

## Just Another Day on Twitter: A Complete 24 Hours of Twitter Data

Jürgen Pfeffer<sup>1</sup>, Daniel Matter<sup>1</sup>, Kokil Jaidka<sup>2</sup>, Onur Varol<sup>3</sup>, Afra Mashhadi<sup>4</sup>, Jana Lasser<sup>5, 15</sup>, Dennis Assenmacher<sup>6</sup>, Siqi Wu<sup>7</sup>, Diyi Yang<sup>8</sup>, Cornelia Brantner<sup>9</sup>, Daniel M. Romero<sup>7</sup>, Jahna Otterbacher<sup>10</sup>, Carsten Schwemmer<sup>11</sup>, Kenneth Joseph<sup>12</sup>, David Garcia<sup>13</sup>, Fred Morstatter<sup>14</sup>

<sup>1</sup>School of Social Science and Technology, Technical University of Munich, Germany

<sup>2</sup>Centre for Trusted Internet and Community, National University of Singapore, Singapore

<sup>3</sup>Computer Science Department, Sabanci University, Turkey

<sup>4</sup>School of Science, Technology, Engineering Mathematics, University of Washington (Bothell), USA

<sup>5</sup>Faculty of Computer Science and Biomedical Engineering, Graz University of Technology, Austria

<sup>6</sup>GESIS – Leibniz Institute for the Social Sciences, Germany

<sup>7</sup>School of Information, University of Michigan, USA

<sup>8</sup>Computer Science Department, Stanford University, USA

<sup>9</sup>Department of Geography, Media and Communication, Karlstad University, Sweden

<sup>10</sup>Faculty of Pure and Applied Sciences, Open University of Cyprus & CYENS CoE, Cyprus

<sup>11</sup>Department of Sociology, Ludwig Maximilian University of Munich, Germany

<sup>12</sup>Department of Computer Science and Engineering, University at Buffalo, USA

<sup>13</sup>Department of Politics and Public Administration, University of Konstanz, Germany

<sup>14</sup>Information Sciences Institute, University of Southern California, USA

<sup>15</sup>Complexity Science Hub Vienna, Austria

### Abstract

At the end of October 2022, Elon Musk concluded his acquisition of Twitter. In the weeks and months before that, several questions were publicly discussed that were not only of interest to the platform’s future buyers, but also of high relevance to the Computational Social Science research community. For example, how many active users does the platform have? What percentage of accounts on the site are bots? And, what are the dominating topics and sub-topical spheres on the platform? In a globally coordinated effort of 80 scholars to shed light on these questions, and to offer a dataset that will equip other researchers to do the same, we have collected *all* 375 million tweets published within a 24-hour time period starting on September 21, 2022. To the best of our knowledge, this is the first complete 24-hour Twitter dataset that is available for the research community. With it, the present work aims to accomplish two goals. First, we seek to answer the aforementioned questions and provide descriptive metrics about Twitter that can serve as references for other researchers. Second, we create a baseline dataset for future research that can be used to study the potential impact of the platform’s ownership change.

### Introduction

On March 21, 2006, Twitter’s first CEO Jack Dorsey sent the first message on the platform. In the subsequent 16 years, close to 3 trillion tweets have been sent.<sup>1</sup> Roughly two-thirds of these have been either removed from the platform because the senders deleted them or because the accounts (and all

their tweets) have been banned from the platform, have been made private by the users, or are otherwise inaccessible via the historic search with the v2 API endpoints. By utilizing Twitter’s *count/all API* and the approaches described in this article, we estimate that about 900 billion public tweets were on the platform when Elon Musk acquired Twitter in October 2022 for \$44B<sup>2</sup>.

Besides its possible economic value, Twitter has been instrumental in studying human behavior with social media data and the entire field of Computational Social Science (CSS) has heavily relied on data from Twitter. At the AAAI International Conference on Web and Social Media (ICWSM), in the past two years alone (2021-2022), over 30 scientific papers analyzed a subset of Twitter for a wide range of topics ranging from public and mental health analyses to politics and partisanship. Indeed, since its emergence, Twitter has been described as a digital *socioscope* (i.e., social telescope) by researchers in fields of social science (Mejova, Weber, and Macy 2015), “a massive antenna for social science that makes visible both the very large (e.g., global patterns of communications) and the very small (e.g., hourly changes in emotions)”. Beyond CSS, there is increasing use of Twitter data for training large pre-trained language models in the field of natural language processing and machine learning, such as Bernice (DeLucia et al. 2022), where 2.5 billion tweets are used to develop representations for Twitter-specific languages, and TwHIN-BERT (Zhang et al. 2022) that leverages 7 billion tweets covering over 100 distinct languages to model short, noisy, and user-generated text.

Although Twitter data has fostered interdisciplinary re-

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>While we do not have an official source for this number, it represents an educated guess from a collaboration of dozens of scholars of Twitter.

<sup>2</sup><https://www.nytimes.com/2022/10/27/technology/elon-musk-twitter-deal-complete.html>

search across many fields and has become a “model organism” of big data, scholarship using Twitter data has also been criticized for various forms of bias that can emerge during analyses (Tufekci 2014). One major challenge giving rise to these biases is getting access to data and knowing about data quality and possible data biases (Ruths and Pfeffer 2014; González-Bailón et al. 2014; Olteanu et al. 2019). While Twitter has long served as one of the most collaborative big social media platforms in the context of data-sharing with academic researchers, there nonetheless exists a lack of transparency in sampling procedures and possible biases created from technical artifacts (Morstatter et al. 2013; Pfeffer, Mayer, and Morstatter 2018). These unknown biases may jeopardize research quality. At the same time, access to unfiltered/unsampled Twitter data is nearly impossible to access, and thus the above-mentioned studies, as well as thousands of others, still retain unknown and potentially significant biases in their use of sampled data. Only a few studies have tried to collect complete Twitter data samples. For instance, Geenen et al. (2016) have utilized a third-party data provider to collect all Dutch-speaking tweets over a period of one week. In 2010, Kwak et al. (2010) have crawled the entire Twitter site to get 41.7 million user profiles and 1.47 billion following connections. At around the same time, Cha et al. (2010) found ~55 million accounts and 1.96 billion following connections.

**Contributions.** The data collection efforts presented in this paper were driven by a desire to address these concerns about sampling bias that exist because of the lack of a complete sample of Twitter data. Consequently, the main contribution of this article is to create the first *complete* dataset of 24 hours on Twitter and make these tweets available via future collaborations with the authors and contributors of this article. The dataset collected and described here can be used by the research community to:

- Promote a better understanding of the communication dynamics on the platform. For example, it can be used to answer questions like, how many active (posting) accounts are on the platform? And, what are the dominating languages and topics?
- Create a set of descriptive metrics that can serve as references for the research community and provide context to past and present research papers on Twitter.
- Provide a baseline for the situation *before* the recent sale of Twitter. With the new ownership of Twitter, platform policies as well as the company structures are under significant change, which will create questions about whether previous Twitter studies will be still valuable references for future studies.

In the following sections, we describe the data collection process and provide some descriptive analyses of the dataset. We also discuss ethical considerations and data availability.

## Data

**Data Collection.** We have collected 24 hours of Twitter data from September 20, 15:00:00 UTC to September 21 14:59:59 UTC. The data collection was accomplished by

utilizing the Academic API (Pfeffer et al. 2022) that is free and openly available for researchers. The technical setup of the data collection pipeline was dominated by two major challenges: First, how can we avoid—at least to a satisfying extent—a temporal bias in data collection? Second, how can we get a *good* representation of Twitter? In the following, these two aspects are discussed in more detail.

**What is a complete dataset?** What does *complete* mean when we want to collect a day’s worth of Twitter data? It has been shown previously that the availability of tweets fluctuates, especially in the first couple of minutes (Pfeffer et al. 2022)—people might delete their tweets because of typos, tweets might be removed because of violations of terms of service, etc. To reduce this initial uncertainty, we have decided to collect the data 10 minutes after the tweets were sent. Consequently, this dataset does not include all tweets that were sent on the collection day but instead tries to create a somewhat *stable* representation of Twitter.

**Avoiding temporal collection bias.** We wanted to collect a set of tweets close to the time when they were created. However, collecting data takes time, which can introduce possible temporal bias, e.g., if we want to collect data from the previous hour and the data collection job takes three hours, then the data that is collected at the end of the collection job will be much older (with potentially more tweet removals) than the data that is collected at the beginning. To tackle this challenge, we have split the day into 86,400 collection tasks, each consisting of 1 second of Twitter activity. The collection of every second of data started exactly 10 minutes after the data creation time. Because the data collection of a second took more than a minute during peak times, we have distributed the workload to 80 collection processes, i.e., Academic API tokens, in order to avoid backlogs.

**Data collection queries.** The backbone of our data collection effort is a query that—to the best of our knowledge—is able to collect *ALL* tweets within a specified time frame. The query is based on the following three aspects. First, the Academic API allows for negative selectors to limit a search query, e.g., “indictment -trump” will return a tweet including the term “indictment” only if the tweet does not include the term “trump”. While simple “A and not A” queries are not allowed to collect all possible Tweets, a negative selection is possible when combined with a language selector. The query “-trump lang:de” will return all German tweets that do not include the term “trump”. If we now replace the term “trump” with a long random string, we will receive *all* German tweets. Second, Twitter assigns a language code to *every* tweet<sup>3</sup> (including “und” for undefined). Consequently, we can construct a long *OR*-condition with all possible languages: (lang:am OR lang:ar OR ...). We can get the list of all currently used Twitter languages by requesting a non-existing language, e.g., “-trump lang:test” will return an error including a list of all possible languages—currently, there are 74 possible languages. Combined with the exclusion of the long random string, this will query all tweets on the platform. Third, it is possible to limit the time frame

<sup>3</sup>We have verified this with 1% Sample API data.

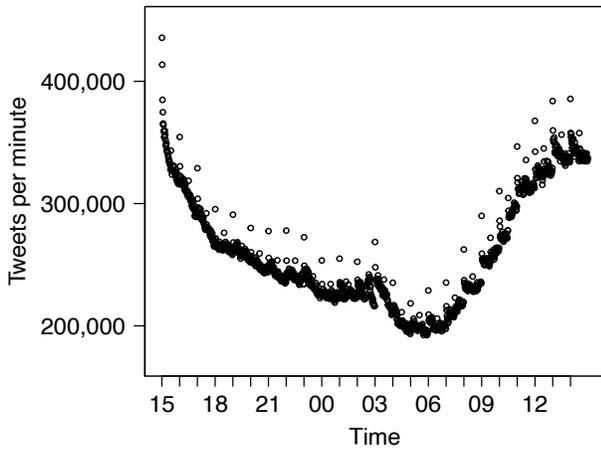


Figure 1: Tweets per minute over the 24-hour collection period, time in UTC.

for data collection to a single second, e.g., by setting the start to “2022-09-21T12:00:00Z” and the end to “2022-09-21T12:00:01Z”.

**Number of tweets.** With the above-described process, we have collected 374,937,971 tweets within the 24-hour time span. On average, this amounts to 4,340 [2,989 – 8,955] tweets per second. Fig. 1 plots the number of tweets per Minute (avg=260,374, min=192,322, max=435,721). The data collection started at 15:00 UTC, when almost the entire Twitter world is awake. Then, we can see from Japan to Europe time zone after time zone getting off the platform. While Europe and the Americas are sleeping, Asia keeps the number of tweets at around 200,000. Starting at 7:00 UTC, Europe is getting active again, followed by the Americas from East to West. Another astonishing observation of this time series is that the first minute of every hour has on average 15.5% more tweets than the minute before—most likely due to bot activities and other timed tweet releases, e.g., news.

## Descriptive Analyses

### Active Users

The 375 million tweets in our dataset were sent by 40,199,195 accounts. While the publicly communicated numbers of users of a platform are often based on the number of active and passive visitors, we can state that Twitter has (or at least had on our observed day) 40 million active contributors who have sent at least one tweet. Less than 100 accounts have created about 1% (=3.5M) tweets. ~175,000 accounts (0.44%) created 50% of all tweets.

These numbers are not surprising when we consider that > 95% of active accounts have sent one or two tweets. However, these numbers lend more nuance to recent reports from the Pew Research Center, which reported that while the majority of Americans use social media, approximately 97% of all tweets were posted by 25% of the users (McClain 2021). In fact, our dataset suggests that worldwide, the numbers may be more skewed than previously suggested.

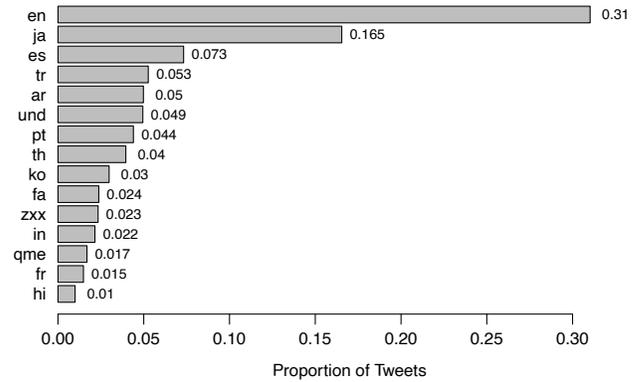


Figure 2: All languages occurring in at least 1% of the tweets.

### User Metrics

**Followers.** The active accounts on our day of Twitter data have a mean of 2,123 followers (median=99). We can find six accounts with more than 100 million followers (max=133,301,854), and 427/8,635 accounts with more than 10/1 million followers. Exactly 50% of accounts that were active on our collection day have less than 100 followers.

**Following.** These accounts follow much fewer other accounts: mean=547, median=197, range: 0–4,103,801. Interestingly, there are 2,377 accounts that follow more than 100,000 other accounts. One-third of accounts follow less than 100 accounts, but only 1.7% of accounts follow zero other accounts.

**Listed.** Lists are a Twitter feature for users to organize accounts around topics and filter tweets. While there is little evidence that lists are used widely on the platform, this feature might be useful for getting an impression about the number of *interesting* content creators on the platform. The 40 million active accounts in our dataset are listed (i.e., number of lists that include a user) in 0 to 3,086,443 lists (mean=10.1, median=0). 1,692/46,139 accounts are in lists of at least 10,000/1,000 accounts.

**Tweets sent.** The user information of the tweet metadata also includes the number of tweets that a user has sent—or at least how many of those tweets are still available on Twitter. The sum of the sent tweets variable of all 40 million accounts is ~404 billion (mean=9,704, median=1,522). If we assume that our initial estimate of having 900 billion tweets on the platform at the time of data collection is somewhat

% Total Tweets	% Total Users	Min. no. of Tweets
1%	0.00023%	2,267
10%	0.01199%	465
25%	0.07284%	152
50%	0.43526%	39
75%	1.70955%	11
90%	4.18836%	3

Table 1: Distribution of user activity

correct, the accounts active in our dataset have contributed ~45% of all of the available tweets over the entire lifetime of Twitter.

**Verified accounts.** At the time of our data collection, we can identify 221,246 verified accounts among the 40 million active users.

### Tweets and Retweets

79.2% of all tweets refer to other tweets, i.e. they are retweets or quotes of or replies to other tweets. Consequently, 20.8% of the tweets in our dataset are original tweets. The tweets with references are of the following types: 50.7% retweets, 4.3% quotes, 24.2% replies, i.e. half of all tweets are retweets and a fourth are replies.

**Retweeted and liked.** Studying the retweet and like numbers from the tweets' metadata has created little insight since the top retweeted tweets are very old tweets that have been retweeted by chance on our collection day. Furthermore, we can see the number of likes only for tweets that have been tweeted and retweeted. In any case, the retweeted number is interesting—the 374 million tweets have been retweeted 401 billion times. In other words, significant parts of historic Twitter get retweeted on a daily basis.

### Languages

Twitter annotates a language variable for every tweet. Fig. 2 shows those languages that were annotated on at least 1% of our dataset. Together, these 15 languages make up 92.5% of all tweets. Besides the most common languages on Twitter, we can also find interesting language codes in this list: *und* stands for undefined and represents tweets for which Twitter was not able to identify a language; *qme* and *zxx* seem to be used by Twitter for tweets consisting of only media or a Twitter card.

### Media

There are 112,779,266 media attachments in our data collection (76.9% photos, 20.7% videos, 2.4% animated GIFs), of which 37,803,473 have unique media keys (83.8% photos, 10.0% videos, 6.2% animated GIFs).

### Geo-Tags

We found only 0.5% of tweets to be geo-tagged. This is not surprising as previous works have shown that the percentage of geo-tagging in Twitter has been declining (Ajao, Hong, and Liu 2015). Fig. 3 shows the distribution of the geo-tagged tweets across the world, with USA (20%), Brazil (11%), Japan (8%), Saudi Arabia (6%) and India (4%) being the top five countries.

### Estimating Prevalence of Bot Accounts

Twitter has a pivotal role in public discourse and entities that are after power and influence often utilize this platform through social bots and other means of automated activities. Since the early days of Twitter, researchers have been studying bot behavior, and it has become an active research area (Ferrara et al. 2016; Cresci 2020). The first estimation

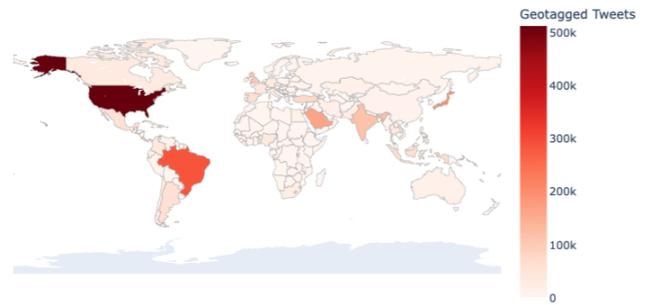


Figure 3: Choropleth map of the geo-tagged tweets across the world.

of bot prevalence on Twitter indicates that 9-15% of Twitter accounts exhibit automated behavior (Varol et al. 2017), while others have observed significantly higher percentages of tweets produced by bot-likely accounts on specific discourses (Uyheng and Carley 2021; Antenore, Camacho Rodriguez, and Panizzi 2022). One major challenge in estimating bot prevalence is the variety of definitions, datasets, and models used for detection (Varol 2022).

In this study, we employed BotometerLite (Yang et al. 2020), a scalable and light-weight version of the Botometer (Sayyadiharikandeh et al. 2020), for computing bot scores for unique accounts in our collection. In Fig. 4a, we present the distribution of bot scores and nearly 20% of the 40 million active accounts have scores above 0.5 suggesting bot-likely behavior.

While identification of bots is a complex and possibly controversial challenge, plotting the distributions of BotometerLite scores grouped by account age in Fig. 4b suggests the proportions of accounts that show bot-like behavior has increased dramatically in recent years. This result may also suggest that the longevity of simpler bot accounts is limited and they are no longer active on the platform. In Fig. 4c, we also present the distribution of bot scores for different rates of activities in our dataset. Accounts that have over 1,000 posts exhibit higher rates of bot-like behaviors.

It is important to mention that accounts studied in this paper were identified due to their content creation activities. Our collection cannot capture passive accounts that are simply used to boost follower counts without visible activity on tweet streams. Fair assessment of bot prevalence is only possible with complete access to Twitter's internal database; since activity streams, network data, and historical tweet archives can capture different sets of accounts (Varol 2022).

### Content on Twitter

The top 500 hashtags occurred 81,468,508 times in the tweets. Via manual inspection, we were able to identify the meaning of 95% of these top hashtags. They can be aggregated into ten the categories.

Table 2 suggests that a large proportion of tweets referred to entertainment, which together comprised about 30% of tweets. These included mentions of celebrities (25.5%) and other entertainment-related tweets (5.4%) such as mentions of South Korean boy band members, and other references

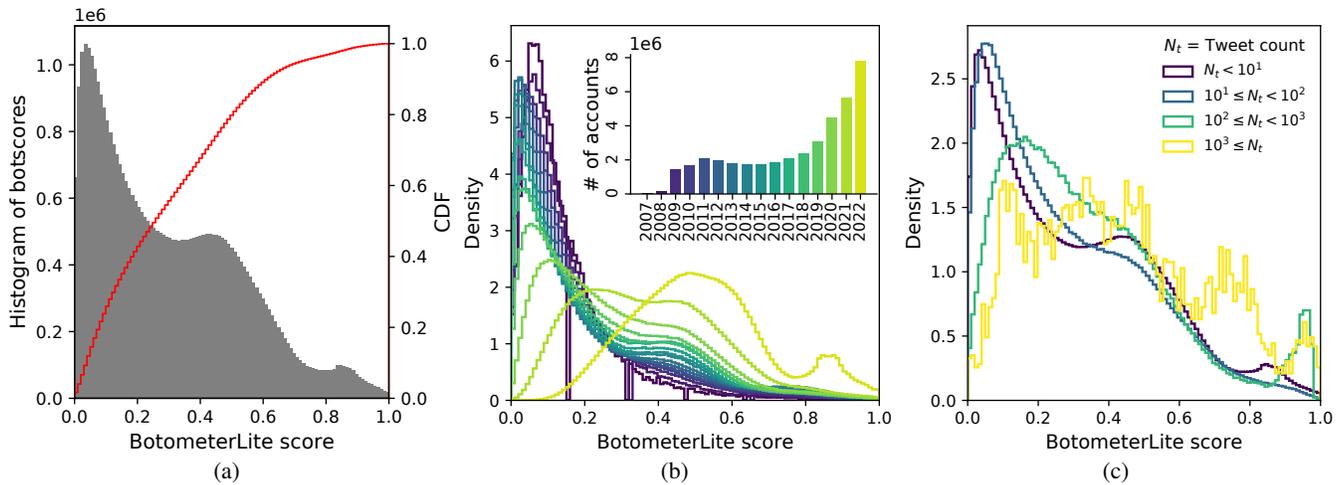


Figure 4: BotometerLite scores distribution: (a) histogram and cumulative distribution, (b) by account age, (c) by tweet counts in our dataset.

to music, movies, and TV shows. Our data collection time window occurred during Fall/Winter 2022, when the world was discussing the protests in Iran after the death of Mahsa Amini. Therefore, the Iranian protests also comprised a large proportion of the hashtag volume at 16.6%.

Finally, and perhaps surprisingly, the category *sex* comprised over a quarter of all content covered by the top hashtags, and was almost completely related to escorts. “Other” topics reflect that on “regular” Twitter days, sports, tech, and art may take up only about 3.3% of Twitter volume.

Fig. 5 is a hashtag visualization that attempts to provide an overview of the entire content on Twitter. We first removed all tweets from accounts with more than 240 tweets (=10% of the max. daily allowed number of tweets) to reduce the noise from bots using random trending hashtags. From the remaining tweets, we extracted the 10,000 most often used hashtags in our dataset and created a hashtag similarity matrix with the number of accounts that have used a pair of two hashtags on the day of data collection. Every element in Fig. 5 represents a hashtag. The position is the result of Multidimensional Scaling (MDS) and the color shows the dominant language that was used in the tweets with the particular hashtag. In this figure, we can see how languages separate the Twitter universe but that there are also topical sub-communities within languages.

## Discussion and Potential Applications

Twitter is a social media platform with a worldwide user-base. Open access to its data also makes it attractive to a large community of researchers, journalists, technologists, and policymakers who are interested in examining social and civic behavior online. Early studies of Twitter explored who says what to whom on Twitter (Wu et al. 2011), characterizing its primary use as a communication tool. Other early work mapped follower communities through ego networks (Gruzd, Wellman, and Takhteyev 2011). However,

Twitter has since expanded into its own universe, with a plethora of users, uses, modalities, communities, and real-life implications. Twitter is increasingly the source of breaking news, and many studies from the U.S. and Europe have reported that Twitter is one of the primary sources of news for their citizens. Twitter has been used for political engagement and citizen activism worldwide. During the COVID-19 pandemic, Twitter even assumed the role of the official mouthpiece and crisis communication tool for many governments to contact their citizens, and from which citizens could seek help and information.

Fig. 3 confirms prior reports that geotagging practices are limited in many low- and middle-income countries (Malik et al. 2015); however, this should not deter scholars from exploring alternative methods of triangulating the location of users (Schwartz et al. 2013), and creating post-stratified estimates of regional language use (Jaidka et al. 2020; Giorgi et al. 2022). In prior studies, the difficulties in widespread data collection and analyses have so far implied that most answers are based on smaller samples (usually constrained by geography, for convenience) of a burgeoning Twitter population. Fig. 5 and Table 2 also impressively illustrate that Twitter is about so much more than US politics.

We hope that our dataset is the first step in creating alternatives for conducting a representative and truly inclusive analysis of the Twitterverse. Temporal snapshots are invaluable to map the national and international migration patterns that increasingly blur geopolitical boundaries (Zagheni et al. 2014).

The increasing popularity of Twitter has led it into issues of scale, where its moderation can no longer check the large proportion of bots on the platform. Our findings in Fig. 4 indicate that the infestation of bots may be more pernicious than previously imagined. We are especially concerned that the escalation of the war on Ukraine by Russia may reflect a spike (in our dataset) in the online activity of bots from Russia operated either by the Russian government or its allied

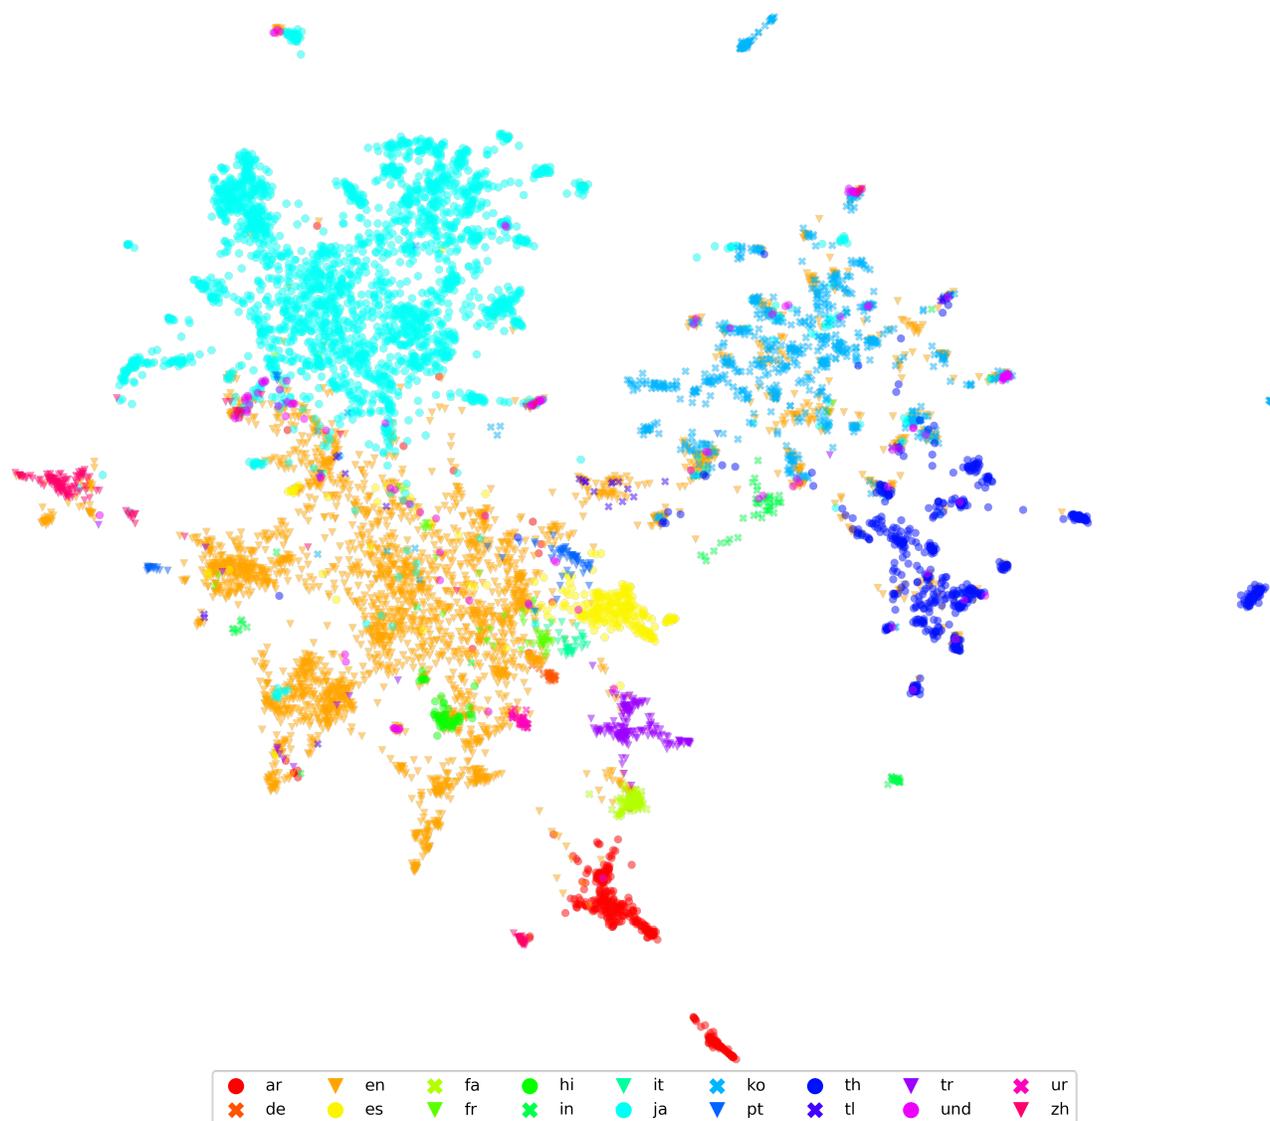


Figure 5: MDS of top 10,000 hashtags based on co-usage by same accounts; colors represent dominant language in tweets using a hashtag.

intelligence agencies (Badawy, Ferrara, and Lerman 2018). These and other bots serve to amplify trending topics and facilitate the spread of misinformation (though, perhaps, at a rate less than humans do (Vosoughi, Roy, and Aral 2018)). They may also misuse hashtags to divert attention away from social or political topics (Earl, Maher, and Pan 2022; Broniatowski et al. 2018) or strategically target influential users (Shao et al. 2018; Varol and Uluturk 2020). We hope that our work will spur more studies on these topics, and we welcome researchers to explore our data.

By observing bursts of discussions around politically charged events and characterizing the temporal spikes in Twitter topics, we can better rationalize how our experience of Twitter as a political hotbed differs from the simplified understanding of the American Twitter landscape reported

in Mukerjee, Jaidka, and Lelkes (2022), which suggested that politics is largely a sideshow on Twitter. It is worth considering that these politically active users may not be representative of social media users at large (McClain 2021; Wojcieszak et al. 2022).

Twitter is also under scrutiny for how its platform governance may conflict with users' interests and rights (Van Dijck, Poell, and De Waal 2018). Concerns have been raised about alleged biases in the algorithmic amplification (and deamplification) of content, with evidence from France, Germany, Turkey, and the United States, among other countries (Majó-Vázquez et al. 2021; Tanash et al. 2015; Jaidka, Mukerjee, and Lelkes 2023). Other scholars have also criticized Twitter's use as a censorship weapon by governments and political propagandists worldwide (Varol 2016; Elmas,

Overdorf, and Aberer 2021; Jakesch et al. 2021). They, and others, may be interested in examining the trends in the enforcement of content moderation policies by Twitter.

Besides answering questions of data, representativeness, access, and censorship, we anticipate that our dataset is suited to explore the temporal dynamics of online (mis)information in the following directions:

- **Content characteristics:** We have provided a high-level exploration of the topics on Twitter. However, more can be done with regard to understanding users’ concerns and priorities. While hashtags act as signposts for the broader Twitter community to find and engage in topics of mutual interest (Cunha et al. 2011), tweets without hashtags may offer a different understanding of Twitter discourse, where users may engage in more interpersonal discussions of news, politics, and sports than the numbers suggest (Rajadesingan, Budak, and Resnick 2021).
- **Patterns of information dissemination:** Informational exchanges occurring on Twitter can overcome spatio-temporal limitations as they essentially reconfigure user connections to create newly emergent communities. However, these communities may vanish as quickly as they are created, as the lifecycle of a tweet determines how long it continues to circulate on Twitter timelines. To the best of our knowledge, no prior research has reported on the average “age” of a tweet, and we hope that a 24-hour snapshot will enable us to answer this question empirically.
- **Content moderation and fake news:** Prior research suggests that 0.1% of Twitter users accounted for 80% of all fake news sources shared in the lead-up to a US election (Grinberg et al. 2019). However, we expect there to be cross-lingual differences in this distribution, especially for low- or under-resourced languages with fewer open tools for fact-checking. Similarly, we expect that the quality of moderation and hate speech will vary by geography and language, and recommend the use of multilingual large language models to explore these trends (with attention to persisting representativeness caveats (Wu and Dredze 2020)).
- **Mass mobilization:** Twitter is increasingly the hotbed of protest, which has led to some activists donning the role

Category	Hashtags	Occurrence	
Celebrities	159	20,809,742	25.5%
Sex	104	20,529,196	25.2%
Iranian Protests	15	13,488,295	16.6%
Entertainment	45	4,392,227	5.4%
Advertisement	32	4,644,540	5.7%
Politics	38	3,858,550	4.7%
Finance	30	3,549,107	4.4%
Games	21	3,348,128	4.1%
Other	31	2,672,291	3.3%
Unknown	25	4,176,432	5.1%
Sum	500	81,468,508	100.0%

Table 2: The categories of the top 500 hashtags in the dataset

of “movement spillovers” (Zhou and Yang 2021) or serial activists (Bastos and Mercea 2016) who broker information across different online movements, thereby acting as key coordinators, itinerants, or gatekeepers in the exchange of information. Such users, as well as the constant communities in which they presumably reside (Chowdhury et al. 2022), may be easier to study through temporal snapshots, as facilitated by this dataset.

- **Echo chambers and filter bubbles:** On Twitter, algorithms can affect the information diets of users in over 200 countries, with an estimated 396.5 million monthly users (Kemp 2022). Recent surveys of the literature have considered the evidence on how platforms’ designs and affordances influence users behaviors, attitudes, and beliefs (González-Bailón and Lelkes 2022). Studies of the structural and informational networks based on snapshots of Twitter can offer clues to solving these puzzles without the constraints of data selection.

## Ethics Statement and Data Availability

**Ethics statement.** We acknowledge that privacy and ethical concerns are associated with collecting and using social media data for research. However, we took several steps to avoid risks to human subjects since participants no longer opt into being part of our study, in a traditional sense (Zimmer 2020). In our analysis, we only studied and reported population level, and aggregated observations of our dataset. We share publicly only the tweet IDs with the research community to account for privacy issues and Twitter’s TOS. For this purpose, we use a data sharing and long-term archiving service provided by GESIS - Leibniz Institute for the Social Sciences, a German infrastructure institute for the social sciences <sup>4</sup>.

With regards to data availability, this repository adheres to the FAIR principles (Wilkinson et al. 2016) as follows:

- **Findability:** In compliance with Twitter’s terms of service, only tweet IDs are made publicly available at DOI: <https://doi.org/10.7802/2516>. A unique Document Object Identifier (DOI) is associated with the dataset. Its metadata and licenses are also readily available.
- **Accessibility:** The dataset can be downloaded using standard APIs and communications protocols (the REST API and OAI-PMH).
- **Interoperability:** The data is provided in raw text format.
- **Reusability:** The CC BY 4.0 license implies that researchers are free to use the data with proper attribution.

In light of the recent changes to Twitter’s APIs, we expect significant limitations when accessing tweets. Consequently, we want to invite the broader research community to approach one or more of the authors and collaborators (see Acknowledgments) of this paper with research ideas about what can be done with this dataset. We will be very happy to collaborate with you!

<sup>4</sup><https://www.gesis.org/en/data-services/share-data>

## Acknowledgments

The data collection effort described in this paper could not have been possible without the great collaboration of a large number of scholars, here are some of them (in random order): Chris Schoenherr, Leonard Husmann, Diyi Liu, Benedict Witzemberger, Joan Rodriguez-Amat, Florian Angermeir, Stefanie Walter, Laura Mahrenbach, Isaac Bravo, Anahit Sargsyan, Luca Maria Aiello, Sophie Brandt, Wienke Strathern, Bilal Çakir, David Schoch, Yuliia Holubosh, Savvas Zannettou, Kyriaki Kalimeri.

## References

- Ajao, O.; Hong, J.; and Liu, W. 2015. A survey of location inference techniques on Twitter. *Journal of Information Science*, 41(6): 855–864.
- Antenore, M.; Camacho Rodriguez, J. M.; and Panizzi, E. 2022. A Comparative Study of Bot Detection Techniques With an Application in Twitter Covid-19 Discourse. *Social Science Computer Review*, 08944393211073733.
- Badawy, A.; Ferrara, E.; and Lerman, K. 2018. Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. In *2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)*, 258–265. IEEE.
- Bastos, M. T.; and Mercea, D. 2016. Serial activists: Political Twitter beyond influentials and the twitterariat. *New Media & Society*, 18(10): 2359–2378.
- Broniatowski, D. A.; Jamison, A. M.; Qi, S.; AlKulaib, L.; Chen, T.; Benton, A.; Quinn, S. C.; and Dredze, M. 2018. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American journal of public health*, 108(10): 1378–1384.
- Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, K. 2010. Measuring User Influence in Twitter: The Million Follower Fallacy. *Proceedings of the International AAAI Conference on Web and Social Media*, 4(1): 10–17.
- Chowdhury, A.; Srinivasan, S.; Bhowmick, S.; Mukherjee, A.; and Ghosh, K. 2022. Constant community identification in million-scale networks. *Social Network Analysis and Mining*, 12(1): 1–17.
- Cresci, S. 2020. A decade of social bot detection. *Communications of the ACM*, 63(10): 72–83.
- Cunha, E.; Magno, G.; Comarela, G.; Almeida, V.; Gonçalves, M. A.; and Benevenuto, F. 2011. Analyzing the dynamic evolution of hashtags on twitter: a language-based approach. In *Proceedings of the workshop on language in social media (LSM 2011)*, 58–65.
- DeLucia, A.; Wu, S.; Mueller, A.; Aguirre, C.; Dredze, M.; and Resnik, P. 2022. Bernice: A Multilingual Pre-trained Encoder for Twitter. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 6191–6205.
- Earl, J.; Maher, T. V.; and Pan, J. 2022. The digital repression of social movements, protest, and activism: A synthetic review. *Science Advances*, 8(10): eab18198.
- Elmas, T.; Overdorf, R.; and Aberer, K. 2021. A Dataset of State-Censored Tweets. In *ICWSM*, 1009–1015.
- Ferrara, E.; Varol, O.; Davis, C.; Menczer, F.; and Flammini, A. 2016. The rise of social bots. *Communications of the ACM*, 59(7): 96–104.
- Giorgi, S.; Lynn, V. E.; Gupta, K.; Ahmed, F.; Matz, S.; Ungar, L. H.; and Schwartz, H. A. 2022. Correcting Sociodemographic Selection Biases for Population Prediction from Social Media. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 228–240.
- González-Bailón, S.; and Lelkes, Y. 2022. Do social media undermine social cohesion? A critical review. *Social Issues and Policy Review*.
- González-Bailón, S.; Wang, N.; Rivero, A.; Borge-Holthoefer, J.; and Moreno, Y. 2014. Assessing the bias in samples of large online networks. *Social Networks*, 38: 16–27.
- Grinberg, N.; Joseph, K.; Friedland, L.; Swire-Thompson, B.; and Lazer, D. 2019. Fake news on Twitter during the 2016 US presidential election. *Science*, 363(6425): 374–378.
- Gruzd, A.; Wellman, B.; and Takhteyev, Y. 2011. Imagining Twitter as an imagined community. *American Behavioral Scientist*, 55(10): 1294–1318.
- Jaidka, K.; Giorgi, S.; Schwartz, H. A.; Kern, M. L.; Ungar, L. H.; and Eichstaedt, J. C. 2020. Estimating geographic subjective well-being from Twitter: A comparison of dictionary and data-driven language methods. *Proceedings of the National Academy of Sciences*, 117(19): 10165–10171.
- Jaidka, K.; Mukerjee, S.; and Lelkes, Y. 2023. Silenced on social media: the gatekeeping functions of shadowbans in the American Twitterverse. *Journal of Communication*.
- Jakesch, M.; Garimella, K.; Eckles, D.; and Naaman, M. 2021. Trend alert: A cross-platform organization manipulated Twitter trends in the Indian general election. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–19.
- Kemp, S. 2022. Digital 2022: Global overview report. Technical report, DataReportal.
- Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a Social Network or a News Media? In *Proceedings of the 19th International Conference on World Wide Web*, 591–600. Association for Computing Machinery. ISBN 9781605587998.
- Majó-Vázquez, S.; Congosto, M.; Nicholls, T.; and Nielsen, R. K. 2021. The Role of Suspended Accounts in Political Discussion on Social Media: Analysis of the 2017 French, UK and German Elections. *Social Media+ Society*, 7(3): 20563051211027202.
- Malik, M.; Lamba, H.; Nakos, C.; and Pfeffer, J. 2015. Population bias in geotagged tweets. In *proceedings of the international AAAI conference on web and social media*, volume 9, 18–27.
- McClain, C. 2021. 70% of U.S. social media users never or rarely post or share about political, social issues. Technical report, Pew Research Center.

- Mejova, Y.; Weber, I.; and Macy, M. W. 2015. *Twitter: a digital socioscope*. Cambridge University Press.
- Morstatter, F.; Pfeffer, J.; Liu, H.; and Carley, K. M. 2013. Is the Sample Good Enough? Comparing Data from Twitter’s Streaming API with Twitter’s Firehose. In *Seventh International AAAI Conference on Weblogs and Social Media*, 400–408.
- Mukerjee, S.; Jaidka, K.; and Lelkes, Y. 2022. The Political Landscape of the US Twittersphere. *Political Communication*, 1–31.
- Olteanu, A.; Castillo, C.; Diaz, F.; and Kıcıman, E. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2: 13.
- Pfeffer, J.; Mayer, K.; and Morstatter, F. 2018. Tampering with Twitter’s Sample API. *EPJ Data Science*, 7(50).
- Pfeffer, J.; Mooseder, A.; Lasser, J.; Hammer, L.; Stritzel, O.; and Garcia, D. 2022. This Sample seems to be good enough! Assessing Coverage and Temporal Reliability of Twitter’s Academic API.
- Rajadesingan, A.; Budak, C.; and Resnick, P. 2021. Political discussion is abundant in non-political subreddits (and less toxic). In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media*, volume 15.
- Ruths, D.; and Pfeffer, J. 2014. Social Media for Large Studies of Behavior. *Science*, 346(6213): 1063–1064.
- Sayyadiharikandeh, M.; Varol, O.; Yang, K.-C.; Flammini, A.; and Menczer, F. 2020. Detection of novel social bots by ensembles of specialized classifiers. In *Proceedings of the 29th ACM international conference on information & knowledge management*, 2725–2732.
- Schwartz, H.; Eichstaedt, J.; Kern, M.; Dziurzynski, L.; Lucas, R.; Agrawal, M.; Park, G.; Lakshminanth, S.; Jha, S.; Seligman, M.; et al. 2013. Characterizing geographic variation in well-being using tweets. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7, 583–591.
- Shao, C.; Ciampaglia, G. L.; Varol, O.; Yang, K.-C.; Flammini, A.; and Menczer, F. 2018. The spread of low-credibility content by social bots. *Nature communications*, 9(1): 1–9.
- Tanash, R. S.; Chen, Z.; Thakur, T.; Wallach, D. S.; and Subramanian, D. 2015. Known unknowns: An analysis of Twitter censorship in Turkey. In *Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society*, 11–20.
- Tufekci, Z. 2014. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *Eighth international AAAI conference on weblogs and social media*.
- Uyheng, J.; and Carley, K. M. 2021. Computational Analysis of Bot Activity in the Asia-Pacific: A Comparative Study of Four National Elections. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 727–738.
- Van Dijck, J.; Poell, T.; and De Waal, M. 2018. *The platform society: Public values in a connective world*. Oxford University Press.
- van Geenen, D.; Schaefer, M. T.; Boeschoten, T.; Hekman, E.; Bakker, P.; and Moons, J. 2016. Mining One Week of Twitter. Mapping Networked Publics in the Dutch Twittersphere.
- Varol, O. 2016. Spatiotemporal analysis of censored content on twitter. In *Proceedings of the 8th ACM Conference on Web Science*, 372–373.
- Varol, O. 2022. Should we agree to disagree about Twitter’s bot problem? *arXiv preprint arXiv:2209.10006*.
- Varol, O.; Ferrara, E.; Davis, C.; Menczer, F.; and Flammini, A. 2017. Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the international AAAI conference on web and social media*, volume 11, 280–289.
- Varol, O.; and Uluturk, I. 2020. Journalists on Twitter: self-branding, audiences, and involvement of bots. *Journal of Computational Social Science*, 3(1): 83–101.
- Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true and false news online. *science*, 359(6380): 1146–1151.
- Wilkinson, M. D.; Dumontier, M.; Aalbersberg, I. J.; Appleton, G.; Axton, M.; Baak, A.; Blomberg, N.; Boiten, J.-W.; da Silva Santos, L. B.; Bourne, P. E.; et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3(1): 1–9.
- Wojcieszak, M.; Casas, A.; Yu, X.; Nagler, J.; and Tucker, J. A. 2022. Most users do not follow political elites on Twitter; those who do show overwhelming preferences for ideological congruity. *Science advances*, 8(39): eabn9418.
- Wu, S.; and Dredze, M. 2020. Are All Languages Created Equal in Multilingual BERT? In *Proceedings of the 5th Workshop on Representation Learning for NLP*, 120–130.
- Wu, S.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Who says what to whom on twitter. In *Proceedings of the 20th international conference on World wide web*, 705–714.
- Yang, K.-C.; Varol, O.; Hui, P.-M.; and Menczer, F. 2020. Scalable and generalizable social bot detection through data selection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 1096–1103.
- Zagheni, E.; Garimella, V. R. K.; Weber, I.; and State, B. 2014. Inferring international and internal migration patterns from twitter data. In *Proceedings of the 23rd international conference on world wide web*, 439–444.
- Zhang, X.; Malkov, Y.; Florez, O.; Park, S.; McWilliams, B.; Han, J.; and El-Kishky, A. 2022. TwHIN-BERT: A Socially-Enriched Pre-trained Language Model for Multilingual Tweet Representations. *arXiv preprint arXiv:2209.07562*.
- Zhou, A.; and Yang, A. 2021. The Longitudinal Dimension of Social-Mediated Movements: Hidden Brokerage and the Unsung Tales of Movement Spilloverers. *Social Media+ Society*, 7(3): 20563051211047545.
- Zimmer, M. 2020. “But the data is already public”: on the ethics of research in Facebook. In *The Ethics of Information Technologies*, 229–241. Routledge.