

From Hesitancy Framings to Vaccine Hesitancy Profiles: A Journey of Stance, Ontological Commitments and Moral Foundations

Maxwell A. Weinzierl, Sanda M. Harabagiu

Human Language Technology Research Institute, The University of Texas at Dallas
maxwell.weinzierl@utdallas.edu, sanda@utdallas.edu

Abstract

While billions of COVID-19 vaccines have been administered, too many people remain hesitant. Twitter, with its substantial reach and daily exposure, is an excellent resource for examining how people frame their vaccine hesitancy and to uncover vaccine hesitancy profiles. In this paper we expose our processing journey from identifying Vaccine Hesitancy Framings in a collection of 9,133,471 original tweets discussing the COVID-19 vaccines, establishing their ontological commitments, annotating the Moral Foundations they imply to the automatic recognition of the *stance* of the tweet authors toward any of the COVAXFRAMES that we have identified. When we found that 805,336 Twitter users had a stance towards some COVAXFRAMES in either the 9,133,471 original tweets or their 17,346,664 retweets, we were able to derive nine different Vaccine Hesitancy Profiles of these users and to interpret these profiles based on the ontological commitments of the frames they evoked in their tweets and on value of their stance towards the evoked frames.

1 Introduction

Social media microblogging platforms, specifically Twitter, have become highly influential and relevant to shaping attitudes towards vaccination. With 206 million daily active users as of 2021, Twitter has substantial reach and daily exposure, being the most popular social network for news consumption (Auxier and Anderson 2021). Since Twitter allows people to express their beliefs about vaccines and their hesitancy to vaccinate, their trust or mistrust in vaccines as well as their stance on vaccination mandates, it is an excellent resource for investigating how vaccine hesitancy is framed. While vaccine hesitancy is mostly believed to be fueled by misinformation (Garett and Young 2021), in our study of the Twitter discourse focusing on the COVID-19 vaccines, we have found that misinformation is not the only explanation for vaccine hesitancy. Vaccine hesitancy is also driven by the erosion of trust in vaccines, even when no misinformation is referred, or by lack of health literacy and even by the interaction between civil rights and vaccination mandates.

Social science stipulates, according to Chong and Druckman (2007); Entman (2004), that discourse almost inescapably involves framing – a strategy of highlighting cer-

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Vaccine Hesitancy Framing 1: *Governments hide vaccine safety information.*

STANCE: Accept

Tweet: Why would the government block the Office for National Statistics from publishing side effects and deaths after taking covid vaccine? What are they hiding?

STANCE: Reject

Tweet: @USER I am not talking about protection, but prima facie what i have heard from many in government and private hospital doctors is, 2 doses are very much effective in stopping mortality or high organ damage. Vaccine can't stop covid, we are talking more about mortality reduction

Vaccine Hesitancy Framing 2: *It is not known if the COVID-19 vaccines will provide protection against future variants.*

STANCE: Accept

Tweet: It's not a vaccine. The COVID-19 mRNA vaccine does not provide immunity to Covid or it's variants so you can still catch Covid and transmit it to others making you asymptomatic. You will likely need a booster shot every 6 months, so get ready to roll up that sleeve every six mon

STANCE: Reject

Tweet: @USER @USER Reasons to get the vaccine: 1. It can protect you in case your immune system can't fight the virus. 2. It can help protect your community and vulnerable people. 3. It can help to prevent the spread of variants. 4. We don't know the long-term effects of COVID-19 yet #GetTheShot

Table 1: Examples of Vaccine Hesitancy Framings and tweets evoking them, while the tweet authors accept or reject the framing.

tain issues to promote a certain interpretation or attitude. **Vaccine Hesitancy Framings** (VHFs), highlight issues regarding confidence in the safety of vaccines by using *specific* misinformation, as exemplified by the first VHF listed in Table 1, or erode the trust in vaccines by demotivating people from vaccination, as exemplified in the second VHF listed in Table 1. While VHFs are not directly expressed in tweets discussing vaccines, they can be inferred from the discourse spanning these tweets. Furthermore, tweet authors do not only evoke these VHFs, but they also express their *stance* towards them. Table 1 illustrates a tweet whose author agrees with the first VHF shown in the Table, thus *adopting* the framing and another tweet whose author disagrees with the

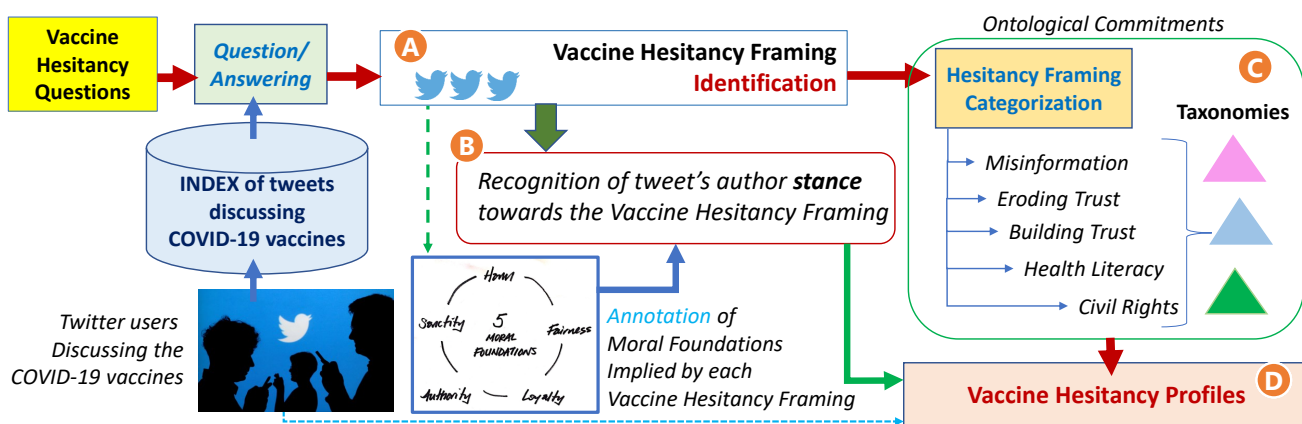


Figure 1: Recognizing Vaccine Hesitancy Profiles (VHPs) by taking into account (1) the identification of Vaccine Hesitancy Framings (VHFs); (2) the *stance* of tweet authors towards the framings; and (3) the ontological commitments of the identified framings.

same framing, thus *rejecting* it. Tweets with different stance for the second VHF are listed as well in the table.

Vaccine hesitancy is a continuum between those that accept vaccines with no doubts, to those that absolutely refuse vaccines, with vaccine hesitant individuals in heterogeneous groups between these two extremes, according to (Larson et al. 2015). To uncover the **Vaccine Hesitancy Profiles** (VHPs), which was the main objective of our study, we focused on the Twitter discourse regarding the COVID-19 vaccines. Intuitively, all tweet authors sharing the same VHP are expected to also share some commonalities through the way they frame their hesitancy towards COVID-19 vaccines. Therefore, as illustrated in Figure 1, our methodology considered (A) the identification of VHFs from an index of 5,865,046 tweets discussing COVID-19 vaccines and (B) the recognition of the stance towards the VHFs evoked by 805,336 users. Furthermore, we relied on previous work in social psychology considered the Moral Foundations Theory (MFT) (Haidt and Graham 2007; Haidt and Joseph 2004) as a theoretical framework for analyzing moral framing, using the same five key values of human morality, emerging from evolutionary, social, and cultural origins. Each VHF was annotated with the Moral Foundation (MF) they imply. For example, for the first VHF illustrated in Table 1, the annotated MFs are: *Harm*, *Betrayal* and *Authority* while for the second VHF illustrated in the same table, the only annotated MFs is *Betrayal*. The annotated MFs informed the recognition of the tweet author stance towards the VHFs that are evoked in their tweets. These annotations have contributed to the automatic recognition of stance.

Figure 1 shows that our journey, from identifying VHFs (revealed as answers to questions about vaccine hesitancy), to finally recognizing the VHPs, involved also (C) the derivation of the ontological commitments of the VHFs, by categorizing them into misinformation, framings eroding trust in vaccines, framing building trust in vaccines, framings showcasing the health literacy of the tweet author or framings in which civil rights are brought up. This categorization enabled us to iden-

tify the common themes and concerns of the VHFs in each category, and to organize them into taxonomies, completing the ontological organization of the VHFs. These taxonomies, together with the stance information, allowed us to create a representation of each Twitter user that we recognized to have a stance towards the VHFs, and (D) to reveal the VHPs. The ontological information along with the stance information enables us to interpret the VHPs. For example, the tweet authors mostly evoking the first VHF illustrated in Table 1 belong to the profile of UNDECIDED, whereas the authors accepting the second VHF from the same table are DEMOTIVATED in their hesitancy, whereas those rejecting it belong to the profile of those that are MOTIVATED to vaccinate.

2 Identification of Vaccine Hesitancy Framings through Question/Answering

Current NLP methods (Luo, Zimet, and Shah 2019; Du et al. 2017) used for recognizing vaccine hesitancy assume that a neutral sentiment detected in a tweet is equivalent to hesitancy, while a positive sentiment is interpreted as acceptance of vaccination and a negative sentiment as refusal. However, as defined in (Larson et al. 2015), vaccine hesitancy refers to the delay in acceptance or refusal of vaccines despite availability of vaccination services. Moreover, vaccine hesitancy is informed by factors such as complacency, convenience, and confidence, which are framed in complex ways in language (Macdonald 2015). For example, when misinformation is used in framing vaccine confidence, it typically results in vaccine hesitancy. Similarly, when civil rights are highlighted in a particular framing, it promotes vaccine refusal, while when trust in vaccines is increased, it leads to vaccine acceptance, and eventual uptake. Therefore, developing novel NLP techniques capable of discovering framings of vaccine hesitancy in social media discourse is essential, especially for uncovering the various vaccine hesitancy profiles of users.

Recent work in NLP concerning automatic recognition of framings targeted the study of political bias and polarization in social and news media (Field et al. 2018; Roy and

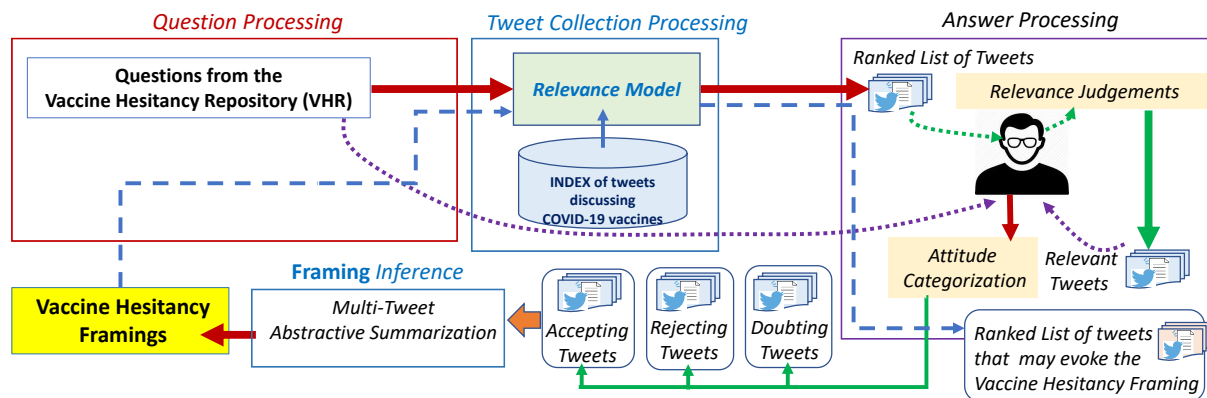


Figure 2: A Question Answering Framework for (a) identifying Vaccine Hesitancy Framings as well as (b) tweets that potentially evoke a Vaccine Hesitancy Framing.

Goldwasser 2020), mainly addressing the recognition of 15 cross-cutting dimensions of political framings e.g., economic dimensions, fairness and equality or policy prescription and evaluation. Although recent Twitter content analysis (Rao et al. 2020) revealed that there is significance correlation between polarized attitudes towards vaccines and political dimensions, to our knowledge, no NLP methods have yet been developed to identify vaccine hesitancy framings, although vaccine hesitancy is often discussed in social/news media.

In a novel approach that uses Question Answering, illustrated in Figure 2, we have found that VHF’s focusing on the COVID-19 vaccines can be successfully identified as answers to questions from the Vaccine Confidence Repository (VCR) (Rossen et al. 2019), which is a set of 18 questions targeting hesitancy, informed by the anti-vaccine content analysis reported in (Kata 2010). The same questions were used in the study reported in (Rossen et al. 2019) to discern hesitancy profiles from the answers returned on survey links available from Facebook pages and parenting forums. The questions from VCR targeted the Human Papillomavirus (HPV) vaccine. We adapted the questions by asking about the COVID-19 vaccine. These questions covered five different themes stipulating that: (T1) vaccines are unsafe and unnatural; (T2) vaccination is ineffective; (T3) redundant vaccinations; (T4) people should be free to decide if they want to vaccinate; and (T5) vaccination is a conspiracy. Instead of soliciting answers from Twitter users, we decided to (a) automatically find tweets that answer the same questions and (b) infer the framings evoked by the answers.

As shown in Figure 2 each question was processed, transforming it into a query that can be handled by a relevance model implementing the BM25 vector relevance model (Beaulieu et al. 1997). In addition, as in (Weinzierl and Harabagiu 2020b), we considered the BERT-RERANK (Nogueira and Cho 2020) scoring function to re-rank the tweets provided by the BM25 relevance model. But, finding the answers to the questions required an index of tweets discussing the COVID-19 vaccines.

In order to obtain a collection of tweets discussing the COVID-19 vaccine, we started by obtaining approval from

the Institutional Review Board at the University of Texas at Dallas. IRB-21-515 stipulated that our research met the criteria for exemption #8(iii) of the Chapter 45 of Federal Regulations Part 46.101.(b). Afterwards, tweets discussing the COVID-19 vaccines were obtained by using the query “(covid OR coronavirus) vaccine lang:en”. A collection of 9,133,471 original tweets and 17,346,664 retweets was obtained from the Twitter streaming API. These tweets were authored between December 18th, 2019, and July 21st, 2021. To detect duplications in the original tweets, we performed Locality Sensitive Hashing (LSH) (Das et al. 2007) with term trigrams, 100 permutations, and a Jaccard threshold of 50%, obtaining 5,865,046 unique original tweets discussing COVID-19 vaccines. We first build an index for 5,865,046 unique original tweets using Lucene (Foundation 1999), informing the relevance model.

Relevance judgements were produced on the 300 best ranked tweets from the first index, language experts categorizing afterwards the attitude against the predication of the question of each relevant tweet. Less than 60% of the tweets were judged relevant by two language experts from the University of Texas at Dallas. As shown in Figure 2, relevant tweets were categorized as accepting the predication, rejecting it and doubting it, with a Cohen Kappa score of 0.81, which indicates strong agreement between annotators (0.8-0.9) (McHugh 2012). From tweets sharing the same attitude towards the question predication, the Pyramid method (Nenkova and Passonneau 2004) was used to infer a query-focused multi-tweet summary, which was considered a VHF. In this way, we identified a set of 113 VHF’s targeting the COVID-19 vaccine, which we assembled in COVAXFRAMES. Then, each VHF from COVAXFRAMES was reused as a question and processed by the relevance model against a new index containing all the 9,133,471 original tweets and 17,346,664 retweets. From the best-ranked 400,000 tweets retrieved for each VHF, only the tweets that had a relevance score above a threshold $T_r = 2.0$, selected from initial experiments, were considered to potentially evoke the VHF, which resulted in 19,233,144 tweets potentially evoking a VHF from COVAXFRAMES.

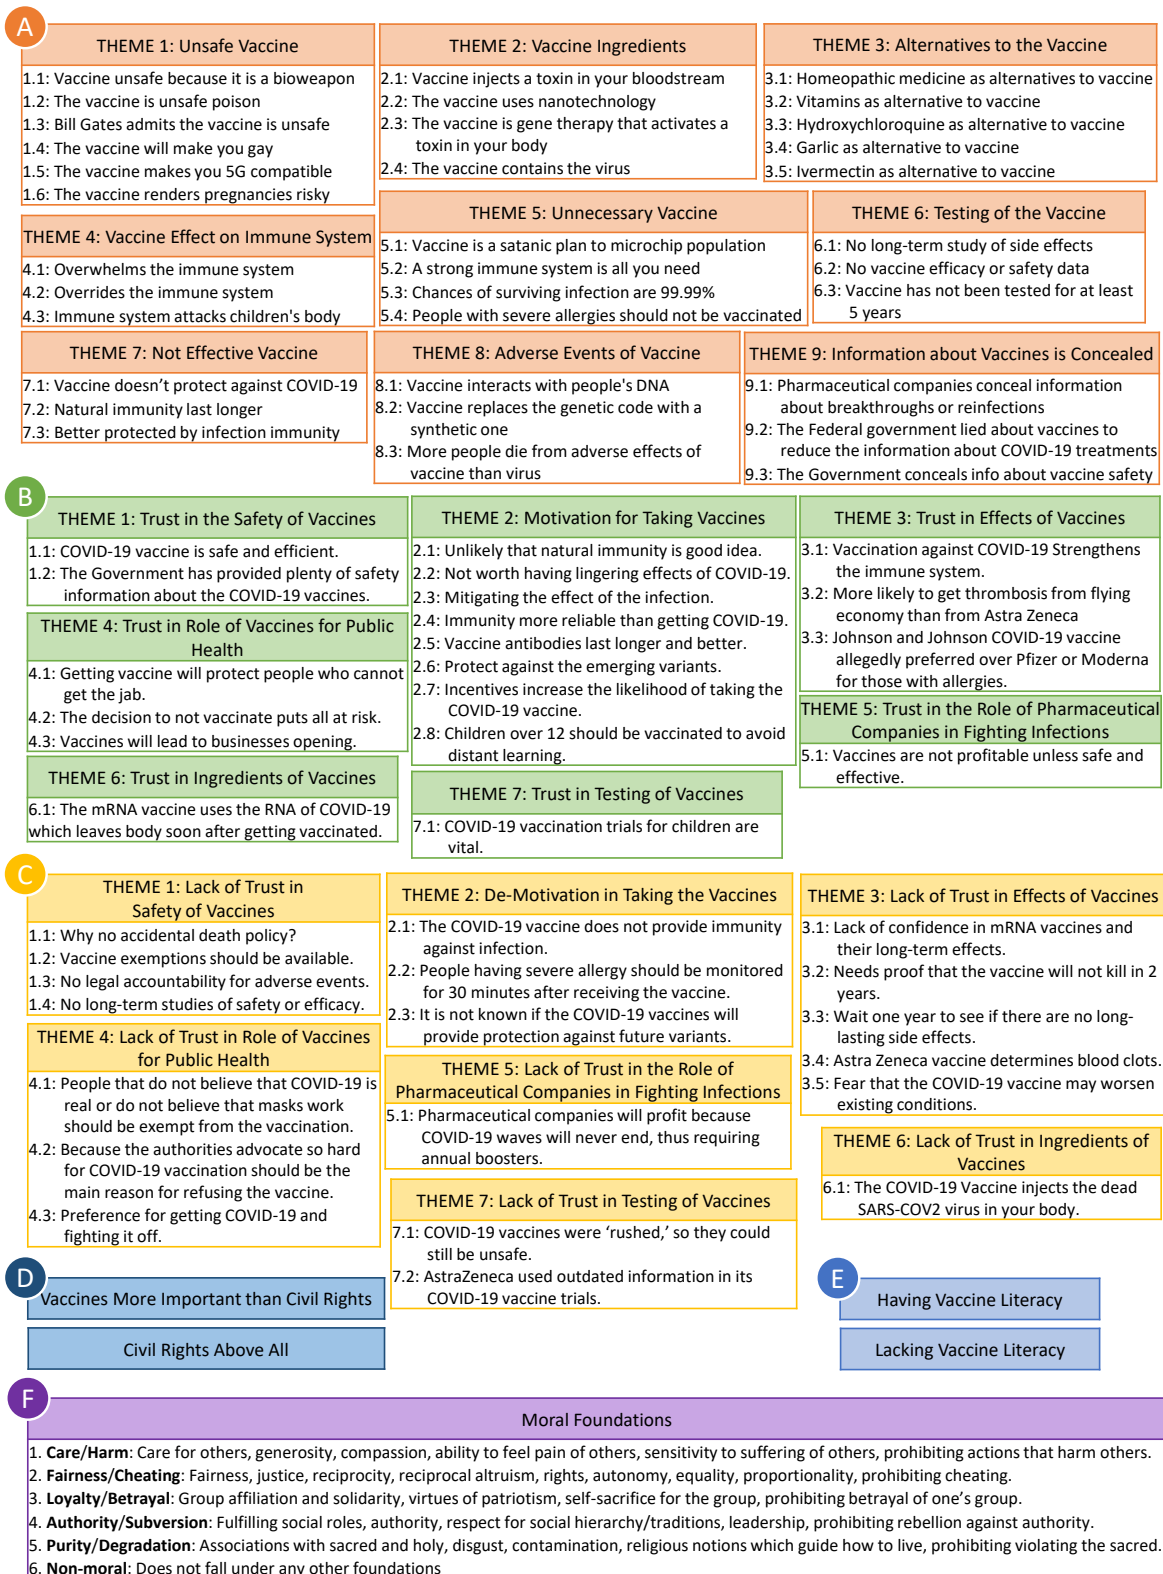


Figure 3: Taxonomy of Themes and Concerns of (A) Misinformation, (B) Building Trust, (C) Eroding Trust, (D) Civil Rights, (E) Vaccine Literacy, and (F) Moral Foundations defined by the Moral Foundation Theory, annotated when implied in Vaccine Hesitancy Framings from COVAXFRAMES.

3 Ontological Commitments of Vaccine Hesitancy Framings

The VHF’s were inspected first to distinguish which ones contain misinformation and which do not. Of the 113 VHF’s from COVAXFRAMES, 38 VHF’s were categorized as Misinformation framings. The remaining VHF’s were categorized in the following way: 52 VHF’s addressed issues of trust in vaccines, but without using misinformation; 32 VHF’s addressed issues of Civil Rights. In addition, we inspected all VHF’s to decide if they exhibit health literacy or lack of it: 28 VHF’s were categorized as Literacy framings. For the VHF’s addressing issues of trust in vaccines, we found that 27 VHF’s are Eroding Trust in vaccines while 25 VHF’s are Building Trust framings. This categorization allowed us to organize a Misinformation Taxonomy, which encodes the common themes and concerns of the Misinformation framings. Similarly, we have organized a taxonomy of Building Trust and a taxonomy of Eroding Trust. These taxonomies are illustrated in Figure 3. For the VHF’s addressing Civil Rights issues, we have created only two themes, as we have done for the VHF’s exhibiting Vaccine Literacy or lack thereof, in their respective ontologies.

In addition, all VHF’s from COVAXFRAMES were annotated with as many Moral Foundations as they implied. Figure 3 (F) lists the definitions of each MF from the Moral Foundation Theory (Haidt and Joseph 2004) that we have used. A computational linguist and an expert in public health have independently assigned MFs to all the VHF’s, and the inter-judge agreement was a Cohen’s Kappa score of 0.85, where disagreements were resolved between annotators. The most common Moral Foundations were *Harm* and *Subversion*, occurring in 45 VHF’s and 44 VHF’s respectively out of the 113 VHF’s, while the least common Moral Foundations were *Cheating* and *Loyalty*, occurring in 8 VHF’s and 9 VHF’s respectively. These ontological commitments that were organized in the COVAXFRAMES contributed to the discovery of hesitancy profiles, along with the information about the stance towards any of the VHF’s evoked in tweets.

4 Stance Recognition

The recognition of the stance of a tweet author towards any of the COVAXFRAMES is made possible by the STANCEID-MORALITY system, illustrated in Figure 4. Given any VHF f from COVAXFRAMES and any tweet t that may evoke a VHF, produced by the Q/A framework illustrated in Figure 2, we hypothesize that if an *Accept* or *Reject* stance towards VHF f is recognized automatically, then the tweet t is recognized as evoking VHF f , otherwise tweet t does not evoke the VHF f . We also believed that the stance of the author of tweet t is revealed not only by the interactions between lexical, semantic and emotion information expressed in the tweet t , but also by their interactions with the Moral Foundations (MFs) implied by the VHF f . Therefore, we designed for the STANCEID-MORALITY system a novel neural architecture that combines the advantages of the contextual embeddings learned by COVID-Twitter-BERT-v2 (Müller, Salathé, and Kummervold 2020) with Graph Attention Networks (GATs) (Veličković et al. 2018), where lexical, emotion, and seman-

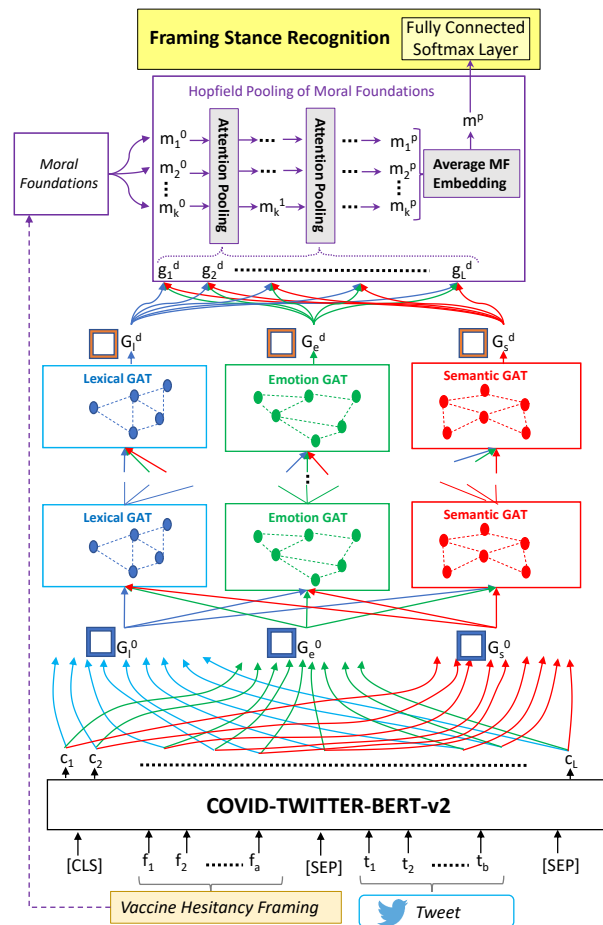


Figure 4: Neural architecture of STANCEID-MORALITY.

tic information can be processed (Weinzierl, Hopfer, and Harabagiu 2021), and a special case of Continuous Hopfield Networks (Ramsauer et al. 2020), namely Hopfield Pooling, where the MFs can be processed as well.

The processing of a VHF f from COVAXFRAMES and the tweet t that potentially evokes f starts in the STANCEID-MORALITY system with joint word-piece tokenization (Devlin et al. 2019), producing the sequence of word-piece tokens $[CLS], f_1, f_2, \dots, f_a, [SEP], t_1, t_2, \dots, t_b, [SEP]$. This sequence of tokens is provided to COVID-Twitter-BERT-v2 for generating the corresponding contextualized embeddings.

COVID-Twitter-BERT-v2 is a language model which was pre-trained on 97M COVID-19 tweets, providing domain-specific language modeling for tasks concerning COVID-19. All contextualized embeddings c_1, c_2, \dots, c_L are organized in a matrix $G \in \mathbb{R}^{L \times 1024}$, which is provided as input to either the stacked lexical GATs, the stacked semantic GATs or the stacked emotion GATs. Each of these GATs operate on graphs in which the nodes are words from f or t . The lexical graph relies on dependency parse edges between words, the emotion edges connect words which share emotion tags from SenticNet 5 (Cambria et al. 2018), and the semantic edges are between pairs of words identified as semantically similar

in SenticNet 5.

Each GAT operates on one of these graphs, refining the representations of each word from each graph through self-attention between adjacent connections. There are d layers of GATs in STANCEID-MORALITY. A GAT at layer $n \in \{1, \dots, d\}$ computes a hidden representation for every contextual embedding $g_i^{n-1} \in G^{n-1}$ as $h_i^n = W^n g_i^{n-1}$, where W^n is a learned weight matrix. This hidden representation is required for computing the self-attention weights of each GAT:

$$\alpha_{i,j}^n = \frac{\exp(\text{LeakyReLU}((a^n)^T [h_i^n, h_j^n]))}{\sum_{k \in \text{adj}(i)} \exp(\text{LeakyReLU}((a^n)^T [h_i^n, h_k^n]))} \quad (1)$$

where a^n is a learned weight vector of size $2F$, $[\dots, \dots]$ represents concatenation; $\text{LeakyReLU}(x) = \max(0.2x, x)$; and $\text{adj}(\dots)$ produces the list of adjacent nodes for a given node from the Lexical, Emotion, or Semantic graphs. The attention weights $\alpha_{i,j}^n$ determine the output of each GAT at layer n :

$$g_i^n = \sigma \left(\sum_{j \in \text{adj}(i)} \alpha_{i,j}^n h_j^n \right) \quad (2)$$

where σ is an exponential linear unit (ELU) nonlinearity (Clevert, Unterthiner, and Hochreiter 2016).

Each of the d layers of GATs have a hidden size F , producing a graph representation G_l^n , G_e^n , and G_s^n respectively of size $L \times F$. The GAT hidden size F and the number of layers d are selected from experiments on the development collection, outlined in Section 7. These Lexical, Emotion, and Semantic Graph representations are concatenated together to form $G^n = [G_l^n, G_e^n, G_s^n]$, with $G^n \in \mathbb{R}^{L \times 3F}$, which is provided as input to all three Lexical, Emotion, and Semantic GATs for the next layer, producing G^{n+1} . This allows each Lexical, Emotion, and Semantic GAT to consider previous Lexical, Emotion, and Semantic Graph representations jointly, learning graph node embeddings which consider interactions between different graphs. The output of the final GAT layers $G^d = [G_l^d, G_e^d, G_s^d]$ is provided to the Hopfield Pooling of Moral Foundation (HP-MF) module.

Each of the 10 MFs m_i are each assigned a unique Moral Foundation Embedding (MFE) $m_i^0 \in \mathbb{R}^{3F}$, initialized randomly and learned throughout the training process of the STANCEID-MORALITY system. The MFEs of the k MFs annotated for the VHF f , $m_1^0, m_2^0, \dots, m_k^0$, are used as initial query embeddings for performing independent Hopfield Pooling (Weinzierl and Harabagiu 2020a) on G^d . Hopfield pooling performs attention pooling p times, where each iteration j refines the query MFE m_i^j from performing attention pooling on the outputs of the final GAT layers G^d utilizing the previous MFE m_i^{j-1} as the query. For each of the k MFs judged within the VHF f , we perform attention pooling at each step j , from 1 to p , independently for each MF m_i . Attention weights are computed using m_i^{j-1} as the query against the final lexical, emotion, and semantic word embeddings $G^d = \{g_1^d, g_2^d, \dots, g_L^d\}$:

$$\beta_x^{i,j} = \frac{\exp(g_x^d \cdot m_i^{j-1})}{\sum_{y=1}^L \exp(g_y^d \cdot m_i^{j-1})} \quad (3)$$

Where \cdot represents the dot product. These attention weights $\beta_x^{i,j}$, which range from 0 to 1, represent how closely the MFE m_i^{j-1} aligns with each of the concatenated lexical, emotion, and semantic word embeddings $g_x^d \in G^d$. The updated MFE m_i^j is then computed as a weighted sum, using $\beta_x^{i,j}$ as the weights, over the lexical, emotion, and semantic word embeddings $g_x^d \in G^d$:

$$m_i^j = \sum_{x=1}^L \beta_x^{i,j} g_x^d \quad (4)$$

Hopfield pooling is therefore performed for each of the k initial MFEs $m_1^0, m_2^0, \dots, m_k^0$ found in VHF f by attention pooling p times over G^d to iteratively construct $m_1^p, m_2^p, \dots, m_k^p$. These final k MFEs are summarized into a single fixed-length representation by taking the average MFE as $z = \frac{1}{k} \sum_{i=1}^k m_i^p$:

This embedding z is provided to the stance recognition layer, which employs a fully connected layer with a softmax activation function to produce final probabilities $P(\text{Accept} | f, t)$, $P(\text{Reject} | f, t)$, and $P(\text{No_Stance} \vee \neg \text{Evoke} | f, t)$, where we merge the probabilities $P(\text{No_Stance} | f, t)$ and $P(\neg \text{Evoke} | f, t)$ into a single probability output, as a tweet with *No Stance* towards VHF f and a tweet which does not evoke VHF f are both ignored when we perform vaccine hesitancy profiling. The STANCEID-MORALITY system is trained end-to-end on the cross-entropy loss function:

$$\mathcal{L} = - \sum_{(s,f,t) \in D} \log P(s | f, t; \theta) \quad (5)$$

where $s \in \{\text{Accept}, \text{Reject}, \text{No_Stance} \vee \neg \text{Evoke}\}$, D is a set of all training examples of labeled [tweet, VHF] pairs, and θ is a set of all trainable parameters from STANCEID-MORALITY. These parameters are optimized with ADAM (Kingma and Ba 2015), a variant of gradient descent, to minimize \mathcal{L} .

Because the STANCEID-MORALITY system implements a supervised method for stance recognition, we relied on a training dataset, a development dataset as well as a testing dataset that allowed us to perform experiments and collect results. Generating these datasets was made possible by the annotations performed on the tweets deemed relevant for each of the VHF questions used in the QA framework presented in Section 2. Researchers from the University of Texas at Dallas judged (a) whether a tweet evokes any of the VHFs from COVAXFRAMES; and (b) if so, they annotated the *stance* of the tweet’s author towards the VHF. 14,180 tweets were judged, with 11,616 tweets evoking one or more VHFs from COVAXFRAMES. They were organized in [tweet, VHF] *pairs*, annotated with a stance value that could be *Accept*, *Reject* or *No Stance*. Statistics for the number of tweets evoking a VHF, as well as of the stance their authors have towards the VHF, are provided in Table 2. To evaluate the quality of judgements, we randomly selected a subset of 1,000 tweets (along with the VHF against which they have been judged a *stance* value), which have been judged by at least two different language experts. Inter-judge agreement was computed using Cohen’s Kappa score, yielding a score of 0.67 for the *stance* of tweets

Split	Evoked	Accepted	Rejected	No Stance	Total
train	8,390	5,241	1,668	1,481	10,250
dev	941	567	211	163	1,115
test	2,285	1,461	448	376	2,815
Total	11,616	7,269	2,327	2,020	14,180

Table 2: Distribution of stance values for VHF in the Training, Development, and Test splits of COVAXFRAMES.

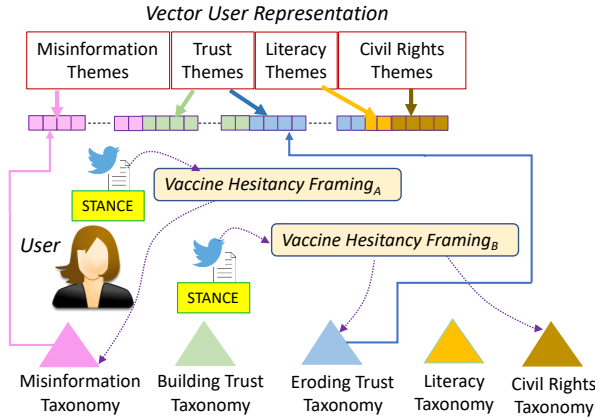


Figure 5: Representation of the Twitter users discussing vaccine hesitancy and having a stance towards hesitancy framings through Vector User Representations.

for COVID-19 VHF, which indicates moderate agreement between annotators (0.60-0.79).

When we performed stance recognition with the STANCEID-MORALITY system on all the tweets that may evoke a VHF from COVAXFRAMES, produced by the QA system responding to any COVAXFRAMES (discussed in Section 2), we identified a total of 1,741,269 tweets from 805,336 users which held an *Accept* or *Reject* stance towards one or more VHF from COVAXFRAMES.

5 Deriving Vaccine Hesitancy Profiles

Revealing the VHPs from the 805,336 users having a stance towards any of the COVAXFRAMES requires first to produce a representation of each of these users that encodes knowledge about the way the users frame their vaccine hesitancy as well as the stance they have regarding it. To encode the knowledge regarding vaccine hesitancy, we relied on the ontological commitments we have produced for COVAXFRAMES. More specifically, we decided to use the themes encoded in the taxonomies illustrated in Figure 3 and considered the stance each user had towards VHF within each theme. As shown in Figure 5, we produced for each user a vector $u \in \mathbb{R}^H$, where $H = 27$ is the total number of themes across all taxonomies illustrated in Figure 3 (i.e. 9 themes from the Misinformation taxonomy; 7 themes from the taxonomy of Building Trust; 7 themes from the taxonomy of Eroding Trust; and 2 themes each for Civil Rights and Vaccine Literacy). The values of the vector u were computed by:

$$u_h = \sum_{(t,f,s) \in U(u,h)} \frac{SV(s)}{|U(u,h)|} \quad (6)$$

where $U(u, h)$ is the set of tweets t authored by u , evoking a VHF f belonging to the theme h ; while s represents the stance value of u towards f , with $s \in \{Accept, Reject\}$. $SV(s) = 1$ if the user u accepted f and $SV(s) = -1$ if u rejected f .

In this way, we obtained 805,336 sparse vectors representing the users, that enabled us to cast the recognition of VHPs as a clustering task, which could reveal the groups of users that manifest vaccine hesitancy with similar stance towards VHF that share the same knowledge, encoded at the level of the hesitancy theme. For this purpose, we performed sparse k-means (Lloyd 1982) clustering on these user vector representations, varying the number of clusters k from 2 to 12. The final number of profiles, k was selected following the Elbow method (Thorndike 1953), i.e. by (a) computing the L2-distance between the each VHP’s centroid vector and all user vectors in the profile; and (b) computing the average distance of these distances for each VHP. We found that the number of VHPs for $k = 9$ as satisfying the Elbow method, as it obtained the minimum average distance from each user vector to the centroids of the VHPs, with an L2-distance of 1.05.

6 Interpretation of Vaccine Hesitancy Profiles

The 9 VHPs were manually inspected by exploring the tweets of the 50 users closest to the VHP centroid vectors, and each VHP was assigned a name based on the interpretation. Figure 6 supports our interpretations by illustrating how the predominant stances of profile users are interacting with (A) the ontology commitments; and (B) the MFs.

The UNDECIDED VHP includes 177, 836 users (22%) who are on the fence about the COVID-19 vaccines. These users are characterized by a 50/50 split in acceptance and rejection of the misinformation VHF that the COVID-19 vaccine is an unsafe poison, while also having a 60/40 split in trust in the government to provide accurate COVID-19 vaccine safety information. They tend to pick-and-choose which VHF they *Accept* and *Reject*, leading to theme-level inconsistencies in their beliefs. Users from this VHP tend to both adopt and reject VHF with MFs of *Subversion*, *Harm*, *Authority*, and *Care*. They are the primary target of those that propagate COVID-19 vaccine misinformation, such as the MISINFORMERS, and adopt misinformation nearly as often as they reject it.

The DEMOTIVATED VHP includes 89, 827 users (11%) who have largely lost their motivation to vaccinate against COVID-19. These users overwhelmingly accept demotivating VHF, e.g. that the COVID-19 vaccine does not provide immunity, that you can still get infected even after getting vaccinated, and that breakthrough cases after getting fully vaccinated are common. The users in this profile accept some COVID-19 vaccine misinformation about the vaccine ingredients, but primarily they are complacent, considering that the perceived risks of the COVID-19 vaccine do not justify the uptake. The predominant MFs of the VHF they evoke

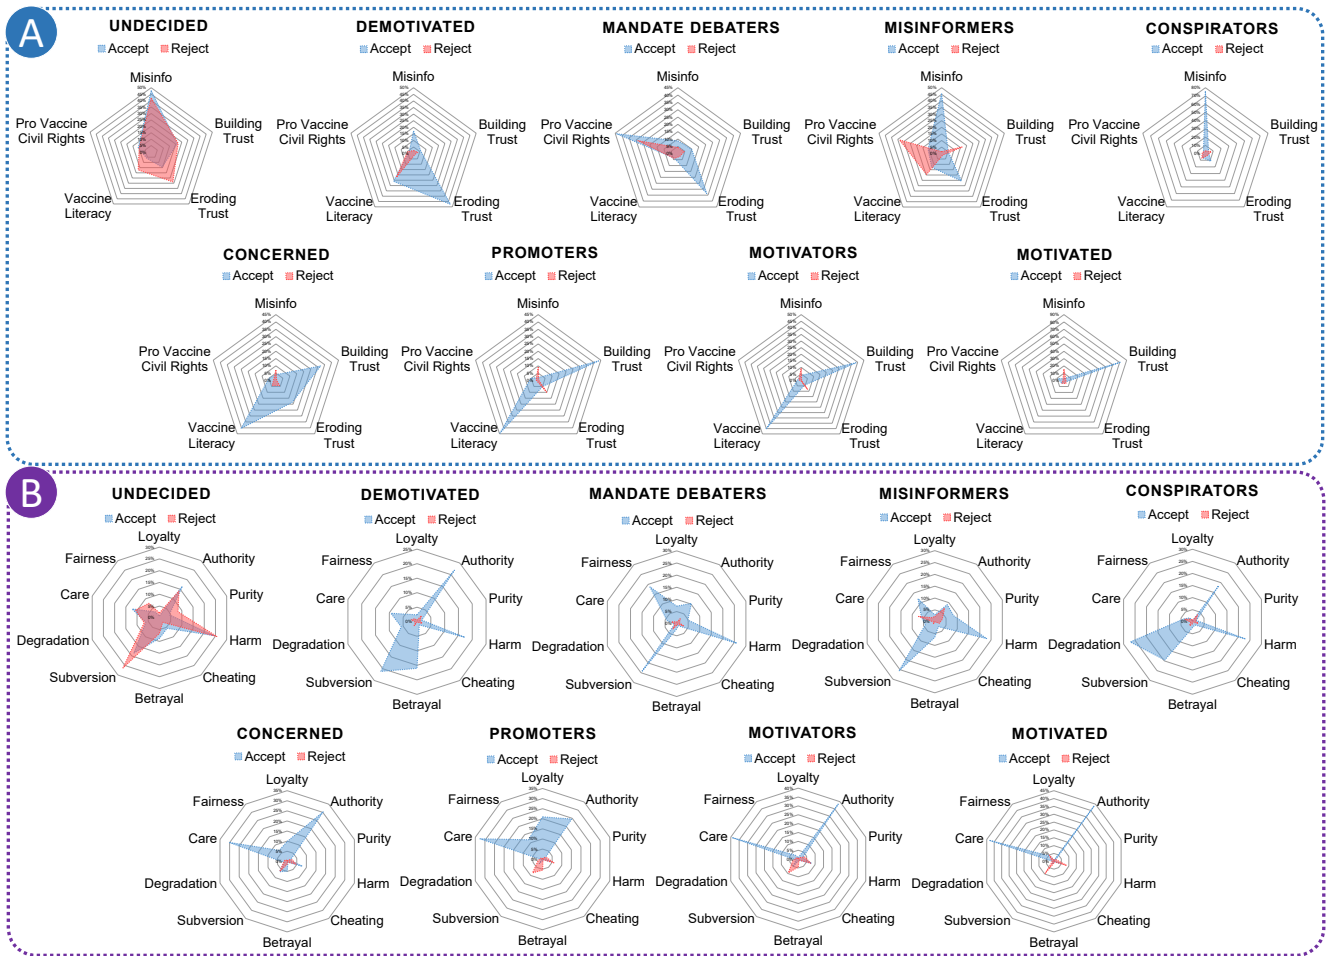


Figure 6: COVID-19 Vaccine Hesitancy Profiles (VHPs) informed by *stance* and (A) ontological commitments of Vaccine Hesitancy Framings (VHF) or (B) Moral Foundations (MF) implied by Vaccine Hesitancy Framings (VHF).

are: *Authority* and *Subversion* with an undertone of *Harm*, *Betrayal*, and *Care*. Both pairs of *Authority* and *Subversion*, and *Care* and *Harm*, are diametrically opposed MFs, which may explain why they believe that vaccination is necessary for others, but not themselves.

The MANDATE DEBATERS VHP includes 86,306 users (11%) who discuss the civil rights issues surrounding mandating vaccination against COVID-19. These users heavily adopt the VHF that everyone should make their own informed decisions about COVID-19 vaccines and that people should not be chastised on whether they decide to avoid the vaccine, but otherwise debate whether vaccine mandates are ever appropriate. They overwhelmingly adopt the VHF that all healthcare workers should be vaccinated against COVID-19, and that refusing the COVID-19 vaccine puts the lives of others at risk, but also adopt erosion of trust framings which surround perceived issues with COVID-19 vaccines, such as their concern that the AstraZeneca vaccine may cause blood clots. The predominant MFs evoked by accepted VHF are *Harm* and *Subversion*, with an additional focus on *Fairness* and *Authority*. They have the highest adoption of *Fairness* of

all the VHPs, which aligns with their focus on what is fair with regard to COVID-19 vaccine mandates.

The MISINFORMERS VHP includes 37,906 users (5%) who aggressively propagate COVID-19 vaccine misinformation. These users intensely adopt VHF containing misinformation or eroding trust in vaccines, and completely reject VHF about vaccine mandates. They are entirely demotivated to take the COVID-19 vaccine because they believe that it is unnecessary, since the survival rate of COVID-19 is 99.99%, and that the vaccine does not provide immunity. They also believe that the COVID-19 vaccine is an unsafe poison, that the vaccine is experimental and should not be used for children, and that the vaccine has not been sufficiently tested. These users evoke VHF with MFs that clearly align with those evoked by the UNDECIDED, and entirely adopted by the DEMOTIVATED. Espousing these MFs corresponds to the undue influence of the MISINFORMERS, leading to the propagation of COVID-19 vaccine misinformation to a wider audience, which we can see is often adopted by both the UNDECIDED and the DEMOTIVATED.

The CONSPIRATORS VHP includes 26,822 users (3%)

who believe in COVID-19 vaccine conspiracy theories. These users solely adopt VHF referring to conspiracy theories, e.g. the vaccine contains a neurotoxin; or the mRNA vaccine is gene therapy which will change your DNA; or the government hides vaccine safety information, and that the vaccine itself contains the virus. The VHF they adopt have MFs such as *Degradation*, *Harm*, *Subversion*, and *Authority*, and they belong to the only VHP which has heavy focus on *Degradation*.

The CONCERNED VHP includes 67,845 users (8%) who follow closely the science regarding the COVID-19 vaccines, maintaining high vaccine literacy, but still have some minor concerns with the COVID-19 vaccines, manifested in their adoption of VHF that erode trust in vaccines. These users adopt VHF which focus on the ability to reduce mortality by the COVID-19 vaccine, believing that the vaccines protect against grave forms of COVID-19 and are absolutely necessary for those at-risk. They believe that the COVID-19 vaccines will protect against emerging variants, but are generally concerned with some specific VHF, such as that the AstraZeneca vaccine may cause blood clots, the belief that the vaccines do not provide immunity, or the concern that the vaccine may not have been tested for long enough time yet. They adopt VHF with MFs such as *Authority*, *Care*, and *Loyalty*.

The PROMOTERS VHP includes 135,933 users (17%) who actively promote the role of the COVID-19 vaccines for public health. These users are characterized by overwhelming adoption and propagation of VHF that build trust in vaccines, and the framings they evoke highlight their vaccine literacy. These users overwhelmingly accept that vaccination is key in protecting yourself and others against COVID-19, that vaccination protects against severe COVID-19, that the vaccines will protect against emerging variants, and that the government has provided plenty of vaccine safety information. They also have a secondary focus on mythbusting through rejecting misinformation and trust-eroding framings, such as rejecting that the COVID-19 vaccine has not been sufficiently tested. These users also adopt many VHF that build trust in vaccines, while the VHF they evoke have the MFs of *Care*, *Authority*, and *Loyalty*.

The MOTIVATORS VHP includes 131,717 users (16%) who specifically try to motivate users to get vaccinated. These users are similar to the PROMOTERS, widely adopting of VHF for building trust in vaccines, e.g. the COVID-19 vaccines protect against the emerging variants, or the vaccines trigger your body to naturally create immunity more reliably than getting COVID-19, and that the COVID-19 vaccines have been tested, tracked, and are safe. They also reject misinformation and VHF eroding trust in vaccines. These users have high vaccine literacy, and the VHF they evoke have the MFs of *Care* and *Authority*.

The MOTIVATED VHP includes 51,144 users (6%) who share stories of themselves and others getting vaccinated against COVID-19, including reassurance that the side effects are very minor and that the vaccines are extremely safe. They adopt trust-building framings such as that the lingering effects and risks of COVID-19 are much worse than the minor side effects of getting vaccinated. The VHF they evoke have the

MFs of *Care* and *Authority*.

7 Experimental Results

To evaluate the VHPs, we performed an external evaluation by sampling pairs of users and judging if these users should belong in the same or a different cluster based on the content of their tweets. We selected the 50 users for each VHP having user vector representations that are closest to the VHP’s centroid vector. For each of the 9 VHPs, we then sampled from these 50 users without replacement 20 pairs of users from the same VHP and 20 pairs of users from different VHPs, obtaining a total of 360 pairs of users. This approach ensured that bias was removed from the manual evaluation, such that each pair of users had a 50% chance of being from the same VHP and a 50% chance of being from a different VHP. Language experts were tasked with judging if each pair was from the *same* VHP cluster or from a *different* VHP cluster by inspecting the content of the tweets of each pair of users. We then compared these judgements to the VHP assignments obtained by k-means clustering. K-means clustering of $k = 9$ produced a Rand index of 0.840 and a Fowlkes-Mallows index of 0.847 (Fowlkes and Mallows 1983), with Precision of 0.815 and Recall of 0.880. The random-clustering baseline had 0.5 for the Rand index, the Fowlkes-Mallows index, Precision, and Recall. The comparison with the random-clustering baseline demonstrates quantitatively that the 9 VHP clusters are better-than-random. Moreover, our results show that the 9 VHPs are of high quality, since users were 7.3 times more likely to be identified as similar if they were from the same VHP and 4.0 times more likely to be identified as different if they were from different VHPs.

The discovery of the VHPs was made possible by the recognition of the tweets that evoke the VHF that we identified and the stance of the tweet authors. Therefore, we also evaluated the quality of stance recognition. Stance recognition performance towards VHF in the test collection of COVAXFRAMES was evaluated on three systems: (1) the STANCEID-BASELINE system; (2) the STANCEID system; and (3) the STANCEID-MORALITY system. The STANCEID-BASELINE system utilizes the “[CLS]” embedding from COVID-Twitter-BERT-v2 as the framing stance recognition input embedding z . The STANCEID system utilizes Lexical, Emotion, and Semantic Graph Attention Networks to produce the framing stance recognition input embedding z (Weinzierl, Hopfer, and Harabagiu 2021). The STANCEID-MORALITY system, described in Section 4 and illustrated in Figure 4, utilizes Lexical, Emotion, and Semantic Graph Attention Networks along with Hopfield Pooling of Moral Foundations to perform framing stance recognition. Hyperparameters were selected based on initial experiments on the training and development collections of COVAXFRAMES. All system hyperparameters follow those of Weinzierl, Hopfer, and Harabagiu (2021), while STANCEID-MORALITY also performs Hopfield Pooling with Moral Foundations with $p = 6$, has a GAT hidden size $F = 32$, and $d = 3$ stacked GAT layers. All systems follow the same training schedule: 10 epochs, a linearly decayed learning rate of $5e - 4$ with a warm-up for 10% of training steps, and an attention drop-out

System	Macro F1	Macro P	Macro R	Accept F1	Accept P	Accept R	Reject F1	Reject P	Reject R
STANCEID-BASELINE	69.1	68.8	69.5	81.0	79.3	82.8	57.2	58.2	56.2
STANCEID	72.4	69.6	75.4	80.6	77.1	84.5	64.1	62.1	66.3
STANCEID-MORALITY	75.2	73.0	77.9	83.6	77.8	90.5	66.8	68.3	65.4

Table 3: Framing Stance Recognition results on the COVAXFRAMES test collection.

rate of 10%. Results are provided in Table 3.

Performance was determined based on Precision (P), Recall (R), and F_1^1 score for detecting the *Accept* and *Reject* values of stance. We also compute a Macro averaged Precision, Recall, and F_1 score. The STANCEID-BASELINE system produced a Macro F_1 score of 69.1, which demonstrates the advantage of pre-training BERT on domain-specific COVID-19 tweets and fine-tuning stance recognition systems. The STANCEID system produced a Macro F_1 score of 72.4, which indicates that integrating Lexical, Emotional, and Semantic Graphs improves stance recognition. The STANCEID-MORALITY system produced a Macro F_1 score of 75.2, supporting our hypothesis that MFs play a key role in detecting tweets which evoke VHF along with recognizing *acceptance* and *rejection* of VHF. The results also show that detecting *rejection* of VHF is more difficult than detecting *acceptance*.

Improvements in stance recognition for the STANCEID system are driven by results for the *Reject* stance. The *Reject* stance has the fewest number of [tweet, VHF] pairs, with only 2,327 instances in a dataset of 14,180 [tweet, VHF] pairs. The STANCEID system overcomes this resource constraint by integrating additional Lexical, Emotion, and Semantic information. Stance recognition is further improved by the STANCEID-MORALITY system for both the *Accept* and *Reject* stance values. The STANCEID-MORALITY system clearly benefits from integrating MF resources with the Hopfield pooling approach, which provides the best results on recognizing both *acceptance* and *rejection*.

8 Conclusion

In this paper we described a novel methodology for recognizing Vaccine Hesitancy Profiles (VHPs), applied to the COVID-19 vaccines. This methodology relies on the identification of how people frame their vaccine hesitancy, what Moral Foundations are implied by their Vaccine Hesitancy Framings (VHFs), and what stance the tweet authors have towards COVAXFRAMES. By considering the ontological commitments of the VHF from COVAXFRAMES we derived nine VHPs of 805,336 Twitter users having a stance towards some COVAXFRAMES. The interpretation of the VHPs revealed that 22% of these users are UNDECIDED; 11% are DEMOTIVATED; 11% are MANDATE DEBATERS; 5% are MISINFORMERS; 3% are CONSPIRATORS; 8% are CONCERNED; 17% are PROMOTERS; 16% are MOTIVATORS; and 6% are MOTIVATED.

¹ F_1 is defined as $F_1 = 2 \times P \times R / (P + R)$

Ethics Statement

Accurate vaccine hesitancy profiling at scale has the potential to enable public health researchers to design customized interventions to target users most likely to be convinced to vaccinate. Public health outreach could become much more personalized, directly addressing the themes, concerns, and moral priorities held by users on Twitter. Potential downsides to the approach outlined in this paper include the mistaken assignment of Twitter users to hesitancy profiles, due to sarcasm, jokes, or untruthful postings, which may be difficult for our system to recognize. Additionally, many users may change their stance and remove or rebut their own tweets over time. The downside of mistaken user hesitancy profiling is minimal, as we expect our system to be used by public health practitioners when developing their interventions. Questionnaires that precede the application of an intervention would filter out Twitter for incorrect hesitancy profiles.

Our data collection process was reviewed and approved by the Institutional Review Board at the University of Texas at Dallas. All tweets collected were public, and only the tweet IDs and annotations will be shared, such that others must go through the approval process to use the data. We expect COVAXFRAMES would become a valuable resource for identifying Vaccine Hesitancy Framings, and recognizing the *stance* each Twitter user has towards those framings. While we believe that the annotation quality of COVAXFRAMES is high (0.67 Cohen’s Kappa score), mistakes in judgements of *stance* are likely due to difficult complex phenomena, such as sarcasm. We believe that such potential misjudgments are rare and thus minimally impact the quality of the hesitancy profiles.

References

- Auxier, B.; and Anderson, M. 2021. Social media use in 2021. *Pew Research Center*.
- Beaulieu, M. M.; Gatford, M.; Huang, X.; Robertson, S.; Walker, S.; and Williams, P. 1997. Okapi at TREC-5. In *The Fifth Text REtrieval Conference (TREC-5)*, 143–165.
- Cambria, E.; Poria, S.; Hazarika, D.; and Kwok, K. 2018. SenticNet 5: Discovering Conceptual Primitives for Sentiment Analysis by Means of Context Embeddings. In *AAAI Conference on Artificial Intelligence*.
- Chong, D.; and Druckman, J. N. 2007. Framing Theory. *Annual Review of Political Science*, 10: 103–126.
- Clevert, D.-A.; Unterthiner, T.; and Hochreiter, S. 2016. Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). arXiv:1511.07289.
- Das, A. S.; Datar, M.; Garg, A.; and Rajaram, S. 2007. Google News Personalization: Scalable Online Collaborative

- Filtering. In *Proceedings of the 16th International Conference on World Wide Web*, WWW '07, 271–280. New York, NY, USA: Association for Computing Machinery.
- 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. Association for Computational Linguistics.
- Du, J.; Xu, J.; Song, H.-Y.; and Tao, C. 2017. Leveraging machine learning-based approaches to assess human papillomavirus vaccination sentiment trends with Twitter data. *BMC Medical Informatics and Decision Making*, 17.
- Entman, R. M. 2004. *Projections of Power: Framing News, Public Opinion, and U.S. Foreign Policy*. University of Chicago Press.
- Field, A.; Kliger, D.; Wintner, S.; Pan, J.; Jurafsky, D.; and Tsvetkov, Y. 2018. Framing and Agenda-setting in Russian News: a Computational Analysis of Intricate Political Strategies. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3570–3580. Association for Computational Linguistics.
- Foundation, A. S. 1999. Apache Lucene. <https://github.com/apache/lucene>. Accessed: 2022-04-18.
- Fowlkes, E. B.; and Mallows, C. L. 1983. A Method for Comparing Two Hierarchical Clusterings. *Journal of the American Statistical Association*, 78(383): 553–569.
- Garett, R.; and Young, S. D. 2021. Online misinformation and vaccine hesitancy. *Translational Behavior Medicine*, 11: 2194–2199.
- Haidt, J.; and Graham, J. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20: 98–116.
- Haidt, J.; and Joseph, C. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Social Justice Research*, 133: 55–66.
- Kata, A. 2010. A postmodern Pandora’s box: Anti-vaccination misinformation on the Internet. *Vaccine*, 28: 1709–16.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Larson, H. J.; Jarrett, C.; Schulz, W. S.; Chadhuri, M.; Zhou, Y.; Dube, E.; Shuster, M.; MacDonald, N. E.; Wilson, R.; and on Vaccine Hesitancy, S. W. G. 2015. Measuring Vaccine Hesitancy, the Development of a Survey Tool. *Vaccine*, 33: 4165–4175.
- Lloyd, S. 1982. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2): 129–137.
- Luo, X.; Zimet, G.; and Shah, S. 2019. A natural language processing framework to analyze the opinions on HPV vaccination reflected in twitter over 10 years (2008–2017). *Human Vaccines and Immunotherapeutics*, 33: 1496–1504.
- Macdonald, N. 2015. Vaccine hesitancy: Definition, scope and determinants. *Vaccine*, 32.
- McHugh, M. L. 2012. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3): 276–282.
- Müller, M.; Salathé, M.; and Kummervold, P. E. 2020. COVID-Twitter-BERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter. <https://arxiv.org/abs/2005.07503>.
- Nenkova, A.; and Passonneau, R. 2004. Evaluating Content Selection in Summarization: The Pyramid Method. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, 145–152. Association for Computational Linguistics.
- Nogueira, R.; and Cho, K. 2020. Passage Re-ranking with BERT. arXiv:1901.04085.
- Ramsauer, H.; Schäfl, B.; Lehner, J.; Seidl, P.; Widrich, M.; Gruber, L.; Holzleitner, M.; Pavlovic, M.; Sandve, G. K.; Greiff, V.; Kreil, D. P.; Kopp, M.; Klambauer, G.; Brandstetter, J.; and Hochreiter, S. 2020. Hopfield Networks is All You Need. *CoRR*.
- Rao, A.; Morstatter, F.; Hu, M.; Chen, E.; Burghardt, K.; Ferrara, E.; and Lerman, K. 2020. Political Partisanship and Anti-Science Attitudes in Online Discussions about COVID-19 (Preprint). *Journal of Medical Internet Research*, 23.
- Rossen, I.; Hurlstone Angeli, M. J.; Dunlop, P. D.; and Lawrence, C. 2019. Accepters, fence sitters, or rejecters: Moral profiles of vaccination attitudes. *Social Science Medicine*, 23–27.
- Roy, S.; and Goldwasser, D. 2020. Weakly Supervised Learning of Nuanced Frames for Analyzing Polarization in News Media. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7698–7716. Online: Association for Computational Linguistics.
- Thorndike, R. L. 1953. Who belongs in the family? *Psychometrika*, 18(4): 267–276.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Weinzierl, M.; and Harabagiu, S. 2020a. HLTRI at W-NUT 2020 Shared Task-3: COVID-19 Event Extraction from Twitter Using Multi-Task Hopfield Pooling. In *Proceedings of the Sixth Workshop on Noisy User-generated Text (W-NUT 2020)*, 530–538. Association for Computational Linguistics.
- Weinzierl, M.; and Harabagiu, S. M. 2020b. The University of Texas at Dallas HLTRI’s Participation in EPIC-QA: Searching for Entailed Questions Revealing Novel Answer Nuggets. In *Thirteenth Text Analysis Conference*, volume 13. Text Analysis Conference.
- Weinzierl, M.; Hopfer, S.; and Harabagiu, S. 2021. Misinformation Adoption or Rejection in the Era of COVID-19. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*. AAAI Press.