Understanding the Anti-Mask Debate on Social Media Using Machine Learning Techniques

Luca Cerbin*, Jason DeJesus*, Julia Warnken*, and Swapna S. Gokhale*
University of Connecticut, Storrs, CT 06269 USA

Abstract

Masks are believed to slow the spread of Covid-19, and can prevent many deaths, yet this inexpensive, common sense public health measure has ignited a fierce debate in the United States. Opponents of masks or anti-maskers have resorted to measures such as organizing protests and marches to make their views public. They have also taken to social media platforms to vigorously argue against the use of masks. Even with the advent of vaccines, masks are still likely to be recommended for a long time. It then becomes important to mine the debate around masks to understand the concerns of the detractors and the arguments used by the proponents to counter these concerns. This paper analyzes the mask dialogue on Twitter, using the data collected in July and August 2020, which coincided with the time when the stay-at-home orders were being relaxed, and the opening of schools and other activities was being contemplated. These tweets are explored in three ways – informal opinion mining is used to reveal the reasons for concerns and support, social parameters of the tweets and tweeters are analyzed to expose the dynamics of the two communities, and classification framework is built to distinguish between pro- and anti-mask tweets so that the latter can be tagged to prevent the spread of discordant information. Our results indicate that the concerns of anti-maskers are more political and ideological rather than related to the adverse health impacts of masks. Members of the close-knit, small anti-mask community promote each other's views compared to the pro-maskers, although the anti-maskers themselves are not fringe by any means. The classification framework can detect anti-mask tweets with excellent accuracy of over 90%, and hence, it can be used to label tweets that sow misinformation about masks before they spread through the ether and influence people.

Key Words: Masks, anti-mask, pro-mask, twitter, classification, machine learning.

1 Introduction and Motivation

The coronavirus pandemic has upended every single tenet and ritual of our modern society. Discussion and practice of measures such as masks, social and physical distancing, vaccines, hand hygiene, and disinfectants have now become a part of our daily routines. Of these, one of the most contentious issues that has bitterly divided the U.S. society is the wearing of masks. A seemingly simple act of wearing a facial covering that covers both the mouth and the nose serves as a stark reminder of the pandemic, and has also been the topic of a fierce debate. Proponents of masks point to several studies that recommend their use to slow the spread of Covid-19 [19]. Opponents, however, contend that most of the studies have looked at the use of face masks in health care, and not community settings. They further claim that these studies were observational, not the gold standard of science, a randomized controlled trial. It does not help that early in the pandemic public health officials in the U.S. discouraged the use of masks by the general public. At the time "mass masking" was not recommended either by the CDC or the WHO, perhaps to conserve them for healthcare and other front-line workers [9]. Later, however, they backtracked from this initial position and vigorously advocated the use of masks to blunt the spread of the virus and prevent deaths. The u-turn regarding masks and the subsequent political divide over them has come to symbolize the chaos of the U.S. response to the still-raging pandemic [50].

Expressions of pro-mask and anti-mask opinions are plentiful and varied in the physical, offline world. In some counties, where the coronavirus has surged out of control, mask mandates have been imposed and this has further outraged their residents. Those opposed to mask mandates have staged protests, and one local health official had to even quit her job after receiving a death threat for a mask order [33]. In addition to expressing their views through their actions by either wearing or not wearing masks in public spaces and/or organizing protests, people have often turned to social media platforms such as Twitter and Facebook to express their support or opposition to masks. These social media platforms havenot only been woven tightly into the fabric of our society, but sharing on these platforms has skyrocketed especially during the pandemic, because a number of people are either in self-imposed or government-mandated isolation and lockdown. Therefore, in addition to the offline expression of the pro- and anti-mask opinions, this debate over masks has been playing out vociferously over these platforms as well.

Compliance with masks has been spotty at best through the U.S., even though the CDC and other public health experts have repeatedly indicated, on multiple occasions, that wearing

* Computer Science & Engg. Email: {swapna.gokhale} @uconn.edu.

masks could save a significant number of lives [13, 23]. Furthermore, the use of masks is likely to continue despite the approval and roll out of vaccines. In fact, masks and social distancing will probably be recommended at least for a while, because a lot is still unknown about what protections vaccines can afford in terms of preventing infection, its severity and its spread [44]. It is thus believed that masks are and will continue to be an effective tool against fighting the pandemic. Given the usefulness of masks, it is then imperative to understand the public outlook towards their use. Based on such understanding we can launch educational and public awareness initiatives to dispel the myths and misinformation and encourage their adoption broadly. Moreover, understanding the drivers and spread of misinformation can be valuable during future pandemics.

The novelty of the paper lies in understanding the debate over masks through social media dialogue. Using the data collected in July 2020 and August 2020 from Twitter, just when the "stay home, stay safe" orders were beginning to be relaxed, and the opening of schools was being contemplated, this paper seeks to answer the following research questions:

(i) Will masks be embraced by the community at large, or are there a significant number of detractors and skeptics (anti-maskers) who will continue to defy the simple, inexpensive and most innocuous of the public health guidelines? What misgivings do the detractors and skeptics express? What misinformation about masks is circulating on social media, which if left unchecked will make an eventual broad scale acceptance of masks by the public almost impossible? (ii) How socially cohesive and tight knit is the community of anti-maskers, compared to the group of supporters? (iii) To curb the spread of discordant information, is it feasible to automatically detect the tweets that carry misinformation and express skepticism about masks before they make their way through the ether? This is especially important as social media users are more likely to believe false information about Covid-19 and ignore public health advice [43].

Our results expose the culture wars associated with the use of masks. Concerns of anti-maskers appear to be more motivated by politics and ideology, rather than driven by actual health, convenience or any other pragmatic reasons. Benefits to public health, advocated by pro-maskers to counter this anti-mask rhetoric is weak and unlikely to be convincing and per-suasive. Although the anti-mask views are not fringe, the group of anti-mask users is small, tight knit, and very supportive and encouraging of each other. Despite the small size of their network, anti-maskers have effectively spread their opinions and views widely. Separating the anti-mask tweets from the pro-mask ones is feasible, and can be accomplished with high accuracy by employing a combination of linguistic, auxiliary, and social features to train machine learning models. Most ML classifiers, including Support Vector Machines, RandomForest, Gradient Boosting, achieve an accuracy of over 90% in separating the anti-mask tweets from the pro-mask ones. Importance analysis shows that a bulk of the contribution towards classification comes from the text of the tweets, and from the social parameters that indicate the reach and popularity of the tweets and the tweeters.

The rest of the paper is organized as follows: Section 2 explains the process of collecting and preparing the data. Section 3 describes opinion mining. Section 4 summarizes the findings of social analysis. Section 5 presents the sequence of steps involved in building the classification framework. Section 6 discusses the results. Section 7 compares and contrasts related research. Section 8 offers concluding remarks and directions for future research.

2 Data Preparation

This section discusses three steps in the preparation of data: data collection, data labeling, and data pre-processing.

2.1 Data Collection

Data was collected twice, one month apart, using the crawling seeds *#wearadamnmask*, *#nomaskforme*, *#maskupamerica*, *#masksareforsheep*, *#nomasks*, *#nomaskmandate*, *#antimaskers*, *#maskitorcasket* in July 2020 and August 2020. These two-time frames were chosen as they represented two significant epochs in the mask debate. In July 2020, as the country was emerging from the lockdown, masks were viewed as a way to restore a sense of normalcy. Furthermore, masks came into sharp spotlight in this one-month period because of the tussle surrounding the reopening of schools, and students returning to college campuses. Masks also became a hot button issue during this period when the Democratic presidential candidate Joe Biden suggested that if elected he will issue a national mask mandate [39]. In the same period, leading public health experts, including the CDC promoted the use of masks as "life-saving", highlighting that if everyone committed to wearing masks, we could save a significant number of American lives [13]. Thus, the two data collection epochs one month apart occurred during an eventful period for the fate of the masks and their acceptance. Both data sets were collected using the using the rtweet library in R [28]. The following represent examples of pro-mask and anti-mask tweets from the July 2020 and August 2020 data sets.

July 2020:

Erry time I ride @trimet to work there's always a couple ppl not wearing masks... like really? #WearADamnMask (**P**)

History REPEATING itself #NoMasks #NoMaskOnMe #BLMIsADomesticTerroristGroup https://t.co/oJjoPA9WW9 (**A**)

August 2020:

At least my mask hides my pimple #moreimportantly issaveslives #WearADamnMask https://t.co/hbxquonf8R (**P**)

@healthvermont @CDCgov We are tired of the government telling us to stay safe. Freedom Trumps safety in America. Fellow Vermonters tell the government. We will not comply! #NoMasks #freedom #Vermont (**A**)

2.2 Data Labeling

This set of crawling seeds was harvested because it included both the anti-mask and pro-mask perspectives. For example, we expected that hashtags such as *#maskupamerica* and *#maskitorcasket* would be used in tweets that support masks, whereas hashtags such as *#nomasksforme* and *#masksareforsheep* would be used to show opposition. We anticipated that the tweets would neatly separate according to support and opposition, consistent with the corresponding hashtags. Such clear, neat separation would obviate the need for manual labeling and facilitate weak supervised learning with the hashtags serving as labels. Skimming through the tweets, however, invalidated this assumption and many hashtags were creatively embedded in both supporting and opposing tweets. For example, the hashtag *#nomask* is used in the following two tweets, the first one is clearly pro-mask whereas the second one is anti-mask. In fact, the use of anti-mask hashtags in tweets that express pro-mask opinion has been found to be prevalent. It is believed that such use inadvertently boosts the anti-mask movement, making it difficult to automatically separate such tweets [12]. Such mocking may also fuel the anti-maskers.

Save lives - wear a mask, clean your hands keep a safe distant. #nomask (**P**)

I feel fine cause I dont wear one! #nomask (**A**)

Manual annotation of the tweets seemed inevitable, and was undertaken to classify each tweet into one of two groups –'A' for anti-mask, and 'P' for pro-mask. The entire data set was labeled twice, independently, with a gap of about one week between the two labelings. Duplicates were eliminated before the labeling. Only those tweets where the labels matched on two independent occasions were included in the final corpus, which consisted of 4042 tweets.

About 500 tweets were eliminated because of mismatch of labels. In the corpus, about 57% of the tweets are pro-mask, and 43% are anti-mask. This data also contained a number of public safety announcements (PSAs) from schools, colleges and sports teams. There were tweets that expressed political opinion regarding the conventions, wildfires in California, and the BLM protests without the express mention of masks other than the hashtag. In the manual labeling process, we eliminated these tweets to build a high-quality data set that truly reflects the public opinion about masks instead of other peripheral and allied political issues.

2.3 Data Pre-processing

The labeled data was pre-processed in the following steps shown in Figure 1. It was converted to UTF-8 encoding, and transformed to lower case. Then, numbers, punctuation and stop words were removed. After word stemming and stripping white space, domain specific words that occur in both pro-mask and anti-mask tweets with a similar frequency were removed as they are likely to be uninformative.

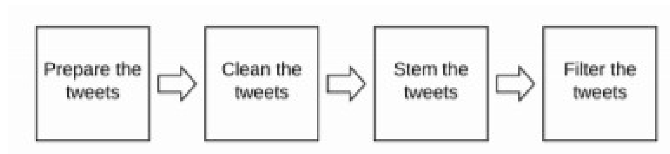

Figure 1: Pre-processing steps

3 Opinion Mining

We represent the remaining words in each of the pro-mask and anti-mask categories into word clouds as shown in Figures 2a and 2b. We read words from these word clouds, and find associations between them to reveal interesting insights into the opinions of supporters and detractors.

Proponents point to the life-saving benefits associated with the use of masks. They promote the use of masks through several phrases such as "masks save lives", "wear a mask save a life". They advocate the covering of both the mouth and nose. Medical terms such as doctor, hospital, patient appear which points to the vital role played by medical and front-line workers during the pandemic. Pro-maskers are also believers of allied public health recommendations such as staying home, washing hands, and practicing social distancing to curb the spread of the virus. The role of masks in bringing

Figure 2: Word clouds of pro-mask and anti-mask tweets

the lives of students and children back to normal is highlighted. Blame for mismanaging the Covid-19 pandemic is laid at President Trump's feet with the phrase "trumpvirus". Some pro-maskers also appear to be supporting a more aggressive stance over the use of masks through a mandate. It can be imagined that some pro-maskers may have had encounters with anti-maskers, through the use of the words no mask, walk, guy, masker.

Anti-maskers persistently and vigorously oppose masks, as evident through the repeated use of phrases such as nomasks, nomasksforme, masksoffamerica, nomaskmandate. The overlap or intersection between anti-maskers and anti-vaxxers is on display through the terms novaccine and vaccine. This group is of the fervent belief that Covid is a hoax, as expressed through the terms covidhoax, scamdemic, and plandemic. Pandemic is Dr. Milkovitz's documentary that claims to have exposed Dr. Fauci's fraud [15]. Views such as masks take away freedom, and amount to tyranny are voiced. Calls to open businesses, end the lockdown, especially directed at Republican politicians such as Gov. Abbott and Gov. Mike

Devine who opposed mask mandates are found in these clouds. The sentiment that the government is fear mongering and withholding the truth are also expressed. Joe Biden and his proposed mask mandate after inauguration also appear in the list of concerns.

This mining and analysis of associations exposes that the concerns against the use of masks are more political and ideological, rather than being rooted in any health or convenience matters. This divide may have largely been fueled by the Trump administration's defiance towards the wearing of masks. Proponents and opponents can be seen to be split along the political divide. Those leaning left view the virus as a serious threat, and those leaning right tend to downplay its seriousness. In viewing a mask as a political symbol, proponents may be viewed as trying to hype the seriousness of the virus, whereas, opponents may be viewed as purposely trivializing the virus and prioritizing the economy over health and safety. Interestingly, both proponents and opponents argue that masks block their respective happiness, to the proponents of masks, blocking of happiness is tantamount to threatening their safety, whereas to the opponents the blocking of happiness arises from hijacking their liberty. Many medical reasons are cited for not wearing masks, these include claustrophobia, panic attacks, autism spectrum disorder, and sensory processing issues.

4 Social Analysis

Social media platforms are conducive to a viral spread of rumors and misinformation. Tweet data collected using the rtweet library also records a number of parameters that may indicate the reach of the tweets and tweeters, which may further offer insights into how a tweet may circulate over the platform. In this section, we explore the social parameters of the tweets [26] to examine the conjecture that the anti-mask community is close knit and much better organized and connected compared to the pro-mask com-munity. A brief description of these parameters follows. Values of these parameters, along with their classification into user- or tweet-level are reported in Table 1.

Table 1: Social features

Parameter	Pro-Mask	Anti-Mask	User/Tweet
Percentage	57.74	42.26	Tweet
Avg. Tweet Length	152.20	154.83	Tweet
Avg. Retweet Count	1.11	4.95	Tweet
Avg. Favorite Count	4.43	11.42	Tweet
Avg. Follower Count	5596.72	3099.60	User
Avg. Friend Count	2624.66	2475.14	User
Avg. Status Count	30976.34	25587.25	User
Avg. Favorites Count	67	3	User
Avg. List Count	102.72	28.54	User
Avg. Quote Retweet Count	5586.10	3342.84	Tweet
Avg. Quoted Favorite Count	21599.03	12764.23	Tweet
Avg. Quoted Follower Count	2163732.13	12677596.95	Tweet
Avg. Quoted Friend Count	7906.06	9343.33	Tweet
Avg. Quoted Status Count	76204.73	47290.24	Tweet
Percent Verified	2.18	0.18	Tweet
Percent Mentions	30.21	48.36	Tweet
Percent Replies	28.28	43.03	Tweet
Percent Quoted	22.41	24.47	Tweet

- **Avg. Tweet Length**: Number of characters in each tweet.
- **Avg. Favorite Count**: Number of times a tweet has been liked by Twitter users.
- **Avg. Quoted Favorite Count**: Number of likes the quoted tweet received.
- **Avg. Quoted Retweet Count**: Number of retweets the quoted tweet received.
- **Avg. Quoted Followers Count**: Number of followers of the user who tweeted the quoted tweet.
- **Avg. Quoted Statuses Count**: Number of status updates of the user who created the quoted tweet.
- **Avg. Quoted Friends Count**: Number of friends of the user who tweeted the quoted tweet currently has.
- **Avg. Retweet Count**: Number of times a tweet is retweeted.
- **Avg. List Count**: Number of public lists in which the tweeter claims membership.
- **Avg. Statuses Count**: Number of tweets posted by the tweeter.
- **Avg. Followers Count**: Number of followers of the tweet owner.
- **Avg. Friends Count**: Number of friends of the tweet owner.
- **Percent verified**: Percentage of tweets that were shared from verified accounts.
- **Percent quoted**: Percentage of tweets that quote other tweets.
- **Percent replies**: Percentage of tweets that were replies to existing tweets.
- **Percent mentions**: Percentage of tweets that mentioned other users.

The average length for pro-mask and anti-mask tweets is similar, however, the average retweet and favorite counts of these two classes of tweets are significantly different. Roughly,these counts are about 3-4 times higher for anti-mask compared to pro-mask tweets. If retweeting and favoriting (liking) is viewed as akin to endorsing the content of the tweets, then a potential explanation for this discrepancy could be the strength of the passion regarding anti-mask opinions, and a lack of similar passion when expressing pro-mask opinions. When a pro-mask tweeter shares a tweet supporting opinion. On the anti-mask side, however, because the nature of the position is not popular or mainstream, someone who agrees with the opposing view is more likely to make it known masks, that opinion may be assumed to be more mainstream, and it may not be considered valuable to favor a normal

Comparing the metrics at the user-level, the average follower count is higher for pro-mask users compared to anti-mask users. This skew could be explained by the much higher percentage of pro-mask tweets being posted from verified accounts compared to anti-mask tweets. Verified accounts usually belong to the more famous, elite and educated people, and they tend to have many more followers than the ordinary users. These accounts could also belong to public health authorities and organizations, who tweet to encourage people to wear masks. In terms of absolute numbers, a total of 67 verified users have shared pro-mask tweets, whereas only 3 verified users have shared anti-mask tweets. This suggests that more prominent users are in favor of masks compared to the few known ones that oppose them.

An interesting difference is in the average status count, which is slightly higher for pro-mask compared to anti-mask users. This may be possible because pro-mask users may be more extroverted and comfortable sharing their mainstream opinion. A user who shares many statuses is more likely to desire to keep their friends/followers abreast of what is happening, which is a social characteristic. Pro-mask users also have a higher favorites count, perhaps showing they are generally more active on Twitter. As further evidence of this, the average quote retweet count is significantly higher for pro-mask users than anti-mask users. Average quote favorite count and average quote status count are also higher for pro-mask users. The average list count, which indicates the number of lists or groups of which a user is a member is substantially higher for pro-mask users compared to anti-mask users. These social parameters indicate that pro-mask users appear to be more active on the platform and they are more likely outgoing, because many metrics that require active participation on the platform are significantly higher for pro-mask users. On the other hand, the percent of tweets that are replies to existing tweets, and the percent of tweets that mention other users is higher for anti-mask users. Replies always have user mentions, but not all tweets that contain user mentions are replies. The percent of tweets that contain quoted tweets is very similar for both pro- and anti-mask users. This suggests that the anti-mask community although small is deeply engaged in supporting and promoting the anti-mask view.

In summary, the pro-mask community is open and active in encouraging the use of masks (perhaps through the sharing of useful, public health benefits), but pro-maskers are less engaged with each other. Anti-mask users, on the other hand, form a tight-knit group and appear strongly interested in endorsing and propagating the anti-mask view. However, it is important to remember that about 40% of the tweets support the anti-mask view, so although the number of anti-maskers may be few and close knit, the view itself is not fringe.

5 Tweet Classification

While masks may not be completely effective, they certainly do not amount to a “dangerous waste of time”. They can be at least partially effective, and prevent a significant number of deaths, as indicated by the CDC [13, 23]. Anti-maskers seem just like many other conspiracy cultists including anti-vaxxers and flat earthers, and they share misinformation about masks. Such discordant information spreads virally over these platforms, and provides impetus in the use of these platforms to organize marches, rallies and demonstrations. Such anti-mask rhetoric can persuade people on the fence to further denounce the use of masks. Pro-maskers seem to be appealing to the collective goodwill to adhere to public health measures such as washing hands, staying home, and maintaining a safe distance, which is likely to be ineffective. It then becomes imperative to identify and label the anti-mask dialogue, in the hope of limiting its persuasive power. Given the excessive volume of content that gets sharedon these platforms, manual separation of anti-mask tweets is impossible, highlighting the need for automated detection. This section presents a classification approach to distinguish between pro-mask and anti-mask tweets, labeled as 'P' and 'A' respectively.

5.1 Feature Extraction

The first step is to extract features that abstract away the important properties of the tweets while ignoring the unnecessary details. We considered linguistic, auxiliary, social, psycho-linguistic, and sentiment features as discussed below in the classification framework.

5.1.1 Linguistic Features. Tweets were processed using natural language techniques so that the key features including the semantic relationship between the words and the contextual information of the words and sentences were numerically encoded in high-dimensional vectors. We considered a number of vector representations such as bag-of-words [52], *n*-grams, Term Frequency-Inverse Document Frequency (TF-IDF) [35] and word2vec and doc2vec [36] that are commonly used for classification. Of these, we used the *n*-grams/TF-IDF and word embeddings.

In the *n*-grams method, a sample of text is represented by the most frequent instances of every unique *n* continuous words as a dimension. The most frequent word grams are selected from the entire corpus. The tweets were represented through unigram (1-gram) vectors, and the weight for each unigram is its TF-IDF score which is given by:

$$TF\text{ - }IDF = tf * log(\frac{}{df}) \qquad (1)$$

In Equation (1), *tf* is the number of times a particular term occurs in a tweet, *T* is the total number of tweets, and *df* is the number of tweets containing that particular term. The main advantage of a $TF-IDF$ score over the simple frequency counts of the *n*-gram method is that it assigns a higher weight to the terms that occur more frequently through the entire data set. Thus, the $TF-IDF$ score should assign a higher weight to those phrases that are the most important in determining whether a tweet is anti-mask or pro-mask. After pre-processing, the size of our corpus (number of unique words) is over 9000. Of these, we calculated the TF-IDF vector representations of the top 2000 most relevant unigrams. We used the TF-IDF implementations from the NLTK library to extract these features [31].

Although the TF-IDF score provides a differentiated representation of the words based on their frequency of occurrence, it does not preserve any relationship between the words. Word embeddings are a powerful technique that represent semantically related words as closely related vectors. Words with similar meanings are mapped to low-dimensional, non-sparse vectors that exist near each other in a predefined vector space. A good word embedding can preserve the contextual information behind words in a tweet that a *n*-gram/TF-IDF scheme cannot. We use Word2Vec, which is a popular technique to create distributed numerical representations of word features using a two-layer neural network with back propagation [36]. Word2vec trains words against other words that neighbor them in the input corpus. Word2Vec allows us to encode the context of a given word by including information about preceding and succeeding words in the vector that represents a given instance of a word. Therefore, the results obtained from using Word2vec may result in a much better classification.

We implemented Word2Vec using the gensim library [45]. From the preprocessed tweets, we generated a list of tokens, and built a model to represent each word by a 10-dimensional vector, where the parameter *min count* is 1. The number of workers, which is the number of partitions during testing is 8. The model considers all the words in the corpus. We created the vector representations for all the tokens, and the total number of epochs used is 25. We used the continuous bag of words (CBOW) model to generate the representations. The other option was to use the skip gram model. Skip gram works well with a small amount of data and is found to represent rare words well. On the other hand, CBOW is faster and has better representations for more frequent words [27, 40]. We chose CBOW based on earlier success with this model to classify the anti-vaxx dialogue [42].

We also included POS (part-of-speech) tagging using the NLTK library [31]. The NLTK library provides the ability to classify each word as one of 35 parts of speech. POS tagging occurred before removing stop words to capture any differences in the raw text. The occurrences of each part of speech is counted for each tweet and fed as input to our models.

5.1.2 Psycho-Linguistic Features. Some studies show that refusal to wear masks may be linked to sociopathic, narcissistic and psychopathic tendencies [51]. These leanings are reflected in an excessive use of first-person pronouns "I" and "me", in written and spoken language. Therefore, we considered the use of these first-person pronouns in the anti-mask and pro-mask tweets. The use of these pronouns, however, did not appear significantly different in these two groups. In total, the pro-mask tweets used "I" 8 times compared to the use of "I" 14 times in the anti-mask tweets. Counting the instances of both "I" and "me", the pro-mask tweets had 103 occurrences, while the anti-mask tweets had 102. Because the differences appeared insignificant, these first-person pronouns were not considered further in the classification.

5.1.3 Auxiliary Features. Written texts including social media feeds do not carry with them clues that can be gathered from facial expressions and body language that accompany face-to-face or spoken communication. Therefore, in social media texts, users may use a variety of punctuation marks and other means such as hashtags and emoticons to emphasize their point. These auxiliary features are believed to somewhat substitute the clues that can be learned from communicating in the physical space, and are known to improve classification accuracy [14]. Therefore, we included numbers of hashtags, mentions, punctuations, links, words in upper case letters, question marks, exclamation marks, periods, quotations, and all punctuations as features.

5.1.4 Social Features: We used the social features listed in Table 1 in the classification framework. Because their values differed widely, we transformed each feature using the MinMaxScaler in sklearn library [1]. This function scales and translates each feature individually such that it lies in the range of 0 and 1. This transformation is often used as an alternative to zero mean, unit variance scaling [1].

5.1.5 Sentiment Features: Textblob [32] and Vader [25] sentiment scores, computed for each original tweet (before preprocessing) were used in the classification. TextBlob calculates the sentiment polarity for each tweet, which ranges from −1 to +1, where −1, 0 and +1 indicate negative, positive and neutral sentiment respectively. Vader computes a compound score as a normalized and weighted composite score obtained by analyzing each word in a tweet for its direction of sentiment - a negative (positive) valency for negative (positive) sentiment. It therefore ranges from -1 to +1 depending on the net sentiment of the tweet. The compound score provides a single unitary measure for the sentiment analysis of the tweet.

5.2 ML Models

We considered the following popular supervised machine learning models for classification. Implementations of these models in the Scikit package, were used [7], and the parameters chosen for implementation are listed below.

- **Random Forests**: Random Forests is an ensemble learning classification technique based on Decision Trees [30]. By using bagging to reduce variance, the method generates a number of decision trees with different training sets and parameters. Random Forests is easy to apply and a flexible approach. To a certain degree it eliminates the overfitting problem that often occurs when using decision trees. The number of trees was 100, the number of features in each tree was equal to the square-root of the number of total features by default, and each decision tree was allowed to grow fully up to its leaves.

- **Support Vector Machines (SVMs)**: Support Vector Machines (SVMs) is a powerful classification technique that estimates the boundary (called hyper-plane) with the maximum margin [49]. We used SVMs with RBF kernel, the regularization parameter C is set to 1000, and kernel coefficient gamma is set to 0.01. The remaining parameters are set to their default values.

- **Multi-Layer Perceptron (MLP)**: Multi-Layer Perceptron (MLP) is one of the feed-forward Artificial Neural Networks (ANN) that consists of input, hidden, and output layers [11]. The numbers of neurons in these four layers were 10, 8, 5, and 2. We used rectifier linear unit (ReLu) activation function to minimize the vanishing problem that the gradients of the loss function goes toward zero that usually occurs in deep neural networks.

- **Gradient Boosting (GB)**: Gradient Boosting is another ensemble learning technique which builds classifier trees in a forward stagewise fashion [16]. Each stage takes a small step towards the minimization of classification error from the previous step. The algorithm continues until a maximum number of trees are built or there is no significant improvement in minimizing the error. Finally, predictions for the test data are obtained by combining predictions of the trees built in each stage using a weighted sum to obtain the final prediction. The parameters of the gradient boosting algorithm can be classified into tree specific, boosting and miscellaneous parameters. The number of trees is set to 1600, the fraction of observations to be selected for each tree (subsample) is set to 0.55, the maximum depth of each tree is set to 5, the minimum samples in each leaf is set to 1, the learning rate which determines the impact of each tree on the final outcome is set to 0.05.

- **Long Short-Term Memory (LSTM)**: LSTM is an artificial recurrent neural network architecture used in deep learning [24]. We used Keras library to implement the model [10]. Keras computations require vectors of the same length, we truncated and pad the input sequences to 360. The model knows that zero values carry no information. In the LSTM model, the first layer is the embedded layer that uses vectors of length 100 to represent each word. The next layer is the LSTM layer with 100 memory units (smart neurons). Finally, because this is a classification problem, we use a dense output layer with a single neuron and a sigmoid activation function to predict either 0 or 1 for the anti-mask and pro-mask classes. Binary cross entropy is used as the loss function. The efficient ADAM optimization algorithm was used, and the model is batch sizes of 64 and 100 epochs.

5.3 Performance Metrics

Our objective is to identify anti-mask tweets, and hence, to define the performance metrics, we designate the anti-mask and pro-mask classes as positive and negative respectively.

Tweets can thus be classified into four groups – true positive (TP) (anti-mask labeled anti- mask), true negative (TN) (pro-mask labeled pro-mask), false positive (FP) (pro-mask labeled anti-mask), and a false negative (FN) (anti-mask labeled pro-mask). These four groups lead to the following metrics to compare classifier performance:

- **Accuracy**: Accuracy is defined as the percentage of tweets that are labeled correctly:

$$\text{Accuracy} = \frac{\text{TP+TN}}{\text{TP+FP+TN+FN}} \quad (2)$$

- **Precision**: Precision measures the percentage of the tweets that are actually anti-mask out of all the tweets that are predicted as anti-mask:

$$\text{Percision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (3)$$

- **Recall**: Recall measures how many of the anti-mask tweets are actually labeled as anti-mask:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (4)$$

- **F-score**: F-score seeks a balance between Precision and Recall:

$$\text{F1} = 2\text{x}\,\frac{\text{Prescision}*\text{Recall}}{\text{Precision+Recall}} \quad (5)$$

Precision is the percentage of relevant from the set detected and recall is the percentof relevant from within the global population [34]. Precision is an important measure to determine when the costs of a false positive is high. Applying symmetrical logic, recall would be the metric of significance when the cost of a false negative is high. In the context of detecting anti-mask tweets, false positive labeling implies that a pro-mask tweet is labeled as anti-mask, whereas a false negative labeling implies that an anti-mask tweet is labeled as pro-mask. In false positive labeling, because a pro-mask tweet may be labeled as anti-mask it may be subject to actions such as being censored or tagged for misinformation. However, any additional stringent punitive actions such as removing the tweet altogether may lead to freedom of speech violations. In false negative labeling, an anti-mask tweet will slip through the cracks and will not be tagged for carrying misinformation. While such mislabeling may cause damage by spreading discordant information, it will not lead to any violations of people's individual rights. Therefore, in this problem, precision may be a more important metric than recall. A balance may also be sought between precision and recall to trade off in fringing freedom of speech against the spread of discordant information.

6 Results and Discussion

We split the entire corpus using stratified sampling into two partitions; the training partition consisted of 75% and the testing partition contained 25% of the tweets. All the models listed in Section 5.2, except for LSTM, were trained and tested on a combination of TF- IDF, word embedding, POS tags, auxiliary and social features. LSTM was fed pre-processed text directly along with auxiliary and social features. We combined all the features for model training, guided by its success in their use in detecting tweets that spread vaccine misinformation [42]. The results of the performance metrics for all the models are noted in Table 2.

The table shows that all the classifiers except for SVM can distinguish between anti- mask and pro-mask tweets of accuracy and F1-score over 90%. Moreover, the accuracy of the SVM is only slightly lower than 90%. For some models, the accuracy reaches as high as 96%. These results show that anti-mask tweets that can sow discordant information about masks, and promote non-compliance can be accurately separated from social media dialogue. They also show that this accuracy can be achieved even after data from different time periods is combined. Each time period presents a different context or a backdrop against which this dialogue played out, in July it was lifting the lockdown, and in August it was reopening schools and restarting the sports and other activities. However, without regard to the underlying background information, pro- and anti-mask sentiment can be detected.

We use the Random Forest model to determine the importance of scores of the various types of features. The relative scores are summarized in Table 3. The table indicates that the bulk of the contribution, around 82%, which includes TF-IDF plus word embeddings plus POS tags, comes from the text of the tweets. Social features which determine the reach of the tweet and the popularity and level of activity of the tweeters contribute about 10%. Sentiment scores have very little contribution, around 3%. This could be because we found that the sentiment scores were not sufficiently different

Table 2: Performance of ML models

Model	Accuracy	Precision	Recall	F1-Score
RF	95.64	0.9359	0.9913	0.9628
LSTM	93.66	0.9305	0.9640	0.9458
SVM	89.81	0.9056	0.9224	0.9139
GB	95.71	0.9647	0.9441	0.9780
MLP	94.46	0.9382	0.9672	0.9525

between the pro- and anti-mask tweets, with anger and aggressiveness being the most dominant emotion in both, as illustrated in the two examples below:

Table 3: Importance scores for feature types

Feature Type	Importance Score
TF-IDF	0.4666
Embeddings	0.2598
Social Features	0.1398
POS Tags	0.1036
Sentiment	0.0310
Auxiliary	0.000

If masks are so effective then why did the mandatory rule not apply to shop staff? So COVID will kill me, the customer but not the shop worker? The insanity is breathtaking in its stupidity, incomprehension and indefensibility. #NoMasks (**P**)

#WearADamnMask with over 140k #COVID deaths Passengers cheer as 'Karen' is kicked off flight for refusing to wear mask https://t.co/wRd0iaJ1WF via @nypost. (**P**)

The first tweet is anti-mask and expresses anger towards the hypocrisy surrounding the use of masks, and the second tweet is pro-mask expressing anger and cynicism towards those who choose not to wear masks. Because anger was the dominant sentiment in both pro-mask and anti-mask tweets, sentiment scores may not have been effective in the classification. There is a good degree of sarcasm and irony expressed on both sides as well, the first tweet above notes the hypocrisy of imposing a mask mandate on the customers but not the staff. Detecting of emotions [38] could shed light on which emotions are expressed in the data set.

7 Related Research

Social media conversations are spontaneous and unfiltered, and hence, can offer genuine insights into people's opinions on a variety of offline events, topics, and policies. Because the donning of masks is relatively recent and controversial, efforts that have analyzed social media conversations around masks are gaining prominence. Ahmed *et. al.* [4] build a network of users from mask-related conversations on Twitter, and analyze this network using centrality measures to find the most influential users. Even when face masks were recommended, there remained widespread confusion about who should be wearing a mask – whether healthy people should be wearing it, and for whose protection [47]. A geographical analysis of anti- mask activity based on Twitter content has been conducted [48]. Lang *et al.* [29] examine the uses of pro- and anti-mask hashtags, and find that an increase in the volume of these hashtags is correlated with an increase in the cases. The further classify pro-mask hashtags into those urging the use of masks and issuance of mask mandates, and assertions of the efficacy, altruistic value, and positive masculinity associated with mask wearing. Anti-mask hashtags are further sorted into rejection of mask wearing, insults to mask wearers, and disinformation that asserts the negative effects of mask wearing. Al-Ramahi *et. al.* [5] also find the volume of anti-mask hashtags correlated with the volume of Covid cases. They identify three themes in anti-mask tweets, namely, constitutional rights, conspiracy theories, and fake news, pandemic, and data. Pascual-Ferra [41] analyze the toxic speech in the mask debate, and find that the tweets that included anti-mask hashtags were more likely than tweets with pro-mask hashtags to contain toxic language. He *et. al.* [22] understand the common attitudes and reasons for resistance towards the wearing of masks, and corroborate some of the reasons for the opposition as found by Al-Ramahi [5]. They use supervised machine learning to separate tweets that are relevant to the wearing of masks, and further filter those that do not express any personal opinions.

Our work can be distinguished from these contemporary efforts in that it tries to automatically separate anti-mask tweets from pro-mask tweets using machine learning. As is shown in these efforts, anti-mask tweets pedal misinformation and incite anger and hatred against government restrictions designed to curb the spread of the virus. Misinformation may discourage the adoption of this common-sense public health measure, whereas, provoking people may ultimately lead to violence and bloodshed in the physical world. Identifying anti-mask tweets may prevent the damage that they may cause, but due to the sheer volume of content shared on social media platforms, it is impossible to do so manually. Our approach is therefore valuable to automatically separate and tag such anti-mask tweets.

Overall, social media feeds have been mined to understand the public outlook on hot button medical and other health-related issues, the most notable topic that is related to masks is vaccines. The issue of masks and vaccines are inextricably linked together in the Covid world, especially, because it is believed that there is a significant overlap between anti-vaxxers and anti-maskers. Therefore, we also review the work on identifying anti-vaxx dialogue on social media as closely related to this work. Research at the intersection of vaccines and social media use both unsupervised and supervised learning for harnessing informal opinions, and also classify these perceptions into support or opposition. Some works consider specific vaccines such as Dengavaxia [2], MMR [3], Flu [8], and Zika [18], while some mine general attitudes about vaccines (anti-vaxx opinions, adversity and safety signals, fake news and rumors and interference from trolls) without reference to any particular vaccine [21, 37, 6, 35, 17, 53], and recently the Covid-19 vaccine [46, 42].

8 Conclusions and Future Research

This paper analyzes the debate around masks on Twitter using the tweets collected during the months of July and August 2020, just as many states were beginning to lift their stay home, stay safe orders, and plans were being conceived to reopen schools. Our initial analysis mines the opinions of anti-mask and pro-mask groups and compares their social features. A classification framework is then built which can differentiate between the two groups of tweets with an accuracy over 90%. Our research reveals that concerns of anti-maskers are mostly centered around politics and ideology, rather than on pragmatic issues of convenience and health. The anti-mask group is small, close-knit and supportive of each other's opinions, and hence, may be surprisingly effective at spreading the anti-mask hysteria. The benefits for public health, advocated by pro-maskers is unlikely to convince the politically motivated anti-maskers to change their views and habits. The classification framework, by the virtue of separating anti-mask tweets from pro-mask ones accurately can label tweets that sow such incorrect information about masks. Such labeling can warn other users that the views promoted by these tweets are not mainstream, and detrimental to public health.

Longitudinal analysis of the mask dialogue, with data collected at several other points during the pandemic, especially after President Trump was hospitalized due to Covid-19 is a topic of the future. A detailed topic modeling [20] framework to discover both the pro- and anti-mask themes, similar to pro-vaxx and anti-vaxx themes is also underway. Finally, collecting data from other social media platforms such as Facebook, and incorporating it in the analysis is also ongoing.

References

[1] "sklearn.preprocessing.MinMaxScaler,".https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.MinMaxScaler.html, 2007.

[2] A. B. C. Abrigo and M. R. J. E. Estuar, "A Comparative Analysis of N-Gram Deep Neural Network Approach to Classifying Human Perception on Dengavaxia," *Proc. of 2nd Intl. Conf. on Information and Computer Technologies*, pp. 46-51, 2019.

[3] A. Addawood, "Usage of Scientific References in MMR Vaccination Debates on Twitter," *Proc. of IEEE/ACM Intl. Conf. on Advances in Social Network Analysis and Mining (ASONAM)*, pp. 971-979, 2018.

[4] W. Ahmed, J. Vidal-Alaball, F. L. Segui, and P. A. Moreno-Sanchez, "A Social Network Analysis of Tweets Related to Masks during the Covid-19 Pandemic," *Intl. Journal on Environ Res Public Health*, 17(21):8235, doi: 10.3390/ijerph17218235, November 2020.

[5] M. Al-Ramahi, A. Elnoshokaty, O. El-Gayae, T. Nasralah, and A. Wahbeh, "Public Discourse Against Masks in the COVID-19 Era: Infodemiology Study of Twitter Data," *MIR Public Health Surveillance*, 7(4):e26780, doi: 10.2196/26780, April 2021.

M. Asghari, D. Sierra-Sosa, and A. Elmagharby, "Trends on Health in Social Media: Analysis using Twitter Topic Modeling," *Proc. of Intl. Symposium on Signal Processing and Information Technology*, pp. 558-563, December 2018.

[6] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. VanderPlas, A. Joly, B. Holt, and Gäel Varoquaux, "API Design for Machine Learning Software: Experiences from the Scikit-Learn Project," *Proc. of ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, pp. 108-122, 2013.

[7] A. Celesti, A. Galletta, F. Celesti, M. Fazio, and M. Villari, "Using Machine Learning to Study Flu Vaccines Opinions of Twitter Users," *Proc. of Fourth Intl. Workshop on ICT Solutions for Health*, pp. 1103-1106, 2019.

[8] K. K. Cheng, T. H. Lam, and C. C. Leung, "Wearing Face Masks in the Community during the Covid-19 Pandemic: Altruism and Solidarity," *Lancet*, April 2020.

[9] Francois Chollet, Keras blog, (Keras is a deep learning library), 2015.

[10] W. H. Delashmit and M. T. Manry, "Recent Developments in Multilayer Perceptron Neural Networks," *Proc. of the 7th Annual Memphis Area Engineering and Science Conference*, 2005.

[11] C. Dotto and L. Morrish. "Coronavirus: How Pro-Mask Posts Boost the Anti-Mask Movement," https://firstdraftnews.org/latest/coronavirus-how-pro-mask-posts-boost-the-anti-mask-movement/, August 2020. Accessed: 2021-01-21.

[12] Centers for Disease Control and Prevention, "CDC Calls on Americans to Wear Masks to Prevent Covid-19 Spread," https://www.cdc.gov/media/releases/2020/p0714-americans-to-wear-masks.html, July 2020. Accessed: 2021-01-21.

[13] E. Forslid and N. Wiken, *Automatic Irony- and Sarcasm Detection in Social Media*, Master's Thesis, Uppsala Univeritet, August 2015.

[14] S. Frenkel, B. Decker, and D. Alba, "How the 'Plandemic' Movie and its False-hoods Spread Widely Online," https://www.nytimes.com/2020/05/20/technology/plandemic-movie-youtube-facebook-coronavirus.html, May 2020. Accessed: 2021-01-21.

[15] J. H. Friedman, "Stochastic Gradient Boosting," *Computational Statistics & Data Analysis*, 38(4):367-378, 2002.

[16] B. Gambini, "Study: Fake Russian Twitter Accounts Politicized Discourse About Vaccines," http://www.buffalo.edu/news/releases/2020/03/024.html, March 2020. Accessed: 2020-06-28.

[17] A. Ghenai and Y. Mejova, "Catching Zika Fever: Application of Crowdsourcing and Machine Learning for Tracking Health Misinformation on Twitter," *Proc. of Intl. Conf. on Healthcare Informatics*, 2017.

[18] M. Godoy, "Yes, Wearing Masks Helps, Here's Why," https://www.npr.org/sections/health-shots/2020/06/21/880832213/, June 2020. Accessed: 2021-01-21.

[19] S. Gokhale, "Monitoring the Perception of Covid-19 Vaccine using Topic Models," *Proc. of the 13th IEEE Intl. Conf. on Social Computing and Networking*, Virtual,

Due to Covid-19, December 2020.

[20] S. K. Habibabadi and P. D. Haghighi, "Topic Modelling for Identification of Vaccine Reactions in Twitter," *Proc. of ACSW*, Sydney, Australia, January 2019.

[21] L. He, C. He, T. L. Reynolds, Q. Bai, Y. Huang, C. Li, K. Zheng, and Y. Chen, "Why Do People Oppose Mask Wearing? A Comprehensive Analysis of U.S. Tweets During COVID-19 Pandemic," *Journal of the American Medical Informatics Association*, 28(7):1564-1573, July 2021.

[22] Healthline.com, "As Many As 130,000 Lives Could Be Saved the Next 3 Months If Everyone Wore a Mask," https://www.healthline.com/health-news/as-many-as-130000-lives-could-be-saved-the-next-3-months-if-every one-wore-a-mask, November 2020. Accessed: 2021-01-21.

[23] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, 9(8):1735-1780, November 1997.

[24] C. J. Hutto and C. Gilbert, "Vader: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text," *Proc. of Intl. AAAI Conference on Weblogs and Social Media*, 2014.

[25] E. Israfilova, A. Arslan, N. Yildrim, and T. Kaya, "Influencer Identification System Design Using Machine Learning Techniques," C. Kahraman, S. Cebi, S. Cevik Onar, B. Oztaysi, A. C. Tolga, and I. U. Sari, Editors, *Intelligent and Fuzzy Techniques in Big Data Analytics and Decision Making*. Springer, January 2021.

[26] D. Karani, "Introduction to Word Embedding and Word2Vec," https://towardsdatascience.com/introduction -to-word-embedding-and-word2vec, September 2018.

[27] M. Kearney, "Collecting Twitter Data," https://cran.r-project.org/web/packages/rtweet/rtweet.pdf, January 2020.

[28] J. Lang, W. W. Erickson, and Z. Jing-Schmidt, "#maskon! #maskoff! Digital Polarization of Mask-Wearing in the United States During COVID-19," *Public Library of Science ONE*, 16(4):e0250817, doi: 10.1371/journal.pone.0250817. eCollection, April 2021.

[29] A. Liaw and M. Wiener, "Classification and Regression by Random Forest," *R news*, 2(3):18-22, 2002.

[30] E. Loper and S. Bird, "NLTK: The Natural Language Toolkit," *CoRR*, cs.CL/0205028, 2002.

[31] S. Loria, P. Keen, M. Honnibal, R. Yankovsky, D. Karesh, E. Dempsey, W. Childs, J. Schnurr, A. Qalieh, L. Ragnarsson, J. Coe, A. L. Calvo, N. Kulshrestha, J. Eslava, *et al*., "Textblob: Simplified Text Processing," *Secondary TextBlob: Simplified Text Processing*, 3, 2014.

[32] V. Luckerson, "How Mask Mandates were Beaten Down in Rural Oklahoma," https://www.newyorker.com/news/dispatch/how-mask-mandates-were-beaten-down-in-rural-Oklahoma, May 2020. Accessed: 2021-01-21.

[33] S. MacAvaney, H. R. Yao, E. Yang, K. Russell, N. Goharian, and O. Frieder, "Hate Speech Detection: Challenges and Solutions," *PLOS ONE*, August 2019.

[34] R. Mahajan, W. Romine, M. Miller, and T. Banerjee, "Analyzing Public Outlook towards Vaccination using Twitter," *Proc. of Intl. Conf. on Big Data*, pages 2763-2772, December 2019.

[35] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed Representations of Words and Phrases and Their Compositionality," *Neural Information Processing Systems*, 2013.

[36] T. Mitra, S. Counts, and J. W. Pennebaker, "Understanding Anti-Vaccination Attitudes in Social Media," *Proc. of the Tenth Intl. AAAI Conf. on Web and Social Media*, pp. 269–278, 2016.

[37] A. Mondal and S. Gokhale, "Mining Emotions on Plutchik's Wheel," *Proc. of Intl. Workshop on Sentiment Analysis and Mining of Social Networks*, Virtual, Due to Covid-19, December 2020.

[38] S. Morrison, "Biden Wants a National Mask Mandate, Can He Do That?" https://www.vox.com/2020/8/21/21395570/biden-mask-mandate-for-all-national-states, August 2020. Accessed: 2021-01-21.

[39] C. Nicholson, "A Beginner's Guide to Word2Vec and Neural Word Embeddings," https://pathmind.com/wiki/word2vec, 2019.

[40] P. Pascual-Ferra, N. Alperstein, D. J. Barnett, and R. N. Rimal, "Toxicity and Verbal Agression on Social Media: Polarized Discourse on Wearing Face Masks During the COVID-19 Pandemic," *Big Data & Society*, pp. 1-17, January-July 2021.

[41] N. Paul and S. Gokhale, "Analysis and Classification of Vaccine Dialogue in the Coronavirus Era," *Proc. of IEEE Big Data Workshop on Smart and Connected Communities*, Virtual, Due to Covid-19, December 2020.

[42] S. Perry, "Social Media Users are More Likely to Believe False Information About Covid-19 and to Ignore Public Health Advice, Study Suggests," https://www.minnpost.com/second-opinion/2020/07/social-media-users-are-more-likely-to-believe-false-infor-mation-about-covid-19-and-to-ignore-public-health-advice-study-suggests/, July 2020. Accessed: 2021-01-21.

[43] The Associated Press, "Can I Stop Wearing a Mask After Getting a Covid-19 Vaccine," https://apnews.com/article/mask-wearing-after-coronavirus-vaccine-f69b720444bd08565bf1f3a42e4a24ef, December 2020. Accessed: 2021-01-21.

[44] R. Řehůřek and P. Sojka, "Software Framework for Topic Modelling with Large Corpora," *Proc. of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pp. 45-50, May 2010.

[45] R. F. Sear, N. Velasquez, R. Leahy, N. Johnson Restrepo, S. E. Oud, N. Gabriel, Y. Lupu, and N. F. Johnson, "Quantifying Covid-19 Content in the Online Health Opinion War using Machine Learning," *IEEE Access*, 8:91886-91893, May 2020.

[46] Y. Shen, "Covid-19 Outbreak: Tweet Analysis on Face Masks," https://towardsdatascience.com/covid-19-out break-tweet-analysis-on-face-masks-27ef5db199dd, March 2020. Accessed: 2021-01-21.

[47] Knau Staff, "Twitter Analysis Shows Arizona is #1 in Anti-Face Mask Activity," https://www.knau.org/post/twitter-analysis-shows-arizona-1-anti-face-mask-activity, July 2020. Accessed: 2021-01-21.

[48] J. Suykens and J. Vandewalle, "Least Squares Support Vector Machine Classifiers," *Neural Processing Letters*, 9(3):293-300, 1999.

[49] A. Taylor, "How the Split Over Masks Sums Up America's Chaotic Coronavirus Response," https://www.washingtonpost.com/world/2020/06/25/face-masks-america-divided/, June 2020. Accessed: 2021-01-21.

[50] A. Whalen, "Narcissists and Psychopaths are More Likely to Refuse to Wear Masks, Says New Research," https://www.newsweek.com/narcissists-psychopaths-face-mask-requirement-mandate-social-distancing-covid-19-coronavirus-1519732, July 2020. Accessed: 2021-01-21.

[51] R. Zafarani, M. A. Abbasi, and H. Liu, *Social Media Mining: An Introduction*, Cambridge University Press, New York, NY, USA, 2014.

[52] W. Zhao, "Misinformation Correction Across Social Media Platforms," *Proc. of Intl. Conf. on Computational Science and Computational Intelligence*, pp. 1371-1376, 2019.

Luca Cerbin (photo not available) is a senior at Swarthmore college. Luca is majoring in Computer Science with a double minor in Applied Math and English. He works in a biology lab on image processing and machine learning on planarians. His research interests are in machine learning and understanding the limitations associated with it. This work was done when he was a visiting student at the Univ. of Connecticut.

Jason DeJesus (photo not available) is a graduate student in the Electrical and Computer Engineering Department at the University of Connecticut. He works as a Graduate Assistant in a cybersecurity lab, with a focus on hardware security.

Julia Warnken (photo not available) is pursuing her Masters of Engineering in Computer Science and Engineering at the University of Connecticut. Her research interests include machine learning, data mining, and genetic or biomedical applications.

Swapna S. Gokhale (photo not available) is an Associate Professor of Computer Science and Engineering at the University of Connecticut. She received her PhD in Electrical and Computer Engineering from Duke University in 1998. Her research interests lie in the areas of performance and dependability analysis of computer systems, software engineering education, and mining and analysis of social media, epidemiological and sustainability data. She has published over 150 conference and journal papers on these topics. She is a recipient of a NSF CAREER award and a Senior Member of the IEEE.