

Social-media insights into US mental health amid the COVID-19 global pandemic: a Longitudinal analysis of publicly available Twitter data (January 22- April 10, 2020)

Danny Valdez, Marijn ten Thij, Krishna Bathina, Lauren Alexandra Rutter, Johan Bollen

Submitted to: Journal of Medical Internet Research
on: June 14, 2020

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 5

Supplementary Files..... 34

 Figures 35

 Figure 3 36

 Figure 4 37

 Figure 1 38

 Figure 2 39

 Figure 5 40

Social-media insights into US mental health amid the COVID-19 global pandemic: a Longitudinal analysis of publicly available Twitter data (January 22- April 10, 2020)

Danny Valdez¹ PhD; Marijn ten Thij² PhD; Krishna Bathina² MS; Lauren Alexandra Rutter³ PhD; Johan Bollen² PhD

¹Indiana University School of Public Health Department of Applied Health Science Bloomington US

²Indiana University Luddy School of Informatics, Computing, and Engineering Bloomington US

³Indiana University Psychological and Brain Sciences Bloomington US

Abstract

Background: The COVID-19 pandemic led to unprecedented mitigation efforts that disrupted the daily lives of millions. Beyond the general health repercussions of the pandemic itself, these measures also present a significant challenge to the world's mental health and healthcare systems. Considering traditional survey methods are time-consuming and expensive, we need timely and proactive data sources to respond to the rapidly evolving effects of health policy on our population's mental health. Significant pluralities of the US population now use social media platforms, such as Twitter, to express the most minute details of their daily lives and social relations. This behavior is expected to increase during the COVID-19 pandemic, rendering social media data a rich field from which to understand personal wellbeing.

Objective: Broadly, this study answers three research questions: RQ1: What themes emerge from a corpus of US tweets about COVID-19?; RQ2: To what extent does social media use increase during the onset of the COVID-19 pandemic?; and RQ3: Does sentiment change in response to the COVID-19 pandemic?

Methods: We analyzed 86,581,237 public domain English-language US tweets collected from an open-access public repository (Chen, Lerman, & Ferrara, 2020) in three steps. First, we characterized the evolution of hashtags over time using Latent Dirichlet Allocation (LDA) topic modeling. Second, we increased the granularity of this analysis by downloading Twitter timelines of a large cohort of individuals (n = 354,738) in 20 major US cities to assess changes in social media use. Finally, using this timeline data, we examined collective shifts in public mood in relation to evolving pandemic news cycles by analyzing the average daily sentiment of all timeline tweets with the Valence Aware Dictionary and sEntiment Reasoner (VADER) (Hutto & Gilbert, 2014) sentiment tool.

Results: LDA topics generated in the early months of the dataset corresponded to major COVID-19 specific events. However, as state and municipal governments began issuing stay-at-home orders, latent themes shifted towards US-related lifestyle changes rather than global pandemic-related events. Social media volume also increased significantly, peaking during stay-at-home mandates. Finally, VADER sentiment analysis sentiment scores of user timelines were initially high and stable, but decreased significantly, and continuously, by late March.

Conclusions: Our findings underscore the negative effects of the pandemic on overall population sentiment. Increased usage rates suggest that, for some, social media may be a coping mechanism to combat feelings of isolation related to long-term social distancing. However, in light of the documented negative effect of heavy social media usage on mental health, for many social media may further exacerbate negative feelings in the long-term. Thus, considering the overburdened US mental healthcare structure, these findings have important implications for ongoing mitigation efforts.

(JMIR Preprints 14/06/2020:21418)

DOI: <https://doi.org/10.2196/preprints.21418>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

✓ **Please make my preprint PDF available to anyone at any time (recommended).**

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

No, I do not wish to publish my submitted manuscript as a preprint.

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to the public.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in <http://www.jmir.org/>, I will be able to make my manuscript PDF available to the public.



Original Manuscript

Abstract

Background. The COVID-19 pandemic led to unprecedented mitigation efforts that disrupted the daily lives of millions. Beyond the general health repercussions of the pandemic itself, these measures also present a significant challenge to the world's mental health and healthcare systems. Considering traditional survey methods are time-consuming and expensive, we need timely and proactive data sources to respond to the rapidly evolving effects of health policy on our population's mental health. Significant pluralities of the US population now use social media platforms, such as Twitter, to express the most minute details of their daily lives and social relations. This behavior is expected to increase during the COVID-19 pandemic, rendering social media data a rich field from which to understand personal wellbeing.

Purpose. Broadly, this study answers three research questions: RQ1: What themes emerge from a corpus of US tweets about COVID-19?; RQ2: To what extent does social media use increase during the onset of the COVID-19 pandemic?; and RQ3: Does sentiment change in response to the COVID-19 pandemic?

Methods. We analyzed 86,581,237 public domain English-language US tweets collected from an open-access public repository (Chen, Lerman, & Ferrara, 2020) in three steps. First, we characterized the evolution of hashtags over time using Latent Dirichlet Allocation (LDA) topic modeling. Second, we increased the granularity of this analysis by downloading Twitter timelines of a large cohort of individuals ($n = 354,738$) in 20 major US cities to assess changes in social media use. Finally, using this timeline data, we examined collective shifts in public mood in relation to evolving pandemic news cycles by analyzing the average daily sentiment of all timeline tweets with the Valence Aware Dictionary and sEntiment Reasoner (VADER) (Hutto & Gilbert, 2014) sentiment tool.

Results. LDA topics generated in the early months of the dataset corresponded to major COVID-19 specific events. However, as state and municipal governments began issuing stay-at-home orders, latent themes shifted towards US-related lifestyle changes rather than global pandemic-related events. Social media volume also increased significantly, peaking during stay-at-home mandates. Finally, VADER sentiment analysis sentiment scores of user timelines were initially high and stable, but decreased significantly, and continuously, by late March.

Discussion & Conclusion. Our findings underscore the negative effects of the pandemic on overall population sentiment. Increased usage rates suggest that, for some, social media may be a coping mechanism to combat feelings of isolation related to long-term social distancing. However, in light of the documented negative effect of heavy social media usage on mental health, for many social media may further exacerbate negative feelings in the long-term. Thus, considering the overburdened US mental healthcare structure, these findings have important implications for ongoing mitigation efforts.

Key Words: Social media analytics; COVID-19; US mental health; informatics

Introduction

Beyond the obvious physical health ramifications of the COVID-19 pandemic, public health, and the greater medical community, is also bracing for a mental health crisis[1]. Within the span of four months, 45% of Americans indicated that the COVID-19 pandemic had taken a toll on their mental health, reporting higher levels of sadness and worsening of chronic psychiatric conditions[2]. Yet, despite an abundance of anecdotal evidence and peer-reviewed editorials that identify the potential mental-health fallout of this public health crisis, the extent of these effects is empirically unknown.

Scientists mobilized quickly to measure many facets of the pandemic, including potential mental-health effects. However, the time-consuming and costly nature of survey development[3] and instrument validation make it difficult to draw real-time conclusions[4], especially amidst rapidly evolving news cycles that shift pandemic related discourse. In the absence of survey data, social media represents a potentially valuable data source for studying emergent social issues, including the effect of those issues on behaviors and social mood[5]. Repeated tracking of social media data can provide a diachronic perspective on public morale and collective changes in sentiment, as participants voluntarily contribute to narratives, providing unprompted and diverse understandings of various issues[6–8]. Numerous scholars have successfully used social media data to identify trends and nuances in public mood using a combination of machine learning and artificial intelligence approaches. Some examples include comparing the happiness of users to their online social networks [9,10], identifying detailed predictors of mood through social media feeds[5], predicting cognitive distortions expressed among groups at-risk of mental health disorders[11], tracking the emotions of social media users at high resolution [12,13], and mapping negative affectivity among users with internalizing disorders[14]. Collectively, these studies demonstrate the feasibility and value of using sentiment analysis on social media data to study societal mood and well-being, as well as biomedical signals among social media users that can provide useful proxies for mental health[12,15–17]. In fact, these approaches may be especially useful considering the speed with which the pandemic became an acute socio-economic phenomenon, the pervasiveness of COVID-19 related content available online, and the natural reaction of many to post on social media about pandemic-related events.

Indeed, throughout the COVID-19 pandemic, individuals have sought out crisis-related news at increased capacities[18], leading to a collective increase in global social media use[19]. This renders social media data about the COVID-19 pandemic a powerful source of information from which to draw real-time conclusions about aggregate social wellbeing during an unprecedented

public health event. However, we must remember that just as survey data are prone to biases, so are data derived from social media[20]. Therefore, to draw accurate inferences about mood/sentiment and mental health, we must remain cognizant of the type of analysis performed and what the analysis represents to measure nuanced aspects of sentiment. Here we contrast (1) topics discussed, (2) topic-related sentiment, and (3) personal sentiment to arrive at a more comprehensive and accurate assessment of changes in sentiments expressed through social media and their relevance to public health in the United States.

Broadly, this study answers three research questions: (1) RQ1: What themes emerge from a corpus of US tweets about COVID-19? (2) RQ2: To what extent does social media use increase during the onset of the COVID-19 pandemic? And (3) RQ3: What patterns emerge from longitudinal tracking of sentiment during the onset of the COVID-19 pandemic? To address these research questions, we analyzed a large-scale set of Twitter data that are strictly relevant to the topic of COVID-19 in the United States from January 22, 2020 onwards. Using this data, we also compiled a second corpus of individual, geolocated social media timeline data from the same period to understand changes in personal sentiment as a proxy for mental health and evolving US perceptions of the COVID-19 pandemic.

Methods

Data

We collected two distinct data sets, each reflecting different aspects of changes in social media behavior before and during the COVID-19 pandemic. The first dataset of tweets, collected from an open-access repository containing all COVID related tweets published in the United States[21], was designed to capture *topical differences* (i.e. themes) in the Twitter discussion during the events that marked the onset of the pandemic. The repository provides a list of Tweet IDs, which we used to extract tweet content from Twitter's API (endpoint: GET statuses/show/id). We downloaded each tweet as well as the standard meta-data provided by Twitter. Specifically, we retrieved COVID-19 related tweets posted between January 22, 2020 (first day of data collection and

roughly one week prior to the first confirmed US COVID-19 case) through April 9, 2020 (the middle of social distancing efforts). Hereafter, we refer to this set of tweets as the “*COVID-19 corpus*”, ($n = 86,581,237$ tweets). Please refer to Figure 1 for a visual representation of this dataset and how it was retrieved.

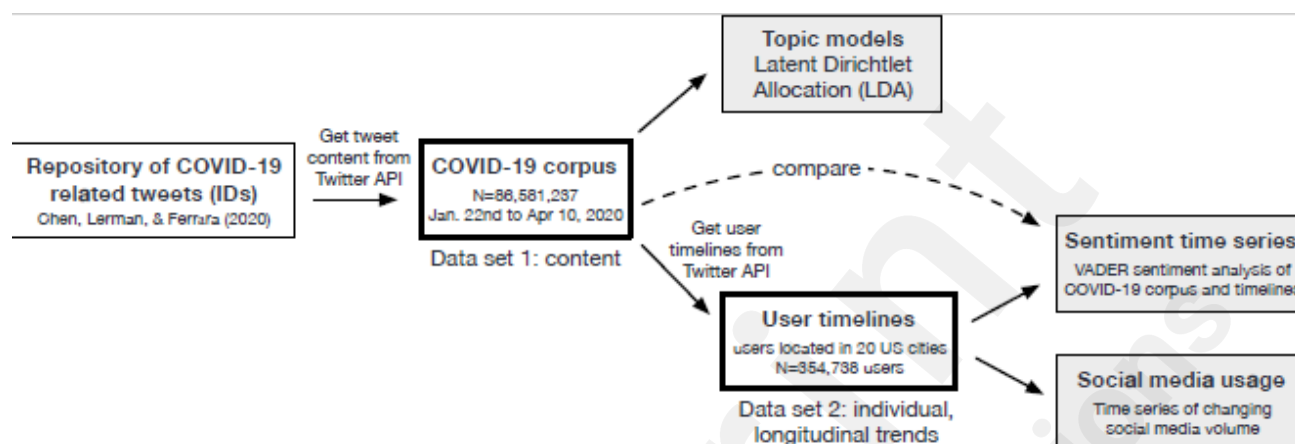


Figure 1: This diagram illustrates the procedure undertaken to procure Tweet IDs from an open-access COVID-19 repository. These tweet IDs were run through Twitter’s API to create two distinct data sets. The COVID-19 corpus, which contains all COVID-19 related English-language tweets published in the United States between January-April 2020. The User time-line data contains the 3200 most recent tweets of users residing in the 20 cities most-affected by COVID-19.

Second, to gauge fluctuations of personal activity and mood at the individual, rather than topical, level, we downloaded the twitter timelines (i.e. the 3,200 most recent tweets) of individual social media users who contributed to the COVID-19 corpus and *also* resided within the 20 US cities with the most COVID-19 cases per 100,000 people from the Twitter API (endpoint: GET statuses / user_timeline). These timelines capture the changing behavior and emotions of individual Twitter users during the COVID-19 pandemic, but do not strictly pertain to tweets exclusively related to COVID-19. We refer to this data as *user-timeline data*; $n = 354,738$ users, ($n = 69,349,479$ tweets), as shown in Figure 1. All tweets in either dataset were scrubbed of any personally identifiable information in accordance with ethical social media use practices.

To ensure that we are measuring expressed sentiment in our data we exclude non-English tweets and, within the user-timeline data, specifically, retweets, and biasing key-words, including: (1) “coronavirus”, (2) “COVID-19”, (3) “pandemic”, among others. These words were removed because they inherently carry a negative connotation and their inclusion would artificially decrease sentiment

given that the corpus itself is composed of COVID-19 related content. In other words, because users are naturally tweeting about ‘coronavirus’, ‘virus’ and the ‘pandemic’, the inclusion of those words may not necessarily reflect the individuals’ well-being. Note, the resulting sample sizes for each corpora exceeds the mean observed in a recent scoping survey of the literature on social media analytics for public health by several orders of magnitude ($n = 20,000$)[22], resulting in ample representation with which to conduct our analyses. Additionally, previous studies have used large-scale sentiment analysis to accurately predict social mood[23], and how sentiment expressed on social media correlates with psychological wellbeing[24]. Thus, the use of sentiment analysis for this study is appropriate.

Analyses

LDA Topic Models. Latent Dirichlet Allocation (LDA) topic models are unsupervised machine-learning tools that perform probabilistic inferences to consolidate large volumes of text data into manageable themes[25]. Simply, words with high probabilities of association (i.e. a high likelihood to appear in proximity with other words) are grouped together to form a latent theme, or topic, that qualitatively represents a content area within the collection of text. These methods have been applied in myriad ways, e.g. to determine common themes in product reviews[26], to map themes within bodies of scientific literature[27], and to identify themes in social media data[28]. Thus, these tools are appropriate for exploratory analyses that seek to consolidate dense text data.

Sentiment Analysis. Sentiment analysis refers to a set of supervised and/or unsupervised Machine Learning (ML) and Natural Language Processing (NLP) techniques that extract affective or emotional indicators from text, e.g. to determine whether a tweet expresses a negative or positive emotion about policy[23]. Here we used the Valence Aware Dictionary and sEntiment [sic] Reasoner (VADER)[29] to gauge the emotional valence of tweets. VADER is a rule-based, open-source tool that recognizes common terms, idioms, abbreviations, and jargon, while accounting for grammatical structures, such as punctuation, negation, hedging, and magnification, that are commonly employed in the vernacular of social media platforms. The VADER lexicon is one of the largest of its kind

containing over 7,500 common terms that are each rated for their emotional valence by 10 independent human raters. However, the word ‘virus’ and its many variations (e.g. viruses, viral) are not part of the VADER lexicon, meaning changes in the frequency of these words will not bias VADER scores. VADER has been extensively validated for Twitter content[30], showing some of the highest accuracy and coverage for tweets in a benchmark of more than 20 sentiment analysis tools[31].

Change-point detection. We applied the Pruned Exact Linear Time (PELT) change-point detection algorithm to identify significant changes in tweet volume and sentiment[32]. Change-point detection algorithms perform a set of mathematical operations over a time series (a series of time-based observations) to identify points in time at which the statistical properties of the time-series data change significantly[33]. The PELT algorithm specifically attempts to find a set of change points for a given time series, such that their number and location in time minimizes a given segmentation cost. We chose the PELT algorithm over other similar change detection algorithms because it is considered to be a more conservative estimate (preferring not to identify change points unless strict conditions are satisfied), thus yielding more accurate detection of statistical changes[34]. Additionally, PELT uses an off-line approach to change detection[34], meaning it can consider all possible data points when identifying significant changes, regardless of the type of data, while maintaining high levels of performance.

Procedure

RQ1: What themes emerge from a corpus of US tweets about COVID-19? We divided the COVID-19 corpus into daily segments and generated one topic model per day consisting of 20 topics each. We chose this number to reflect the widest possible span of themes while summarizing the major themes of the online discussion, which is a process used in previous studies[28,35]. We then looked at the top 20 associated words per topic and collapsed similar words into general themes, taking into account similarities of words (i.e. United States and US), and potential misspellings, which are common in social media posts. As an example, ‘hubei’ and ‘wuhan’ were collapsed into

the theme 'China'. We then found the frequency ratio (i.e. the number of occurrences of a certain word divided by the total number of words) of COVID-19 related themes (China, US, pandemic, social distancing, Trump, home, lockdown and deaths) and plotted them on a daily basis to show the evolution of topics over time, indicating both the contribution of the theme to all content of that day and the relative ranking of these terms amongst these themes. We used inter-coder agreement methods to arrive at mutually agreeable interpretations of collapsed themes[36].

RQ2: To what extent does social media use increase during the onset of the COVID-19 pandemic? For this analysis we used the user-timeline data instead of the general COVID-19 corpus (See Fig 1.) because within-subject individual posting frequency is a better marker for tracking changes in social media use behavior[37]. Since Twitter's API limits us to only the 3200 *most recent postings* per individual, we only selected individuals who posted on Twitter before January 22, 2020, retaining 354,738 users. This ensures that our analytic sample captures individual behavior in the 20-most affected US cities throughout the interval of interest (January 22, 2020 to April 9, 2020). We performed a seasonal decomposition— a method that separates the baseline, trend, and seasonal components of a time series— to determine whether we can observe increased Twitter use during the pandemic relative to events just prior. We then detected significant points of change using the PELT algorithm[38].

RQ3: What patterns emerge from longitudinal tracking of sentiment during the onset of COVID-19 pandemic? Sentiment can change because individuals discuss different topics, e.g. using pejorative terms such as “virus” more frequently, or because of personal, individual changes in how people actually feel. We therefore compared daily VADER sentiment scores for (1) the COVID-19 corpus (to gauge topical sentiment) and (2) user-timeline within-subjects data (to assess personal changes in sentiment) from January 22, 2020 (the first official day of data collection), through April 9, 2020. We determine change-points in the time-series of daily averaged VADER sentiment with the PELT detection algorithm so we can identify significant changes in sentiment throughout this time

period.

Results

RQ1: What themes emerge from a corpus of US tweets about COVID-19?

We consolidated the COVID-19 corpus (n= 86,581,237 tweets) into themes using LDA topic models. Figure 1 highlights the eight most salient topics and how their prominence (width of bars and rank) changed over time relative to major COVID-19 related milestones identified by the WHO.

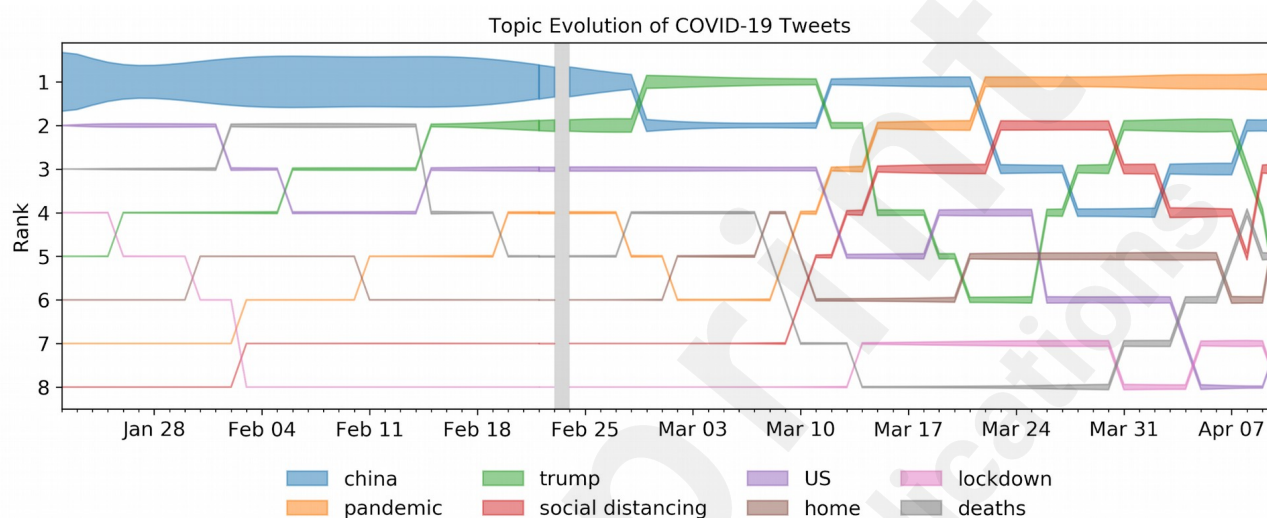


Figure 2: Topic-group rankings over time in the COVID-19 corpus. Each topic is ranked by its frequency ratio. The width of the bars indicates the fraction of words in the topic on a given day. The colors of the areas indicate which theme the corresponding area belongs to.

As shown in Figure 2, topics continued to rise and fall in prominence relative to emerging news cycles throughout the time period under consideration. Indeed, the majority of COVID-19 related Twitter activity focused on China in February. However, from March through April, as the novel coronavirus increasingly began to affect the US population, ‘China’ became less prominent as more US-centered topics such as “lockdown” and “social distancing” emerged. Although China remained a prominent theme throughout the duration of interest, US-centered topics gradually came to dominate social media spaces.

RQ2: To what extent does social media use increase during the onset of the COVID-19 pandemic?

Using user-timeline data, we compared the frequency of social media posts before and during the pandemic of 354,738 users individuals in 20 of the most-affected US cities (n= 69,349,479 tweets). Figure 3 highlights changes in posting volume between January 22, 2020 and April 9, 2020.

The peaks and troughs (dashed line) of this graph show how seasonal and weekly cycles shape tweet volume. The solid line plots the trend in the time-series after removing cycles and seasonal effects (through seasonal decomposition), and the progressively darker shades of brown denote the number of cities that imposed mandatory lockdowns.

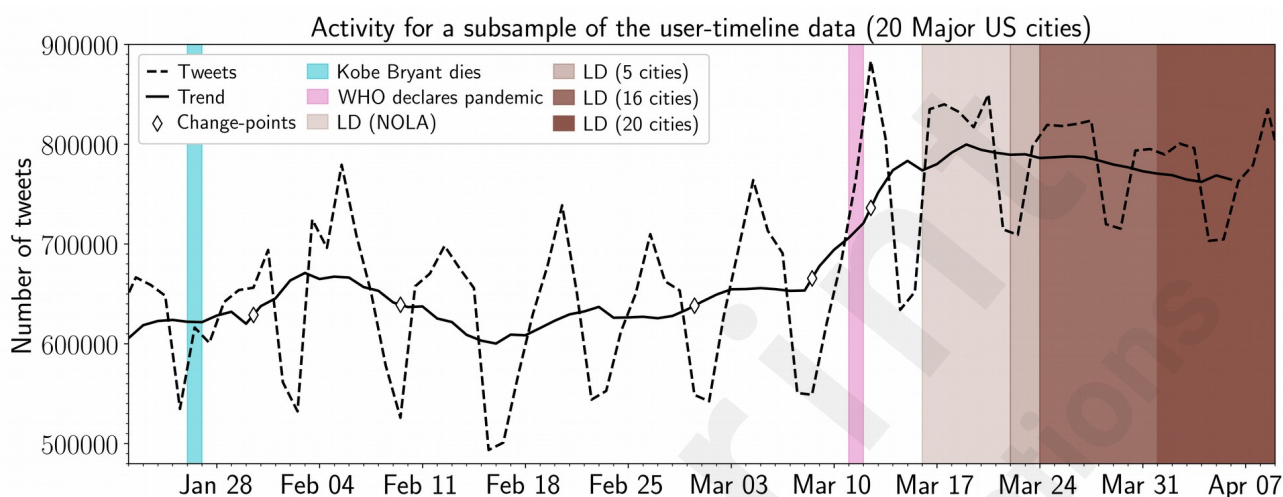


Figure 3: Number of daily tweets in a subsample of the user-timeline data ($n=292,000$). The dashed line displays the daily number of tweets in the subsample of the user-timeline data. The solid line indicates the underlying trend in the daily number of tweets (after removing its seasonal effects). The diamond markers indicate significant change-points indicated by the PELT-algorithm. The light-blue and pink annotations denote the death of Kobe Bryant (Jan 26, 2020) and the WHO declared COVID-19 as a global pandemic (March 11, 2020), respectively. The brown annotations indicate the dates at which lockdowns were enforced in the 20 considered cities (ranging March 16, 2020 to April 1, 2020). The opacity of the brown annotation indicates how many cities had enforced a lockdown at that date.

Generally, we observed a consistent upward trend in total Twitter volume from early to late March. The PELT change-point algorithm identified two significant volume changes on March 8 and 12 — around the time COVID-19 was declared a global pandemic (March 11) and President Trump's national emergency declaration in the United States (March 14). The upward trend stabilized thereafter, albeit at higher observed volumes than prior to the onset of the COVID-19 pandemic. This supports the notion that individuals in our sample were more engaged with social media and made more use of it, possibly to discuss or obtain further information relevant to the news cycle.

RQ3: What patterns emerge from longitudinal tracking of sentiment during the onset of the COVID-19 pandemic?

We applied the VADER sentiment tool to (1) the COVID-19 corpus (to assess sentiment of all US tweets about COVID-19; and (2) the user-timeline data (to track changes in user sentiment using their most recent 3,200 tweets). Figure 4 tracks sentiment relative to major COVID-19 related

milestones for both datasets, with the orange line tracking the COVID-19 corpus, and blue line tracking the user-timeline data, respectively. In the COVID-19 corpus, there was an unmistakable increase in sentiment with two PELT-identified significant changes on March 9 (just before WHO classified COVID-19 as a pandemic) and March 19 (shortly after President Donald Trump declared a national emergency). Figure 5 further shows the percentage of positively scored COVID-19 tweets increased over time, reinforcing the positive trend.

Conversely, the user-timeline data (which again contained the most recent 3,200 tweets of a given user) shows decreases in sentiment over the same period. The user-timeline data has one PELT-identified significant change a pandemic). Of note, there was a notable but short-lived drop in user sentiment before the WHO point on March 9 (just classified COVID-19 as sentiment on January 28-- the day NBA player Kobe Bryant died in a helicopter crash).

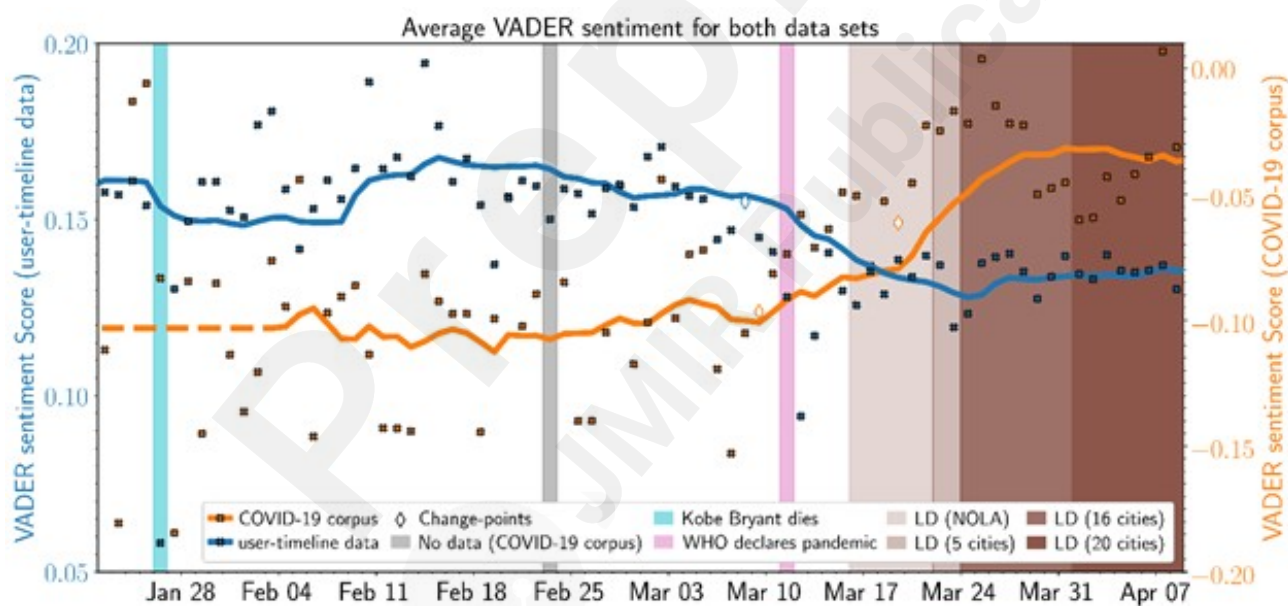


Figure 4: Distribution of non-zero VADER sentiment in the COVID-19 corpus & user-timeline data. The solid lines display the 14-day moving average and the crosses display the actual values of the daily average VADER sentiment (blue for the user-timeline data and orange for the COVID-19 corpus, respectively). The diamond markers indicate change-points indicated by the PELT-algorithm, the color of the edge of the diamonds refers to the time-series to which this change-point belongs. The light-blue, gray, and pink annotations denote the days of the death of Kobe Bryant (Jan 26, 2020), the day of missing data in the COVID-19 corpus (February 23, 2020), and the WHO declaring COVID-19 a global pandemic (March 11, 2020), respectively. The brown annotations indicate the dates at which lockdowns were enforced in the 20 considered cities (ranging March 16, 2020 to April 1, 2020). The opacity of the brown annotation indicates how many cities had enforced a lockdown at that date.

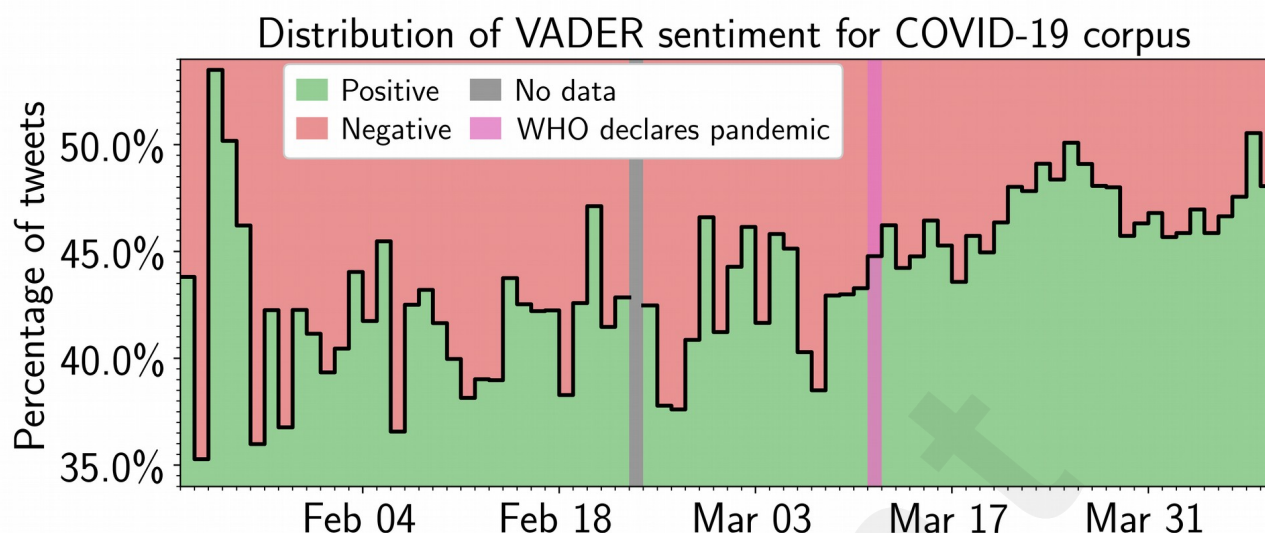


Figure 5: Distribution of positive and negative VADER sentiment in COVID-19 corpus. The solid line displays the fraction of tweets that hold positive or negative VADER sentiment in the user-timeline data (colored green for positive and red for negative sentiment). The gray and pink annotations denote the day of missing data (February 23, 2020) and the WHO declared COVID-19 as a global pandemic (March 11, 2020), respectively.

Discussion

The purpose of this study was to draw conclusions about US mental health amid the COVID-19 pandemic using computational social media analytics of a publicly available repository containing all COVID-19 related tweets. Using the COVID-19 tweet repository released by Chen and colleagues, we expanded on their original topical analysis of COVID-19 content by examining both topics and sentiment specific to the US population to understand how the COVID-19 pandemic may be impacting social wellbeing. We used social media data as the medium for analysis for several reasons, in particular its usefulness to gauge real-time changes in mental health/social wellbeing before, during, and after rapid socio-economic changes¹⁵. Below we discuss the findings of this study in detail and highlight public health implications about social wellbeing during the onset of the COVID-19 pandemic.

COVID-19 Social Media Themes: Content is Reactionary to News Cycles

Bento and colleagues predicted that crisis-related information seeking would increase during the COVID-19 pandemic[18]. Numerous studies have supported that prediction by identifying increased panic within social media spaces as users react to COVID-19 news-related content on their feeds[39,40]. Through our LDA topic modeling analysis, we echo much of those findings as well; the topics uncovered by their LDA frequently correspond to the life cycles of COVID-19 related news

and current events. For example, the name used to identify the virus on social media evolved to match changes in viral terminology as presented on news outlets-- e.g. the *novel coronavirus* became *coronavirus*, *COVID-19*, and *COVID-19 pandemic*, respectively. The locations that emerged within the topics also corresponded to the movement of COVID-19 from mainland China to the United States. For example, in late January, 'China' was a dominant theme in Twitter content (shown by the width of the bars in Figure 2) indicating the US public may have considered the then-classified epidemic a foreign matter (e.g. *TWEET: OMG China just shut down trading... Still think it's just a flu?*). However, in later weeks/months, US-centered topics/tweets replaced 'China' in prominence as more Americans became preoccupied with pandemic control measures at home (e.g. *TWEET: lol it's wild seeing every single other country grappling with this virus in ways we KNOW the US will never do*). These findings suggest that as the COVID-19 pandemic became more prominent in the US, social media content changed to topics of more direct impact to the US, and with greater frequency of use (evidenced by a PELT-identified significant increase in Twitter volume during that time). This volume increase remained consistently above levels observed prior to the COVID-19 pandemic, which is possibly explained by several factors, including increased anxiety as the COVID-19 pandemic reached the United States, the isolating effects of state-wide stay at home orders, or other social fallouts driven by the pandemic. These findings collectively support those of Castillo and colleagues[41], who contend that social media content follows the life-cycle of news stories. Thus, news, as a vehicle for crisis-related social media communication, should be studied more intently.

COVID-19 Topical Sentiment: Increased sentiment may indicate a priming effect

Given the expected negative physical and emotional outcomes of a major pandemic, we were surprised to observe that the sentiment of the COVID-19 corpus trended positively. However, this increase in sentiment is likely not indicative of actual changes in population sentiment, but rather the effect of a common *topical* bias in social media analytics[42]. Our COVID-19 data set was specifically selected to contain tweets that related to the topic of COVID-19. Hence, its sentiment

scores will reflect the language used to discuss this particular topic, and not necessarily the underlying emotion of the population. Topically selected social media posts are likely primed by news cycles[43] and shows of false optimism—a phenomena where individuals tend to post content that is more positive or optimistic than their true emotions[44]. Previous research has furthermore shown that language in general is biased towards positivity, especially when posts are collected for a random topic[30]. For example, language used to compose a post about COVID-19 may contain trending verbiage or framing devices, including showcasing support of groups frequently part of news cycles (e.g. TWEET *Great news...it's a welcome burden lifted off our incredible nurses and medical first responders*). Similarly, users may just be conveying positivity through carefully selected 'popular' words around a trending topic (e.g. TWEET: *No matter how hard the situation nowadays during the pandemic outbreak, we should keep being positive and optimistic*). Within the VADER lexicon, many of the words commonly used in this context are scored positively (great, welcome, incredible, positive, optimistic), which also artificially inflates sentiment ratings. Thus, we posit that topically driven tweet samples may not validly reflect actual changes in population mood, but rather topic-driven language sentiment. This justifies our approach to analyze within-subject timelines of individual posts which are not necessarily bound by the criterion of strictly being COVID-19 related, thereby increasing the odds of reflecting personal changes in mood trajectories.

User-timeline sentiment: Lower timeline sentiment may indicate decreased social wellbeing

By contrast, we found a negative trajectory in sentiment scores for the user-timeline data. This means that, while content in the COVID-19 corpus trended positively (possibly due to priming), relative to the totality of their timelines, our sample mood is *lower* than it once was. Thus, this comparison gives us deeper insight into *underlying* mood and sentiment. This finding further supports the assertion that the positivity conveyed in social media posts may not validly reflect what people are feeling at a given moment. Indeed, that positivity may be serving as a veneer posted 'in the moment' to convey positivity during a time of uncertainty. To obtain an accurate assessment of

mood and well-being, other reference points (in this case social media posts prior to the pandemic) are needed to examine in the moment sentiment relative to their prior histories. Within our timeline data, we captured one event that occurred prior to the pandemic that also affected sentiment scores relative to timeline histories—the passing of NBA basketball player Kobe Bryant. The effect of Bryant’s passing led to a sharp decline in sentiment, which lasted approximately 24 hours before returning to levels observed previously. Regarding the pandemic, after the PELT-identified shift on March 8, sentiment scores (relative to timelines) not only became lower, but consistently lower, and did not return to levels observed before the pandemic reached the United States. This trend may hold implications into the longitudinal effects of the pandemic and subsequent impacts on US mood/social sentiment, as this continued trend may be indicative of a long-lasting shift in mood and wellbeing.

Social Media Use in Times of Crisis

Analyses of aggregated social media feeds are shown to adequately predict other phenomena, including the stock market[45], political leanings[46], and when analyzed through a time-series, collective shifts in general mood[47]. Our study contributes to this body of literature by highlighting the disparity between how Americans portray themselves on social media, versus the latent sentiment they may be experiencing during times of crisis. Generally, Americans were not posting social media content about COVID-19 prior to the first documented US case. Once COVID-19 became a reality in the United States, however, there was a continued increase in the total number of US-based tweets about the novel coronavirus, indicative of growing social media use in our sample. By analyzing the COVID-19 corpus and user-timeline data separately, but with the same analytic procedures, we see divergent findings that reinforce the difference between in-the-moment portrayals versus the longitudinal information that can be gleaned from individual timeline analyses. For example, tweets and posts about COVID-19 may attempt to be lighthearted or convey optimism; however, individually, social media users may not be as optimistic as they were prior to the pandemic. During an unprecedented public health crisis it is therefore important to look beyond the topical focus of messages on social media that reference the crisis itself as a proxy of public mood, as they are likely

to be affected by other influences (e.g., political framing, projecting hope). Ultimately, our findings exemplify COVID-19 as a case-study in social media behavior, whose outcomes should be generalized to other crisis-related events.

Concluding Remarks and Implications

At the time of writing, the US COVID-19 death toll was just over 136,000, with millions of people disrupted by the various effects of the pandemic. This study elucidates the possible mental health effects of the COVID-19 pandemic among Twitter users using a computational approach to analyze a corpus of all archived, US-based COVID-19 tweets from January-April 2020. These analyses revealed, to varying extents, how the pervasiveness of COVID-19 content available on social media and abrupt shift in lifestyles may be negatively affecting social sentiment relative to points just prior to the pandemic. Given that sentiments expressed on social media have been used as a proxy for mental wellbeing[48] these findings support calls from public health and medical scholars who contend that mental health is an urgent concern during the COVID-19 pandemic, especially as our findings illustrate a declining trend in sentiment. Thus, we encourage further research on US mental health status amid the pandemic using survey methods or other primary data collections to substantiate our findings with testable outcomes. We also call for more research on mental-health interventions amid the COVID-19 pandemic, with particular attention to modality (i.e. in-person vs virtual), and the efficacy of those efforts.

Limitations

Our study is subject to limitations. First, Twitter requires users to opt-in to geo-tagging features. Consequently, any information inferred about the user's city of residence is often limited to self-reported data as specified on profile pages. This may mean that some of our timeline data may not originate from the city specified by a user, as this information can easily be misrepresented[49] (i.e. stating they live in New York, NY, but actually residing in Newark, NJ). We also acknowledge a likely bias regarding key demographic information, including age, gender, and socioeconomic status among likely social media users[50], in addition to the temporal, spatial, and geographic patterns that

may affect how sentiment is expressed on social media (e.g. older adults posting early in the morning versus younger adults posting late at night; or urban-based versus rural users)[51]. It is also, as of yet, not possible to accurately diagnose someone with a mental health condition through social media feeds alone, although research has indicated social media content contains important signals with respect to mental health and biomedical signals. Thus, we relied on trend data to draw inferences about the *possibility* of mental health decline based on average sentiment scores[52]. However, these limitations do not diminish the importance or validity of this work. Rather, they create avenues for additional research that expand on the findings of this paper and leverage the limitations inherent to social media data-- such as measuring cognitive distortions on Twitter based on posting time, or approaches that measure diagnosable mental health conditions through social media data, particularly during times of increased panic and crisis. And, because social media has been widely used to draw conclusions about public mood through large-scale sentiment analysis procedures[10,18], we contend our approach is appropriate to draw the conclusions discussed herein.

BIBLIOGRAPHY

1. Auerbach J, Miller BF. COVID-19 Exposes the Cracks in Our Already Fragile Mental Health System. *Am J Public Health*. Published online April 9, 2020:e1-e2. doi:10.2105/AJPH.2020.305699
2. Pfefferbaum B, North CS. Mental Health and the Covid-19 Pandemic. *New England Journal of Medicine*. 2020;0(0):null. doi:10.1056/NEJMp2008017
3. Gualano MR, Lo Moro G, Voglino G, Bert F, Siliquini R. Effects of Covid-19 Lockdown on Mental Health and Sleep Disturbances in Italy. *International Journal of Environmental Research and Public Health*. 2020;17(13):4779. doi:10.3390/ijerph17134779
4. Coughlan M, Cronin P, Ryan F. Survey research: Process and limitations. *International Journal of Therapy and Rehabilitation*. 2009;16(1):9-15. doi:10.12968/ijtr.2009.16.1.37935
5. Berry N, Emsley R, Lobban F, Bucci S. Social media and its relationship with mood, self-esteem and paranoia in psychosis. *Acta Psychiatrica Scandinavica*. 2018;138(6):558-570. doi:10.1111/acps.12953
6. Cho SE, Jung K, Park HW. Social Media Use during Japan's 2011 Earthquake: How Twitter Transforms the Locus of Crisis Communication. *Media International Australia*. 2013;149(1):28-40. doi:10.1177/1329878X1314900105

7. Page RE. *Stories and Social Media: Identities and Interaction*. Routledge; 2013.
8. Stieglitz S, Dang-Xuan L. Emotions and Information Diffusion in Social Media—Sentiment of Microblogs and Sharing Behavior. *Journal of Management Information Systems*. 2013;29(4):217-248. doi:10.2753/MIS0742-1222290408
9. Bollen J, Goncalves B, Ruan G, Mao H. Happiness is assortative in online social networks. *Artificial Life*. 2011;17(3):237-251. doi:10.1162/artl_a_00034
10. Bollen J, Gonçalves B, Leemput I van de, Ruan G. The happiness paradox: your friends are happier than you. *EPJ Data Sci*. 2017;6(1):1-10. doi:10.1140/epjds/s13688-017-0100-1
11. Simms T, Ramstedt C, Rich M, Richards M, Martinez T, Giraud-Carrier C. Detecting Cognitive Distortions Through Machine Learning Text Analytics. In: *2017 IEEE International Conference on Healthcare Informatics (ICHI)*. ; 2017:508-512. doi:10.1109/ICHI.2017.39
12. Correia RB, Wood IB, Bollen J, Rocha LM. Mining Social Media Data for Biomedical Signals and Health-Related Behavior. *Annual Review of Biomedical Data Science*. 2020;3(1):null. doi:10.1146/annurev-biodatasci-030320-040844
13. Golder SA, Macy MW. Diurnal and Seasonal Mood Vary with Work, Sleep, and Daylength Across Diverse Cultures. *Science*. 2011;333(6051):1878-1881. doi:10.1126/science.1202775
14. Social media use by adolescents linked to internalizing behaviors. Accessed July 16, 2020. <https://medicalxpress.com/news/2019-09-social-media-adolescents-linked-internalizing.html>
15. Bollen J, Pepe A, Mao H. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. *arXiv:09111583 [cs]*. Published online November 8, 2009. Accessed May 31, 2020. <http://arxiv.org/abs/0911.1583>
16. Dodds PS, Clark EM, Desu S, et al. Reply to Garcia et al.: Common mistakes in measuring frequency-dependent word characteristics. *PNAS*. 2015;112(23):E2984-E2985. doi:10.1073/pnas.1505647112
17. Bekalu MA, McCloud RF, Viswanath K. Association of Social Media Use With Social Well-Being, Positive Mental Health, and Self-Rated Health: Disentangling Routine Use From Emotional Connection to Use. *Health Educ Behav*. 2019;46(2_suppl):69S-80S. doi:10.1177/1090198119863768
18. Bento AI, Nguyen T, Wing C, Lozano-Rojas F, Ahn Y-Y, Simon K. Information Seeking Responses to News of Local COVID-19 Cases: Evidence from Internet Search Data. *arXiv:200404591 [physics]*. Published online April 6, 2020. Accessed May 19, 2020. <http://arxiv.org/abs/2004.04591>
19. Koeze E, Popper N. The Virus Changed the Way We Internet. *The New York Times*. <https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html>. Published April 7, 2020. Accessed May 31, 2020.
20. Olteanu A, Castillo C, Diaz F, Kıcıman E. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *Front Big Data*. 2019;2. doi:10.3389/fdata.2019.00013
21. Chen E, Lerman K, Ferrara E. COVID-19: The First Public Coronavirus Twitter Dataset.

- arXiv:200307372 [cs, q-bio]*. Published online March 16, 2020. Accessed May 19, 2020. <http://arxiv.org/abs/2003.07372>
22. Edo-Osagie O, De La Iglesia B, Lake I, Edeghere O. A scoping review of the use of Twitter for public health research. *Comput Biol Med.* 2020;122:103770. doi:10.1016/j.combiomed.2020.103770
 23. Liu B. *Sentiment Analysis and Opinion Mining.*; 2012.
 24. Jaidka K, Giorgi S, Schwartz HA, Kern ML, Ungar LH, Eichstaedt JC. Estimating geographic subjective well-being from Twitter: A comparison of dictionary and data-driven language methods. *PNAS.* 2020;117(19):10165-10171. doi:10.1073/pnas.1906364117
 25. Blei DM, Ng AY, Jordan MI. Latent Dirichlet Allocation. *Journal of Machine Learning Research.* 2003;3(Jan):993-1022.
 26. Wang H, Ding Y, Tang J, et al. Finding Complex Biological Relationships in Recent PubMed Articles Using Bio-LDA. *PLoS One.* 2011;6(3). doi:10.1371/journal.pone.0017243
 27. Moghaddam S, Ester M. ILDA: interdependent LDA model for learning latent aspects and their ratings from online product reviews. In: *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '11. Association for Computing Machinery; 2011:665–674. doi:10.1145/2009916.2010006
 28. Barry AE, Valdez D, Padon AA, Russell AM. Alcohol Advertising on Twitter—A Topic Model. *American Journal of Health Education.* 2018;49(4):256-263. doi:10.1080/19325037.2018.1473180
 29. Hutto CJ, Gilbert E. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. In: *ICWSM.* ; 2014.
 30. Elbagir S, Yang J. Twitter Sentiment Analysis Using Natural Language Toolkit and VADER Sentiment. *Hong Kong*. Published online 2019:5.
 31. Ribeiro MT, Singh S, Guestrin C. Model-Agnostic Interpretability of Machine Learning. *arXiv:160605386 [cs, stat]*. Published online June 16, 2016. Accessed May 19, 2020. <http://arxiv.org/abs/1606.05386>
 32. Killick R, Fearnhead P, Eckley IA. Optimal Detection of Changepoints With a Linear Computational Cost. *Journal of the American Statistical Association.* 2012;107(500):1590-1598. doi:10.1080/01621459.2012.737745
 33. Liu S, Yamada M, Collier N, Sugiyama M. Change-point detection in time-series data by relative density-ratio estimation. *Neural Networks.* 2013;43:72-83. doi:10.1016/j.neunet.2013.01.012
 34. Wambui GD, Waititu GA, Wanjoya A. The Power of the Pruned Exact Linear Time(PELT) Test in Multiple Changepoint Detection. *American Journal of Theoretical and Applied Statistics.* 2015;4(6):581. doi:10.11648/j.ajtas.20150406.30
 35. Valdez D, Pickett AC, Goodson P. Topic Modeling: Latent Semantic Analysis for the Social Sciences*. *Social Science Quarterly.* 0(0). doi:10.1111/ssqu.12528

36. Kuckartz U, Rädiker S. Analyzing Inter-coder Agreement. In: Kuckartz U, Rädiker S, eds. *Analyzing Qualitative Data with MAXQDA: Text, Audio, and Video*. Springer International Publishing; 2019:267-282. doi:10.1007/978-3-030-15671-8_19
37. Lee K, Agrawal A, Choudhary A. Real-time disease surveillance using Twitter data: demonstration on flu and cancer. In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '13. Association for Computing Machinery; 2013:1474–1477. doi:10.1145/2487575.2487709
38. Gabrielli L, Rinzivillo S, Ronzano F, Villatoro D. From Tweets to Semantic Trajectories: Mining Anomalous Urban Mobility Patterns. In: Nin J, Villatoro D, eds. *Citizen in Sensor Networks*. Lecture Notes in Computer Science. Springer International Publishing; 2014:26-35. doi:10.1007/978-3-319-04178-0_3
39. Ahmad AR, Murad HR. The Impact of Social Media on Panic During the COVID-19 Pandemic in Iraqi Kurdistan: Online Questionnaire Study. *Journal of Medical Internet Research*. 2020;22(5):e19556. doi:10.2196/19556
40. Depoux A, Martin S, Karafillakis E, Preet R, Wilder-Smith A, Larson H. The pandemic of social media panic travels faster than the COVID-19 outbreak. *J Travel Med*. 2020;27(3). doi:10.1093/jtm/taaa031
41. Castillo C, El-Haddad M, Pfeffer J, Stempeck M. Characterizing the life cycle of online news stories using social media reactions. In: *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. CSCW '14. Association for Computing Machinery; 2014:211–223. doi:10.1145/2531602.2531623
42. Ferrara E, Yang Z. Measuring Emotional Contagion in Social Media. *PLOS ONE*. 2015;10(11):e0142390. doi:10.1371/journal.pone.0142390
43. Schleuder J, White A, Cameron G. Priming Effects of Television News Bumpers and Teasers on Attention and Memory. *Journal of Broadcasting & Electronic Media*. 1993;37:437-452. doi:10.1080/08838159309364234
44. Bala DK. SOCIAL MEDIA AND CHANGING COMMUNICATION PATTERNS. :6.
45. Bollen J, Mao H, Zeng X. Twitter mood predicts the stock market. *Journal of Computational Science*. 2011;2(1):1-8. doi:10.1016/j.jocs.2010.12.007
46. Makazhanov A, Rafiei D, Waqar M. Predicting political preference of Twitter users. *Soc Netw Anal Min*. 2014;4(1):193. doi:10.1007/s13278-014-0193-5
47. Coppersmith G, Dredze M, Harman C, Hollingshead K. From ADHD to SAD: Analyzing the Language of Mental Health on Twitter through Self-Reported Diagnoses. In: *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. Association for Computational Linguistics; 2015:1–10. doi:10.3115/v1/W15-1201
48. Derks D, Fischer AH, Bos AER. The role of emotion in computer-mediated communication: A review. *Computers in Human Behavior*. 2008;24(3):766-785. doi:10.1016/j.chb.2007.04.004

49. Gore RJ, Diallo S, Padilla J. You Are What You Tweet: Connecting the Geographic Variation in America's Obesity Rate to Twitter Content. *PLOS ONE*. 2015;10(9):e0133505. doi:10.1371/journal.pone.0133505
50. Pew Research. Twitter Use 2012. Pew Research Center: Internet, Science & Tech. Published May 31, 2012. Accessed September 17, 2020. <https://www.pewresearch.org/internet/2012/05/31/twitter-use-2012/>
51. Padilla JJ, Kavak H, Lynch CJ, Gore RJ, Diallo SY. Temporal and spatiotemporal investigation of tourist attraction visit sentiment on Twitter. *PLOS ONE*. 2018;13(6):e0198857. doi:10.1371/journal.pone.0198857
52. Choudhury MD, De S. Mental Health Discourse on reddit: Self-Disclosure, Social Support, and Anonymity. :10.

APPENDIX A: STUDY SYNTAX

RQ1:

```
import matplotlib
import matplotlib.pyplot as plt
import glob
import pandas as pd
import numpy as np
import re
import gzip as gz
from collections import Counter
from collections import defaultdict
from scipy.special import comb
from scipy import stats

import matplotlib.ticker as ticker
from matplotlib.ticker import (MultipleLocator, FormatStrFormatter,
                               AutoMinorLocator)
from sklearn.cluster import KMeans
from multiprocessing import Pool

from matplotlib.dates import WeekdayLocator, DayLocator, DateFormatter

## estimating the flow by fitting a curve

def bernstein_poly(i, n, t):
    """
    The Bernstein polynomial of n, i as a function of t
    """

    return comb(n, i) * (t**(n-i)) * (1 - t)**i

def bezier_curve(points, nTimes=1000):
    nPoints = len(points)
    xPoints = np.array([p[0] for p in points])
    yPoints = np.array([p[1] for p in points])
```

```

t = np.linspace(0.0, 1.0, nTimes)

polynomial_array = np.array([ bernstein_poly(i, nPoints-1, t) for i in range(0, nPoints) ])

xvals = np.dot(xPoints, polynomial_array)
yvals = np.dot(yPoints, polynomial_array)

return xvals, yvals

# replacing similar words to a common word
to_replace = {'(?i)(.*)corona(.*)': 'corona',
              '(?i)covid(.*)': 'corona',
              '(?i)ncov(.*)': 'corona',
              '(?i)virus(.*)': 'corona',
              '(?i)chin(.*)': 'china',
              '(?i)wuhan(.*)': 'china',
              '(?i)hubei(.*)': 'china',
              '(?i)province(.*)': 'china',
              '(?i)pence(.*)': 'trump',
              '(?i)(.*)donald(.*)': 'trump',
              '(?i)case(?i)': 'cases',
              '(?i)distancing': 'SD',
              '(?i)social': 'SD',
              '(?i)(.*)american(.*)': 'US',
              '(?i)(.*)american(.*)': 'US',
              '(?i)(.*)death(.*)': 'deaths',
              '#(.*)': '',
              '^[-+]?[0-9]+$': '',
              '"', "'", '\.': '',
              'us': 'US'}

def replace(token):
    for word,rep in to_replace.items():
        token = re.sub(word, rep, token)
    return token

# read in LDA data
# list of top 20 words in each topic per day
topics = defaultdict(list)
with open('LDA/LDA.csv') as f:
    data = [x.strip().split(',') for x in f.readlines()]
for d in data:
    topics[d[0]].extend([replace(word) for word in d[2:]])

cleaned = {}
# for each day, get the ratio of word usage
for d,values in topics.items():
    c = Counter(values) # get a counter
    s = sum(c.values()) # get total sum of counter
    c = c.most_common() # order counter by frequency
    cleaned[d] = {x:y/s for x,y in c if len(x) > 1} # calculate ratio

## convert to dataframe
## replace SD to social distancing (for graph label)
df = pd.DataFrame(cleaned).fillna(0).T

```

```

df.index = pd.to_datetime(df.index)
df.rename({'SD': 'social distancing'}, axis=1, inplace = True)

df = df[df.index < '2020-04-11'] # filter data to april 10

def create_data(words, init, spacing = 0.005):
    """
    calculate the top and bottom of the flow for each word
    """
    t = df[words].T # subset data

    tops = {date:{x:None for x in words} for date in t.columns} # initialize dictionary to store location of words in graph
    bottoms = {date:{x:None for x in words} for date in t.columns} # initialize dictionary to store location of words in graph

    spacing = spacing # set a spacing value between flows

    for date in t.columns: # for each day
        inds = t[date].rank(method='min').sort_values().index # get a rank ordering of words
        areas = t[date].sort_values().values # get an ordering of the frequency ratio
        bottoms[date][inds[0]] = init # for the plot, the bottom value of the first word is the initial starting point
        tops[date][inds[0]] = init+areas[0] # for the plot, the top value of the first word is the initial point plus its ratio

        for i in range(1,len(inds)): # for the rest of the words, update the same way but add a spacing in between each one
            bottoms[date][inds[i]] = tops[date][inds[i-1]] + spacing
            tops[date][inds[i]] = bottoms[date][inds[i]] + areas[i]

    return t, bottoms, tops

# choose words of interest for LDA
words = ['china', 'pandemic', 'trump', 'social distancing', 'US', 'home', 'lockdown', 'deaths']
t, bottoms, tops = create_data(words, init = 0, spacing = 0.005) # get placements for flow diagram

# store the hashtag values in an array
hashtag_ratios = np.array(list(maxs_dict.values()))
# add hashtag symbol
word_labels = ['#' + x for x in np.array(list(maxs_dict.keys()))]
# replace character with a dash - the original dash character can't be viewed on matplotlib
word_labels = [x if x != '#covid 19' else '#covid-19' for x in word_labels]

l = []

for word in t.index:
    points = list(zip(range(len(t.columns)), t.loc[word].values))
    xbez, ybez = bezier_curve(points, nTimes=79 * 15)
    l.append(ybez[:-1])

temp = pd.DataFrame(l)
temp.index = t.index

temp = temp.T.groupby(pd.qcut(temp.T.index, 79)).mean()
temp.index = t.columns
temp = temp.T
temp.columns = pd.to_datetime(temp.columns)

```

```

temp[pd.Timestamp('2020-02-23 00:00:00')] = np.nan

data = temp[sorted(temp.columns)]

ranks = data.rank(ascending=False, method="dense")
fig = plt.figure(figsize=(13, 4))
ax = fig.add_subplot(111)
for x, word in enumerate(words):
    ax.fill_between(data.columns, ranks.loc[word, :] - 4 * data.loc[word, :], ranks.loc[word, :] + 4 * data.loc[word, :],
                    color="C{}".format(x), lw=1, alpha=0.5, zorder=2, label = word)
    ax.fill_between([pd.Timestamp('2020-02-22'), pd.Timestamp('2020-02-23')],
                    [ranks.loc[word, '2020-02-22'] - 4 * data.loc[word, '2020-02-22'],
                     ranks.loc[word, '2020-02-22'] + 4 * data.loc[word, '2020-02-24']],
                    [ranks.loc[word, '2020-02-22'] - 4 * data.loc[word, '2020-02-22'],
                     ranks.loc[word, '2020-02-22'] + 4 * data.loc[word, '2020-02-24']],
                    color="C{}".format(x), lw=1, alpha=0.5, zorder=2)

legend = ax.legend(loc='lower center', bbox_to_anchor=(0.5, -0.35),
                  fontsize = 12, ncol = 4, frameon = False)

ax.set_xlim(pd.Timestamp("2020-01-22"), pd.Timestamp("2020-04-10"))
ax.xaxis.set_major_locator(WeekdayLocator())
ax.xaxis.set_minor_locator(DayLocator())
ax.xaxis.set_major_formatter(DateFormatter("%b %d"))
ax.set_ylabel("Rank", fontsize=12)

ax.set_ylim(0.1, 8.5)
ax.yaxis.set_major_locator(matplotlib.ticker.MultipleLocator(1))

ax.invert_yaxis()
# ax.grid(which="minor", b=True, alpha=0.3, zorder=1)
# ax.grid(b=True, zorder=1)
ax.tick_params(labelsize=12)

ax.set_title("Topic Evolution of COVID-19 Tweets", fontsize = 12);
plt.axvspan('2020-02-23', '2020-02-24', facecolor="0.7", alpha=0.5, label="No data")

plt.savefig('Topic_evolution.png', dpi=300,
          bbox_inches='tight', bbox_extra_artists=[legend])

```

RQ2, RQ3:

```

import matplotlib as mpl
import matplotlib.dates as mdates
import matplotlib.pyplot as plt
import matplotlib.ticker as mticker

import pandas as pd

mpl.rcParams['figure.figsize'] = (16, 4)
mpl.rcParams['text.usetex'] = True
mpl.rcParams['text.latex.preamble'] = [r'\usepackage{amsmath,amssymb,amsthm,bbm}']

tab10 = plt.get_cmap('tab10').colors
fontsize = 28
labelsize = 22
tiny_labelsize = 16

```

```
ms = 6
```

```
_CITY_COLOR = tab10[0]
_COVID_COLOR = tab10[1]
_NEG_COLOR = tab10[3]
_POS_COLOR = tab10[2]
_ND_COLOR = tab10[7]
_KOBE_COLOR = tab10[9]
_WHO_COLOR = tab10[6]
_LD_COLOR = tab10[5]
```

```
_WHO_LABEL = r"WHO declares pandemic"
_KOBE_LABEL = r"Kobe Bryant dies"
```

```
_START = "2020-01-01"
_MISSING_DATE = "2020-02-23"
_START_COVID = "2020-01-22"
_END = "2020-04-09"
_WHO = "2020-03-11"
_KOBE = "2020-01-26"
_DAY = pd.Timedelta("1 day")
```

```
_FIG_FLDR = "figures/"
_DATA_FLDR = "results/"
```

```
_PD = pd.read_csv(_DATA_FLDR + "public_health_paper_plot_data.tsv", sep="\t", parse_dates=True, index_col=0)
```

```
fig = plt.figure(figsize=(16, 6))
ax = fig.add_subplot(111)
```

```
ax.plot(_PD.index, _PD["COVID_tweet_quantity"].values, lw=2, color="k", linestyle="--", label=r"Tweets")
ax.plot(_PD.index - pd.Timedelta(hours=12), _PD["COVID_tweet_trend"], lw=2, color="k", label=r"Trend")
ax.scatter(_PD.index, _PD["COVID_tweet_CP"], color="k", s=2 * ms ** 2, marker="d",
           edgecolor="k", facecolor="w", label=r"Change-points", zorder=3)
```

```
ax.set_title("Activity for a subsample of the user-timeline data (20 Major US cities)",
            fontsize=labelsize)
```

```
ax.fill_between([pd.Timestamp(_KOBE), pd.Timestamp(_KOBE) + _DAY], 480000, 900000,
               color=_KOBE_COLOR, alpha=0.5, label=_KOBE_LABEL, zorder=1)
ax.fill_between([pd.Timestamp(_WHO), pd.Timestamp(_WHO) + _DAY], 480000, 900000,
               color=_WHO_COLOR, alpha=0.5, label=_WHO_LABEL, zorder=1)
ax.fill_between([pd.Timestamp("2020-03-16"), pd.Timestamp("2020-03-22")], 480000, 900000,
               color=_LD_COLOR, alpha=0.24, label=r"LD (NOLA)", zorder=1)
ax.fill_between([pd.Timestamp("2020-03-22"), pd.Timestamp("2020-03-24")], 480000, 900000,
               color=_LD_COLOR, alpha=0.4, label=r"LD (5 cities)", zorder=1)
ax.fill_between([pd.Timestamp("2020-03-24"), pd.Timestamp("2020-04-01")], 480000, 900000,
               color=_LD_COLOR, alpha=0.84, label=r"LD (16 cities)", zorder=1)
ax.fill_between([pd.Timestamp("2020-04-01"), pd.Timestamp(_END)], 480000, 900000,
               color=_LD_COLOR, alpha=1, label=r"LD (20 cities)", zorder=1)
```

```
ax.set_xlim(pd.Timestamp(_START_COVID), pd.Timestamp(_END))
ax.xaxis.set_major_locator(mdates.WeekdayLocator())
ax.xaxis.set_minor_locator(mdates.DayLocator())
ax.xaxis.set_major_formatter(mdates.DateFormatter("%b %d"))
```

```
ax.set_ylabel(r"Number of tweets", fontsize=labelsize)
ax.set_ylim(480000, 900000)
ax.yaxis.set_major_locator(mticker.MultipleLocator(100000))
```

```

ax.yaxis.set_minor_locator(mticker.MultipleLocator(10000))
ax.legend(fontsize=tiny_labelsize, ncol=3, loc="upper left",
          handlelength=1, handletextpad=0.5, framealpha=1, columnspacing=0.5, labelspacing=0.25)
ax.tick_params(labelsize=labelsizesize)

fig.savefig(_FIG_FLDR + "overall_tweet_quantity.png",
            dpi=600, papertype='letter', orientation='portrait',
            facecolor='w', edgecolor='w', bbox_inches='tight')

fig = plt.figure(figsize=(16, 8))
ax = fig.add_subplot(111)
ax2 = fig.add_subplot(111)
tax = ax.twinx()
tax2 = ax.twinx()

tax.plot(_PD.index, _PD["COVID_VADER_smooth"].values,
         color=_COVID_COLOR, lw=5, zorder=2)
tax.plot(_PD.index, _PD["COVID_VADER_smooth_filled"].values,
         color=_COVID_COLOR, lw=5, linestyle="--", zorder=2)

ax.set_zorder(2)
ax.patch.set_visible(False)
ax.plot(_PD.index, _PD["CITY_VADER_smooth"].values, color=_CITY_COLOR, lw=5, zorder=3)

tax2.set_zorder(3)
tax2.patch.set_visible(False)
tax2.plot(_PD.index, _PD["COVID_VADER_avg"].values, marker="X",
          ms=ms, mec="k", color=_COVID_COLOR, lw=0, zorder=4)
tax2.scatter(_PD.index, _PD["COVID_VADER_CP"].values, s=2 * ms**2, marker="d",
             edgecolor=_COVID_COLOR, facecolor="w", zorder=5)

ax2.set_zorder(6)
ax2.patch.set_visible(False)
ax2.plot(_PD.index, _PD["CITY_VADER_avg"].values, color=_CITY_COLOR, marker="X", lw=0,
         ms=ms, mec="k", zorder=6)
ax2.scatter(_PD.index, _PD["CITY_VADER_CP"].values, s=2 * ms**2,
            marker="d", edgecolor=_CITY_COLOR, facecolor="w", zorder=7)

tax.fill_between([pd.Timestamp(_MISSING_DATE), pd.Timestamp(_MISSING_DATE) + _DAY], -1, 1,
                 color=_ND_COLOR, alpha=0.5, zorder=1)
tax.fill_between([pd.Timestamp(_KOBE), pd.Timestamp(_KOBE) + _DAY], -1, 1,
                 color=_KOBE_COLOR, alpha=0.5, zorder=1)
tax.fill_between([pd.Timestamp(_WHO), pd.Timestamp(_WHO) + _DAY], -1, 1,
                 color=_WHO_COLOR, alpha=0.5, zorder=1)
tax.fill_between([pd.Timestamp("2020-03-16"), pd.Timestamp("2020-03-22")], -1, 1,
                 color=_LD_COLOR, alpha=0.24, zorder=1)
tax.fill_between([pd.Timestamp("2020-03-22"), pd.Timestamp("2020-03-24")], -1, 1,
                 color=_LD_COLOR, alpha=0.4, zorder=1)
tax.fill_between([pd.Timestamp("2020-03-24"), pd.Timestamp("2020-04-01")], -1, 1,
                 color=_LD_COLOR, alpha=0.84, zorder=1)
tax.fill_between([pd.Timestamp("2020-04-01"), pd.Timestamp(_END)], -1, 1,
                 color=_LD_COLOR, alpha=1, zorder=1)

ax.set_title("Average VADER sentiment for both data sets", fontsize=labelsizesize)

ax.set_xlim(pd.Timestamp(_START_COVID), pd.Timestamp(_END))
ax.xaxis.set_major_locator(mdates.WeekdayLocator())
ax.xaxis.set_minor_locator(mdates.DayLocator())
ax.set_ylim(0.05, 0.2)

```

```

ax.set_ylabel(r"VADER sentiment Score (user-timeline data)", fontsize=labelsize, color=_CITY_COLOR)
ax.tick_params(axis='y', colors=_CITY_COLOR)

tax2.set_ylim(-0.2, 0.01)
tax2.yaxis.set_major_formatter(mticker.NullFormatter())

tax.set_ylim(-0.2, 0.01)
tax.set_ylabel(r"VADER sentiment Score (COVID-19 corpus)", fontsize=labelsize, color=_COVID_COLOR)
tax.tick_params(axis='y', colors=_COVID_COLOR)

ax2.set_ylim(0.05, 0.2)

for a in [ax, ax2, tax, tax2]:
    a.set_xlim(pd.Timestamp(_START_COVID), pd.Timestamp(_END))
    a.xaxis.set_major_locator(mdates.WeekdayLocator())
    a.xaxis.set_minor_locator(mdates.DayLocator())
    a.xaxis.set_major_formatter(mdates.DateFormatter("%b %d"))

    a.yaxis.set_major_locator(mticker.MultipleLocator(0.05))
    a.yaxis.set_minor_locator(mticker.MultipleLocator(0.005))
    a.tick_params(labelsize=labelsize)

tax.plot([], [], color=_COVID_COLOR, lw=5, marker="X", ms=ms, mec="k", label=r"COVID-19 corpus")
tax.plot([], [], color=_CITY_COLOR, lw=5, marker="X", ms=ms, mec="k", label=r"user-timeline data")
tax.scatter([], [], marker="d", s=2 * ms ** 2, edgecolor="k", facecolor="w", label=r"Change-points")
tax.fill_between([], -1, 1, facecolor=_ND_COLOR, alpha=0.5, label=r"No data (COVID-19 corpus)")
tax.fill_between([], -1, 1, facecolor=_KOBE_COLOR, alpha=0.5, label=_KOBE_LABEL)
tax.fill_between([], -1, 1, facecolor=_WHO_COLOR, alpha=0.5, label=_WHO_LABEL)
tax.fill_between([], -1, 1, facecolor=_LD_COLOR, alpha=0.24, label=r"LD (NOLA)")
tax.fill_between([], -1, 1, facecolor=_LD_COLOR, alpha=0.40, label=r"LD (5 cities)")
tax.fill_between([], -1, 1, facecolor=_LD_COLOR, alpha=0.84, label=r"LD (16 cities)")
tax.fill_between([], -1, 1, facecolor=_LD_COLOR, alpha=1, label=r"LD (20 cities)")

tax.legend(fontsize=tiny_labelsize, ncol=5, loc="lower right",
           handlelength=1, handletextpad=0.5, framealpha=1, columnspacing=0.5, labelspacing=0.25)

fig.savefig(_FIG_FLDR + "VADER_average.png",
           dpi=600, papertype='letter', orientation='portrait',
           facecolor='w', edgecolor='w', bbox_inches='tight')

fig = plt.figure(figsize=(10, 4))
ax = fig.add_subplot(111)

ax.set_title("Distribution of VADER sentiment for COVID-19 corpus", fontsize=labelsize)

ax.step(_PD.index - pd.Timedelta(hours=12), _PD["COVID_VADER_perc_pos"].values,
       color="k", linestyle="-", where="post", lw=2, zorder=3)
ax.fill_between(_PD.index - pd.Timedelta(hours=12), _PD["COVID_VADER_perc_pos"].fillna(0),
               step="post", facecolor=_POS_COLOR, alpha=0.5, zorder=1, label="Positive")
ax.fill_between(_PD.index - pd.Timedelta(hours=12), _PD["COVID_VADER_perc_pos"].fillna(100), 100,
               step="post", facecolor=_NEG_COLOR, alpha=0.5, zorder=1, label="Negative")

ax.fill_between([pd.Timestamp(_MISSING_DATE), pd.Timestamp(_MISSING_DATE) + _DAY], 0, 100,
               facecolor=_ND_COLOR, alpha=0.8, zorder=2, label="No data")
ax.fill_between([pd.Timestamp(_WHO), pd.Timestamp(_WHO) + pd.Timedelta("1 day")], 0, 100, lw=0,
               color=_WHO_COLOR, alpha=0.8, label=_WHO_LABEL, zorder=2)

ax.set_xlim(pd.Timestamp(_START_COVID), pd.Timestamp(_END))
ax.xaxis.set_major_locator(mdates.WeekdayLocator(byweekday=1, interval=2))

```

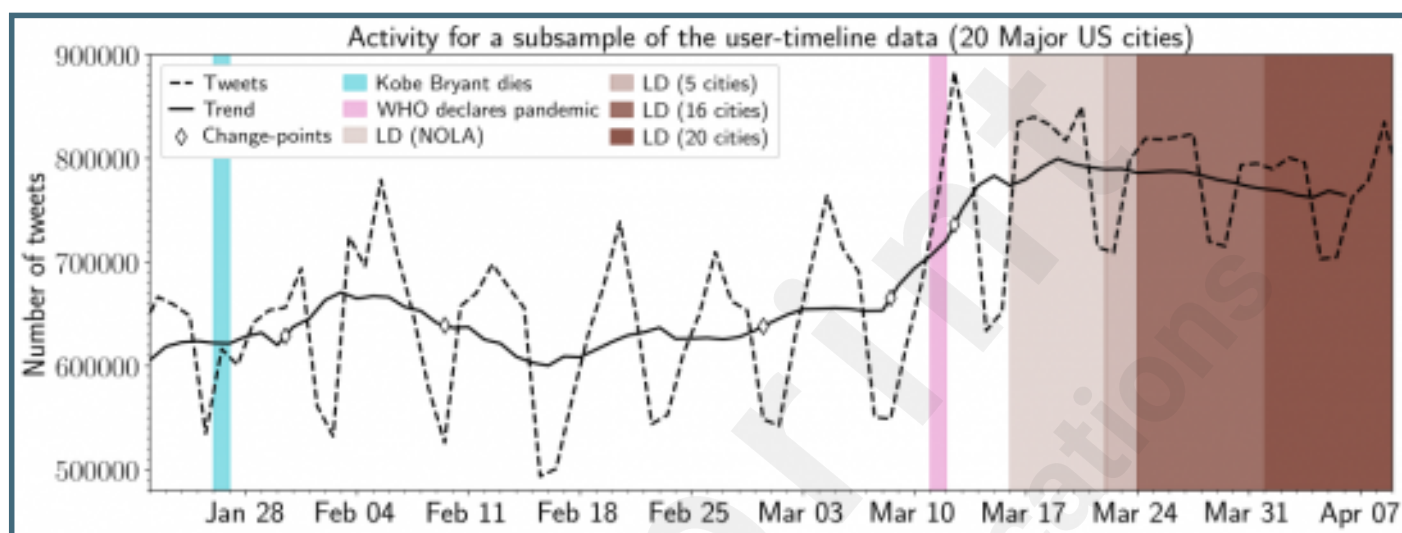
```
ax.xaxis.set_minor_locator(mdates.DayLocator())
ax.xaxis.set_major_formatter(mdates.DateFormatter("%b %d"))
ax.set_ylabel(r"Percentage of tweets", fontsize=22)
ax.set_ylim(34, 54)
ax.yaxis.set_major_locator(mticker.MultipleLocator(5))
ax.yaxis.set_minor_locator(mticker.MultipleLocator(0.5))
ax.yaxis.set_major_formatter(mticker.PercentFormatter())
ax.legend(fontsize=tiny_labelsize, ncol=2, bbox_to_anchor=(0.35, 1.025), loc="upper center",
          handlelength=1, handletextpad=0.5, framealpha=1, columnspacing=0.5, labelspacing=0.25)
ax.tick_params(labelsize=labelsizesize)

fig.savefig(_FIG_FLDR + "VADER_pos_neg.png",
            dpi=600, papertype='letter', orientation='portrait',
            facecolor='w', edgecolor='w', bbox_inches='tight')
```

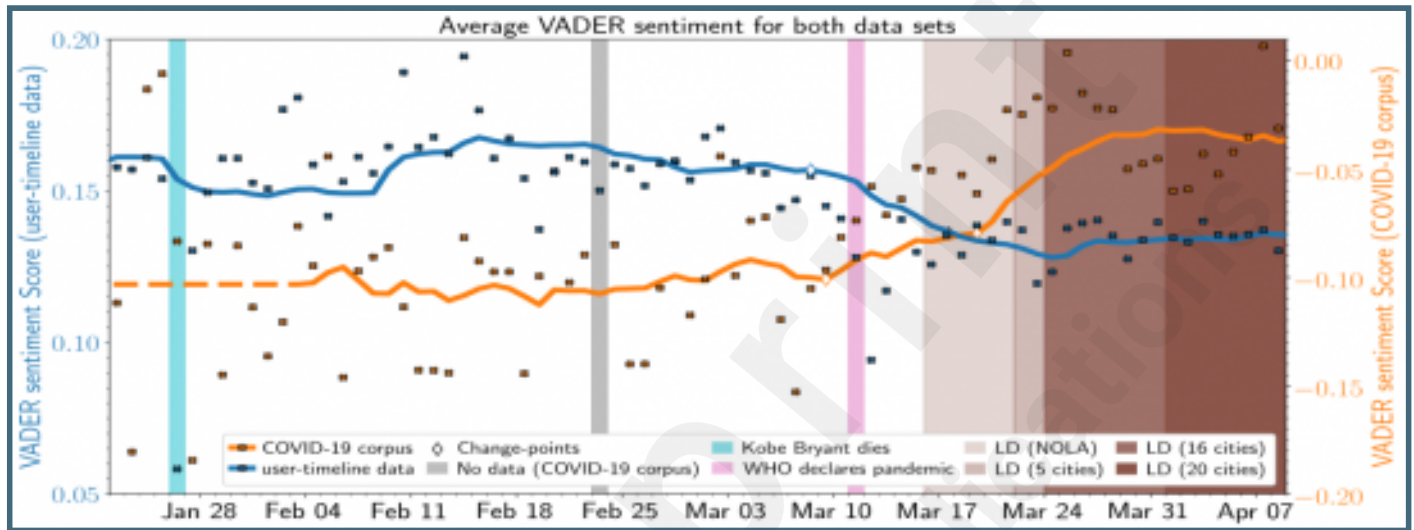
Supplementary Files

Figures

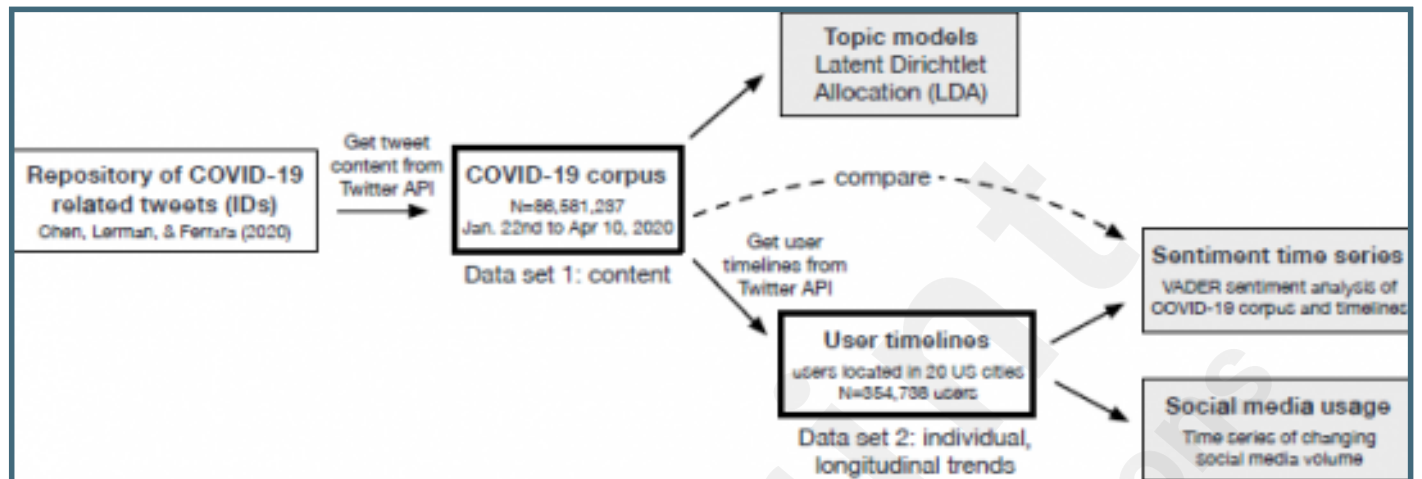
Number of daily tweets in a subsample of the user-timeline data ($n=292,000$). The dashed line displays the daily number of tweets in the subsample of the user-timeline data. The solid line indicates the underlying trend in the daily number of tweets (after removing its seasonal effects). The diamond markers indicate significant change-points indicated by the PELT-algorithm. The light-blue and pink annotations denote the death of Kobe Bryant (Jan 26, 2020) and the WHO declared COVID-19 as a global pandemic (March 11, 2020), respectively. The brown annotations indicate the dates at which lockdowns were enforced in the 20 considered cities (ranging March 16, 2020 to April 1, 2020). The opacity of the brown annotation indicates how many cities had enforced a lockdown at that date.



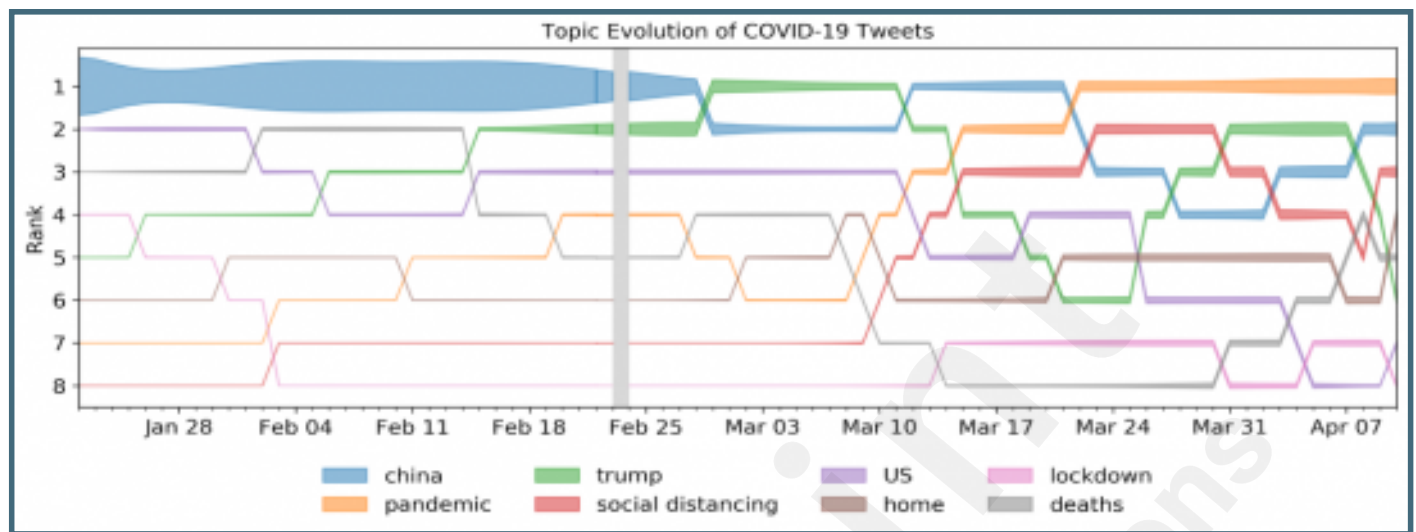
Distribution of non-zero VADER sentiment in the COVID-19 corpus & user-timeline data. The solid lines display the 14-day moving average and the crosses display the actual values of the daily average VADER sentiment (blue for the user-timeline data and orange for the COVID-19 corpus, respectively). The diamond markers indicate change-points indicated by the PELT-algorithm, the color of the edge of the diamonds refers to the time-series to which this change-point belongs. The light-blue, gray, and pink annotations denote the days of the death of Kobe Bryant (Jan 26, 2020), the day of missing data in the COVID-19 corpus (February 23, 2020), and the WHO declaring COVID-19 a global pandemic (March 11, 2020), respectively. The brown annotations indicate the dates at which lockdowns were enforced in the 20 considered cities (ranging March 16, 2020 to April 1, 2020). The opacity of the brown annotation indicates how many cities had enforced a lockdown at that date.



This diagram illustrates the procedure undertaken to procure Tweet IDs from an open-access COVID-19 repository. These tweet IDs were run through Twitter's API to create two distinct data sets. The COVID-19 corpus, which contains all COVID-19 related English-language tweets published in the United States between January-April 2020. The User time-line data contains the 3200 most recent tweets of users residing in the 20 cities most-affected by COVID-19.



Topic-group rankings over time in the COVID-19 corpus. Each topic is ranked by its frequency ratio. The width of the bars indicates the fraction of words in the topic on a given day. The colors of the areas indicate which theme the corresponding area belongs to.



Distribution of positive and negative VADER sentiment in COVID-19 corpus. The solid line displays the fraction of tweets that hold positive or negative VADER sentiment in the user-timeline data (colored green for positive and red for negative sentiment). The gray and pink annotations denote the day of missing data (February 23, 2020) and the WHO declared COVID-19 as a global pandemic (March 11, 2020), respectively.

