Investigating Influential COVID-19 Perspectives: A Multifaceted Analysis of Twitter Discourse

Shadaab Kawnain Bashir[1], Hossein Shirazi[2], Noushin Salek Faramarzi[3], Thomas Harris[1], Ashmita Shishodia[2], Hajar Homayouni[2], and Indrakshi Ray[1]

[1] Colorado State University, Fort Collins, CO, USA
{shadaab,indrakshi.ray}@colostate.edu,tvharris123@gmail.com
[2] San Diego State University, San Diego, CA, USA
{hshirazi,ashishodia2546,hhomayouni}@sdsu.edu
[3] Stony Brook University, Stony Brook, NY, USA
nsalekfarama@cs.stonybrook.edu

Abstract. Social media influencers, those with verified accounts or with more than 10,000 followers, played a crucial role in the propagation of narratives during the COVID-19 pandemic. We investigate their impact by characterizing and contrasting the differences in content patterns between influential individuals versus public organizations during the pandemic, analyzing emotions, sentiments, and scientific claims expressed in their Tweets. Advanced machine learning approaches, including customized transformer models, few-shot learning, and large language models such as GPT-3.5, were used. The findings reveal a stark contrast in sentiment usage across sub-domains like vaccines and lockdowns, with organizations predominantly employing neutral tones while individuals displaying a significant negative sentiment bias. Individuals often conveyed more negative emotions, whereas organizations exhibited greater optimism. However, many claims from both groups were not verified, highlighting the need to combat misinformation.

Keywords: COVID-19 · Emotion Detection · Sentiment Analysis · Scientific Claim Identification · User Profiling.

1 Introduction

The emergence and pervasiveness of social media have transformed the way information is disseminated and consumed globally [18]. Particularly, in the context of the COVID-19 pandemic, social networks have played a vital role in informing public sentiment and influencing behaviors [8]. It has acted as an indispensable conduit for disseminating health advisories and updates on the virus's progress. Twitter [4] has become a hotbed for spreading rumors and false claims about the virus, exacerbating public confusion and anxiety. However, the rapid spread of

[4] Now known as X, we collected and executed the experiments while it was known as Twitter.

information and the omnipresent nature of the platform has also fueled misinformation, creating a "digital infodemic" parallel to the pandemic [29] and making it difficult for the public to make informed decisions [14]. Influencers, defined by individuals or organizations with a substantial following base, can amplify messages, influence opinions, and mold social narratives [13]. The sharing of questionable information by influential figures on Twitter has played a substantial role in shaping misconceptions and distorting public understanding of the pandemic.

While the influence of social media is undeniable, there exists a nuanced divergence in the approach and impact of individual versus organizational influencers. Individual influencers, comprising celebrities, thought leaders, and other public figures, often present personal and subjective points of view. Their content tends to reflect personal opinions, experiences, or endorsements, which can vary widely in terms of accuracy and reliability [4]. On the other hand, organizational influencers, such as official accounts of entities like the World Health Organization (WHO), focus on disseminating verified information and official updates. Their content is typically more factual, objective, and aligned with public health guidelines [7]. This distinction is critical in understanding the landscape of social media discourse during the pandemic, as the type of influencer can significantly influence the nature and impact of the information being disseminated.

The current situation prompts a study of how individual and organizational influencers uniquely shape the social media landscape during the COVID-19 pandemic. There exists a notable research void in thoroughly examining and contrasting the content, sentiment, and impact of these distinct influencer categories. Understanding this discrepancy is essential to understanding the broader effects of social networks on public opinions and behaviors during health crises. Initial studies have delved into characteristics of COVID-19 Twitter data, tracking the emotional evolution expressed in Tweets during the pandemic by jointly analyzing both the types of emotion and the overall polarity of sentiment [20, 24, 33, 41, 43]. However, as far as our knowledge extends, existing studies have not undertaken a thorough analysis of the distinctions between organizational and individual influencers for COVID-19-related Tweets.

This research categorizes influencers into individual and organizational entities, examining content, sentiment, and emotional responses. It also evaluates the influence of these Tweets on public sentiment and trust in scientific information. This will help in understanding influencers' roles and communication strategies during the pandemic. In this study, we examine the following research questions.

- **RQ1:** *How do individual influencers and organizational influencers express their opinions and communicate their perspectives about the COVID-19 pandemic on Twitter and what are the differences?* This inquiry aims to investigate the differences between individuals and organizations in terms of the emotional appeals they use, the overall sentiment they portray, and the degree to which they rely on scientific evidence.

- **RQ2:** *To what extent deep learning algorithms can effectively profile individuals' and organizations' accounts on social media?* This question focuses on the identification and categorization of influencers in the dataset as either individuals or organizations, based on their characteristics and communication styles using deep learning algorithms.

To bridge the gap between these research questions and actionable analysis, four specific tasks were designed: *user profiling*, *emotion detection*, *sentiment analysis*, and *scientific content identification*. We need these tasks to understand the nuanced dynamics of Twitter discourse and also to leverage computational methods to systematically categorize and evaluate the content. User profiling is essential for distinguishing between individual users and organizations, facilitating an analysis that considers the source of each Tweet. This differentiation is crucial for RQ1, as it allows for an investigation into the unique communication styles and content preferences of these two groups. Emotion detection and sentiment analysis further enrich this understanding by quantifying the emotional and sentimental dispositions conveyed in Tweets, enabling a nuanced analysis of the emotional appeals and overall sentiment portrayed by different users. These tasks directly support the exploration of how opinions and perspectives are expressed, as outlined in RQ1. Finally, the scientific content identification task aligns closely with both RQ1 and RQ2 by evaluating the extent to which Tweets rely on scientific evidence, a factor that can significantly influence the credibility and reception of the information shared. Together, these tasks form a comprehensive framework for addressing the posed research questions, leveraging deep learning and computational linguistics to uncover insights into the digital discourse surrounding the COVID-19 pandemic on Twitter.

Due to the lack of a ground-truth labeled dataset, we developed a labeled dataset by collecting 1875 influential Twitter posts related to COVID-19 and annotating them using a multi-step pipeline. Each Tweet was manually labeled by human annotators for user profile (*individual* or *organization*), emotion (*joy, sadness, fear, surprise, anger*), sentiment (*positive, negative,* or *neutral*), and presence of scientific claims (*scientifically verifiable claims, scientific knowledge reference, general scientific research*) as illustrated in Fig. 1.

In our study, we used an automated approach to analyze the annotated dataset using NLP algorithms. The core of our methodology involved applying Bidirectional Transformer Encoder Representations (BERT) [9] for general text encoding, providing crucial foundational contextual embeddings for nuanced language interpretation. To address the specificity of scientific content, we integrated SciBERT, a BERT variant trained on scientific corpora, enhancing the accuracy of scientific claim detection. Furthermore, we used large language models like LLAMA2 and GPT-3.5 for their conversational understanding capabilities, helping to analyze emotions and sentiments. These models were instrumental in interpreting communicative intentions and subtleties within Tweets. This multifaceted NLP framework allowed for a comprehensive analysis of influencer categorization, sentiment, and scientific claim detection in the context of COVID-19, leveraging the strengths of each model to fulfill our specific research goals.

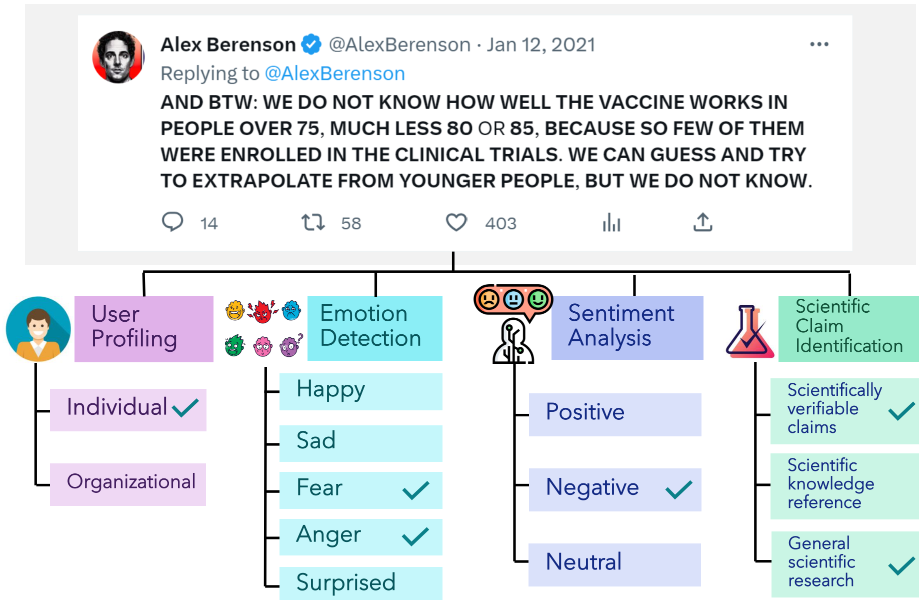

Fig. 1: Example Annotated Tweet From Our Annotated Dataset Demonstrating Multilayer Analysis Across Dimensions of User Profiling, Emotion Detection, Sentiment Analysis, and Scientific Claim Identification.

Our Contribution

- We created a unique dataset of manually annotated influential COVID-19 Tweets labeled for user profiles, emotions, sentiments, and scientific claims.
- We leveraged this new dataset in conjunction with advanced machine learning approaches including customized transformer models, few-shot learning, and Large Language Models. The GPT-3.5 model showcased impressive F1 scores, excelling with 0.91 for sadness, 0.89 for surprise, and SetFit performed best for fear with an F1 of 0.85, 0.91 for anger and 0.92 for joy. In sentiment analysis, RoBERTa$_{\text{Twitter}}$ emerged as a standout performer, achieving leading F1 scores of 0.72 for negative, 0.89 for neutral, and 0.78 for positive sentiment. Domain-specific language models like SciBERT demonstrated notable proficiency, with F1 scores of 0.87 for references, and 0.92 for general scientific claims and SetFit achieved an F1 score of 0.97 for verifiable claims.
- Key findings reveal individuals exhibit more subjective sentiment compared to organizations' impartiality; however, organizations disseminate higher overall volumes of pandemic-related content. Across both groups, many assertions lack rigorous verification, highlighting the imperative for validation to combat misinformation. Additionally, we observed a marked polarity in individual Tweets (positive or negative sentiments), in contrast to the more neutral stance of organizations. The sheer volume of organizational Tweets, however, outweighs individual Tweets. By empowering nuanced investigation

across these facets, this research significantly enhances the ability to decode complex perspectives and narratives that shape the discourse on social media crises.

– Our study advances beyond traditional sentiment analysis and topic modeling by integrating a multi-dimensional approach. This includes user profiling, emotion detection, sentiment analysis, and scientific claim identification, providing a holistic understanding of COVID-19-related discourse on social media.

The rest of this paper is structured as follows: Section 2 reviews relevant literature. Section 3 describes the data collection methodology and the annotation process. Sections 4, 5, 6, 7 present the experimental approach and results for each dimension of the analysis: user profiling, emotion detection, sentiment analysis, and identification of scientific claims respectively. Finally, Section 8 concludes the paper.

2 Related Work

Related work pertaining to user profiling, emotion detection, sentiment analysis, and scientific claim identification are described below.

2.1 User Profiling

User profiling involves creating detailed profiles of users based on their social media activity and behavior. Several related works have addressed various aspects of user profiling, including the identification of individuals and organizations and the clustering of Twitter users.

Kappus *et al.* [19] leveraged network analysis and hashtag clustering to group Twitter users based on their pandemic-related Tweets. Their study revealed highly insular communities on the platform. Furthermore, Egger and Yu [11] evaluated topic modeling approaches for COVID-19 Tweets related to cluster travel, demonstrating the effectiveness of BERTopic. However, in contrast to their approach of analyzing individual Tweets, our work aims to utilize user profile attributes rather than focusing solely on their Tweets. Our research extends existing methodologies by leveraging the power of transformer models, specifically BERT and BERTopic, to profile Twitter users. By examining their overall behavior and engagement patterns, we gain a deeper understanding of users' interests, preferences, and potential connections within the Twitter ecosystem.

One notable work is the development of the "demographer" Python package [39], which focuses on identifying gender and ethnicity, and distinguishes between accounts belonging to individuals or organizations. Another relevant study by Liang *et al.* [23] introduced DUWE, a dynamic user profiling model that tracks the interests of Twitter users over time using their Tweet history.

2.2 Emotion Detection & Sentiment Analysis

The COVID-19 pandemic has led to several studies on the emotional impact of the pandemic on individuals. Prior studies [1, 26, 30] have employed natural language processing techniques to analyze emotions and sentiments expressed in COVID-19-related Twitter data. While sentiment and emotion are commonly examined in tandem, it is noteworthy that there exist research endeavors dedicated to investigating these phenomena individually. For example, Nandal *et al.* [27] leveraged lexicon-based approaches to identify specific emotions like fear, sadness, and anger in Tweets. Similarly, Manguri *et al.* [25] used TextBlob for analysis and observed that over 50% of the Tweets showed neutral sentiment. Ainapure *et al.* [2] categorized Tweet sentiment as positive, negative, or neutral using VADER sentiment analysis. Yu *et al.* [42] found increased engagement with negative sentiment about COVID-19. Wrycza and Jacek [40] analyzed the topic of working from home which highlighted a predominantly positive response.

Several works have tracked emotional evolution over the timeline of the pandemic. Jalil *et al.* [17] revealed fluctuating positive and negative sentiments towards vaccines using supervised learning classifiers. Storey and O'Leary [33] demonstrated a shift from predominantly negative linguistic sentiments in early 2020 to more positive expressions by 2021 using the NRC Emotion Lexicon.

Previous research [6, 10, 36] focused on sentiment and emotion recognition utilizing lexicon-based Python tools. Our technique, on the other hand, employs transformers and LLMs to efficiently solve both sentiment and emotion detection tasks. Our study specifically aims to analyze the sentiment conveyed within Twitter, employing gold-standard data for this analysis.

2.3 Scientific Claim Identification

Recent works have sought to develop methods for the identification of scientific claims, with a particular focus on claims related to COVID-19. Saakyan *et al.* [32] presented COVID-Fact, a system designed specifically for verifying COVID-19 claims using the CORD-19 corpus. Wadden *et al.* [37], on the other hand, proposed a general framework for scientific claim identification using semantics to model claims and evidence. While promising, current scientific claim identification methods have limitations to address moving forward. Wadden *et al.* [38] improved claim identification in LONGCHECKER by considering full abstract context rather than isolated sentences. However, as noted by Landers *et al.* [22], existing frameworks still face challenges with scientific terminology and reasoning.

The SciTweets annotation framework by Hafid *et al.* [16] provides a valuable foundation for our scientific claim identification task. Their taxonomy categorizes Tweets as direct scientific claims, references to sources, or general terminology. We adopt this multi-label approach in our COVID-19 Twitter dataset, annotating for scientific claims, sources, and terms.

While previous works have predominantly concentrated on general Tweets across the platform, our research places emphasis on COVID-19 Tweets. In contrast to most prior studies [32,37,38] that rely on transformer based models such as BERT and its variations, our work encompasses a diverse set of models by leveraging LLMs in addition to transformers.

3 Data Collection & Annotation

The data collection from Twitter was tailored to capture a diverse array of discussions related to COVID-19 which was subsequently annotated.

3.1 Data Preparation

We used a dataset created by Zuo *et al.* [44] that contains more than 447 million Tweets collected from September 2020 to October 2021.
Influencers Accounts. We defined *influencers* as individuals and organizations that have verified accounts or those who have accounts with more than 10,000 followers.
Data Sampling. To create manageable samples for annotation and analysis, we employed iterative, targeted sampling strategies. After initial random sampling, we used keyword filtering to focus on Tweets about key topics like "face mask", "vaccine", and "quarantine", compiling corpus subsets specific to different analysis tasks. Using BERT and cosine similarity, we removed semantically similar Tweets to avoid annotation overlap. Through this multi-stage sampling approach, we arrived at final samples of 1875 Tweets only from influencers' accounts.

3.2 Manual Data Annotation

The study used a manual annotation process for Tweets, with three Computer Science undergraduate students proficient in English working together. For each task, annotators were given clear and step-by-step instructions. Each Tweet was randomly assigned to two, with a third adjudicated if needed. The consistency of annotations was evaluated using an inter-annotator agreement score.
User Profiling. To classify influencers into two distinct groups of *individuals* and *organizations*, we extracted user biographies and usernames associated with each Tweet. Subsequently, we tasked our annotators with determining the classification of each account into either group. Individual profiles are identified by personal narratives, experiences, milestones, and interests, while organizational accounts focus on professional accomplishments, formal tone, and references to collective entities or roles within organizations. The analysis yielded a Cohen's Kappa score of 0.96, indicating near-perfect agreement and reliability among annotators.
Emotion Detection. Emotion detection, aimed to identify the emotions expressed in Tweets. Emotion detection is split into two sub-tasks: whether a Tweet contains emotion or not, and if yes, what is the type of emotion? Ekman *et al.* [12] categorized human emotions into anger, surprise, disgust, enjoyment, fear, and sadness. We noticed many occurrences within these categories, except for disgust, where there were fewer samples available. As a result, we adjusted

our framework to cover five primary emotional categories: *joy, sadness, surprise, fear,* and *anger* for investigating emotions expressed in Tweets. The consistency of the annotation task was assessed using a Jaccard Index score of 0.47 which represents moderate agreement among annotators for multi-categorical annotations. Also, to address the challenges of annotating emotions in a large dataset, we employed text augmentation techniques. We used back-translation, translating Tweets into foreign languages and back into English. This approach resulted in an augmented dataset with over 11,000 Tweets.

Sentiment Analysis. Our team of annotators analyzed each Tweet from the author's perspective to identify clear expressions of sentiment polarity. They classified Tweets into three categories: *positive* (hope, support, optimism), *negative* (fear, frustration, sadness), and *neutral* (no emotional expressions). We categorize Tweets based on the author's intended emotion rather than solely isolated linguistic components. For instance, we labeled the following Tweet as neutral in sentiment:

> *"We have jobs at serious risk NOW* ! *67 of our industry cannot survive another lockdown*! *That is 32k high at salons amp*; *251k jobs*! *amp*; *21 billion lost to gdp So businesses are on their knees trading at 40*! *So please listen uk."*

Despite negative phrasing like "at serious risk" and "cannot survive", our annotation process deemed the Tweeter's perspective as more objective than emotional. This exemplifies why we categorize based on holistic authorial intent rather than isolated linguistic components. By annotating based on intended emotion rather than isolated syntax, our methodology enables richer sentiment analysis. The Cohen's Kappa value of 0.46 indicates moderate alignment in sentiment annotation among our annotators. This level of agreement is considered reasonable, given that sentiment analysis involves a certain degree of subjectivity as a result of the inherent complexities of human emotions and language.

Scientific Claim Identification. To identify the type of scientific information present in Tweets, we adopted the annotation scheme established in prior work by Hafid *et al.* scheme [16] to identify and label Tweets for (i) Scientific claims or assertions of direct knowledge about COVID-19, (ii) References or citations to external scientific sources such as journals, reports, or news articles related to COVID-19, (iii) General scientific lexicon (e.g. "study", "data", "researchers") without specific factual claims. For each Tweet, our annotators assigned binary indicator values denoting the presence or absence of each scientific content type. We found that Tweets can display multiple categories simultaneously, and a score of 0.88 was obtained for Krippendorff's alpha, indicating strong agreement considering multi-categorical annotations.

Table 1 and Fig. 2 demonstrate the distribution of labels in different tasks in our dataset. While individual and organizational accounts are nearly equally represented, individuals exhibit higher rates of detected emotions and subjective positive/negative sentiment compared to organizations' prevailing neutrality. This evinces a contrast between the more emotional, subjective perspectives of individuals versus the objective stances organizations maintain.

Table 1: An Overview of Annotations and Labels in the Dataset Across 4 Tasks, Including a Breakdown of [Ind]ividual and [Org]anizational Profiling.

Task	**Label**	**Total (%)**	**Ind.(%)**	**Org.(%)**
User Profiling	Individual	893 (48.0%)	–	–
	Organization	964 (52.0%)	–	–
Emotion Detection	*Emotion Existence*	709 (38.2%)	28.1	10.1
	Joy	243 (13.1%)	8.7	4.4
	Sad	138 (7.4%)	5.6	1.8
	Fear	143 (7.7%)	5.4	2.3
	Anger	278 (15.0%)	12.7	2.3
	Surprised	97 (5.2%)	3.9	1.3
Sentiment Detection	Negative	364 (19.6%)	15.3	4.3
	Neutral	1194 (64.3%)	23.1	41.2
	Positive	299 (16.1%)	9.6	6.5
Scientific Claim	*Claim Existence*	922 (49.6%)	25.1	24.5
	Verifiable	729 (39.2%)	19.9	19.3
	Reference	205 (11.0%)	5.3	5.7
	General	143 (7.7%)	4.4	3.3

Additionally, despite nearly half of Tweets containing declarative claims, most claims resist categorical verification, underscoring the need for rigorous validation to prevent misinformation, especially invoking scientific credibility. Further examination shows unreferenced claims are common among individual and organizational users. Individuals exhibit a higher percentage of general claims lacking citations. However, organizations have a slightly higher rate of referenced claims. This highlights the imperative of verifying unchecked assertions invoking scientific credibility across all user groups.

Our analysis reveals revealing contrasts in how individuals and organizations engage on Twitter. As seen in Table 1, individuals are more emotionally expressive overall, conveying greater anger while organizations project more joy. This pattern extends to sentiment as well, with organizations maintaining an impartial, neutral stance while individuals adopt more negative tones. The greater subjectivity of individuals is juxtaposed with the detached objectivity organizations pursue. Additionally, our claim analysis flags a concerning trend - a substantial number of assertions across both groups are categorized as verifi-

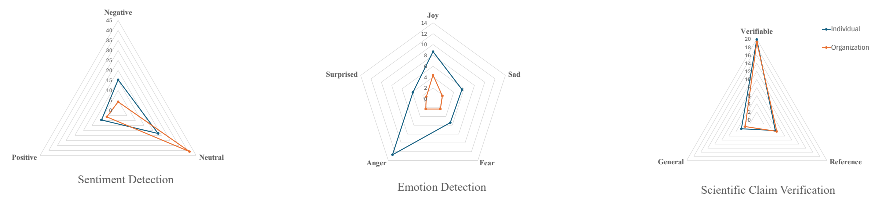

Fig. 2: Percentage Breakdown of Individual and Organizational Accounts in Sentiment Detection (Left), Emotion Detection (Middle), and Scientific Claim Identification (Right) Tasks.

able, yet lack citations or factual references. Organizations hold a slight edge in propounding these unsubstantiated but potentially verifiable claims.

4 Task 1: User Profiling

Identification of Twitter Account Ownership. Profiling the ownership of Twitter accounts allows us to distinguish between official accounts representing companies, institutions, or brands, and influencers with personal views and experiences.

To classify Twitter accounts into distinct categories of individuals and organizations, we leverage Tweet metadata. Our account identification process is based on three key features: *Twitter name*, *Twitter handle*, and *Twitter biography*. Although this approach has its limitations, such as not utilizing the actual Tweets posted by the account, it offers the benefit of requiring minimal information for the model to generate accurate predictions. Concatenating these features provides a condensed input for our BERT classifier to categorize accounts as either individual users or organizations. The data is divided into training and testing sets, with a split of 80% for training and 20% for testing. To convert the Tweet content into textual embeddings, a BERT layer is employed. Subsequently, these embeddings are fed into a single dense output layer to facilitate the classification process.

In addition to the BERT model, we also incorporated the Demographer [21] in our experiments. As an established approach for demographic inference, the Demographer extracts user metadata, linguistic cues, and network patterns to predict age, gender, and location. However, empirical results revealed BERT's overwhelming advantage. BERT outperformed Demographer in key metrics, achieving an F1 of 0.95 and an accuracy of 0.96, highlighting its proficiency in encoding semantic and contextual information from sparse text. Its self-attention mechanism can effectively model complex user attributes, despite lacking extensive feature engineering.

Topic of Occupation of the Users. User bios on Twitter offer a chance to understand the interests and occupations of the Twitter population. Occupation provides valuable insight into a user's knowledge base, credibility, authority, and

Table 2: Top 5 Topics and Key Terms Identified Through Biography Clustering of Individual and Organization Accounts.

	Identified Topics	**# Accounts**	**% Accounts**
Individual	KAG, Trump, Trump2020	15048	4.5
	18, NSFW, Onlyfans	14730	4.2
	Periodista, Politica, Cuenta	10769	3.1
	Film, Actor, Director	8977	2.6
	Resist, BLM, Bidenharris2020	5015	1.4
Organization	Marketing, Business, Digital Tech.	7985	7.1
	Content, 18, Onlyfans Promo	6112	5.4
	Sports, Basketball, Athletics	4333	3.8
	Football, Club, League	4210	3.7
	Financial, Investment, Trading	4152	3.7

sphere of influence, thereby contributing to a nuanced understanding of how Twitter influences vary across different professional segments. If an influencer's occupation is known, for instance, their Tweets might be given more weight if they are considered experts in their field, influencing not only individual perspectives but also shaping collective societal discourse. Similarly, detecting occupation can also shed light on potential bias or vested interests, providing an additional layer of contextual information that is important in analyzing the authenticity and reliability of an influencer's content. From a socio-economic standpoint, knowing a user's occupation can help discern how information and influence diffusion patterns on Twitter correspond to various job sectors, thereby enhancing our understanding of societal impact. Moreover, this knowledge could be used to optimize targeted communication strategies, drive professional engagement, and develop policy-making processes that effectively address the real-time societal implications of Twitter-based influences.

By applying clustering techniques to textual descriptions in the user's biography, we can extract the most salient topics. Our approach leverages BERTopic [15] to cluster Twitter bios. First, we categorize accounts as individuals or organizations using BERT. This allows specialized clustering based on how each group crafts bios. For individuals, we extract professions, hobbies, and interests. For organizations, we identify their core purpose.

To address the issue of BERTopic producing an excessive number of outlier results, we used a technique that reduces outliers using probabilities obtained from the HBDSCAN [5]. Twitter biographies are classified into a certain topic cluster or an outlier group using the BERTopic algorithm, then the HBDSCAN model groups the data based on the most relevant subject and determines the likelihood that a Twitter bio belongs to any other cluster. Although these probabilities were not included in the initial topic clustering procedure, we reduced the number of outlier bios by choosing the topic with the highest likelihood and allocating the outlier to that category. To evaluate our optimized BERTopic approach, we compared the silhouette and coherence metrics [31] to a baseline. Silhouette quantifies cluster cohesion and separation. Our score of 0.52 confirms distinct, tightly clustered topics. Coherence measures topic interpretability. Our high score of 0.80 indicates coherent, semantically aligned topics.

Our topic clustering analysis revealed distinct trends in interests and occupations between individuals and organizations on Twitter. Table 2 displays the top 5 topics identified for each group along with selected representative words and the number of accounts per cluster. For individuals, politically-oriented topics like "KAG, Trump, Trump2020" and "Resist, BLM, Bidenharris2020" feature prominently, underscoring Twitter's role in civic discourse. Entertainment interests like "Film, Actor, Director" also rank highly. Organizations showcase more commercial themes, including "Marketing, Business, Digital Tech" and "Financial, Investment, Trading." Sports-related accounts also comprise a major cluster, consistent with the prevalence of teams and leagues on the platform.

The bifurcation between individuals and organizations is clearly evidenced in these topic profiles. By categorizing accounts before clustering bios, we can extract key themes unique to each group's priorities and intentions on Twitter. This table summarizes our core findings on the primary spheres of interest emerging from biography text analysis.

5 Task 2: Emotion Detection

For this task, we developed a pipeline system to categorize emotions in Tweets. First, a preliminary model filters Tweets to detect any emotions, while those with detectable emotions are passed to the next phase. In the second stage, five specialized classifiers evaluate each distinct emotion, identifying *joy*, *sadness*, *surprise*, *fear*, or *anger* specifically.

To train models for our pipeline architecture, we partitioned the Tweet dataset into three subsets - 70% for training, 20% for testing, and 10% for validation. With the goal of classifying individual emotions, we implemented six distinct binary classifiers. Five of these models focus on identifying the presence of one specific emotion: joy, sadness, surprise, fear or anger. The sixth classifier serves as a preliminary filter, detecting if a Tweet contains any emotion at all before passing it to the specific emotion classifiers.

In our quest to determine the best classifier, we employed a BERT model consisting of a BERT layer followed by a dense output layer. We evaluated the performance of this model using both augmented and non-augmented training sets. Moreover, to overcome the limitations posed by limited training data, we harnessed the power of SetFit [35] as a few-shot learning technique. Additionally, we have used LLMs like GPT-3.5 [28] and LLaMA-2-70B [34] and to probe these models we have used the following prompt:

"Classify the emotion of the Tweet. Does this Tweet contain any emotion? Yes or No. Does this Tweet contain Joy emotion? Yes or No. Does this Tweet contain Sad emotion? Yes or No. Does this Tweet contain Surprise emotion? Yes or No. Does this Tweet contain Anger emotion? Yes or No. Does this Tweet contain Fear emotion? Yes or No."

This comprehensive approach enabled us to thoroughly explore and compare various classifier configurations. Table 6 presents a comparative evaluation of emotion detection performance using BERT, SetFit, GPT-3.5, and LLaMA-2-70B models. Experiments were conducted using the both original dataset and an augmented version expanded via back-translation. The results demonstrate SetFit's effectiveness as a few-shot learning technique, achieving the top F1 scores for the existence of emotion. The model maintains consistent performance across all categories as it uses sentence transformers to generate semantically meaningful sentence embedding preserving the semantic integrity of entire sentences, unlike traditional models such as BERT which uses individual tokens. BERT, on the original dataset, exhibits high variance, for the F1 score. Its performance is lowest for sadness and surprise. Augmenting the training data substantially improves BERT's performance, reducing its gap with SetFit. The augmented BERT model attains a comparable F1 score to that of SetFit for emotion existence and for specific emotions. Furthermore, GPT-3.5 achieved the highest F1 score for detecting emotions like sadness and surprise.

Analysis indicates joy is detected most accurately by all models. This emotion likely has a more recognizable linguistic pattern. Intrinsically more complex emotions like sadness and surprise pose greater challenges for the BERT-based model but the chat-based models were able to make better inferences of these emotions as it is trained on a diverse range of text data. However, the chat-based models were not able to detect the emotions from Tweets efficiently, and hence by using better prompts the performance of these models can be improved. Overall, GPT-3.5 demonstrates proficiency in multi-label Tweet emotion detection and good choice for emotion-based tasks.

6 Task 3: Sentiment Analysis

Our main focus in the analysis was on assigning sentiment labels to the Tweets, classifying them into three categories: *positive*, *negative*, and *neutral*. We performed pre-processing steps to replace emojis and remove hashtags from Tweets. Additionally, we replace usernames with "USER_HANDLE", URLs with "URL_TOKEN", and substitute special characters with NULL characters. To ensure consistency, the entire text of the Tweets was converted to lowercase. These preprocessing steps aimed to clean the data and standardize the text for further analysis. We use BERT, $\text{RoBERTa}_{\text{Twitter}}$ [3], SetFit, GPT-3.5, and LLaMA-2-70B to conduct sentiment analysis on our dataset. $\text{RoBERTa}_{\text{Twitter}}$ is a RoBERTa-base model trained on 58 million Tweets and fine-tuned for sentiment analysis using the TweetEval benchmark. The BERT model was fine-tuned by adding a single dense layer with three outputs representing negative, neutral, and positive sentiments. Additionally, SetFit was employed, which is designed to train models with limited amounts of data. For the LLMs we have used the following prompt for sentiment detection:

Classify the sentiment of the author behind the Tweet as "Negative", "Neutral", or "Positive".

Table 6 shows that our task-specific model, RoBERTa$_{\text{Twitter}}$ achieved the highest F1 scores for identifying negative sentiment, and positive sentiment, and performed competitively on neutral sentiment compared to SetFit. In contrast, LLaMa-2-70B achieved the lowest F1 score in neutral sentiment. Conversely, GPT-3.5 achieved the lowest score in negative sentiment.

Our findings demonstrate the effectiveness of using a model fine-tuned on informal, social media text for accurate classification of sentiment polarity and intensity. The social media-optimized RoBERTa architecture was best able to capture the nuances and variances in how sentiment is expressed on such platforms. This highlights the importance of using task and domain-specific models for optimized performance on sentiment analysis versus general pre-trained models like BERT, SetFit, or LLMs alone.

Nevertheless, it is worth noting that our experimentation with TextBlob and VADER lexicon-based models had subpar results in Tweet sentiment analysis due to their limitations in contextual learning. These models were less suitable for sentiment analysis than contextual models like BERT and RoBERTa. The analysis also revealed that neutral Tweets were classified most accurately, while negative Tweets posed the greatest challenges.

Table 3: COVID-19 Keywords for 6 Topics that were used for Sentiment and Scientific Claim Analysis.

Topic	Keywords
Face Mask	mask, face, N95, KN95, covering, nose wire, adjustable strap, breathable fabric
Travel	travel, border, closures, hotel, passport
WFH	wfh, home, remote, hotel, telecommuting, virtual, video, zoom
Social Distancing	social, distancing, six feet, physical, close contact
Lockdown	lockdown, stay home, curfew, restricted movement, quarantine, isolation
Vaccine	vaccine, vaccination, vaccinated, pfizer, moderna, astraZeneca, johnson, sinopharm, sputnik

Topic-Targeted Sentiment Analysis. To analyze the polarity of the sentiments for specific discussion topics on Twitter, we filtered the full dataset to extract relevant subsets of Tweets. As shown in Table 3, we focused on six key topics and curated custom datasets for each one. For example, the "Lockdown" topic, we searched for Tweets containing related terms like "stay home", "curfew". This coarse-grained, keyword-based approach produced subsets containing only relevant Tweets for each discussion theme. The sentiment expressed in these subsets was classified using our top-performing model.

Fig. 3 depicts the polarity of sentiments expressed by influencers over a 220-day period on six crucial themes. The graph reveals that influencers, including individuals and organizations, predominantly express neutral feelings for all key topics, indicating that influencers are primarily focused on distributing information to their followers during a crisis. Individuals tend to display more negative attitudes than organizations, as organizations represent a collective entity with goals, objectives, and brand images. They often maintain a professional demeanor to avoid negativity in discussions. On the other hand, individuals are more subjective towards their feelings. Furthermore, it can be noted that the topic of face masks is rated least positively by both individuals and organizations. There were instances of misinformation and conspiracy theories involving face masks throughout the pandemic and influencers may have been afraid to market face masks for fear of endorsing or spreading inaccurate information.

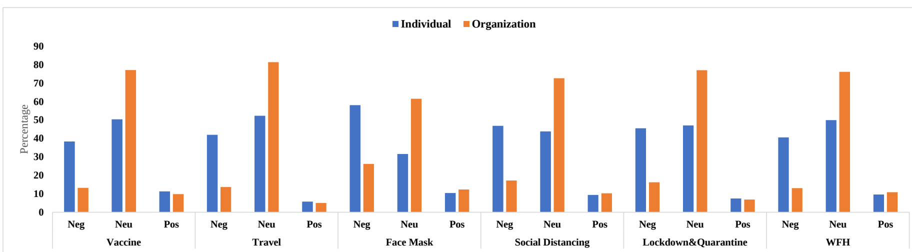

Fig. 3: Individual vs. Organization Sentiment Polarity on Significant Topics Related to COVID-19. Individuals are more subjective towards their opinions whereas organization tries to maintain a professional demeanor by posting neutral Tweets.

7 Task 4: Scientific Claim Identification

The task of identifying scientific content on Twitter involves categorizing them based on their scientific information. The annotation scheme labels Tweets for *scientifically verifiable claims*, *scientific knowledge references*, and *general scientific claim*. Each Tweet was assigned binary values representing the presence or absence of these scientific content types, and multiple categories could be present in a single Tweet.

In this work, we utilize a SciBERT-based [16] classifier for categorizing the diversity of scientific content on Twitter. We implemented a multi-label system to categorize scientific Tweet content. Tweets were assigned to one or more of the three categories (similar to our annotation). Similar to the prior tasks, we

Table 4: Comparative Performance Analysis of BERT, SetFit, Task Specific, and LLMs in Emotion Detection, Sentiment Analysis, and Scientific Claim Identification. Task Specific Models, for Emotion Detection, we employed BERT with augmented data and for Sentiment Analysis, $\text{RoBERTa}_{\text{Twitter}}$ was utilized and SciBERT was employed for Scientific Claim Identification.

Task	Label	BERT				SetFit				Task Specific Model				GPT-3.5				LLaMA-2-70B			
		P	R	A	F1	P	R	A	F1	P	R	A	F1	P	R	A	F1	P	R	A	F1
Emotion Detection	*Emotion Existence*	0.85	0.77	0.77	0.81	0.83	0.82	0.83	0.83	0.85	0.79	0.84	0.83	0.75	0.70	0.70	0.70	0.76	0.59	0.59	0.56
	Joy	0.79	0.81	0.93	0.80	0.93	0.91	0.94	0.92	0.85	0.79	0.90	0.82	0.91	0.91	0.91	0.91	0.91	0.81	0.81	0.84
	Sadness	0.67	0.61	0.89	0.63	0.76	0.81	0.86	0.78	0.74	0.63	0.89	0.67	0.91	0.93	0.93	0.92	0.93	0.66	0.66	0.75
	Surprised	0.45	0.26	0.88	0.33	0.79	0.79	0.87	0.79	0.77	0.48	0.86	0.59	0.91	0.87	0.87	0.89	0.91	0.66	0.66	0.75
	Fear	0.64	0.46	0.88	0.53	0.87	0.83	0.89	0.85	0.80	0.64	0.87	0.71	0.88	0.80	0.80	0.83	0.90	0.56	0.56	0.64
	Anger	0.84	0.86	0.93	0.84	0.90	0.91	0.91	0.91	0.84	0.92	0.89	0.88	0.89	0.89	0.89	0.89	0.89	0.78	0.78	0.80
Sentiment Detection	Negative	0.68	0.53	0.53	0.60	0.64	0.66	0.66	0.65	0.77	0.68	0.67	0.72	0.47	0.73	0.73	0.57	0.35	0.69	0.46	0.69
	Neutral	0.80	0.93	0.93	0.86	0.86	0.86	0.87	0.86	0.86	0.91	0.91	0.89	0.88	0.54	0.54	0.67	0.75	0.39	0.52	0.39
	Positive	0.76	0.46	0.46	0.58	0.75	0.72	0.72	0.74	0.81	0.75	0.75	0.78	0.48	0.88	0.88	0.62	0.40	0.74	0.54	0.74
Scientific Claim	*Claim Existence*	0.85	0.69	0.77	0.76	0.88	0.80	0.85	0.84	0.85	0.76	0.81	0.80	0.60	1.00	0.60	0.67	0.68	0.88	0.67	0.67
	Verifiable	0.85	0.83	0.85	0.84	0.95	0.99	0.97	0.97	0.86	0.83	0.82	0.80	0.70	0.72	0.71	0.70	0.72	0.71	0.72	0.71
	Reference	0.77	0.51	0.87	0.49	0.64	0.78	0.89	0.70	0.88	0.84	0.88	0.87	0.67	0.71	0.72	0.67	0.77	0.74	0.71	0.73
	General	0.80	0.66	0.93	0.70	0.53	0.70	0.92	0.61	0.94	0.92	0.93	0.92	0.71	0.71	0.72	0.67	0.61	0.70	0.71	0.64

conducted experiments with both BERT, SetFit, and LLMs as well. Tweets received a 0/1 label for each category indicating exclusion/inclusion. In another branch of experiments, we evaluated the model on a binary scientific classification task. Rather than multi-label categorization, this setup predicted whether Tweets were related to science or not.

Table 6 presents a comparative evaluation of task-specific model SciBERT, BERT, SetFit, and LLMs on the task of categorizing Tweets based on the type of scientific content present. The models are assessed on multi-label classification for detecting Tweets scientific claims. containing: (1) verifiable scientific claims, (2) references to scientific sources, and (3) general scientific terminology. Our study indicates that SetFit excels in recognizing scientific claims, particularly in identifying those that are verifiable. Nonetheless, for categorizing reference and general claim types, the SciBERT model, tailored specifically for scientific text, demonstrates superior efficacy. Conversely, GPT-3.5 and LLaMA-2-70B exhibit limited capability in detecting scientific claims.

Topic-Targeted Scientific Claim Analysis. To investigate the dissemination of scientific information about COVID-19 issued by individuals and organizations based on Table 3 across six topics, we conducted an analysis and the results are provided in Table 5. Results showed that individuals and organizations had the lowest percentage of Tweets without factual assertions. On the other hand, organizations shared the most scientifically verifiable Tweets, complete with authentic references, outnumbering individual Tweets across these six topics except for Travel. As a result, it is evident that organizations propagated more credible claims than individuals since they have established reputations to sustain.

Table 5: Scientific Claims classified into [Ver]ifiable, [Ref]erences to Scientific sources, and [Gen]eral categories on Topics Related to COVID-19 for Individual vs. Organization.

Topic	**Ver**		**Ref**		**Gen**	
	Ind(%)	Org(%)	Ind(%)	Org(%)	Ind(%)	Org(%)
Vaccine	41.00	48.91	46.73	53.24	28.17	34.28
Travel	45.96	39.30	55.86	44.12	30.22	25.65
Face Mask	30.47	58.00	34.78	65.20	19.21	34.91
Social Distancing	37.24	44.81	41.99	47.99	25.03	29.71
Lockdown & Quarantine	39.90	47.97	47.01	52.97	26.76	32.32
Work From Home	42.75	47.75	48.54	51.44	28.88	33.52

8 Conclusion

Analysis of COVID-19 Twitter discourse in this paper offers valuable insights into how influential voices shape public narratives during global crises. By developing an integrated framework to characterize Tweets across multiple facets—including user profiles, emotions, sentiment, and scientific claims—this research enables a nuanced investigation of influencer's role amidst the pandemic. Using BERT embeddings, we improved account classification between individuals and organizations. Through biography clustering with BERTopic, we revealed influencer communities' Tweets about COVID-19, benefiting marketing research and dataset curation. Additionally, we pioneered an emotion detection model for Tweets, released on HuggingFace. We further constructed a tailored sentiment classifier to track attitude trajectories throughout the pandemic. Uniquely, our approach identifies author sentiment, while prior works relied solely on textual polarity. This nuanced technique provides richer insights into Twitter's emotional landscape. Scientific claim verification involved training a custom SciBERT-based classifier on domain-specific COVID-19 data. Key findings reveal notable differences between individuals and organizations in emotion expression, engagement patterns, and claim verification. The techniques developed in this work can inform communication strategies to promote healthy public discussions as online platforms grow increasingly influential.

Limitations include dataset constraints like language and scope. The dataset is restricted to English Tweets from influencers above 10,000 followers, limiting generalizability. Inherent subjectivity in annotating textual emotions also poses challenges. Additionally, the models rely solely on textual content without contextual cues. Nonetheless, the toolkit offers unprecedented capacity to investigate the complex, far-reaching impacts of COVID-19 across the digital landscape over time. Moving forward, we plan investigations into follower influence dynamics stemming from influencer content. We intend to develop a quantifiable metric assessing the degree of impact influencers have on followers. Formally evaluating the sway held by influencers via this proposed index would significantly advance comprehension of how impactful their messaging proves. Our forthcoming work seeks to unravel the mechanisms and effects at play following the publication of questionable influencer posts. Pinpointing the relationships between influencer content and audience groups represents an open research question warranting rigorous inspection using computational methods we aim to pioneer.

Table 6: Comparative Performance Analysis of BERT, SetFit, Task Specific, and LLMs in Emotion Detection, Sentiment Analysis, and Scientific Claim Identification. Task Specific Models, for Emotion Detection, we employed BERT with augmented data and for Sentiment Analysis, $\text{RoBERTa}_{\text{Twitter}}$ was utilized and SciBERT was employed for Scientific Claim Identification.

Task	Label	BERT				SetFit				Task Specific Model				GPT-3.5				LLaMA-2-70B			
		P	R	A	F1	P	R	A	F1	P	R	A	F1	P	R	A	F1	P	R	A	F1
Scientific Claim	*Claim Existence*	0.85	0.69	0.77	0.76	0.88	0.80	0.85	0.84	0.85	0.76	0.81	0.80	0.60	1.00	0.60	0.67	0.68	0.88	0.67	0.67
	Verifiable	0.85	0.83	0.85	0.84	0.95	0.99	0.97	0.97	0.86	0.83	0.82	0.80	0.70	0.72	0.71	0.70	0.72	0.71	0.72	0.71
	Reference	0.77	0.51	0.87	0.49	0.64	0.78	0.89	0.70	0.88	0.84	0.88	0.87	0.67	0.71	0.72	0.67	0.77	0.74	0.71	0.73
	General	0.80	0.66	0.93	0.70	0.53	0.70	0.92	0.61	0.94	0.92	0.93	0.92	0.71	0.71	0.72	0.67	0.61	0.70	0.71	0.64

Acknowledgement

This work was partially supported by the U.S. National Science Foundation under Grant No. 1822118 and 2226232, Award Numbers DMS 2123761, the member partners of the NSF IUCRC Center for Cyber Security Analytics and Automation – AMI, NewPush, Cyber Risk Research, NIST and ARL, by NIST under Award No. 60NANB23D152, the State of Colorado (grant #SB 18-086) and the authors' institutions. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, or other organizations and agencies.

References

1. Abd-Alrazaq, A., Alhuwail, D., Househ, M., Hamdi, M., Shah, Z.: Top Concerns of Tweeters During the COVID-19 Pandemic: Infoveillance Study. Journal of medical Internet research **22**(4), e19016 (2020)
2. Ainapure, B.S., Pise, R.N., Reddy, P., Appasani, B., Srinivasulu, A., Khan, M.S., Bizon, N.: Sentiment Analysis of COVID-19 Tweets Using Deep Learning and Lexicon-Based Approaches. Sustainability **15**(3), 2573 (2023)
3. Barbieri, F., Camacho-Collados, J., Anke, L.E., Neves, L.: TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In: Findings of the Association for Computational Linguistics: EMNLP 2020. pp. 1644–1650 (2020)
4. Bruns, A., Harrington, S., Hurcombe, E.: 'Corona? 5G? or Both?': the Dynamics of COVID-19/5G Conspiracy Theories on Facebook. Media International Australia **177**(1), 12–29 (2020)

5. Campello, R.J.G.B., Moulavi, D., Sander, J.: Density-based Clustering Based on Hierarchical Density Estimates. In: Advances in Knowledge Discovery and Data Mining. PAKDD 2013. pp. 160–172 (2013)
6. Chakraborty, K., Bhatia, S., Bhattacharyya, S., Platos, J., Bag, R., Hassanien, A.E.: Sentiment Analysis of COVID-19 Tweets by Deep Learning Classifiers—A study to show how popularity is affecting accuracy in social media. Applied Soft Computing **97**, 106754 (2020)
7. Chau, M., Xu, J.: Business Intelligence in Blogs: Understanding Consumer Interactions and Communities. MIS Quarterly **36**(4), 1189–1216 (2012)
8. Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C.M., Brugnoli, E., Schmidt, A.L., Scala, A.: The COVID-19 Social Media Infodemic. Scientific Reports **10**(1), 1–10 (2020)
9. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. CoRR **abs/1810.04805** (2018)
10. Dubey, A.D.: Twitter Sentiment Analysis During COVID-19 Outbreak. Available at SSRN 3572023 (2020)
11. Egger, R., Yu, J.: A Topic Modeling Comparison Between LDA, NMF, Top2Vec, and BERTopic to Demystify Twitter Posts. Frontiers in Sociology **7**, 886498 (2022)
12. Ekman, P., et al.: Basic emotions. Handbook of Cognition and Emotion **98**(45-60), 16 (1999)
13. Freberg, K., Graham, K., McGaughey, K., Freberg, L.A.: Who are the Social Media Influencers? A Study of Public Perceptions of Personality. Public Relations Review **37**(1), 90–92 (2011)
14. Ghebreyesus, T.A.: Munich Security Conference. World Health Organization **15** (2020)
15. Grootendorst, M.: BERTopic: Neural Topic Modeling with a Class-based TF-IDF Procedure. arXiv preprint arXiv:2203.05794 (2022)
16. Hafid, S., Schellhammer, S., Bringay, S., Todorov, K., Dietze, S.: Scitweets - A Dataset and Annotation Framework for Detecting Scientific Online Discourse. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Managemen. pp. 3988–3992 (2022)
17. Jalil, Z., Abbasi, A., Javed, A.R., Badruddin Khan, M., Abul Hasanat, M.H., Malik, K.M., Saudagar, A.K.J.: Covid-19 Related Sentiment Analysis Using State-of-the-art Machine Learning and Deep Learning Techniques. Frontiers in Public Health **9**, 2276 (2022)
18. Kaplan, A.M., Haenlein, M.: Users of the World, Unite! The Challenges and Opportunities of Social Media. Business Horizons **53**(1), 59–68 (2010)
19. Kappus, P., Groß, P.: Finding Clusters of Similar-minded People on Twitter Regarding the Covid-19 Pandemic. arXiv preprint arXiv:2203.04764 (2021)
20. Kaur, H., Ahsaan, S.U., Alankar, B., Chang, V.: A Proposed Sentiment Analysis Deep Learning Algorithm for Analyzing COVID-19 Tweets. Information Systems Frontiers pp. 1–13 (2021)
21. Knowles, R., Carroll, J., Dredze, M.: Demographer: Extremely Simple Name Demographics. In: Proceedings of the First Workshop on NLP and Computational Social Science. pp. 108–113 (2016)
22. Landers, E., et al.: An Assessment of Scientific Claim Verification Frameworks: Final Presentation. Computer & Information Science: Research Experiences for Undergraduates in Disinformation Detection and Analytics **8** (2022)
23. Liang, S., Zhang, X., Ren, Z., Kanoulas, E.: Dynamic Embeddings for User Profiling in Twitter. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. p. 1764–1773 (2018)

24. Lopez, C.E., Vasu, M., Gallemore, C.: Understanding the Perception of COVID-19 Policies by Mining a Multilanguage Twitter Dataset. CoRR **abs/2003.10359** (2020)
25. Manguri, K.H., Ramadhan, R.N., Amin, P.R.M.: Twitter Sentiment Analysis on Worldwide COVID-19 Outbreaks. Kurdistan Journal of Applied Research pp. 54–65 (2020)
26. Mansoor, M., Gurumurthy, K., U, A.R., Prasad, V.R.B.: Global Sentiment Analysis of COVID-19 Tweets Over Time. CoRR **abs/2010.14234** (2020)
27. Nandal, N., Tanwar, R., Pathan, A.S.K.: Sentiment Analysis based Emotion Extraction for COVID-19 Using Crawled Tweets and Global Statistics for Mental Health. Procedia Computer Science **218**, 949–958 (2023)
28. OpenAI: Models. https://platform.openai.com/docs/models/gpt-3-5 (2023), (Accessed on 11/18/2023)
29. Pennycook, G., McPhetres, J., Zhang, Y., Lu, J.G., Rand, D.G.: Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention. Psychological Science **31**(7), 770–780 (2020)
30. Rajput, N.K., Grover, B.A., Rathi, V.K.: Word frequency and sentiment analysis of twitter messages during Coronavirus pandemic. CoRR **abs/2004.03925** (2020)
31. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. Journal of Computational and Applied Mathematics **20**, 53–65 (1987)
32. Saakyan, A., Chakrabarty, T., Muresan, S.: COVID-fact: Fact Extraction and Verification of Real-World Claims on COVID-19 Pandemic. In: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). pp. 2116–2129 (2021)
33. Storey, V.C., O'Leary, D.E.: Text Analysis of Evolving Emotions and Sentiments in COVID-19 Twitter Communication. Cognitive Computation **16**(4), 1834–1857 (2024)
34. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv preprint arXiv:2307.09288 (2023)
35. Tunstall, L., Reimers, N., Jo, U.E.S., Bates, L., Korat, D., Wasserblat, M., Pereg, O.: Efficient few-shot learning without prompts. arXiv preprint arXiv:2209.11055 (2022)
36. Vijay, T., Chawla, A., Dhanka, B., Karmakar, P.: Sentiment Analysis on COVID-19 Twitter Data. In: 2020 5th IEEE International Conference on Recent Advances and Innovations in Engineering (ICRAIE). pp. 1–7 (2020)
37. Wadden, D., Lin, S., Lo, K., Wang, L.L., van Zuylen, M., Cohan, A., Hajishirzi, H.: Fact or fiction: Verifying scientific claims. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 7534–7550. Association for Computational Linguistics (2020)
38. Wadden, D., Lo, K., Wang, L.L., Cohan, A., Beltagy, I., Hajishirzi, H.: Longchecker: Improving Scientific Claim Verification by Modeling Full-Abstract Context. arXiv preprint arXiv:2112.01640 (2021)
39. Wood-Doughty, Z., Mahajan, P., Dredze, M.: Johns Hopkins or johnny-hopkins: Classifying Individuals versus Organizations on Twitter. In: Proceedings of the Second Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media. pp. 56–61. Association for Computational Linguistics (2018)

40. Wrycza, S., Maślankowski, J.: Social Media Users' Opinions on Remote Work During the COVID-19 Pandemic. Thematic and Sentiment Analysis. Information Systems Management **37**(4), 288–297 (2020)
41. Xavier, T., Lambert, J.: Sentiment and Emotion Trends in Nurses' Tweets About the COVID-19 Pandemic. Journal of Nursing Scholarship **54**(5), 613–622 (2022)
42. Yu, H., Yang, C.C., Yu, P., Liu, K.: Emotion Diffusion Effect: Negative Sentiment COVID-19 Tweets of Public Organizations Attract More Responses from Followers. PloS one **17**(3), e0264794 (2022)
43. Zhou, J., Yang, S., Xiao, C., Chen, F.: Examination of Community Sentiment Dynamics due to COVID-19 Pandemic: A Case study from a State in Australia. SN Computer Science **2**, 1–11 (2021)
44. Zuo, C., Banerjee, R., Chaleshtori, F.H., Shirazi, H., Ray, I.: Seeing Should Probably Not be Believing: the Role of Deceptive support in COVID-19 Misinformation on Twitter. ACM Journal of Data and Information Quality **15**(1), 1–26 (2022)