



What Are People Concerned About During the Pandemic? Detecting Evolving Topics about COVID-19 from Twitter

Chia-Hsuan Chang^{1,2} · Michal Monselise¹ · Christopher C. Yang¹

Received: 4 November 2020 / Revised: 4 November 2020 / Accepted: 23 November 2020 /

Published online: 17 January 2021

© The Author(s), under exclusive licence to Springer Nature Switzerland AG part of Springer Nature 2021

Abstract

With the novel coronavirus (COVID-19) pandemic affecting the lives of the citizens of over 200 countries, there is a need for policy makers and clinicians to understand public sentiment and track the spread of the disease. One of the sources for gaining valuable insight into public sentiment is through social media. This study aims to extract this insight by producing a list of the most discussed topics regarding COVID-19 on Twitter every week and monitoring the evolution of topics from week to week. This research will propose two topic mining that can handle a large-scale dataset—rolling online non-negative matrix factorization (Rolling-ONMF) and sliding online non-negative matrix factorization (Sliding-ONMF)—and compare the insights produced by both techniques. Each algorithm produces 425 topics over the course of 17 weeks. However, topics that have not evolved from one week to the next beyond a certain evolution threshold are consolidated into a single topic. Since the topics produced by the Rolling-ONMF algorithm each week depend on the topics from the previous week, we find that the Sliding-ONMF algorithm produces more varied topics each week; however, the topics produced by the Rolling-ONMF algorithm contain keywords that appear more consistent with each other when reviewing the terms manually. We also observe that the Sliding-ONMF algorithm is able to capture events that have shorter time frames rather than ones that last throughout many months while the Rolling-ONMF algorithm detects more general themes due to a higher average evolution score which leads to more topic consolidation. We have also conducted a qualitative analysis and grouped the detected topics into themes. A number of important themes such as government policy, economic crisis, COVID-19-related updates,

This article belongs to the Topical Collection: *Special Collection on COVID-19*
Guest Editors: Christopher C. Yang, Julio Cesar Facelli, David Buckeridge, Fei Wang

Editor: Julio Cesar Facelli

✉ Chia-Hsuan Chang
sham82503@gmail.com

Extended author information available on the last page of the article.

COVID-19-related events, prevention, vaccines and treatments, and COVID-19 testing are identified. These reflected the concerns related to the pandemic expressed in social media.

Keywords Coronavirus · Twitter analysis · Topic analysis · Topic evolution detection · Social monitoring

1 Introduction

Since the coronavirus disease (COVID-19) was first detected in November 2019, cases of infection by the virus have been reported in over 200 countries and caused hundreds of thousands of deaths. In addition to the death and illness that the virus has caused globally, it has also caused the disruption of daily life in many countries. As different countries began to implement measures to combat the pandemic, we witnessed different phenomena that were documented by the media but also discussed on social media. Hospitals began to fill up with patients in many countries. The scarcity of healthcare resources brought about the need for healthcare workers to make tough choices and decide which patients should be prioritized. Such cases have been documented in Italy [2], Spain, the Netherlands, and France [27]. Furthermore, as the pandemic progressed, healthcare experts had more information regarding what makes some people more vulnerable than others to the virus. As it became apparent that the elderly and those with pre-existing conditions were the most vulnerable, the world observed a sad trend of deadly outbreaks in nursing homes [3]. Starting in late 2019, there has been a constant stream of information in both news outlets and social media regarding the virus. Since the pandemic originated in China and then quickly spread to South Korea and Italy, there are a number of studies that have already managed to observe public behavior on social media during the initial months of the pandemic in these countries. A study by Li et al. identified a correlation between posts on Weibo related to COVID-19 and the number of reported cases in Wuhan [13]. This study shows that social media can be used to augment information from news media since it is correlated with official government reports. A study by Park et al. examining Twitter activity in South Korea during the beginning of the pandemic found that monitoring social media for activity can help health practitioners in uncovering information that affects their decisions regarding health policy [18]. By monitoring social media, and particularly Twitter, we can gather useful information about the most discussed news stories in real time as well as public sentiment regarding these stories and about the pandemic in general. Previous research has shown that crowd-sourced information from the Internet may be used to predict and explain disease outbreaks [9]. To assist in the efforts of epidemiologists and public health experts, we may use crowd-sourced data to understand public sentiment and concerns in real time. Mining social media may provide health practitioners and researchers with access to information that may not always be presented in the news media.

2 Related Work

We provide a brief review of research related to public sentiment regarding COVID-19 on Twitter, topic detection in social media, and detection of evolving events in social media.

2.1 Recent Literature on COVID-19 and Twitter

Since the start of the pandemic, multiple research projects have focused on studying Twitter to learn information about public sentiment and trends regarding COVID-19. Most of these research projects focused on mining tweets related to COVID-19 using a list of keywords and looking for a general trend in public sentiment regarding the pandemic. The study by Abd-Alrazaq et al. [1] aimed to uncover the top concerns of Twitter users during the pandemic by mining the social media platform for tweets related to COVID-19 and performing an analysis of the tweets. Likewise, the study by Chen et al. [5] aimed to develop a dataset of COVID-19-related tweets and perform an analysis of the tweets for general trends. Another study that performed an exploratory analysis of tweets is the study by Ordun et al. [17]. This research project used topic modeling, UMAP, and digraphs to study general trends in tweets. The study by Chen et al. [6] performed an analysis of tweets but focused on classifying the data into controversial and non-controversial terms discussing the pandemic. Kouzy et al. [11] performed an analysis of COVID-19-related tweets to detect misinformation regarding the pandemic. These research projects look at the entire time period in the tweet dataset and detect topics without separating the data into smaller time periods like weeks or examining the evolution of topics over time. However, our study aims to track the different trends in tweets over time and learn about the most discussed topics in near real time.

2.2 Algorithms for Topic Detection in Social Media

There are a number of algorithms that have been used for detecting emerging topics in social media and in Twitter in particular. Ibrahim et al. provided a survey of machine learning techniques for topic detection in Twitter [10]. They categorized detection algorithms into five types, including clustering, frequent pattern mining, exemplar-based technique, matrix factorization, and probabilistic models. Among these different kinds of algorithms, both probabilistic models and matrix factorization attract the most attention [7, 12, 14, 16, 24, 25, 29].

The most representative probabilistic model is latent Dirichlet allocation (LDA) [4], which assumes that one document consists of multiple topics, however with the fact that most tweets are made up of few sentences and tend to convey one topic per tweet. LDA may not work well in determining topics from tweets due to the limited information presented in individual tweets. One feasible solution is to pool several similar tweets into a long synthetic document [14, 24]. Then, the synthetic documents are served as targets for extracting hidden topics using LDA. The defect of this strategy is that the long pooled documents make it unscalable. On the other hand, an alternative approach is to relax the assumption. Several models, including

Dirichlet-multinomial mixture model (DMM) [29], GPU-DMM [12], and LF-DMM [16], generate each document by a topic rather than by a topic distribution, which extract more coherent topics from tweets. Nevertheless, the iterative sampling process behind these models is also the main obstacle to applying them to large-scale tweet collections.

Consequently, with the demand for a scalable and efficient topic detection algorithm, matrix decomposition is proposed as an optimal solution. Non-negative matrix factorization (NMF) [22] is one of the most influential topic detection methods, which extracts topical patterns by decomposing a bag-of-words matrix V . NMF also demonstrates its effectiveness in determining topics from several short-text contexts, including tweets [7]. Likewise, Suri and Roy compared the use of LDA and NMF in topic detection in tweet collections and found that LDA takes longer to run. Therefore, in real-time scenarios, NMF is preferred for Twitter topic detection [25].

2.3 Detection of Topic Evolution in Social Media

While the Internet did not exist in the last big pandemic of 1918, there have been smaller pandemics over the last few years that were tracked on social media. In our research, we have examined the methods used by these research projects to track emerging global health events using the Internet and social media. The study by Ginsberg et al. followed influenza epidemics by monitoring search engine queries [9] in different geographic areas over time. The research by Szomszor et al. used Twitter to track and understand the public response to the swine flu pandemic [26]. This research examines whether Twitter users preferred official news sources over untrustworthy sources like blogs. During the H1N1 pandemic, Signorini et al. [23] tracked disease activities and public concerns, e.g., case count, travel, and hygiene, from massive collected tweets using a keyword counting method. Compared to previous studies, we aim at extracting highly discussed topics and their evolution in an unsupervised manner so that the extracted topics can better reflect the common interests of users on Twitter.

To model the temporal topic patterns through NMF, Time Evolving Non-negative Matrix Factorization (TENMF) [21] was introduced. Given a tweet collection, all tweets will be separated into a set of sequential bins by their timestamps. The topic dictionary at time t is learned using NMF by taking the bin of time t as input. For transferring the learned topic dictionary across time, the resultant decomposed matrices at time $t - 1$ will be treated as initial matrices for NMF at time t . Hence, the learned topics are smoothly transformed in the process of time. Likewise, Dynamic NMF (DNMF) [20] was another feasible model proposed to address the temporal topic extraction. DNMF also requires time-sliced tweet bins. The difference is that DNMF takes one synthetic bin as input at a time. Each synthetic bin is constructed by grouping consecutive bins in a sliding window as input, for example, as the size of window is set on 2, the first synthetic bin contains tweets from first bin and second bin. DNMF determines two types of topics: evolving topics and emerging topics. The evolving topics are smooth evolution from previously extracted topics, and this works by limiting the variations of previous topics. The emerging topics are then extracted by finding the optimal word distributions that best describe the rising word trends. Both TENMF and DNMF demonstrate the strategy to model temporal topic patterns.

However, when we set 1 week as the granularity to create time-sliced tweet sets, the volume of tweets in each bin impede us to fit into memory and execute the following decomposition process. Considering collecting over 4 million tweets a week and having 100,000 word features, it requires about 372 GB to store the whole bag-of-words matrix in the float type. As a result, in this study, we will introduce an online NMF to process tweets streamingly, and its two temporal extensions inspired by TENMF and DNMF in a memory-efficient way.

3 Data Collection

We collected the COVID-19-related tweets using the Streaming API from Twitter. Using the API, we retrieved a massive amount of tweets that contained either “coronavirus” or “covid19” from March to June of 2020. To prevent from including duplicated tweets, we only retained the original English tweets in our dataset. That is, when we found any retweet and quoted tweet, we searched its original tweet and only stored the original one. Figure 1 shows the volume of tweets we collected. The volume of COVID-19-related tweets skyrocketed after March 8, when the first outbreak started in the USA. We then had 60.32M tweets in total for further topic analysis. Before that, we first pre-processed all tweets by removing stop words, numbers, links, mentions, emojis, and symbols. Since hashtags may carry on the hidden topical patterns, we kept them in tweets.

4 Topic Detection Model

In this section, we first introduce an online matrix factorization approach that is capable of detecting topics from a large number of tweets. Then, we propose two extensions that help in learning temporal topic patterns.

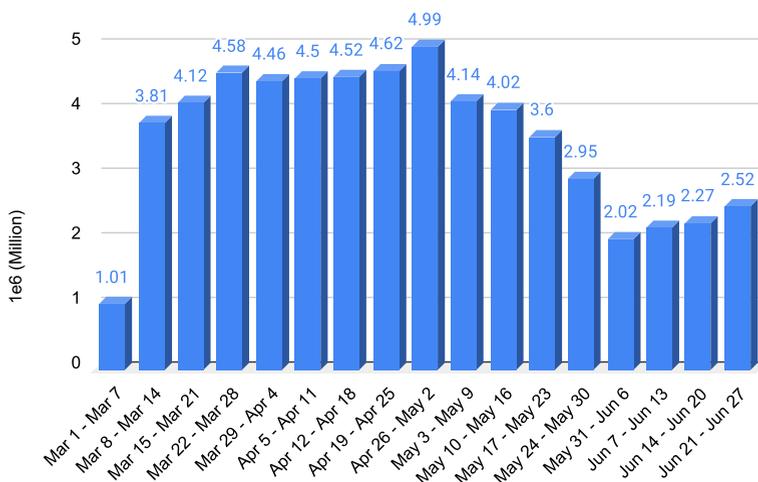


Fig. 1 The volume of tweets

4.1 Online NMF

A recent study by Chen et al. [7] exhibits that non-negative matrix factorization (NMF) prevails latent Dirichlet allocation, one of the most influential topic models, in several short text contexts. Hence, we choose NMF as our base scheme to analyze our collected tweets. Specifically, given the corpus $V \in \mathbb{R}^{F \times N}$ with F word features and N documents, the goal of NMF is to decompose V into the multiplication between two low-rank matrices as shown in (1).

$$\arg \min_{W, H} \|V - WH\|_F^2, \text{ s.t. } W, H > 0, \tag{1}$$

where $W \in \mathbb{R}^{F \times K}$, $H \in \mathbb{R}^{K \times N}$, and K is the number of topics defined manually. Usually, $K \ll F$. Through the decomposition, the resultant W encodes the weight of each word for each topic, and H reveals the topic tendency of each document. Yet, the canonical NMF is not tailored for analyzing topics from tweets. In particular, the large volume of tweets (both F and N increase) makes it very difficult to fit V into memory and its following decomposition operations. To illustrate this issue, we use the tweets collected from April 26 to May 2 to indicate this memory exhaustion problem. During that period, we collect 4.99M tweets ($N = 4.99 \times 1e6$) and use the most frequent 50K words as word features ($F = 50,000$), which requires about 232 GB of memory to fit the float typed matrix. Therefore, we need a model that has the capability of batch learning for extracting topical themes among texts. For example, if we split 4.99M into a set of batches with document size 2000, fitting each sub-matrix with the same word features only requires 0.18 GB at a time.

Considering the limitation of canonical NMF, we split V into Q data batches $\{V^q\}_{q=1}^Q$ as shown in Fig. 2, where each V^q has the same number of data samples, denoted s . Furthermore, we reformulate the (1) into sum of loss of a set of batches:

$$f(W) = \min_{W>0} \frac{1}{N} \sum_{q=1}^Q \sum_{i=1}^{V^q} \ell(v_i, W), \tag{2}$$

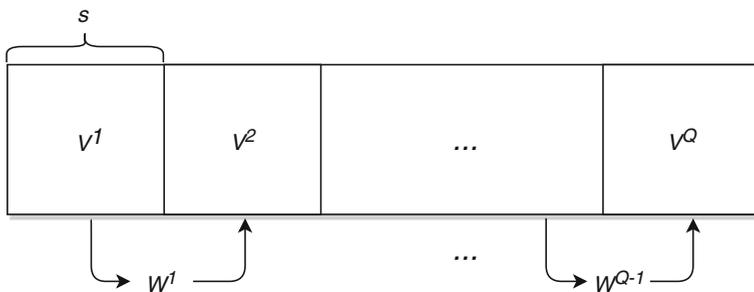


Fig. 2 The illustration of ONMF

where N is the total number of data samples, and v_i is one sample from V^q . We then define loss function of each sample $\ell(v_i, W)$ as (3) and (4).

$$\ell(v_i, W) \triangleq \min_{h_i \geq 0} \hat{\ell}(v_i, W, h_i), \quad (3)$$

where

$$\hat{\ell}(v_i, W, h_i) \triangleq \frac{1}{2} \|v_i - Wh_i\|_2^2 \quad (4)$$

To solve (2) by considering each batch sequentially, we leverage the stochastic majorization-minimization framework [31], in which solving (2) is separated into two steps: (1) learning H , namely non-negative encoding, and (2) updating W , namely dictionary update. Specifically, given a current batch V^q and previous optimized dictionary W^{q-1} , we first learn coefficient matrix H^q by

$$H^q = \arg \min_{H \geq 0} \hat{\ell}(V^q, W^{q-1}, H) \quad (5)$$

Note that when $q - 1 = 0$, i.e., the first batch, we randomly initialize a non-negative matrix W^0 as the input. Subsequently, given $\{V^q, H^q\}$, we update the basis matrix W^q by

$$W^q = \arg \min_{W > 0} \frac{1}{s} \sum_{v_i \in V^q, h_i \in H^q} \frac{1}{2} \|v_i - Wh_i\|_2^2 \quad (6)$$

Equation (6) is then can be rewritten as

$$W^q = \arg \min_{W > 0} \frac{1}{2} \text{tr}(W^T W A^q) - \text{tr}(W^T B^q), \quad (7)$$

where $A^q = \frac{1}{s} H^q (H^q)^T$ and $B^q = \frac{1}{s} V^q (H^q)^T$. To search the minimal results of the (5) and (7), we use the coordinate descent algorithm and steepest descent algorithm, respectively. Moreover, we only store the topic dictionary W^q , i.e., topic-word coefficients, during the training time. For inferring topic coefficients of each tweet, we further apply (5) after getting the final topic dictionary. As a result, we show the algorithm of this online NMF in the Algorithm 1, where we sequentially apply (5) (line 4) and (7) (line 6) on each batch, and consider the last topic dictionary to represent the whole corpus.

Algorithm 1 The algorithm of ONMF.

Input: data batches: $\{V^q\}_{q=1}^Q$, number of topics: K

1 **foreach** $q \in [1, Q]$ **do**

2 **if** $q - 1 = 0$ **then**

3 Randomly initialize $W^{q-1} \in \mathbb{R}^{F \times K}$

4 $H^q = \arg \min_{H \geq 0} \hat{\ell}(V^q, W^{q-1}, H)$

5 $s = |V^q|$; $A^q = \frac{1}{s} H^q (H^q)^T$; $B^q = \frac{1}{s} V^q (H^q)^T$

6 $W^q = \arg \min_{W > 0} \frac{1}{2} \text{tr}(W^T W A^q) - \text{tr}(W^T B^q)$

Output: The last dictionary $W^{q=Q}$

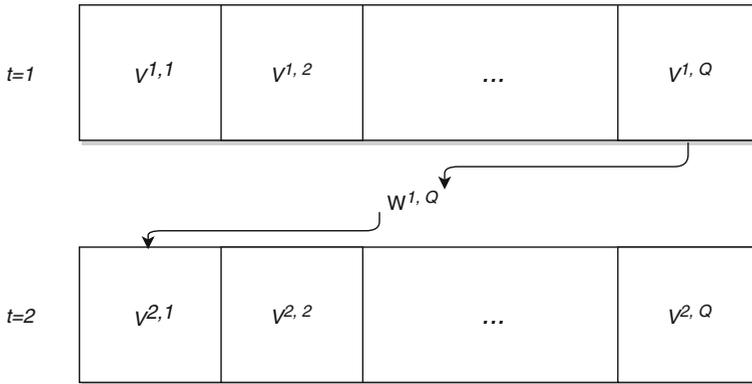


Fig. 3 The illustration of Rolling-ONMF

4.2 Rolling-ONMF

To extend the ONMF for catching the temporal patterns of topics, we discretize the whole corpus of tweets according to their timestamps into T time slices. The granularity of the time slice can be a day, a week, or any time span, which depends on the purpose. We then apply ONMF on each slice to extract the topical patterns of each week. For transferring the learned topics from $t - 1$ to t , as shown in Fig. 3, we use the last topic dictionary of $t - 1$, $W^{t-1,Q}$, as prior knowledge to guide the ONMF to learn topics from t . In this way, we only need to change the (5) into (8) at the adjacent batches between two time slices. The idea behind (8) is inspired by TENMF [21], which aims at learning topics from previous time to current time. With this objective, the topic i at time t is assumed as an evolution from the topic i in time $t - 1$.

$$H^{t,1} = \arg \min_{H \geq 0} \hat{\ell}(V^{t,1}, W^{t-1,Q}, H) \tag{8}$$

We also note that the shape of W should keep the same across time slices for applying the ONMF. Hence, we need to scan all possible vocabularies across all time slices to settle down the number of word features F at first¹.

4.3 Sliding-ONMF

With the requirement of fixed vocabularies of Rolling-ONMF, we propose Sliding-ONMF to release the constraint. Given the discretized time slices, Sliding-ONMF combines textual data across time slices based on a sliding window w into synthesized slices and apply ONMF on them. Take Fig. 4 as an example, as $w = 2$, Sliding-ONMF hereby has three different topic models as results. In each model, Sliding-ONMF learns the topic dictionary, W , by considering the whole tweets within a sliding window, but only infers the topic coefficient, H , for tweets of the later week in a sliding window. For example, in Model1, topic dictionary is determined from

¹We also keep the number of topics k invariant across times.

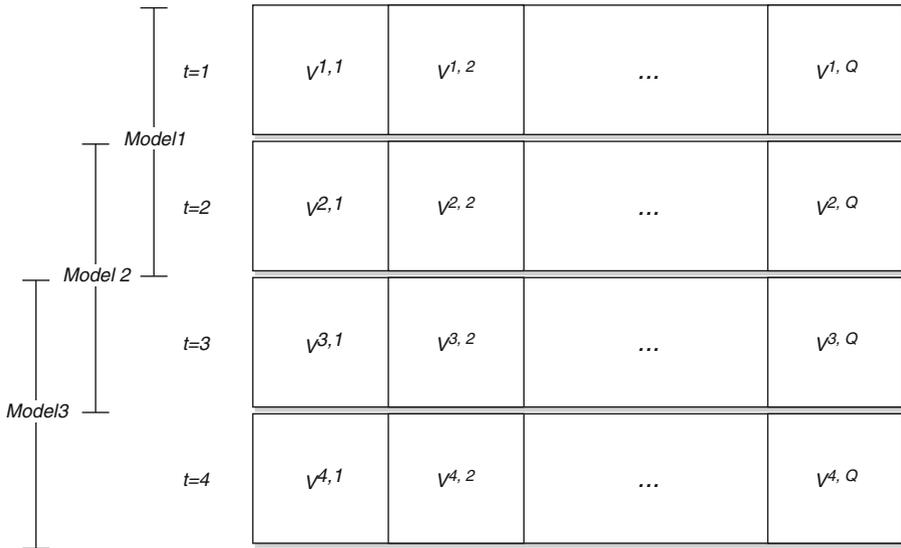


Fig. 4 The illustration of Sliding-ONMF

tweets at $t = 1$ and $t = 2$, and only infers the topic coefficient of tweets at $t = 2$. In this manner, we can infer temporal topic variations by comparing the determined topics between adjacent models (e.g., Model1 versus Model2 and Model2 versus Model3).

5 Topic Trend Extraction

After applying the Rolling-ONMF and Sliding-ONMF, we extract two topical information W^t and H^t that encode topic-word dictionary and topic-tweet coefficient at time slice t , respectively. Note that we strip batch index q because only the topic dictionary of the last batch at each time slice, $W^{t,q=Q}$, will be used. After optimizing the topic dictionary, we can infer the topic coefficient for all tweets at the time slice. To analyze the topic trends across time slices, we rely on both W^t and H^t to determine the topic evolution and construct topic-volume tuples over time.

5.1 Determining the Topic Evolution

This step aims to calculate the evolution score of each topic from $t - 1$ to t . For measuring the value, we transform the topic dictionary into a set of discrete word distributions by (9), where each column represents a word distribution of a topic.

$$\hat{W}_{i,j}^t = \frac{W_{i,j}^t}{\sum_{i=1}^F W_{i,j}^t}, \forall i \in F, \forall j \in K \quad (9)$$

With the word distributions of each topic at time $t - 1$ and time t , we can measure the level of evolution by considering the word overlap between topics across two

time slices. We measure the overlap using the Jaccard Index and propose the topic evolution detection algorithm in Algorithm 2, where $\hat{W}_{1:C,k}^{t-1}$ is a set of top- C probable words of topic k in \hat{W}^{t-1} , and $\text{JaccardIndex}(A, B) = \frac{A \cap B}{A \cup B}$.

Algorithm 2 The algorithm for detecting topic evolution.

Input: two matrices containing topic-word distributions at time $t - 1$ and t : \hat{W}^{t-1} , \hat{W}^t , number of topics K , and number of top probable words C

- 1 Initialize an evolution matrix $J \in \mathbb{R}^{K \times K}$ with 0 for each entry
- 2 **foreach** $k \in [1, K]$ **do**
- 3 **foreach** $k' \in [1, K]$ **do**
- 4 $J[k, k'] = \text{JaccardIndex}(\hat{W}_{1:C,k}^{t-1}, \hat{W}_{1:C,k'}^t)$

Output: J

Consequently, we can analyze J for inferring the topic evolution between $t - 1$ and t . For example, the higher score of $J[1, 2]$ indicates the more likely that topic 2 of time t continues the topic 1 of time $t - 1$.

5.2 Constructing Topic-Volume Tuples Over Time

To catch the topic trend over time, we construct a set of tuples $\{(k, t, r)\}$, where k is a topic, t indexes a time slice, and r is the volume of the topic k at t . Specifically, with J over time, we consolidate topics across adjacent time slices when their evolution scores surpass the threshold τ , leading to a single topic k . In this process, we also record the time slice index t of each involved topic. Subsequently, given k and t , we can easily populate the topic volume r by averaging the score in k^{th} row in H^t . As a result, for each topic k , we will filter out the low-volume topics whose ratio below average value $1/K$ (e.g., if $K = 25$, the average value is 4%). With the tuples, we then manually assign the topic label for each topic k and draw its topic ratio over time.

6 Result

6.1 Parameter Setting

To apply both Rolling-ONMF and Sliding-ONMF ($w=2$) on our collected tweets, we discretized every tweet by each week, resulting in 17 time slices ($T = 17$) ranging from the first week of March to the last week of June. This also represents that we learned the topics at the week level and monitored their trends from week by week. The batch size s was set on 2000, which only processes 2000 tweets at a time to control memory usage. Since the tweets in different time slices might carry a different number of topics, it was hard and impractical to use perplexity or coherence measurement to determine the best number like in previous studies. For this reason, we experimentally set the number of topics K as 25.

6.2 Model Performance

To measure the performance of the extracted topics, we adopted two measurements: coherence score and diversity. The coherence score, proposed by Mimno et al. [15], determines the coherent level of a topic model by the following steps: (1) each topic is represented by its top- N contributed words, (2) the words are then paired up by the combination C_2^N , (3) co-occurrence statistics, such as pointwise mutual information, of all word pairs are applied, and (4) averaging co-occurrence statistics is used as a coherence score of a topic model. When a topic model performs a higher coherence score, it should be a better model for catching the word co-occurrence patterns into topics. On the other hand, when a topic model extracts more diverse topics, it potentially provides more meaningful and recognizable information. Hence, diversity is another key measure of the extracted topics. To measure this, each topic is represented by the top- N contributed words. Then, we adopt the Jaccard index to measure similarities between topics. As a result, the diversity of a topic model is computed by $1 - \text{average Jaccard index}$. The higher diversity means the more diverse topics extracted from a topic model.

Table 1 presents the performances of Rolling-ONMF and Sliding-ONMF.

Number of Topics With 25 topics per week, each model determined 425 topics through the 17 weeks. To connect the evolved topics between weeks, we used 0.2 for τ for consolidating topics across time slices. Insignificant topics were removed using average ratio 4% (i.e., $1/K = 1/25$) as criteria. As a result, there were 49 out of 425 topics and 97 out of 425 topics left for Rolling-ONMF and Sliding-ONMF, respectively. There are roughly double the topics remaining after consolidation in Sliding-ONMF compared to Rolling-ONMF. We can consolidate more topics due to a higher mean evolution score (0.2996 for Rolling-ONMF vs. 0.2612 for Sliding-ONMF). We also present the different evolution characteristics of both models in Fig. 5, where a brighter color means more similarity shared between topics. Figure 5a shows that most topics in Week12 of Rolling-ONMF are evolved from the same topics in Week11, so the brighter values center on the diagonal. On the contrary, in Fig. 5b, Sliding-ONMF determines the topics in Week12 without considering topics in Week11, which explains why Sliding-ONMF can extract more topics with different themes.

Coherence Score Due to the unequal sample size (i.e., number of topics), we adopted Welch's unequal variance t test for testing the statistical significance. The coherence score of Sliding-ONMF is significantly higher (p value = 0.049) than that of Rolling-ONMF. We attributed the worse coherence score of Rolling-ONMF to the constraint

Table 1 Model performance

	Number of topics	Coherence score (std)	Diversity (std)
Rolling-ONMF	49	-6.12 (1.65)	0.84 (0.07)
Sliding-ONMF	97	-5.61 (1.46)	0.82 (0.09)

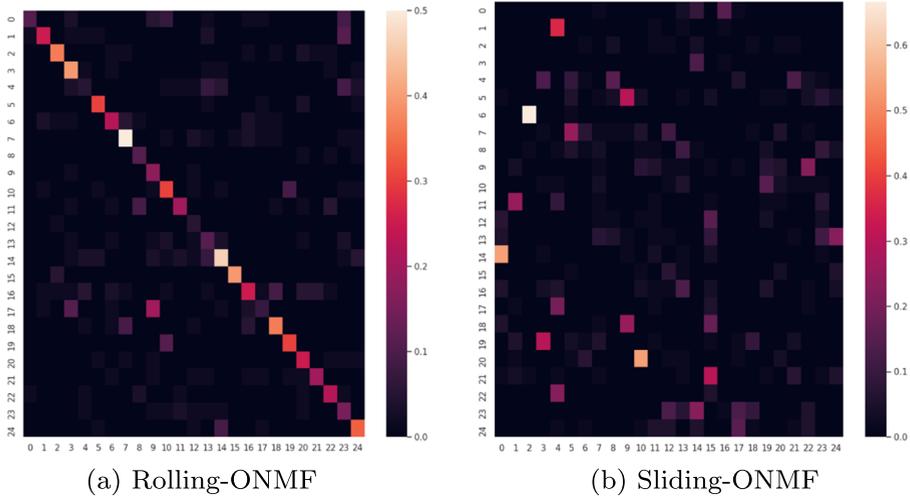


Fig. 5 The evolution matrix of Rolling-ONMF and Sliding-ONMF. The x-axis is the topics from Week12, and the y-axis is the topics from Week 11. Each cell represents the evolution score estimated by Algorithm 2

that it needs to learn the topics based on learned topics from previous time slices. Sliding-ONMF is free from such a constraint, so it can better catch the topics from the corpus. In previous studies [20], the continuity of topics helps topics to evolve smoothly and helps investigators to evaluate the topics. Yet, this constraint harms the coherence score in our study.

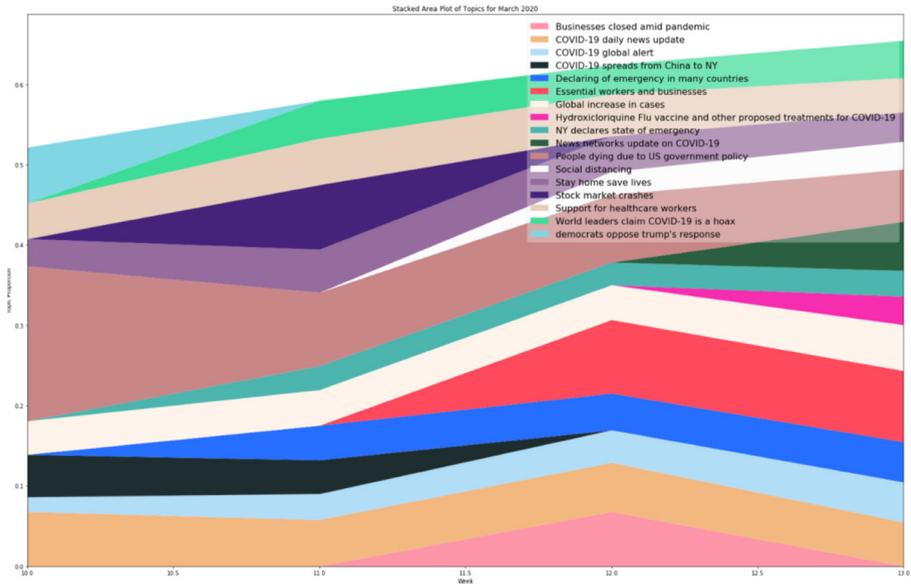
Diversity The diversity of Rolling-ONMF is higher than that of Sliding-ONMF with p value 0.07. As a result, the two models have different strengths and weaknesses. When the goal is to extract diversified topics, one may use Rolling-ONMF. When the number of coherent topics is the main consideration, one should instead adopt Sliding-ONMF.

7 Discussion

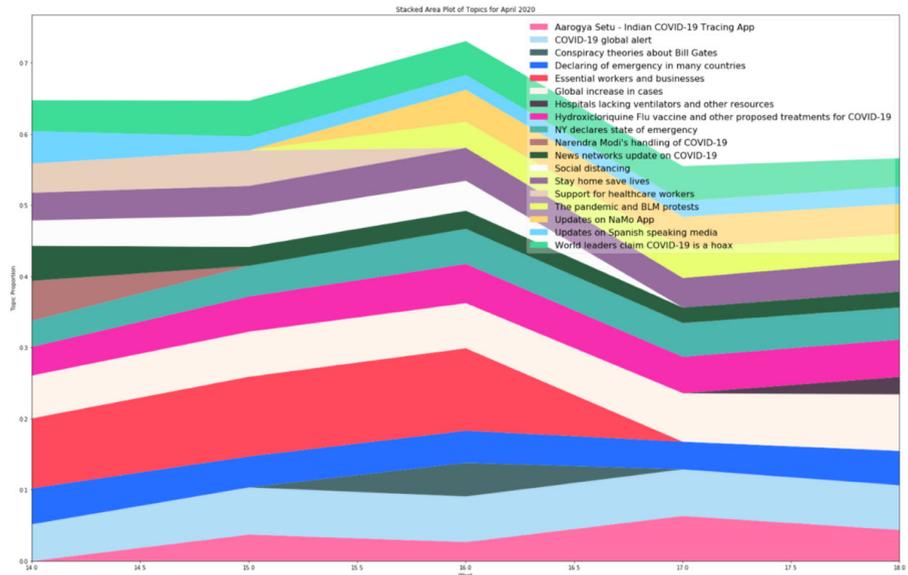
As discussed in the experimental result section, Rolling-ONMF and Sliding-ONMF have their own strengths and weaknesses. The topics detected by these two methods are not the same. The total number of topics and the duration of the topics vary between the two methods. In this section, we discuss the difference between them. In addition, we categorize the detected topics and discuss the themes that are most concerning during the pandemic.

7.1 Rolling-ONMF

The topics in Rolling-ONMF seem to capture more general themes and fewer topics that discuss an event specific to a single week. The topics and their proportion of the entire tweet dataset are described in the area plots in Fig. 6.

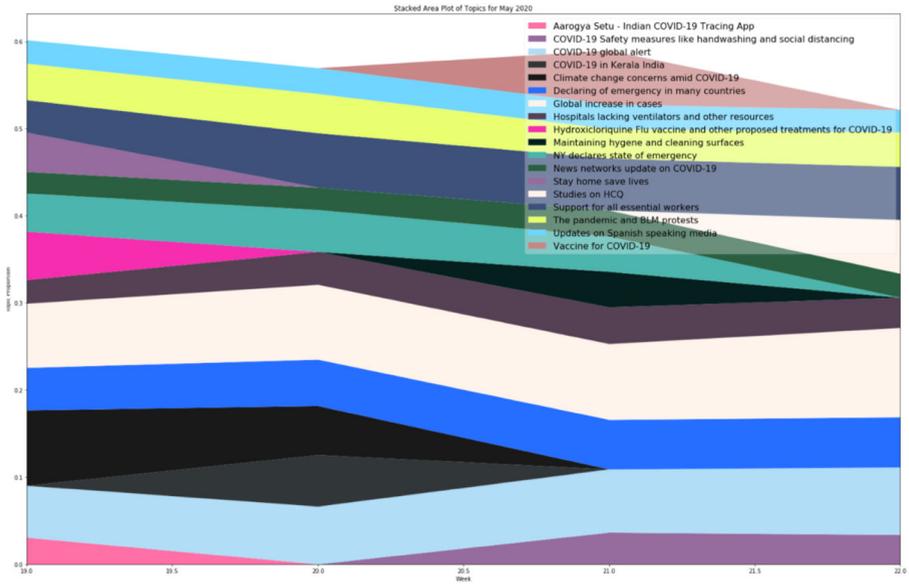


(a) March 2020 Rolling-ONMF Stacked Area Plot

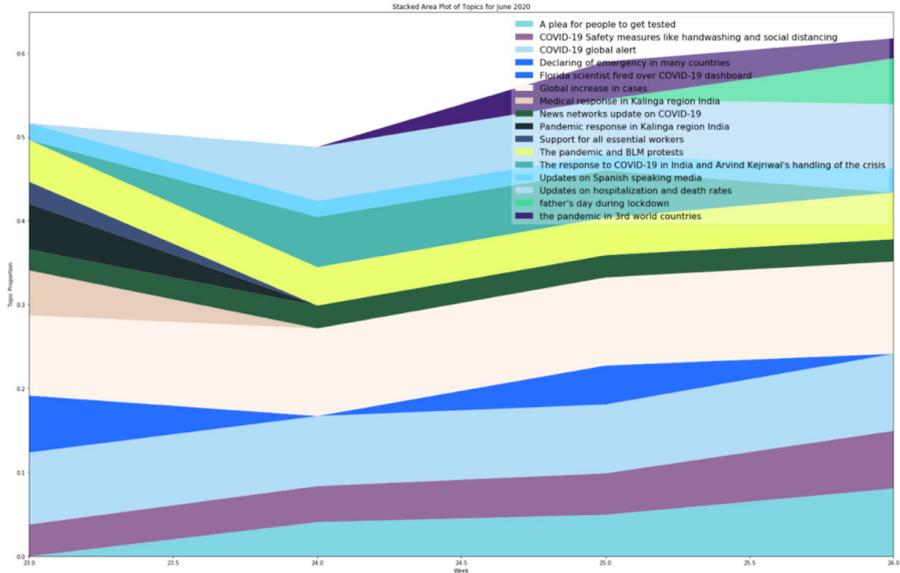


(b) April 2020 Rolling-ONMF Stacked Area Plot

Fig. 6 The stacked area plots for months March–June 2020 describing the persistent and short-term topics observed in the Rolling-ONMF



(c) May 2020 Rolling-ONMF Stacked Area Plot



(d) June 2020 Rolling-ONMF Stacked Area Plot

Fig. 6 (continued)

The topics that persisted the most weeks throughout the 4 months are Global increase in cases and COVID-19 global alert. These topics are more general and not specific to an event that happened in a single week which is why they persisted for

17 weeks each. On average, a topic persisted for 5 weeks with 14 topics persisting for only 1 week.

7.2 Sliding-ONMF

The Sliding-ONMF algorithm produced more topics after consolidating using the evolution score. The topics and their proportion of the entire tweet dataset are described in the area plots in Fig. 7.

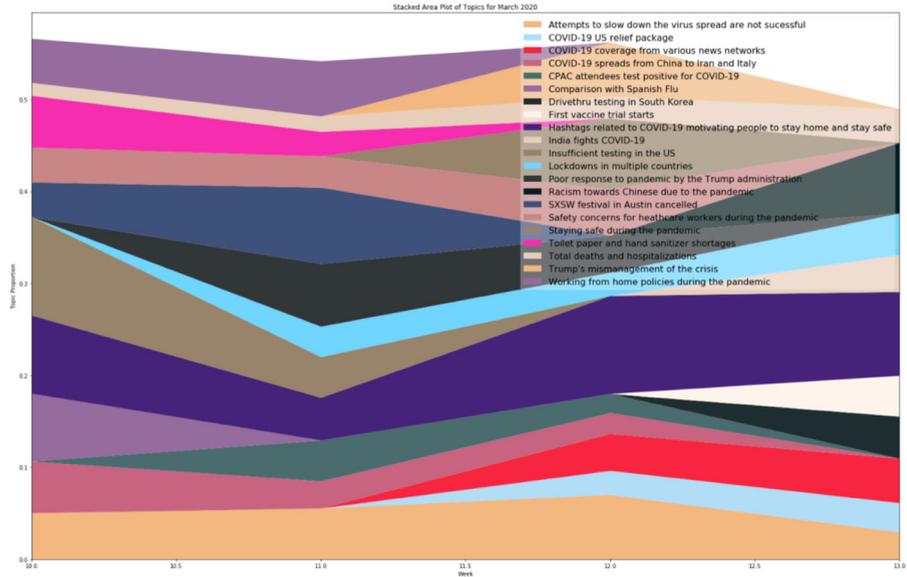
The most persistent topic using this algorithm is Total Deaths and Hospitalization that persisted for 17 weeks. Unlike the Rolling-ONMF algorithm, the mean number of weeks that a topic persisted here is only 2.67 with 37 topics only persisting for 1 week. This is consistent with the fact that we were able to capture more information unique to each week rather than more general trends. For example, in the Sliding-ONMF, there are topics related to the SXSW festival canceled in March and the Together at Home Concert in April.

7.3 Themes Detected in the Topics

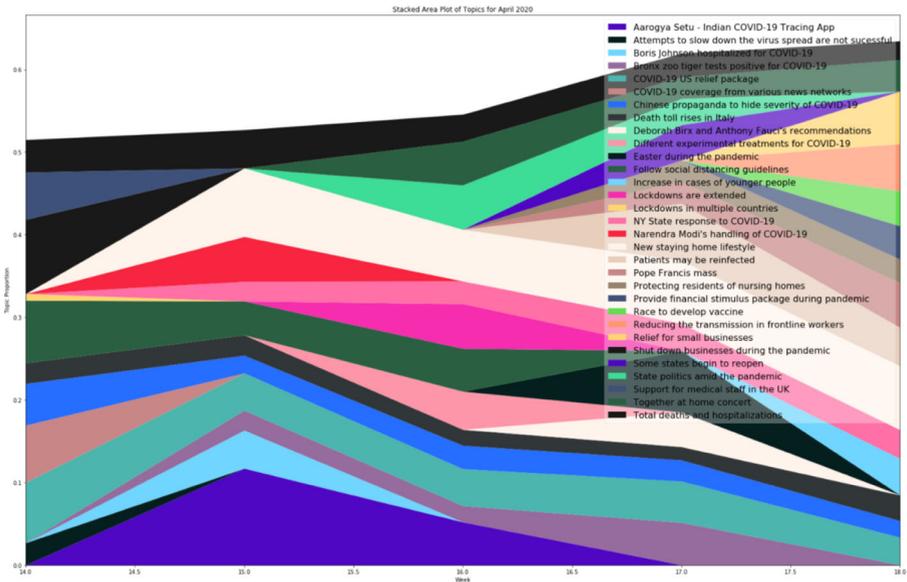
After performing a step of automated consolidation and filtering of all topics using the evolution score and the proportion of each topic out of the weekly total, all topics were manually reviewed for qualitative analysis. All three authors reviewed the topics in search of common themes. The authors grouped the topics into 11 distinct themes in both the Rolling-ONMF and the Sliding-ONMF algorithms. Themes were selected in a way that reflected whether a topic occurred in 1 or 2 weeks or lasted throughout many weeks. Themes were also reflective of the different types of tweets and separated between tweets discussing hand washing and social distancing and those tweets related to vaccine development and potential COVID-19 treatments. Additionally, incoherent topics were manually filtered out. For example, a topic that centered around the word Dr and contained information about Dr Fauci, Dr Birx, and other doctors from other countries was deemed to be incoherent in the qualitative analysis since it did not contain any information about an event that happened during a certain week or a theme that lasted for a number of weeks. After the manual examination, 49 topics were reduced to 39 from the Rolling-ONMF algorithm and 97 topics were reduced to 74 from the Sliding-ONMF algorithm. We observe that while there was more consolidation of topics in Rolling-ONMF since the themes were more consistent from one week to the other, the topics that did remain after consolidation in Rolling-ONMF were more coherent in their themes. On the other hand, the Sliding-ONMF algorithm produced more distinct topics from week to week but a larger proportion of these topics ended up being removed from the analysis since the keywords did not have a common coherent theme related to something that occurred during that week. We observe that the volume of the themes in the rolling method is higher because of the continuity that is a part of this algorithm. In the next section, we present the qualitative results of the two models for comparison.

After all three authors examined the topics, there were a number of themes that emerged both in the Rolling-ONMF topics and the Sliding-ONMF topics. The

stacked area chart for these themes is detailed in Fig. 8, examples of selected theme are detailed in Table 2, and the weekly average proportion for each theme is detailed in Table 3.

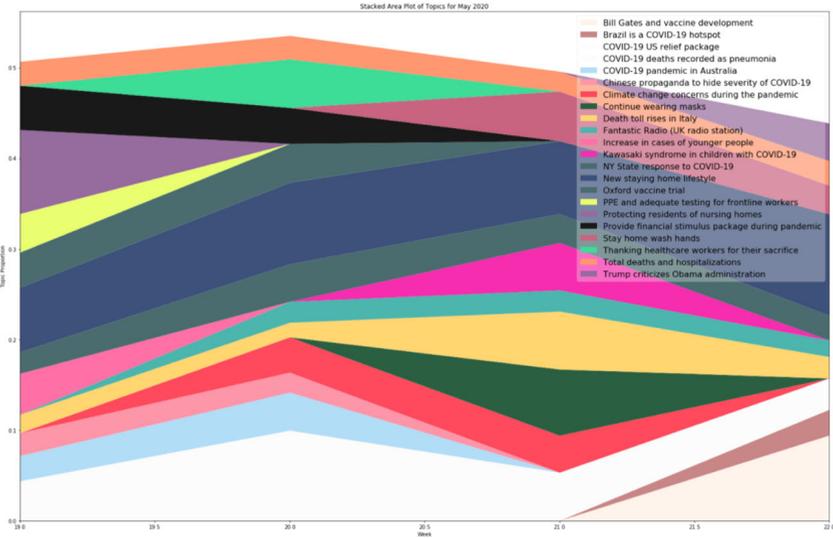


(a) March 2020 Sliding-ONMF Stacked Area Plot

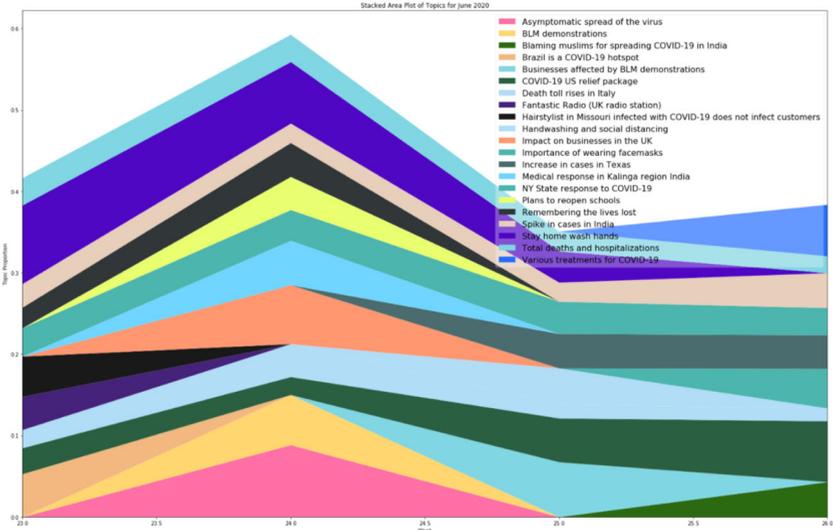


(b) April 2020 Sliding-ONMF Stacked Area Plot

Fig. 7 The stacked area plots for months March–June 2020 describing the persistent and short-term topics observed in the Sliding-ONMF



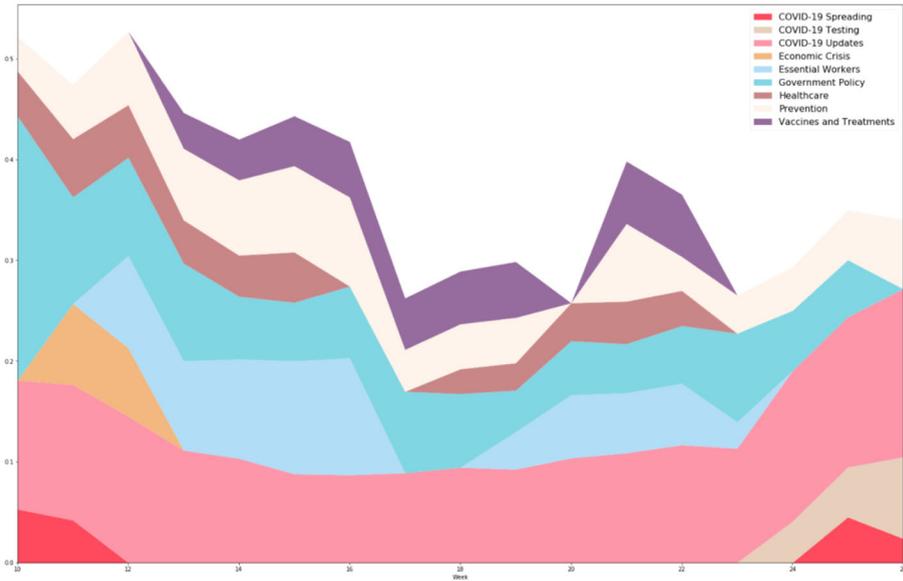
(c) May 2020 Sliding-ONMF Stacked Area Plot



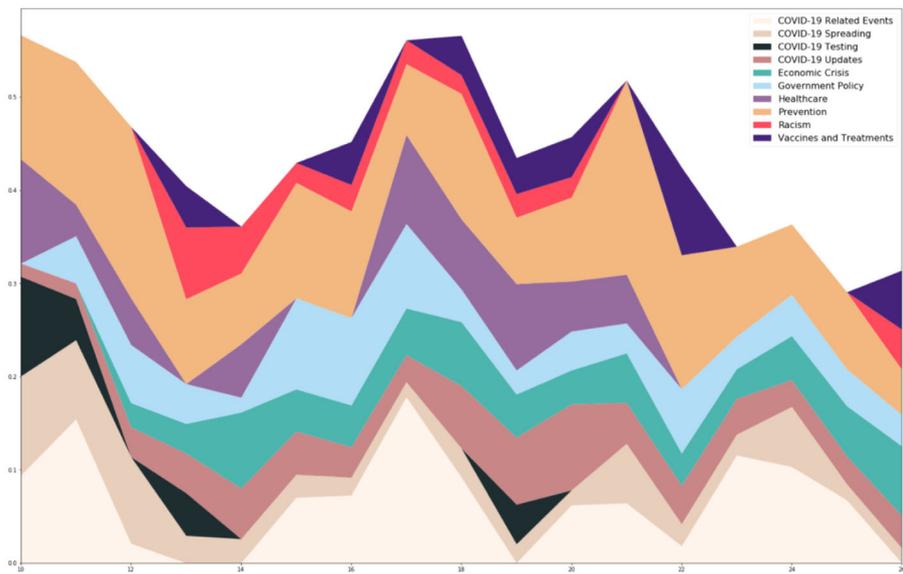
(d) June 2020 Sliding-ONMF Stacked Area Plot

Fig. 7 (continued)

When examining Table 2, we observe a result that is consistent with our previous observations regarding the differences between the Rolling-ONMF and the Sliding-ONMF algorithms. The topics detected by the Rolling-ONMF algorithm are more general and mostly span over longer periods of time, while the topics in the Sliding-ONMF algorithm are more specific and span over shorter periods of time. In fact, the average number of weeks that each topic persists in the Rolling-ONMF algorithm is



(a) Overall Themes Rolling-ONMF Stacked Area Plot



(b) Overall Themes Sliding-ONMF Stacked Area Plot

Fig. 8 The stacked area plots for months March–June 2020 describing the overall themes observed in the Sliding-ONMF and Rolling-ONMF

5.1 weeks while they persist an average of 2.67 weeks in the Sliding-ONMF algorithm. The topics in the two algorithms do not have many weeks of overlap since they capture different types of topics. Since the topics in the Sliding-ONMF algorithm

persist for much less time, they tend to discuss specific events that occurred mostly in a single week or pertain to a specific news event. For example, if we compare the topics in the healthcare theme, in the Rolling-ONMF algorithm, we see more general topics like support for healthcare workers or ventilator shortages. However, in the Sliding-ONMF algorithm, we see more specific topics that lasted between 1 and 3 weeks each and discuss more specific events like Kawasaki syndrome in children and cases of reinfection that circulated in the news.

7.3.1 COVID-19-Related Events

COVID-19-related events are defined as news events that were reported due to COVID-19 or events happening directly related to COVID-19. Topics discussing related events typically only showed up for 1 or 2 weeks at a time. Since there was more consolidation of topics in the Rolling-ONMF algorithm, there are no event topics in this algorithm. The detection of more distinct topics related to events in the Sliding-ONMF was most likely due to the effect in the rolling method that the new topics are not easily identified if they are very different from the topics in the previous week. In the Sliding-ONMF algorithm, we see a number of events including the cancellation of the SXSW festival in Austin appearing in the first 2 weeks of March, CPAC attendees testing positive for COVID appearing in the second and third weeks of March, the Together at Home Concert appearing in April, Easter during the pandemic in April, climate change-related headlines in May, a hairstylist testing positive in Missouri but not infecting her customers in May, and Black Lives Matter (BLM) demonstrations starting in late May. However, we observe that due to more consolidation in the Rolling-ONMF algorithm, the BLM topic appears in April since it has been algorithmically consolidated with previous topics regarding the mistreatment of African Americans during the pandemic. In the Sliding-ONMF algorithm, however, this topic only appears in early June.

7.3.2 COVID-19-Related Updates

This theme differs from the related events theme by mainly discussing updates about the death toll and containing hashtags like #coronavirusupdate. The COVID-19 updates theme had the highest weekly average proportion in the Sliding-ONMF algorithm (11.5%). This theme lasted throughout the period and appeared prominently in both the Rolling-ONMF and Sliding-ONMF algorithms. This is likely due to more consolidation occurring in the Rolling-ONMF algorithm which led to topics that were more likely to center around more general keywords and hashtags. In the Rolling-ONMF algorithm, topics carry on from week to week as they are very similar to those being detected in the previous week. Therefore, topics that contained more general themes were more prominent.

7.3.3 Prevention

Prevention-related topics have the highest average proportion in Sliding-ONMF (11.12%) while in Rolling-ONMF they have a weekly average proportion of only

Table 2 Selected themes and topics for Rolling-ONMF and Sliding-ONMF

Selected themes	Rolling-ONMF topics	Rolling-ONMF Weeks range	Sliding-ONMF topics	Sliding-ONMF Weeks range
COVID-19-related Events	Conspiracy theories about Bill Gates	19, 20	Hairstylist in Missouri infected with COVID-19 does not infect customers Pope Francis mass SXSW festival in Austin cancelled Together at home concert Toilet paper and hand sanitizer shortages	23 17, 18 10, 11 16, 17, 18 10, 11
	Florida scientist fired over COVID-19 dashboard	16 25		
	The pandemic and BLM protests father's day during lockdown	16–26 26		
COVID-19 spreading	COVID-19 spreads from China to NY the pandemic in 3rd world countries	10, 11 25, 26	Asymptomatic spread of the virus Attempts to slow down the virus spread are not successful COVID-19 spreads from China to Iran and Italy	24 10–14 10, 11, 12
	COVID-19 testing	A plea for people to get tested	Insufficient testing in the USA PPE and adequate testing for frontline workers Drivethru testing in South Korea	10, 11 19 13
COVID-19 updates	COVID-19 daily news update	10–13	Increase in cases of younger people Total deaths and hospitalizations Brazil is a COVID-19 hotspot	18, 19 10–26 22, 23
	COVID-19 global alert	10–26		
	Global increase in cases	10–26		
Economic crisis	Businesses closed amid pandemic	12	Relief for small businesses	18
	Stock market crashes	11	Shut down businesses during the pandemic COVID-19 US relief package	14 12–26

Table 2 (continued)

Selected themes	Rolling-ONMF topics	Rolling-ONMF Weeks range	Sliding-ONMF topics	Sliding-ONMF Weeks range
Essential workers	Essential workers and businesses	12–16		
	Support for all essential workers	19–23		
Government policy	Narendra Modi's handling of COVID-19	14	Lockdowns are extended	16
	Pandemic response in Kalinga region India	23	NY State response to COVID-19	15–26
	NY declares state of emergency	11–21	Some states begin to reopen	17
	Declaring of emergency in many countries	11–23	Trump criticizes Obama administration	22
	World leaders claim COVID-19 is a hoax	11–18	Lockdowns in multiple countries	11–14
Healthcare	Hospitals lacking ventilators and other resources	18–22	Kawasaki syndrome in children with COVID-19	21
	Support for healthcare workers	10–15	Patients may be reinfected	17, 18
			Protecting residents of nursing homes	17, 18, 19
Prevention	COVID-19 safety measures like hand washing and social distancing	21–26	Handwashing and social distancing	25
	Maintaining hygiene and cleaning surfaces	21	Hashtags related to COVID-19 motivating people to stay home and stay safe	
	Social distancing	10–13	Importance of wearing facemasks	26
	Stay home save lives	12–16	New staying home lifestyle	15–22
		10–19		

Table 2 (continued)

Selected themes	Rolling-ONMF topics	Rolling-ONMF Weeks range	Sliding-ONMF topics	Sliding-ONMF Weeks range
Racism			Blaming Muslims for spreading COVID-19 in India	26
			Chinese propaganda to hide severity of COVID-19	14–20
			Racism towards Chinese due to the pandemic	13
Vaccines and treatments	Hydroxycloquine, flu vaccine and other proposed treatments for COVID-19	13–19	Bill Gates and vaccine development	22
	Studies on HCQ	22	Different experimental treatments for COVID-19	16
	Vaccine for COVID-19	21	First vaccine trial starts	13
			Oxford vaccine trial	19, 20

Table 3 Average weekly theme proportion

Theme	Average weekly proportion Sliding-ONMF	Average weekly proportion Rolling-ONMF
COVID-19 updates	0.115	0.04
Essential workers	0.075	0
Government policy	0.075	0.054
Economic crisis	0.074	0.05
COVID-19-related events	0.069	0.086
Prevention	0.058	0.112
COVID-19 testing	0.057	0.06
Vaccines and treatments	0.051	0.053
Healthcare	0.041	0.069
COVID-19 spreading	0.041	0.04
Racism	0	0.035

5.76%. Topics related to prevention of COVID-19 include tweets that center around the themes of social distancing, maintaining hygiene and hand washing, wearing masks, and staying home like the hashtag “stay home save lives.” This topic persisted throughout the entire period examined in this research. We observe that tweets regarding staying home and preventing COVID started from the first week of March even though quarantine had not been in place in many places in the world during that time. In fact, the restaurant booking website reports a decline in year over year bookings in the USA starting March 2nd even though the first stay home order issued in the USA was in California effective March 20th.²

7.3.4 Economic Crisis

Topics related to the economic crisis discuss shutdowns of businesses due to the crisis, stock market crashes, government relief packages to small businesses, providing financial stimulus packages due to the pandemic, and the impact to business in different countries. The topic started a week into the pandemic and lasted throughout in the Sliding-ONMF algorithm but in the Rolling-ONMF algorithm, it was only discussed in the second and third weeks of the pandemic. Some of the topics in this theme are more news related (i.e., the stock market crashes that peaked in late March and early April), while others have a general theme that persists throughout the entire period. In particular, we see a discussion of the stock market crash begins in the second week of March in the Rolling-ONMF algorithm. Additionally, many tweets center around the topic of the COVID-19 relief package in the USA starting in the 3rd week of March and up until June.

²<https://www.opentable.com/state-of-industry>

7.3.5 Government Policy

Government policy is a theme that encompasses responses from different governments around the world. The theme lasted throughout the 17-week period with both algorithms. Topics in this theme include the declaration of a state of emergency in multiple countries as well as US states, the response of different local-level and country-level leaders and statements from different leaders regarding the pandemic. Many of the tweets are from individuals reacting to the new regulations. Over time, we see signs of quarantine fatigue. This phenomenon has also been documented in other research [30].

7.3.6 COVID-19 Spreading

Topics related to the rapid spread of the disease appeared throughout the entire period in the Sliding-ONMF algorithm but only in the first two and last two weeks in the time period in the Rolling-ONMF algorithm. The topics in this theme discuss the unsuccessful efforts to mitigate the virus as it spread initially from China to Italy and Iran, the discussion of rising death tolls, and spread from asymptomatic carriers of the virus. When observing the individual tweets, we see concern by individuals over the spread of the virus. In the early weeks of our tracking efforts, there are many tweets from Italy containing concern from individuals over the state of the country in dealing with the pandemic.

7.3.7 Healthcare

This theme contained a mix of news story-based topics regarding healthcare as well as general support for healthcare workers around the world that was expressed by Twitter users. Some examples of topics are the discovery that children who are infected with COVID-19 may get Kawasaki syndrome as a result, the concern that people who have recovered from COVID-19 may be reinfected, and concerns about infections in nursing homes. Since this topic was driven by news stories, it did not appear consistently throughout the entire time period. These patterns match the research that has emerged regarding nursing homes [3] and the tough choices that healthcare workers have had to make [27].

7.3.8 Vaccines and Treatments

This theme did not appear throughout the entire time period and was heavily driven by news stories and updates about the progress made in developing vaccines and discovering various treatment options for COVID-19. There are discussions of updates from the different vaccine trials like the Oxford and AstraZeneca trial, but also discussions of conspiracy theories like tweets related to Bill Gates and vaccines. It is important to track these mentions of conspiracy theories and the public sentiment towards, as has been shown in other research tracking controversial tweets during the COVID-19 pandemic [1, 6]. Research has shown that exposure to social media can

affect an individual's willingness to vaccinate [8]; therefore, it is crucial for clinicians and policy makers to be aware of public sentiment regarding the COVID-19 vaccine.

7.3.9 Racism

This theme only appeared in the Sliding-ONMF algorithm and had very few topics. Examples of topics were tweets discussing Muslims in India infecting others intentionally with COVID-19, and tweets with hashtags containing racist terms related to China and COVID-19. These topics represented a relatively small proportion of all tweets during the weeks that they appeared.

7.3.10 Essential Workers

The theme discussing essential workers appeared only in the Rolling-OMNF algorithm and was mainly related to the praise of essential workers. The two topics in this theme are essential workers and businesses and support for all essential workers. This theme lasted throughout most of the time period. Though there were some changes in the keywords representative of this theme, the theme continued to contain hashtags like #frontlineheroes and keywords discussing donation, businesses, and nurses.

8 Conclusion

Social media is steadily becoming the primary communication platform for people to convey their thoughts in almost every aspect. Therefore, analyzing the contents from tweets can help better understand people's reflections towards a specific event. During the COVID-19 pandemic, several studies used topic modeling for extracting the main concerns and demands of society. These previous studies [1, 5, 17] only focus on tweets accumulated in a short period (e.g., few weeks) and overlook the potential evolved topics. To track the topics through time, we first collected COVID-19-related tweets from the first week of March to the last week of June. We also developed two matrix decomposition methods—Rolling-ONMF and Sliding-ONMF—for identifying the hidden topics and their temporal evolution from the collected tweets. Compared to existing temporal topic models [20, 21], our proposed models could deal with the large scale of tweets and infer the hidden topics in a reasonable time. Two models have different strengths. Rolling-ONMF tends to catch high-volume topics and keep continuity. On the contrary, Sliding-ONMF is able to detect more emerging topics that appear in shorter time frames. We found that two types of topics can be identified using our methods. The first type is the emerging COVID-19-related event, which helps in recognizing the prompt event. For instance, we identified a topic related to the SXSW festival cancellation in March and the Together at Home Concert in April. The other type is the recurring topic, which helps the decision-maker and governor monitor the trend of a specific topic, such as reflections of government policy or vaccination development.

There are two potential future directions worth investigating further. Since our proposed methods only catch the continuity of topics across different times, the first

direction is how to learn more topic dynamics such as merge, split, and diminish [28]. The second direction is incorporating other informative dimensions. For instance, geolocation plays a vital role in tracking disease diffusion. It may additionally explain the relationship between the topic and its associated location [19]. Likewise, we observed that the co-occurrence pattern of hashtags is also as an essential indicator of a topic. We may propose a better mechanism to model hashtags rather than modeling them with plain texts.

References

1. Abd-Alrazaq A., Alhuwail D., Househ M., Hamdi M., Shah Z. (2020) Top concerns of tweeters During the COVID-19 pandemic: infoveillance study. *Journal of Medical Internet Research* 22(4):e19016. <https://doi.org/10.2196/19016>. <http://www.jmir.org/2020/4/e19016/>
2. Armocida B., Formenti B., Ussai S., Palestra F., Missoni E. (2020) The Italian health system and the COVID-19 challenge. *The Lancet Public Health* 5(5):e253. [https://doi.org/10.1016/S2468-2667\(20\)30074-8](https://doi.org/10.1016/S2468-2667(20)30074-8). [https://www.thelancet.com/journals/lanpub/article/PIIS2468-2667\(20\)30074-8/abstract](https://www.thelancet.com/journals/lanpub/article/PIIS2468-2667(20)30074-8/abstract)
3. Barnett M. L., Grabowski D. C. (2020) Nursing homes are ground zero for COVID-19 pandemic. *JAMA Health Forum* 1(3):e200369–e200369. <https://doi.org/10.1001/jamahealthforum.2020.0369>. <https://jamanetwork.com/channels/health-forum/fullarticle/2763666>
4. Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet allocation. *J Mach Learn Res* 3(4-5):993–1022. <https://doi.org/10.1016/b978-0-12-411519-4.00006-9>
5. Chen E, Lerman K, Ferrara E (2020) Tracking social media discourse about the COVID-19 pandemic: development of a public coronavirus twitter data set. *JMIR Public Health and Surveillance* 6(2):e19273. <https://doi.org/10.2196/19273>
6. Chen L, Lyu H, Yang T, Wang Y, Luo J (2020) In the eyes of the beholder: sentiment and topic analyses on social media use of neutral and controversial terms for COVID-19. arXiv:2004.10225[cs]
7. Chen Y, Zhang H, Liu R, Ye Z, Lin J (2019) Experimental explorations on short text topic mining between LDA and NMF based schemes. *Knowl-Based Syst* 163:1–13. <https://doi.org/10.1016/j.knosys.2018.08.011>
8. Dunn AG, Leask J, Zhou X, Mandl KD, Coiera E (2015) Associations between exposure to and expression of negative opinions about human papillomavirus vaccines on social media: an observational study. *Journal of Medical Internet Research* 17(6):e144. <https://doi.org/10.2196/jmir.4343>. <http://www.jmir.org/2015/6/e144/>
9. Ginsberg J., Mohebbi M. H., Patel R. S., Brammer L., Smolinski M. S., Brilliant L. (2009) Detecting influenza epidemics using search engine query data. *Nature* 457(7232):1012–1014. <https://doi.org/10.1038/nature07634>. <https://www.nature.com/articles/nature07634>
10. Ibrahim R, Elbagoury A, Kamel MS, Karray F (2018) Tools and approaches for topic detection from Twitter streams: survey. *Knowl Inf Syst* 54(3):511–539. <https://doi.org/10.1007/s10115-017-1081-x>
11. Kouzy R., Abi Jaoude J., Kraitem A., El Alam M. B., Karam B., Adib E., Zarka J., Traboulsi C., Akl E. W., Baddour K. (2020) Coronavirus goes viral: quantifying the COVID-19 misinformation epidemic on twitter. *Cureus* 12(3) pp 7255 <https://doi.org/10.7759/cureus.7255>. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7152572/>
12. Li C, Wang H, Zhang Z, Sun A, Ma Z (2016) Topic modeling for short texts with auxiliary word embeddings. In: *Proceedings of the 39th international ACM SIGIR conference on research and development in information retrieval*, pp 165–174. <https://doi.org/10.1145/2911451.2911499>
13. Li J., Xu Q., Cuomo R., Purushothaman V., Mackey T. (2020) Data mining and content analysis of the Chinese social media platform Weibo during the early COVID-19 outbreak: retrospective observational infoveillance study. *JMIR Public Health Surveill* 6(2):e18700. <https://doi.org/10.2196/18700>. <http://publichealth.jmir.org/2020/2/e18700/>
14. Mehrotra R, Sanner S, Buntine W, Xie L (2013) Improving LDA topic models for microblogs via tweet pooling and automatic labeling. In: *Proceedings of the 36th international ACM SIGIR conference on research and development in information retrieval*, pp 889–892. <https://doi.org/10.1145/2484028.2484166>

15. Mimmo D, Wallach HM, Talley E, Leenders M, McCallum A (2011) Optimizing semantic coherence in topic models. In: Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, pp 262–272
16. Nguyen DQ, Billingsley R, Du L, Johnson M (2015) Improving topic models with latent feature word representations. *Trans Assoc Comput Linguistics* 3:299–313. <https://doi.org/10.1162/tacl.a.00140>. <https://www.aclweb.org/anthology/Q15-1022>
17. Ordun C, Purushotham S, Raff E (2020) Exploratory analysis of COVID-19 tweets using topic modeling, umap, and digraphs. *arXiv:2005.03082* [cs]
18. Park HW, Park S, Chong M (2020) Conversations and medical news frames on twitter: infodemiological study on COVID-19 in South Korea. *J Med Internet Res* 22(5):e18897. <https://doi.org/10.2196/18897>. <https://www.jmir.org/2020/5/e18897/>
19. Qazi U, Imran M, Ofli F (2020) GeoCoV19: a dataset of hundreds of millions of multilingual COVID-19 tweets with location information. *SIGSPATIAL Spec* 12(1):6–15. <https://doi.org/10.1145/3404820.3404823>
20. Saha A, Sindhvani V (2012) Learning evolving and emerging topics in social media: a dynamic NMF approach with temporal regularization. In: Proceedings of the fifth ACM international conference on web search and data mining, pp 693–702. <https://doi.org/10.1145/2124295.2124376>
21. Saito S, Hirata Y, Sasahara K, Suzuki H (2015) Tracking time evolution of collective attention clusters in Twitter: time evolving nonnegative matrix factorisation. *PLOS ONE* 10(9):e0139085. <https://doi.org/10.1371/journal.pone.0139085>
22. Shahnaz F, Berry MW, Pauca V, Plemmons RJ (2006) Document clustering using nonnegative matrix factorization. *Inf Proc Manag* 42(2):373–386. <https://doi.org/10.1016/j.ipm.2004.11.005>. <https://linkinghub.elsevier.com/retrieve/pii/S0306457304001542>
23. Signorini A, Segre AM, Polgreen PM (2011) The use of Twitter to track levels of disease activity and public concern in the U.S. during the influenza A H1N1 pandemic. *PLOS ONE* 6(5):e19467. <https://doi.org/10.1371/journal.pone.0019467>. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0019467>
24. Steinskog A, Therkelsen J, Gambäck B (2017) Twitter topic modeling by tweet aggregation. <https://www.aclweb.org/anthology/W17-0210>
25. Suri P, Roy NR (2017) Comparison between LDA NMF for event-detection from large text stream data. In: 2017 3Rd International conference on computational intelligence communication technology (CICT), pp 1–5. <https://doi.org/10.1109/CICT.2017.7977281>
26. Szomszor M, Kostkova P, Louis CS (2011) Twitter informatics: tracking and understanding public reaction during the 2009 swine flu pandemic. In: 2011 IEEE/WIC/ACM International conferences on web intelligence and intelligent agent technology, vol 1, pp 320–323. <https://doi.org/10.1109/WI-IAT.2011.311>
27. Verelst F, Kuylen E, Beutels P (2020) Indications for healthcare surge capacity in European countries facing an exponential increase in coronavirus disease (COVID-19) cases, March 2020. *Eurosurveillance* 25:13. <https://doi.org/10.2807/1560-7917.ES.2020.25.13.2000323>. <https://www.eurosurveillance.org/content/10.2807/1560-7917.ES.2020.25.13.2000323>
28. Cui W, Liu S, Tan L, Shi C, Song Y, Gao Z, Qu H, Tong X (2011) TextFlow: towards better understanding of evolving topics in text. *IEEE Trans Vis Comput Graph* 17(12):2412–2421. <https://doi.org/10.1109/TVCG.2011.239>. <http://ieeexplore.ieee.org/document/6065008/>
29. Yin J, Wang J (2014) A Dirichlet multinomial mixture model-based approach for short text clustering. In: Proceedings of the ACM SIGKDD international conference on knowledge discovery and data mining, pp 233–242. <https://doi.org/10.1145/2623330.2623715>
30. Zhao J, Lee M, Ghader S, Younes H, Darzi A, Xiong C, Zhang L (2020) Quarantine fatigue: first-ever decrease in social distancing measures after the covid-19 pandemic outbreak before reopening united states. *arXiv:2006.03716*
31. Zhao R, Tan VY (2017) Online nonnegative matrix factorization with outliers. *IEEE Trans Signal Process* 65(3):555–570. <https://doi.org/10.1109/TSP.2016.2620967>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Chia-Hsuan Chang^{1,2}  · **Michal Monselise**¹ · **Christopher C. Yang**¹

Michal Monselise
mm4837@drexel.edu

Christopher C. Yang
chris.yang@drexel.edu

¹ College of Computing and Informatics, Drexel University, Philadelphia, PA, USA

² Department of Information Management, National Sun Yat-sen University,
Kaohsiung City, Taiwan