

Mega-COV: A Billion-Scale Dataset of 100+ Languages for COVID-19

Muhammad Abdul-Mageed, AbdelRahim Elmadany, El Moatez Billah Nagoudi,
Dinesh Pabbi, Kunal Verma, Rannie Lin

Natural Language Processing Lab
University of British Columbia

{muhammad.mageed, a.elmadany, moatez.nagoudi}@ubc.ca,
{dinesh09, vkunal196}@ece.ubc.ca, krieyalam@gmail.com

Abstract

We describe Mega-COV, a billion-scale dataset from Twitter for studying COVID-19. The dataset is diverse (covers 268 countries), longitudinal (goes as back as 2007), multilingual (comes in 100+ languages), and has a significant number of location-tagged tweets (~ 169 M tweets). We release tweet IDs from the dataset. We also develop two powerful models, one for identifying whether or not a tweet is related to the pandemic (best $F_1=97\%$) and another for detecting misinformation about COVID-19 (best $F_1=92\%$). A human annotation study reveals the utility of our models on a subset of Mega-COV. Our data and models can be useful for studying a wide host of phenomena related to the pandemic. Mega-COV and our models are publicly available.

1 Introduction

The seeds of the coronavirus disease 2019 (COVID-19) pandemic are reported to have started as a local outbreak in Wuhan (Hubei, China) in December, 2019, but soon spread around the world (WHO, 2020). As of January 24, 2021, the number of confirmed cases around the world exceeded 99.14M and the number of confirmed deaths exceeded 2.13M.¹ In response to this ongoing public health emergency, researchers are mobilizing to track the pandemic and study its impact on all types of life in the planet. Clearly, the different ways the pandemic has its footprint on human life is a question that will be studied for years to come. Enabling scholarship on the topic by providing relevant data is an important endeavor. Toward this goal, we collect and release Mega-Cov, a billion-scale multilingual Twitter dataset with geo-location information.

¹Source: The Center for Systems Science and Engineering, Johns Hopkins University. Dashboard: <https://www.arcgis.com/apps/opsdashboard/index.html#/bda7594740Fd40299423467b48e9ecf6>.

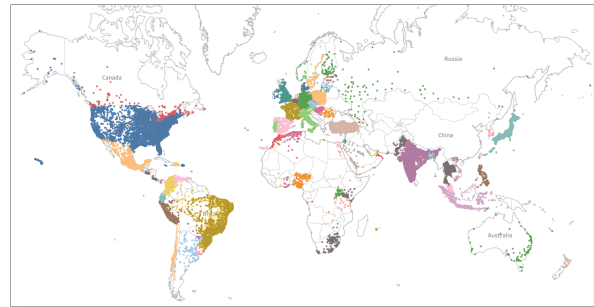


Figure 1: Global coverage of Mega-COV based on our geo-located data. Each dot is a city. Contiguous cities of the same color belong to the same country.

As a result of the pandemic, most countries around the world went into lockdown and the public health emergency has restricted physical aspects of human communication considerably. As hundreds of millions of people spend more time sheltering in place, communication over social media became more important than ever. In particular, the content of social media communication promises to capture significant details about the lives of tens of millions of people. Mega-Cov is intended as a repository of such a content.

There are several ongoing efforts to collect Twitter data, and our goal is to complement these. More specifically, we designed our methods to harvest a dataset that is unique in multiple ways, as follows: **Massive Scale:** Very large datasets lend themselves to analyses that are not possible with smaller data. Given the global nature of COVID-19, we realize that a large-scale dataset will be most useful as the scale allows for slicing and dicing the data across different times, communities, languages, and regions that are not possible otherwise. For this reason, we dedicated significant resources to harvesting and preparing the dataset. Mega-COV has solid international coverage and brings data from 1M users from 268 countries (see Section 3.1). Overall, our dataset has ~ 1.5 B tweets (Section 2). This is one order of magnitude larger than #COVID-19 (Chen et al., 2020), the largest dataset we know

of ($\sim 144\text{M}$ tweets as of June 1, 2020).²

Topic Diversity: We do not restrict our collection to tweets carrying certain hashtags. This makes the data general enough to involve content and topics directly related to COVID-19, regardless of existence of accompanying hashtags. This also allows for investigating themes that may not be directly linked to the pandemic but where the pandemic may have some bearings which should be taken into account when investigating such themes. This is important because users can, and indeed do, post about activities impacted by the health crisis without using any hashtags. In fact, users may not mention COVID-19 at all, even though what they are posting about could be affected by the pandemic one way or another (e.g., “eating habits”, “shopping behavior”). Section B and Section C in the Appendix provide a general overview of issues discussed in the dataset.

Longitudinal Coverage: We collect multiple data points (up to 3,200) from each user, with a goal to allow for comparisons between the present *and* the past across the same users, communities, and geographical regions (Section 3.2). Again, this is desirable since without data from pre-COVID-19 time it will be challenging to hold any such comparisons. For example, some users may have stopped posting about “exercising” during the pandemic but we cannot definitely identify this without access to these users’ previous data where they may have been posting about their physical activities.

Language Diversity: Since our collection method targets users, rather than hashtag-based content, Mega-COV is *linguistically diverse*. In theory, any language posted to Twitter by a user whose data we have collected should be represented. Based on Twitter-assigned language codes, we identify a total of 65 languages. However, applying two different language detection tools to the whole dataset, we identify more than 100 languages. (Section 3.3).

No Distribution Shift: Related to the two previous points, but from a machine learning perspective, by collecting the data without conditioning on existence of specific (or any) hashtags we avoid introducing distribution bias. In other words, the data can be used to study various phenomena in-the-

²Both our own dataset and that of Chen et al. (2020) are growing over time. All our statistics in the current paper are based on our collection as of May 15, 2020. As of October 6, 2020, authors of #COVID-19 report 649.9M tweets on their GitHub (<https://github.com/euchen102/COVID-19-TweetIDs>), and our own dataset has exceeded 5B tweets.

wild. This warrants more *generalizable* findings and models.

A dataset as large as Mega-COV can be hard to navigate. In particular, an informative description of the dataset is necessary for navigating it. In this paper, we provide an explanation of a number of global aspects of the dataset, including its geographic, temporal, and linguistic coverage. We also provide a high-level content analysis of the data, and explore user sharing of content from particular web domains with a focus on news media. In the context of our investigation of Mega-COV, we make an array of important discoveries. For example, we strikingly discover that, perhaps for the first time in Twitter history, *users address one another and retweet more than they post tweets*. We also find a noticeable rise in ranks for news sites (based on how frequent their URLs are shared) during 2020 as compared to 2019, with a *shift toward global (rather than local) news media*. A third finding is how *use of the Twitter platform surged in March*, perhaps making it the busiest time in the history of the network.

Furthermore, we develop two groups of effective neural models: **(1) COVID-relevance models** (for detecting whether a tweet is related to COVID-19 or not). **(2) COVID-misinformation models** (for detecting whether a text carries fake information or not). In addition to releasing our best models, we also apply them to a total of 30M tweets from Mega-COV and release our tags to accelerate further research on the topic.

The rest of the paper is organized as follows: In Section 2, we describe our data collection methods. Section 3 is where we investigate geographic, linguistic, and temporal dimensions of our data. We describe our models for detecting COVID-19 tweets and COVID-misinformation in Section 4. Section 5 is where we apply our relevance and misinformation models to a large sample of Mega-COV. Section 6 is about data release and ethics. We provide a literature review in Section 7, and conclude in Section 8.

2 Data Collection

To collect a sufficiently large dataset, we put crawlers using the Twitter streaming API³ on Africa, Asia, Australia, Europe, North America, and South America starting in early January, 2020. This allows us to acquire a diverse set of tweets

³API link: <https://github.com/tweepy/>

Data	Tweets	Retweets	Replies	All
2007-2020	612M	507M	369M	1.5B
2020	122M	174M	129M	425M
Users	1M	976K	994K	1M

Table 1: Distribution of tweets, retweets, and replies in Mega-COV (numbers rounded).

from which we can extract a random set of user IDs whose timelines (up to 3,200 tweets) we then iteratively crawl every two weeks. This gave us data from July 30th, 2020 backwards, depending on how prolific of a poster a user is (see Table 4a for a breakdown.). In this paper, we describe and analyze the version of Mega-COV collected up to May 15, 2020 and use the term Mega-COV to refer it. Mega-COV comprises a total of 1,023,972 users who contribute 1,487,328,805 tweets. For each tweet, we collect the whole json object. This gives us access to various types of information, such as user location and the language tag (including “undefined”) Twitter assigns to each tweet. We then use the data streaming and processing engine, Spark, to merge all user files and run our analyses. To capture a wide range of behaviors, we keep *tweets*, *retweets*, and *responses* (i.e., direct user-to-user interactions) as independent categories. Table 1 offers a breakdown of the distribution of the different types of posts in Mega-COV. Tweet IDs of the dataset are publicly available at our GitHub⁴ and can be downloaded for research. To the extent it is possible, we intend to provide semi-regular updates to the dataset repository.

3 Exploring Mega-COV

3.1 Geographic Diversity

A region from which a tweet is posted can be associated with a specific ‘point’ location or a Twitter place with a ‘bounding box’ that describes a larger area such as city, town, or country. We refer to tweets in this category as *geo-located*. A smaller fraction of tweets are also *geo-tagged* with longitude and latitude. As Table 2 shows, Mega-COV has $\sim 187\text{M}$ geo-located tweets from $\sim 740\text{K}$ users and $\sim 31\text{M}$ geo-tagged tweets from $\sim 267\text{K}$ users. Table 2 also shows the distribution of tweets and users over the top two countries represented in the dataset, the U.S. and Canada (North America), and other locations (summed up as one category, but see also Table 3 for countries in the data by con-

⁴Accessible at: <https://github.com/UBC-NLP/megacov>.




Figure 2: World map coverage of Mega-COV. Each dot is a point co-ordinate (longitude and latitude) from which at least one tweet was posted. Clearly, users tweet while traveling, whether by air or sea.

tinents). As explained, to allow comparisons over time (including behavioral changes during COVID-19), we include pre-2020 data in Mega-COV. For the year 2020, Mega-COV has $\sim 66\text{M}$ geo-located tweets from $\sim 670\text{K}$ users and $\sim 3\text{M}$ geo-tagged tweets from $\sim 109\text{K}$ users.⁵ We note that significant parts from the data could still belong to the different countries but just not geo-located in the original json files retrieved from Twitter. Figure 2 shows actual point co-ordinates of locations from which the data were posted. Figure 3 shows the geographical diversity in Mega-COV based on *geo-located* data. We show the distribution in terms of the number of cities over the 20 countries from which we retrieved the highest number of locations in the dataset, broken by all-time and the year 2020. Overall, Mega-COV has data posted from a total of 167,202 cities that represent 268 countries. Figure A.2 in Appendix A shows the distribution of data over countries. The top 5 countries in the data are *the U.S., Canada, Brazil, the U.K., and Japan*. As we mention earlier, other top countries in the data across the various continents are shown in Table 3.

3.2 Temporal Coverage

Our goal is to make it possible to exploit Mega-COV for comparing user social content over time. Since we crawl user timelines, the dataset comprises content going back as early as 2007. Figure 4a shows the distribution of data over the period 2007-2020. Simple frequency of user posting shows a surge in Twitter use in the period of Jan-April 2020 compared to the same months in 2019 (see Figure A.1 in Appendix A). Indeed, we identify 40.53% more posting during the first 4 months

⁵The dataset has $\sim 134\text{K}$ “locations” which we could not resolve to a particular country using only the json information.

	Geolocated	Tweeted From			Geotagged	Tweeted From		
		Canada	U.S.	Other		Canada	U.S.	Other
All-Time	186,939,854	16,459,655	70,756,282	99,723,917	31,392,563	3,600,952	11,449,400	16,342,211
All-Users	739,645	102,388	327,213	463,673	266,916	47,622	117,096	165,860
2020	65,584,908	3,331,720	24,259,973	37,993,215	2,942,675	246,185	1,187,131	1,509,359
2020-Users	670,314	61,205	254,067	392,627	109,348	14,525	43,486	62,837

Table 2: Mega-COV geolocated and geotagged users and their tweets from North America vs. Other locations. See also Table 3 for statistics from top countries by continent.

Continent	Country	All		2020	
		Geo-Located	Users	Geo-Located	Users
Africa	Nigeria	1,876,879	16,220	1,057,742	14,872
	South Africa	1,503,181	9,751	692,367	6,373
	Egypt	873,079	8,840	452,738	5,900
	Ghana	373,996	3,942	202,470	3,089
	Kenya	373,667	4,480	172,796	3,026
Asia	Japan	7,646,901	32,038	2,752,890	23,773
	Indonesia	4,540,286	22,893	1,871,154	18,056
	Spain	4,327,475	43,236	1,431,567	20,902
	Philippines	4,078,410	15,477	1,636,265	11,011
	India	3,107,917	33,931	1,576,549	27,940
	Saudi Arabia	2,158,584	18,402	833,634	15,087
Australia	Australia	1,179,205	12,090	352,215	5,454
Europe	UK	11,714,012	70,787	2,970,848	44,420
	Turkey	5,067,118	32,589	1,463,550	25,477
	France	2,030,523	36,017	729,500	12,497
	Italy	1,829,369	27,071	527,648	8,308
	Germany	1,272,339	24,215	385,306	7,412
North America	US	69,515,949	327,213	23,578,430	254,067
	Canada	16,066,337	102,388	3,200,804	61,205
	Mexico	3,665,791	36,190	1,106,352	17,406
South America	Brazil	15,879,664	48,339	8,060,537	41,277
	Argentina	3,142,778	14,576	1,298,381	10,901
	Colombia	1,612,765	10,319	629,426	6,884
	Chile	1,003,459	6,212	378,770	3,674
	Ecuador	447,250	3,435	170,098	2,221

Table 3: Distribution of data over top countries per continent in Mega-COV (all data vs. 2020).

of 2020 compared to the same period in 2019. This is expected, both due to physical distancing and a wide range of human activity (e.g., “work”, “shopping”) moving online. More precisely, moving activities online causes users to be on their machines for longer times and hence have easier access to social media. The clear spike in the month of March 2020 is striking. It is particularly so given a *shifted pattern of use: retweeting and replying (to others) are both observably more frequent than tweeting itself*. This especially takes place during the month




Figure 3: Geographical diversity in Mega-COV based on geo-located data.

of March, and somewhat continues in April, as shown in Figure 4b. Figure 4a and Figure 4b also show a breakdown of tweets, retweets, and replies. **A striking discovery is that, for 2020, users are engaged in conversations with one another more than tweeting directly to the platform.** This may be the first time this pattern exists, perhaps in the history of the network. At least based on our massive dataset, this conclusion can be made. In addition, for 2020, we also see users retweeting more than tweeting. Based on Mega-COV, this is also happening for the first time.

3.3 Linguistic Diversity

We perform the language analysis based on tweets ($n \sim 1.5B$), including retweets and replies. Twitter assigns 65 language ids to $\sim 1.4B$ tweets, while the rest are tagged as “und” (for “undefined”). Mega-COV has $\sim 104M$ ($\sim 7\%$) tweets tagged as “und”. We run two language identification tools, langid (Lui and Baldwin, 2012) and Compact Language Detector (Ooms and Sites, 2018)⁶ langid (Lui and Baldwin, 2012),⁷ on the whole dataset (including tweets tagged “und” by Twitter).⁸ After merging language tags from Twitter and the 2 tools, we acquire a total of 104 labels. This makes Mega-COV very linguistically rich. Table 4 shows the top 20 languages identified by Twitter (left) and the top 20 languages tagged by one of the two tools, langid (Lui and Baldwin, 2012), after removing the 65 Twitter languages (right).


4 Models

We develop two groups of models suited for answering important questions related to COVID-19, including making use of Mega-COV. These are (1) **COVID-relevance**, where a classifier will

⁶<https://code.google.com/p/cld2>.

⁷<https://github.com/saffsd/langid.py>

⁸As (Lui and Baldwin, 2014) point out, langid makes error on Twitter data. For this reason, we opted for adding predictions from CLD2.



(a) Twitter user activity for Jan-May, 2020.

(b) Distribution of Mega-COV (2007-2020)

Figure 4: Data distribution and user activity.

Lang	Freq	Lang	Freq
English (en)	900M	Hebrew (he)	1.2M
Spanish (es)	122M	Croatian (hr)	685K
Portuguese (pt)	79.5M	Maltese (mt)	325K
Japanese (ja)	46.6M	Slovak (sk)	246K
Arabic (ar)	45M	BI (id)	208K
Indonesian (in)	37M	Latin (la)	183K
French (fr)	29.5M	Bosnian (bs)	143K
Turkish (tr)	28.5M	Dzongkha (dz)	137.8K
Tagalog (tl)	19M	Swahili (sw)	92K
Italian (it)	8.8M	Azerbaijani (az)	68.9K
Thai (th)	7.7M	Quechua (qu)	61K
Hindi (hi)	7M	Albanian (sq)	61K
Dutch (nl)	6.9M	Malay (ms)	59K
Russian (ru)	6.2M	Kinyarwanda (rw)	56.8K
German (de)	6M	Esperanto (eo)	55K
Catalan (ca)	3.5M	Javanese (jv)	53K
Korean (ko)	2.9M	Xhosa (xh)	47.7K
Haitian Creole (ht)	2.8M	Irish (ga)	44.6K
Polish (pl)	2.4M	Kurdish (ku)	43K
Estonian (et)	2.1M	Volapük (vo)	41K

Table 4: Top 20 languages assigned by Twitter (left) and top 20 languages assigned by langid (right) in Mega-COV. **BI**: Bahasa Indonesia.

label a tweet as *relevant* to COVID-19 or *not* and (2) **COVID-misinformation**, where a model predicts text veracity pertaining COVID-19 (i.e., whether a text carries *true* or *false* information related to the pandemic). We now describe our methods.

4.1 Methods

For *all* our models, we fine-tune 3 popular pre-trained language models: (1) Multilingual cased BERT (mBERT) (Devlin et al., 2018) and (2-3) XLM Roberta base and large (XLM-R_{Base}, XLM-R_{Large}) (Conneau et al., 2020). The mBERT and XLM-R_{Base} models have similar architectures, with 12 layers each with 12 attention heads, and 768 hidden units. XLM-R_{Large} has 24 layers each with 16 attention heads, and 1,024 hidden units. While all

the 3 models use a masking objective, the XLM-R models do not include the next sentence prediction objective used in BERT.

4.2 Hyper-Parameters and Optimization

For each model, we use the same pre-processing in the respective code released by the authors. For *all* models, we typically use a sequence length of 50 tokens. We use a learning rate of $5e-6$ and a batch size of 32. We train each model for 20 epochs and identify the best epoch on a development set. We report performance on both development and test sets. We describe our baseline for each of the relevance and misinformation models in the respective sections below. We now introduce each of these two model groups.

4.3 COVID-Relevance Models

Data. Our COVID-relevance models predict whether a tweet is *related* to COVID-19 or *not* (i.e., not related). To train the models, we sample $\sim 2.3M$ multilingual tweets (65 languages) collected with COVID-19 hashtags from (Chen et al., 2020) and use them as our positive class (i.e., related to COVID-19). Examples of hashtags include *#Coronavirus*, *#covid-19*, and *#pandemic*. That is, we use the hashtags as a proxy for labels. This type of *distant supervision* has been validated and widely used in many NLP models (Go et al., 2009; Mohammad and Kiritchenko, 2015; Abdul-Mageed and Ungar, 2017). For the negative class (i.e., not related to COVID-19), we use a random sample of $\sim 2.3M$ from the 2019 part (Jan-Nov) of Mega-COV. More description of the dataset we created for training the relevance models and the distribution of the data over the various languages is in Table E.1 (Appendix E).

Splits and Training. We split the data into 80%

Model	DEV		TEST	
	Acc	F ₁	Acc	F ₁
Baseline I	55.10	71.05	54.99	70.96
Baseline II	79.88	88.81	75.33	85.93
mBERT	97.35	97.33	97.39	97.37
XLM-R _{Base}	97.72	97.70	97.71	97.69
XLM-R _{Large}	97.92	97.90	97.95	97.93

Table 5: Performance of COVID-relevance models. **Baseline I:** Majority class in TRAIN. **Baseline II:** A model that chooses the majority class (related class) 75% of the time.

TRAIN (n=3,146,334), 10% DEV (n=393,567), and 10% TEST (n=392,918). We then remove all hashtags which were used by (Chen et al., 2020) for collecting the data and fine-tune each of the 3 language models on TRAIN.

Results. As shown in Table 5, XLM-R_{Large} acquires best results with 97.95 *acc* and 97.93 macro F_1 on TEST. These results are significantly better than a majority class baseline (based on TRAIN) and another arbitrarily-chosen (yet quite competitive) baseline model that chooses the related class (majority class in TRAIN) 75% of the time.

Model Generalization. Our COVID-relevance models are trained with distant supervision (hashtags as surrogate labels). It is conceivable that content related to COVID-19 would still occur in real world without accompanying hashtags. To test the extent to which our best model would perform on external data, we evaluate it on two external Twitter datasets, CoAID (Cui and Lee, 2020) and ReCOVery (Zhou et al., 2020), both of which are claimed by the authors to be completely (100%) related to COVID-19.⁹

As Tabel 6 shows, We do observe a drop in model performance as compared to our best model on our own TEST set in Table 5 (acc drops on average by 15.5% and 7.6% F_1). However, the best model is still highly effective. It acquires an average acc of 82.46% and F_1 of 90.38% on the CoAID and ReCOVery datasets. We now introduce our misinformation models.

⁹Each of the two datasets are also labeled for fake news (*true vs. fake*) focused on COVID-19, but our focus here is exclusively on using the two datasets as gold-labeled TEST sets for evaluating our COVID-relevance model. Note that we will use these two datasets again as explained in Section 4.4 as well.

Data	Acc	F ₁
COAID	76.25	86.52
ReCOVery	89.46	94.44
Average	82.46	90.38

Table 6: Performance of our COVID-relevance models on the Twitter data in CoAID, ReCOVery, and CoAID+ReCOVery.

4.4 COVID-Misinformation Models

To train models for detecting the veracity of news related to COVID-19, we exploit two recent and publicly available fake news datasets (in English): CoAID (Cui and Lee, 2020), and ReCOVery (Zhou et al., 2020). We now describe each of these datasets:

	Fake			True		
	Claims	News	Tweets	Claims	News	Tweets
CoAID	839	837	10,900	376	2716	149,343
ReCOVery	-	665	26,418	-	1,364	114,402
Total	839	1,502	37,318	376	4,080	263,745

Table 7: COVID-19 Misinformation Datasets.

	Tweets					
	Fake			True		
	TRAIN	DEV	TEST	TRAIN	DEV	TEST
CoAID	8,072	1,009	1,009	110,076	13,759	13,759
ReCOVery	18,272	2,284	2,284	86,437	10,805	10,805
CoAID*	8,072	163	171	110,076	6,314	6,388
ReCOVery*	18,272	154	139	86,437	1,218	1,263

Table 8: Statistics of CoAID and ReCOVery datasets across the data splits. CoAID* and ReCOVery* are de-duplicated versions.

CoAID. Cui and Lee (2020) present a Covid-19 heAthcare mIsinformation Dataset (CoAID), with diverse COVID-19 healthcare misinformation, including fake news on websites and social platforms, along with related user engagements (i.e., tweets and replies) about such news. CoAID includes 3, 235 news articles and claims, 294, 692 user engagement, and 851 social platform posts about COVID-19. The dataset is collected from December 1, 2019 to July 1, 2020. Table 7 shows class distribution of news articles and tweets in CoAID. More information about CoAID is in Appendix F. **ReCOVery.** Zhou et al. (2020) choose 60 news publishers with ‘extreme’ levels of credibility (i.e., *true vs. fake* classes) from an original list of $\sim 2, 000$ to collect a total of 2, 029 news articles on COVID-19, published between January and May

2020. They also collect 140, 820 tweets related to the news articles, considering those tweets related to true articles to be true and vice versa. Table 7 shows class distribution of news articles and tweets in ReCOVeRY.

Splits and Cleaning. Table 8 shows the distribution of tweets in CoAID and Recovery *before* and *after* the de-duplication process. As Table 8 shows, de-duplication results in significantly reducing the sizes of DEV and TEST sets in the two resources. The distribution of news article is shown in Table F.1 (Appendix F).

Training. We use both CoAID and ReCOVeRY after de-duplication for training neural models to detect fake news related to Covid-19. Using the same hyper-parameters and training setup as the COVID-relevance models, we fine-tune the pre-trained language models on the Twitter dataset and the news dataset, independently.¹⁰ Since Mega-COV is a social media dataset, we only focus on training Twitter models here and provide the news models in Appendix F. For the *Twitter models*, we develop one model on CoAID, another on ReCOVeRY, independently, and a third model for CoAID+ReCOVeRY (concatenated). Again, for each of these 3 datasets, we fine-tune on TRAIN and identify the best model on DEV. We then report the best model on both DEV and TEST.

Results. Since our focus is on detecting *fake* texts, we show results on the positive class only in Table 9. We report results in terms of *precision*, *recall* and F_1 . Our baseline is a small LSTM with 2 hidden layers, each of which has 50 nodes. We add a dropout of 0.2 after the first layer and arbitrarily train the LSMT for 3 epochs. As Table 9 shows, our best results for fake tweet detection on TEST for CoAID is at 90% F_1 (mBERT/XLM-R_{Large}), for ReCOV 68% (mBERT), and for these two combined is 92%. All results are above the LSTM baseline. We show results of the COVID-misinformation *news* models in Table F.2 (Appendix F).

5 Applications on Mega-COV

Now that we have developed two highly effective models, one for COVID-relevance and another for COVID-misinformation, we can employ these models to make discoveries using Mega-COV. Since

¹⁰Even though we could have used the monolingual versions of the transformer-based language models (i.e., BERT and RoBERTa), we stick to the multilingual versions for consistency.

Data	Model	DEV			TEST		
		Precision	Recall	F_1	Precision	Recall	F_1
CoAID	LSTM	81.00	91.00	86.00	95.00	78.00	86.00
	mBERT	91.00	84.00	87.00	94.00	87.00	90.00
	XLM-R _{Base}	93.00	87.00	90.00	87.00	88.00	88.00
	XLM-R _{Large}	98.00	86.00	92.00	97.00	93.00	90.00
ReCOV	LSTM	60.00	56.00	58.00	54.00	57.00	55.00
	mBERT	81.00	59.00	68.00	87.00	55.00	68.00
	XLM-R _{Base}	72.00	58.00	64.00	75.00	55.00	64.00
	XLM-R _{Large}	89.00	52.00	66.00	89.00	51.00	65.00
CoAID+ReCOV	LSTM	79.00	58.00	67.00	66.00	70.00	68.00
	mBERT	94.00	89.00	91.00	94.00	89.00	92.00
	XLM-R _{Base}	88.00	88.00	88.00	88.00	88.00	88.00
	XLM-R _{Large}	86.00	94.00	90.00	85.00	93.00	89.00

Table 9: Performance of our COVID-misinformation *Twitter* models on the fake class only across the 3 settings CoAID, ReCOVeRY, and CoAID+ReCOVeRY. LSTM is our baseline.

our misinformation models are focused only on English (due to the external gold data we used for training being English only), we will restrict this analysis to the English language.¹¹ We were curious whether model predictions will have different distributions on the different types of Twitter posts (i.e., tweets, retweets, and replies). Hence, to enable such comparisons, we extract a random sample of 10M samples from each of these post types (for a total of 30M) from the year 2020 in Mega-COV. We then apply the XLM-R_{Large} *relevance* and *misinformation* models on the extracted samples. Table 10 shows the distribution of predicted labels from each of the two models across the 3 posting types (tweets, retweets, and replies). Strikingly, as the top half of the table shows, while only 7.77% of tweets are predicted as *COVID-related*, almost all retweets (99.84%) are predicted as *related*. This shows that users’ retweets were focused almost exclusively on COVID-19. The table (bottom half) also shows that retweets are highest carriers of content predicted as *fake* (3.67%), followed by tweets (2.3%). From the table, we can also deduce that only 2.45% of all English language Twitter content (average across the 3 posting types) are predicted as fake. ***Given the global use of English, and the large volume of English posts Twitter receives daily, this percentage of fake content is still problematically high.***

5.1 Annotation Study

We perform a human annotation study on a small sample of 150 random posts from those the model predicted as both COVID-related *and* fake. Two annotators labeled the 150 samples for two types of

¹¹But we emphasize the multilingual capacity of our COVID-relevance model.

tags, relevance and veracity. For relevance, all the 150 posts were found relevant by the two annotators (perfect agreement). For veracity, since some posts can be very challenging to identify, we asked annotators to assign one of the 3 tags in the set $\{true, fake, unknown\}$. We did not ask annotators to consult any outside sources (e.g., Wikipedia or independent fact-checking sites) to identify veracity of the samples. Inter-annotator agreement is at $Kappa (K)=77.81\%$, thus indicating almost perfect agreement. On average, annotators assigned the *fake* class 39.39% of the time, the *true* class 3.02%, and the *unknown* class 57.05%. While these findings show that it is hard for humans to identify data veracity without resorting to external sources, it also demonstrates the utility of the model in detecting actual fake stories in the wild. We provide a number of samples from the posts that were automatically tagged as *COVID-related* and either *true* or *false* by our misinformation/veracity model in Table 11.

Model	Data	Prediction	Percentage
COVID Relevance	Tweets	Related	7.77
		Unrelated	92.23
	Retweets	Related	99.84
		Unrelated	0.16
	Replies	Related	12.94
		Unrelated	87.06
COVID Misinfo.	Tweets	Fake	2.3
		True	97.10
	Retweets	Fake	1.38
		True	98.33
	Replies	Fake	3.67
		True	96.62

Table 10: Distribution of predicted labels from our COVID-relevance and COVID-misinformation models on randomly selected 30M English samples from Mega-COV data.

6 Data Release and Ethics

Data Distribution. The size of the data makes it an attractive object of study. Collection and exploration of the data required significant computing infrastructure and use of powerful data streaming and processing tools. To facilitate use of the dataset, we organize the tweet IDs we release by time (month and year) and language. This should enable interested researchers to work with the exact parts of the data related to their research questions even if they do not have large computing infrastructure.

Ethical Considerations. We collect Mega-COV from the public domain (Twitter). In compliance with Twitter policy, we do not publish hydrated tweet content. Rather, we only publish publicly available tweet IDs. All Twitter policies, including respect and protection of user privacy, apply. We decided not to assign geographic region tags to the tweet IDs we distribute, but these already exist on the `json` object retrievable from Twitter. Still, location information should be used with caution. Twitter does not allow deriving or inferring, or storing derived or inferred, potentially sensitive characteristics about users. Sensitive user attributes identified by Twitter include health (e.g., pregnancy), negative financial status or condition, political affiliation or beliefs, religious or philosophical affiliation or beliefs, sex life or sexual orientation, trade union membership, and alleged or actual commission of a crime. If they decide to use Mega-COV, we expect researchers to review Twitter policy¹² and applicable laws, including the European Union’s General Data Protection Regulation (GDPR)¹³, beforehand. *We encourage use of Mega-COV for social good, including applications that can improve health and well-being and enhance online safety.*

7 Related Works

Twitter Datasets for COVID-19. Several works have focused on creating datasets for enabling COVID-19 research. To the best of our knowledge, all these works depend on a list of hashtags related to COVID-19 and focus on a given period of time. For example, Chen et al. (2020) started collecting tweets on Jan. 22nd and continued updating by actively tracking a list of 22 popular keywords such as *#Coronavirus*, *#Corona*, and *#Wuhancoronavirus*. As of May 30, 2020 (Chen et al., 2020) report 144M tweets. Singh et al. (2020) collect a dataset covering January 16 2020-March 15 2020 using a list of hashtags such as *#2019nCoV*, *#ChinaPneumonia* and *#ChinesePneumonia*, for a total of 2.8M tweets, $\sim 18M$ re-tweets, and $\sim 457K$ direct conversations. Using location information on the data, authors report that tweets strongly correlated with newly identified cases in these locations. Similarly, Alqurashi et al. (2020) use a list of keywords and hashtags related to Covid-19 with Twit-

¹²<https://developer.twitter.com/en/developer-terms/policy>

¹³<https://gdpr-info.eu>

Post	Prediction
Vatican confirms Pope Francis and two aides test positive for Coronavirus - MCM Whoaa 🤔 URL	Fake
🤔~CDC recommends men shave their beards to protect against coronavirus – USER URL	Fake
COVID - 19 : Chinese health authorities confirm patient zero ' had sex with bats ' URL	Fake
Royal Palace confirms Queen Elizabeth tests positive for coronavirus URL	Fake
Is COVID - 19 airborne contagious ? New study shows that coronavirus may be caught from the air * 3 - hours * after it has been exposed .	True
A close relative of SARS-CoV - 2 found in bats offers more evidence it evolved naturally URL	True
Antiviral remdesivir prevents disease progression in monkeys with COVID - 19 — National Institutes of Health (NIH) URL	True
COVID Surges Among Young Adults URL	True

Table 11: Sample *Mega-COV* posts predicted as *COVID-related*, and either *true* or *fake* by our models.

ter’s streaming API to collect a dataset of Arabic tweets. The dataset covers the period of March 1 2020-March 30 2020 and is at 4M tweets. The authors’ goal is to help researchers and policy makers study the various societal issues prevailing due to the pandemic. In the same vein, Lopez et al. (2020) collect a dataset of ~ 6.5 M in multiple languages, with English accounting for $\sim 63.4\%$ of the data. The dataset covers January 22 2020-March 2020. Analyzing the data, authors observe the level of retweets to rise abruptly as the crisis ramped up in Europe in late February and early March.

Twitter in emergency and crisis. Social media can play a useful role in disaster and emergency since they provide a mechanism for wide information dissemination (Simon et al., 2015). Examples include use of Twitter information for the Typhoon Haiyan in the Philippines (Takahashi et al., 2015), Tsunami in Padang Indonesia (Carley et al., 2016), the Nepal 2015 earthquakes (Verma et al., 2019), Harvey Hurricane (Marx et al., 2020). A number of works have focused on developing systems for emergency response. An example is McCreadie et al. (2019). Other works focused on developing systems for detecting misuse of social media (Alshehri et al., 2018, 2020; Nagoudi et al., 2020; Elmadany et al., 2020).

Misinformation About COVID-19. Misinformation can spread fast during disaster. Social data have been used to study rumors and various types of fake information related to the Zika (Ghenai and Mejova, 2017) and Ebola (Kalyanam et al., 2015) viruses. In the context of COVID-19, a number of works have focused on investigating the effect of misinformation on mental health (Rosenberg et al., 2020), the types, sources, claims, and responses of a number of pieces of misinformation about COVID-19 (Brennen et al., 2020), the propagation

pattern of rumours about COVID-19 on Twitter and Weibo (Do et al., 2019), check-worthiness (Wright and Augenstein, 2020), modeling the spread of misinformation and related networks about the pandemic (Cinelli et al., 2020; Osho et al., 2020; Pierri et al., 2020; Koubaa, 2020), estimating the rate of misinformation in COVID-19 associated tweets (Kouzy et al., 2020), the use of bots (Ferrara, 2020), predicting whether a user is COVID19 positive or negative (Karisani and Karisani, 2020), and the quality of shared links Singh et al. (2020). Other works have focused on detecting racism and hate speech (Devakumar et al., 2020; Schild et al., 2020; Shimizu, 2020; Lyu et al., 2020) and emotional response (Kleinberg et al., 2020).

8 Conclusion

We presented *Mega-COV*, a billion-scale dataset of 104 languages for studying COVID-19 pandemic. In addition to being large and highly multilingual, our dataset comprises data pre-dating the pandemic. This allows for comparative and longitudinal investigations. We provided a global description of *Mega-COV* in terms of its geographic and temporal coverage, over-viewed its linguistic diversity, and provided analysis of its content based on hashtags and top domains. We also provided a case study of how the data can be used to track global human mobility. The scale of the *Mega-COV* has also allowed us to make a number of striking discoveries, including (1) the shift toward retweeting and replying to other users rather than tweeting in 2020 and (2) the role of international news sites as key sources of information during the pandemic. In addition, we developed effective models for detecting COVID relevance and COVID misinformation and applied them to a large sample of our dataset. Our dataset and models are publicly available.

References

- Muhammad Abdul-Mageed and Lyle Ungar. 2017. EmoNet: Fine-grained emotion detection with gated recurrent neural networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 718–728.
- Sarah Alqurashi, Ahmad Alhindi, and Eisa Alanazi. 2020. Large arabic twitter dataset on covid-19. *arXiv preprint arXiv:2004.04315*.
- Ali Alshehri, El Moatez Billah Nagoudi, and Muhammad Abdul-Mageed. 2020. Understanding and detecting dangerous speech in social media. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 40–47, Marseille, France. European Language Resource Association.
- Ali Alshehri, El Moatez Billah Nagoudi, Alhuzali Hassan, and Muhammad Abdul-Mageed. 2018. Think before your click: Data and models for adult content in Arabic twitter. In *The 2nd Text Analytics for Cybersecurity and Online Safety (TA-COS-2018), LREC*.
- J Scott Brennan, Felix M Simon, Philip N Howard, and Rasmus Kleis Nielsen. 2020. Types, sources, and claims of covid-19 misinformation. *Reuters Institute*.
- Kathleen M Carley, Momin Malik, Peter M Landwehr, Jürgen Pfeffer, and Michael Kowalchuck. 2016. Crowd sourcing disaster management: The complex nature of twitter usage in padang indonesia. *Safety science*, 90:48–61.
- Emily Chen, Kristina Lerman, and Emilio Ferrara. 2020. Covid-19: The first public coronavirus twitter dataset. *arXiv preprint arXiv:2003.07372*.
- Matteo Cinelli, Walter Quattrociocchi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoni, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. 2020. The covid-19 social media infodemic. *arXiv preprint arXiv:2003.05004*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. **Unsupervised cross-lingual representation learning at scale**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Limeng Cui and Dongwon Lee. 2020. Coaid: Covid-19 healthcare misinformation dataset. *arXiv preprint arXiv:2006.00885*.
- Delan Devakumar, Geordan Shannon, Sunil S Bhopal, and Ibrahim Abubakar. 2020. Racism and discrimination in covid-19 responses. *Lancet (London, England)*, 395(10231):1194.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Tien Huu Do, Xiao Luo, Duc Minh Nguyen, and Nikos Deligiannis. 2019. Rumour detection via news propagation dynamics and user representation learning. *arXiv preprint arXiv:1905.03042*.
- Tien Huu Do, Duc Minh Nguyen, Evaggelia Tsiliogianni, Bruno Cornelis, and Nikos Deligiannis. 2018. Twitter user geolocation using deep multiview learning. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6304–6308. IEEE.
- AbdelRahim Elmadany, Chiyu Zhang, Muhammad Abdul-Mageed, and Azadeh Hashemi. 2020. Leveraging affective bidirectional transformers for offensive language detection. In *The 4th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT4), LREC*, pages 102–108.
- Emilio Ferrara. 2020. # covid-19 on twitter: Bots, conspiracies, and social media activism. *arXiv preprint arXiv:2004.09531*.
- Amira Ghenai and Yelena Mejova. 2017. Catching zika fever: Application of crowdsourcing and machine learning for tracking health misinformation on twitter. *arXiv preprint arXiv:1707.03778*.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *CS224N project report, Stanford*, 1(12):2009.
- Mark Graham, Scott A Hale, and Devin Gaffney. 2014. Where in the world are you? geolocation and language identification in twitter. *The Professional Geographer*, 66(4):568–578.
- Bo Han, Afshin Rahimi, Leon Derczynski, and Timothy Baldwin. 2016. Twitter geolocation prediction shared task of the 2016 workshop on noisy user-generated text. In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 213–217.
- Janani Kalyanam, Sumithra Velupillai, Son Doan, Mike Conway, and Gert Lanckriet. 2015. Facts and fabrications about ebola: A twitter based study. *arXiv preprint arXiv:1508.02079*.
- Negin Karisani and Payam Karisani. 2020. Mining coronavirus (covid-19) posts in social media. *arXiv preprint arXiv:2004.06778*.
- Bennett Kleinberg, Isabelle van der Vegt, and Maximilian Mozes. 2020. Measuring emotions in the covid-19 real world worry dataset. *arXiv preprint arXiv:2004.04225*.

- Anis Koubaa. 2020. Understanding the covid19 outbreak: A comparative data analytics and study. *arXiv preprint arXiv:2003.14150*.
- Ramez Kouzy, Joseph Abi Jaoude, Afif Kraitem, Molly B El Alam, Basil Karam, Elio Adib, Jabra Zarka, Cindy Traboulsi, Elie W Akl, and Khalil Baddour. 2020. Coronavirus goes viral: Quantifying the covid-19 misinformation epidemic on twitter. *Cureus*, 12(3).
- Christian E Lopez, Malolan Vasu, and Caleb Gallemore. 2020. Understanding the perception of covid-19 policies by mining a multilanguage twitter dataset. *arXiv preprint arXiv:2003.10359*.
- Marco Lui and Timothy Baldwin. 2012. langid. py: An off-the-shelf language identification tool. In *Proceedings of the ACL 2012 system demonstrations*, pages 25–30. Association for Computational Linguistics.
- Marco Lui and Timothy Baldwin. 2014. Accurate language identification of twitter messages. In *Proceedings of the 5th workshop on language analysis for social media (LASM)*, pages 17–25.
- Hanjia Lyu, Long Chen, Yu Wang, and Jiebo Luo. 2020. Sense and sensibility: Characterizing social media users regarding the use of controversial terms for covid-19. *arXiv preprint arXiv:2004.06307*.
- Julian Marx, Milad Mirbabaie, and Christian Ehnis. 2020. Sense-giving strategies of media organisations in social media disaster communication: Findings from hurricane harvey. *arXiv preprint arXiv:2004.08567*.
- Richard McCreadie, Cody Buntain, and Ian Soboroff. 2019. Trec incident streams: Finding actionable information on social media. *Proceedings of the 16th International Conference on Information Systems for Crisis Response and Management, Valencia, Spain*.
- Saif M Mohammad and Svetlana Kiritchenko. 2015. Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 31(2):301–326.
- El Moatez Billah Nagoudi, AbdelRahim Elmadany, Muhammad Abdul-Mageed, Tariq Alhindi, and Hasan Cavusoglu. 2020. Machine Generation and Detection of Arabic Manipulated and Fake News. In *Proceedings of the Fifth Arabic Natural Language Processing Workshop*, pages 69–84.
- J Ooms and D Sites. 2018. cld2: Google’s Compact Language Detector 2. Retrieved February, 7:2019.
- Abiola Osho, Caden Waters, and George Amariuca. 2020. An information diffusion approach to rumor propagation and identification on twitter. *arXiv preprint arXiv:2002.11104*.
- Francesco Pierri, Carlo Piccardi, and Stefano Ceri. 2020. Topology comparison of twitter diffusion networks effectively reveals misleading information. *Scientific reports*, 10(1):1–9.
- Stephen Roller, Michael Speriosu, Sarat Rallapalli, Benjamin Wing, and Jason Baldrige. 2012. Supervised text-based geolocation using language models on an adaptive grid. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1500–1510. Association for Computational Linguistics.
- Hans Rosenberg, Shahbaz Syed, and Salim Rezaie. 2020. The twitter pandemic: The critical role of twitter in the dissemination of medical information and misinformation during the covid-19 pandemic. *Canadian Journal of Emergency Medicine*, pages 1–7.
- Leonard Schild, Chen Ling, Jeremy Blackburn, Gianluca Stringhini, Yang Zhang, and Savvas Zannettou. 2020. ” go eat a bat, chang! ”: An early look on the emergence of sinophobic behavior on web communities in the face of covid-19. *arXiv preprint arXiv:2004.04046*.
- Kazuki Shimizu. 2020. 2019-ncov, fake news, and racism. *The Lancet*, 395(10225):685–686.
- Tomer Simon, Avishay Goldberg, and Bruria Adini. 2015. Socializing in emergencies—a review of the use of social media in emergency situations. *International Journal of Information Management*, 35(5):609–619.
- Lisa Singh, Shweta Bansal, Leticia Bode, Ceren Budak, Guangqing Chi, Kornraphop Kawintiranon, Colton Padden, Rebecca Vanarsdall, Emily Vraga, and Yanchen Wang. 2020. A first look at covid-19 information and misinformation sharing on twitter. *arXiv preprint arXiv:2003.13907*.
- Bruno Takahashi, Edson C Tandoc Jr, and Christine Carmichael. 2015. Communicating on twitter during a disaster: An analysis of tweets during typhoon haiyan in the philippines. *Computers in Human Behavior*, 50:392–398.
- Rakesh Verma, Samaneh Karimi, Daniel Lee, Omprakash Gnawali, and Azadeh Shakery. 2019. Newswire versus social media for disaster response and recovery. *arXiv preprint arXiv:1906.10607*.
- WHO. 2020. Who statement regarding cluster of pneumonia cases in wuhan, china. *Beijing: WHO*, 9.
- Dustin Wright and Isabelle Augenstein. 2020. Fact check-worthiness detection as positive unlabelled learning. *arXiv preprint arXiv:2003.02736*.
- Xinyi Zhou, Apurva Mulay, Emilio Ferrara, and Reza Zafarani. 2020. Recovery: A multimodal repository for covid-19 news credibility research. *arXiv preprint arXiv:2006.05557*.

Appendices

A Data Distribution




Figure A.1: Frequency of tweeting during Jan-May (10⁷) 2020 vs. Jan-May 2019.

B Hashtag Content Analysis

Hashtags usually correlate with the topics users post about. We provide the top 30 hashtags in the data in Table B.1. As the table shows, users tweet heavily about the pandemic using hashtags such as *COVID19*, *coronavirus*, *Coronavirus*, *COVID19*, *Covid19*, *covid19* and *StayAtHome*. Simple word clouds of hashtags from the various languages (Figure B.1 provides clouds from the top 10 languages) also show COVID-19 topics trending. We also observe hashtags related to **gaming** (e.g., *NowPlaydo*, *PSshare*, and *NintendoSwitch*). This reflects how users may be spending part of their newly-found home time. We also note frequent occurrence of **political** hashtags in languages such Arabic, Farsi, Indian, and Urdu. This is in contrast to discussions in European languages where politics are not as visible. For example, in Urdu, discussions involving the army and border issues show up. This may be partly due to different political environments, but also due to certain European countries such as Italy, Sweden, Spain, and the U.K. being hit harder (and earlier) than many countries in the Middle East and Asia. In Indian languages such as Tamil and Hindi, posts also focused on movies (e.g., *Valimai*), TV




Figure A.2: Geographical diversity in Mega-COV. We show the distribution of our geo-located data over the top 20 countries with most tweets and responses. Overall, 268 countries are represented in the data.

2019		2020	
Hashtag	Freq	Hashtag	Freq
NewProfilePic	64,922	COVID19	260,024
love	41,964	coronavirus	219,615
Repost	39,128	NewProfilePic	102,724
art	35,825	BBB20	91,775
music	28,335	Covid-19	70,106
travel	28,236	COVID-19	67,737
GameofThrones	21,484	Coronavirus	53,251
nature	18,563	covid19	47,940
instagood	18,491	StayHome	44,165
photooftheday	18,032	NintendoSwitch	42,812
tbt	17,332	love	42,497
realestate	17,255	bbb20	39,974
shopmycloset	16,760	NowPlaying	39,069
GameOfThrones	16,127	Repost	37,036
peing	15,930	AnimalCrossing	36,369
fitness	15,623	ACNH	35,528
food	15,358	photography	35,209
BellLetsTalk	14,853	COVID2019	33,512
NowPlaying	14,849	shopmycloset	31,428
family	14,060	music	30,537
style	14,041	StayAtHome	30,313
SoundCloud	13,904	QuedateEnCasa	30,194
WeTheNorth	13,579	stayhome	27,540
GOT	13,458	PS4share	27,487
np	13,335	SocialDistancing	27,376
MyTwitterAnniv.	12,965	lockdown	27,344
Toronto	12,964	TikTok	27,287

Table B.1: Top 30 hashtags in Mega-COV for 2019 vs. 2020.

shows (e.g., *Big Boss*), doctors, and even fake news along with the pandemic-related hashtags.

An interesting observation from the **Chinese language** word cloud is the use of hashtags such as *ChinaPneumonia* and *WuhanPneumonia* to refer to the pandemic. We did not observe these same hash-




Figure B.1: Word clouds for hashtags in tweets from the top 10 languages in the data. We note that tweets in non-English can still carry English hashtags or employ Latin script.

tags in any of the other languages. Additionally, for some reason, **Apple** seems to be trending during the first 4 months of 2020 in China owing to hashtags such as *appledaily* and *appledailytw*. Some languages, such as Romanian and Vietnamese, involve discussions of **bitcoin** and crypto-currency. This was also seen in the Chinese language word cloud, but not as prominently.

C Domain Sharing Analysis

Domains in URLs shared by users also provide a window on what is share-worthy. We perform an analysis of the top 200 domains shared in each of 2019 and 2020. The major observation we reach is the *surge in tweets involving news websites, and the rise in ranks for the majority of these websites compared to 2019*. Table C.1 shows the top 40 news domains in the 2020 data and their change in rank compared to 2019. Such a heavy sharing of news domains reflects users’ needs: Intuitively, at times of global disruption, people need more frequent updates on ongoing events. Of particular importance, especially relative to other ongoing political polarization in the U.S., is the striking rise of the conservative news network *Fox News*, which has moved from a rank of 118 in 2019 to 67 in 2020 with a swooping 51 positions jump. We also note the rank of some news sites (e.g., *The Globe and Mail* and *The Star* going down. This is perhaps due to people resorting to international (and more diverse) sources of information to remain informed about countries other than their own.

Other domains: Other noteworthy domain activities include those related to gaming, video and music, and social media tools. Ranks of these domains have not necessarily shifted higher than 2019 but remain prominent. This shows these themes still being relevant in 2020. In spite of the eco-

Domain	Rank	Domain	Rank
theguardian.com	↑ 3	thehill.com	↑ 51
nytimes.com	↑ 10	globeandmail.com	↓ -38
cnn.com	↑ 18	businessinsdr.com	↑ 31
apple.news	↑ 4	theatlantic.com	↑ 27
washingtonpost.com	↑ 16	newsbreakapp.com	↑ 472
cbc.ca	↓ -13	eldiario.es	↑ 62
bbc.co.uk	↓ -4	apnews.com	↑ 48
bbc.com	↑ 3	abc.es	↑ 89
nyti.ms	↓ -11	reuters.com	↑ 59
foxnews.com	↑ 51	thestar.com	↓ -64
forbes.com	↓ -14	francebleu.fr	↑ 424
nbcnews.com	↑ 39	globalnews.ca	↑ -78
wsj.com	↑ 11	independent.co.uk	↓ -10
bloomberg.com	↑ 13	elmundo.es	↑ 21
ctvnews.ca	↑ 2	indiatimes.com	↑ 0
nypost.com	↑ 100	radio-canada.ca	↓ -66
cnbc.com	↑ 43	lavanguardia.com	↑ 96
usatoday.com	↑ 6	dailymail.co.uk	↑ 23
latimes.com	↑ 23	politico.com	↑ 403
huffpost.com	↑ 66	sky.com	↑ 114

Table C.1: Top 40 domains in 2020 data and their rank change relative to their rank in 2019.

conomic impact of the pandemic, shopping domains such as *etsy.me* and *poshmark.com* have markedly risen in rank as people moved to shopping online in more significant ways. We now introduce a case study as to how our data can be used for mobility tracking.

D Case Study: Mapping Human Mobility with Mega-COV

Geolocation information in Mega-COV can be used to characterize and track human mobility in various ways. We investigate some of these next.

Inter-Region Mobility. Mega-COV can be exploited to generate *responsive maps* where end users can check mobility patterns between different regions over time. In particular, geolocation information can show mobility patterns between regions. As an illustration of this use case, albeit

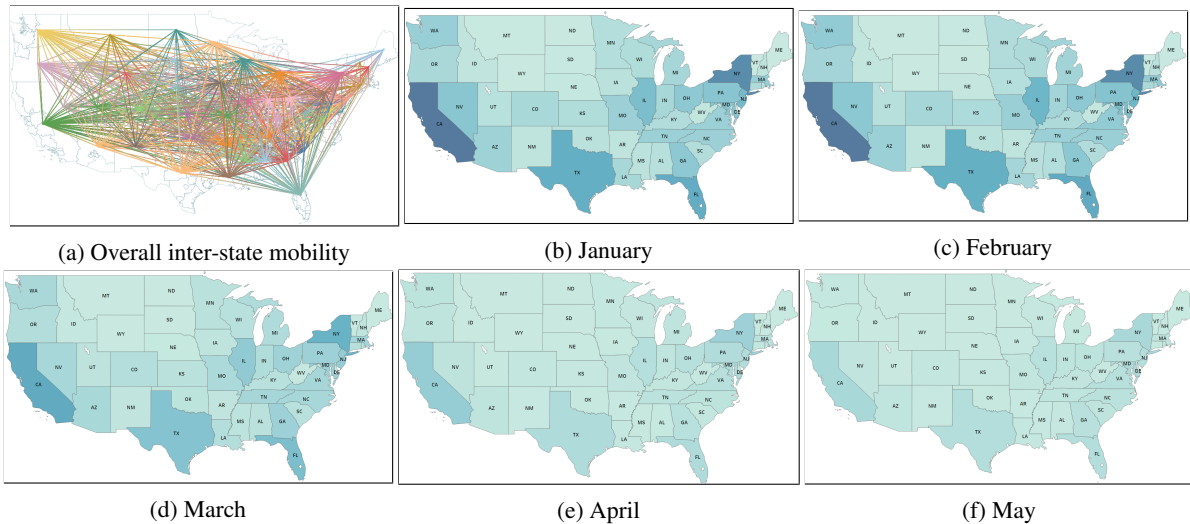


Figure D.1: Inter-state user mobility in the U.S. for Jan-May, 2020.

in a static form, we provide Figure D.1a where we show how users move between U.S. states. We can also exploit *Mega-COV* to show inter-state mobility during a given window of time.¹⁴ Figures D.1b- D.1e present user mobility between U.S. states. The figure shows a clear change from higher mobility in January and February to much less activity in March, April, and May. Clear differences can be seen in key states where the pandemic has hit hard such as New York (NY), California (CA), and Washington State (WA). We provide visualizations of mobility patterns for a number of countries where the pandemic has hit (sometimes hard), as follows: Brazil, Canada, Italy, Saudi Arabia, and the United Kingdom.

Intra-Region Mobility. We also use information in *Mega-COV* to map each user to a single home region (i.e., city, state/province, and country). We follow Geolocation literature (Roller et al., 2012; Graham et al., 2014; Han et al., 2016; Do et al., 2018) in setting a condition that a user must have posted at least 10 tweets from a given region. However, we also condition that at least 60% of all user tweets must have been posted from the same region. We use the resulting set of users whose home location we can verify to map user weekly mobility *within* their own city, state, and country exclusively for both Canada and the U.S. as illustrating examples. We provide the related visualization in supplementary material under “**User Weekly Intra-Region Mobility**”.

¹⁴Here, due to increased posting in 2020, we normalize the number of visits between states by the total number of all tweets posted during a given month.

D.1 User Weekly Intra-Region Mobility

We can also visualize user mobility as a distance from an average mobility score on a weekly basis. Namely, we calculate an *average weekly mobility score* for the year 2019 using geo-tag information (longitude and latitude) and use it as a *baseline* against which we plot user mobility for each week of 2019 and 2020 up until April. In general, we observe a drop in user mobility in Canada starting from mid-March. For U.S. users, we notice a very high mobility surge starting around end of February and early March, only waning down the last week of March and continuing in April as shown in Figure D.8. For both the U.S. and Canada, we hypothesize the surge in early March (much more noticeable in the U.S.) is a result of people moving back to their hometowns, returning from travels, moving for basic needs stocking, etc.

E COVID-Relevance Model

E.1 Dataset

We randomly sample 200K tweets from the English data in Chen et al. (2020) and a maximum of 100K from each of the rest of languages. For languages where there is < 100K tweets, we take all data. For the negative class, we extract data from Jan-Nov, 2019 from *Mega-COV*. For each language, we take roughly the same number of tweets we sampled for the positive class. Table E.1 shows the distribution of the positive class data from Chen et al. (2020).




Figure D.2: User monthly mobility within New York State.

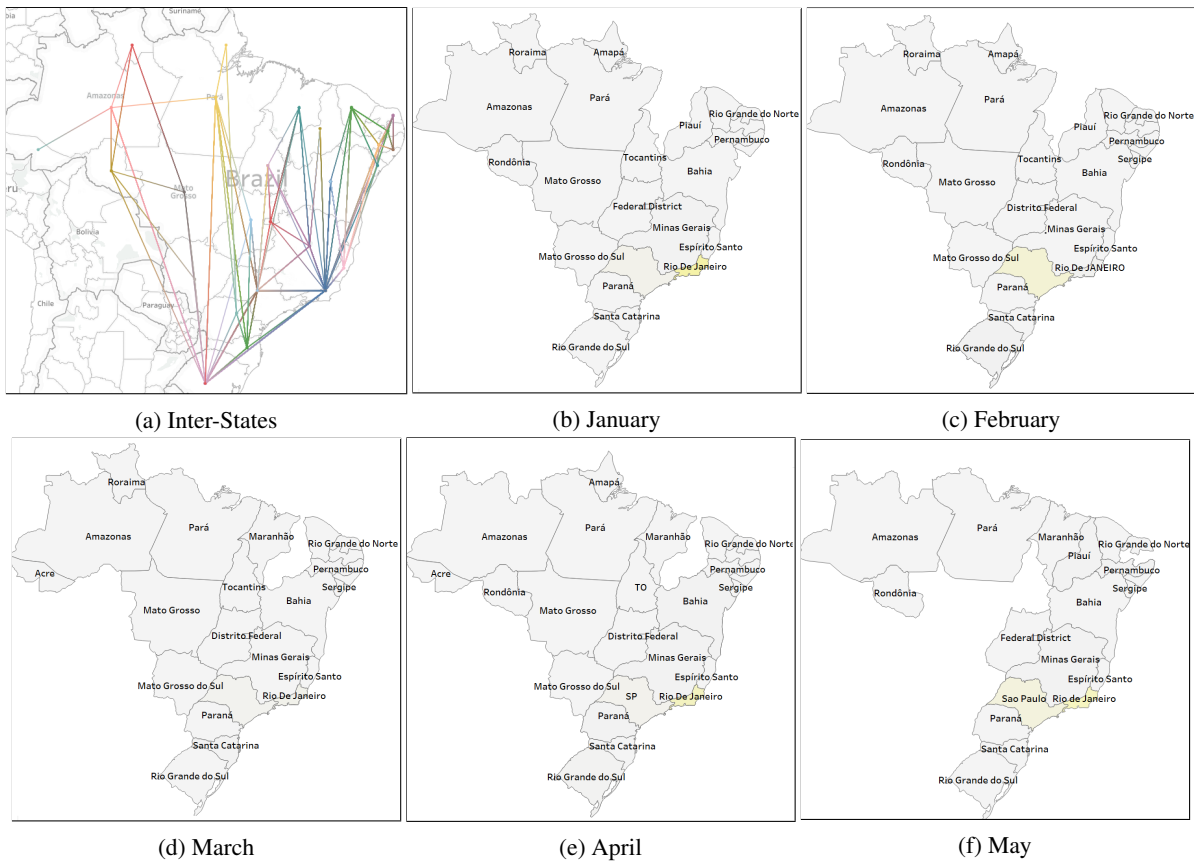


Figure D.3: User mobility between Brazil states (estados) during Jan-May 2020.

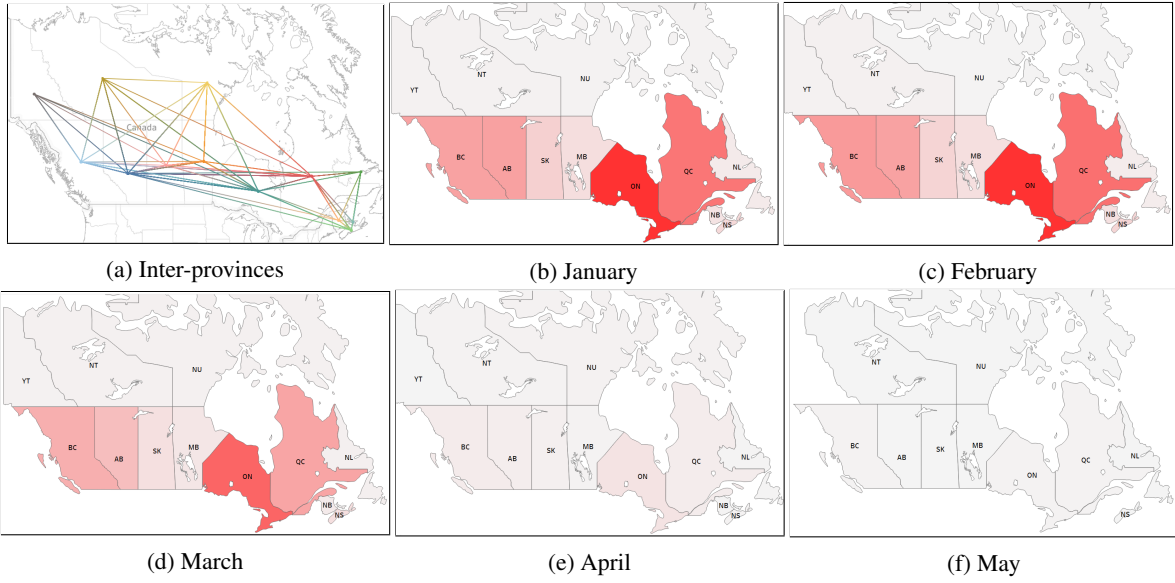


Figure D.4: User mobility between Canada Provinces during Jan-May 2020




Figure D.5: User mobility between Italy regions (regioni) during Jan-May 2020.




Figure D.6: User mobility between Saudi Arabia regions during Jan-May 2020.

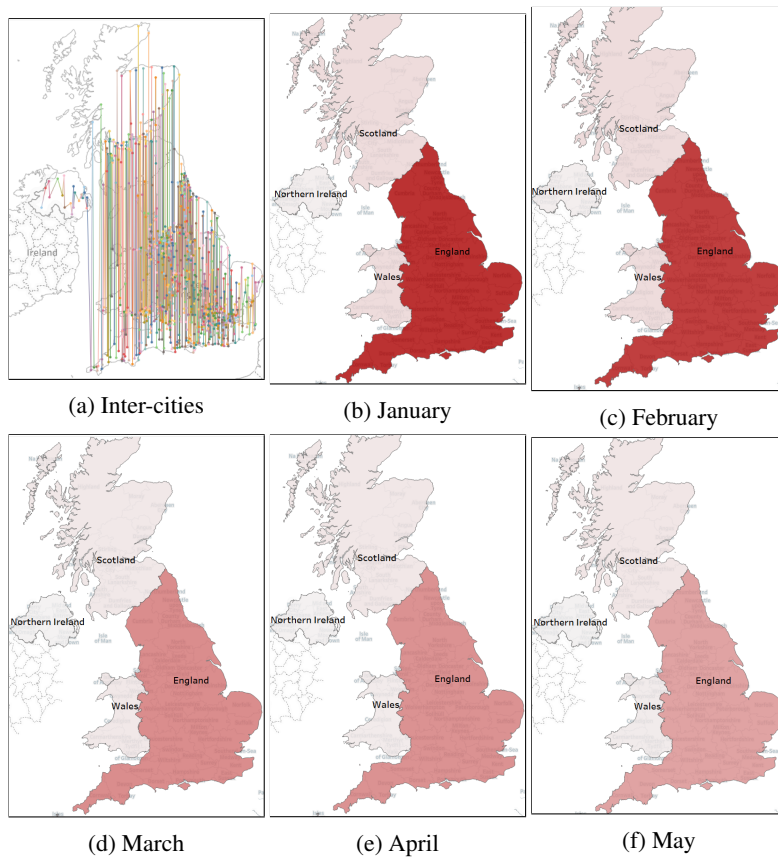




Figure D.7: User mobility between United Kingdom counties during Jan-May 2020.

Replies		RT		Tweets	
Lang	Frequency	Lang	Frequency	Lang	Frequency
Hebrew (he)	337,237	Hebrew (he)	322,351	Hebrew (he)	504,327
Croatian (hr)	242,772	Croatian (hr)	198,764	Croatian (hr)	243,601
Maltese (mt)	94,695	Maltese (mt)	85,054	Maltese (mt)	145,395
Dzongkha (dz)	64,063	Slovak(sk)	77,846	Slovak(sk)	131,544
Bahasa Indonesia (id)	46,463	Latin (la)	66,061	Latin (la)	104,488
Bosnian (bs)	43,066	Bahasa Indonesia (id)	61,295	Bahasa Indonesia (id)	100,561
Slovak(sk)	36,662	Bosnian (bs)	45,276	Bosnian (bs)	54,200
Swahili (sw)	21,803	Swahili (sw)	28,122	Dzongkha (dz)	46,950
Azerbaijani (az)	20,242	Dzongkha (dz)	26,853	Swahili (sw)	42,076
Latin (la)	13,030	Quechua (qu)	22,559	Malay (ms)	32,967
Albanian (sq)	12,878	Malay (ms)	19,511	Quechua (qu)	31,175
Xhosa (xh)	11,936	Esperanto (eo)	19,397	Albanian (sq)	30,361
Irish (ga)	8,607	Kinyarwanda (rw)	19,371	Kinyarwanda (rw)	30,080
Malagasy (mg)	7,727	Azerbaijani (az)	19,182	Azerbaijani (az)	29,507
Quechua (qu)	7,449	Javanese (jv)	18,180	Javanese (jv)	29,121
Kinyarwanda (rw)	7,427	Albanian (sq)	17,904	Esperanto (eo)	29,019
Esperanto (eo)	6,755	Xhosa (xh)	14,886	Kurdish (ku)	24,259
Malay (ms)	6,683	Irish (ga)	14,807	Afrikaans (af)	22,871
Assamese (as)	6,442	Kurdish (ku)	14,475	Volapük (vo)	21,840
Volapük (vo)	6,245	Galician (gl)	13,337	Irish (ga)	21,151

Table D.1: Top 20 languages detected by langid in Mega-COV V0.1 which were not detected by twitter, broken by tweets, retweets, and replies.



(a) Canada users



(b) U.S. users

Figure D.8: Canadian and American user *weekly* mobility during 2019-2020. Each point (a week) is modeled as a mobility distance from weekly average mobility in 2019.

lang	#tweets	lang	#tweets	lang	#tweets	lang	#tweets
en	200K	ar	76K	uk	6.6K	sr	838
es	100K	ru	50K	no	5.5K	bg	739
th	100K	lt	44.7K	eu	4.8K	dv	634
fr	100K	pl	40.6K	cy	4.3K	pa	450
in	100K	fa	32.6K	ne	4K	my	277
ja	100K	ro	32.5K	lv	3.2K	ps	244
pt	100K	sv	24.2K	mr	3K	am	229
it	100K	fi	24K	iw	2.6K	ckb	190
und	100K	vi	22.7K	ml	2.4K	sd	144
tr	100K	et	21.3K	hu	2.2K	km	128
tl	100K	ur	20.3K	te	2.1K	lo	47
de	100K	ht	16.2K	gu	1.8K	hy	34
zh	100K	da	16.2K	bn	1.5K	ka	23
ca	100K	sl	13.5K	kn	1.4K	bo	15
nl	100K	cs	13.3K	or	1.2K	ug	5
ko	100K	ta	13.1K	is	1.2K		
hi	97.4K	el	10.7K	si	1.2K		

Table E.1: Distribution of language in the COVID-Relevance training data for the positive (i.e., *related*) classe.

F COVID-Misinformation Detection

Cui and Lee (2020) present a Covid-19 healthcare misinformation Dataset (CoAID), with diverse COVID-19 healthcare misinformation, including fake news on websites and social platforms, along with related user engagement (i.e., tweets and replies) about such news. CoAID includes 3,235 news articles and claims, 294,692 user engagement, and 851 social platform posts about COVID-19. The topics of CoAID include: $\{COVID-19, coronavirus, pneumonia, flu9, lock\ down, stay\ home, quarantine\ and\ ventilator\}$. The dataset is collected from December 1, 2019 to July 1, 2020 and is organized as follows:

- **News Articles.** To collect the *true* news (not fake), 9 reliable media outlets were identified. These include World Health Organization¹⁵ and the U.S. National Institute of Health¹⁶, for example. To collect *fake* news, 6 fact-checking websites were used (e.g. LeadStories¹⁷, PolitiFact¹⁸).
- **Claims.** The true and fake claims (i.e., news with one or two sentences) were collected using: (1) the official WHO website,¹⁹ (2) WHO official Twitter account,²⁰ and (3) the medical news today website²¹.

¹⁵<https://www.who.int/>

¹⁶<https://www.nih.gov/>

¹⁷<https://leadstories.com/hoax-alert/>

¹⁸<https://www.politifact.com/coronavirus/>

¹⁹<https://www.who.int/>

²⁰<https://twitter.com/who>

²¹<https://www.medicalnewstoday.com>

	News					
	Fake			True		
	TRAIN	DEV	TEST	TRAIN	DEV	TEST
CoAID*	669	84	84	2,172	272	272
ReCOVery	532	66	66	1,091	136	136
	Tweets					
	Fake			True		
	TRAIN	DEV	TEST	TRAIN	DEV	TEST
CoAID	8,072	1,009	1,009	110,076	13,759	13,759
ReCOVery	18,272	2,284	2,284	86,437	10,805	10,805

Table F.1: Statistics of CoAID, ReCOVery, and FakeCovid datasets across the data splits. For CoAID*, we merge the claim and news.

- **User Engagement.** Queries based on the true and fake articles and claims were used to build a dataset of user engagement from Twitter where the goal was to acquire the tweets discussing the news in question and related Twitter replies.

Data	Model	DEV		TEST	
		Acc.	F1	Acc.	F1
CoAID	mBERT	98.88	98.45	97.47	96.48
	XLM-R _{Base}	98.31	97.64	96.35	94.74
	XLM-R _{Large}	99.16	98.84	96.91	95.66
ReCOVery	mBERT	86.76	84.14	85.64	82.73
	XLM-R _{Base}	85.78	83.56	87.13	85.01
	XLM-R _{Large}	88.73	86.36	88.12	85.91
CoAID+ReCOV.	mBERT	93.39	91.41	92.11	89.66
	XLM-R _{Base}	92.50	89.90	91.04	87.79
	XLM-R _{Large}	93.21	90.86	92.83	90.37

Table F.2: Results of our fake news detector models on the DEV and TEST splits of CoAID and ReCOVery news articles datasets.