for which we do not achieve the best mean squared error is writer trust, where we still remain competitive a difference of 0.001 MSE from the best performing model. For all other emotions for both writers and readers, emotion analysis BERT outperforms both of the other models. We achieve 0.111 MSE and 0.106 MSE improvements on writer and reader emotion prediction respectively. For overall emotional intensity prediction, we outperform the next best model by 0.090 MSE.

Table 5 presents inferred emotions for sample test entries from WRIME. We qualitatively examined these results, which showed that emotion analysis BERT is capable of detecting emotion even when emotions are not explicitly stated, like they would be in a sentence such as “I am joyful.” The model was able to infer anger and disgust from texts including phrases like “I waited for an hour but no one came!!”, and joy and anticipation from “Haven’t had back-to-back holidays in a while. Tomorrow is a weekday break!!”

4.2 Emotion inference results

Figure 1 presents the distributions of emotions for both writers and readers predicted by emotion anal-




Figure 1: Box plots of distributions of predicted emotion scores for writers and readers. Whiskers are drawn at $1.5 \times$ inter-quartile range from the first and third quartiles. Mean predicted scores are indicated by black points.

ysis BERT. The highest average predicted score was surprise for writers and fear for readers. With the exception of fear, inferred writer scores were more intense, especially in sadness, anger, and trust, where median inferred writer scores were greater than the third quartile of inferred reader scores. Why the emotion intensities predicted for writers differed from those assumed for readers could be of interest to future studies.

The distributions of predicted writer and reader emotions are compared in a Q-Q plot in Figure 2. Anticipation and fear followed similar distributions for both writer and reader inferences, while other emotions showed lower values for readers, especially for trust.

Tweet	一時間待ってもこない！！ガス会社なんなの！！急いで帰ってきたのに？時間守らないひときらい！！！！なんか食べ物でも持ってきたら許そう。(I waited for an hour but no one came!! What’s up with this gas company!! Even though I hurried back? I hate people who can’t keep track of time!!! Maybe I’ll forgive them though if they bring me food.)								
Annotator	Joy	Sadness	Anticipation	Surprise	Anger	Fear	Disgust	Trust	
Writer	0.019	1.206	-0.069	0.785	2.424	0.070	2.224	-0.106	
Reader	0.049	0.648	0.120	0.657	1.573	0.369	2.139	-0.060	
Tweet	久しぶりの2連休。明日平日休み！！(Haven’t had back-to-back holidays in a while. Tomorrow is a weekday break!!)								
Annotator	Joy	Sadness	Anticipation	Surprise	Anger	Fear	Disgust	Trust	
Writer	2.989	0.171	2.111	0.469	0.107	0.069	0.126	0.827	
Reader	2.396	0.049	1.567	0.241	-0.011	0.022	-0.005	0.023	

Table 5: Inferred emotion scores for data points from the WRIME test set. English translations are provided in parentheses after each Japanese Tweet. The two most intense emotions per annotator for each Tweet are in bold.




Figure 2: Q-Q plots of predicted writer and reader emotions.

The percentage distributions of emotions for writers and readers inferred by the model for all keywords and per keyword are shown in Figure 3. While predicted writer emotion scores tended to be higher, this did not necessarily mean the same for the proportion each emotion constituted of the sum of all emotion scores. Although writers had higher predicted average anticipation, surprise, and disgust, these emotions comprised a larger percentage of inferred reader emotions than they did writer emotions. Fear, which was predicted at higher levels for readers, also constituted a larger proportion of total inferred reader emotion.

Predicted emotions of Tweets referencing Moderna were compared to those of Tweets referencing Pfizer to identify any differences in emotions towards the two brands of vaccines. The results of a Kolmogorov-Smirnov test for the intensities of emotions towards the two brands are presented in Table 6. Only differences in predictions for writer joy (Pfizer higher), writer sadness (Moderna higher), and reader sadness (Moderna higher) could be considered attributable to differences in




Figure 3: Distributions of predicted emotions for writers and readers for all keywords and for each keyword. Emotion proportions were calculated by dividing the total score of each emotion by the sum of all emotion scores for the emotion’s assigned subject (writer or reader).

underlying distributions.

Sample emotion inferences can be seen in Table 7. Like with the qualitative results from WRIME, emotion analysis BERT was capable of detecting emotion even in a lack of explicitly mentioned emotions. The model could identify anger and disgust from Tweets containing phrases such as “I can’t call it anything other than foolish,” and trust from “Today, I finally got vaccinated. My arm hurts a little bit, but other than that there are no problems.”

Emotion	KS statistic	p-value
Writer Joy	0.099	0.008
Writer Sadness	0.098	0.008
Writer Anticipation	0.069	0.135
Writer Surprise	0.067	0.153
Writer Anger	0.054	0.379
Writer Fear	0.073	0.101
Writer Disgust	0.039	0.765
Writer Trust	0.065	0.188
Reader Joy	0.073	0.099
Reader Sadness	0.124	0.0003
Reader Anticipation	0.066	0.165
Reader Surprise	0.065	0.187
Reader Anger	0.063	0.207
Reader Fear	0.069	0.134
Reader Disgust	0.055	0.350
Reader Trust	0.066	0.163

Table 6: Kolmogorov-Smirnov test results for emotions between Tweets referencing Moderna and Pfizer. Significant KS statistics and p-values are in bold.

4.3 Comparison to vaccinations in Japan

Vaccination data was taken from Our World in Data (Mathieu et al., 2021)⁴. The vaccination data was comprised of a set of periodic vaccination measures for several countries. We filtered the dataset to only include entries from Japan in December 2021, matching the collection period of our vaccine-related Tweet dataset. For each date in the dataset, we focused on vaccinations, or the number of vaccinations administered on that date; new people vaccinated, or the number of people who received their first dose on said date; and new people fully vaccinated, or the number of people who received their second dose on that date. One thing to note is that vaccinations have plateaued starting November 2021, resulting from a lower number of vaccinations and a slower uptake of boosters (Mathieu et al., 2021).

Figure 4 compares the average predicted score for each emotion to the number of new vaccinations, people vaccinated, and people fully vaccinated across December 2021. No easily discernible trend common to both emotion scores and vaccination metrics was found.

For each writer and reader emotion, we performed a correlation analysis with each vaccination

⁴Vaccination data is available at <https://github.com/owid/covid-19-data/tree/master/public/data/vaccinations>.

measure by taking the Pearson correlation coefficient between the sums of the particular emotion’s predicted intensities and their corresponding vaccination metric for each date recorded in the vaccination dataset. The results are presented in Figure 5. The only results with satisfactory p-values were under vaccinations, which were positive correlated with writer joy ($r = 0.36, p = 0.047$), writer anticipation ($r = 0.40, p = 0.027$), writer trust ($r = 0.40, p = 0.025$), reader anticipation ($r = 0.44, p = 0.013$), and reader trust ($r = 0.39, p = 0.031$). We did not observe any significant correlations between the predicted emotion scores and people vaccinated or people fully vaccinated, as all p-values for these results were greater than or equal to 0.157, which is above the alpha of 0.05.

5 Broader Impact

Research into understanding the emotions of a population towards vaccines could hold both positive and negative societal impacts. Any relationship discovered between emotions and vaccinations could be leveraged in campaigns aimed at convincing citizens to receive vaccinations, reducing the spread of COVID-19. On the other hand, knowledge of what emotions could affect vaccine acceptance can be used in efforts to increase vaccination hesitancy and slow down vaccination rates, which could prolong risks of infection.

6 Conclusion

We fine-tuned a BERT on the task of emotion analysis, and used the emotion analysis BERT to infer emotion intensities of writers and readers from a corpus of Tweets containing keywords related to COVID-19 and COVID-19 vaccination. Our results revealed that surprise and fear were respectively the most intensely predicted emotions for writers and readers. Furthermore, vaccinations were weakly positively correlated with writer joy, writer anticipation, writer trust, reader anticipation, and reader trust.

Future works can extend this study by designing the emotion analysis to be aspect-based with respect to the keywords, as it is possible that the keywords are not the objects of the inferred emotions. Another possible area for further research could be correlation analyses with other COVID-19 metrics, such as infections and reproduction rate.

Tweet	バカとしか言いようがない。アメリカ帰りの女が自宅待機期間に男と会って、その男がサッカー場で撒き散らす。あああ、バカだなあ、オミクロン株接触者 天皇杯観戦 2021年12月16日 (I can't call it anything other than foolish. A woman who returned from the United States met with a man while she was still in quarantine, then that man went to a soccer game and infected others. Ah, it really is foolish. There are close contacts with the Omicron strain among spectators at the Emperor's Cup on December 16, 2021.)							
Annotator	Joy	Sadness	Anticipation	Surprise	Anger	Fear	Disgust	Trust
Writer	0.180	0.675	0.214	1.250	2.138	0.448	2.077	0.259
Reader	-0.030	0.648	0.104	0.762	1.147	0.575	1.848	0.090
Tweet	皆様、お疲れ様です。今日はようやくワクチン接種しました。腕がちょっと痛いですが、それ以外に問題はありません。神宮球場にもちょこっと顔を出してノムさんを偲んできました。皆様、身の安全確保と体調管理に気をつけて、お過ごし下さい。今日も一日お疲れ様でした。(Thank you everyone for your hard work. Today, I finally got vaccinated. My arm hurts a little bit, but other than that there are no problems. I also swung by Jingū Stadium for a bit to think about Mr. Nomu. Everyone, please look after your safety and health. Thank you for working hard today.)							
Annotator	Joy	Sadness	Anticipation	Surprise	Anger	Fear	Disgust	Trust
Writer	2.020	0.428	1.321	0.594	0.075	0.145	0.014	2.069
Reader	1.582	0.100	0.827	-0.089	-0.038	0.239	0.046	0.419

Table 7: Inferred emotion scores for data points from the constructed vaccine-related Tweet dataset. English translations are provided in parentheses after each Japanese Tweet. Links are removed. The two most intense emotions per annotator for each Tweet are in bold.

Acknowledgements

This work was supported by JST SICORP Grant Number JPMJSC2107, Japan.

References

- Nora Al-Twairish. 2021. [The evolution of language models applied to emotion analysis of arabic tweets](#). *Information*, 12(2):84.
- İrfan Aygün, Buket Kaya, and Mehmet Kaya. 2021. [Aspect based twitter sentiment analysis on vaccination and vaccine types in covid-19 pandemic with deep learning](#). *IEEE Journal of Biomedical and Health Informatics*.
- Rakesh C Balabantaray, Mudasir Mohammad, and Nibha Sharma. 2012. [Multi-class twitter emotion classification: A new approach](#). *International Journal of Applied Information Systems*, 4(1):48–53.
- Md Khayrul Bashar. 2021. [Exploration of public emotion dynamics in japan from twitter data during covid-19](#). In *2021 6th International Conference on Intelligent Informatics and Biomedical Sciences (ICIIBMS)*, volume 6, pages 310–315. IEEE.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Sakun Boon-Itt and Yukolpat Skunkan. 2020. [Public perception of the covid-19 pandemic on twitter: Sentiment analysis and topic modeling study](#). *JMIR Public Health and Surveillance*, 6(4):e21978.
- Rajesh Bose, Raktim Kumar Dey, Sandip Roy, and Debabrata Sarddar. 2019. [Analyzing political sentiment using twitter data](#). In *Information and communication technology for intelligent systems*, pages 427–436. Springer.
- Wen-Ying Sylvia Chou and Alexandra Budenz. 2020. [Considering emotion in covid-19 vaccine communication: addressing vaccine hesitancy and fostering vaccine confidence](#). *Health communication*, 35(14):1718–1722.
- Jyoti Choudrie, Shruti Patil, Ketan Kotecha, Nikhil Matta, and Ilias Pappas. 2021. [Applying and understanding an advanced, novel deep learning approach: A covid 19, text based, emotions analysis study](#). *Information Systems Frontiers*, 23(6):1431–1465.
- Alexandra Ciobotaru and Liviu P. Dinu. 2021. [RED: A novel dataset for Romanian emotion detection from tweets](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 291–300, Held Online. INCOMA Ltd.
- Niko Colnerič and Janez Demšar. 2018. [Emotion recognition on twitter: Comparative study and training a](#)




Figure 4: Average score for each emotion and the number of new vaccinations, people vaccinated, and people fully vaccinated for several days in December 2021. Vaccination data missing due to gaps in the dataset is represented with broken lines.

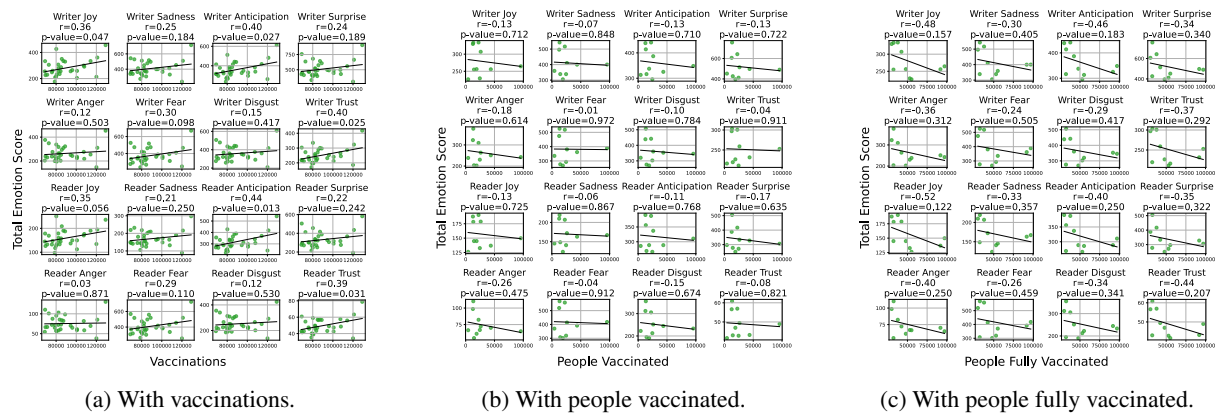


Figure 5: Correlation analysis of emotion scores and vaccination statistics. Emotion scores are scatter plotted against vaccination numbers with best fit lines. No significant correlations are found for emotions with people vaccinated and people fully vaccinated.

unison model. *IEEE transactions on affective computing*, 11(3):433–446.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Victoria Jane Hall, Sarah Foulkes, Ayoub Saei, Nick Andrews, Blanche Oguti, Andre Charlett, Edgar Wellington, Julia Stowe, Natalie Gillson, Ana Atti, et al. 2021. **Covid-19 vaccine coverage in health-care workers in england and effectiveness of bnt162b2 mrna vaccine against infection (siren): a prospective, multicentre, cohort study**. *The Lancet*, 397(10286):1725–1735.

Jonathas GD Harb, Régis Ebeling, and Karin Becker. 2020. **A framework to analyze the emotional reactions to mass violent events on twitter and influential factors**. *Information Processing & Management*, 57(6):102372.

Nick K Jones, Lucy Rivett, Shaun Seaman, Richard J Samworth, Ben Warne, Chris Workman, Mark Ferris, Jo Wright, Natalie Quinnell, Ashley Shaw, et al. 2021. **Single-dose bnt162b2 vaccine protects against asymptomatic sars-cov-2 infection**. *Elife*, 10:e68808.

Tomoyuki Kajiwar, Chenhui Chu, Noriko Take-mura, Yuta Nakashima, and Hajime Nagahara. 2021. **WRIME: A new dataset for emotional intensity estimation with subjective and objective annotations**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2095–2104, Online. Association for Computational Linguistics.

- Taku Kudo, Kaoru Yamamoto, and Yuji Matsumoto. 2004. [Applying conditional random fields to Japanese morphological analysis](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 230–237, Barcelona, Spain. Association for Computational Linguistics.
- Hocheol Lee, Eun Bi Noh, Sea Hwan Choi, Bo Zhao, and Eun Woo Nam. 2020. [Determining public opinion of the covid-19 pandemic in south korea and japan: social network mining on twitter](#). *Healthcare Informatics Research*, 26(4):335–343.
- 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*.
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#). *arXiv preprint arXiv:1711.05101*.
- Edouard Mathieu, Hannah Ritchie, Esteban Ortiz-Ospina, Max Roser, Joe Hasell, Cameron Appel, Charlie Giattino, and Lucas Rod s-Guirao. 2021. [A global database of covid-19 vaccinations](#). *Nature human behaviour*, 5(7):947–953.
- Qian Niu, Junyu Liu, Masaya Kato, Yuki Shinohara, Natsuki Matsumura, Tomoki Aoyama, and Momoko Nagai-Tanima. 2022. [Public opinion and sentiment before and at the beginning of covid-19 vaccinations in japan: Twitter analysis](#). *medRxiv*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). *Advances in neural information processing systems*, 32.
- Robert Plutchik. 1980. [A general psychoevolutionary theory of emotion](#). In *Theories of emotion*, pages 3–33. Elsevier.
- M Rathan, Vishwanath R Hulipalled, KR Venugopal, and LM Patnaik. 2018. [Consumer insight mining: aspect based twitter opinion mining of mobile phone reviews](#). *Applied Soft Computing*, 68:765–773.
- Kirk Roberts, Michael A. Roach, Joseph Johnson, Josh Guthrie, and Sanda M. Harabagiu. 2012. [EmpaTweet: Annotating and detecting emotions on Twitter](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 3806–3813, Istanbul, Turkey. European Language Resources Association (ELRA).
- Andi Muhammad Tri Sakti, Emma Mohamad, and Arina Anis Azlan. 2021. [Mining of opinions on covid-19 large-scale social restrictions in indonesia: public sentiment and emotion analysis on online media](#). *Journal of medical Internet research*, 23(8):e28249.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in neural information processing systems*, pages 5998–6008.
- Peerapon Vateekul and Thanabhat Koomsubha. 2016. [A study of sentiment analysis using deep learning techniques on thai twitter data](#). In *2016 13th International joint conference on computer science and software engineering (JCSSE)*, pages 1–6. IEEE.
- Merryn Voysey, Sue Ann Costa Clemens, Shabir A Madhi, Lily Y Weckx, Pedro M Folegatti, Parvinder K Aley, Brian Angus, Vicky L Baillie, Shaun L Barnabas, Qasim E Bhorat, et al. 2021. [Single-dose administration and the influence of the timing of the booster dose on immunogenicity and efficacy of chadox1 ncov-19 \(azd1222\) vaccine: a pooled analysis of four randomised trials](#). *The Lancet*, 397(10277):881–891.
- Wenbo Wang, Lu Chen, Krishnaprasad Thirunarayan, and Amit P Sheth. 2012. [Harnessing twitter’ big data’ for automatic emotion identification](#). In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*, pages 587–592. IEEE.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *arXiv preprint arXiv:1609.08144*.