NETWORK SCIENTIFIC AND INFORMATION THEORETIC APPROACHES TO SOCIAL MEDIA DURING EXTREME CLIMATE EVENTS

A Thesis Presented

by

Benjamin Freixas Emery

 to

The Faculty of the Graduate College

of

The University of Vermont

In Partial Fulfillment of the Requirements for the Degree of Master of Science Specializing in Complex Systems and Data Science

May, 2019

Defense Date: March 19th, 2019 Thesis Examination Committee:

Christopher M. Danforth, Ph.D., Advisor Peter Sheridan Dodds, Ph.D., Advisor Gregory Warrington, Ph.D. Meredith T. Niles, Ph.D., Chairperson Cynthia J. Forehand, Ph.D., Dean of Graduate College

Abstract

In addition to the tragedy they cause, major natural hazard and disaster events place a large cost on the governments and aid organizations who help people prepare for and recover from them. Such organizations are in constant need of strategies for distributing aid efficiently and comprehensively. The emergence of social media in everyday life has provided a platform for such organizations to coordinate relief efforts and communicate with people affected by disasters. It also has allowed affected individuals to communicate with one another on a large scale. The present thesis examines the dynamics of Twitter during extreme climate events and their aftermath in order to shed light on potential strategies for aid providers.

We begin by looking at the five most expensive natural disasters in the United States between 2011 and 2016. We isolate Twitter users for each disaster who are likely tweeting about food security or other basic needs during the event and its aftermath. We examine the follower count distributions of these users for each event. We then narrow focus to Hurricane Sandy, and look at the relationship between follower counts and relative increase in tweeting rate during the event. We find that users with fewer than 100 followers were more likely to increase their rate of tweet publication than influentials with many followers.

We also use a synthetic model of Twitter's communication network to mimic the way Twitter stores and samples tweet data. We quantify the sensitivity of three measures of network centrality to these mechanisms. This provides insight relevant to those who build network representations of Twitter communication using the data Twitter provides. We see differences in the sensitivity of the centrality measures studied, differences in sensitivity to the different mechanisms, and a dependence between measure and mechanism.

Finally, we construct a network representation of Puerto Rican Twitter users surrounding Hurricane María and its aftermath. We examine the evolution of this network over time, and communities present within the aggregate network. Using information theoretic tools, we discern differences in the body of tweets between different communities in the network and different periods of time surrounding the hurricane's landfall. We observe many differences between communities, with more focus on Puerto Rico in the community containing most local government figures, whereas major celebrities tended to talk about more general Latin American issues. We also hand-categorize Twitter users in the network as news outlets, politicians, citizens, weather stations, meteorologists, or journalists, finding that the distribution of user type has a temporal dependence.

DEDICATION

For my tía, tío, abuela, prima, y primo, who, During and after the hurricane, demonstrated That they are the strongest people I know.

Acknowledgements

Colossal thanks to the Computational Story Lab and the Pizza People, to Meredith, Chris, and Peter for the ideas and guidance, to Sam Rosenblatt and Sarah Shugars for peer-ushering the wonderful world of network science, to Brooke Foucault-Welles for inspiring the methods of Chapters 2 and 3, and to Juniper Lovato, Laurent Hébert-Dufresne, and Antoine Allard for organizing the greatest workshop of all time.

TABLE OF CONTENTS

	Dedica	tion	
	Acknow	vledgements	
	List of	Figures	
	List of	Tables	
1	Social	Media Usage Patterns During Natural Hazards	
	1.1 Ir	ntroduction	
	1.	1.1 Background	
	1.	1.2 Focus	
	1.2 M	Iethods	
	1.	2.1 Identifying tweets	
	1.	2.2 Distribution of Network Sizes and Relation to tweet Volume .	
	1.	2.3 Tweet Volume Increase by Network Size	
	1.3 R	esults	
	1.4 D	viscussion	
2	A Sim	nulation of the Limitations of Network Analysis on Twitter	
	2.1 In	ntroduction	
	2.2 M	lethods	
	2.	2.1 Synthesizing Toy Network	
	2.	2.2 Artificial Noise	
	2.	2.3 Sensitivity Analysis	
	2.3 R	esults	
	2.	3.1 Structure and Degree Distribution	
	2.	3.2 Sensitivity Coefficients	
	2.4 D	Discussion	
3	Twitte	er's Communication Network of Hurricane María Victims	
	3.1 Ir	ntroduction	
	3.	1.1 Background	
	3.	1.2 Existing Research	
	3.	1.3 Focus of Study	
	3.2 M	$\mathbf{I} ethods \ . \ . \ . \ . \ . \ . \ . \ . \ . \ $	
	3.	2.1 Data Collection	
	3.	2.2 Network Specification	
	3.	2.3 Shannon Entropy	
	3.	2.4 Jensen-Shannon Divergence	
	3.3 R	\tilde{s} esults	

e e	3.3.1 Communication Network
e e	3.3.2 Divergence
e	3.3.3 Daily Resolution Network Evolution
3.4]	Discussion
3.5 ($Conclusion \dots \dots$

LIST OF FIGURES

- 1.1 Log-log plot of the distribution of individuals observed to tweet during each disaster as a function of their follower count across the five disasters: a) Hurricane Irene; b) Hurricane Sandy; c) Louisiana flooding; d) Ohio Valley/Midwest tornadoes; e) Midwest/Southeast tornadoes. Most accounts have a small number of followers (e.g. less than 100), and a few accounts have many followers (e.g. more than 10,000). . . .
- 1.2 Log-log plot of the fractional change in tweet rate as a function of follower count for (a) before and during Hurricane Sandy and (b) the pairs of times collected for the null distribution. The increased density observed above 0 suggests that most individuals tweet more frequently during the disaster. In addition, the rate increase is largest for "average" individuals, i.e. those with 100 followers or fewer. This is of notable contrast to the null distribution, which is roughly symmetric about the zero-axis. Note that white pixels indicate one or zero individuals exhibiting the corresponding rate change.
- 1.3 Violin plots showing the distributions of fractional tweet rate change of the users found to be tweeting about Hurricane Sandy as it occurred (a) before and during Hurricane Sandy and (b) the pairs of times collected for the null distribution. Separate violin diagrams are drawn for users whose follower counts fall into each order of magnitude from 10^o to 10⁵. On each violin, a black bar indicates a Bayesian 95% confidence interval for the mean of the population distribution given the sample. For the Hurricane Sandy data, the intervals for 10^o and 10¹ are both notably higher than, and don't overlap with the intervals for any of the higher orders of magnitude in follower count. The same is not true for the null distributions, for which most of the intervals overlap and are generally closer together. The values of endpoints of the intervals are given in Table 1.3.
- 2.1 Schematic for how Twitter's database storage deviates the from the true underlying communication structure by only recording the original author of a Tweet that has been Retweeted. When user C is shown user A's Tweet by user B, the underlying communication network should really look as it does on the left. Twitter's method of recording data, however, leaves out information that results in the network representation looking as it does on the right.

9

13

2.	.2 Diagrams of the network structure and in-degree distribution histograms for the original synthesized network (a,b), the sampled network (c,d), the rewired network, (e,f), and the rewired, sampled network (g,h).	17
3. 3.	 The in-degree distribution of the network, plotted on a log-log scale. The networks of tweets using one or more of the keywords in Table 3.1 from users likely to be affected by Hurricane María or to have a direct connection to someone affected. The tweets were authored from September 16 to October 15. 	29 30
3.	.3 The tweet network before hurricane landfall on September 20	31
3	.4 The tweet network after hurricane leaves the island.	32
3.	.5 The contributions of individual words to the Jensen-Shannon Diver- gence between the collection of tweets before and after the period of landfall for Hurricane María. 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 Shan- non entropy of the tweets that contained that word, hashtag, or emoji. Darker colors indicate that tweets using the word, hashtag, or emoji are more diverse in content, where colors closer to white indicate that the word, hashtag, or emoji appeared mostly in retweets of the same tweet	25
ર	6 (a) The word-level contribution to the divergence between the collec-	30
0.	tion of tweets in the largest and fourth-largest community. (b) The word-level contribution to the divergence between the collection of	
	tweets in the largest and second-largest community.	36
3.	.7 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 3.8 in the appendix, where we show the same timeseries separated by	
~	individual accounts.	38
3.	.8 The same timeseries as Figure 3.7, but separated by individual user	41

LIST OF TABLES

1.1	Weather and Climate Billion-Dollar Hazard Events to affect the U.S.	
	from 2011-2016 (CPI-Adjusted)	3
1.2	Words used in the Twitter analysis	4
1.3	Bounds of the Bayesian 95% confidence intervals for the mean of the logged fractional increase in tweet rate in our Sandy and null datasets.	9
2.1	The coefficients of sensitivity for three measures of centrality to the sampling, rewiring, and the composition of the two	18
3.1	Words used to flag tweets in our study (top) and their English trans- lations (bottom).	25

CHAPTER 1

Social Media Usage Patterns During Natural Hazards

1.1 INTRODUCTION

1.1.1 BACKGROUND

In recent years, natural disaster relief has been the subject of much scrutiny and reform, as the frequency and severity of extreme events increases[42] and is projected to continue increasing [19]. Concurrent with this increased attention to extreme climate events has been the rise in prevalence of mobile communication technology and social media. In 2017 it was estimated that 77% of Americans own and use smartphones[1], 24% of Americans use Twitter, and 80% of social media activity happens on mobile devices [45].

The use of social media via mobile devices for practical communication during natural crisis events dates back to as early as 2010 when Haiti was struck by a magnitude-7 earthquake [30]. Since then, an ever-growing body of work has looked at what types of information are propagated by social media during these events, and how such information spreads [3, 6, 7, 26, 27].

1.1.2 Focus

This chapter is comprised of my own contributions to a manuscript published in the Public Library of Science by Niles et al in 2019[39]. Here we investigate the local networks of users (their followers in particular) and the relationship between the local network of a user and their activity during a natural hazard event. We do this in order to uncover useful information about whom to focus on when attempting to disseminate useful relief information on social media.

1.2 Methods

1.2.1 IDENTIFYING TWEETS

Table 1.1 provides statistics and descriptions of the five disaster events, provided by the National Oceanic and Atmospheric Administration (NOAA) [40]. For each of these events, we identified keywords indicating that a user is likely to be referring to the specific disaster. We include general terms as well as terms specific to the event in question. Table 1.2 shows the union of the sets of words used for all the events.

Event	Dates	Direct Summary (from NOAA)	CPI- Adjusted Estimated Cost (in Billions of Dollars)	Deaths
Hurricane Sandy	10/30/2012 to 10/31/2012	"Extensive damage across several northeastern states (MD, DE, NJ, NY, CT, MA, RI) due to high wind and coastal storm surge, particularly NY and NJ. Damage from wind, rain and heavy snow also extended more broadly to other states (NC, VA, WV, OH, PA, NH), as Sandy merged with a developing Nor'easter. Sandy's impact on major population centers caused widespread interruption to critical water / electrical services and also caused 159 deaths (72 direct, 87 indirect). Sandy also caused the New York Stock Exchange to close for two consecutive business days, which last happened in 1888 due to a major winter storm."	70.9	159
Hurricane Irene	8/26/2011 to 8/28/2011	"Category 1 hurricane made landfall over coastal NC and moved northward along the Mid-Atlantic Coast (NC, VA, MD, NJ, NY, CT, RI, MA, VT) causing torrential rainfall and flooding across the Northeast. Wind,damage in coastal NC, VA, and MD was moderate with considerable damage resulting from falling trees and power lines, while flooding caused extensive flood damage across NJ, NY, and VT. Over seven million homes and businesses lost power during the storm. Numerous tornadoes were also reported in several states further adding to the damage."	15.1	45
Southeast/ Ohio Valley/ Midwest Tornadoes	4/25/2011 to 4/28/2011	"Outbreak of tornadoes over central and southern states (AL, AR, LA, MS, GA, TN, VA, KY, IL, MO, OH, TX,OK) with an estimated 343 tornadoes. The deadliest tornado of the outbreak, an EF-5, hit northern Alabama, killing 78 people. Several major metropolitan areas were directly impacted by strong tornadoes including Tuscaloosa, Birmingham, and Huntsville in Alabama and Chattanooga, Tennessee, causing the estimated damage costs to soar."	11.4	321
Louisiana Flooding	8/12/2016 to 8/15/2016	"A historic flood devastated a large area of southern Louisiana resulting from 20 to 30 inches of rainfall over several days. Watson, Louisiana received an astounding 31.39 inches of rain from the storm. Two- day rainfall totals in the hardest hit areas have a 0.2% chance of occurring in any given year: a 1 in 500 year event. More than 30,000 people were rescued from the floodwaters that damaged or destroyed over 50,000 homes, 100,000 vehicles and 20,000 businesses. This is the most damaging U.S. flood event since Superstorm Sandy impacted the Northeast in 2012."	10.4	13
Midwest/ Southeast Tornadoes	5/22/2011 to 5/27/2011	"Outbreak of tornadoes over central and southern states (MO, TX, OK, KS, AR, GA, TN, VA, KY, IN, IL, OH, WI, MN, PA) with an estimated 180 tornadoes. Notably, an EF-5 tornado struck Joplin, MO resulting in at least 160 deaths, making it the deadliest single tornado to strike the U.S. since modern tornado record keeping began in 1950."	10.2	177

Table 1.1: Weather and Climate Billion-Dollar Hazard Events to affect the U.S. from 2011-2016 (CPI-Adjusted)

1.2.2 DISTRIBUTION OF NETWORK SIZES AND RELA-TION TO TWEET VOLUME

We examined statistics associated with the follower network of individuals who authored tweets in the collection described in Section 1.2.1. Specifically, for each tweet we use the author's user ID, as well as the number of accounts that followed the author. Individuals with multiple messages in the two-week window were assigned the follower count associated with their first tweet. The number of messages posted by each user during the interval of interest was aggregated by user as well, representing roughly 10% of their total number of posts. We plot the base-10 logarithm of user count for a matrix of binned message frequencies and follower counts. Using account data accumulated for all disasters, we plot the total number of tweets per account against the number of followers of the account using both linear and logarithmic scales. While the follower count is not a proxy for meaningful interaction, it is a first order approximation of the size of an account's audience.

canned	food assistance	food security	fridge	hurricane	rain	snow	unprepared
drinks	food bank	food shelf	generator	irene	sandy	store	water
emergency	food insecurity	food stamps	groceries	power	shelter	supermarket	watson
farm	food market	food store	grocery store	prepare	shock	supplies	wind
flood	food pantry	foods	help	preparing	SNAP	tornado	

Table 1.2: Words used in the Twitter analysis.

1.2.3 TWEET VOLUME INCREASE BY NETWORK SIZE

We estimated changes in individual behavior observed during Hurricane Sandy, compared to a baseline reference, as a function of network size. To do this, we used the

"total tweet count" field in the Decahose JSON metadata, which represents the exact number of messages posted to the account up to that moment. For each user found to have tweeted one of the keywords surrounding Hurricane Sandy's landfall, we collected the first and last tweet authored during the month of September 2012. We used these two tweets to compute a baseline tweet rate, found by taking the difference in total tweet count and dividing by the number of days between the two tweets. We repeated this process for October 21 through November 4, 2012 to compute the tweet rate during the disaster and its aftermath. We required at least two tweets during each period for a user to be included. We used the quotient of the tweet rate during Sandy (numerator) and the baseline tweet rate (denominator) to compute the estimated change in tweet volume for each user. Despite being restricted to a random 10% of messages, and therefore not being able to observe most tweets, the message rate calculation is exact for the period of observation. Work by Barabasi has shown that the rates of human activities such as emailing follow a Pareto distribution of lag time between events [10]. Although tweeting likely follows a similar distribution, our sample does not allow us to accurately measure the lag time between user tweets, so we are limited to this homogeneous approach.

We also generated a null-model of this change in tweet volume by using the same method to estimate the change for the same users between every month of the year and the following 16-day period, analogous to the dates we sampled for Hurricane Sandy. These pairs of time periods were all observed in 2012, except for those that overlapped with Hurricane Sandy, which were instead drawn from 2011.

1.3 Results

We sought to understand the relationship between tweet frequency and follower count. While the follower count associated with an individual is not a perfect reflection of their influence, it does serve as a proxy for the size of their audience. In looking at the follower counts associated with individuals tweeting about disaster events, we are seeking an understanding of the role various stakeholders play in the spread of information.



Figure 1.1: Log-log plot of the distribution of individuals observed to tweet during each disaster as a function of their follower count across the five disasters: a) Hurricane Irene; b) Hurricane Sandy; c) Louisiana flooding; d) Ohio Valley/Midwest tornadoes; e) Midwest/-Southeast tornadoes. Most accounts have a small number of followers (e.g. less than 100), and a few accounts have many followers (e.g. more than 10,000).

In Figure 1.1, we plot the distribution of follower counts, which appear typical for social networks. To explore user behavior further, we establish a baseline tweet rate for each account, and observe the increase (or decrease) in activity during the disaster event. We find a consistent trend that the individuals who tweet the most during disaster events tend to have "average" sized networks, as shown if Figures 1.2 and 1.3. Goncalves et al. found that social networks reflect Dunbar's number,

leading an individual's set of meaningful relationships to be limited to between 100 and 200 accounts [23]. It is these accounts in which we see the largest increase in activity during disasters. Previous work has found that "hidden influentials" in social networks, which are users with average-sized audiences, are key to allowing systemwide information-cascades and therefore play a major role facilitating protests online [9, 14, 24].

Our analysis also suggests that individuals were tweeting more frequently during Sandy than during other disaster events that we studied. Further analysis explores how the tweet rate changed as compared to baseline during Hurricane Sandy. In Figure 1.2 we see that, while the distribution of tweet rate change between two time periods is normally symmetric about 0 for users of all follower counts, this distribution for Hurricane Sandy is shifted upwards for users with 100 followers or fewer. In Figure 1.3, the same tweet rate increase distribution is shown, but the follower counts of the users are discretized by order of magnitude (0-9,10-99,100-999,...). This demonstrates that while users of all follower counts tend to have only a very small change in tweet rate during a typical baseline period (right panel), during Hurricane Sandy a positive average tweet rate change is observed for all users (left panel). More notably, the average tweet rate change is, slightly but significantly, higher during Sandy for users in the first and second groups: those with follower counts between 0 and 100. These results suggest that people with average-sized networks were more likely to tweet with a higher relative frequency during Hurricane Sandy than those with larger networks.

We note that even in the null comparison, nearly all of the groups show an increase in activity, which is seemingly paradoxical for this null distribution. We expect, however, that this is due to a subtle sampling bias in the fact that the "before" time period is longer than "during". Users were only included if at least two of their tweets were in our tweet database for both periods. Because the "during" period is shorter, we are more likely to include a user with an increase than one with a decrease. This is accounted for in our comparison between different groups and the comparison to the null distribution.



Figure 1.2: Log-log plot of the fractional change in tweet rate as a function of follower count for (a) before and during Hurricane Sandy and (b) the pairs of times collected for the null distribution. The increased density observed above 0 suggests that most individuals tweet more frequently during the disaster. In addition, the rate increase is largest for "average" individuals, i.e. those with 100 followers or fewer. This is of notable contrast to the null distribution, which is roughly symmetric about the zero-axis. Note that white pixels indicate one or zero individuals exhibiting the corresponding rate change.



Figure 1.3: Violin plots showing the distributions of fractional tweet rate change of the users found to be tweeting about Hurricane Sandy as it occurred (a) before and during Hurricane Sandy and (b) the pairs of times collected for the null distribution. Separate violin diagrams are drawn for users whose follower counts fall into each order of magnitude from 10^0 to 10^5 . On each violin, a black bar indicates a Bayesian 95% confidence interval for the mean of the population distribution given the sample. For the Hurricane Sandy data, the intervals for 10^0 and 10^1 are both notably higher than, and don't overlap with the intervals for any of the higher orders of magnitude in follower count. The same is not true for the null distributions, for which most of the intervals overlap and are generally closer together. The values of endpoints of the intervals are given in Table 1.3.

follower count order	lower (Sandy)	upper (Sandy)	lower (null)	upper (null)
10^{0}	0.136	0.225	0.033	0.111
10^{1}	0.173	0.187	0.045	0.054
10^{2}	0.108	0.114	0.051	0.054
10^{3}	0.094	0.107	0.030	0.036
10^{4}	0.062	0.096	0.016	0.031
10^{5}	0.031	0.120	-0.003	0.031

Table 1.3: Bounds of the Bayesian 95% confidence intervals for the mean of the logged fractional increase in tweet rate in our Sandy and null datasets.

1.4 DISCUSSION

Our results indicate a useful difference in the behavior of different types of Twitter users during crises. In particular, people with more modest followings (less than 100) seem to have more substantial reactions to a crisis event than those with larger followings. In 1992 the anthropologist Robert Dunbar theorized the number of people humans could maintain meaningful relationships with ranged from 100 to 250 [18]. We find that the users who react the most to disasters have followings that are users who could potentially have meaningful relationships with all of their followers. We suggest that the difference in increase is due to the greater likelihood of these connections being adjacent to meaningful relationships.

This has particularly important implications for events where aspects of relief efforts, like the type of relief provided, where to get help, the length of a wait, change rapidly with time. In these situations, it's important to provide new information to people willing to share it at the same frequency as it becomes available. In the situation where relief organizations have little access to specific information about individual Twitter accounts, the best course of action may be to target individuals with these modest followings, as they'll be likely to share more frequently, relative to their own baseline, in such an event.

The present study is limited in a number of ways. Notably, the interpretations we present cannot be taken as definite without at least conducting similar studies on a range of different disaster events. We also acknowledge that our approach to measuring change in Twitter activity is not the most appropriate due to heavy-tailed "bursty" nature of social media activity [38]. We are, however, confined to our approach in this study due to our use of the Gardenhose, as recognizing the heavy tail would require access to the timestamps of every tweet for each user.

CHAPTER 2

A SIMULATION OF THE LIMITATIONS OF NETWORK ANALYSIS ON TWITTER

2.1 INTRODUCTION

The collection of problems that are approached from a networks perspective is everdiversifying, as contributions from network science become more widely recognized by the greater scientific community, and the data to employ a networks perspective becomes more accessible. At an alarmingly high rate, however, when network science is employed in an applied setting, the effect of often inevitable data imperfection on the analysis is either briefly acknowledged or entirely ignored.

Problems exist not only with research that fails to address noisy data, but with research that neglects to fully acknowledge the mechanisms generating such noise. Research often tests noise types independently, assumes it to be random, or considers it outside the context of the data in question [13, 16, 21, 43, 44]. Methods have been developed to estimate full-network statistics, such as the degree distribution, from incomplete network data [11]. The collection of mechanisms that generate noise in network data can be visualized as points in a two-dimensional space with a node/edge error axis and an added/missing data axis. Boundary specification problems, edge misrouting, response bias, and other common network data noise can be represented as unique points in this space [31, 34]. Additionally, the variations in the data type also play a role in the way noise affects the network representation. For instance, data may be reported or observed, inferred or measured, part of a small or large system, and no single approach can adequately address the issues that could occur in all the possible scenarios.



Figure 2.1: Schematic for how Twitter's database storage deviates the from the true underlying communication structure by only recording the original author of a Tweet that has been Retweeted. When user C is shown user A's Tweet by user B, the underlying communication network should really look as it does on the left. Twitter's method of recording data, however, leaves out information that results in the network representation looking as it does on the right.

The third chapter of this thesis will rely heavily on building a picture of the communication network in a segment of Twitter from the Gardenhose API sample containing 10% of all Tweets. In this network, the nodes of this will be Twitter accounts, and the links will be tweets where the source is the author of the tweet and the target is a retweeted or mentioned user. There are many other ways to generate a network from tweet data, such as mapping the co-occurance of hashtags in tweets

[22] or the reciprocal replies to a tweet [12].

There are two major mechanisms of noise that prevent a fully accurate picture from being realized. One issue is the small sample size of the messages sent. The other and more subtle issue has to do with Retweets specifically. Consider three Twitter users: A, B, and C, where user B follows user A, and user C follows user B, but not user A. User B may Retweet a Tweet authored by user A, allowing it to be shown in user C's feed. If user C Retweets that Tweet, which they see due to user B's Retweet of it, Twitter only records that user C Retweeted user A, and leaves out all information about how user C really Retweeted something shown to them by user B. This process is shown graphically in Figure 2.1. This recording mechanism causes the resulting network to be more star-shaped than the underlying true communication network, which we explore.

2.2 Methods

2.2.1 Synthesizing Toy Network

We begin by generating a 100-node Watts-Strogatz small-world network with k = 2and p = 0.05 [49]. We then cycle through 30 rounds of random edge addition. Each round, we cycle through every ordered pair of nodes i, j and add an edge from i to jwith probability $p = 0.5 \frac{d_j}{m}$, where d_j is the current in-degree of node j, and m is the current number of edges in the network. The resulting hybrid of a small-world and scale-free network serves as a reasonable toy model for an online social network such as Twitter.

2.2.2 Artificial Noise

To simulate the sampling procedure of the Gardenhose, we randomly delete 90% of the edges in the network. To simulate Twitter's misleading Retweet recording, we iterate through all directional paths of length 2 from node *i* to *j* to *k*, and with probability $p = 0.5 \frac{d_k}{m}$, we rewire the *ij* edge to an *ik* edge.

2.2.3 Sensitivity Analysis

To determine the sensitivity of different centrality measurements to a transformation of the network, we use a measurement defined by Martin and Niemeyer [35]. This "sensitivity coefficient" of a centrality measure c to a transformation of network G to G' is defined as

$$\rho_c(G, G') = \frac{n_c}{n_c + n_d},$$
(2.1)

where n_c is the number of node pairs whose order is preserved in a node-list ranked by the centrality measure, and n_d is the number of node pairs whose order changed in such a node list. This coefficient is thus simply the fraction of node pairs whose order changes. The measure is optimally robust if $\rho_c(G, G') = 1$, and very sensitive if $\rho_c(G, G') = 0$.

2.3 Results

2.3.1 Structure and Degree Distribution

Figure 2.2 shows the structure of all versions of the synthetic network described in Section 2.2.1, as well as their respective in-degree distributions. The initial synthetic distribution shown in Figure 2.2(b) has a substantial right skew, with most nodes having in-degree 0 or 1, with the count rapidly decaying as the degree rises, with one solitary node having an in-degree higher than 60. The sampling of the edges yields a network whose degree distribution maintains a similar shape, as seen in Figure 2.2(d), but with the total number of nodes decreasing, and the range of in-degrees represented shrinking to less than one sixth the original range.







(b)







Figure 2.2: Diagrams of the network structure and in-degree distribution histograms for the original synthesized network (a,b), the sampled network (c,d), the rewired network, (e,f), and the rewired, sampled network (g,h).

The rewiring of the network results in a substantially deviant in-degree distribution. The weight of the tail has not only increased, but the vast majority of this increase has been due to a major increase in the largest node's in-degree. This reflects a mechanism similar to the first-mover advantage found in Simon's rich-get-richer model [17]. When this network is sampled in the manner of Twitter's Gardenhose, the shape of the distribution is not substantially changed.

2.3.2 Sensitivity Coefficients

Table 2.1 shows the sensitivity of betweenness, closeness, and degree centrality to the sampling of the edges, the rewiring of the network, and the rewiring *then* sampling. We find that all the centrality measures are more robust to the rewiring than the sampling. Degree centrality is the most robust to all the transformations, and is no more sensitive to the composition of rewiring and sampling than to sampling alone. Closeness is more sensitive to all the transformations and betweenness is the most sensitive. Both closeness and betweenness are more sensitive to sampling than rewiring but most sensitive to their composition.

centrality measure	ρ (sampled)	ρ (rewired)	ρ (rewired & sampled)
betweenness	0.53	0.70	0.37
closeness	0.62	0.77	0.54
degree	0.59	0.79	0.59

Table 2.1: The coefficients of sensitivity for three measures of centrality to the sampling, rewiring, and the composition of the two.

2.4 DISCUSSION

In this chapter, we briefly explore some of the biases encountered when constructing a communication network from Twitter data. The findings are particularly important to consider in conjunction with Chapter 3, which does exactly that. We find that when working with such data, it's best to use degree centrality if the option is available. Additionally, we find that it's largely advantageous to have access to Twitter's Firehose, a feed of every Tweet, as the centrality measures are much more sensitive to the sampling than to Twitter's implicit rewiring that occurs as the result of their recording. Betweenness centrality is particularly sensitive to sampling, and is the only measure for which more than half of node pairs had their rank-orders flipped by the rewiring and sampling.

There are several limitations to consider in this toy model. The synthesis of the toy Twitter network is not a perfectly accurate representation of the Twitter's network production mechanism, as it initializes as a basic small-world network, and only introduces new edges, not new nodes. The size of the network, at only 100 nodes, likely introduces some error. In particular, the biggest node, which demonstrated the strong first-mover advantage in the mechanistic rewiring of edges, only represents 1% of the "users". As of 2013, Twitter had about 2 billion accounts, and while the 99th percentile by follower count was about 3,000, the 99.9th percentile was about 25,000 [15]. There are therefore some considerable liberties taken in assuming the top 1% of users have the same follower count.

Additionally, the rewiring algorithm used is a simplified one that only considers chains of two Retweets. On Twitter, Retweet chains of any length are shortened in the database, so it's possible that the error caused by this is more extreme than that presented.

CHAPTER 3

TWITTER'S COMMUNICATION NETWORK OF HURRICANE MARÍA VICTIMS

3.1 INTRODUCTION

3.1.1 BACKGROUND

In the early morning of September 20th 2017, Puerto Rico was struck by its first category-4 cyclone in 85 years [37]. Hurricane María 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 [4], was entirely wiped out. Most of the island was left without power for nearly two months, and some areas didn't get power back for nearly a year [46].

Although estimates vary, the Puerto Rican government accepts an independent study approximating that 2,975 deaths resulted from the storm and its aftermath [5]. An internal report at the Federal Emergency Management Agency acknowledges a failure to adequately prepare for the 2017 hurricane season, resulting in the agency being unable to properly support the victims of Hurricane María [47].

Despite 89% of cell phone towers being offline [20], many Puerto Ricans were connecting with loved ones 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 [36].

At the time, Hurricane María was the latest in a sequence of natural crisis events spanning nearly a decade in which social media was utilized for practical communication. As the example with the highest death toll in recent memory, we hope to uncover specific properties of the communication via social media of María victims using tools from network science and information theory.

3.1.2 EXISTING RESEARCH

Work surrounding previous major natural hazard events 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. 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 [39]. 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 [2]. Researchers have

also found that during hurricanes Twitter can be used to predict how to best focus recovery efforts [25], and that Twitter activity correlates with hurricane damage [32]. Algorithms have been developed to identify flood victims asking for help on Twitter [41].

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 [8, 28, 29]. 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: after the Deepwater Horizon oil spill, Twitter networks were used to study the online conversation as the recovery efforts unfolded [48].

3.1.3 Focus of Study

We seek to identify how Twitter could potentially have been leveraged to increase the effectiveness of aid to Puerto Rico during Hurricane María. More specifically, we hope to uncover who the best people on Twitter would be to provide with information about aid for such information to spread maximally to the people who need it. Are such accounts different when dealing with information about different types of aid? Do they change over time as the cycle from anticipation to aftermath unfolds? We hope that a network topology for communication and information theoretic tools provide answers that would be potentially useful to humanitarian agencies.

3.2 Methods

3.2.1 DATA COLLECTION

Our goal 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. By doing this we collected 93 users who sent geolocated tweets from Puerto Rico during the hurricane. We next used Twitter's free Premium API to search for all tweets that mentioned or retweeted a user from our original list between the dates of September 16 and October 15 (inclusive).

Our time interval is chosen to maximize preparatory time while avoiding overlap with the recovery from Hurricane Irma, which had hit only days before. We choose to extend 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 didn't enable geolocation, as well as people elsewhere who are personally connected to people on the island during the hurricane.

It's important that we gather accounts with geolocation turned off, as geolocated tweets comprise approximately 1% of all content on Twitter. This extended our list of account names to 1,511 users. With our fully-generated account list, we scanned Twitter's Gardenhose, a random 10% sample of public tweets, and recorded every 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 restrict our corpus to Spanish tweets as an additionally stringent criterion, 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 keep only those using hand-selected vocabulary that indicates a probable discussion of the hurricane's effects or need for supplies, shown in Table 3.1. These words are adapted from the list of words used in Chapter 1, with the addition of words

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

Table 3.1: Words used to flag tweets in our study (top) and their English translations (bottom).

3.2.2 NETWORK SPECIFICATION

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 Gardenhose as described in Section 3.2.1 constitutes a directed link from the author to the mentioned or retweeted user.

We identify communities in the network topology with a Girvan-Newman clustering algorithm. Inspection of the network's tweet content revealed that the second largest community consisted almost entirely of people involved in the Venezuelan uprising, and several smaller communities shared this property. In order to get the network to represent the focus of the present study, we counted the proportion of tweets in the biggest Venezuelan network that contained the words "venezuela", "vzla", or "maduro", and computed the measurement's Wilson interval lower bound at 99% confidence. We 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 user list due to the heavy-tailed nature of online social networks [33]. We hand-identify the fifty largest nodes by in-degree as 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

$$d = \frac{m}{n(n-1)},\tag{3.1}$$

where n is the number of nodes and m is the number of 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 path-length 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 path lengths, we compute the average shortest path-length as

$$a = \sum_{s,t \in V} \frac{p(s,t)}{n(n-1)},$$
(3.2)

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

3.2.3 Shannon Entropy

We heavily rely on information theoretic measurements in our analysis of actual 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

$$H = -\sum_{i=1}^{n} p_i \log_2 p_i,$$
(3.3)

where the *i*th word appears with probability p_i . This entropy is maximized when every word occurs with equal probability, and approaches 0 as the corpus becomes dominated by one unique token.

3.2.4 JENSEN-SHANNON DIVERGENCE

To compare two subsets of our corpus, we compute the Jensen-Shannon divergence between them. This is a symmetric measurement of the difference between the respective probability mass functions of the corpora being compared. If we consider two corpora, P and Q, and their combined distribution $M = \pi_p P + \pi_q Q$, where π_p and π_q are the relative sizes of the two corpora such that $\pi_p + \pi_q = 1$, the Jensen-Shannon Divergence D is given by

$$D(P||Q) = \pi_p \sum_{i=1}^k p_i \log_2 \frac{p_i}{m_i} + \pi_q \sum_{i=1}^k q_i \log_2 \frac{q_i}{m_i}.$$
(3.4)

In this definition, the mixed distribution M has k unique words, and p_i , q_i , m_i are the probabilities of encountering M's *i*th word in P, Q, or M respectively. This expression may be recognized as the weighted average of each corpus's Kullback-Leibler divergence from the mixed corpus, although it can be expressed more succinctly as

$$D(P||Q) = \sum_{i=1}^{k} \pi_p p_i \log_2 p_i + \pi_q q_i \log_2 q_i - m_i \log_2 m_i.$$
(3.5)

It's evident from its form in Equation 3.5 that the Jensen Shannon Divergence is simply a sum of contributions from each word in the mixed corpus, where the contribution from word i is

$$D(P||Q)_i = \pi_p p_i \log_2 p_i + \pi_q q_i \log_2 q_i - m_i \log_2 m_i.$$
(3.6)

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

3.3 Results

3.3.1 Communication Network



Figure 3.1: The in-degree distribution of the network, plotted on a log-log scale.

The majority of nodes in the aggregate network described in Section 3.2.2 have in-degree of 0 or 1, while only four have an in-degree greater than 15. The full in-degree distribution is shown in Figure 3.1. This is consistent with the heavy tails usually seen in online social network degree distributions [33]. Among the fifty highest ranked Twitter accounts by in-degree of this network, sixteen are news outlets, eight are journalists, six are agencies (government and NGO), six are musicians, four are weather stations or meteorologists, and three are citizens.

The topology of the communication networks is shown in Figures 3.2, 3.3, and 3.4. The size of each node is proportional to the node's in-degree, or the number of tweets tagging or retweeting that user. Communities detected in the network topology by the Girvan-Newman clustering algorithm are colored accordingly. Figure 3.2 shows the aggregate network for September 16 to October 15 2017. Figures 3.3 and 3.4 show



Figure 3.2: The networks of tweets using one or more of the keywords in Table 3.1 from users likely to be affected by Hurricane María or to have a direct connection to someone affected. The tweets were authored from September 16 to October 15.



Figure 3.3: The tweet network before hurricane landfall on September 20.



Figure 3.4: The tweet network after hurricane leaves the island.

the network during anticipation of the hurricane (September 16 – September 19) and during the aftermath (September 22 – October 15).

The aggregate network has one major community of more than 300 nodes, which holds the majority of the highest ranked nodes by in-degree. Alongside this community, there are three communities of around 100 nodes. The remainder of the detected communities have fewer than 100 nodes, and most have fewer than 20. The four users with in-degree greater than 15 are all in the largest community. Those are Puerto Rican Governor Ricardo Rossello, one weather station, and two news outlets. The network has a density of $d = 6.1 \times 10^{-4}$ and the average shortest path-length is a = 1.3. These measurements indicate that although the 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 Figure 3.1.

The anticipation network 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.

3.3.2 DIVERGENCE

We use Jensen-Shannon divergence Wordshifts 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, such 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 Figure 3.5, 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 Figure 3.6(a)'s Wordshift are the tweets from the central community containing Governor Rossello's account to those from the fourth-largest community, which contains several famous Puerto Rican musicians such as Chayanne, Luis Fonsi, and Ricky Martin. Although this community is dominated by famous Puerto Ricans, discussion within this community seems to give space to the Mexican earthquake as well, as indicated by the prominence of "méxico" and the relative infrequency of "maría" and "luz", which directly translates to "light", but colloquially means "electricity" or "power".



Figure 3.5: The contributions of individual words to the Jensen-Shannon Divergence between the collection of tweets before and after the period of landfall for Hurricane María. 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. Darker colors indicate that tweets using the word, hashtag, or emoji are more diverse in content, where colors closer to white indicate that the word, hashtag, or emoji appeared mostly in retweets of the same tweet.



Figure 3.6: (a) The word-level contribution to the divergence between the collection of tweets in the largest and fourth-largest community. (b) The word-level contribution to the divergence between the collection of tweets in the largest and second-largest community.

3.3.3 Daily Resolution Network Evolution

We separate the aggregate network into its daily subnetworks, and look at the topfifty users in the aggregate network at the daily resolution. Figure 3.7 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 network, but he appears to mostly speak with smaller nodes, and not as much with other major figures.



Figure 3.7: 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 3.8 in the appendix, where we show the same timeseries separated by individual accounts.

3.4 DISCUSSION

Previous work found that, in response to a major natural crisis event, average users with small followings increased their social media activity more than influentials with large followings [39]. Those results spoke toward who was responsible for the propagation of information throughout the Twitter network. The results we present here begin to address the naturally following question: where is such 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.

This helps paint one more piece of the complete picture of crisis communication. We now have reason to believe that information is spread throughout the network by average individuals, the first of which get such information from "local" influencers, such as the politicians and journalists from Puerto Rico. Meanwhile, major celebrities seem mostly removed from the conversation with the conversation centered on local influencers and average affected individuals, touching the Puerto Rico crisis along with a variety of issues in Latin America.

There are several important limitations to acknowledge surrounding this work and its possible implications. 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. Additionally, most people of interest had no cell-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.

3.5 CONCLUSION

Time and time again people prove incredibly resilient under conditions of extreme diversity. In the case of Hurricane María, 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 important implications in 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.

In any natural crisis situation, there's 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 blockages that keep information to those who need it, however, can render this help not useful to overwhelming masses of crisis victims. It's 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, at least on Twitter, best plant information so that it most efficiently propagates through the network, getting to as many people as possible, as quickly as possible.

APPENDIX



Figure 3.8: The same timeseries as Figure 3.7, but separated by individual user.

BIBLIOGRAPHY

- [1] (2018). Mobile fact sheet. *Pew Research Center*.
- [2] Abdullah, N. A., Nishioka, D., Tanaka, Y., and Murayama, Y. (2017). 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.*
- [3] Alexander, D. E. (2013). Social media in disaster risk reduction and crisis management. Science and Engineering Ethics, 20(3):717–733.
- [4] Allen, G. and Peñaloza, M. (2015). Power problems: Puerto rico's electric utility faces crippling debt. *National Public Radio*.
- [5] Andrade, E., Barrett, N., Colon-Ramos, U., Edberg, M., Garcia-Meza, A., Goldman, A., Goldman, L., Roess, A., Sandberg, J., Santos-Burgoa, C., Zeger, S., Merly, N. E., Nazario, C. M. N., Pérez, C., Pérez, E. S., Latinovic, L., and Rivera, I. (2018). Ascertainment of the estimated excess mortality from hurricane maría in puerto rico.
- [6] Bagrow, J. P. (2018). Information spreading during emergencies and anomalous events. In *Complex Spreading Phenomena in Social Systems*, pages 269–286. Springer.

- [7] Bagrow, J. P., Wang, D., and Barabasi, A.-L. (2011). Collective response of human populations to large-scale emergencies. *PloS one*, 6(3):e17680.
- [8] Bailey, M., Jackson, S. J., and Foucault Welles, B. (2019). Women tweet on violence: From #yesallwomen to #metoo. Ada: A Journal of Gender, New Media, and Technology, (15).
- [9] Baños, R. A., Borge-Holthoefer, J., and Moreno, Y. (2013). The role of hidden influentials in the diffusion of online information cascades. *EPJ Data Science*, 2(1):6.
- [10] Barabasi, A.-L. (2005). The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039):207.
- [11] Bliss, C. A., Danforth, C. M., and Dodds, P. S. (2014). Estimation of global network statistics from incomplete data. *PloS one*, 9(10):e108471.
- [12] Bliss, C. A., Kloumann, I. M., Harris, K. D., Danforth, C. M., and Dodds,
 P. S. (2012). Twitter reciprocal reply networks exhibit assortativity with respect to happiness. *Journal of Computational Science*, 3(5):388–397.
- [13] Borgatti, S. P., Carley, K. M., and Krackhardt, D. (2006). On the robustness of centrality measures under conditions of imperfect data. *Social networks*, 28(2):124– 136.
- [14] Borge-Holthoefer, J., Perra, N., Gonçalves, B., González-Bailón, S., Arenas, A., Moreno, Y., and Vespignani, A. (2016). The dynamics of information-driven coordination phenomena: A transfer entropy analysis. *Science advances*, 2(4):e1501158.

- [15] Bruner, J. (2013). Tweets loud and quiet. O'Reilly.
- [16] Costenbader, E. and Valente, T. W. (2003). The stability of centrality measures when networks are sampled. *Social networks*, 25(4):283–307.
- [17] Dodds, P. S., Dewhurst, D. R., Hazlehurst, F. F., Van Oort, C. M., Mitchell, L., Reagan, A. J., Williams, J. R., and Danforth, C. M. (2017). Simon's fundamental rich-get-richer model entails a dominant first-mover advantage. *Physical Review E*, 95(5):052301.
- [18] Dunbar, R. I. (1992). Neocortex size as a constraint on group size in primates. Journal of human evolution, 22(6):469–493.
- [19] Field, C. B., Barros, V., Stocker, T. F., and Dahe, Q. (2012). Managing the risks of extreme events and disasters to advance climate change adaptation: special report of the intergovernmental panel on climate change. Cambridge University Press.
- [20] FitzGerald, D. and Hernandez, D. (2017). A week after Maria's strike, Puerto Rico remains offline. Wall Street Journal.
- [21] Frantz, T. L., Cataldo, M., and Carley, K. M. (2009). Robustness of centrality measures under uncertainty: Examining the role of network topology. *Computational and Mathematical Organization Theory*, 15(4):303.
- [22] Gallagher, R. J., Reagan, A. J., Danforth, C. M., and Dodds, P. S. (2018). Divergent discourse between protests and counter-protests: #blacklivesmatter and #alllivesmatter. *PloS one*, 13(4):e0195644.

- [23] Gonçalves, B., Perra, N., and Vespignani, A. (2011). Modeling users' activity on twitter networks: Validation of dunbar's number. *PloS one*, 6(8):e22656.
- [24] González-Bailón, S., Borge-Holthoefer, J., and Moreno, Y. (2013). Broadcasters and hidden influentials in online protest diffusion. *American Behavioral Scientist*, 57(7):943–965.
- [25] Guan, X. and Chen, C. (2014). Using social media data to understand and assess disasters. *Natural hazards*, 74(2):837–850.
- [26] Gurman, T. and Ellenberger, N. (2015). Reaching the global community during disasters: Findings from acontent analysis of the organizational use of twitter after the 2010 haiti earthquake. J. Health Commun, 20:687–696.
- [27] Houston, J. B. and Franken, N. J. (2015). Disaster interpersonal communication and posttraumatic stress following the 2011 joplin, missouri, tornado. *Journal of Loss and Trauma*, 20(3):195–206.
- [28] Jackson, S. J. and Foucault Welles, B. (2015). Hijacking #mynypd: Social media dissent and networked counterpublics. *Journal of Communication*, 65(6):932–952.
- [29] Jackson, S. J. and Foucault Welles, B. (2016). #ferguson is everywhere: initiators in emerging counterpublic networks. *Information, Communication & Society*, 19(3):397–418.
- [30] Keim, M. (2011). Emergent use of social media: A new age of opportunity for disaster resilience. *Prehospital and Disaster Medicine*, 26(S1).
- [31] Kossinets, G. (2006). Effects of missing data in social networks. Social networks, 28(3):247–268.

- [32] Kryvasheyeu, Y., Chen, H., Obradovich, N., Moro, E., Van Hentenryck, P., Fowler, J., and Cebrian, M. (2016). Rapid assessment of disaster damage using social media activity. *Science advances*, 2(3):e1500779.
- [33] Kumar, R., Novak, J., and Tomkins, A. (2010). Structure and evolution of online social networks. In *Link mining: models, algorithms, and applications*, pages 337– 357. Springer.
- [34] Laumann, E. O., Marsden, P. V., and Prensky, D. (1989). The boundary specification problem in network analysis. *Research methods in social network analysis*, 61:87.
- [35] Martin, C. and Niemeyer, P. (2017). Estimating the sensitivity of centrality measures w.r.t. measurement errors. CoRR, abs/1704.01045.
- [36] Mazzei, P. (2017). No cell service in puerto rico? drive to a highway in san juan and pull over. *Miami Herald*.
- [37] Meyer, R. (2017). What's happening with the relief effort in puerto rico? a timeline of the unprecedented catastrophe of hurricane maria. *The Atlantic*.
- [38] Myers, S. A. and Leskovec, J. (2014). The bursty dynamics of the twitter information network. In Proceedings of the 23rd international conference on World wide web, pages 913–924. ACM.
- [39] Niles, M. T., Emery, B. F., Reagan, A. J., Dodds, P. S., and Danforth, C. M. (2019). Social media usage patterns during natural hazards. *PLOS ONE*, 14(2):1–16.

- [40] Oceanic, N. and Administration, A. (2016). Billion-dollar weather and climate disasters.
- [41] Singh, J. P., Dwivedi, Y. K., Rana, N. P., Kumar, A., and Kapoor, K. K. (2017). Event classification and location prediction from tweets during disasters. *Annals of Operations Research*, pages 1–21.
- [42] Smith, A. B. (2018). 2017 us billion-dollar weather and climate disasters: a historic year in context. NOAA Climate. gov [available online at: https://www. climate. gov/news-features/blogs/beyond-data/2017-us-billion-dollar-weather-andclimate-disasters-historic-year, accessed 26 May 2018].
- [43] Smith, J. A. and Moody, J. (2013). Structural effects of network sampling coverage i: Nodes missing at random. *Social networks*, 35(4):652–668.
- [44] Smith, J. A., Moody, J., and Morgan, J. H. (2017). Network sampling coverage ii: The effect of non-random missing data on network measurement. *Social networks*, 48:78–99.
- [45] Smith A, Anderson M, C. T. (2018). Social media user in 2018. Pew Research Center.
- [46] Sullivan, E. (2018a). Nearly a year after maria, puerto rico officials claim power is totally restored. *National Public Radio*.
- [47] Sullivan, L. (2018b). Fema report acknowledges failures in puerto rico disaster response. National Public Radio.
- [48] Sutton, J., Spiro, E., Butts, C., Fitzhugh, S., Johnson, B., and Greczek, M. (2013). Tweeting the spill: Online informal communications, social networks,

and conversational microstructures during the deepwater horizon oilspill. International Journal of Information Systems for Crisis Response and Management (IJISCRAM), 5(1):58–76.

[49] Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *nature*, 393(6684):440.