

# VaxxHesitancy: A Dataset for Studying Hesitancy towards COVID-19 Vaccination on Twitter

Yida Mu, Mali Jin, Charlie Grimshaw, Carolina Scarton, Kalina Bontcheva, Xingyi Song

Department of Computer Science, The University of Sheffield, UK  
{y.mu, m.jin, cgrimshaw1, c.scarton, k.bontcheva, x.song}@sheffield.ac.uk

## Abstract

Vaccine hesitancy has been a common concern, probably since vaccines were created and, with the popularisation of social media, people started to express their concerns about vaccines online alongside those posting pro- and anti-vaccine content. Predictably, since the first mentions of a COVID-19 vaccine, social media users posted about their fears and concerns or about their support and belief into the effectiveness of these rapidly developing vaccines. Identifying and understanding the reasons behind public hesitancy towards COVID-19 vaccines is important for policy makers that need to develop actions to better inform the population with the aim of increasing vaccine take-up. In the case of COVID-19, where the fast development of the vaccines was mirrored closely by growth in anti-vaxx disinformation, automatic means of detecting citizen attitudes towards vaccination became necessary. This is an important computational social sciences task that requires data analysis in order to gain in-depth understanding of the phenomena at hand. Annotated data is also necessary for training data-driven models for more nuanced analysis of attitudes towards vaccination. To this end, we created a new collection of over 3,101 tweets annotated with users' attitudes towards COVID-19 vaccination (stance). Besides, we also develop a domain-specific language model (VaxxBERT) that achieves the best predictive performance (73.0 accuracy and 69.3 F1-score) as compared to a robust set of baselines. To the best of our knowledge, these are the first dataset and model that model vaccine hesitancy as a category distinct from pro- and anti-vaccine stance.

## Introduction

When the first COVID-19 vaccine was publicly administered in the UK in the end of 2020,<sup>1</sup> most counties also begun to promote vaccination to the public. A high vaccination uptake is considered the most reliable and effective way to contain the COVID-19 pandemic and protect high-risk groups, since these vaccines have been shown to prevent serious illnesses caused by the SARS-CoV-2 virus.<sup>2</sup>

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup><https://www.england.nhs.uk/2020/12/landmark-moment-as-first-nhs-patient-receives-covid-19-vaccination/>

<sup>2</sup><https://www.pfizer.com/news/press-release/press-release-detail/pfizer-and-biontech-confirm-high-efficacy-and-no-serious>

Nevertheless, despite the significant impact of COVID-19 on peoples lives, including restrictions on free movement and the requirement to use masks, the vaccines raised also concerns in a significant number of citizens worldwide with corresponding negative impact on countries' vaccination rates. One of the reasons is vaccine disinformation online (Muric et al. 2021; Gisondi et al. 2022). Examples can be found easily in social media, where an "anti-vaxxer" (i.e. a user promoting narratives against vaccination) posts disinformation (e.g., false statistics or conspiracies) about vaccines to discourage their followers from getting vaccinated (Loomba et al. 2021; Jennings et al. 2021). Common topics of such disinformation are usually related to politics and religion in an attempt to deliberately deceive the public (Rzym-ski et al. 2021; van der Linden et al. 2021). Some real-world examples appear in Table 1 (see Tweets 2 & 3) and there are already research datasets (Cotfas et al. 2021; Chen, Chen, and Pang 2022) that have collected and categorized such anti-vaxx posts.

On the other hand, medical professionals and some governments<sup>3</sup> recognised that some citizens can be hesitant towards vaccination, but through suitably targeted public health information campaigns they might be persuaded to take it up. Crucially also vaccine hesitant citizens are usually not actively trying to stop others from being vaccinated by spreading anti-vaxx content. The most common reasons behind hesitancy tend to arise from limited understanding of the way COVID-19 vaccines work, e.g. new biotechnologies such as mRNA, side effects, natural immunity (Poddar et al. 2022a,b).

Although previous work has also classified such posts (e.g., Tweets 5 & 6 in Table 1) as anti-vaxx (Cotfas et al. 2021), we argue that the hesitant stance in these posts is different from that of anti-vaxx posts, which are spreading false information. The ability to monitor and understand citizens' unanswered questions and concerns about vaccines is critical for policymakers and vaccine developers, as it enables them to undertake better targeted information campaigns and other actions aimed at improving trust in government policies. In addition, the separation of vaccine hesitant posts

<sup>3</sup>For example, the UK one: <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthandwellbeing>

#	Definitions	Examples
<b>Pro</b>	<i>The tweet expresses opinions and actions supporting COVID-19 vaccination use.</i>	<i>T1: Coronavirus vaccine developed by Oxford University appears safe and trains the immune system. HTTPURL</i>
<b>Anti</b>	<i>The tweet expresses opinions and actions against COVID-19 vaccination, with the aim of, persuading others to refuse vaccination. Please be aware that tweets expressing the person’s own intention to refuse vaccination themselves belong to the “Vaccine Hesitant” category.</i>	<i>T2: Will it alter DNA &amp; remote control ppl thru 5g ? Yes T3: @USER BREAKING NEWS: #Pfizer vaccine Impact on Fertility... T4: the vaccine was a set up for the micro chips, Trump going be president again. After Bidens 4yrs up</i>
<b>Hesitancy</b>	<i>The tweet is centered on the person’s intention to delay and/or refuse the vaccine and is often in first person.</i>	<i>T5: Why is there even a vaccine for an illness with a such high survival rate? T6: I think it came out about 8 months after the pandemic hit and I was like absolutely no, that’s too quick. I would I would not be happy with that.</i>
<b>Irrelevant</b>	<i>Any tweet that does not express stance towards the COVID-19 vaccine, e.g., tweets about other kinds of vaccines or tweets primarily about COVID-19 or other aspects of the pandemic.</i>	<i>T7: @USER What are the recent news of COVID vaccine?</i>

Table 1: Category definitions and tweet examples.

Dataset	Time	Tweets	Labels	Language
Cotfas et al. (2021)	Nov 2020 ~Dec 2020	2,792	in favor, against, neutral	en
Poddar et al. (2022a)	Jan 2020 ~March 2021	1,700	in favor, against, neutral	en
Chen, Chen, and Pang (2022)	Jan 2020 ~March 2021	17,934	positive, negative, neutral, off-topic	fr, de, etc.
Di Giovanni et al. (2022)	Nov 2020 ~ June 2021	3,000	pro, anti, neutral, out-of-context	es, de, it.
Delcea et al. (2022)	July 2021 ~ Aug 2021	4,636	in favor, against, neutral	en
<b>VaxxHesitancy</b>	<b>Nov 2020 ~April 2022</b>	3,101	<b>pro, anti, hesitancy, irrelevant</b>	en

Table 2: Differences in specifications between previous datasets v.s. ours. Note that the ‘off-topic’ and ‘out-of-context’ labels denote the tweet is irrelevant to COVID-19 vaccines or vaccination in Chen, Chen, and Pang (2022) and Di Giovanni et al. (2022); while the ‘irrelevant’ label in our dataset is the same as the ‘Neutral’ label in existing datasets, i.e., no explicit attitude or intention is expressed.

from anti-vaxx disinformation benefits the content moderation and debunking efforts of fact-checkers and social media platforms.

Specifically within the field of computational social science, prior data-driven research on user attitudes towards COVID-19 vaccination (Müller and Salathé 2019; Cotfas et al. 2021; Poddar et al. 2022a) has considered vaccine stance of social media posts as a three-way classification task, namely: pro-vaxx, anti-vaxx and neutral.<sup>4</sup>

However, we argue that a more nuanced understanding of online debates on vaccination is needed, especially with respect to detecting and tracking the concerns voiced by vaccine hesitant citizens. To this end, we introduce a new, separate stance category of **vaccine hesitancy** which is centered on the person’s intention to delay or refuse the vaccine and is usually expressed in first person.

To the best of our knowledge, we are the first to consider

<sup>4</sup>Some of these categories may be named differently by different researchers. For consistency, we use pro-vaxx to encompass all similar labels used previously, e.g., *in favor*, *pro*, *positive*.

COVID-19 **vaccine hesitancy** as a separate stance category. The main contributions of this paper include:

- We re-frame the task of vaccine stance classification into a **four-way classification** setup, by considering vaccine hesitancy as a separate stance category.
- For this new, more nuanced classification task, we create a **publicly available dataset**<sup>5</sup> of over 3,100 COVID-19 vaccine-related tweets labelled as belong to one of four stance categories: *pro-vaxx*, *anti-vaxx*, *vaxx-hesitant*, or *irrelevant*, which also spans the longest time period. The characteristics of our dataset are compared to those existing three-way classification ones in Table 2.
- We make available a **domain-specific language models (VaxxBERT)**, pre-trained on 175 million unlabelled COVID-19 vaccine-related tweets.
- We also provide a **linguistic analysis**, which highlights the difference in language patterns between *anti-vaxx* and *vaxx-hesitant* posts.

<sup>5</sup><https://github.com/GateNLP/VaxxHesitancy>

## Related Work

The development of social networks brought an advent in natural language processing (NLP) research that aimed to automatically detect users' attitudes towards a topic (Augenstein et al. 2016; Derczynski et al. 2017; Gorrell et al. 2019). Particularly relevant to our paper, previous work studied the automatic detection of users' attitudes towards vaccination and vaccines (Skeppstedt, Kerren, and Stede 2017; Bello-Orgaz, Hernandez-Castro, and Camacho 2017; Müller and Salathé 2019). Additionally, social science research has employed survey-based methods (e.g., self-reporting questionnaires) to analyze people's stance towards vaccination on a limited scale (Funk and Tyson 2020; Bonnevie et al. 2021). In this section, however, we focus our review on recent papers and benchmarks related to the **COVID-19 vaccination**.

### Existing COVID-19 Vaccine Stance Dataset

To analyze Twitter users' stance towards COVID-19 vaccination, Cotfas et al. (2021) introduce the first open-source<sup>6</sup> dataset containing English tweets published after the first month of the announcement of the Pfizer & BioNTech COVID-19 vaccine (i.e., from Nov 2020 to Dec 2020). Based on the data collection pipeline provided by Cotfas et al. (2021), Poddar et al. (2022a) and Delcea et al. (2022) develop two extended datasets which cover a limited time span (i.e., 3-month and 1-month after Dec 2020, respectively). In addition to monolingual datasets in English, Chen, Chen, and Pang (2022) release the first collection of multilingual tweets published in Luxembourg related to COVID-19 vaccines. However, we observe that the majority of tweets in Chen, Chen, and Pang (2022) are in French, which is the official language of Luxembourg. In general, all datasets mentioned above are developed for a standard three-way classification experimental setup, i.e., applying supervised methods to map a tweet into one of the stance categories including *anti-vaxx*, *pro-vaxx*, and *irrelevant*. We also observed a number of open-source unlabelled datasets (DeVerna et al. 2021) related to COVID-19 vaccination that is used for statistical analysis in the field of social sciences. We further display more specifications of these datasets in Table 5.

### COVID-19 Vaccine Stance Classification

Given the user-generated content, previous studies have used standard supervised classifiers (such as SVM and BERT) to identify users' stance towards COVID-19 vaccination (Müller and Salathé 2019; Cotfas et al. 2021; Chen, Chen, and Pang 2022). Similar to (Gururangan et al. 2020), we also observe that domain adaptive pre-trained language models (PLMs) can significantly perform better than vanilla PLMs. For example, COVID-BERT (Müller, Salathé, and Kummervold 2020) shows an improvement compared to the original bert-large checkpoint (Devlin et al. 2019) in various COVID-19 downstream tasks especially automatic detection of users' attitudes towards COVID-19 events (Müller,

<sup>6</sup>Here, we only discuss publicly available datasets as we need access to the dataset specifications.

Salathé, and Kummervold 2020; Cotfas et al. 2021; Poddar et al. 2022a). This suggests that the second phase of pre-training (i.e. the domain & task adaptive strategy) makes the COVID-BERT to capture more temporal & topical information for COVID-specific tasks.

As for the user-level analysis, Poddar et al. (2022a) shed light on the temporal concept drift of the percentage of users who were against the vaccine before and after COVID-19 pandemic by using tweets stance classifier to map the original publisher (Twitter users) into *pro-vaxxer* or *anti-vaxxer*. Besides, other studies have focused on the difference in user demographic features between people who support and those who are against COVID-19 vaccination (Wang and Liu 2021; Almadan et al. 2022; Aw et al. 2021). There are several reasons for the COVID-19 vaccine hesitancy, including potential side effects, distrust of the government (health system), and the vaccine developers (Johnson et al. 2020; Praveen, Ittamalla, and Deepak 2021; Germani and Biller-Andorno 2021; Poddar et al. 2022b). First Draft has also reported on widespread, high engagement vaccine disinformation narratives, as well as on the presence of "data deficits" where media and government sources are failing to provide the right information (Dodson, Mason, and Smith 2021).<sup>7</sup> In particular, disinformation about vaccines is known to impact negatively on citizens trust in COVID-19 vaccination, being a direct cause of vaccine hesitancy (Jennings et al. 2021; Loomba et al. 2021; Gisondi et al. 2022).

## Our Work

Previous studies have focused on analyzing user stance towards COVID-19 vaccination using standard stance categories (i.e., three-way setup). However, there is a need to distinguish between tweets that express personal concerns and hesitation about vaccination from the ones that share disinformation. For this purpose, we develop a four-category dataset enabling the analysis of vaccine hesitation on Twitter. We show details about previously published datasets and ours in Table 2. In comparison, our dataset differs from existing ones in a number of points: (i) we further divide user stance towards COVID-19 vaccination into fine-grained categories including vaccine hesitancy which is overlooked in previous works; (ii) our dataset covers a longer temporal span (i.e., 18-month) encompassing the period from the announcement of the first COVID-19 vaccine officially administered to the second or even booster doses being made publicly available; and (iii) we propose efficient data annotation strategies that use less time and human resources.

## Data

In general, we frame our dataset development pipeline into three steps:

- (i) **COVID-19 Vaccine-Related Tweets Collection.** We first collect a set of COVID-19 vaccine-related tweets ( $D$ ) through keywords searching;
- (ii) **Data Sampling.** Given  $D$ , we then filter out a representative subset of tweets  $T$  to annotate;

<sup>7</sup><https://firstdraftnews.org/long-form-article/covid-19-vaccine-research-reports/>



**Annotation based on the Tweet**

Please select the label based on the tweet text above.

Pro-vaccine  
 Anti-vaccine  
 vaccine Hesitancy  
 Irrelevant

**Confidence**

Please select the confident level of your annotation.

1  
 2  
 3  
 4  
 5

**Comment: (Required if your confidence score is below or equal 3)**

Please fill this section for the annotation comments. For example: Sarcasm tweet.

Figure 3: User Interface in GATE Teamware

The final subset  $T$  contains 3,101 tweets in English and is of similar size to other vaccine-related datasets (see Table 2). The only exception is Chen, Chen, and Pang (2022), which however contains only a tiny fraction of English tweets, as it is predominantly in French (over 60%) and German (over 30%).

## Data Annotation

**Data Annotation Workflow** The manual data annotation workflow consisted of three separate steps, as follows: (i) annotator training; (ii) a quality test session, and a final (iii) independent data annotation. All these were carried out using a collaborative web-based corpus annotation tool (GATE Teamware)<sup>13</sup> (Karmakharm et al. 2023):

- (i) **Annotator Training.** First annotators were trained during in-person tutorials which introduced the GATE Teamware platform and worked through the definitions of the four vaccine stance categories, including a detailed set of real-world tweet examples (a selection of these is included in Table 1);
- (ii) **Annotator Test Sessions.** To ensure that all volunteer annotators correctly applied the annotation guidelines and produced work of good quality, a test session was organised for each annotator. This consisted of 10 tweets covering all four stance categories. A subsequent question & answer session was also provided to explain mistakes and answer questions. In this way we ensure that annotators understand well the label definitions and can distinguish between them.
- (iii) **Dataset Annotation.** Once the preparatory stages are completed, the annotators are ready to start annotating the data independently. Annotators are shown one tweet at a time and are asked to select the tweet stance, a confidence level in their annotation and an optional comment. Table 4 shows the definitions of each confidence

<sup>13</sup><https://github.com/GateNLP/gate-teamware>

Groups	Annotators	All	Confi. $\geq 3$ <sup>†</sup>	Confi. = 5 <sup>†</sup>
Group 1	U1 & U2	0.40	0.54	0.56
	U1 & U3	0.61	0.69	0.85
	U2 & U3	0.44	0.50	0.54
Group 2	U4 & U5	0.47	0.54	0.77
	U4 & U6	0.54	0.88	0.90
	U5 & U6	0.59	0.81	0.87
Group 3	U7 & U8	0.33	0.49	0.83
	U7 & U9	0.31	0.46	0.53
	U8 & U9	0.62	0.67	0.83
Group 4	U10 & U11	0.52	0.51	0.66
	U10 & U12	0.51	0.60	0.71
	U11 & U12	0.53	0.53	0.86
Group 5	U13 & U14	0.49	0.69	0.81
	U13 & U15	0.22	0.34	0.41
	U14 & U15	0.31	0.41	0.44
Group 6	U16 & U17	0.58	0.62	0.74
	U16 & U18	0.54	0.61	0.72
	U17 & U18	0.57	0.64	0.86

Table 3: Cohen’s kappa coefficient ( $K$ ) between every two annotators in each group. <sup>†</sup> denotes that the Cohen’s Kappa coefficient ( $K$ ) in column ‘Confi.=5’ is significantly higher than the values in columns ‘All’ and ‘Confi. $\geq 3$ ’ ( $t$ -test,  $p < 0.001$ ).

level. Annotators can also leave an optional comment for each tweet. However, a comment is compulsory if the confidence score is three or lower. Figure 3 shows the corresponding GATE Teamware user interface.

**Annotation Methodology and Quality Assurance** Prior work has used triple annotation, where each tweet is labelled by three annotators and the final label is determined through a simple majority (i.e., a minimum of two annotators need to have selected the same label).

Since manual annotation is time consuming and expensive, we opted for a different, more efficient strategy.

A total of 18 volunteers were recruited to manually annotate the tweets in the subset  $T$ . These 18 participants were divided into six separate groups (i.e., three annotators per group). In each group 220 tweets were assigned to each annotator (660 tweets per group), which included 60 tweets also assigned to other annotators in the group. In this way we obtained 180 double annotated and 300 single-annotated tweets from each group.

This methodology maximizes our capacity to annotate more tweets with fewer annotators. Impact on annotation quality was measured within each group by calculating the Cohen’s kappa coefficient ( $K$ ) which measures inter-rater reliability for annotator pairs (see Table 3).

Data annotation quality is improved further based on the confidence level of the tweet annotations. In particular, tweet annotations with a low confidence score were discarded. As can be seen in Table 3, this improves Cohen’s kappa  $K$  significantly ( $t$ -test,  $p < 0.05$ ).

## The VaxxHesitancy Dataset

The complete VaxxHesitancy dataset has 3,101 tweets in total, including a training set with 2,670 tweets and a golden

Confidence	Definitions
5	Extremely confident about the annotation (I'm certain about the annotation without a doubt.)
4	Fairly confident about the annotation (I'm confident about the annotation, but might be in small chance other annotators may label it in a different category.)
3	Pretty confident about the annotation (I'm pretty sure about the annotation, but might be in high chance other annotators may label it in a different category.)
2	Not confident about the annotation (I'm not sure about the annotation, it seems it also belongs to other categories, but you can still include this instance as a "silver standard instance" in training.)
1	Extremely unconfident about the annotation (I'm really unsure about the annotation. It may belong to another category as well, you may wish to discard this instance from the training.)

Table 4: Confidence Scores and Definitions

#	Train 2+	Train	Test	Sum
Pro	784	791	176	967
Anti	562	571	76	647
Hesitant	341	344	28	372
Irrelevant	916	964	151	1,115
Sum	2,603	2,670	431	3,101

Table 5: Dataset statistics. Train set 2+ indicates that we filtered out tweets with a confidence score of 1.

test set (double agreed) with 431 tweets (See Table 5).

In more detail, the VaxxHesitancy dataset is split into two parts:

- **Test set:** obtained from tweets with where both annotators agree with confidence score higher than three.
- **Training set:** consists of all single annotated tweets (2,318), as they can be more noisy, and the remaining 352 double annotated tweets. For double-annotated tweets, we retain the labels from each of the annotators, their respective confidence scores, and any comments. This gives flexibility in the way this information is used during training, as discussed in our experimental Section ).

## Data Characterization

### Linguistic Analysis

In order to investigate the differences between tweets in the *anti-vaxx* and *vaxx-hesitant* categories, we conduct a comparative linguistic analysis. We opted to only investigate the differences between these two categories since they were conflated in previous datasets under a common anti-vaccination category.

We perform an univariate Pearson's correlation test to characterize which linguistic patterns (i.e., BOW and LIWC Dictionary<sup>14</sup>) are highly correlated with each of the two categories (i.e., *anti-vaxx* and *vaxx-hesitant*) following the approach from Schwartz et al. (2013).

<sup>14</sup>We use Linguistic Inquiry and Word Count: LIWC2015 <https://www.liwc.app/>

### BOW

We first use the bag-of-words model (BOW) to represent each post as a TF-IDF weighted distribution over a 3,000-sized vocabulary consisting of the most frequent uni-grams (i.e., word-level tokens). We only extract tokens appearing in more than 10 and no more than 30% of the total number of tweets. To better display the differences in BOW features associated with each category, we created a word cloud (see Figure 4) that shows the top 100 BOW features for each of the two categories which are indicated by the different font types (i.e., *anti-vaxx* (*thin*) in blue and *vaxx-hesitant* (*bold*) in red) and font size (the larger the font, the higher the Pearson correlation  $r$ ) with the respective category.

In Figure 4 we observe that *anti-vaxx* tweets contain more external links (denoted by a common HTTPURL token) and fear-inducing words such as 'death', 'killing', 'illegal', etc. Furthermore, some of these tokens are linked to some widely spread false narratives related to the COVID-19 vaccine, e.g., '5G Micro Chip in Vaccine'<sup>15</sup> and 'Bill Gates talked about using vaccines to control population growth'.<sup>16</sup> Below we show some examples from our dataset:

Tweet 1: *'Mind Control and 5G Bill gates will insert micro chips with vaccine HTTPURL'*

Tweet 2: *'The vaccine was made by the government ... and you don't want the government to control you. Other medicines are not from the government ... and doctors are not from the government so use that'*

On the other hand, tweets belonging to the *vaccine-hesitant* category tend to have a more prevalent use of words related to first-person pronouns (e.g., 'I'm', 'me', 'my', etc) and self-disclosure (e.g., 'feel', 'suspicious', 'scared', etc.). This suggests that the *vaxx-hesitant* tweets in the dataset are well aligned with our definition (see Table 1) and can be used to shed light on the posters' personal intentions for delaying or refusing the vaccine.

### LIWC

We also represent each tweet using the 93 psycho-linguistic categories from the LIWC 2015 dictionary (Pennebaker,

<sup>15</sup><https://www.politifact.com/factchecks/2021/aug/10/viral-image/link-between-covid-19-variant-and-firm-works-5g-ju/>

<sup>16</sup><https://www.politifact.com/factchecks/2021/oct/11/blog-posting/bill-gates-didnt-say-he-wanted-use-vaccines-reduce/>




Figure 4: Top 100 BOW features associated with **anti-vaxx** (thin) and **vaccine-hesitant** (bold) categories. The larger the font, the higher the Pearson correlation value, and vice versa. HTTPURL indicates external links in anti-vaxx tweets.

LIWC			
Anti-vaxx	<i>r</i>	Vaccine-hesitant	<i>r</i>
Clout	0.133	1st pers singular	0.294
Analytical thinking	0.105	Pronoun	0.188
Other punctuation	0.103	Authentic	0.187
All Punctuation	0.096	Question marks	0.185
Quotation marks	0.092	Function	0.175
Death	0.092	Personal pronouns	0.175
Word Count	0.089	Dictionary words	0.159
Negations	0.083	Interrogatives	0.132
Religion	0.077	Prepositions	0.124
Social	0.074	Conjunctions	0.124

Table 6: LIWC categories associated with anti-vaxx and vaccine hesitant tweets sorted by Pearson’s correlation (*r*) between the normalized frequency and the labels ( $p < .001$ ).

Francis, and Booth 2001). Table 6 shows the top 10 most correlated LIWC categories with each of the two stance categories (anti-vaxx and hesitant).

The results are similar to those obtained using BOW. Namely, LIWC categories related to first-person pronouns (e.g., ‘1st pers singular’, ‘Pronoun’, ‘personal pronouns’, etc) are more prevalent in tweets belonging to the *vaccine-hesitant* category. Furthermore, some LIWC features such as ‘Interrogatives’ and ‘Question marks’ are also highly correlated with that category. This indicates that vaccine-hesitant users usually raise questions about the COVID-19 vaccine, e.g. concerning vaccine safety.

Here are two examples from the dataset:

Tweet 3: ‘Would it be possible to know the ingredients of the vaccine? Just like food packaging - it is wise to know the ingredients to prevent allergies reaction.’

Tweet 4: ‘#AstraZeneca should my mom get it or not I really worry she’s 60 with asthma I really worry...’

In comparison, we notice that frequent LIWC features in tweets belonging to the *anti-vaxx* category are ‘Quote’, ‘Death’, ‘Negations’, and ‘Religion’. This indicates that anti-vaxx tweets are more likely to use external links that refer to misleading articles or websites and/or to refer to religious reasons (e.g. ‘Vatican permits use of COVID-19 vaccines made using aborted fetal tissue’<sup>17</sup>), aiming to raise fear and influence citizens against vaccination.

## COVID-19 Vaccine Stance Prediction

### Baseline Models

Following previous work (Cofas et al. 2021; Poddar et al. 2022a; Chen et al. 2022), we train two strong baselines to classify posts into our four categories:

- **BERT**: Following Devlin et al. (2019), we directly fine-tune BERT by adding a fully connected layer on the top of the bert-large<sup>18</sup> model. We consider the special token (i.e., [CLS]) as the tweet-level representation.

<sup>17</sup><https://www.reuters.com/business/healthcare-pharmaceuticals/vatican-permits-use-covid-19-vaccines-made-using-aborted-foetal-tissue-2020-12-21/>

<sup>18</sup><https://huggingface.co/bert-large-uncased>

Models	Accuracy	Precision	Recall	F1-score
<b>All Confidence (Set 1)</b>				
BERT	57.73 ± 0.77	52.08 ± 0.45	54.99 ± 0.94	52.82 ± 0.47
COVID BERT	69.65 ± 1.43	62.70 ± 1.23	62.97 ± 1.61	62.65 ± 1.40
VaxxBERT	72.5 ± 1.20	68.02 ± 1.24	69.35 ± 1.30	68.55 ± 1.44
<b>Confidence &gt;1 (Set 2)</b>				
BERT	58.52 ± 0.58	52.77 ± 0.63	55.53 ± 1.15	53.10 ± 0.84
COVID BERT	69.74 ± 1.43	63.25 ± 1.09	63.65 ± 1.10	63.29 ± 1.09
VaxxBERT	72.91 ± 1.23	68.53 ± 1.11	69.45 ± 1.71	68.92 ± 1.34
<b>Higher confidence score (Set 3)</b>				
BERT	59.07 ± 0.74	55.04 ± 1.04	56.51 ± 1.64	54.95 ± 1.32
COVID BERT	71.14 ± 0.94	64.54 ± 1.29	66.61 ± 1.03	65.05 ± 1.17
VaxxBERT	<b>73.04 ± 0.94</b>	<b>68.71 ± 1.59</b>	<b>70.22 ± 1.36</b>	<b>69.29 ± 1.31</b>

Table 7: Model predictive performance.

- **COVID-BERT** (Müller, Salathé, and Kummervold 2020) is a domain-adapted uncased BERT model which is pretrained on COVID-19-related tweets (i.e., ‘covid-twitter-bert-v2’).<sup>19</sup> We fine-tune the COVID-BERT model using the same strategy as in Devlin et al. (2019).

### Domain Specific PLM (Our VaxxBERT Model)

Following Gururangan et al. (2020), we develop the VaxxBERT domain-specific language model, which is based on the uncased COVID-BERT (Müller, Salathé, and Kummervold 2020), which has been trained on 160 million unannotated (raw) tweets related to the COVID-19 virus.

In VaxxBERT, we continue pretraining the COVID-BERT checkpoint on over 175 million **unlabelled domain-specific tweets** (i.e., we use all remaining tweets in  $D$ ). Following Müller, Salathé, and Kummervold (2020) and Gururangan et al. (2020), we set the sequence length to 256, batch size to 96, validation ratio to 10% of the training set and learning rate to  $2e-5$ . We randomly mask tokens across epochs (2 epochs in total) with the default token masking rate (15%). Model checkpoints are saved every 50,000 training steps and the best checkpoint is selected by the lowest validation loss. For domain-specific pretraining, we use open-source PyTorch scripts from the transformers library<sup>20</sup> (Wolf et al. 2020). VaxxBERT was trained for 240 hours on two NVIDIA GeForce RTX Graphic cards with 24GB memory.

### Experimental Setup

All tweets are pre-processed by replacing URLs and user @mentions with special tokens (i.e., HTTPURL and @USER, respectively). The maximum sequence length is set to 256 tokens. We use CrossEntropyLoss as the loss function and the Adam (Kingma and Ba 2015) optimizer. All models are trained using an early-stopping strategy with a learning rate  $l = 3e - 6$ , and batch size of 32. We train each

<sup>19</sup><https://huggingface.co/digitalepidemiologylab/covid-twitter-bert-v2>

<sup>20</sup><https://github.com/huggingface/transformers/tree/main/examples/pytorch/language-modeling>

model five times with different random seeds, and report the mean accuracy, precision, recall and F1-score in Table 7.

### Results

We design a battery of controlled experiments to test the quality of our dataset using the following subsets:

- **Set 1** There are 352 double-annotated tweets and 2318 single-annotated tweets in the training set. We keep all single-annotated ones and randomly choose a label from two annotations for these double-annotated ones.
- **Set 2** We retain only tweets with annotations with confidence score of 2 or higher (285 tweets) from **Set 1**. This is motivated by the definition of confidence score of 1, which indicates that the annotator is extremely unsure about the stance category and wishes to drop the tweet from training. In addition, all remaining double-annotated tweets have one of their two labels chosen at random as the category used for training.
- **Set 3** is similar to Set 2, except that the labels of all double-annotated tweets are selected based on which one of the two has a higher confidence score. In cases where both annotations have the same confidence score but different categories, then one of them is selected at random.

**Set 1 v.s. Set 2** We first compare the performance of models trained on all tweets in the training set (Set 1) vs those trained on Set 2, i.e. annotations with confidence scores higher than 1.

We notice that most of the models trained on set 2 perform slightly better than those trained on set 1 according to all evaluation metrics. This indicates that dropping tweets with low confidence annotations helps improve the overall quality of both the dataset and the models.

**Set 2 v.s. Set 3** We also compare the performance of models trained on tweets where labels of double annotated tweets are either chosen at random (Set 2) or on the basis of a higher confidence score (Set 3).

In this case, models trained on Set 3 perform better than those trained on Set 2. Combined with the results above,



Column	Field	Type	Description
0	id	str	Unique Tweet ID
1	user1_stance1	str	Stance by Annotator 1
2	user1_comment	str	Comment by Annotator 1
3	user1_confidence	int	Confidence Score by Annotator 1
4	user1_time	int	Time (i.e., number of seconds) used by Annotator 1
5	user1_annotator_id	str	Annotator 1 ID
6	user2_stance1	str	Stance by Annotator 2
7	user2_comment	str	Comment by Annotator 2
8	user2_confidence	int	Confidence Score by Annotator 2
9	user2_time	int	Time (i.e., number of seconds) used by Annotator 2
10	user2_annotator_id	str	Annotator 2 ID

Table 8: Dataset Columns. Note that single annotated tweets are only linked to columns from 0 to 5, the remaining columns are filled with ‘N/A’.

we argue that considering annotator confidence helps to discard low-quality annotations and helps select better labels in cases where annotators disagree.

**VaxxBERT v.s. Baselines** To evaluate the performance of the domain adaptive pretrained model, we also compare VaxxBERT against the two strong baselines introduced above (i.e., bert-large and COVID-BERT). In general, our VaxxBERT model significantly ( $t$ -test,  $p < 0.001$ ) outperforms the two baselines on all three sets (i.e., Set 1, Set 2, Set 3). Among them, VaxxBERT trained on Set 3 achieves the best performance (i.e., 73.0% accuracy and 69.3% F1-measure). This demonstrates that the extra domain-specific information helps VaxxBERT to improve its predictive performance, which aligns with the findings of Gururangan et al. (2020) and Müller, Salathé, and Kummervold (2020).

## Conclusion

This research was motivated by the need of media and governments to better understand the reasons behind COVID-19 vaccine hesitancy, since automatic tools for monitoring and analysis of vaccine sentiment can help policymakers design better-informed and targeted information aimed at addressing key concerns of vaccine hesitant citizens.

To this end, we introduced a new open-source dataset that tracks users’ attitudes toward COVID-19 vaccines on Twitter. It is focused in particular on capturing vaccine hesitant tweets. Our linguistic analysis revealed significant differences between tweets belonging to the anti-vaxx and vaccine hesitant categories. Our second contribution is in using our un-annotated collection of 175 million tweets related to COVID-19 vaccination, in order to train a domain-specific PLM (VaxxBERT) that outperforms other competitive baselines on our finer-grained vaccine stance categorisation task.

In the future, we plan to enrich our dataset with multi-modal information such as images and videos. Additionally, we intend to increase the number of annotated samples in our dataset by implementing advanced NLP techniques such as active learning.

## Applications

Our work has several practical implications:

- First, our dataset and the new VaxxBERT domain-specific language model can be easily reused by researchers as they are released via widely used open science platforms (i.e., Zenodo and HuggingFace (Wolf et al. 2020)). Furthermore, our VaxxBERT model achieves the best predictive performance, which can be used as a strong baseline in future research.
- We re-frame the previous task of predicting user stance towards COVID-19 vaccination by considering a separate vaccine hesitancy category, which we demonstrate is different from anti-vaxx tweets. This can help shed light on understanding vaccine hesitancy in other cases (e.g. MMR vaccines). Also, we propose a data collection methodology that can be adopted and applied to multilingual and multi-platform datasets.
- Our dataset captures the various reasons behind vaccine hesitancy (see Figure 2) and it can be used to develop interpretable classification models that generate faithful rationales.
- Our dataset can be repurposed for a standard three-way setting by integrating the categories of *anti-vaxx* and *vaxx-hesitant*. This integration will allow for cross-dataset evaluation with other benchmarks.
- Finally, our dataset can be used for qualitative research by social scientists and psychologists in order to understand better the demographic features and personality traits of vaccine hesitant and anti-vaxx users of Twitter.

## Dataset Availability and Ethics Statement

Our dataset is publicly available in compliance with the FAIR principles (Wilkinson et al. 2016):

- **Findable:** Our dataset has been published in the Zenodo dataset sharing service with a unique digital object identifier (DOI: 10.5281/zenodo.7601328). We also share the VaxxBERT model via the Huggingface platform<sup>21</sup>.
- **Accessible:** Original tweets are retrievable based on their tweet IDs using the standard Twitter API<sup>22</sup>.

<sup>21</sup><https://huggingface.co/GateNLP/covid-vaccine-twitter-bert>

<sup>22</sup><https://developer.twitter.com/en/docs/twitter-api/tweets/lookup/api-reference/get-tweets-id>

- **Interoperable:** Table 8 summarises the dataset structure in CSV format and the description of each column (11 columns in total). CSV datasets are easily imported and processed by most widely used data processing tools.
- **Re-usable:** Anyone with a Twitter developer account can re-use our dataset. Using the transformer library (Wolf et al. 2020) researchers can also fine-tune VaxxBERT for other NLP tasks.

This work has ethical approval #037567 from our University Research Ethics Committee. Our data collection protocol complies with the Twitter data policies for research.<sup>23</sup> We only share tweets IDs following the Twitter API policy and replace annotator names with identifiers.

### Acknowledgements

This research is supported by a UKRI grant EP/W011212/1 and an EU Horizon 2020 grant (agreement no.871042) (“So-BigData++: European Integrated Infrastructure for Social Mining and BigData Analytics” (<http://www.sobigdata.eu>)). We would like to thank George Chrysostomou and all the anonymous reviewers for their valuable feedback.

### References

Almadan, A.; Maher, M. L.; Pereira, F. B.; and Guo, Y. 2022. Will You Be Vaccinated? A Methodology for Annotating and Analyzing Twitter Data to Measure the Stance Towards COVID-19 Vaccination. In *Future of Information and Communication Conference*, 311–329. Springer.

Augenstein, I.; Rocktäschel, T.; Vlachos, A.; and Bontcheva, K. 2016. Stance Detection with Bidirectional Conditional Encoding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 876–885.

Aw, J.; Seng, J. J. B.; Seah, S. S. Y.; and Low, L. L. 2021. COVID-19 vaccine hesitancy—A scoping review of literature in high-income countries. *Vaccines*, 9(8): 900.

Bello-Orgaz, G.; Hernandez-Castro, J.; and Camacho, D. 2017. Detecting discussion communities on vaccination in twitter. *Future Generation Computer Systems*, 66: 125–136.

Bonnevie, E.; Gallegos-Jeffrey, A.; Goldbarg, J.; Byrd, B.; and Smyser, J. 2021. Quantifying the rise of vaccine opposition on Twitter during the COVID-19 pandemic. *Journal of communication in healthcare*, 14(1): 12–19.

Chen, N.; Chen, X.; and Pang, J. 2022. A multilingual dataset of COVID-19 vaccination attitudes on Twitter. *Data in Brief*, 44: 108503.

Chen, N.; Chen, X.; Pang, J.; Borga, L. G.; D’Ambrosio, C.; and Vögele, C. 2022. Measuring COVID-19 Vaccine Hesitancy: Consistency of Social Media with Surveys. In *International Conference on Social Informatics*, 196–210. Springer.

Cotfas, L.-A.; Delcea, C.; Roxin, I.; Ioanăș, C.; Gherai, D. S.; and Tajariol, F. 2021. The longest month: analyzing COVID-19 vaccination opinions dynamics from tweets in the month following the first vaccine announcement. *Ieee Access*, 9: 33203–33223.

Delcea, C.; Cotfas, L.-A.; Crăciun, L.; and Molănescu, A. G. 2022. New Wave of COVID-19 Vaccine Opinions in the Month the 3rd Booster Dose Arrived. *Vaccines*, 10(6): 881.

Derczynski, L.; Bontcheva, K.; Liakata, M.; Procter, R.; Hoi, G. W. S.; and Zubiaga, A. 2017. SemEval-2017 Task 8: RumourEval: Determining rumour veracity and support for rumours. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, 69–76.

DeVerna, M. R.; Pierri, F.; Truong, B. T.; Bollenbacher, J.; Axelrod, D.; Loynes, N.; Torres-Lugo, C.; Yang, K.-C.; Menczer, F.; and Bryden, J. 2021. CoVaxxy: A collection of English-language Twitter posts about COVID-19 vaccines. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 992–999.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186.

Di Giovanni, M.; Pierri, F.; Torres-Lugo, C.; and Brambilla, M. 2022. VaccinEU: COVID-19 vaccine conversations on Twitter in French, German and Italian. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 1236–1244.

Dodson, K.; Mason, J.; and Smith, R. 2021. Covid-19 vaccine misinformation and narratives surrounding Black communities on social media. *First Draft*.

Funk, C.; and Tyson, A. 2020. Intent to get a COVID-19 vaccine rises to 60% as confidence in research and development process increases. *Pew Research Center*, 3.

Germani, F.; and Biller-Andorno, N. 2021. The anti-vaccination infodemic on social media: A behavioral analysis. *PloS one*, 16(3): e0247642.

Gisondi, M. A.; Barber, R.; Faust, J. S.; Raja, A.; Strehlow, M. C.; Westafer, L. M.; and Gottlieb, M. 2022. A deadly infodemic: social media and the power of COVID-19 misinformation. *Journal of Medical Internet Research*, 24(2): e35552.

Gorrell, G.; Kochkina, E.; Liakata, M.; Aker, A.; Zubiaga, A.; Bontcheva, K.; and Derczynski, L. 2019. SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, 845–854.

Gururangan, S.; Marasović, A.; Swayamdipta, S.; Lo, K.; Beltagy, I.; Downey, D.; and Smith, N. A. 2020. Don’t Stop Pretraining: Adapt Language Models to Domains and Tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8342–8360.

Jennings, W.; Stoker, G.; Bunting, H.; Valgarðsson, V. O.; Gaskell, J.; Devine, D.; McKay, L.; and Mills, M. C. 2021. Lack of trust, conspiracy beliefs, and social media use predict COVID-19 vaccine hesitancy. *Vaccines*, 9(6): 593.

Johnson, N. F.; Velásquez, N.; Restrepo, N. J.; Leahy, R.; Gabriel, N.; El Oud, S.; Zheng, M.; Manrique, P.; Wuchty, S.; and Lupu, Y. 2020. The online competition between pro- and anti-vaccination views. *Nature*, 582(7811): 230–233.

<sup>23</sup><https://developer.twitter.com/en/products/twitter-api>

- Karmakharm, T.; Wilby, D.; Roberts, I.; and Bontcheva, K. 2023. GATE Teamware (Version 0.1.4) [Computer software]. <https://github.com/GateNLP/gate-teamware>.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Levenshtein, V. I.; et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, 707–710. Soviet Union.
- Loomba, S.; de Figueiredo, A.; Piatek, S. J.; de Graaf, K.; and Larson, H. J. 2021. Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature human behaviour*, 5(3): 337–348.
- Müller, M.; Salathé, M.; and Kummervold, P. E. 2020. COVID-Twitter-BERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter. *arXiv preprint arXiv:2005.07503*.
- Müller, M. M.; and Salathé, M. 2019. Crowdbreaks: tracking health trends using public social media data and crowdsourcing. *Frontiers in public health*, 7: 81.
- Muric, G.; Wu, Y.; Ferrara, E.; et al. 2021. COVID-19 vaccine hesitancy on social media: building a public twitter data set of antivaccine content, vaccine misinformation, and conspiracies. *JMIR public health and surveillance*, 7(11): e30642.
- Pennebaker, J. W.; Francis, M. E.; and Booth, R. J. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates*, 71.
- Poddar, S.; Mondal, M.; Misra, J.; Ganguly, N.; and Ghosh, S. 2022a. Winds of Change: Impact of COVID-19 on Vaccine-related Opinions of Twitter users. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 782–793.
- Poddar, S.; Samad, A. M.; Mukherjee, R.; Ganguly, N.; and Ghosh, S. 2022b. CAVES: A Dataset to Facilitate Explainable Classification and Summarization of Concerns towards COVID Vaccines. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 3154–3164.
- Praveen, S.; Ittamalla, R.; and Deepak, G. 2021. Analyzing the attitude of Indian citizens towards COVID-19 vaccine—A text analytics study. *Diabetes & Metabolic Syndrome: Clinical Research & Reviews*, 15(2): 595–599.
- Rzymiski, P.; Borkowski, L.; Drąg, M.; Flisiak, R.; Jemielity, J.; Krajewski, J.; Mastalerz-Migas, A.; Matyja, A.; Pyrc, K.; Simon, K.; et al. 2021. The strategies to support the COVID-19 vaccination with evidence-based communication and tackling misinformation. *Vaccines*, 9(2): 109.
- Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; and Seligman, M. E. 2013. Personality, Gender, and Age in the Language of Social Media: The Open-vocabulary Approach. *PLoS ONE*, 8(9).
- Skeppstedt, M.; Kerren, A.; and Stede, M. 2017. Automatic detection of stance towards vaccination in online discussion forums. In *Proceedings of the International Workshop on Digital Disease Detection using Social Media 2017 (DDDSM-2017)*, 1–8. Taipei, Taiwan: Association for Computational Linguistics.
- van der Linden, S.; Dixon, G.; Clarke, C.; and Cook, J. 2021. Inoculating against COVID-19 vaccine misinformation. *EClinicalMedicine*, 33.
- Wang, Y.; and Liu, Y. 2021. Multilevel determinants of COVID-19 vaccination hesitancy in the United States: A rapid systematic review. *Preventive medicine reports*, 101673.
- Wilkinson, M. D.; Dumontier, M.; Aalbersberg, I. J.; Appleton, G.; Axton, M.; Baak, A.; Blomberg, N.; Boiten, J.-W.; da Silva Santos, L. B.; Bourne, P. E.; et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3(1): 1–9.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, 38–45.