"Double vaccinated, 5G boosted!": Learning Attitudes towards COVID-19 Vaccination from Social Media

NINGHAN CHEN, Department of Computer Science, University of Luxembourg, Esch-sur-Alzette, Luxembourg
XIHUI CHEN, Department of Computer Sciencesadf, St. Pölten University of Applied Sciences, St. Pölten, Austria
ZHIQIANG ZHONG, Faculty of Natural Sciences, Aarhus Universitet, Aarhus, Denmark
JUN PANG, Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Esch-sur-Alzette, Luxembourg

The sudden onset of the recently concluded COVID-19 pandemic has driven substantial progress in various scientific fields. One notable example is the comprehension of public vaccination attitudes and the timely monitoring of their fluctuations through social media platforms. This approach can serve as a cost-effective means to supplement surveys in gathering public vaccine hesitancy levels. In this article, we propose a deep learning framework leveraging textual posts on social media to extract and track users' vaccination stances in near real time. Compared to previous works, we integrate into the framework the recent posts of a user's social network friends to collaboratively detect the user's genuine attitude towards vaccination. Based on our annotated dataset from X (formerly known as Twitter), the models instantiated from our framework can increase the performance of attitude extraction by up to 23% compared to the state-of-the-art text-only models. Using this framework, we successfully confirm the feasibility of using social media to track the evolution of vaccination attitudes in real life. In addition, we illustrate the generality of our framework in extracting other public opinions such as political ideology. We further show one practical use of our framework by validating the possibility of forecasting a user's vaccine hesitancy changes with information perceived from social media.

CCS Concepts: • **Information systems** → **Collaborative and social computing systems and tools**

Additional Key Words and Phrases: Social media, vaccine hesitancy, text mining, graph neural networks, COVID-19, dataset

ACM Reference Format:
Ninghan Chen, Xihui Chen, Zhiqiang Zhong, and Jun Pang. 2025. "Double vaccinated, 5G boosted!": Learning Attitudes towards COVID-19 Vaccination from Social Media. *ACM Trans. Web* 19, 1, Article 8 (February 2025), 24 pages. https://doi.org/10.1145/3702654

This research was funded in whole, or in part, by the Luxembourg National Research Fund (FNR), grant reference C21/IS/16281848 (HETERS) and PRIDE17/12252781 (DRIVEN). For the purpose of open access, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.
Authors' Contact Information: Ninghan Chen, Department of Computer Science, University of Luxembourg, Esch-sur-Alzette, Luxembourg, Luxembourg; e-mail: ninghan.chen@uni.lu; Xihui Chen, Department of Computer Sciencesadf, St. Pölten University of Applied Sciences, St. Pölten, Austria; e-mail: xihui.chen@fhstp.ac.at; Zhiqiang Zhong, Faculty of Natural Sciences, Aarhus Universitet, Aarhus, Midtjylland, Denmark; e-mail: zzhong@ca.au.dk; Jun Pang, Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Esch-sur-Alzette, Luxembourg, Luxembourg; e-mail: jun.pang@uni.lu.

ACM 1559-1131/2025/02-ART8
https://doi.org/10.1145/3702654

1 Introduction

The recently concluded COVID-19 pandemic has propelled significant advances, not only in biology and medicine-related domains but also in areas of public health surveillance [55, 57]. The progress will get us better prepared for future similar crises. As one notable example, social media has become an imperative source of data to understand the public's reactions to the unprecedented global health crisis, e.g., mental health and well-being [11, 69] and social problems such as ageism [32]. Due to social distancing and fear of the unknown, people spent more time than ever on social media. Even before the pandemic, social media has been exploited by data analysts as an auxiliary data source to complement public health surveillance [55, 57] and understand social events such as natural disasters [1, 44, 59], despite its inherent bias, such as socio-demographical bias [3, 12]. In this article, we aim to leverage the enormous daily posts during the COVID-19 pandemic to extract users' vaccination attitudes and track their temporal changes.

Vaccination has shown its critical role in combatting the global COVID-19 pandemic [43]. Despite the decreased efficacy against the infection of continuously evolving variants, a high-level uptake of the currently available vaccines is still a key to restrain the number of severe diseases, deaths and particularly hospitalisation, which is crucial for medical systems to remain operating as normal [6]. Regrettably, similar to the vaccines of other infectious diseases, not everyone is willing to be vaccinated [22]. A high level of *vaccine hesitancy* will lead to stagnant uptake ratios, and its impact varies across countries and between different subgroups of people such as healthcare workers [60] and immigrants [2]. It is thus imperative to grasp the public's level of vaccine hesitancy in time for healthcare departments to take effective and proactive measures. Compared to social surveys with pre-defined questions, examining vaccine hesitancy through social media is cost-effective and able to catch its evolution [12], especially with the rapid development of a crisis like COVID-19. However, we still face challenges. First, the globalism of the pandemic requires a method that can deal with multilingual posts. This fails most feature-based methods which classify posts with textual features, e.g., through keywords, and usually focus on one single language. Second, users' attitudes are usually ambiguously expressed in vaccination-related discourses. This significantly impairs the performance of existing models. Consider the following example: "I wouldn't do it for their vaccine, I'm waiting for the 6G". The user expresses his/her support for vaccination by making fun of the conspiracy theory that chips are implanted with vaccine injection. After experimenting with the state-of-the-art text feature-based classification methods which are adopted as benchmarks including the transformer-based ones such as CT-BERT [53, 54], we only get an accuracy of 0.65 on our collected data from X, which is formerly known as Twitter[1]. This is not reliable enough for trustworthy analyses. Last but not least, the aim of continuous tracking excludes the methods exploiting community-based methods [17, 33] due to the relative stability of community affiliations.

To address the preceding challenges, we take advantage of the recent advances of deep learning technologies in **natural language processing (NLP)** and social networks, and propose a framework that can accurately classify a textual post according to the vaccination stance expressed by its originator. To deal with the multilingual challenge, we make use of the transformer models instead of dictionary-based ones. To mitigate language ambiguity, we adopt two methods. The first method is to exploit the recent posts of a user because social media users will stick to the same stance in a relatively short time period. The second method is based on recent studies which reveal that a user's vaccination attitudes correlate with those of their neighbours in social networks, e.g., friends and friends of friends. For example, online social network users with

[1] As our data were collected when X was still called *Twitter*, we choose to use Twitter in this article.

negative attitudes often have social relations with users of positive attitudes [33, 45]. Inspired by *collaborative filtering* in recommender systems [21], we integrate the recent posts of a user's social network neighbours in our framework.

To train and test models instantiated from our framework, we collect 9,135,393 tweets from Twitter generated by 69,936 users, and create the first annotated dataset of 18,246 tweets manually labelled with affective vaccination stances, e.g., positive and negative. In addition to the experimental evaluation, we draw the temporal evolution of vaccination attitudes extracted from our collected tweets. We cross-validate with published social studies and manually analyse the popular social events occurring around significant changes in vaccine hesitancy levels. Our model has also been adopted and applied in measuring individual users' vaccine hesitancy in a recent study [14] which successfully cross-validates the consistency between Twitter and social surveys. Moreover, to test the generality of our framework in extracting public opinions, we apply our framework to obtain individuals' political ideology from a publicly available social network dataset. All these validation results successfully illustrate the effectiveness of our framework, as well as the power of social media as a data source to grasp public vaccine hesitancy in practice in near real time.

Newsagents, governments, healthcare professionals and even anti-vaccine activists use social media to spread news, knowledge and suggestions to persuade or dissuade people from getting vaccinated [62]. To showcase the practical use of our post-based attitude learning framework, to the best of our knowledge, we are the first to demonstrate that the information users perceive from social media can be used as predictors of vaccine hesitancy changes.

Contributions. Our contributions are as follows:

— We propose a framework to extract vaccination stances from textual social media posts. Our framework integrates recent posts of a user's social network neighbours to help reduce the interference of linguistic features, e.g., sarcasm and irony.
— We design models instantiating our framework. Based on our annotated dataset from Twitter, the best model can increase the performance of attitude extraction by up to 23% compared to state-of-the-art text-only models.
— We validate the generality of our framework in public opinion extraction by understanding individual users' political ideology.
— Using the model with the best performance, we track the evolution of vaccination attitudes. The utility of the extracted vaccination attitudes is further validated by the consistency with published statistics and explainable significant fluctuations of vaccine hesitancy in terms of social events such as wide propagation of misinformation and negative news.
— We show a practical use of our framework by validating the possibility of predicting a user's vaccination hesitancy changes with the information he/she perceives from social media.

Throughout this article, we (re-)establish the power of social media as a complementary data source in public health surveillance despite its inherent biases. Specifically, when exploited properly, it can provide healthcare bodies with useful information to guide or support their decision-making processes. Note that despite our target at COVID-19 vaccination attitudes, our framework is actually general to be used or extended to extract other pubic opinions from social media posts.

2 Related Work

Vaccine hesitancy is believed to be a major cause of stagnant vaccine coverage and contributor to vaccine program failure [22]. Despite the lack of a unified definition, one widely accepted representation of vaccine hesitancy is a continuum, ranging from complete rejection of vaccines to varying degrees of scepticism [67]. In this section, we concentrate on the vaccine hesitancy studies after the onset of the COVID-19 pandemic. A considerable amount of literature has been published investigating the state of vaccine hesitancy and the influential factors in different regions [34, 61] for specific groups of people such as healthcare employees [10, 29], immigrants [5] and college students [7]. Although surveys are still the most adopted method to collect sampled populations' attitudes or stances towards vaccination [4], some recent works leverage social media as a new dimension [30, 33]. Compared to self-reporting questionnaires, social media data are cost-effective to access and, more importantly, allow analysis over large populations which was not previously feasible [3, 12, 33].

The methods extracting vaccination attitudes from social media fall into two categories: community based and post based. Cossard et al. [17] found that pro- and anti-vaccine users naturally cluster into communities and calculated community partitions of various communication graphs to infer users' vaccination stances. Johnson et al. [33] made use of the topics of fan pages (similar to discussion groups) on Facebook to approximate users' attitudes and analysed the communities formed by 100 million users across the world in terms of their vaccination attitudes. Post-based methods benefit from the various types of information encoded in social media posts such as texts, labels and pictures. Gunaratne et al. [30] relied on the hashtags in tweets to approximate the vaccination attitudes in tweets. Sentiment analysis [65], as part of NLP, aims to derive multiple types of subjective opinions expressed in texts such as political stances [18, 25]. The introduction of deep learning leads to more powerful models that can process posts at the sentence or paragraph levels such as word2vec [42] and BERT [20, 49]. The sentiments extracted from texts have been used as references to study vaccine hesitancy [28]. For instance, Gbashi et al. [28] detected the opinions of media towards vaccines in Africa through Twitter and Google News. A few previous works, with similar goals to ours, leveraged transformers such as CT-BERT [46] to extract COVID-19 vaccination attitudes from tweets [53, 54]. Despite the promising performance in English tweets with more than 90% in terms of accuracy, such methods do not work well in multilingual datasets, as we will show in our experiments.

Discussion. The community-based methods cannot capture the fast development of public vaccination attitudes due to the relatively stable connections between users. Moreover, community memberships are effective for analysis on the level of populations but fail to accurately derive individual users' attitudes. The post-based methods in previous studies are not specifically designed and trained for COVID-19 vaccines. As a result, they cannot capture the special linguistic characteristics of the online discourses during the COVID-19 pandemic. This is partially because of the lack of social media posts which are related to COVID-19 vaccines and annotated with vaccination stances. In this article, we propose a framework that not only benefits from state-of-the-art post-based methods but also deals with the interference of linguistic features such as sarcasm and irony in discourses related to COVID-19 vaccination. We also create the first annotated dataset of tweets which can facilitate developing future models on subjective opinion extraction. Intuitively, our framework fuses the recent online discourses of the originating user's friends (including his/her own past posts) as part of the input to help mitigate the impact of sarcasm and irony.

The previous works that are closest to our work include the work of Cotelo et al. [18], Ebrahimi et al. [25] and Wang et al. [65], which leverage auxiliary information on social media in addition to texts to classify social media posts, e.g., structure of social connections [18, 25] and sentiment diffusion patterns [65]. However, they still rely on manually selected features either for texts under analysis or auxiliary information. For instance, Wang et al. [65] used features of information diffusion trees and diffusion networks to predict the likelihood of sentiment changes when social media posts propagate among users. This likelihood is then used to adjust the classification results of text-only classifiers. Cotelo et al. [18] represented social structural information with calculated community partitions and combined it with the structure embedding with word frequencies in the targeted text to infer political stances. One problem is that the features effective for one type of opinions may not result in the same level of performance for other opinions. Moreover, the existing works still use word-level representations for texts and are only tested effective in a single language. Although they can be retrained with new training data in other languages, it is not clear whether they can result in similar performance when dealing with multiple languages simultaneously. This prevents them from being applied in regions/countries with populations of various origins.

We also notice that other works exist to classify social media messages but for different purposes from ours, e.g., misinformation detection [23]. Auxiliary information that can contribute to the tasks is exploited to improve classification performance. For instance, Ducci et al. [23] proposed a new LSTM-based method for misinformation detection by leveraging the information diffusion process modelled by cascades. In this article, we will focus on vaccination attitudes and propose methods and discover auxiliary information that is not only effective in ensuring performance but also easy to access.

Different from the current literature, in this work we propose an end-to-end deep learning framework specifically for inferring vaccination attitudes which can automatically extract a useful representation of the inputs without manual intervention. This allows our framework to be easily adapted to infer other subjective opinions when properly retrained with new data. Considering the global requirement of monitoring vaccine hesitancy, we ensure that our framework can deal with multilingual texts.

3 Extracting Vaccination Attitudes from Social Media Posts

3.1 Problem Definition

Extracting the vaccination stance of a social media post can be technically formulated as an NLP task [39], i.e., classifying texts according to given class labels.In this article, we focus on the affective stance towards COVID-19vaccination. Thus, the set of labels is $\mathcal{L} = \{NE, PO, NG\}$, where *NE*, *PO* and *NG* correspond to neutral, positive and negative, respectively.Note that the ‘stances’ towards vaccination are different from ‘affective sentiments’ that are frequently or commonly applied in NLP. The commonly adopted concept of sentiments indicates the polarity of the affect expressed in a piece of text about something. Stances towards vaccination can also be interpreted as a type of subjective opinion, but they are specified in the context of whether people would like to get vaccinated. For instance, the poster of a tweet can be affectively negative about available vaccines and concerns about their safety, but he is still inclined to get vaccinated considering the consequences of being infected. This difference also helps explain our motivation to design a new sophisticated model to learn vaccination stances instead of referring to existing NLP sentiment analysis methods.As far as we know, there does not exist an available dataset with annotations of vaccination stances. As a result, we conductour own annotation and publish our dataset for subsequent similar research. To empirically show the difference between these two concepts, we also compare the sentiments extracted with state-of-the-art NLP methods to vaccination stances we manually annotate in Section 4.

The basic idea of text classification in NLP is first to calculate a representation of the given text which summarises its linguistic features and then output the most likely class label. Classification methods differ from each other in terms of the formats of text representation and the mapping from representations to class labels. Text classification confronts the challenge that the attitude or emotion expressed by the same words varies according to the context. For instance, the figurative usage of symptom words is found to mislead dictionary-based classification methods and significantly deteriorate classification accuracy [9]. Additional information is thus introduced to mitigate the ambiguity. To counterpart figurative inference of symptom words, text sentiments are used.

Similarly, we notice that during the COVID-19 pandemic, Twitter users tend to express their stances towards vaccination implicitly.This makes it difficult to capture their true attitudes through short messages of a maximum of 140 words. Inspired by Biddle et al. [9], we intend to see whether we can exploit auxiliary information which is cost-effective to collect and effective to reduce ambiguity. In this article, we consider two types of information accessible to social media platforms and data analysts.First, considering that people are unlikely to change their stances in a short period, given a target post, we combine the user's recent messages and use the attitudes expressed in them as a reference to infer his/her attitude. Second, inspired by collaborative filtering, a technique widely used in recommender systems which assumes that friends share similar preferences [21], we further leverage the user's friends' recent messages as the second reference. To achieve this, we require users' social connections which areusually captured by a *social graph*.

We use $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to represent the social graph recording the social relations between users, where $\mathcal{V}$ is the set of nodes and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges between nodes. A node $v \in V$ corresponds to a social media user,and an edge (v, v') indicates the existence of a relationship between two users v and v'. Note that we ignore the direction of relationships in this articleto take into account all the neighbours of a user, e.g., both followers and followees on Twitter.Thus, $(v, v') \in \mathcal{E}$ implies $(v', v) \in \mathcal{E}$.We abuse the terms *user* and *node* in the following discussion when it is clear from the context. Let $\mathcal{N}_i^k$ be the set of neighbours of node v_i within k hops,i.e., $\{v \mid d_{\mathcal{G}}(v, v_i) \leq k\}$, where $d_{\mathcal{G}}(v, v_i)$ is the shortest distance between v and v_i in the graph $\mathcal{G}$. Note that node v_i is also in $\mathcal{N}_i^k$ as $d_{\mathcal{G}}(v_i, v_i) = 0$. We use x_i^t to denote the textual message posted by user v_i at time t. We use $\mathcal{M}_i^{<t}$ to denote the list of posts originated by user v_i before time t chronologically ordered by their post time, and $\mathcal{M}_{\mathcal{N}_i^k}^{<t}$ to represent the set of post lists of the neighbours of user v_i within k hops.

Our COVID-19 vaccination attitude classification problem can be defined as calculating the probability distribution of all labels in $\mathcal{L}$. The final vaccination stance of x_i^t, i.e., $f(x_i^t)$, is determined by the label with the largest probability. Formally, we have

$$f(x_i^t) = \underset{stance \in \mathcal{L}}{\operatorname{argmax}} \Pr(stance \mid x_i^t, \mathcal{G}, \mathcal{M}_{\mathcal{N}_i^k}^{<t}).$$

3.2 A Vaccination Attitude Learning Framework

To solve the classification problem formulated previously, we propose a framework which takes advantage of the recent success of adopting deep learning in NLP and graph analysis such as transformer-based text embedding [46] and **graph neural networks (GNNs)** [36]. Figure 1 depicts an overview of our framework by labelling its three main components in different background colours: (i) a text-level information embedding module, (ii) a GNN-enhanced module and (iii) a classification module. The first module is used to learn the linguistic representation of the targeted post, whereas the second module summarises the pre-held attitudes of the originating user and those users who are close to him/her on the social media platform. We concatenate the outputs of these two modules as the input of the classification module which outputs a probability distribution over the labels.

The second module is critical for our framework in the sense that it prepares the two references we select as auxiliary information to reduce ambiguity. Its ultimate goal is to calculate an overall generalisation of the pre-held attitudes of the originating user and the users in his/her neighbourhood in the social graph. Specifically, it aims to achieve two tasks. The first task is to combine the recent messages of the originating user and those of the users in his/her neighbourhood. The output of this task is a representation vector for each concerned user that captures the linguistic features which are most representative of his/her pre-held vaccination attitude. We call this repre-

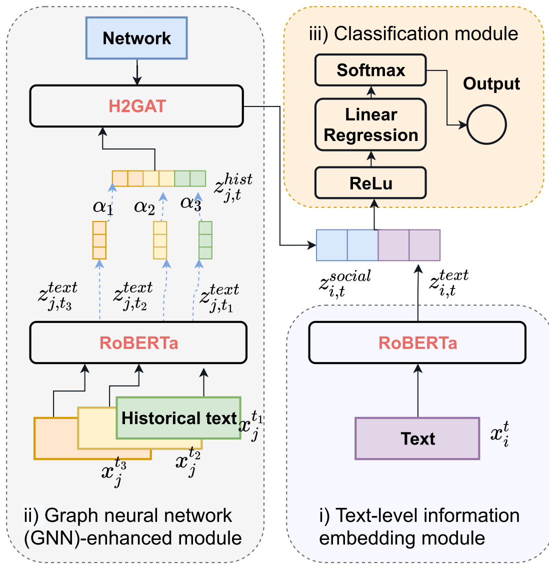

Fig. 1. An overview of our attitude classification framework consisting of three main components.

sentation vector a user's *pre-held attitude* representation. The second task is to efficiently combine the originating user's pre-held attitude representation with those of the users in his/her neighbourhood and calculate the final representation vector. In this work, we use a recent advance in graph deep learning, i.e., GNN, to implement the module.

GNN [36] has shown its advantage in transforming graph information, including structures and attributes of nodes and edges. Intuitively, it employs a message-passing scheme to integrate the attributes of a node's neighbourhood into the representation of the node. The calculated representation is a vector which is subsequently used for various downstream applications such as node classification and link prediction. Variants of GNN differ from each other in terms of the implementation of their message-passing schemes.

3.3 Our Model

Our framework can be instantiated by assembling various methods that can achieve the corresponding tasks of the modules. In the following, we present how we implement every module of the framework and justify our choices.

Text-Level Information Embedding. The goal of this module is to calculate a representation vector that can accurately summarise the linguistic features of the given text message. Due to the popularity of big models, recent works leverage pre-trained NLP transformers to calculate the representation for short texts [40, 49]. NLP transformers have been empirically evaluated in the work of Zhao et al. [70], where RoBERTa [41] is shown to outperform the competing models. As a result, we adopt RoBERTa to learn text representations in our model. The model takes a textual post, e.g., x_i^t, as input, and outputs a sentiment-oriented representation vector $z_{i,t}^{text} \in \mathbb{R}^d$, where d is the pre-defined dimension of the vector. RoBERTa will be fine-tuned with the posts in the training set.

GNN-Enhanced Module. As we emphasised previously, we make use of GNN to summarise the pre-held vaccination attitudes of the originating user and the users who are close to him/her in the social graph. Formally, given the target post x_i^t, we utilise this module to capture the representative linguistic features of the recent messages posted by v_i and his friends within k hops before t. The definition of *recent* is flexible and depends on application scenarios. In this work, we select the last λ tweets before t as a user's recent tweets. Its output will be subsequently used as complementary contextual information to ameliorate classification performance. Therefore, the input of this module consists of the social network $\mathcal{G}$ and the historical messages of user v_i's k-hop neighbours, i.e., $\mathcal{M}^{<t}_{\mathcal{N}_i^k}$. Note that the post originator's pre-held attitude is also part of the calculation as v_i is included in $\mathcal{N}_i^k$.We take two successive steps to calculate the final output vector $z_{i,t}$. We first integrate the recent posts of each user in v_i's k-hop neighbourhood as a summary of his/her vaccination discourse.In the second step, we propose a new GNN-based model named *H2GAT* to aggregate the discourse of all v_i's k-hop neighbours.

Step 1: Text-Level Encoding. For each user $v_j \in \mathcal{N}_i^k$, we use his/her last λ posts in $\mathcal{M}_j^{<t}$,denoted by the list $(x_j^{t_\lambda}, x_j^{t_{\lambda-1}}, \ldots, x_j^{t_1})$ where t_m $(1 \le m \le \lambda)$ is the timestamp of v_j's last m-th post. We apply the pre-trained RoBERTa model and then obtain the corresponding list of text-level representations, i.e., $(z^{text}_{j,t_\lambda}, z^{text}_{j,t_{\lambda-1}}, \ldots, z^{text}_{j,t_1})$. There are many ways to integrate v_j's past text-level representations while distinguishing their various temporal importance, e.g., Hawkes [31] and GRU [15]. In our implementation, we adopt the dynamic-aware *position encoding* by assigning a fixed importance factor to each past post according to its position in the list. We also test other methods as shown in Section 5. The results show that position encoding is the most effective despite its simplicity. The reason behind this may be because users' stances are relatively stable and the importance of a message is not determined by the previous ones as much as in the scenarios where Hawkes and GRU apply.

Formally, the integrated text-level representation of user v_j is calculated as follows:

$$z^{hist}_{j,t} = \sum_{m \le \lambda} \alpha_m \cdot z^{text}_{j,t_m},$$

where α_m $(1 \le m \le \lambda)$ is trainable and describesthe positional relation between the past posts. Note that k and λ are pre-defined hyper-parameters that should be tuned manually.

Step 2: H2GAT. It is pointed out that the *heterophily* phenomenon widely exists among online social network users [51]. This phenomenon also exists in vaccination-related discourses as the attitudes and linguistic features can be significantly different between users, although they may hold similar attitudes towards vaccination. Considering the heterophily of vaccination discussion in a user's neighbourhood, we adopt and extend a recent GNN-based method called *H2GCN* [71]. The same as other GNNs, it also has multiple layers, the ℓ-th of which can be formulated as follows:

$$H_i^\ell = \mathsf{Combine}(\{\mathsf{Aggregate}\{H_j^{\ell-1} : j \in \widehat{\mathcal{N}}_i^{k'}\} : k' \in \{1, \ldots, k\}\}),$$

where $\widehat{\mathcal{N}}_i^{k'}$ represents node v_i's k'-order neighbours, i.e., nodes that have an exact distance of k' from v_i in $\mathcal{G}$. Formally, $\widehat{\mathcal{N}}_i^{k'} = \{v_j \mid d_{\mathcal{G}}(v_j, v_i) = k'\}$. Note the difference of $\widehat{\mathcal{N}}_i^{k'}$ from $\mathcal{N}_i^{k'}$. As $\mathcal{G}$ is connected, when $k' > 0$, we have $\widehat{\mathcal{N}}_i^{k'} \subset \mathcal{N}_i^{k'}$. The initial H_j^0 for each $v_j \in \mathcal{V}$ is set to $z^{hist}_{j,t}$.

Note that different from the work of Zhu et al. [71], which adopts the Aggregate function of GCN [36], we use GAT [64] to distinguish the contributions of friends for better performance. The output representation of v_i, denoted by $z^{social}_{i,t}$, is calculated by combining node representations of all layers:

$$z^{social}_{i,t} = \mathsf{Combine}(\{H_i^0, H_i^1, \ldots, H_i^k\}).$$

Many ways exist to implement the function Combine. We adopt the one in H2GCN [71] in our model which outputs the concatenation of all inputs.

Table 1. Statistics of Our Data Collection and Annotation

Social network	**Number of nodes**	France	35,081
		Germany	16,304
		Belgium	15,647
		Luxembourg	2,904
		total	**69,936**
	Number of edges		8,909,985
	Average degree		127.40
Tweets	**Number of tweets**	France	5,925,354
		Germany	2,669,875
		Belgium	530,885
		Luxembourg	9,279
		total	**9,135,393**
	Number of tweets/user		130.62
Annotated tweets	**Number of tweets**		18,246
	Number of users		4,157

By concatenating $z_{i,t}^{social}$ and $z_{i,t}^{text}$, we obtain the text-level representation vector for classification. Formally,

$$z_{i,t} = z_{i,t}^{social} \parallel z_{i,t}^{text}.$$

Attitude Classification. We implement a simple two-layer structure to conduct the classification. Note the output vector has a length of $|\mathcal{L}|$. Recall that $\mathcal{L}$ is the set of class labels. The first layer applies the ReLU function to each element in $z_{i,t}$ and at the second layer, we use a linear regression where $W \in \mathbb{R}^{2d\times|\mathcal{L}|}$ and $b \in \mathbb{R}^{|\mathcal{L}|}$ is the bias vector. The softmax function is applied to calculate the probability distribution over the class labels. Formally, the distribution p is calculated as follows:

$$p = \mathsf{softmax}\left(\mathsf{ReLU}(z_{i,t}) \cdot W + b\right).$$

The class label with the largest probability will be chosen as the output. We use CrossEntropyLoss as the objective function to optimise the entire model.

4 Data Curation

To train and evaluate our vaccination attitude learning framework, we collect a dataset from Twitter focusing on four adjacent Western European countries: Germany, France, Luxembourg and Belgium. Reasons to select these countries include their importance to European economy and their similar and almost synchronous pandemic management policies. They can also well represent the first group of countries receiving and administering COVID-19 vaccines. Due to the lack of publicly available data on vaccination stances, we create the first public set of annotated tweets for training and testing our framework. The statistics about our dataset are summarised in Table 1. This dataset and our annotation are publicly available.[2]

[2]The dataset is publicly accessible at https://doi.org/10.5281/zenodo.5851407.

4.1 Data Collection and Preprocessing

Our dataset consists of two types of data: (i) a social network composed of active Twitter users, and (ii) the tweets of selected users related to COVID-19 vaccine or vaccination. By 'active users', we mean users who are active in vaccination-related discussions and frequently interact with others.

Step 1: Social Graph Construction. We start with identifying the Twitter users in our targeted region who actively participated in vaccine-related discourse. Instead of directly searching tweets by keywords, we refer to a publicly available dataset which contains the IDs of COVID-19 related tweets [13]. We extract the tweet IDs spanning between January 22, 2020, and March 15, 2021, covering the beginning of the vaccination campaign. Through these IDs, we download the corresponding tweets. Each downloaded tweet is associated with meta-information which includes the location of the originator, either self-provided by the originator or attached by the device's positioning services such as GPS. Due to the ambiguity of the self-reported locations, we use the geocoding APIs, Geopy and ArcGis Geocoding to regularise their formats. For example, a user input location *Moselle* is transformed to a precise and machine-parsable location: *Mosselle, Lorraine, France*. Based on the regularised locations, we filter the downloaded tweets and remove those posted by users out of the region. In total, we obtain 990,448 tweets from 767,583 users.

To find the users with frequent interaction with other Twitter users, we construct a retweeting weighted graph. An edge is created between two users if one user retweeted a tweet from the other user or mentioned him/her. The edge weight is assigned as the number of mentions or retweets between them. We remove all edges with weights smaller than 2 and calculate the largest weakly connected component of the graph which consists of 72,960 active users. As retweeting or mentioning a user does not mean that these two users have a following relationship, we crawl the remaining users' followers, with which a graph is constructed with the remaining users and the relationships between them as edges. In the end, we take the largest weakly connected component of the resulting graph as the final social graph. This graph consists of 69,936 nodes and 8,909,985 edges. On average, each user has 127.4 followers. This indicates that our selected users are sufficiently active on Twitter.

Step 2: Vaccine-Related Tweet Collection. In this step, we crawl the tweets originated or retweeted by the users in the social graph. Note that we are only interested in the tweets related to COVID-19 vaccination. We use a list of keywords to filter out the irrelevant ones. The keywords should be general enough to cover tweets in German, French and English. With our observation and several trials, we select the keywords containing the following strings: *vax*, *vaccin*, *covidvic*, *impfstoff*, *vacin*, *vacuna* and *impfung*. We use the Twitter Academic Research API to download tweets. Due to the limitations of the API, each request can only download a maximum of 500 tweets. To be efficient and ensure the coverage ratio, we create a download request for every user in each month. In total, we collect 1,626,472 tweets, and each user has about 130 tweets on average. We clean the downloaded tweets by removing mentions of other users with '@', quoted hyperlinks and 'RT'.

In Figure 2, we show the distribution of the languages of the collected tweets. We can see that as official languages in our target region, French and German are the major languages in our dataset. This makes our dataset rather different from other published datasets with English as the dominant language [54].

4.2 Data Annotation

According to the best of our knowledge, no datasets of social media posts are publicly available with users' COVID-19 vaccination attitudes annotated. As a result, we select a subset of our downloaded tweets and manually attach them with attitude labels. In this article, we focus on users'

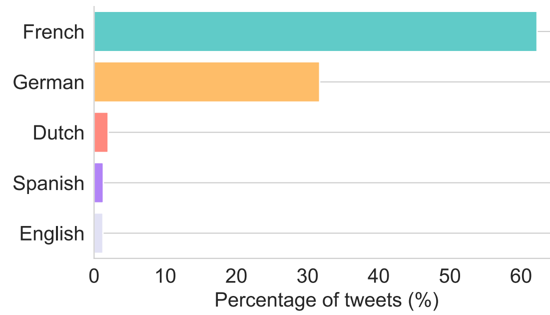

Fig. 2. Distribution of tweets in the top five most populated languages in our collection.

Table 2. Annotation Labels with Small Example Tweets

Label	Examples (translated to English)
positive	vaccines: why there are no long-term side effects.
	We have a new weapon against the virus: the vaccine. Hold together, again.
negative	this nurse gets covid-19 vaccine. then she talks to media how great it is. then passes out. watch!
	Corona mass vaccination aka disability manufacturing
neutral	How safe is the Covid 19 vaccine for people with diabetes?
	If there is a covid vaccine, what will you do?
positive but dissatisfied with government management	It's bad enough for individuals to refuse #COVID19 #vaccines for themselves. But forcing a mass vax site to shut down, knowing it means vaccines may go to waste, is criminal. Call it pandemicide.

affective stances towards COVID-19 vaccination which are *positive*, *negative* and *neutral*. Table 2 lists examples of the attitude labels. After a closer check of the tweets, we notice a relatively large number of tweets in which users express their dissatisfaction or complaints about governments' COVID-19 management policies but possess positive attitudes towards vaccination. Take the last row in Table 2 for an example. It contains a few negative words, such as 'bad', 'criminal' and 'waste', but the originator explicitly expresses his/her support for vaccination. As such tweets deliver a negative emotion, if we do not separate them from those with negative vaccination attitudes, NLP models will be confused and their classification accuracy will be deteriorated. It is obvious that there exist other tweets also expressing combined emotions such as 'positive but doubt about the effectiveness'. However, according to our analysis of collected tweets, the number of such tweets is rather small. Therefore, considering the relatively larger proportion, we add label *PD* indicating 'positive but dissatisfied with government management'. Then the set of labels $\mathcal{L}$, used in the rest of the article, consists of four labels, i.e., $\{PO, NG, NE, PD\}$.

We select tweets to be annotated to cover all the users in our dataset. We order our downloaded tweets in a list according to the number of times being retweeted in descending order. We then iteratively remove the most transmitted tweet from the list and put it into the list of selected tweets until every user in our dataset originally posted or retweeted at least one selected message. After this step, we select 18,246 tweets originating from 4,157 users.

As our tweets are multilingual, we hire 10 bachelor students who can speak at least two of the three most used languages (i.e., German, French and English). One author of this article acts as the coordinator and trains all the hired annotators by explaining the semantics of all labels with examples. To ensure that all annotators hold the same understanding of the labels, they are asked to annotate 200 tweets. The coordinator checks their annotation and communicates to them with extra explanation when necessary. We conduct three rounds to make sure each tweet's label is double validated. In the first round, an annotator annotates all the selected tweets. In the second round, each of the remaining nine annotators is randomly assigned 2,000 tweets and validates the

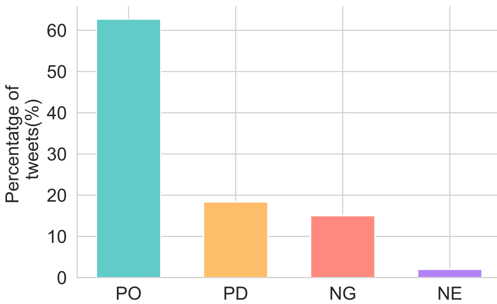

Fig. 3. Distribution of vaccination attitude labels of the annotated tweets. PO, positive; NG, negative; NE, neutral; PD, positive but dissatisfaction.

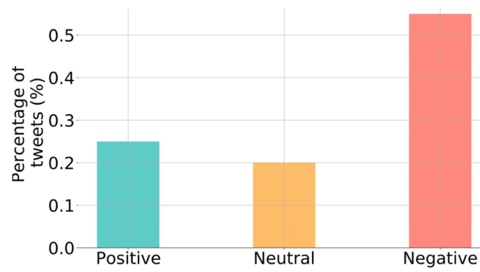

Fig. 4. Distribution of sentiments of annotated tweets.

Table 3. Inter-Annotator Agreement

Label	AOA	Feliss' Kappa	Krippendorf's Alpha
PO	0.72	0.73	0.73
NG	0.82	0.88	0.88
NE	0.74	0.78	0.77
PD	0.61	0.63	0.62

PO, positive; *NG*, negative; *NE*, neutral; *PD*, positive but dissatisfaction.

annotation. In the last round, the coordinator goes through all validated annotations. In Figure 3, we show the distributions of the tweets according to their labels.

Recall that the attitude of a tweet towards vaccination annotated differs from its affective sentiment which is usually adopted in the task of sentiment analysis in NLP. In Figure 4, we show the distribution of sentiments of the annotated tweets which are classified into 'positive', 'neutral' and 'negative'. The sentiments are extracted by a multilingual sentiment analysis model which deploys XLM-R [16] as the text embedding model and is trained on the widely accepted *SemEval-2017 Task 4A* dataset [58]. We can see that the two distributions are significantly different. Considering the concerns about available vaccines' safety, especially with the negative news about their side effects, more than half of the annotated tweets are negative in sentiment analysis. By contrast, more than 60% express posters' inclination to get vaccinated according to our manual annotation.

Annotator Agreement. We leverage three widely accepted measurements to evaluate the inter-annotator reliability for each label: AOA (Average Observed Agreement) [27], Fleiss' kappa [27], and Krippendorff's alpha [37]. The values of all the three measurements range from 0 to 1, where 0 indicates complete disagreement and 1 indicates absolute agreement. Table 3 summarises the inter-annotator agreement for each annotation label. We can see that for labels *PO*, *NG* and *NE*, all three measurements produce scores larger than 0.73, indicating an outstanding agreement level. The annotators' agreement on *PD* falls drastically compared to other labels but still remains moderate according to the ranking criteria of the Fleiss' kappa measurement. This can be explained by our difficulties during annotation in dealing with the special lingual features of *PD* tweets, e.g., frequently used negative terms or sarcastic expressions.

5 Experimental Evaluation

5.1 Evaluation Setup

We set up an evaluation pipeline following the approach for traditional supervised classification [43]. Specifically, we split labelled tweets into training (80%), validation (10%) and testing

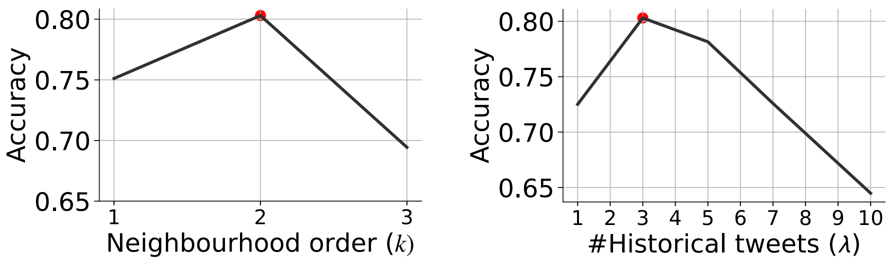

Fig. 5. Parameter tuning for k and λ, performed on a validation dataset.

(10%) sets. The models are optimised with the training set, and the validation set is used to tune hyper-parameters. The model performance is evaluated on the testing set. To reconfirm the overwhelming performance of transformer-based models, we also implement two traditional machine learning models that are widely used in the literature for similar tasks such as sentiment analysis: SVM and Random Forests. We use TfidfVectorizer [52] to convert tweets into the bag of n-gram vectors and use them as the input textual features for these two models. Grid search is utilised as the optimiser to identify the hyper-parameters.

Tweet Preprocessing. The collected tweets contain texts that may be irrelevant to users' subjective opinions, such as mentions of other users with '@', quoted hyperlinks and the keyword "RT" which stands for "retweet" and indicates the states of the tweets. We preprocess the tweets by removing these irrelevant parts to ensure classification accuracy.

In addition to original posts, retweets and quotations take up a large proportion of a user's historical posts. For quotations, a user added some comments which may express opposite opinions to that of the quoted one. Therefore, we only use the texts users upload as valid posts encoding users' vaccination stances. Although retweets cannot fully represent a user's own opinion, the behaviour of retweeting itself indicates some sort of agreement with the ideas expressed in the message retweeted [24]. Based on this idea, we take retweets into account when calculating an individual user's vaccine hesitancy. This is also adopted in the vaccination attitude tracking discussed in the previous section. Note that the same preprocessing criteria are applied for all the analyses presented in this article.

Hyper-Parameter Settings. We train our model for 400 epochs and use Adam [35] for optimisation with the learning rate of 10^{-5} and weight decay of 5×10^{-4}. For the text encoder, i.e., RoBERTa, we adopt the implementation XLM-RoBERTa [50] and follow their default settings where the maximum string length, i.e., parameter d, is 128. For our GNN-enhanced module, we set the embedding dimension as 64. The neighbourhood order k which is also the number of layers and the number of historical tweets λ are important to ensure representation quality. Therefore, we conduct an empirical study to analyse the influence of these two key hyper-parameters to ensure the final performance. In Figure 5, we present the classification accuracy with different values of k (on the left) and λ (on the right) on our validation dataset. We observe that these two hyper-parameters indeed significantly influence classification accuracy. Our model arrives at the best performance with $k = 2$ and $\lambda = 3$.

5.2 Experimental Results

We have two objectives. One is to validate the overwhelming performance of our model over traditional machine learning models that are widely used in the literature, i.e., Random Forests and SVM. The other is to compare other possible implementations of our proposed vaccination

Table 4. Performance Comparison of Selected Baselines and Variants of Our Framework

Model	Precision	Recall	F1	Accuracy
Random Forests	0.4334±0.0775	0.3766±0.0716	0.4491±0.0955	0.4187±0.0606
SVM	0.4033±0.0460	0.3944±0.0196	0.4382±0.0403	0.4168±0.0110
FastText	0.5573±0.0301	0.5665±0.0268	0.5585±0.0202	0.5651±0.0215
BERT-Large	0.6170±0.0031	0.6205±0.0041	0.6191±0.0096	0.6189±0.0078
CT-BERT	0.6385±0.0161	0.6361±0.0105	0.6375±0.0205	0.6372±0.0014
mBERT	0.6622±0.0101	0.5769±0.0153	0.6132±0.0113	0.6466±0.0087
RoBERTa	0.6758±0.0218	0.5848±0.0232	0.6249±0.0210	0.6517±0.0348
GCN+MEAN	0.6936±0.0101	0.6932±0.0153	0.6890±0.0113	0.7033±0.0087
GAT+MEAN	0.7001±0.0091	0.7002±0.0102	0.6983±0.0078	0.7096±0.0099
H2GCN+MEAN	0.7387±0.0021	0.7144±0.0081	0.7286±0.0027	0.7412±0.0015
H2GAT+MEAN	0.7361±0.0038	0.7371±0.0007	0.7331±0.0019	0.7461±0.0008
H2GCN+GRU	0.7813±0.0025	0.7794±0.0048	0.7712±0.0027	0.7829±0.0010
H2GAT+GRU	0.7937±0.0036	0.7968±0.0021	0.7927±0.0011	0.8009±0.0008
H2GCN+Hawkes	0.7843±0.0016	0.7760±0.0038	0.7699±0.0047	0.7831±0.0025
H2GAT+Hawkes	0.7946±0.0036	0.7922±0.0010	0.7903±0.0019	0.7988±0.0023
H2GCN+PE	0.7859±0.0031	0.7813±0.0045	0.7792±0.0022	0.7889±0.0028
H2GAT+PE	**0.7948**±0.0051	**0.7954**±0.0029	**0.7936**±0.0020	**0.8030**±0.0017

attitude learning framework. To distinguish these models, we name them with two parts concatenated with '+'. The first part tells the adopted GNN model, whereas the second part gives the method for handling past tweets' temporal importance. As all models use RoBERTa for text encoding, we do not explicitly put it in the model names. We present their performances in Table 4.

We reconfirmed the existing studies that transformer-based methods outperform traditional machine learning models. They can only guarantee an accuracy of around 0.4, which is at least 0.22 less than all other models. Moreover, we have four major observations that justify the effectiveness of our implementation. First, the consideration of friends' vaccination discourse increases the performance. The text-only classification model with RoBERTa only has an average accuracy of 0.65, whereas the other models, which are implemented with the GNN-enhanced module, achieve at least an accuracy above 0.70. Second, the vaccination discourse between friends on Twitter is heterophily, and the choice of heterophily-aware GNN models, i.e., H2GCN and our H2GAT, can further significantly improve the performance. The the four models following RoBERTa in Table 4 have the same settings except for the GNN methods. Both the application of H2GCN and our H2GAT achieve an increase of about 0.04 compared to the models with GCN [36] and GAT [64]. Third, the consideration of the temporal importance of past tweets leads to another up to 0.06 improvement. This amount of increase is notable for models of subjective opinion extraction due to the flexibility and diversity of users' habits and preferences in composing tweets, especially in a multilingual scenario. We consider four methods to combine a user's last λ tweets: MEAN, GRU, Hawkes and PE (short for positioning encoding). The method denoted by MEAN simply averages the text-level encodings. The PE method adopted in our model generates the best performance. Fourth, our extended H2GAT model outperforms the original H2GCN. Our implementation, i.e., H2GAT+PE, finally improves the text-only RoBERTa model by 23% in terms of accuracy.

Empirical Complexity Analysis. As the RoBERTa model is pre-trained, the models instantiated from our framework have the same complexity as the adopted GNN models. In our experiments, we conduct the training on a server with Xeon E5 CPU and Tesla V100 GPU. On average, the training

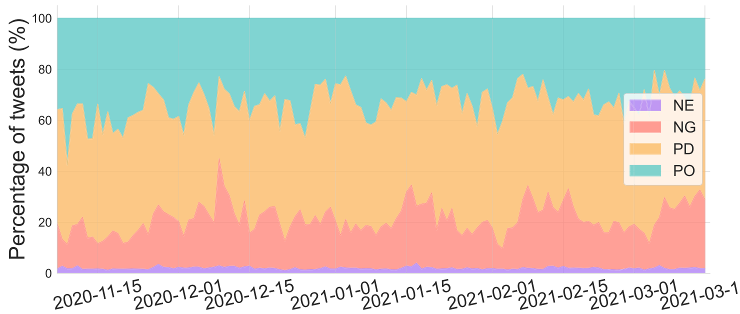

Fig. 6. Temporal distribution of tweets with different vaccination attitude labels.

time for RoBERTa is about 115 seconds for an epoch while 52.5 seconds are needed for an epoch in training the GNN-enhanced module and the classification module. What is more important is the running time of the models when processing a tweet. This will determine the practical utility of our framework in tracking public vaccination attitude in real time. We run four parallel instances of our model H2GAT+PE on the server. On average, it takes 24.68 seconds for every 1,000 tweets, which means that more than 3.5 million can be processed a day. For the regions we target, we collect in total 9 million vaccination-related tweets over 2 years. This implies that our model is sufficiently efficient for processing posts on a daily basis.

Cross-Validation. In addition to experimental evaluation, we also make use of published social studies to cross-validate our model's effectiveness. Lazarus et al. [38] conducted a survey in June 2020, and estimated that the vaccine acceptance rates in France and Germany are 58.9% and 64.5%, respectively. After applying our model to classify the tweets in the same period, we find that the percentages of tweets with positive vaccination attitudes of these two countries are 42.27% and 53.12%, which are similar and retain the relative difference between the two countries. This implies that posts on Twitter can be used as a reference to fast grasp vaccine hesitancy when surveys are not available.

Vaccine Hesitancy Tracking and Manual Analysis. We draw the temporal evolution of the percentage of tweets for each selected label in Figure 6 daily starting from November 8, 2020. More specifically, we apply our well-trained model to all the selected tweets and assign a label indicating the corresponding vaccination attitude. Note that we select a day-level granularity which we believe is appropriate for healthcare departments to monitor the public's vaccination attitudes. We acknowledge that the granularity should be determined by the specific application. Based on previous research reporting that the content of tweets is highly correlated with real-world situations [56], we make a hypothesis that real-world events may contribute to the fluctuating proportion of tweets with different vaccination stances. In vaccine hesitancy monitoring, special attention should be paid to the fluctuations of negative attitudes which may correspond to certain events. For instance, we can see that the first apex of negative tweets in Figure 6 occurred around December 7, 2020, which is the day when the UK started vaccination for COVID-19. We take another three time points that correspond to turning points of the curve of label *NE* as examples and discuss the potential causes. Among them, two correspond to apex points where negative tweets reach local maximum percentages and one corresponds to a base point with local minimum negative tweets. We first plot word clouds in Figure 7 to identify the most frequently used keywords in the week around the selected points. Then we search these keywords on the Internet to identify the events that may contribute to the changes.

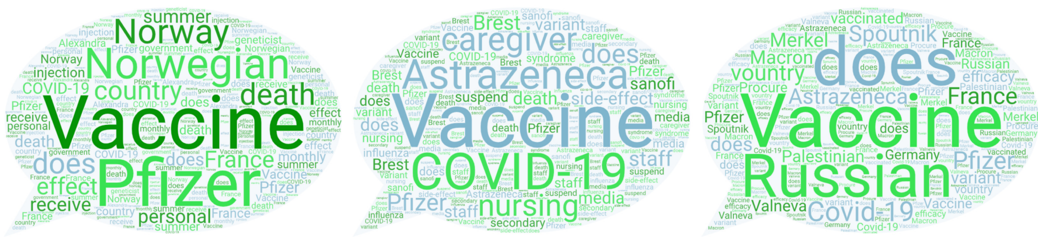

(a) Apex 1 (01/14–01/21) (b) Apex 2 (02/10–02/17) (c) The base (01/30–02/06)

Fig. 7. Word clouds of tweets around selected points.

The first apex occurred around January 16, 2021. We notice that this surge of negative tweets is attributed to the propagation of a large volume of misinformation. Take the two most dominant pieces of misinformation as examples. One said that on January 14, the Norwegian Medicines Agency reported that a total of 29 people had suffered side effects, 13 of which were fatal. The other was about the death of an Indian healthcare worker after receiving COVID-19 vaccines. The second peak happened around February 15, 2021. One piece of negative news was reported that AstraZeneca vaccines were stopped from administration after many health workers of Morlaix Hospital in France suffered from side effects. This news subsequently led to anti-vaccination discussions. The base point occurred between the two peaks around February 3, 2021. From Figure 7(c), we find the dominant positive news that Russia started to offer other countries such as Pakistan with its vaccines.

In addition, we conducted a recent study [14] with the proposed model to measure individual users' vaccine hesitancy levels and successfully confirmed the consistency of Twitter with social surveys in terms of vaccine hesitancy surveillance. The consistency persists across regions and along with time. From the preceding discussion, we can see that our model can enable the use of social media data to track on a daily basis the changes in vaccination attitudes, and capture the impact of social events on public vaccine hesitancy. This may finally help the governments identify the right time to take intervention actions.

5.3 Generality Test of Our Model

In previous sections, we have demonstrated the effectiveness of our model in understanding vaccination attitudes from social media users' textual messages. In the following, we would like to show that our model is general enough to be adapted to extract other public opinions. To test its generality, we apply our model to identify the political ideology of Twitter users from their messages on Twitter. We constructed a new training set from our Twitter dataset using the public Twitter parliamentarians dataset [63]. This dataset consists of verified parliamentarians from 26 countries (including France, Germany and Luxembourg) with their names, political ideology and Twitter IDs. A person's political ideology is estimated and assigned a value between 0 and 10. We extracted a list of parliamentarians from the three countries: France, Germany and Luxembourg. Due to the incompleteness of the dataset [63], we added politicians who joined after the data were released and updated the out-of-date Twitter IDs. In the end, we prepared a training dataset of 1,021 parliamentarians. We classified the political ideology scores (from 1 to 10) into left, centre and right. A politician with a score smaller than 4 is considered as 'left'. Those with scores larger than 6 are labelled as 'right'. Otherwise, they are annotated as 'centre'.

We reduce the political ideology detection task to classifying a user's past discourse on Twitter into left, right or centre. In this case, we simply combine all of a user's past tweets into one large single textual message which is subsequently treated as the message to classify. To

Table 5. Model Performances for Political Ideology Detection with Applicable Baselines and Our Model Variant of the Best Performance

Model	Precision	Recall	F1	Accuracy
Random Forests	0.508	0.503	0.511	0.509
SVM	0.505	0.504	0.508	0.504
FastText	0.526	0.519	0.517	0.522
BERT-Large	0.591	0.591	0.592	0.590
mBERT	0.633	0.631	0.628	0.630
RoBERTa	0.657	0.688	0.662	0.653
TIMME	0.771	0.759	0.788	0.796
H2GAT+PE	**0.790**	**0.776**	**0.798**	**0.804**

ensure a convincing comparison, in addition to the benchmarks used previously, we also include *TIMME*, the multi-task multi-relational embedding model [68] which is specifically designed for political ideology detection. TIMME uses multiple types of relationships between Twitter users including *following*, *retweet*, *mention* and *like* together with tweets to infer users' bipolar political ideologies, i.e., left and right. We extend the original design of TIMME to deal with the triple classification. Moreover, as TIMME is originally implemented for English-only data, we replace the word-level embedding with RoBERTa [50] to handle multilingual texts. Note that CT-BERT is excluded in comparison because it is a specific model trained on messages related to COVID-19.

In Table 5, we show the performance of all the benchmarks and H2GAT+PE. We can see that TIMME and our model outperform all the other benchmarks at a minimum by 18%. Our model performs better than TIMME even if TIMME is designed specifically for the task with carefully catered input. The results show that our model has the potential to be generalised to other analyses based on texts.

6 Use Case: Predicting Vaccination Hesitancy Changes

In this section, we illustrate the use of our vaccination attitude learning framework. Specifically, we analyse the role of the vaccination information widely spread across Twitter in affecting users' attitudes towards vaccination. Considering the comprehensiveness of vaccination discourses, we classify the most popular vaccination-related tweets into *themes* that may correlate with vaccine hesitancy. Based on users' perceived information in these themes, we forecast their vaccination attitude changes with classic machine learning models.

6.1 Period Selection and Theme Labelling

The participation in vaccination discourses fluctuates over time along with the occurrence of social events related to COVID-19 vaccines. We select two time periods after the start of COVID-19 vaccination campaign, in which the volume of tweets experiences significant surges compared to adjacent periods. The first period lasts for 25 days spanning from December 27, 2020, to January 20, 2021, whereas the second period lasts for 15 days between January 25 and February 8, 2021. These two periods involve 161,611 original tweets posted in total, among which 25,449 are retweeted at least once. The total number of times of being retweeted adds up to 242,129. We encounter two challenges to continue our analysis of the impact made by diffused information: the comprehensiveness and large volume of propagated tweets. Due to the huge volume of tweets propagated over Twitter, it is not plausible to consider all of them. Previous studies show that tweets' influence

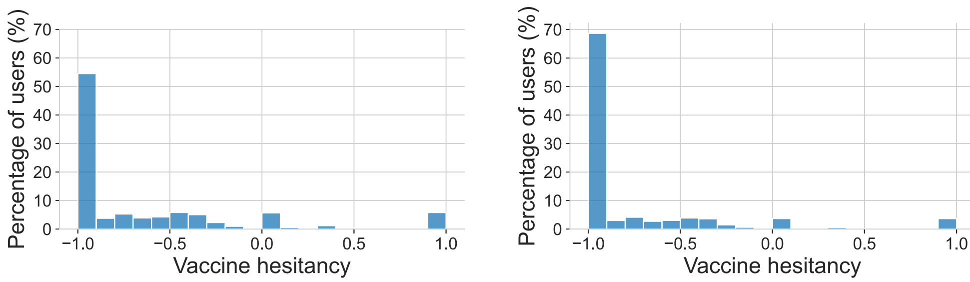

(a) Period 1 (12/27/2020-01/20/2021)

(b) Period 2 (01/25/2021-02/08/2021)

Fig. 8. Distribution of user's vaccine hesitancy for the selected two time periods.

follows the power-law distribution and that 80% of the impacts come from 20% of the most widely spread tweets [26]. Inspired by this result, we leverage the top 25% most widely propagated tweets in every period to approximately represent the themes expressed in the diffused information. In total, we select 501 original tweets that are retweeted 78,891 times from 72.16% of the users. To deal with the comprehensiveness, with a careful examination of the selected tweets, we categorise them into themes that are considered to be responsible for the changes of vaccination attitudes. We refer to previous studies [8, 19], especially the PACV (Parent Attitudes about Childhood Vaccines) survey [47] and the WHO Vaccine Hesitancy Matrix [48], and identify 11 themes that are relevant and can cover the propagated tweets (see Table 7 in Appendix A for explanation and examples). We ask 2 of the 10 hired annotators to manually annotate the selected tweets with their corresponding themes. The Cohen's kappa coefficient $k = 0.82$ implies a high rate of agreement between them.

6.2 Predictability of Vaccine Hesitancy Changes

Quantifying Individual Vaccine Hesitancy. We measure the vaccine hesitancy of an individual user according to the tweets posted or retweeted by the user in a time interval. Formally, it is calculated as $\frac{N_n(v)-N_p(v)}{N_p(v)+N_n(v)}$, where $N_p(v)$ denotes the number of posts with positive vaccination attitudes of user v during the selected interval, and $N_n(v)$ is the corresponding number of tweets with negative attitudes.

Since our purpose is to showcase another usage of our framework, we do not distinguish the various significance of original posts and retweets. Note that it may be unreliable to simply use users' attitudes expressed in their tweets to estimate their vaccine hesitancy. For instance, a user may express negative attitudes on social media but still get vaccinated. We believe that this is still an effective indicator of the level of a user's vaccine hesitancy. We also conduct cross-validation with surveys in prior work [14] with the same approach, and the results show the consistency between them and the feasibility of our approach in practice. For each selected period, we use the tweets posted 14 days before and after the period to evaluate individual users' hesitancy levels and see how they change. To ensure the reliability, we only consider the users who posted or retweeted at least three tweets. If a user's vaccine hesitancy experiences a change smaller than 0.05, we consider the user's attitude *unchanged* and otherwise *increased* or *decreased* depending on the change direction. In Figure 8, we show the distribution of users' vaccine hesitancy in the two selected periods. We can see that a large portion of users support vaccination.

Modelling Perceived Information. A Twitter user perceives information from the tweets retweeted or originally posted by his/her direct friends. As our focus is the information widely

Table 6. Model Performances for Attitude Change Prediction

Model	Precision	Recall	F1	Accuracy
SVM	0.7374	0.7382	0.7345	0.7477
Naive Bayes	0.6468	0.6559	0.6427	0.6658
Random Forest	0.6811	0.6838	0.6795	0.6958
XGBoost	0.7232	0.7229	0.7198	0.7342
GBDT	**0.7533**	**0.7516**	**0.7498**	0.7603

diffused on Twitter, we use a vector $I_u = (c_1, c_2, \ldots, c_m)$ to approximately represent a user u's perceived information, where c_i is the number of popular tweets a user receives from followers in the i-th theme. As we have 11 themes, $m = 11$, in our analysis.

Model Evaluation. We make use of various standard machine learning methods to predict the change of a user u's vaccination attitudes with the input of I_u. The methods consist of SVM with rbf kernel ($C = 1, \gamma = 0.1$), Naive Bayes ($\alpha = 1$), Random Forest (100 trees with maximum tree depth of 5), XGBoost (100 trees with a maximum tree depth of 4) and GBDT (100 trees with a maximum tree depth of 5). Table 6 shows the performance of these methods. All numbers are averaged over five training sessions. We can see that all methods can achieve reasonably good prediction performance and that GBDT outperforms the rest of the models with an accuracy of 0.76. When we consider additional factors such as users' vaccine hesitancy levels before the periods, the accuracy can be improved to 0.86.

Discussion. These results show that we can make accurate predictions with users' perceived popular information. Since we have empirically illustrated the plausibility of using social media posts to track public vaccination attitudes, the results imply that the diffused information on social media like Twitter could be used as indicators to forecast the changes in vaccine hesitancy levels. As repeated many times, although such predictions cannot achieve the same level of trust as social surveys, they provide decision makers with a method to quickly understand and get prepared for the potential damage of misinformation or compare different vaccine hesitancy intervention strategies over social media.

7 Conclusion and Discussion

In this article, we proposed a deep learning framework to learn vaccination attitudes from social media textual posts. Although vaccination attitudes extracted from social media cannot be as accurate and reliable as conventional social surveys, our framework allows for continuous tracking of the fast development of public vaccination attitudes and capturing the changes that deserve specific attention in time. By leveraging friends' vaccination discourse as contextual information, our model successfully reduces the interference of linguistic features such as sarcasm and irony. Our model instantiated from the framework improves the state-of-the-art text-only method by up to 23% in terms of accuracy according to our manually annotated dataset. With cross-validation with published statistics and manual analysis, we further validated the effectiveness of the model to capture public vaccine hesitancy in real life. After identifying 11 themes from widely diffused information on Twitter, with the help of our model, we validated the predictability of users' vaccine hesitancy changes by the information they perceived from social media. This showed a potential use of our model in practice. Through this article, we established again the power of social media data in supplementing public health surveillance, especially in combating infectious virus like COVID-19.

Limitations and Future Work. We have three main limitations to address in the future. First, in our work, we have primarily focused on Twitter, which potentially induces bias in our data and analysis. Thus, it is important to extend our work to other social media platforms such as Facebook and Instagram, and cross-validate our results. Second, we only analysed users' affective vaccination stances, i.e., positive, negative and neutral, which can only be used as an indicator of users' intention to get vaccinated. It will be interesting to look deeper into users' tweets for a longer time and identify underlying determinants that lead to vaccination acceptance. Third, we used only the top 25% most widely spread tweets as representatives to extract the themes of diffused information partly limited by manual annotation. Some information in certain themes may be missed. As an interesting future work, we will develop effective NLP models to learn different tweet themes automatically.

Ethical Considerations. This work is based completely on public data and does not contain private information of individuals. Our dataset is built in accordance with the FAIR data principles [66] and Twitter Developer Agreement and Policy and related policies. Our release of the dataset is also compliant with GDPR (General Data Protection Regulation). To conclude, we have no ethical violation in the collection and interpretation of data in our study.

Appendix

A Diffused Information Themes and Examples

Table 7. Diffused Information Themes and Examples

Theme	Description	Example (Translated to English)
Positive news	Positive news about vaccines and vaccination	Pfizer/BioNTech's vaccine would be effective against the new British variant of COVID19.
Negative news	Negative news about vaccines and vaccination	Portugal: She dies 2 days after the vaccine (at 41 years old). Her family asks for explanations.
Distrust in government management	Doubt about the trustworthiness of governments or medical institutions, e.g., regarding the daily update of statistics	They have lied to us so much about masks, chloroquine, contagion in children, that it will be difficult to trust them the day they will tell us about a harmless vaccine.
Dissatisfaction with politics/policies	Unsatisfactory views of politics/policies, such as ineffective vaccination programs	I am opposed to mandatory vaccination because all of the world's health organizations say that it is not the right way for a vaccine to spread.
Perception of the pharmaceutical industry	Perception that pharmaceutical manufacturers pursue only economic interests rather than public health interests	Pfizer's CEO sold 60 percent of his shares when the Covid vaccine was announced. When the CEO sells, it stinks.
Conspiracy	Content that describes the event as the secret acts of a powerful, malevolent force	18 months they've been on the vaccine ???? When did they know there would be a Covid 19 "pandemic" ????
Beliefs, attitudes about health and prevention	Personal views on vaccines and the immune system, e.g., homeopathy, natural immunity, alternative therapies	There is no point in a generalised vaccine for a disease whose mortality is close to 0.05%.
Positive personal expression	Personal expression of positive attitude towards vaccines	We have a new weapon against the virus: the vaccine. Hold together, again.
Negative personal expression	Personal expression of negative attitude towards vaccines	Why could actually 1.5 billion Chinese get healthy without vaccination, and with us it only works with vaccination...?
Positive information	Positive expressions about vaccines from healthcare professionals	#COVID19 #vaccination How does an mRNA vaccine work?
Negative information	Negative expressions about vaccines from healthcare professionals	My daughter, a nurse at the AP-HP, on the vaccine "Ah ah ah! They don't even dream about it, they start with the old ones so that we can attribute the side effects to age".

References

[1] Muhammad Abulaish, Ashraf Kamal, and Mohammed J. Zaki. 2020. A survey of figurative language and its computational detection in online social networks. *ACM Transactions on the Web* 14, 1 (2020), Article 3, 52 pages.

[2] Shiva Raj Acharya, Deog Hwan Moon, and Yong Chul Shin. 2021. Assessing attitude toward COVID-19 vaccination in South Korea. *Frontiers in Psychology* 12 (2021), 3222.

[3] Allison E. Aiello, Audrey Renson, and Paul N. Zivich. 2020. Social media- and Internet-based disease surveillance for public health. *Annual Review of Public Health* 41, 1 (2020), 101–118.

[4] Rasmieh Al-Amer, Della Maneze, Bronwyn Everett, Jed Montayre, Amy R. Villarosa, Entisar Dwekat, and Yenna Salamonson. 2021. COVID-19 vaccination intention in the first year of the pandemic: A systematic review. *Clinical Nursing* 31, 1-2 (2021), 62–86.

[5] Majid Alabdulla, Shuja Mohd Reagu, Abdullatif Al-Khal, Marwa Elzain, and Roland M. Jones. 2021. COVID-19 vaccine hesitancy and attitudes in Qatar: A national cross-sectional survey of a migrant-majority population. *Influenza and Other Respiratory Viruses* 15, 3 (2021), 361–370.

[6] Nick Andrews, Julia Stowe, Freja Kirsebom, Samuel Toffa, Ruchira Sachdeva, Charlotte Gower, Mary Ramsay, and Jamie Lopez Berna. 2022. Effectiveness of COVID-19 booster vaccines against COVID-19 related symptoms, hospitalisation and death in England. *Nature Medicine* 28 (2022), 831–837.

[7] Serena Barello, Tiziana Nania, Federica Dellafiore, Guendalina Graffigna, and Rosario Caruso. 2020. Vaccine hesitancy among university students in Italy during the COVID-19 pandemic. *European Journal of Epidemiology* 35, 8 (2020), 781–783.

[8] Aida Bianco, Valentina Mascaro, Rossella Zucco, and Maria Pavia. 2019. Parent perspectives on childhood vaccination: How to deal with vaccine hesitancy and refusal? *Vaccine* 37, 7 (2019), 984–990.

[9] Rhys Biddle, Aditya Joshi, Shaowu Liu, Cécile Paris, and Guandong Xu. 2020. Leveraging sentiment distributions to distinguish figurative from literal health reports on Twitter. In *Proceedings of the 2020 International Conference on World Wide Web (WWW '20)*. ACM, New York, NY, USA, 1217–1227.

[10] Nirbachita Biswas, Toheeb Mustapha, Jagdish Khubchandani, and James H. Price. 2021. The nature and extent of COVID-19 vaccination hesitancy in healthcare workers. *Journal of Community Health* 46, 6 (2021), 1244–1251.

[11] Johnna Blair, Chi-Yang Hsu, Ling Qiu, Shih-Hong Huang, Ting-Hao Kenneth Huang, and Saeed Abdullah. 2021. Using tweets to assess mental well-being of essential workers during the COVID-19 pandemic. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 236, 6 pages.

[12] Jobie Budd, Benjamin S. Miller, Erin M. Manning, Vasileios Lampos, Mengdie Zhuang, Michael Edelstein, Geraint Rees, Vincent C. Emery, Molly M. Stevens, Neil Keegan, et al. 2020. Digital technologies in the public-health response to COVID-19. *Nature Medicine* 26, 8 (2020), 1183–1192.

[13] Emily Chen, Kristina Lerman, and Emilio Ferrara. 2020. Tracking social media discourse about the COVID-19 pandemic: Development of a public coronavirus Twitter data set. *JMIR Public Health and Surveillance* 6, 2 (2020), e19273.

[14] Ninghan Chen, Xihui Chen, Jun Pang, Liyousew G. Borga, Conchita D'Ambrosio, and Claus Vögele. 2022. Measuring COVID-19 vaccine hesitancy: Consistency of social media with surveys. In *Social Informatics*. Lecture Notes in Computer Science, Vol. 13618. Springer, 196–210.

[15] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of the 8th Workshop on Syntax, Semantics, and Structure in Statistical Translation (SSST '14).*, 103–111.

[16] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL '20)*. 8440–8451.

[17] Alessandro Cossard, Gianmarco De Francisci Morales, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2020. Falling into the echo chamber: The Italian vaccination debate on Twitter. In *Proceedings of the 2020 International AAAI Conference on Web and Social Media (ICWSM '20)*, Vol. 14. 130–140.

[18] Juan M. Cotelo, Fermin L. Cruz, Fernando Enríquez, and José A. Troyano. 2016. Tweet categorization by combining content and structural knowledge. *Information Fusion* 31 (2016), 54–64.

[19] Giorgia Della Polla, Francesca Licata, Silvia Angelillo, Concetta Paola Pelullo, and Aida Bianco. 2021. Characteristics of healthcare workers vaccinated against influenza in the era of COVID-19. *Vaccines* 9, 7 (2021), 695.

[20] 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 (NAACL-HLT '19)*. 4171–4186.

[21] Yingtong Dou, Hao Yang, and Xiaolong Deng. 2016. A survey of collaborative filtering algorithms for social recommender systems. In *Proceedings of the 12th International Conference on Semantics, Knowledge, and Grids (SKG '16)*. IEEE, 40–46.

[22] Eve Dubé, Caroline Laberge, Maryse Guay, Paul Bramadat, Réal Roy, and Julie A. Bettinger. 2013. Vaccine hesitancy: An overview. *Human Vaccines & Immunotherapeutics* 9, 8 (2013), 1763–1773.

[23] Francesco Ducci, Mathias Kraus, and Stefan Feuerriegel. 2020. Cascade-LSTM: A tree-structured neural classifier for detecting misinformation cascades. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '20)*. ACM, New York, NY, USA, 2666–2676.

[24] Adam G. Dunn, Julie Leask, Xujuan Zhou, Kenneth D. Mandl, and Enrico Coiera. 2015. Associations between exposure to and expression of negative opinions about human papillomavirus vaccines on social media: An observational study. *Journal of Medical Internet Research* 17, 6 (2015), e144.

[25] Javid Ebrahimi, Dejing Dou, and Daniel Lowd. 2016. Weakly supervised tweet stance classification by relational bootstrapping. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP '16)*. 1012–1017.

[26] David Ediger, Karl Jiang, Jason Riedy, David A. Bader, Courtney Corley, Rob Farber, and William N. Reynolds. 2010. Massive social network analysis: Mining Twitter for social good. In *Proceedings of the 2010 International Conference on Parallel Processing (ICPP '10)*. IEEE, 583–593.

[27] Joseph L. Fleiss, Bruce Levin, and Myunghee Cho Paik. 2013. *Statistical Methods for Rates and Proportions*. John Wiley & Sons.

[28] Sefater Gbashi, Oluwafemi Ayodeji Adebo, Wesley Doorsamy, and Patrick Berka Njobeh. 2021. Systematic delineation of media polarity on COVID-19 vaccines in Africa: Computational linguistic modeling study. *JMIR Medical Informatics* 9, 3 (2021), e22916.

[29] Victor Grech, Charmaine Gauci, and Steve Agius. 2020. Vaccine hesitancy among Maltese healthcare workers toward influenza and novel COVID-19 vaccination. *Early Human Development* 1 (2020), 105213.

[30] Keith Gunaratne, Eric A. Coomes, and Hourmazd Haghbayan. 2019. Temporal trends in anti-vaccine discourse on Twitter. *Vaccine* 37, 35 (2019), 4867–4871.

[31] Alan G. Hawkes. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58, 1 (1971), 83–90.

[32] Maria Renee Jimenez-Sotomayor, Carolina Gomez-Moreno, and Enrique Soto-Perez-de Celis. 2020. Coronavirus, ageism, and Twitter: An evaluation of tweets about older adults and COVID-19. *Journal of the American Geriatrics Society* 68, 8 (2020), 1661–1665.

[33] Neil F. Johnson, Nicolas Velásquez, Nicholas Johnson Restrepo, Rhys Leahy, Nicholas Gabriel, Sara El Oud, Minzhang Zheng, Pedro Manrique, Stefan Wuchty, and Yonatan Lupu. 2020. The online competition between pro- and anti-vaccination views. *Nature* 582, 7811 (2020), 230–233.

[34] Jagdish Khubchandani, Sushil Sharma, James H. Price, Michael J. Wiblishauser, Manoj Sharma, and Fern J. Webb. 2021. COVID-19 vaccination hesitancy in the United States: A rapid national assessment. *Journal of Community Health* 46, 2 (2021), 270–277.

[35] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc.eedings of the 2015 International Conference on Learning Representations (ICLR '15)*.

[36] Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 2017 International Conference on Learning Representations (ICLR '17)*.

[37] Klaus Krippendorff. 1970. Estimating the reliability, systematic error and random error of interval data. *Educational and Psychological Measurement* 30, 1 (1970), 61–70.

[38] Jeffrey V. Lazarus, Scott C. Ratzan, Adam Palayew, Lawrence O. Gostin, Heidi J. Larson, Kenneth Rabin, Spencer Kimball, and Ayman El-Mohandes. 2021. A global survey of potential acceptance of a COVID-19 vaccine. *Nature Medicine* 27, 2 (2021), 225–228.

[39] Qian Li, Hao Peng, Jianxin Li, Congying Xia, Renyu Yang, Lichao Sun, Philip S. Yu, and Lifang He. 2022. A survey on text classification: From traditional to deep learning. *ACM Transactions on Intelligent Systems and Technologies* 13, 2 (2022), 41.

[40] Xinlong Li, Xingyu Fu, Guangluan Xu, Yang Yang, Jiuniu Wang, Li Jin, Qing Liu, and Tianyuan Xiang. 2020. Enhancing BERT representation with context-aware embedding for aspect-based sentiment analysis. *IEEE Access* 8 (2020), 46868–46876.

[41] 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. *CoRR abs/1907.11692* (2019).

[42] Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of the 2013 International Conference on Learning Representations (ICLR '13)*.

[43] Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, and Jianfeng Gao. 2021. Deep learning–based text classification: A comprehensive review. *ACM Computing Surveys* 54, 3 (2021), Article 62, 40 pages.

[44] Alexandros Mittos, Savvas Zannettou, Jeremy Blackburn, and Emiliano De Cristofaro. 2020. Analyzing genetic testing discourse on the web through the lens of Twitter, Reddit, and 4chan. *ACM Transactions on the Web* 14, 4 (2020), Article 17, 38 pages.

[45] Gianmarco De Francisci Morales, Corrado Monti, and Michele Starnini. 2021. No echo in the chambers of political interactions on Reddit. *Scientific Reports* 11, 1 (2021), 1–12.
[46] Martin Müller, Marcel Salathé, and Per Egil Kummervold. 2023. COVID-Twitter-BERT: A natural language processing model to analyse COVID-19 content on Twitter. *Frontiers in Artificial Intelligence* 6 (2023), 1023281.
[47] Douglas J. Opel, James A. Taylor, Rita Mangione-Smith, Cam Solomon, Chuan Zhao, Sheryl Catz, and Diane Martin. 2011. Validity and reliability of a survey to identify vaccine-hesitant parents. *Vaccine* 29, 38 (2011), 6598–6605.
[48] World Health Organization. 2014. Report of the SAGE Working Group on Vaccine Hesitancy. Retrieved December 5, 2024 from http://www.who.int/immunization/sage/meetings/2014/october/SAGE
[49] C. Oswald, Sona Elza Simon, and Arnab Bhattacharya. 2022. SpotSpam: Intention analysis-driven SMS spam detection using BERT embeddings. *ACM Transactions on the Web* 16, 3 (2022), Article 14, 27 pages.
[50] Xiaozhi Ou and Hongling Li. 2020. YNU_OXZ @ HaSpeeDe 2 and AMI : XLM-RoBERTa with ordered neurons LSTM for classification task at EVALITA 2020. In *Proceedings of the 2020 Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA '20)*, Vol. 2765.
[51] Shashank Pandit, Duen Horng Chau, Samuel Wang, and Christos Faloutsos. 2007. NetProbe: A fast and scalable system for fraud detection in online auction networks. In *Proceedings of the 2007 International Conference on World Wide Web (WWW '07)*. ACM, New York, NY, USA, 201–210.
[52] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
[53] Soham Poddar, Mainack Mondal, Janardan Misra, Niloy Ganguly, and Saptarshi Ghosh. 2022. Winds of change: Impact of COVID-19 on vaccine-related opinions of Twitter users. In *Proceedings of the 16th International AAAI Conference on Web and Social Media (ICWSM '22)*. 782–793.
[54] Soham Poddar, Azlaan Mustafa Samad, Rajdeep Mukherjee, Niloy Ganguly, and Saptarshi Ghosh. 2022. 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 (SIGIR '22)*. ACM, New York, NY, USA, 3154–3164.
[55] Annamaria Porreca, Francesca Scozzari, and Marta Di Nicola. 2020. Using text mining and sentiment analysis to analyse YouTube Italian videos concerning vaccination. *BMC Public Health* 20, 1 (2020), 259.
[56] Neha Puri, Eric A. Coomes, Hourmazd Haghbayan, and Keith Gunaratne. 2020. Social media and vaccine hesitancy: New updates for the era of COVID-19 and globalized infectious diseases. *Human Vaccines & Immunotherapeutics* 16, 11 (2020), 2586–2593.
[57] Viju Raghupathi, Jie Ren, and Wullianallur Raghupathi. 2020. Studying public perception about vaccination: A sentiment analysis of tweets. *International Journal of Environmental Research and Public Health* 17, 10 (2020), 3464.
[58] Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 Task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval '17)*. 502–518.
[59] Koustav Rudra, Niloy Ganguly, Pawan Goyal, and Saptarshi Ghosh. 2018. Extracting and summarizing situational information from the Twitter social media during disasters. *ACM Transactions on the Web* 12, 3 (2018), Article 17, 35 pages.
[60] Rahul Shekhar, Abu Baker Sheikh, Shubhra Upadhyay, Mriganka Singh, Saket Kottewar, Hamza Mir, Eileen Barrett, and Suman Pal. 2021. COVID-19 vaccine acceptance among health care workers in the United States. *Vaccines* 9, 2 (2021), 119.
[61] Susan M. Sherman, Louise E. Smith, Julius Sim, Richard Amlôt, Megan Cutts, Hannah Dasch, G. James Rubin, and Nick Sevdalis. 2021. COVID-19 vaccination intention in the UK: Results from the COVID-19 vaccination acceptability study (CoVAccS), a nationally representative cross-sectional survey. *Human Vaccines & Immunotherapeutics* 17, 6 (2021), 1612–1621.
[62] Jay J. Van Bavel, Katherine Baicker, Paulo S. Boggio, Valerio Capraro, Aleksandra Cichocka, Mina Cikara, Molly J. Crockett, Alia J. Crum, Karen M. Douglas, James N. Druckman, et al. 2020. Using social and behavioural science to support COVID-19 pandemic response. *Nature Human Behaviour* 4, 5 (2020), 460–471.
[63] Livia van Vliet, Petter Törnberg, and Justus Uitermark. 2020. The Twitter parliamentarian database: Analyzing Twitter politics across 26 countries. *PLoS One* 15, 9 (2020), e0237073.
[64] Petar Velickovic, Arantxa Casanova Guillem Cucurull, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *Proceedings of the 2018 International Conference on Learning Representations (ICLR '18)*.
[65] Lei Wang, Jianwei Niu, and Shui Yu. 2020. SentiDiff: Combining textual information and sentiment diffusion patterns for Twitter sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering* 32, 10 (2020), 2026–2039.
[66] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, et al. 2016. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data* 3, 1 (2016), 1–9.

[67] Kumanan Wilson, Katherine Atkinson, and Shelley Deeks. 2014. Opportunities for utilizing new technologies to increase vaccine confidence. *Expert Review of Vaccines* 13, 8 (2014), 969–977.

[68] Zhiping Xiao, Weiping Song, Haoyan Xu, Zhicheng Ren, and Yizhou Sun. 2020. TIMME: Twitter ideology-detection via multi-task multi-relational embedding. In *Proceedings of the 2020 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '20)*. ACM, New York, NY, USA, 2258–2268.

[69] Senqi Zhang, Li Sun, Daiwei Zhang, Pin Li, Yue Liu, Ajay Anand, Zidian Xie, and Dongmei Li. 2022. The COVID-19 pandemic and mental health concerns on Twitter in the United States. *Health Data Science* 2022 (2022), 9758408.

[70] Lingyun Zhao, Lin Li, Xinhao Zheng, and Jianwei Zhang. 2021. A BERT based sentiment analysis and key entity detection approach for online financial texts. In *Proceedings of the 2021 IEEE International Conference on Computer Supported Cooperative Work in Design (CSCWD '21)*. IEEE, 1233–1238.

[71] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. Beyond homophily in graph neural networks: Current limitations and effective designs. In *Proceedings of the 2020 Annual Conference on Neural Information Processing Systems (NeurIPS '20)*, Vol. 33. 7793–7804.

Received 27 October 2022; revised 2 January 2024; accepted 9 January 2024