

RESEARCH ARTICLE

Local information sources received the most attention from Puerto Ricans during the aftermath of Hurricane Maria

Benjamin Freixas Emery^{1*}, Meredith T. Niles², Christopher M. Danforth¹, Peter Sheridan Dodds¹

1 Computational Story Lab, Vermont Complex Systems Center, The Vermont Advanced Computing Core, Department of Mathematics & Statistics, The University of Vermont, Burlington, VT, United States of America, **2** Department of Nutrition and Food Sciences, The University of Vermont, Burlington, VT, United States of America

* bfemery@uvm.edu



OPEN ACCESS

Citation: Emery BF, Niles MT, Danforth CM, Dodds PS (2021) Local information sources received the most attention from Puerto Ricans during the aftermath of Hurricane Maria. PLoS ONE 16(6): e0251704. <https://doi.org/10.1371/journal.pone.0251704>

Editor: Naoki Masuda, University at Buffalo, UNITED STATES

Received: July 30, 2020

Accepted: April 30, 2021

Published: June 9, 2021

Copyright: © 2021 Emery et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All tweet data is available in a public repository at <https://github.com/dbemerydt/maria-twitter-data>.

Funding: CMD and PSD were supported in part by a gift from MassMutual Insurance. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Abstract

In September 2017, Hurricane Maria made landfall across the Caribbean region as a category 4 storm. In the aftermath, many residents of Puerto Rico were without power or clean running water for nearly a year. Using both English and Spanish tweets from September 16 to October 15 2017, we investigate discussion of Maria both on and off the island, constructing a proxy for the temporal network of communication between victims of the hurricane and others. We use information theoretic tools to compare the lexical divergence of different subgroups within the network. Lastly, we quantify temporal changes in user prominence throughout the event. We find at the global level that Spanish tweets more often contained messages of hope and a focus on those helping. At the local level, we find that information propagating among Puerto Ricans most often originated from sources local to the island, such as journalists and politicians. Critically, content from these accounts overshadows content from celebrities, global news networks, and the like for the large majority of the time period studied. Our findings reveal insight into ways social media campaigns could be deployed to disseminate relief information during similar events in the future.

1 Introduction

1.1 Background

In the early morning of September 20th 2017, Puerto Rico was struck by its first Category 4 hurricane in 85 years [1]. Hurricane Maria made landfall with wind speeds of 155 miles per hour, just 2 miles per hour short of qualifying as a Category 5 hurricane. The island's electrical grid, which had been left in highly vulnerable condition as a result of the territory's extreme debt [2], was entirely wiped out. Most of the island was left without power for nearly two months, and some areas remained without power for nearly a year [3].

The Puerto Rican government accepts an independent study approximating that 2,975 deaths resulted from the storm and its aftermath [4], although estimates from other

researchers place 8,000 excess deaths within the range of plausible realities [5]. An internal report at the Federal Emergency Management Agency (FEMA) acknowledges a failure to adequately prepare for the 2017 hurricane season, resulting in the agency being unable to properly support the victims of Hurricane Maria [6]. Failures at the local level also played a role in the resulting damage. In July 2019, two Puerto Rican officials were arrested on corruption charges for allocating \$15.5 million in relief contracts to favored businesses, despite credible accusations of being unqualified [7]. These arrests occurred among a series of events that brought hundreds of thousands of Puerto Ricans to protest in the streets for two weeks in demonstrations that eventually led to the resignation of Puerto Rican governor Ricardo Roselló [8].

Despite 89% of cell phone towers being offline [9], many Puerto Ricans were connecting with each other in the first few days after the storm. Individuals and families were able to find isolated hot spots of cellphone coverage, and despite only being able to communicate with one another in person or with battery-operated radios, knowledge of these spots spread rapidly. During the first week after the storm, these spots could be found with dozens of cars pulled over, their passengers outside holding smartphones in the air [10].

At the time, Hurricane Maria was the latest in a sequence of natural crisis events spanning nearly a decade in which social media was used for practical communication. Despite causing the highest death toll in 20 years and the largest cumulative blackout in United States history [11], Hurricane Maria received little attention from the mainstream media [12]. Our goal in the present work is to retrospectively investigate and reveal specific details about the nature of public discussions about Hurricane Maria using tools from network science and information theory.

1.2 Existing research

Work surrounding previous major natural hazard events in the contiguous United States has found that the temporal distribution of disaster-related tweeting varies between disaster types, and for hurricanes, the greatest volume of tweeting tends to happen in the anticipation of the event [13]. Additionally, users who demonstrate the greatest increase in activity during the timeline of such an event tend to be those with an average-sized following, rather than highly connected individuals. There has also been work showing that people are largely motivated to spread information on Twitter during disaster events by the belief that the information is important and believable [14]. Researchers have also found that during hurricanes Twitter can be used to predict how to best focus recovery efforts [15], and that Twitter activity correlates with hurricane damage [16]. Algorithms have been developed to identify flood victims asking for help on Twitter [17]. Going beyond publicly available social media, research have found that private SMS messaging groups have been integral to the spread of information among small communities during hurricanes [18].

Twitter in particular lends itself well to a network framework for analysis. Although Twitter does not make follower and following lists easily accessible, public messages between accounts can be easily identified. Research into protest movements has used such interactions between accounts to specify a network representation [19–21]. Other work has analyzed the network topology of hashtag topics and used information theory to identify the differences in discourse between hashtag topics [22]. The use of network analysis on Twitter surrounding disasters is not new. For example, after the Deepwater Horizon oil spill, Twitter networks were used to study the online conversation as the recovery efforts unfolded [23].

1.3 Focus of study

We wish to understand the online conversation on Twitter surrounding Hurricane Maria, beginning more broadly and then narrowing down to those directly affected. We begin by

looking specifically at the differences between English and Spanish tweets. This analysis serves as broader context with which to view our subsequent more granular analysis.

Following our more general analysis, we narrow our focus to those on the island or with close connections to people on the island. We seek to characterize the flow of information on Twitter among people in Puerto Rico during Hurricane Maria. From what types of accounts does most information originate? How large is the audience for these most central accounts? Do different subsets of the information flow network have major differences in the content being spread? The answers to these questions may have potential impact on communication strategies during major natural hazard events.

2 Materials and methods

2.1 Data collection

2.1.1 Between-language analysis. We used Twitter’s Decahose sample, a paid-access API providing a pseudo-random 10% sample of all public tweets. We sought to collect the conversations on Twitter surrounding the hurricane and Puerto Rico around the time that the island was hit. Over the course of 30 days, from September 16 to October 15 in 2017 (inclusive), we collected all tweets that contained the strings “puerto”, “rico”, “hurricane”, “maria”, “huracan”, “maria”, or “huracán”. We expected a collection of tweets related to the hurricane, but acknowledge that the collection inevitably contains some unrelated tweets. We used Twitter’s provided language tag to bin the Tweets into Spanish and English corpora. In either corpus, it’s possible that we inadvertently collected writing in other languages, due to potential multi-lingual tweets. Such cases are ignored for the most part, as words outside of our corpus of all words in the designated language will not be considered in analysis. We record the author of each tweet, it’s unique tweet identification number, the date and time of it’s publication, and the content of the tweet. Our resulting collection contains 1,947,489 Spanish tweets and 3,030,286 English tweets. We show these volumes of tweets over time in Fig 1.

2.1.2 Hashtag analysis. Using the same data source, we gathered all Decahose tweets containing at least two hashtags and constructed a list of co-occurring hashtag pairs on each day: a hashtag topic network. We then extracted the subset of the network that only includes hashtags co-occurring with #PuertoRico. This network subset is also known as the “ego network” of #PuertoRico. The resulting timeline of networks contains 11,633 tweets and 7,392 unique hashtags, allowing for any language.

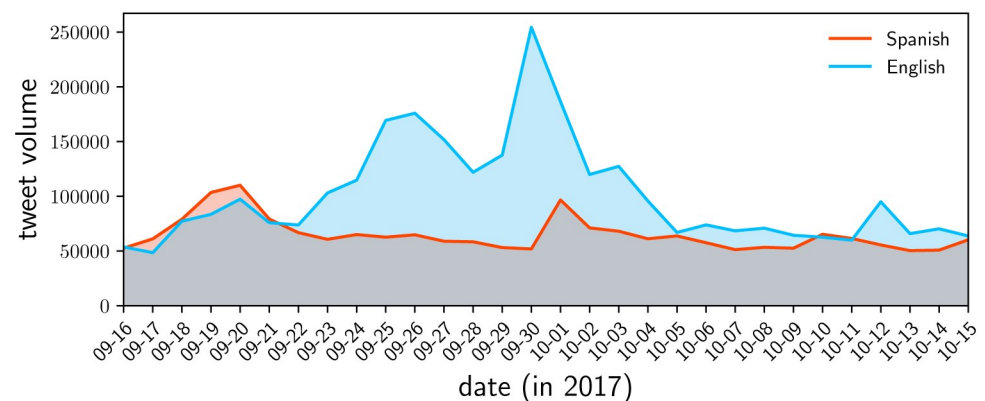


Fig 1. The daily volume of Spanish and English tweets concerning the hurricane or Puerto Rico in our collected subset of Twitter’s Decahose sample, roughly 10% of all public tweets.

<https://doi.org/10.1371/journal.pone.0251704.g001>

2.1.3 Victim proxy. Our goal here was to capture a set of tweets that could serve as a proxy for the conversation among residents of Puerto Rico and people with personal connections to those residents. We began by finding the user names for accounts located on mainland Puerto Rico, or any of the smaller islands of Vieques, Culebra, or Mona during the hurricane. Looking through all geolocated tweets, we saved account names for those who authored tweets from the island on September 20 or 21, 2017. A total of 93 seed users sent geolocated tweets from Puerto Rico during the hurricane. Using Twitter’s free Premium API, we searched for all tweets that mentioned or retweeted a user from this seed list during the date range specified in Section 2.1.1. This extended our list of account names to 1,511 users, hereafter our “user sample”.

We chose our time interval to maximize preparatory time while avoiding overlap with the recovery from Hurricane Irma, which had hit only days before. This choice of preparatory time interval is also consistent with findings by Niles et al., which demonstrate it to be valid for showing tweet increases during a climate emergency [13]. We extended the studied interval for a month, in order to look at the change in communication behavior throughout the development of the situation and the revelations of new information. This allowed us to gather additional users who were on the island but did not enable geolocation, as well as people elsewhere who are personally connected to people on the island during the hurricane.

With our user sample constructed, we recorded every Decahose tweet that (1) was authored by an account on our list, (2) was marked by Twitter as a Spanish tweet, and (3) either mentioned or retweeted another user. We also restricted our corpus to Spanish tweets, as information that’s truly targeted at the Puerto Rican populations will be in Spanish, the island’s primary language. From our list of tweets, we kept only those using hand-selected vocabulary that indicated a probable discussion of the hurricane’s effects or need for supplies, shown in Table 1. These words are adapted from the list of words mentioned in Section 1.2, with the addition of words “puerto rico” and “huracán” and “maria”.

2.2 Hedonometric analysis

For their 2015 work, Dodds et al. collected “happiness” ratings of 10,000 words across ten languages [24]. Words were rated by survey participants on a 9-point, double-Likert scale from 1 (most negative) to 9 (most positive), with 5 representing a neutral word. The resulting collection of word lists and continuous sentiment scores is known as the language assessment by Mechanical Turk or “LabMT” dataset. For these analyses, we assign words their score identified by Dodds et al. In order to improve our ease of analysis, we subtract 5 from the score of

Table 1. Words used to flag tweets in our study (top) and their English translations (bottom).

Spanish	agua	bebida	generador	maria	suministros
	albergue	comida	huracan	maría	supermercado
	alimento	corriente	huracán	prepar	tienda
	asistencia	desprevenido	inund	puerto rico	toldo
	ayuda	en conserva	lluv	refrigera	viento
	banco de alimentos	enlatado	luz	refugio	
English	water	drink	generator	maria	supplies
	hostel	food	hurricane	maria	supermarket
	food	current	hurricane	prepare	store
	assistance	unprepared	flood	puerto rico	tarp
	help	nonperishable	rain	refrigerat	wind
	foodbank	canned	power	refuge	

<https://doi.org/10.1371/journal.pone.0251704.t001>

each word, giving unhappy words a negative score. With this scoring scheme negative words in Spanish like “odiar” (“to hate”) and “guerra” (“war”) respectively have scores of -2.92 and -3.22, where as positive words like “amar” (“to love”) and “beso” (“kiss”) have scores of 3.38 and 3.36. Similarly for English, “joy” has a score of 3.16 and “terror” has a score of -3.24. We compute the sentiment of either language for some time period as

$$\langle h \rangle = \sum_i h_i p_i, \tag{1}$$

where h_i is the sentiment score of a particular word, and p_i is that word’s relative frequency.

For meta-analysis of the happiness of a corpus, we wish to remove words with scores within some radius of zero (neutral). Words with scores close to neutral are removed from consideration to increase the strength of the emotional signal recovered from the corpus. These masked words are largely comprised of either (a) emotionally neutral words or (b) context dependent. Attempts at sentiment analysis prior to the creation of LabMT involved the removal of hand-selected “function” words, but it has been found that many words that individuals might label as “function” have a substantial bias in sentiment [24, 25]. One can therefore deduce that it’s more accurate to remove words that have been empirically determined to have near-neutral net sentiment [26].

To determine the appropriate width for this sentiment mask, we sweep through values and compute the mean and variance of the sentiment scores of all tweets for English and Spanish from Section 2.1 (Fig 2). While the mean for both languages monotonically decreases with increasing neutral mask, the variance has a plateau from 0 to 1, after which it drops. For this reason, we choose a lens radius of 1, which is also consistent with previous work using the labMT dataset [26, 27–31].

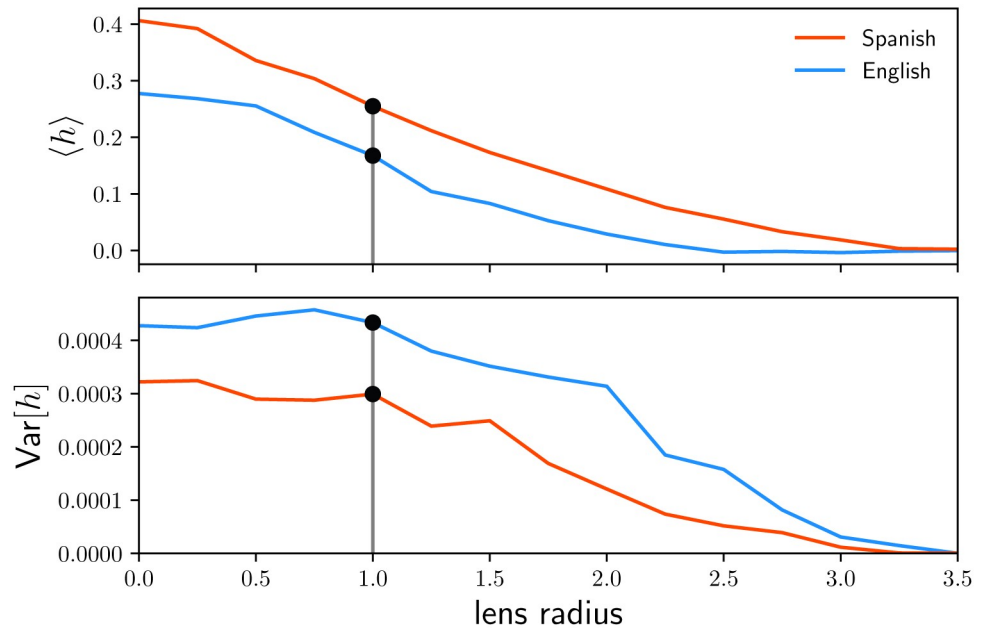


Fig 2. The mean and variance of daily sentiments at each lens radius value for English and Spanish. At all lens values, Spanish has a higher mean happiness score than English, but a smaller spread. We choose out lens size to remove a sufficient amount of neutral words, which in mass can skew the scores and make results less intuitive. A gray vertical line in each plot indicates the selected lens radius. Black dots show the values of the mean and variance of the daily sentiments at the selected lens radius. This value is selected because it is roughly the edge of the variance’s plateau for the two languages (see [26] for further support).

<https://doi.org/10.1371/journal.pone.0251704.g002>

2.3 Network specification

2.3.1 Hashtags. We generate an aggregate network by superimposing the daily ego networks described in Section 2.1.2. We remove the ego node, #PuertoRico, which results in many nodes and small components of nodes becoming disconnected from the largest component of the graph. We remove all components apart from the largest connected component, and conduct the remainder of the analysis on this component of the network. Using the Louvain clustering algorithm, we identify five communities within the network structure [32].

2.3.2 Users. In order to interpret the flow of information among the individuals of interest, and identify the distribution of influence, we generate a network of Twitter users connected by retweets and mentions. Each tweet pulled from the Decahose as described in Section 2.1.3 constitutes a directed link from the author to the mentioned or retweeted user.

We identify communities in the network topology with the Louvain clustering algorithm. Inspection of the network's tweet content revealed that the second largest community consisted almost entirely of people involved in the 2017 Venezuelan protests, and several smaller communities shared this property [33]. To modify the network to better represent the focus of the present study, we counted the proportion of tweets in the largest Venezuelan network that contained the words "venezuela", "vzla", or "maduro", and computed the measurement's Wilson interval lower bound at 99% confidence [34]. We then removed all communities with a proportion of tweets containing those same words at least as high as that lower bound.

The resulting network has 2,011 users (nodes) connected by 2,466 tweets (edges). Although our links are generated from a 10% sample of tweets and we reduce our corpus by content type, we maintain a third of the nodes from our original user list due to the heavy-tailed nature of online social networks [35]. Finally, we hand-identified the fifty largest nodes by in-degree as being news outlets, politicians, citizens, weather stations, meteorologists, or journalists.

We consider three major attributes of the aggregate network: the in-degree distribution, density, and the average shortest path length. The in-degree of a node is the number of links connecting other nodes to it (tweets mentioning or retweeting the corresponding user). We define the density of the network as

$$d = \frac{m}{n(n-1)}, \quad (2)$$

where n is the number of nodes and m is the number of directed links in the network, each counted once regardless of weight. This is in essence the fraction of possible locations for a link where there exists one. We compute the average shortest pathlength by iterating over each ordered pair of nodes s and t and finding the minimum number of links defining a path from s to t . From these pathlengths, we compute the average shortest pathlength as

$$a = \sum_{s,t \in V} \frac{p(s,t)}{n(n-1)}, \quad (3)$$

where V is the set of nodes, and $p(s,t)$ is length of the shortest path from node s to node t .

2.4 Shannon entropy

We rely on information theoretic measurements in our analysis of tweet content. The most basic measurement we use is Shannon's entropy H , which describes the diversity of the probability mass function of nominal variables. Shannon's entropy of a corpus with n unique words

is defined as

$$H = -\sum_{i=1}^n p_i \log_2 p_i, \quad (4)$$

where the i th word appears with probability (or rate) p_i . Shannon entropy is maximized when every word occurs with equal probability, and approaches 0 as the corpus becomes dominated by one word.

2.5 Jensen-Shannon divergence

To compare word distributions for two subsets of our corpus, we compute the Jensen-Shannon divergence between them.

If we consider any two corpora, T_1 and T_2 , with word frequency distributions P_1 and P_2 over k possible unique words, and their combined distribution $M = \pi_1 P_1 + \pi_2 P_2$, where π_1 and π_2 are the relative sizes of the two corpora such that $\pi_1 + \pi_2 = 1$, the Jensen-Shannon Divergence D is given by

$$D(P_1||P_2) = \pi_1 \sum_{i=1}^k p_{1,i} \log_2 \frac{p_{1,i}}{m_i} + \pi_2 \sum_{i=1}^k p_{2,i} \log_2 \frac{p_{2,i}}{m_i}. \quad (5)$$

This expression is the weighted average of each corpus's Kullback-Leibler divergence from the mixed corpus.

We use the linearity of Jensen Shannon Divergence extensively when comparing corpora in order to identify the words and topics that contribute most to their difference.

3 Results

3.1 Daily volume and happiness

We show the time series for daily tweet volume in each language in Fig 1. Both the Spanish and English tweet volume rises in anticipation of the hurricane, reaching a maximum on the date of landfall. The English volume begins steadily rising again two days later, eventually reaching another peak on September 26th. A sudden spike appears in the English tweet volume on September 30, following a string of tweets from Twitter user @realDonaldTrump, blaming the mayor of San Juan for poor leadership. After this spike, English tweet volume decays to a steady 50,000 per day. Spanish Tweet volume sees a much smaller but still prominent spike on October 1.

In Fig 3, we show the daily average happiness of Spanish and English tweets, as well as their difference on each day. Around the date of landfall, English sentiment remains unchanged, while Spanish sentiment spikes upward. We also see downward spikes in the sentiments of both languages on October 1. English also sees spikes on October 2 and October 13 that are unmatched in the Spanish corpus.

3.2 Between-language wordshifts

From here, we examine the contribution of individual words to the difference in average happiness between the two languages over periods in time. For this we use a visualization tool called the "wordshift". Wordshifts for different dates and date ranges are shown in Fig 4. These wordshifts are distinct from wordshifts used in previous work, because the corpora being compared must be treated as having a completely unique sets of words. Because of this, there is no sense in comparing whether or not a certain word is less or more frequent in a certain corpus.

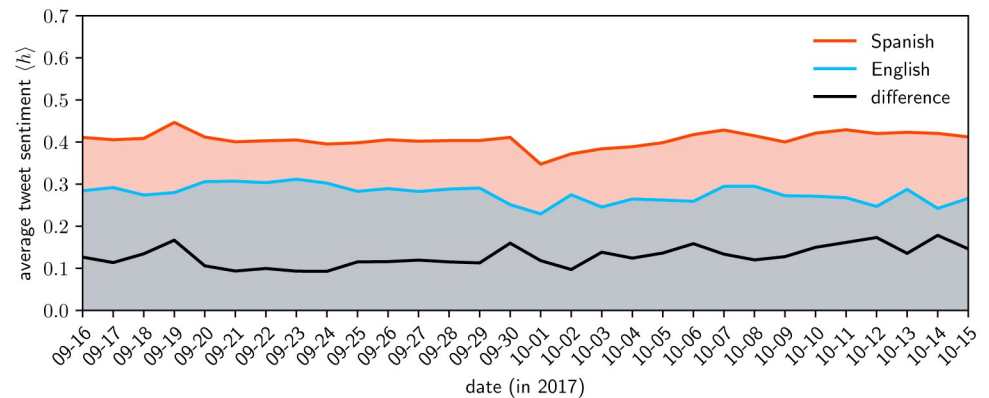


Fig 3. The daily average happiness of Spanish and English tweets in our collected subset of Twitter’s Decahose sample. The difference between the average happiness for each day is shown by the solid black line. Spanish tweets are more positive than English tweets throughout the time period, likely due to the underlying word scores. In previous work studying 10 major languages Dodds et al. found Spanish to be the language containing the highest proportion of happy words [24].

<https://doi.org/10.1371/journal.pone.0251704.g003>

Instead, we look at which words in each corpus contribute the most to bringing the happiness of that corpus up or down. Words on the left of the shift are thus making Spanish look like the happier language, and words on the right make English appear happier. We provide more information on this computation in the caption of Fig 4.

In Fig 4, we present the sentiment wordshift between Spanish and English for the entire studied time interval. We find that both sides are mostly occupied by “happy” words of each language, and the “sad” words that contribute are almost all negations, such as the Spanish “no” and the English “no”, “not”, and “don’t”, all of which translate to “no” in Spanish. Other prominent negative words in the wordshift are “without” and “lost”.

It is also notable that most of the “happy” words contributing to the wordshift are likely being used in a negative context. For example, “power”, “relief”, and “help” all in part refer to services the island has lost or needed. The word “mexico” appears prominently for two reasons. First, the LabMT word list was evaluated by Mexican Spanish speakers, so it is reasonable to suspect “mexico” to receive a higher score than if it had been evaluated by a more globally representative sample of Spanish speakers. Second, an earthquake occurred in Mexico around the same time, prompting people of the Latinx community to talk about the disasters happening to two major geographic parts of Latin America.

Fig 5 highlights the words contributing to the difference in happiness between English and Spanish tweets on 2017–09–19. We choose this date because it is the day before Hurricane Maria made landfall, and because we observe a spike in the sentiment difference between Spanish and English on this date. The words that appear indicate emphasis on the earthquake in Mexico. Some of the words accounting for the upward spike in tweet sentiment at on that day are words such as “héroes” (“heroes”) and “dios” (“God”), which focus on positive aspects of the events.

3.3 Hashtag network

The network of hashtags for the studied date range consists mostly of small nodes with in-degrees in the single-digits (see Fig 6). A small handful of larger nodes dominate, which exist in each of the detected communities. Notable in the network are two closely connected communities containing content about U.S. President Donald Trump. One community (indicated with black nodes) contains slogans attacking Trump such as #Resist, #Resistance, and




Fig 4. Sentiment wordshift for the entire studied time span, 2017–09–16 through 2017–10–15. Each bar indicates the contribution of a word to the difference between the average happiness of Spanish tweets vs English tweets. The length of the bar is the product of the word’s relative sentiment score and its relative frequency. We consider consider words to be “happy” if they have a LabMT score greater than 5 (positive adjusted score), and “sad” if their LabMT score is less than 5 (negative adjusted score). Bars that extend leftward are either happy Spanish words (deep blue) or sad English words (muted yellow), and bars that extend to the right are either happy English words (bright yellow) or sad Spanish words (muted blue).

<https://doi.org/10.1371/journal.pone.0251704.g004>

#ImpeachTrump, as well #AMJoy, a hashtag referring to MSNBC reporter Joy Reid’s news show (Reid visited Puerto Rico on October 14th). The other community contains topics in support of the president, such as #Trump, #MAGA, and #FakeNews, but also contains anti-Trump topics, indicating that some hijacking of those pro-Trump hashtags took place. A third

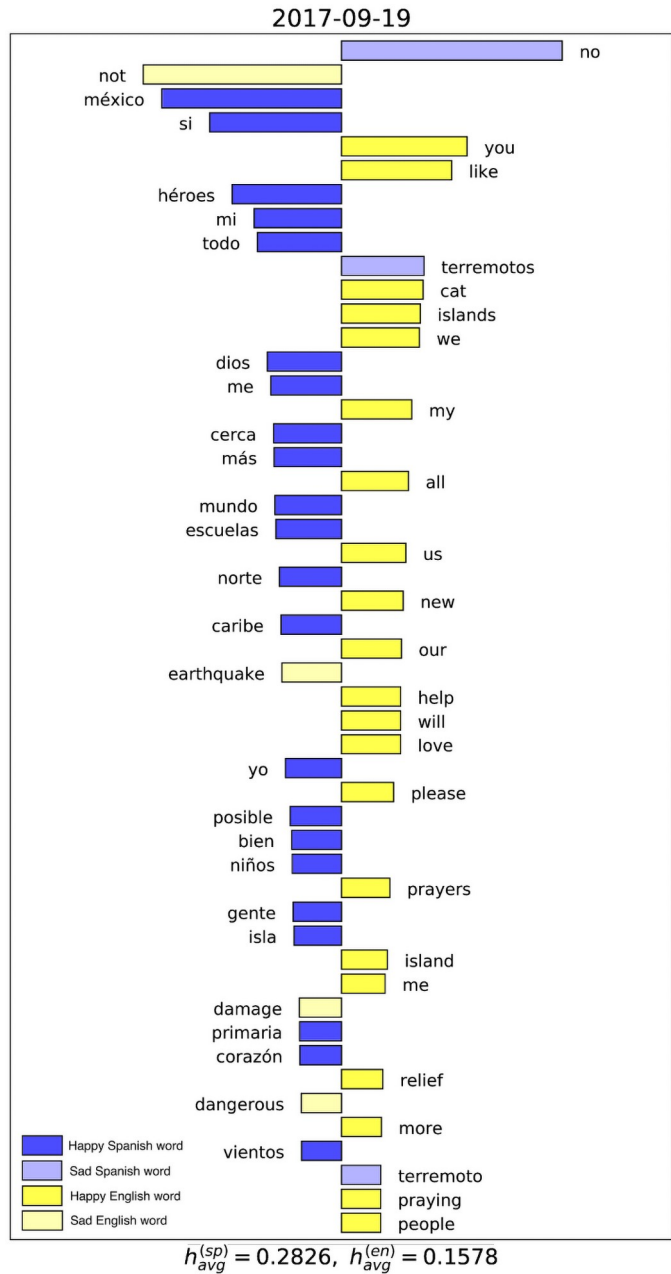


Fig 5. Sentiment wordshift for the individual day on 2017–09–19. See Fig 4 for details on the methodology.

<https://doi.org/10.1371/journal.pone.0251704.g005>

community is dominated by the hashtags #SomosUnaVoz and #OneVoice, both promotional hashtags for a live televised charity concert for disaster relief, which occurred on October 14th as well.

Two further communities, denoted with blue and red nodes, contain hashtags indicating discussion about the hurricane and other recent disaster events, as well as their respective relief efforts.

As a whole the hashtag network demonstrates a large variety of conversation topics, but a generally disproportionate emphasis on Donald Trump that dwarfs conversation about the storm, relief efforts, and other directly related topics.




Fig 7. The in-degree distribution of the communication network (shown in Fig 8(a)), plotted on a log-log scale.

<https://doi.org/10.1371/journal.pone.0251704.g007>

network is very sparse, the expected distance between two nodes is small. This is consistent with the heavy-tailed nature of the degree distribution we see in Fig 7.

The anticipation network in Fig 8(b) has 234 nodes and 261 links, and centers mainly around two accounts: A meteorologist from Florida and a Puerto Rican weather station. Apart from this, the network is made up of many smaller separate connected components between relatively small nodes. The aftermath network has 1,347 nodes and 1,874 links. This subset of the aggregate network strongly resembles the entire aggregate, with the collection of musicians having more prominence, while maintaining a large community of news outlets and weather stations.

3.5 Divergence

We use Jensen-Shannon divergence wordshifts in Figs 9 and 10 to demonstrate differences in subsets of our collection of tweets. The length of a bar corresponds to that word's contributions to the Jensen-Shannon divergence between the two corpora, a measurement of the difference in the two word distributions. The bars are colored according to the Shannon entropy of the collection of tweets that use the word, meaning that darker colors indicate a higher diversity of tweet content and colors closer to white indicate more homogeneous content. Bars that are very close to white indicate that the subcollection of tweets using the word are mostly retweets of one specific tweet. In each wordshift, the fifty words with the highest contribution are shown. Emojis are each considered to be one individual word for this analysis.

In Fig 9, we show the words contributing to the difference between the tweets in anticipation of the hurricane and the tweets from the aftermath. The large majority of the divergence between these two corpora comes from specific words used more in anticipation. Most of these have to do with the severity of the imminent storm. An increase in the frequency of "agua" after the hurricane contributes to this divergence, as well as "#verificando19s", a hashtag used in Mexico to organize a rescue effort in the aftermath of an earthquake on September 19.

The corpora compared in Fig 10(a)'s wordshift are the tweets from the central community containing Governor Rossello's account to those from the fourth-largest community, which

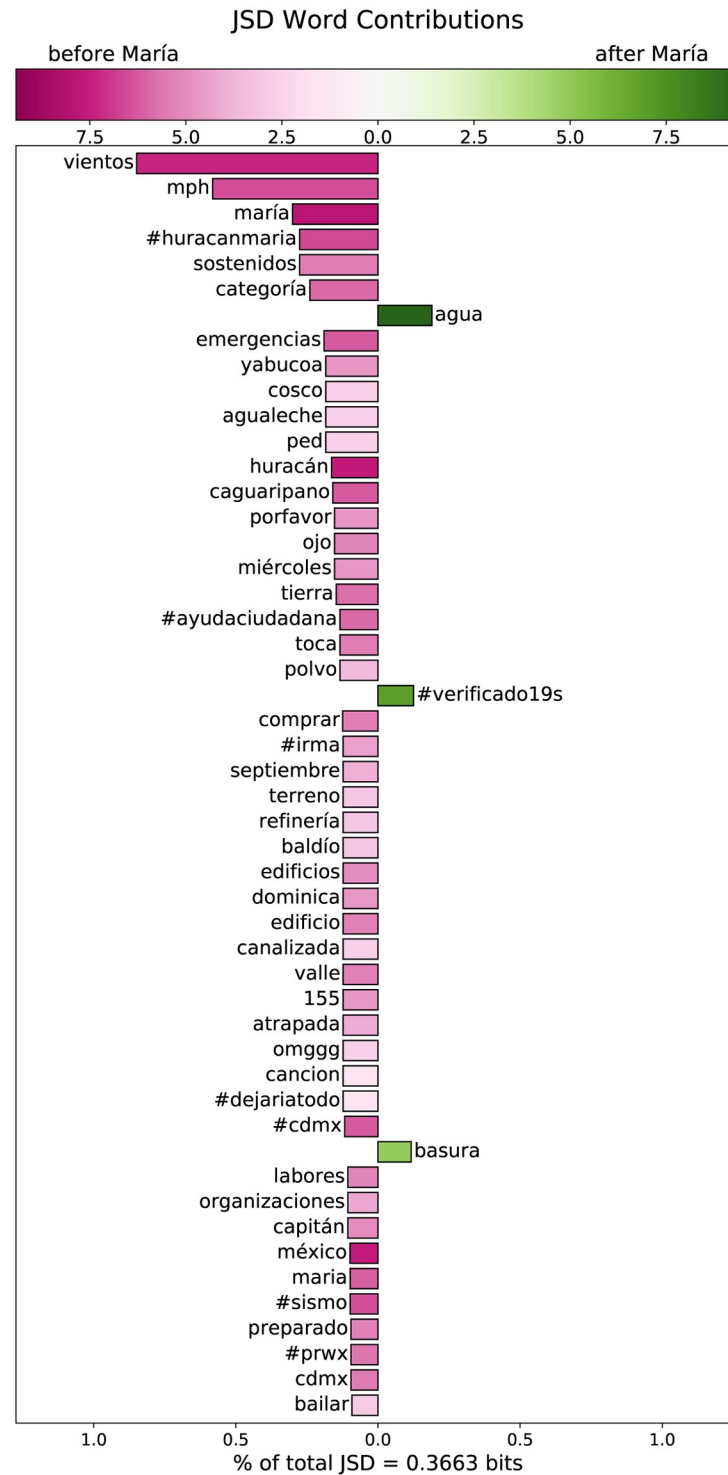


Fig 9. The contributions of individual words to the Jensen-Shannon divergence between the collection of tweets before and after the period of landfall for Hurricane Maria. The length of each bar indicates the contribution of a word, and its direction indicates the corpus in which it was more frequent. The bars are colored according to the Shannon entropy of the Tweets that contained that word, hashtag, or emoji.

<https://doi.org/10.1371/journal.pone.0251704.g009>




Fig 10. (a) The word-level contribution to the Jensen-Shannon divergence (JSD) between the collection of tweets in the largest and fourth-largest communities. (b) The word-level contribution to the JSD between the collection of tweets in the largest and second-largest communities.

<https://doi.org/10.1371/journal.pone.0251704.g010>

3.6 Daily resolution network evolution

We separate the aggregate network into its daily subnetworks, and look at the top fifty users in the aggregate network at the daily resolution. Fig 11 illustrates the change of in-degree over time of these nodes in the daily networks, binned by user type. We see that news outlets are a relatively consistent player in the conversation, whereas weather stations and meteorologists were only prominent before the hurricane and on the first day of October. Government and other aid organizations tend to take up a small portion of the conversation, becoming major players somewhat suddenly on September 24 and October 8–9.

The same figure with every individual user colored separately can be found in the Appendix. In examining individual users with time, we find that in the days surrounding the landfall of the hurricane, the tweets from the meteorology station @NWSSanJuan are heavily propagated. A few days after the end of the rain, Governor Rossello becomes a prominent voice in the network on most days from then on. During the first few days of October, Puerto Rican singers Chayanne, Ricky Martin, and Luis Fonsi occupy much of the conversation space. Later on, starting on October 10, Puerto Rican rapper Residente becomes a major node in the




Fig 11. Stacked bar chart time-series showing the daily in-degrees of nodes in the top-fifty by in-degree of the aggregate network. The nodes are consolidated by manually determined type, shown in the legend. Of these nodes, the news outlets are the most consistently prominent, although politicians and journalists are fairly consistently present in the network as well. We find a spike in activity from musicians starting on October first, and decaying away by the sixth. Musician activity comes back for an interval of four days starting on October 9th, but this is all due to one single account, Residente. This can be seen in Figure S1 Fig, where we show the same timeseries separated by individual accounts.

<https://doi.org/10.1371/journal.pone.0251704.g011>

network, but he appears to mostly connect with smaller nodes, and not as much with other major figures.

4 Discussion

4.1 Global conversations

The happiness levels of each language within the body of tweets about the hurricane is very stable, especially when compared to the same window of time for the collection of all English tweets (<http://hedonometer.org>). There are a handful of incidents where the sentiment trajectories of the two languages track one another, as well as a set of instances where they spike independently. This can be reconciled with the varying scales of impact that different events may have.

An independent spike in Spanish sentiment on September 19 is an intriguing event, as it conveys an important realization about how Spanish speakers spoke about the two major natural disaster events that were happening or were imminent. It can be seen from the wordshift from that day that Spanish speakers are discussing the earthquake in Mexico along with the incoming hurricane. The influential words, however, seem to indicate an emphasis on hope and the inspiring action of helpers. Words that stand out here are “héroes”, “dios”, and “niños”. It appears that there is a focus among Spanish tweeters on religious figures and those demonstrating heroism in the wake of the disaster, including the young boys who were delivering supplies on mountain bikes to the first responders [36].

Our findings point to other methods of investigating the true sources of difference and similarity in the discourse between communities who speak different languages. For instance, one might apply wordshift methods to investigate dips by comparing them to the corpora from previous days. One might also wish to examine information theoretic trends in each of the languages, such as normalized entropy.

4.2 Local conversations

In a previous work we have found that, in response to a major natural crisis event, average users with small followings increased their social media activity more than influential accounts with large followings [13]. While it is challenging to identify those responsible for the propagation of information throughout the Twitter network, our work here begins to address the natural follow on question: Where is the information originating? The consistent dominance of local journalists, news outlets, and politicians in the communication network demonstrates the accounts people turned to were those whose audience is exclusively Puerto Rico on any given day of the year.

Our results suggest that crisis information is spread throughout the network by average individuals, the first of which get such information from “local” Twitter users, such as the politicians and journalists from Puerto Rico. Meanwhile, major celebrities seem mostly removed from a conversation centered on local actors and average affected individuals, touching the Puerto Rico crisis along with a variety of issues in Latin America.

4.3 Limitations, bias, and further research

There are several important limitations to acknowledge surrounding this work and its possible implications. Using Twitter to study communication comes with several inherent biases [37], including skewed representation of users in different racial, ethnic groups, or geographic regions. These biases are only further exacerbated by filtering tweets by key word as we have done. Additionally, the data Twitter provides via its API distorts the network measurements by artificially rewiring retweet chains into stars centered on the original author of the tweet. Network measurements are also altered by the use of only a 10% sample of tweets. Our methods of identifying tweets from Puerto Ricans and those with close connections surely introduced some error, despite being our best resort due to the very low frequency of geolocated tweets. Another compounding problem is that most people of interest had no cellular phone reception for the majority of the studied time interval. We note, however, that this means our network presents a lower bound on representation of regular Puerto Ricans and accounts with a mostly Puerto Rican audience. Our main finding, that the most consistent providers of information were those with more moderate localized audiences, is a conservative one. The true nature of the situation may very well be stronger.

The present study sets a stage for many potential experiments to expand the body of work in this area. Separating our communication network into separate networks for retweets and for mentions may uncover differences in the ways information spread, as these are fundamentally different types of edges [38]. Identifying the content of individual tweets in order to determine whether they convey information relevant to hurricane survivors looking for resources would be critical in revealing dynamics actionable for disaster relief. Lastly, sequential analyses of conversation threads would lead to interesting results for the evolution of topics over time, as has been studied in other situations [39].

5 Conclusion

Time and time again people prove incredibly resilient under conditions of extreme diversity. In the case of Hurricane Maria, Puerto Ricans took care of themselves and their families while facing a barrage of extreme weather, from flooding to heat waves. Further, they searched for ways to communicate with those outside their immediate presence, and found such ways on the sides of highways across the island.

The results we present hold implications for disaster relief, and information dissemination during major crisis events. For the most part we find that the best accounts to provide

important information tend to be local figures: Journalists, news outlets, politicians. While we did see points in time during which celebrities became central, they did not remain reliably central for any major span of time within the studied interval, which included anticipation, event, and aftermath. Meanwhile, content from global tweets indicated that a focus on hope and resilience from Spanish speaking users.

In any natural crisis situation, there is almost never a shortage of helpers. The emergence of massive cooperating teams during times of need is a signature of humanity. Lack of information, or explicit blockages that keep information from those who need it, however, can render this help not useful to the overwhelming masses of crisis victims. It is clear that the transfer of information about where and when to find help is as important as the help itself. We hope our results can inform institutions regarding how to use Twitter to best share information so that it most efficiently propagates through the network, getting to as many people as possible, as quickly as possible, and remaining uncorrupted in content.

Supporting information

S1 Fig.
(PDF)

Acknowledgments

The authors are grateful for support and helpful feedback from the members of the Computational Story Lab and the graduate and undergraduate advisees of Dr. Meredith Niles.

Author Contributions

Conceptualization: Benjamin Freixas Emery.

Data curation: Benjamin Freixas Emery.

Formal analysis: Benjamin Freixas Emery.

Investigation: Benjamin Freixas Emery.

Methodology: Benjamin Freixas Emery, Meredith T. Niles, Christopher M. Danforth.

Supervision: Meredith T. Niles, Christopher M. Danforth, Peter Sheridan Dodds.

Visualization: Benjamin Freixas Emery.

Writing – original draft: Benjamin Freixas Emery.

Writing – review & editing: Benjamin Freixas Emery, Meredith T. Niles, Christopher M. Danforth, Peter Sheridan Dodds.

References

1. Meyer R. What's Happening With the Relief Effort in Puerto Rico? A timeline of the unprecedented catastrophe of Hurricane Maria; 2017. Available from: <https://www.theatlantic.com/science/archive/2017/10/what-happened-in-puerto-rico-a-timeline-of-hurricane-maria/541956/>.
2. Allen G, Peñaloza M. Power Problems: Puerto Rico's Electric Utility Faces Crippling Debt; 2015. Available from: <https://www.npr.org/2015/05/07/403291009/power-problems-puerto-ricos-electric-utility-faces-crippling-debt>.
3. Sullivan E. Nearly A Year After Maria, Puerto Rico Officials Claim Power Is Totally Restored; 2018. Available from: <https://www.npr.org/2018/08/15/638739819/nearly-a-year-after-maria-puerto-rico-officials-claim-power-totally-restored>.
4. Andrade E, Barrett N, Colon-Ramos U, Edberg M, Garcia-Meza A, Goldman A, et al. Ascertainment of the estimated excess mortality from Hurricane María in Puerto Rico; 2018.

5. Kishore N, Marqués D, Mahmud A, Kiang MV, Rodríguez I, Fuller A, et al. Mortality in Puerto Rico after Hurricane Maria. *New England Journal of Medicine*. 2018; 379(2):162–170. <https://doi.org/10.1056/NEJMsa1803972> PMID: 29809109
6. Sullivan L. FEMA Report Acknowledges Failures In Puerto Rico Disaster Response; 2018. Available from: <https://www.npr.org/2018/07/13/628861808/fema-report-acknowledges-failures-in-puerto-rico-disaster-response>.
7. Allyn B. FBI Arrests Former Top Puerto Rico Officials In Government Corruption Scandal. National Public Radio. 2019.
8. Padró Ocasio B, Viglucci A. Gov. Ricardo Rosselló of Puerto Rico, under fire for weeks, agrees to step down; 2019. Available from: <https://www.miamiherald.com/news/nation-world/world/americas/article232949262.html?fbclid=IwAR1eCCZSaShj5XQ8eRod4mfgfSNAeWqp93CmOkwWDFKodDhrH31NgxR-BU>.
9. FitzGerald D, Hernandez D. A Week After Maria's Strike, Puerto Rico Remains Offline; 2017. Available from: https://www.wsj.com/articles/a-week-after-marias-strike-puerto-rico-remains-offline-1506714520?mod=article_inline.
10. Mazzei P. No cell service in Puerto Rico? Drive to a highway in San Juan and pull over; 2017. Available from: <https://www.miamiherald.com/news/weather/hurricane/article175156011.html>.
11. Irfan U. It's been more than 100 days and Puerto Rico is still in the longest blackout in US history; 2018. Available from: <https://www.vox.com/energy-and-environment/2017/10/30/16560212/puerto-rico-longest-blackout-in-us-history-hurricane-maria-grid-electricity>.
12. Chang A. How the media ignored Puerto Rico, in one chart. *Vox*. 2018.
13. Niles MT, Emery BF, Reagan AJ, Dodds PS, Danforth CM. Social media usage patterns during natural hazards. *PLOS ONE*. 2019; 14(2):1–16. <https://doi.org/10.1371/journal.pone.0210484> PMID: 30759111
14. Abdullah NA, Nishioka D, Tanaka Y, Murayama Y. Why I Retweet? Exploring User's Perspective on Decision-Making of Information Spreading during Disasters. In: *Proceedings of the 50th Hawaii International Conference on System Sciences*; 2017.
15. Guan X, Chen C. Using social media data to understand and assess disasters. *Natural hazards*. 2014; 74(2):837–850. <https://doi.org/10.1007/s11069-014-1217-1>
16. Kryvasheyev Y, Chen H, Obradovich N, Moro E, Van Hentenryck P, Fowler J, et al. Rapid assessment of disaster damage using social media activity. *Science advances*. 2016; 2(3):e1500779. <https://doi.org/10.1126/sciadv.1500779> PMID: 27034978
17. Singh JP, Dwivedi YK, Rana NP, Kumar A, Kapoor KK. Event classification and location prediction from tweets during disasters. *Annals of Operations Research*. 2017; p. 1–21.
18. Chu H, Yang JZ. Building Disaster Resilience with Social Messaging Networks: WeChat community in Houston during Hurricane Harvey. *Disasters*. 2019.
19. Jackson SJ, Foucault Welles B. Hijacking #myNYPD: Social media dissent and networked counterpublics. *Journal of Communication*. 2015; 65(6):932–952. <https://doi.org/10.1111/jcom.12185>
20. Jackson SJ, Foucault Welles B. #Ferguson is everywhere: Initiators in emerging counterpublic networks. *Information, Communication & Society*. 2016; 19(3):397–418. <https://doi.org/10.1080/1369118X.2015.1106571>
21. Bailey M, Jackson SJ, Foucault Welles B. Women Tweet on Violence: From #YesAllWomen to #MeToo. *Ada: A Journal of Gender, New Media, and Technology*. 2019;(15). <https://doi.org/10.5399/uo/ada.2019.15.6>
22. Gallagher RJ, Reagan AJ, Danforth CM, Dodds PS. Divergent discourse between protests and counter-protests: #BlackLivesMatter and #AllLivesMatter. *PLOS ONE*. 2018; 13(4):e0195644. <https://doi.org/10.1371/journal.pone.0195644> PMID: 29668754
23. Sutton J, Spiro E, Butts C, Fitzhugh S, Johnson B, Greczek M. Tweeting the spill: Online informal communications, social networks, and conversational microstructures during the Deepwater Horizon oil spill. *International Journal of Information Systems for Crisis Response and Management (IJISCRAM)*. 2013; 5(1):58–76. <https://doi.org/10.4018/jiscrm.2013010104>
24. Dodds PS, Clark EM, Desu S, Frank MR, Reagan AJ, Williams JR, et al. Human language reveals a universal positivity bias. *Proceedings of the National Academy of Sciences*. 2015; 112(8):2389–2394. <https://doi.org/10.1073/pnas.1411678112> PMID: 25675475
25. Garcia D, Garas A, Schweitzer F. The language-dependent relationship between word happiness and frequency. *Proceedings of the National Academy of Sciences*. 2015; 112(23):E2983–E2983. <https://doi.org/10.1073/pnas.1502909112> PMID: 25997447
26. Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PLOS ONE*. 2011; 6(12):e26752. <https://doi.org/10.1371/journal.pone.0026752> PMID: 22163266

27. Reece AG, Reagan AJ, Lix KL, Dodds PS, Danforth CM, Langer EJ. Forecasting the onset and course of mental illness with Twitter data. *Scientific reports*. 2017; 7(1):13006. <https://doi.org/10.1038/s41598-017-12961-9> PMID: 29021528
28. Cody EM, Reagan AJ, Dodds PS, Danforth CM. Public opinion polling with Twitter. arXiv preprint arXiv:160802024. 2016.
29. Reagan AJ, Mitchell L, Kiley D, Danforth CM, Dodds PS. The emotional arcs of stories are dominated by six basic shapes. *EPJ Data Science*. 2016; 5(1):31. <https://doi.org/10.1140/epjds/s13688-016-0093-1>
30. Frank MR, Mitchell L, Dodds PS, Danforth CM. Happiness and the patterns of life: A study of geolocated tweets. *Scientific reports*. 2013; 3:2625. <https://doi.org/10.1038/srep02625> PMID: 24026340
31. Mitchell L, Frank MR, Harris KD, Dodds PS, Danforth CM. The geography of happiness: Connecting Twitter sentiment and expression, demographics, and objective characteristics of place. *PLOS ONE*. 2013; 8(5):e64417. <https://doi.org/10.1371/journal.pone.0064417> PMID: 23734200
32. Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*. 2008; 2008(10):P10008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
33. Wikipedia contributors. 2017 Venezuelan protests—Wikipedia, The Free Encyclopedia; 2019. Available from: https://en.wikipedia.org/wiki/2017_Venezuelan_protests.
34. Wilson EB. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*. 1927; 22(158):209–212. <https://doi.org/10.1080/01621459.1927.10502953>
35. Kumar R, Novak J, Tomkins A. Structure and evolution of online social networks. In: *Link mining: models, algorithms, and applications*. Springer; 2010. p. 337–357.
36. Fountain N. Packs Of Teens On Bikes Join Volunteer Effort After Mexico Earthquake; 2017. National Public Radio. Available from: <https://www.npr.org/2017/09/23/553204495/packs-of-teens-on-bikes-join-volunteer-effort-after-mexico-earthquake>.
37. Olteanu A, Kiciman E, Castillo C. A critical review of online social data: Biases, methodological pitfalls, and ethical boundaries. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*; 2018. p. 785–786.
38. Cha M, Haddadi H, Benevenuto F, Gummadi K. Measuring user influence in twitter: The million follower fallacy. In: *Proceedings of the International AAAI Conference on Web and Social Media*. vol. 4; 2010.
39. Narang K, Nagar S, Mehta S, Subramaniam LV, Dey K. Discovery and analysis of evolving topical social discussions on unstructured microblogs. In: *European Conference on Information Retrieval*. Springer; 2013. p. 545–556.