

The Sound of Silence: Exploring How Decreases in Tweets Contribute to Local Crisis Identification

Rachel Samuels

Georgia Institute of Technology
rachel.samuels@gatech.edu

John Taylor

Georgia Institute of Technology
jet@gatech.edu

Neda Mohammadi

Georgia Institute of Technology
nedam@gatech.edu

ABSTRACT

Recent research has identified a correlation between increasing Twitter activity and incurred damage in disasters. This research, however, fails to account for localized emergencies occurring in areas in which people have lost power, otherwise lack internet connectivity, or are uncompelled to Tweet during a disaster. In this paper, we analyze the correlation between daily Tweet counts and FEMA Building Level Damage Assessments during Hurricane Harvey. We find that the absolute deviation of Tweet counts from steady state is a potentially useful tool for the evolving information needs of emergency responders. Our results show this to be a more consistent and persistent metric for flood damage across the full temporal extent of the disaster. This shows that, when considering the varied information needs of emergency responders, social media tools that seek to identify emergencies need to consider both where Tweet counts are increasing and where they are dropping off.

Keywords

Crisis informatics, emergency response, flooding, hurricanes, social media, Twitter

INTRODUCTION

By the start of August 2017, natural disasters in 2017 had already cost the United States 15 billion dollars and the lives of almost 300 people. The yet-uncalculated costs of the three major hurricanes that hit the East Coast in August and September are expected to catapult that total cost into the hundreds of billions (Lam, 2017). The first of these three, Hurricane Harvey is expected to ultimately cost more than \$183 billion dollars in federal relief (Amadeo, 2017). As the frequency and intensity of hurricanes are only expected to increase, understanding how to improve our crisis infrastructure is of critical importance. Our ability to both minimize initial damage and accelerate recovery is contingent on a wide variety of factors, including primarily the availability of emergency resources and how we structure our emergency response. We will need to improve the efficiency of both factors in order to mitigate future disaster-related costs.

The National Center for Environmental Information (NCEI) reports that tropical cyclones are, on average, the costliest of natural disasters in the United States (NCEI, 2017). This cost can be mitigated by an increase in the resilience of the community infrastructure (Coughlan De Perez et al., 2015), the availability of social, economic, and physical capital (Abramson et al., 2014), and the speed and efficiency of the emergency response network (Conrado, Neville, Woodworth, & O’Riordan, 2016). Much of crisis mitigation research has focused on understanding the importance of and strengthening of these resilience factors. Empowering emergency response organizations by increasing information availability has recently taken center stage in crisis research (Shklovski, Palen, & Sutton, 2008).

Information regarding disaster severity, human locations and mobility, and resource locations tends to be scarce immediately after disaster strikes and in the weeks following. Victims of disasters are one of the main providers

of crisis information, as they are the most intimately familiar with the aftermath. First responders are often operating on eye-witness accounts when they arrive at the scene of local emergencies, such as flooded houses or electrical fires, that comprise the larger disaster. However, it is difficult to acquire and accumulate victims' local knowledge during a massive crisis (Mason, Drew, & Weaver, 2017). There is an unmet need to facilitate the utilization of affected people's information in emergency response. In the past few years, a growing amount of research has sought to utilize humans as sensors—that is, using how victims “sense” damage and danger and how their behavior changes in response—to produce usable data for emergency responders (Hiltz et al., 2014; Kryvasheyeu, Chen, Moro, Van Hentenryck, & Cebrian, 2015; Sakaki, Okazaki, & Matsuo, 2010b; Tien et al., 2016). The conversion of human behavior data to empirical crisis data has been mainly performed through the use of social media; however, that crisis data is not currently in widespread use. We need to improve the accuracy and usability of social media tools before they can be implemented effectively.

LITERATURE REVIEW

Since the explosive increase in Twitter users in 2009 and the increasingly widespread use of other social platforms, researchers have been trying to utilize social media to characterize and improve community resilience to disaster. Initially, researchers used social media data to analyze the details of a specific prior event such as an earthquake (Sakaki, Okazaki, & Matsuo, 2010a); more recently, researchers have been trying to use the real-time flow of information from victims to advise responders during an emergency. The behavioral data recorded through social media has been examined for use in detecting emergencies (Kryvasheyeu et al., 2015), emergency types (Hiltz et al., 2014), damage severity (Kryvasheyeu et al., 2016), detecting resource availability and need (Choe, Park, Han, Park, & Yun, 2017), identifying human sentiment in real time (Caragea, Squicciarini, Stehle, Neppalli, & Tapia, 2014), and characterizing resilience and community recovery through mobility patterns (Wang & Taylor, 2014). Authors have utilized this information to create a variety of tools that are intended to assist professional emergency responders and citizen volunteerism in real time. These tools mainly utilize Twitter and Facebook, and primarily take the form of smart phone applications and websites, such as Ushahidi and XHELP (Reuter, Hughes, & Kaufhold, 2018).

Responders need to know, with high accuracy, the most time- and cost-effective areas in which to focus their efforts. Knowing where to go involves knowing where people and emergencies are. Knowing where to go *first* requires knowing where people are in the most danger relative to the people around them. Social media data is useful for the former, because geo-located Tweets show exactly when and where a single person is. Previous research has concerned itself with connecting social media to the latter. The current theory operates on the idea that people Tweet when they feel endangered, when they feel they can receive aid through social media, or, at the very least, when something abnormal is happening. More Tweets are thought to indicate abnormal occurrences and thus indicate when and where a disaster may be occurring.

In the search for a correlation between Tweeting and the presence of a disaster, Kryvasheyeu et al. performed a rigorous analysis of how well Twitter was able to predict infrastructure costs before, during, and after Hurricane Sandy (Kryvasheyeu et al., 2016). Large numbers of Tweets per day at the ZCTA-scale corresponded well with areas with higher insurance claims, and vice versa. This relationship was identified in twelve other emergency events in 2013 and 2014 at varying strengths. However, despite the strong correlation discovered in the above research, emergency identification and responder advisement systems do not currently have widespread use during disasters (Mason et al., 2017). Many of the inhibiting factors for tool usage are related to a perceived lack of accuracy (Imran, Castillo, Diaz, & Vieweg, 2015).

Accuracy in disaster informatics consists of rumor and completeness. The accuracy of Twitter data with respect to rumor has been addressed through studies that have sought to quantify rumors' propagation through networks (Starbird, Maddock, Orand, Achterman, & Mason, 2014) and to identify rumors in automatic Tweet filtration tools (Zeng, Starbird, & Spiro, 2016). Recent studies have been able to classify Tweets as rumors in emergencies, such as the 2013 Boston Marathon Bombing, with reasonable accuracy (Lu & Brelsford, 2015; Zeng et al., 2016). There is, however, another important aspect to accuracy: data completeness. A lack of data completeness can produce false negatives. A false positive such as a rumor would result in a waste of time and resources; a false negative could result in overlooked victims. Much of existing social media research assumes that if there is a disaster, people are Tweeting about it, but there is little research on the detriments of only using areas that are Tweeting in emergency identification and danger assessment (Reuter et al., 2018). There is also a lack of research on how we can potentially use the gaps in Twitter activity.

Emergency analysis tools that can only function in areas with high volumes of Tweets may be biased against underprivileged people in disasters that can heavily impact energy infrastructure, such as hurricanes. The Rutgers report on the impact of Superstorm Sandy showed that, although lower income families were hit by more than half of the costs of the storm, they received only 27% of the distributed aid. Wealthier communities

with more robust houses and energy infrastructures are also less likely to fear for their livelihoods or fear an extended loss of power, so they would be less likely to take precautions to conserve power. FEMA has encouraged people needing to conserve power to activate their phones' airplane mode (Fugate, 2017), which would remove the phones' data capabilities and the ability to Tweet.

Shelton et al. sought to demonstrate potential flaws in big data by looking at the relationship between Twitter activity and areas in the High-Impact Zone assessment produced by FEMA (Shelton, Poorthuis, Graham, & Zook, 2014). They normalized their data to a steady-state Tweet count value calculated from two weeks prior to the disaster instead of normalizing the activity to the number of active users. This normalization strategy reduces the possibility of areas without any active users during a disaster situation being ignored by the correlation analysis. Their results show that, while some areas' increase in Twitter activity did correlate with being in a High-Impact Zone, other areas lacked Tweets. These areas were termed "data shadows". A large data shadow was particularly evident in Staten Island, in which more than half of all Sandy-related deaths occurred. The majority of social media emergency identification tools developed to date would have been unable to help emergency responders prevent these deaths, as the silent areas would have been ignored.

In order to make social media analysis tools useful for emergency responders, we need to show that the data can be utilized in a way that does not overlook areas that do not Tweet during a disaster because of older demographics or lower socioeconomic status. We theorize that natural disasters can influence human behavior by causing some people to Tweet more, and others to Tweet far less. As such, strong deviations from steady state Twitter activity may be indicative of separate hazards and risks in addition to the spikes in Twitter activity utilized in the majority of social media crisis studies. By accounting for the data shadows identified by Shelton et al., we can increase the usability of social media data for emergency responders while reducing the risk of ignoring areas in need of assistance that do not have the ability or compulsion to Tweet.

In this paper, we seek to determine if there is a significant, useful correlation between hurricane damage and a decline in Twitter intensity that is missed in analyses focused on Twitter activity spikes. We chose to apply our analysis to Hurricane Harvey, which heavily impacted Houston, the fifth-biggest metropolis in the U.S.

RESEARCH DESIGN AND METHODS

Spatial Grid Creation

Based on our data distribution, the size of the Greater Houston Area, the layout of the city's roadway infrastructure, and the findings of Potter et al. (Potter, Koch, Oswalt, & Iannone, 2016), we used ArcGIS' Generate Tessellation function to generate a grid of 10 square kilometer (km²) hexagons. In our grid, approximately 2,000 hexagons overlaid the Greater Houston Area, and approximately 500 of those contained more than three Tweets per day across the study period.

Twitter Data Acquisition

Because we wanted to find a metric of Twitter activity intensity that would consider both users' ability to Tweet *and* their inclination to Tweet, we decided to explore each area's deviation from a defined steady state of Twitter intensity. We defined our steady state using the period of August 7 to August 15, 2017, two weeks before Hurricane Harvey made landfall the night of August 25, 2017. Harvey developed into a tropical storm on August 17, 2017, so this time frame is distant enough from the storm's inception to not be impacted by the storm development and yet close enough to reflect an accurate steady state.

Our Lab has been continuously collecting Tweets through the Twitter Streaming Application Program Interface (API). We are accessing the Twitter API through the use of the python package tweepy. Twitter only provides 1% of real-time global Tweets; however, the 1% can be filtered by certain specifications. We filtered our streamed Tweets to only include geo-referenced data. From these Tweets, we downloaded the associated User ID, originating latitude and longitude, username, and the text of the Tweet. We used this data to extract Tweets generated within the Greater Houston Area during our determined steady state time frame. We then manually filtered out Tweets made by Twitter bots, which are accounts that Tweet pre-programmed messages at automated times throughout the day, through manual inspection of suspicious Tweets and Twitter Users. Identifying Tweets that do not originate with a real person is a continuing problem in crisis informatics; however, we believe that the bots will not change their behavior overtly from the steady state to the crisis, and their impact will be mitigated through our analytical methods.

We repeated this process for the days following August 21, which is two days prior to when Harvey developed into a hurricane. As such, Tweets from August 21 to August 31, 2017 were used to define our perturbed state.

Hazard Assessment Data

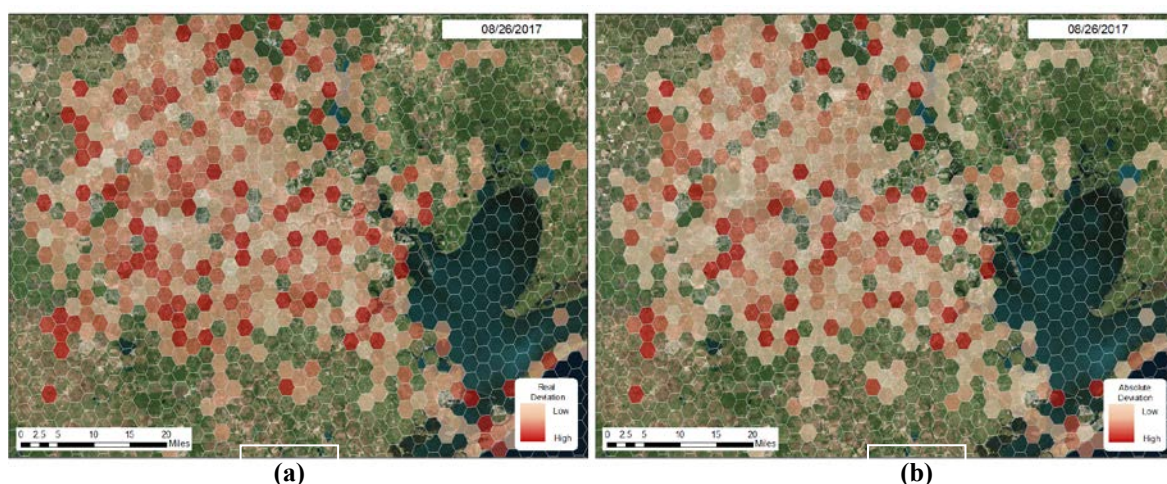
We analyzed five different metrics of hazard indicators related to weather and flooding phenomena in order to avoid the confounding factor of economic influences. Primarily, we used data from the FEMA Building Level Damage Assessments, referred to as FEMA Damage Assessments (The Federal Emergency Management Agency, 2017). These preliminary damage assessments were performed as soon as was safely possible following the hurricane and are available as point data that represent buildings that have been classified as having “No Damage”, being “Affected”, having “Minor Damage”, having “Major Damage”, and being “Destroyed” (The Federal Emergency Management Agency, 2016). We reinterpreted this range into a continuous numeric scale from 0, representing “No Damage”, to 5, representing “Destroyed”. This more subjective scale would represent the overall potential danger to a building’s inhabitants during the hurricane more judiciously than a monetized scale, which might amplify damage to richer communities. Based on the threat to human life represented by flooding, we also used the interpreted flood risk from Risk Management Systems (RMS) for the Greater Houston Area. This flood risk was modeled using elevation, rainfall, and data from the National Oceanic and Atmospheric Administration (Young, 2017). This data was rasterized in ArcMap with a cell size of approximately 30 meters. In another interpretation of flood risk, we mapped historical high water mark data from previous floods provided by the USGS (United States Geological Survey, 2017). Lastly, as one of the greatest dangers with flooding is due to flooded roads, we downloaded the volunteered geographic information (VGI) from the most popular crowdsourced map that identified road closures during Hurricane Harvey (*Flooded streets due to #Harvey*, 2017). Hurricane Harvey was the wettest tropical cyclone in both Texas and US history, and the sudden, massive, unexpected rainfall flooded homes while simultaneously cutting off evacuation routes. We theorized that areas with damaged or blocked escape routes would have a much greater need of emergency rescue. Lastly, we determined whether or not each hexagon was mostly within or without the hurricane evacuation area defined by the Houston Emergency Operations Center (Houston-Galveston Area Council, 2017).

Hexagon Data Attribution

We developed a model in ArcMap that would summarize the Twitter activity statistics for each day within each hexagon. We used the sum total Tweet counts for each day of the steady state to determine the average Tweet count and the standard deviation of that Tweet count for each hexagon. We normalized the Tweet counts within each hexagon for each day of the perturbed state using the standard deviation and average of the Tweet counts in each hexagon across the defined steady state.

We then used ArcMap’s Field Calculator to create two sets of normalized Twitter counts in the attribution field of each hexagon. The first was the real normalized data, which ranged from -1.8 (decreased Twitter activity) to 8.3 (greatly increased Twitter activity). The data is referred to as the real deviation because it preserves the direction, positive or negative, of the deviation. For the second set of Twitter counts, we took the absolute value of this normalized data, which ranged from 0 (no deviation from steady state) to 8.3 (the highest deviation from steady state). We show the difference in deviation spread between the real and absolute normalized Twitter activity data for the day following Hurricane Harvey’s landfall in **Figures 1a and 1b**.

Figure 1. A map of the Greater Houston Area overlaid with a grid of 10 km² hexagons. Each hexagon is colored according to the real (Figure 1a) or absolute (Figure 1b) standard deviation of georeferenced Tweet counts recorded on August 26, 2017 when normalized to the steady state period.



For the hazard assignment, we used a second ArcMap model to classify each hexagon based on:

- 1) The highest classification of any singular FEMA Damage Assessments located within the hexagon.
- 2) The average, maximum, and minimum of the RMS flood hazard cells in the hexagon;
- 3) The average, maximum, and minimum height of the high-water marks in the hexagon;
- 4) The presence (1) or absence (0) of a flooded road in the hexagon; and,
- 5) If the hexagon was mostly inside the evacuation zone (1) or mostly outside of it (0).

ANALYSIS

We determined the Kendall and Spearman rank correlation coefficients for two pairs of variables: 1) the correlation between the highest FEMA Damage Assessment in each hexagon and the real normalized deviation in Twitter intensity from steady state, and 2) the correlation between the highest FEMA Damage Assessment and the absolute normalized deviation in Twitter intensity. Following Kryvasheyev et al., we compare the Kendall and Spearman rank correlation coefficients for the damage assessments and the absolute and real Twitter activity deviations from steady state in **Table 1**.

Table 1. Kendall and Spearman Rank Correlation Coefficients for the maximum FEMA Assessment in each hexagon and the normalized Twitter Activity from August 21 to August 31, 2017.

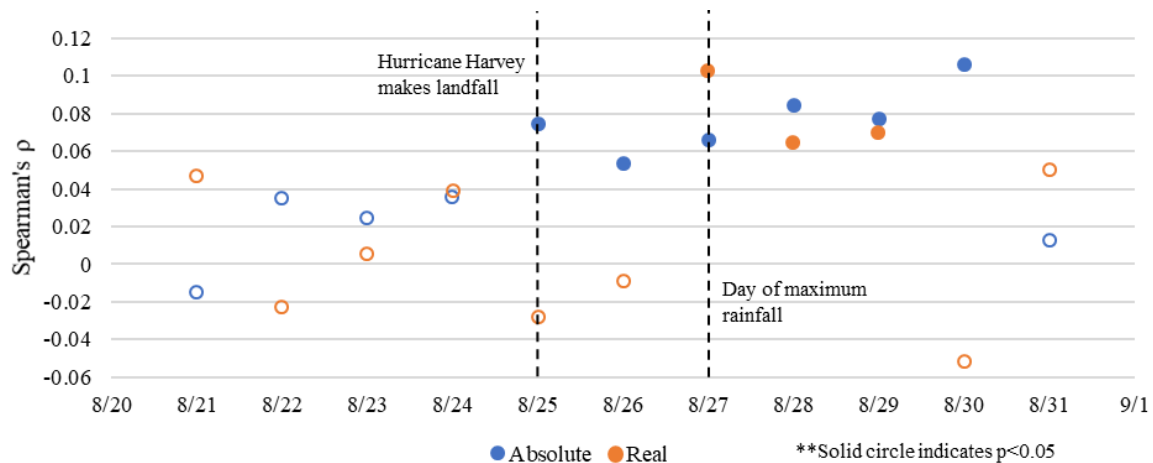
Date	Rainfall (in.)	Absolute Deviation		Real Deviation	
		Spearman's ρ	Kendall's τ	Spearman's ρ	Kendall's τ
21-Aug	0	-0.02	-0.02	0.06	0.05
22-Aug	0	0.04	0.03	-0.03	-0.02
23-Aug	0	0.03	0.02	0.01	0.01
24-Aug	0	0.04	0.04	0.05	0.04
25-Aug	0.28	0.09*	0.07*	-0.03	-0.03
26-Aug	3.83	0.07*	0.05	-0.01	-0.01
27-Aug	11.82	0.08*	0.07*	0.13**	0.10**
28-Aug	1.74	0.10**	0.08*	0.08*	0.06*
29-Aug	2.4	0.09*	0.08*	0.09*	0.07*
30-Aug	0.15	0.13**	0.11**	-0.07	-0.05
31-Aug	0	0.01	0.01	0.07	0.05

* indicates $p < 0.05$; ** indicates $p < 0.01$

FINDINGS

Our first research question concerned identifying the strength of correlation between Twitter data and hurricane hazard at a spatial scale smaller than ZCTAs. The correlation determined in our research is weaker than that found by Kryvasheyev et al. during Hurricane Sandy, but comparable with the correlations found by Kryvasheyev et al. during flooding events. We also used a much smaller spatial area, so a weaker correlation can be expected. Our correlation results are still significant and the real deviation results follow the same general trends seen in prior research. It should be noted that, although the hurricane made landfall the night of August 25th, the hurricane's movements resulted in the heaviest rainfall hitting Houston on August 27th. The fact that the correlations peak on that day—and that they more closely match the prior correlations found during flooding—are likely because the persistent flooding was the biggest danger for the city.

The spatiotemporal behavior of the correlation coefficients for each metric and their significance across the perturbed period is shown in **Figure 2**. The correlation between damage and Twitter activity is significant from August 27th to August 29th for both metrics, and the correlation between real deviation and damage is much stronger on the day of the greatest rainfall. However, the correlation between absolute deviation and damage is significant for the days after the hurricane made landfall (the 25th and the 26th), and for one additional day after (the 30th). Both relationships peak with a Spearman's ρ of 0.13 with a p-value of 0.001, but these peaks occur on disparate days: the real deviation peaks on the day of the most damage, and the absolute deviation peaks later. The absolute deviation is also the stronger relationship for every day except that of maximal flooding.

Figure 2. Spearman's ρ for Twitter Activity Deviation and FEMA Damage for the Perturbed Period

DISCUSSION AND CONTRIBUTIONS

We sought to assess the utility of social media in emergency response by correlating Twitter activity and the extent of infrastructural damage, and to define that correlation through a metric that can encompass areas that have lost the ability to Tweet. To our knowledge, this is the first analysis to supply a method for Twitter data analysis that utilizes a deviation from a steady state to address potential data shadows, as determined by Shelton et al (Shelton et al., 2014). Our results show that, in confirmation of prior literature, Twitter activity spikes are significantly correlated with damage on the day of maximal rainfall in Houston (Kryvasheyeu et al., 2016); they also show that absolute activity deviation is a more significant metric before and after that day. This shows that the absolute change in Twitter activity and the infrastructural damage can provide a secondary hazard indicator to address more persistent informational needs during an evolving natural disaster.

A potential flaw in this analysis is that the decrease in activity primarily indicates areas that people have evacuated, or commercial sectors of the city that have closed in anticipation of danger. In terms of evacuation, the distribution of hexagons that increased and decreased in Twitter activity does not appear to have any correlation with being inside the evacuation zone, and the city of Houston did not have mandatory evacuation orders prior to landfall. The hexagonal grid should minimize the partitioning of our study area by residential and commercial sectors. It is therefore likely that Tweeting behavior deviated from its steady state due to reasons beyond mobility. Future research needs to focus on determining at what spatial and temporal scales the identified relationship is strongest, and at which spatiotemporal scales the relationship breaks down. The effect of broad-scale, enforced evacuation on the correlation needs to be identified and removed such that emergency responders do not look for victims in rows of empty houses.

This analysis needs to be performed on other hurricanes and natural disasters. Based on preliminary further research and the theories identified above, we believe that the potential uses and relevance of our metrics could be dependent on the economic demographics and the severity of the hurricane. Hurricanes that do not heavily impact the energy infrastructure of the city may do very little to a population's ability and inclination to Tweet, and prior research has shown that, in some areas, even power outages do not heavily influence social media access (Jennex, 2012). The situations in which the absolute deviation would be useful needs to be identified through further disaster analysis.

Our research shows that the relationship between human behavior and damage across a disaster's duration is not only because people Tweet more in the face of danger; rather, people are changing their Twitter habits in different ways based on their inclination and their environment. Some are Tweeting more, sharing information with family members or trying to spread awareness about hurricane dangers; some are Tweeting less, saving their phone battery for 911 calls, or are without power. This modulation is the human behavior signal that should be additionally monitored when advising emergency response on an extended timeframe.

CONCLUSIONS AND IMPACTS

There is currently a lack of confirmation that the emergencies identified through Twitter are the most critical emergencies for responders to address, especially in terms of people who become unable to access their phones or are without power. In order for responders to triage the emergencies during a natural disaster, they need to

know the emergency's location, severity, and magnitude throughout the emergency. Existing social media emergency response tools fail this by neglecting the possibility of severe emergencies occurring in areas where Tweets are not occurring (Tapia & Moore, 2014). Through our research, we have shown that absolute Twitter activity deviation from a steady state can be used to identify areas that are experiencing more infrastructural damage during a hurricane. Our method does not ignore areas with fewer Tweets, but rather fully incorporates those areas into its hazard association, and is more particularly useful on days following landfall or after the biggest disaster event.

For emergency responders, the identification of the existence of an emergency through social media is not sufficient; 911 calls and aid requests already inundate emergency agencies with emergency identifications during a major disaster. To be truly useful to responders, social media must be able to point responders in the direction of the greatest and most urgent damage in near-real-time. For tools using only spikes in Twitter activity to identify areas of need, areas without the ability to Tweet are invisible. These tools may be stronger on the day of maximum perturbation, but as the emergency evolves, response information needs change and other populations are more at risk. We need to be able to "hear" distress among these populations as well. We show this is possible by ascertaining the value of listening to where Tweets have stopped occurring. Our paper highlights the importance of considering the presence and absence of data in social media analysis tools designed for use by emergency responders. Ultimately, we find that the most effective analysis tools using humans as sensors must be able to incorporate the sound of humans in danger—and the sound of silence.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1760645. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Abramson, D. M., Grattan, L. M., Mayer, B., Colten, C. E., Arosemena, F. A., Bedimo-Rung, A., & Lichtveld, M. (2014). The Resilience Activation Framework: a Conceptual Model of How Access to Social Resources Promotes Adaptation and Rapid Recovery in Post-disaster Settings. *Journal of Behavioral Health Services and Research*, 42(1), 42–57. <https://doi.org/10.1007/s11414-014-9410-2>
- Amadeo, K. (2017). Hurricane Harvey Facts, Damage and Costs What Made Harvey So Devastating. *The Balance*, 1–7. Retrieved from <https://www.thebalance.com/hurricane-harvey-facts-damage-costs-4150087>
- Caragea, C., Squicciarini, A., Stehle, S., Neppalli, K., & Tapia, A. (2014). Mapping Moods: Geo-Mapped Sentiment Analysis During Hurricane Sandy. *Proceedings of the 11th International ISCRAM Conference*, 642–651.
- Choe, S., Park, J., Han, S., Park, J., & Yun, H. (2017). A study on the real-time management and monitoring process for recovery resources using Internet of Things. *International Research Journal of Engineering and Technology (IRJET)*, 4(3), 2634–2639. Retrieved from <https://irjet.net/archives/V4/i3/IRJET-V4I3673.pdf>
- Conrado, S. P., Neville, K., Woodworth, S., & O'Riordan, S. (2016). Managing social media uncertainty to support the decision making process during Emergencies. *Journal of Decision Systems*, 25(October), 171–181. <https://doi.org/10.1080/12460125.2016.1187396>
- Coughlan De Perez, E., Van Den Hurk, B., Van Aalst, M. K., Jongman, B., Klose, T., & Suarez, P. (2015). Forecast-based financing: An approach for catalyzing humanitarian action based on extreme weather and climate forecasts. *Natural Hazards and Earth System Sciences*, 15(4), 895–904. <https://doi.org/10.5194/nhess-15-895-2015>
- Flooded streets due to #Harvey.* (2017). *Google My Maps*. Retrieved from <https://www.google.com/maps/d/u/0/viewer?mid=1Nzjiw9FHDJPHgNJiVMSHWhBQQd0&hl=en&ll=29.827383837787504%2C-95.27925249999998&z=8>
- Fugate, C. (2017). FCC and FEMA: How to Communicate Before, During and After a Major Disaster. *The Federal Emergency Management Agency Blog*, 1–5. Retrieved from <https://www.fema.gov/blog/2011-09-21/fcc-fema-how-communicate-during-after-major-disaster>
- Hiltz, S., Pfaff, M., Plotnick, L., Shih, P., Ashktorab, Z., Brown, C., ... Culotta, A. (2014). Tweedr: Mining Twitter to Inform Disaster Response. *Proceedings of the 11th International ISCRAM Conference*, 11(May), 354–358. Retrieved from <http://www.iscram.org/legacy/ISCRAM2014/papers/p142.pdf>

- Houston-Galveston Area Council. (2017). Brazoria, Chambers, Galveston, Harris and Matagorda Hurricane Evacuation Zip-Zones. *Houston Emergency Operations Center*, 1.
- Imran, M., Castillo, C., Diaz, F., & Vieweg, S. (2015). Processing Social Media Messages in Mass Emergency. *ACM Computing Surveys*, 47(4), 1–38. <https://doi.org/10.1145/2771588>
- Jennex, M. E. (2012). Social Media – Viable for Crisis Response? *International Journal of Information Systems for Crisis Response and Management*, 4(2), 53–67. <https://doi.org/10.4018/jiscrm.2012040104>
- Kryvasheyev, Y., Chen, H., Moro, E., Van Hentenryck, P., & Cebrian, M. (2015). Performance of Social Network Sensors during Hurricane Sandy. *PLOS ONE*, 10(2), e0117288. <https://doi.org/10.1371/journal.pone.0117288>
- Kryvasheyev, Y., Chen, H., Obradovich, N., Moro, E., Van Hentenryck, P., Fowler, J., & Cebrian, M. (2016). Rapid assessment of disaster damage using social media activity. *Science Advances*, 2(3), 1–11. <https://doi.org/10.1126/sciadv.1500779>
- Lam, L. (2017). 15 Billion-Dollar Disasters Have Impacted the U . S . This Year ; 2017 Tied for Second-Most All-Time. *The Weather Channel*, 1–4. Retrieved from <https://weather.com/news/weather/news/2017-10-17-billion-dollar-weather-disasters-united-states-record-pace>
- Lu, X., & Brelsford, C. (2015). Network Structure and Community Evolution on Twitter: Human Behavior Change in Response to the 2011 Japanese Earthquake and Tsunami. *Scientific Reports*, 4(1), 6773. <https://doi.org/10.1038/srep06773>
- Mason, A. M., Drew, S., & Weaver, D. (2017). Managing Crisis-induced uncertainty: First responder experiences from the 2011 Joplin-Duquesne Tornado. *International Journal of Disaster Risk Reduction*, 23(April), 231–237. <https://doi.org/10.1016/j.ijdr.2017.04.012>
- National Center for Environmental Information. (2017). Calculating the Cost of Weather and Climate Disasters. Retrieved from <https://www.ncei.noaa.gov/news/calculating-cost-weather-and-climate-disasters>
- Potter, K. M., Koch, F. H., Oswalt, C. M., & Iannone, B. V. (2016). Data, data everywhere: detecting spatial patterns in fine-scale ecological information collected across a continent. *Landscape Ecology*, 31(1), 67–84. <https://doi.org/10.1007/s10980-015-0295-0>
- Reuter, C., Hughes, A. L., & Kaufhold, M.-A. (2018). Social Media in Crisis Management: An Evaluation and Analysis of Crisis Informatics Research. *International Journal on Human-Computer Interaction (IJHCI)*, 34(4), accepted. <https://doi.org/10.1080/10447318.2018.1427832>
- Sakaki, T., Okazaki, M., & Matsuo, Y. (2010a). Earthquake shakes Twitter users. In *Proceedings of the 19th international conference on World wide web - WWW '10* (p. 851). New York, New York, USA: ACM Press. <https://doi.org/10.1145/1772690.1772777>
- Sakaki, T., Okazaki, M., & Matsuo, Y. (2010b). Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. *Proceedings of the 19th International Conference on World Wide Web*, 851–860. <https://doi.org/10.1145/1772690.1772777>
- Shelton, T., Poorthuis, A., Graham, M., & Zook, M. (2014). Mapping the data shadows of Hurricane Sandy: Uncovering the sociospatial dimensions of “big data.” *Geoforum*, 52, 167–179. <https://doi.org/10.1016/j.geoforum.2014.01.006>
- Shklovski, I., Palen, L., & Sutton, J. (2008). Finding community through information and communication technology in disaster response. *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*. *ACM*, 127–136. <https://doi.org/10.1145/1460563.1460584>
- Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, R. M. (2014). Rumors, False Flags, and Digital Vigilantes: Misinformation on Twitter after the 2013 Boston Marathon Bombing. *iConference 2014 Proceedings*. <https://doi.org/10.9776/14308>
- Tapia, A. H., & Moore, K. (2014). Good Enough is Good Enough: Overcoming Disaster Response Organizations’ Slow Social Media Data Adoption. *Computer Supported Cooperative Work: CSCW: An International Journal*, 23(4–6), 483–512. <https://doi.org/10.1007/s10606-014-9206-1>
- The Federal Emergency Management Agency. (2016). Damage assessment operations manual - A guide to assessing damage and impact, 121. Retrieved from <http://www.fema.gov/media-library-data/1459972926996-a31eb90a2741e86699ef34ce2069663a/PDAManualFinal6.pdf>
- The Federal Emergency Management Agency. (2017). FEMA Damage Assessments (MapServer). Retrieved

from https://gis.fema.gov/arcgis/rest/services/FEMA/FEMA_Damage_Assessments/MapServer

- Tien, I., Musaev, A., Benas, D., Ghadi, A., Goodman, S., & Pu, C. (2016). Detection of damage and failure events of critical public infrastructure using social sensor big data. *IoTBD 2016 - Proceedings of the International Conference on Internet of Things and Big Data*, (IoTBD), 435–440. <https://doi.org/10.5220/0005932104350440>
- United States Geological Survey. (2017). Harvey Aug 2017 - High-water Mark Data. Retrieved January 1, 2017, from <https://stn.wim.usgs.gov/fev/#HarveyAug2017>
- Wang, Q., & Taylor, J. E. (2014). Quantifying human mobility perturbation and resilience in hurricane sandy. *PLoS ONE*, 9(11), 1–5. <https://doi.org/10.1371/journal.pone.0112608>
- Young, M. (2017). Harvey Now Driving Catastrophic Flooding Across Houston Metropolitan Region. *Risk Management Systems Blog*, 1–5. Retrieved from <http://www.rms.com/blog/2017/08/30/harvey-now-driving-catastrophic-flooding-across-houston-metropolitan-region/>
- Zeng, L., Starbird, K., & Spiro, E. S. (2016). #Unconfirmed: Classifying Rumor Stance in Crisis-Related Social Media Messages. *Proceedings of the International AAAI Conference on Weblogs and Social Media*, 892(ICWSM), 747–750. Retrieved from <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM16/paper/viewFile/13075/12845>