Title
A large-scale COVID-19 Twitter chatter dataset for open scientific research - an international collaboration

Authors

Juan M. Banda[1], Ramya Tekumalla[1], Guanyu Wang[2], Jingyuan Yu[3], Tuo Liu[4], Yuning Ding[5], Gerardo Chowell[6]

Affiliations
1. Department of Computer Science, Georgia State University, Atlanta, Georgia
2. Missouri school of journalism, University of Missouri, Columbia, MO
3. Department of social psychology, Universitat Autònoma de Barcelona, Barcelona, Spain
4. Department of psychology, Carl von Ossietzky Universität Oldenburg, Oldenburg, Germany
5. Language technology lab, Universität Duisburg-Essen, Duisburg, Germany
6. Department of Population Health Sciences, Georgia State University, Atlanta, Georgia
corresponding author(s): Juan Banda (jbanda@gsu.edu), Gerardo Chowell (gchowell@gsu.edu)

Abstract

As the COVID-19 pandemic continues its march around the world, an unprecedented amount of open data is being generated for genetics and epidemiological research. The unparalleled rate at which many research groups around the world are releasing data and publications on the ongoing pandemic is allowing other scientists to learn from local experiences and data generated in the front lines of the COVID-19 pandemic. However, there is a need to integrate additional data sources that map and measure the role of social dynamics of such a unique world-wide event into biomedical, biological, and epidemiological analyses. For this purpose, we present a large-scale curated dataset of over 152 million tweets, growing daily, related to COVID-19 chatter generated from January 1st to April 4th at the time of writing. This open dataset will allow researchers to conduct a number of research projects relating to the emotional and mental responses to social distancing measures, the identification of sources of misinformation, and the stratified measurement of sentiment towards the pandemic in near real time.

Background & Summary

The ongoing COVID-19 pandemic began in the form of a cluster of viral pneumonia patients of unknown etiology in the city of Wuhan, China in December 2019. Unfortunately, the interventions to contain its spread were not implemented soon enough to limit the spread of the virus within China's borders. While transmission has been dramatically reduced in China through strict social distancing interventions, the virus was exported to multiple countries and is now generating sustained transmission in multiple areas of the world, including areas with active hotspots of the disease including the United States, Italy, Spain, and France[1]. As of April 4th, 1,201,453 cases have been recorded including 64,688 deaths according to the worldometer coronavirus pandemic tracker [2].

While the ongoing COVID-19 presents with unprecedented challenges to humanity, the wider scientific community can only advance science when they have access to openly available data. Social media platforms like Twitter and Facebook contain an abundance of text data that can be utilized for research purposes. Over the last decade, Twitter has proven to be a valuable resource during disasters for many-to-many crisis communication [3,4,5]. With Twitter data, it is possible to analyze symptom configurations, risk factors, origin, virus

genetics, and spread patterns can be studied and monitored [6–8]. Recent studies [9,10] prove that data sharing improves quality and strengthens research, with collaborative efforts providing an opportunity for researchers to continually enhance research ideas and avoid redundant efforts [11,12]. We opted to release our data to the public for the greater good when the dataset accumulated 40 million tweets on March 23rd [13]. We have been providing updates every two days [14], with a cumulative update every week, most recently on March 29th [15]. This previous update had over 100 million tweets available for researchers. The community response by word of mouth has led to over 2,401 views and over 1,177 downloads of the resource. Moreover, several international researchers have reached out to contribute data and provide analysis expertise. In this release we have incorporated additional data provided by our co-authors expanding the size of the dataset to over 152,920,832 tweets and added vital data on the early days of the pandemic which was unavailable during the initial data collection effort. This shows the value of this kind of data and the engagement of scientists that come together to create extensive resources for the benefit of society. Aside from providing the full dataset with retweets included, we provide a clean version with no retweets for researchers with limited resources to access a lighter version of the dataset. Furthermore, to assist researchers for NLP tasks we provide the top 1000 frequent terms, 1000 bigrams, and 1000 trigrams. The released dataset adheres with FAIR principles[16]. Due to Twitter's terms of service, tweet text cannot be shared. Therefore, tweet ids are publicly made available using Zenodo [13]. The tweet ids can be hydrated using tools like Social Media Mining Toolkit or twarc [17,18]. The deliverables [13,14] include tweet ids and code to process the tweets. Please note that the code to process the tweets would work only after the tweets are hydrated. We provided the date and time meta-data elements on our dataset for groups wanting to target their research questions to certain days to avoid having to hydrate the whole resource at once. We have automated pipelines to continue collecting tweets as the pandemic runs its course and to provide updates every two days on our Github repository. We are also welcoming any additional data that provides new tweets to our resource.

Methods

The initial versions of this dataset [13,15] only included data collected from the publicly available Twitter Stream API with a collection process that gathered any available tweets within the daily restrictions from Twitter from January to March 11th, filtering them on the following 3 keywords: "coronavirus", "2019ncov", "corona virus". We shifted our focus to collect exclusively COVID-19 tweets on March 12th, 2020 with the following keywords: "COVD19", "CoronavirusPandemic", "COVID-19", "2019nCoV", "CoronaOutbreak", "coronavirus", "WuhanVirus", thus the number of tweets gathered dramatically expanded the dataset. Please note that the Stream API only allows free access to a one percent sample of the daily stream of Twitter. Our methodology relies on Python and the Tweepy package [19], as in our previous work [20]. We recently received another set of 30+ Million tweets collected from January 27th, 2020 to March 27th, 2020 from our co-author, Jingyuan Yu, and his collaborators with the following keywords: "coronavirus", "wuhan", "pneumonia", "pneumonie", "neumonia", "lungenentzündung", "covid19". These tweets were collected in the following languages: English, French, Spanish, and German, while our original collection is done for any language available. We have fully integrated and deduplicated our collaborators' tweet collection with ours, thus the numbers and tweets presented in this dataset are of unique tweet identifiers from January 1st to April 4th (at the time of writing). Table 1 represents the monthly number of tweets included in this dataset.

Table 1: Number of Covid-19 chatter tweets in this dataset.

Month	Full	Clean
January	6,737,875	1,329,579
February	21,070,885	4,745,851
March	110,220,360	21,049,512
April (Last collected 4/5)	14,891,712	3,865,703
Total	**152,920,832**	**30,990,645**

As previously mentioned, the number of collected tweets increased tremendously since starting dedicated collection. All our preprocessing scripts utilize components of the Social Media Mining Toolkit (SMMT) [17]. We make a distinction between our **full** and **clean** versions of the dataset. The **full** dataset consists of both tweets and retweets. There are several practical reasons to leave the retweets; tracing important tweets and their dissemination is one of them. A **clean** version with no retweets is also released, intended for NLP researchers. We also release extracted frequent terms, bigrams, and trigrams for this community. Figure 1 outlines the steps taken to build our dataset.

Figure 1: Dataset gathering and construction steps.
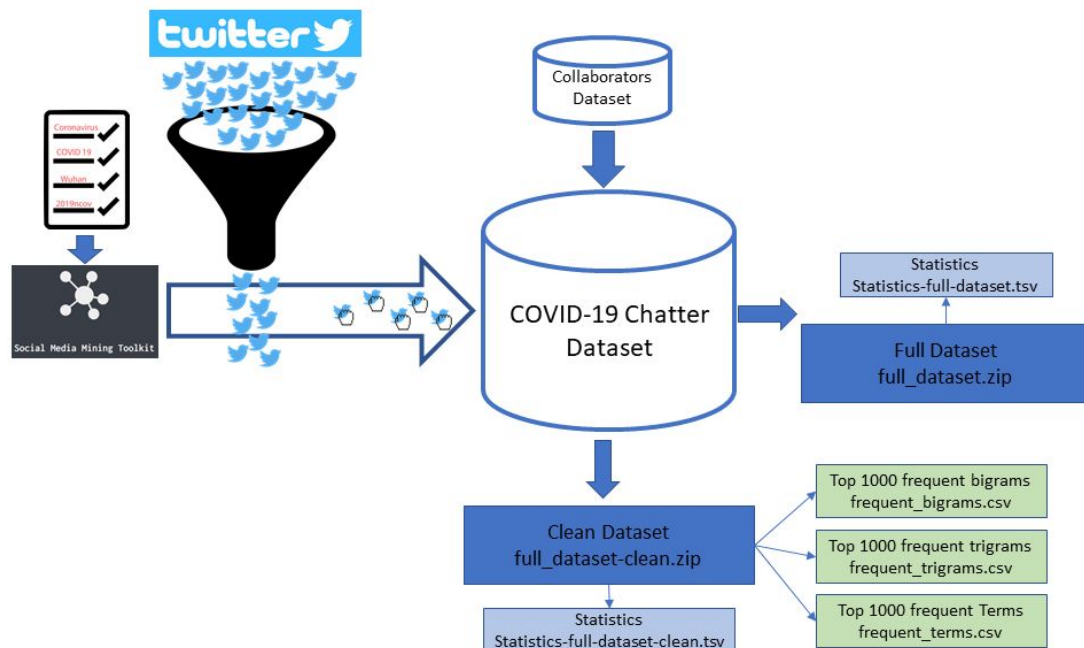
As shown in Figure 1, we used SMMT to listen to the Twitter Stream API for tweets with the described keywords. We then gather all the tweets that have the desired keywords before aggregating them locally. Our contributors used a similar procedure to gather their tweets and provided us with tab delimited files with their data. We processed them to fit our own local format to be able to include them in our dataset after deduplication (removal of tweets we have in common) and only keep unique tweet identifiers between the datasets. We then

3

preprocess the large set of tweets to extract the shareable meta-data of the full dataset (tweet_id, collected date, collected time), preparing the full_dataset.tsv.gz file. At the same time we also remove tweets that are re-tweeted (this is, existing tweets that are re-shared by others) to create the full_dataset-clean.tsv.gz file. Our preprocessing involves cleaning up special characters, such as carriage returns, removing urls and large blank spaces. Our preprocessing is rather relaxed as we are leaving all available languages intact. To generate the frequent terms and ngrams (sets of n-terms that appear constantly together), we remove all stop words in English and Spanish, using the Spacy[21]. These lists are originally quite large, so we only share the top 1000 terms, bigrams, and trigrams. We continue to update our original dataset every two days[14] with major releases every week [13,15] and plan to continue doing this for at least the next 12 months, a period that will likely cover the main pandemic period.

Data Records

The dataset is made available through Zenodo [13] . There are 7 files in this repository. Table 2 details the files, formats, and their utility. The example column consists of a sample line from the files. The tweet ids in the dataset can be hydrated using SMMT. The hydrated tweets would produce a JSON object for each tweet id. It is important to note that when users remove their accounts or individual tweets, these get removed and are no longer available for download. In such cases, we can share the data on request while adhering to the Twitter data sharing policy. The frequent terms, bigrams, and trigrams are retrieved from the cleaned version of the dataset. The full_dataset.tsv consists of all the procured tweet ids. The full_dataset-clean.tsv contains only original tweets with no retweets. While some applications and questions are better served with the full dataset, NLP researchers might prefer a clean dataset to have less inflated counts of the n-grams identified.

Table 2: Details of released COVID-19 dataset. Note that the word TAB is not found, but instead use the special '\t' character for this. We are showing it on the descriptions for illustrative purposes.

File Name	Description	Example
full_dataset.tsv.gz	A zipped, tab separated file which contains all the tweet ids in the format - Tweet ID TAB Date TAB Time	1238315928095297538 TAB 2020-3-13 TAB 4:8:18
full_dataset-clean.tsv.gz	A zipped, tab separated file which does not contain any retweet ids in the format - Tweet ID TAB Date Tab Time	1238315936379293696 TAB 2020-3-13 TAB 4:8:20
statistics-full_dataset-clean.tsv	A tab separated file which contains counts of total tweets each day for the clean dataset in the format - Date TAB Total No of Tweet Ids	2020-3-13 TAB 751804 On March 13,2020 a total of 751,804 clean tweets were collected
statistics-full_dataset.tsv	A tab separated file which contains counts of total tweets each day for full dataset in the format - Date TAB Total No of	3/13/2020 TAB 4160194 On March 13,2020 a total of 4,160,194 tweets were collected

	Tweet Ids	
frequent_terms.csv	A comma separated file which contains the counts of top 1000 frequent terms in the following format - term, Total count	covid19, 1767060 covid 19 term appeared in 1,767,060 tweets
frequent_bigrams.csv	A comma separated file which contains counts of top 1000 bigrams in the format - gram, Total count	covid 19, 1467434 covid 19 bigram appeared in 1,467,434 tweets.
frequent_trigrams.csv	A comma separated file which contains counts of top 1000 trigrams in the format - gram, Total count	coronavirus covid 19, 52143 coronavirus covid 19 appeared in 52,143 tweets

Technical Validation

In order to verify the presence of COVID-19 chatter in the dataset, we performed a text match using Spacy matcher to identify the matched keyword in a randomly sampled set of retrieved tweets. For all the tweets, we performed the following checks:

1) A tweet is checked for its originality. All the original tweets (not retweets) are then separated as a clean dataset.
2) All the tweet texts are preprocessed. Preprocessing includes removal of urls and stopwords. The words of the text are added to a bag of words and then the top 1000 frequent words are identified.
3) In order to validate the total number of tweets, we calculated the unique number of tweets generated each day. All the details are placed in statistics_full_dataset.tsv and statistics_full_dataset_clean.tsv

Usage Notes

In order to use our resource, we have provided all the software tools we utilized to preprocess, clean and parse the Twitter data on our Github repository [14], under the processing code directory. Note that the tweets need to be hydrated first using tools like Social Media Mining Toolkit or twarc [17,18]. Once the tweets are hydrated and a JSON object has been returned, we use the files **parse_json_extreme.py** and **parse_json_extreme_clean.py** to extract the tweet identifier, date of creation, text, language and a few other extra fields. This process can be configured by adding which fields from the tweet json object the user wants to extract in the fields.py file. These utilities produce a full and a clean version of the dataset respectively, on a tab delimited file. This process is optimized to read large files without loading them fully in memory. If the user has a system with very large amounts of RAM memory, we also provide **parse_json_lite.py** to perform the same task. Once the JSON object has been parsed, most users will be able to operate on the tweets directly this way. We additionally provide the **get_1grams.py** and **get_ngrams.py** utilities to generate the most frequent terms and bigrams and trigrams,

respectively. As the hydrated tweet JSON objects are typically quite large, we recommend separating them in daily batches to be able to more efficiently process them. All our previously mentioned tools take a single file as an input parameter for processing and output a new file. In order to combine the results of the ngram generation from multiple files, we prove the following tools that take a folder path as input and iterate through all files present: **combine1grams.py, combineNgrams.py**. In order to share the tweet identifiers with other groups, we provide the getDataset.py, getDataset_clean.py files which generate the equivalent files of **full_dataset.tsv** and **full_dataset-clean.tsv** that are presented in this resource in a compressed (zip) maner. Lastly, dataset statistics can be calculated with **getStats.py**, by passing the full or clean dataset filename to them.

Code Availability

All the code used to preprocess and curate the dataset is available on a public github repository (https://github.com/thepanacealab/covid19_twitter). The code is freely available and the dataset is licensed under Creative Commons Non Commercial Attribution-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, and should only be used for research purposes. The Social Media Mining Toolkit code, which was utilized to download the tweets is also available on a public github repository (https://github.com/thepanacealab/SMMT). In order to use the Social Media Mining toolkit, a twitter developer account is required.

Acknowledgements

Author contributions

JMB and RT wrote the manuscript, performed the data gathering, curation and reconciliation of external data sources. JMB, RT, and GC-P conceptualized the data gathering strategy and reviewed the manuscript. GW, JY, TL,YD acquired, provided additional data, and reviewed the manuscript.

Competing interests

JMB, RT, GC-P, GW, JY, TL, YD have no competing interests.

References

1. World Health Organization. WHO characterizes COVID-19 as a pandemic.

    https://www.who.int/emergencies/diseases/novel-coronavirus-2019/events-as-they-ha

    ppen.

2. Coronavirus Update (Live): 737,575 Cases and 34,998 Deaths from COVID-19 Virus

    Outbreak - Worldometer. https://www.worldometers.info/coronavirus/.

3. Bruns, A. & Liang, Y. E. Tools and methods for capturing Twitter data during natural

disasters. *First Monday* **17**, 1–8 (2012).

4. Zou, L., Lam, N. S. N., Cai, H. & Qiang, Y. Mining Twitter Data for Improved Understanding of Disaster Resilience. *Ann. Assoc. Am. Geogr.* **108**, 1422–1441 (2018).

5. Earle, P. Earthquake Twitter. *Nat. Geosci.* **3**, 221–222 (2010).

6. Gao, J., Tian, Z. & Yang, X. Breakthrough: Chloroquine phosphate has shown apparent efficacy in treatment of COVID-19 associated pneumonia in clinical studies. *Biosci. Trends* **14**, 72–73 (2020).

7. Xu, Z. *et al.* Pathological findings of COVID-19 associated with acute respiratory distress syndrome. *Lancet Respir Med* (2020) doi:10.1016/S2213-2600(20)30076-X.

8. Zhou, F. *et al.* Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study. *Lancet* **395**, 1054–1062 (2020).

9. Warren, E. Strengthening Research through Data Sharing. *N. Engl. J. Med.* **375**, 401–403 (2016).

10. Saez-Rodriguez, J. *et al.* Crowdsourcing biomedical research: leveraging communities as innovation engines. *Nat. Rev. Genet.* **17**, 470–486 (2016).

11. Emmert-Streib, F., Dehmer, M. & Yli-Harja, O. Against Dataism and for Data Sharing of Big Biomedical and Clinical Data with Research Parasites. *Front. Genet.* **7**, 154 (2016).

12. Greene, C. S., Garmire, L. X., Gilbert, J. A., Ritchie, M. D. & Hunter, L. E. Celebrating parasites. *Nature genetics* vol. 49 483–484 (2017).

13. Banda, J. M. & Tekumalla, R. A Twitter Dataset of 40+ million tweets related to COVID-19. (2020) doi:10.5281/zenodo.3723940.

14. Banda Juan M, Tekumalla Ramya. Covid-19 Twitter dataset and pre-processing scripts. https://github.com/thepanacealab/covid19_twitter.

15. Banda, J. M., Tekumalla, R. & Chowell, G. *A Twitter Dataset of 70+ million tweets related to COVID-19*. (2020). doi:10.5281/zenodo.3732460.

16. Wilkinson, M. D. *et al.* The FAIR Guiding Principles for scientific data management and

stewardship. *Sci Data* **3**, 160018 (2016).

17. Tekumalla, R. & Banda, J. Social Media Mining Toolkit (SMMT). Under review. *Genomics and Informatics* (2020).

18. *twarc*. (Github).

19. *tweepy*. (Github).

20. Tekumalla, R., Asl, J. R. & Banda, J. M. Mining Archive.org's Twitter Stream Grab for Pharmacovigilance Research Gold. 859611 (2019) doi:10.1101/859611.

21. Explosion, A. I. spaCy-Industrial-strength Natural Language Processing in Python. *URL: https://spacy. io* (2017).