

THE IMPACT OF GEOGRAPHIC SCALE ON IDENTIFYING DIFFERENT SOCIAL MEDIA BEHAVIOR EXTREMES IN CRISIS RESEARCH

Rachel Samuels
John E. Taylor

Department of Civil and Environmental Engineering
Georgia Institute of Technology
760 Atlantic Drive
Atlanta, GA 30318, USA

ABSTRACT

Our relationship with technology is constantly evolving, and that relationship is adapting even more quickly when faced with disaster. Understanding how to utilize human interactions with technology and the limitations of those interactions will be a crucial building block to contextualizing crisis data. The impact of scale on behavioral change analyses is an unexplored yet necessary facet of our ability to identify relative severities of crisis situations, magnitudes of localized crises, and total durations of disaster impacts. In order to analyze the impact of increasing scale on the identification of extreme behaviors, we aggregated Twitter data from Houston, Texas circa Hurricane Harvey across a wide range of scales. We found inversely related power law relationships between the identification of sharp Twitter activity bursts and sharp activity drop-offs. The relationships between these variables indicate the direct, definable dependence of social media aggregation analyses on the scale at which they are performed.

1 INTRODUCTION

As the supply of data from humans-as-sensors continues to increase, understanding individual data streams in the context of our multi-layered and multi-networked society is becoming more difficult. Social media is increasingly looked to as a potential source of additional information in the notoriously information-scarce environment of crises (Reuter and Kaufhold 2018). The crisis informatics field has continued to flourish and expand alongside the seemingly ever-increasing quantities of available social media data and methods of analyzing that data. As a result, the scope of applications for social media during crises has expanded to include event detection (Sakaki et al. 2010), resource availability and need (Choe et al. 2017), and mobility monitoring (Wang and Taylor 2016a). Analytical methods range from sieving individual posts for information (Hiltz et al. 2014) to analyzing geographic changes in sentiment and behavior to inform gestalt-level decisions (Jongman et al. 2015; Kryvasheyev et al. 2016).

The big data revolution has opened a vast area of possibilities for crisis response (Qadir et al. 2016). One of the greatest strengths of the field is the diversity it contains, which has resulted in a wide range of available techniques for processing this ever-increasing and changing pool of data. Applications from the field are being used currently by international aid organizations (Imran et al. 2014), and strides are being made for local response implementations of social media analysis as well (Kropczynski et al. 2018) despite initial resistance (Tapia and Moore 2014). Social media data analytics techniques range from convolutional neural networks (Caragea et al. 2016) to latent Dirichlet Allocation topic models that combine spatial, temporal, and semantic data to identify disasters (Wang and Taylor 2019). That said, the range of diversity of applications and methods can also impede the process of building a solid foundation. Researchers both external and internal to crisis informatics have noted criticism of social media applications' limits with respect to data bias, social inequality, and lack of confirmed validity (Imran et al. 2015; Jiang 2018).

More data is available to us; however, big data is not complete data. There has been a consistent call for us to critically interrogate the assumptions and capabilities of big data in the context of our political and urban usage (Boyd and Crawford 2012). As the reach and amount of available data increase, holes in that data become both less obvious due to the existing volume and yet more harmful due to the increasing prevalence of that data's use (Morstatter and Liu 2017). Social media, especially when used for crisis response, is not exempt from this call. This is especially true in the case of crisis response, where information availability alone can tip the scales of resource distribution. To ensure more equitable and intelligent use of social media data in crisis response, researchers need to understand the social, spatial, and sociospatial limitations of that data (Jessop et al. 2008). One critical piece of that understanding is understanding the geographic scale at which social media data is capable of identifying disasters and how much information is gained or lost by varying that geographic scale. This is necessary for accurate implementation of our applications and recognition of any dangers of utilizing non-uniform aggregation scales, such as census tracts (Jelinski and Wu 1996; Saib et al. 2014).

The capacity for social media to identify disasters is predicated on both the scale at which disasters are impacting people and the scale at which people are responding to that impact. Disasters are social; a key defining factor of a disaster's scale is how greatly it impacts society's normal functions (Guan and Chen 2014). Most of crisis analytics work operates on the assumption that, because disasters impact people, people will change their behaviors, such as their mobility, purchasing habits, and, yes, social media usage. Social media usage behavior has been defined through posting content, user interactions, sentiment analysis, and posting frequency. The scale at which these behavior changes are identifiable, and how social media usage behavior changes vary across geographic space, is unknown. Previous research has identified a positive correlation between increasing amounts of posting frequency and hurricane damage (Kryvasheyev et al. 2016); however, recent research has also identified a correlation between the magnitude of change, accounting for both increases and decreases in posting activity, and hurricane damage (Samuels et al. 2018). We theorized that a variety of both increasing and decreasing posting behavior is happening at smaller, neighborhood-scales, and thus the identification of different posting behavior trends would be affected by the scale of analysis. Although scale has been identified as a potential impactor on analysis conclusions for crisis data, its effect on the recognition of extremely greater or extremely diminished posting activity changes as crisis identifiers has not been explored nor defined (Knüsel et al. 2019; Shelton et al. 2014). Our conclusions are only as good as our data, and the analyses that utilize other types of social media attributes—be they content, spatiotemporal clusters, or networks—need to incorporate the impact of geographical scale and place into our interpretations of that data.

Our goal within the scope of this paper is to identify the effect of increasing scale on the identification of both extremely high amounts and extremely low amounts of Twitter activity, as compared to a defined baseline. As such, our research question is: what is the effect of increasing scale on the identification of crisis-induced extreme changes in social media activity? In answering this, we contribute insight into not only how well social media posts can represent a change in normal societal function during a disaster, but also into the spatial variances in either accessibility to social media or inclination to post across a geographical space.

To address this question, we chose to focus our efforts on the city of Houston, Texas circa Hurricane Harvey. As the largest city on the Gulf Coast of the U.S. (World Population Review 2018), Houston had a large hurricane-affected population that, based on our analysis of Gulf-based city Tweeting behavior, also has a substantial number of affected Twitter-users.

2 METHODS

2.1 Twitter Data

All of the geolocated Twitter data for the greater metropolitan area of Houston for seven weeks prior to and one week following Hurricane Harvey's landfall were streamed through the Twitter API (Wang and Taylor

2016b). Hurricane Harvey made its first landfall in Houston on August 25th, 2017 in the evening; the hurricane then pivoted and returned on August 27th to deposit torrential, record-breaking rains. For our analysis, we defined our perturbed state—the period of time during which non-normal behavior would be expected—as one day prior to the first landfall through the week following landfall (August 24th-September 1st). To identify non-normal behavior, we needed to select a steady state to act as our baseline for “normal” behavior. We defined this steady state as the five-week period from July 11th and August 16th, following prior research describing the time period length necessary to generate a sufficiently stable analysis (Toepke 2018); a longer period would increase the influence of both seasonality and population flux. The steady state behavior has a left-leaning log normal distribution, matching prior findings, and is explored more in Appendix A. We also allowed for a transitional state, during which the hurricane would broadly impact Twitter posting behavior through anticipation of harm but not through actual hurricane damages or events. This state is defined as the period from the day Harvey was identified as a tropical storm through the day before our perturbed state begins (August 17th-August 23rd).

It should be noted that Houston experienced the most infrastructural damage and flooding on August 27th and not when Hurricane Harvey first made landfall. As such, some of our discussion focuses on behaviors identified on August 27th.

The sets of steady state and perturbed state Tweets were temporally aggregated by day, transformed into individual points through ArcGIS, and plotted using their latitude and longitude attribute information in ArcGIS. The 2010 census tract shapefiles were downloaded from the Harris County GIS data portal.

2.2 Population Data

The census data and census tracts are not at a sufficiently granular resolution to understand the nuances of neighborhood-scale behavior during a crisis. The tracts further from the city center can be as large as 600 km², so we need to find a method of increasing the resolution of the population data. The geographic information systems field has historically utilized National Land Cover Database (NLCD) 2011 data to increase the granularity of the census data with substantial accuracy (Bian and Wilmot 2017; Reibel and Agrawal 2007). Although the NLCD is updated every five years, the 2016 data was not yet available for download at the time of analysis. Additionally, Grubestic et al. found substantial flaws in comparing datasets more than a few years apart (Grubestic and Matisziw 2006). The 2011 NLCD would most closely match the 2010 census data we used as population values. The NLCD contains a raster file with 30 by 30-meter cells that have been classified, through satellite imagery, as one of 16 classes. The classification includes four classes of developed land: open space, low intensity, medium intensity, and high intensity. We downloaded the 2011 dataset and extracted the raster cells that were classified as any type of developed land and were located within the greater metropolitan Houston Area. Using ArcGIS' Raster to Point function, we then transformed each of the raster cells into points located at the center of each cell and spatially joined these points by count into the census tracts for Houston. Using the counts of each type of NLCD class and the population record for each census tract, we performed a multiple linear regression analysis to determine the contributing coefficients of each type of land type with respect to population. The results of the regression analysis are presented in Table 1, with an adjusted R-squared of 0.8317 and a model p-value of <0.001.

We used these model coefficients to determine the weighted averages of each land type within each census tract. As shown, the areas of very intense development have a substantial and significant negative contribution to the residential population of the Houston census tracts. Bian and Wilmot encountered similar results in New Orleans, Louisiana in a study using the same technique to study disadvantaged populations impacted by Hurricane Katrina (Bian and Wilmot 2017). Following ground proofing techniques, they assigned a positive but very small assigned coefficient to the highly developed areas of the city. For the purposes of our population disaggregation, we followed their example and a very small assigned coefficient to areas of low and high development, utilizing the ratio of each model coefficient to

the others. These assigned coefficients were used in the redistribution of the Houston population. The details of the produced model are included in Table 1.

Table 1: Multiple linear regression results for the NLCD land class types and the census data.

NLCD class	Model Coefficient	STD Error	p-value	Assigned Coefficient
Open area	0.30	0.03	<0.001***	0.23
Low intensity	-0.11	0.07	0.123	0.18
Medium intensity	4.50	0.10	<0.001***	0.58
High intensity	-2.34	0.12	<0.001***	0.01

The assigned coefficients were employed in Equation (1) to determine a population quantity to assign to each of the NLCD point shapefiles, based on the distribution of points in each census tract and that tract's population.

$$Pop_{cell_i} = \frac{WA_i * Pop_{totaltract}}{\sum WA_{i \rightarrow l} * Tract_{count_i}} \tag{1}$$

Pop_{cell_i} is the population represented by a single NLCD raster cell point of NLCD type i within a specific census tract; the NLCD types of “Open area”, “Low intensity”, “Medium intensity”, and “High intensity” are represented as $i \rightarrow l$; WA_i is the weighted average for the specific land type; $\sum WA_{i \rightarrow l}$ is the sum of the weighted averages for each land type; $Pop_{totaltract}$ is the total population within the census tract in which the point is located; and $Tract_{count_i}$ is the total number of points of type i within the specific census tract. This was repeated for each of the NLCD types. The resulting dataset was a grid of point files spaced 30 meters apart that were assigned a reasonably accurate population total and could be aggregated into a series of shapefile “nets” spread across the city of Houston.

2.3 Spatial Nets

With the social media data temporally aggregated and mapped and the population data spatially disaggregated, we designed twelve spatial nets to catch the population data and the Twitter stream for each day. We generated the spatial nets of equally sized and shaped hexagonal polygons through ArcGIS’ Generate Tessellation function. Hexagons are better suited for tiling large geospatial areas because they reduce edge effects that can be exacerbated by intersecting rectangles and are more scalable on a curved surface like the globe (Carr et al. 1992; Polisciuc et al. 2016). The twelve hexagonal nets consist of hexagons that have square areas of, respectively, 0.25 km², 0.5 km², 0.75km², 1 km², 2 km², 5 km², 10 km², 15 km², 20 km², 35 km², 50 km², and 80 km². We chose these by beginning with 10 km², the average approximate size of the Houston census tracts, and incorporating logarithmically larger and smaller scales until the analysis appeared to only be catching individual behavior (0.25 km²) or we reached the scale at which information provided very little actionable or useful information in terms of aid distribution or the presence of local disasters (80 km²). A comparison between the 1km² and 80km² nets is shown in Figure 1 for a scalar reference.

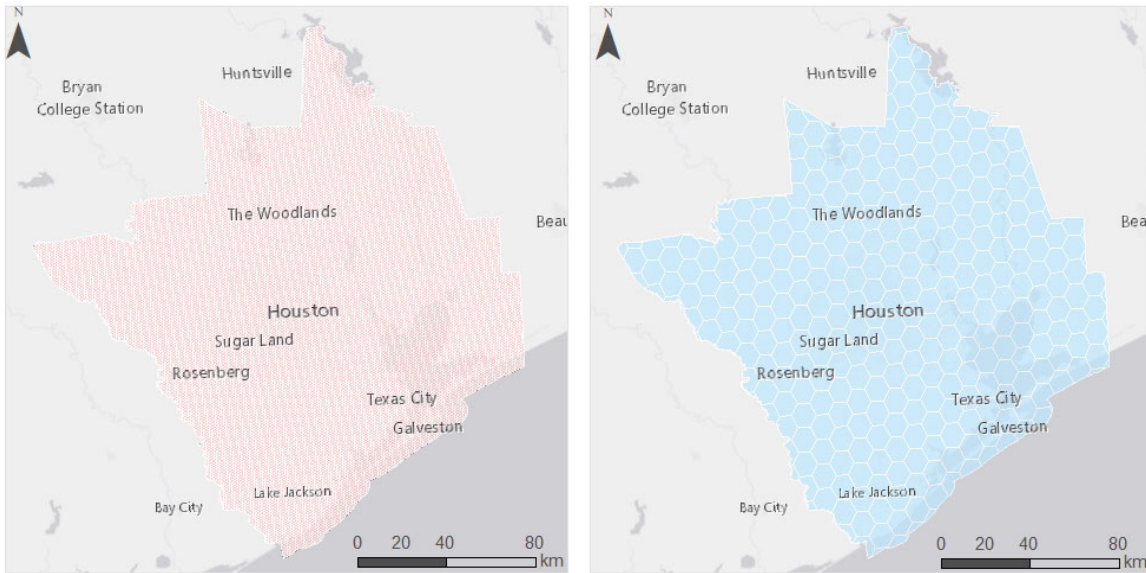


Figure 1: (Left) Net of 1 km² hexagons across Houston. (Right) Net of 80 km² hexagons across Houston.

We summed the number of Twitter posts for each day and the population values of each of the NLCD points within each polygon of each net. Following the methods listed by Kryvasheyeu et al. (2016), areas that did not contain at least one Tweet per day during either the steady state or perturbed state were removed.

2.4 Analytical Methods

In order to understand the distribution of perturbed state Twitter posting counts that were either much higher or much lower than the “normal” behavior observed in the steady state, we used cumulative distribution functions (CDFs) to compare perturbed state Twitter activity to the steady state Twitter activity. We used the steady state post counts for each of the spatial nets to generate a series of CDFs. These CDFs represented the distribution of Twitter activity for a single area across each day of the steady state period. For instance, for a given area *A* within the 10 km² net, we created a CDF from all of area *A*’s steady state Twitter activity counts by day. We then took the perturbed state Twitter activity on a given day, such as August 27th, and used the generated steady state CDF of activity to determine what percentage of steady state days had produced less than the number of Twitter posts produced on August 27th in area *A*. A result of 0.90 would indicate that the perturbed state activity on August 27th was higher than the activity produced on 90% of the days in the steady state, and a result of 0.10 would indicate that the perturbed state activity was only higher than 10% of days in the steady state.

We used each CDF to assess, for each area and each day of the perturbed state, the cumulative likelihood of observing a certain number of Tweets in that area on that day. We categorized this likelihood as being normal, non-normal, or extreme. Although, as stated, we are most interested in extreme values, we included an analysis of non-normal to provide a reference for the impact of how the threshold of “extreme” amounts of activity is defined. Using empirical rule values, non-normal behavior was defined as being less than 16% of steady state values or greater than 84% of steady state values. Extreme behavior was defined as being less than 5% of steady state values or greater than 95% of steady state values. To identify the effect of scale on observing extreme (and non-normal) values and so understand the prevalence and significance of activity bursts, clustering, or drop-offs, we took the distribution of the likelihood of observing the perturbed state Twitter activity and analyzed the distributions of those probabilities across nets and days of the perturbed state.

3 RESULTS

At smaller scales (less than 10 km²), we identified a clear bimodal distribution indicating a large quantity of both extremely high and extremely low Twitter activity across each day of the perturbed state. At scales greater than 20 km², no consistent pattern of posting behavioral tendency was visible. To further interrogate this trend, we assessed the percentage of areas on the day of maximum rainfall and damage, August 27th, that exhibited non-normal or extreme Twitter posting behavior. The percentage of areas for each net that experienced Twitter activity lower than 5% and 16% of its steady state values on August 27th are portrayed in Figure 2. The percentage of areas for each net that experienced Twitter activity higher than 84% and 95% of its steady state value are portrayed in Figure 3. All identified relationships fit well ($R^2 > 0.93$) to a power law distribution, defined as $f(x) = ax^{-k}$. We additionally included the percentage of areas that exhibited normal behavior, within the central 68% of observed Twitter activity behaviors, on both figures for comparison. The power law relationship for normal activity fits less well ($R^2 = 0.77$).

Of most importance for applications, the relationship between increasing scale and the identification of extremely low behaviors is negative, while the relationship between scale and the identification of extremely high behaviors is positive. The minimum value for which this relationship holds true has not been identified. The scaling constants, a , vary between 0.1 and 1.0, but the magnitude of the power k is approximately the same (0.17) for both of the equations for the extreme values. Additionally, the inclusion of non-normal activity analysis was provided as a reference for the impact of the threshold at which the amount of activity could be interpreted as “extreme”. We see the impact of increasing the boundary for the classification from outside the central 68% to outside of the central 90% in the identification of non-normal and extreme perturbed state activity for both sets of behavior trends. Decreasing the threshold at which an observed behavior is classified as noteworthy obviously increases the number of noteworthy observations; however, this effect is slightly larger for the extremely high values, and the effect is more profound at higher scales (greater than 40 km²).

At a gestalt level, the difference between the percentage of areas identified as exhibiting extreme at the smaller scales is much larger (approximately 80% for scales less than 0.5 km²) than the percentage identified at the larger scales (approximately 53%). This relationship is represented in Appendix B. The geographic coverage of those areas, however, is quite similar due to the increased removal of areas without Twitter activity at the smaller scale.

4 DISCUSSION

First and foremost, the analyses presented herein show that the stories our data tell differ when read at different scales. We identify and define the effect of scale on the identification of extreme behaviors, and we show clearly that the effect is different for different extreme behaviors. When we assessed the distribution of the likelihood of seeing each value in the perturbed state using the empirical cumulative distribution function generated in the steady state, we found that few perturbed state values lie close to the average of the steady state, and that the proportion of areas exhibiting extreme behaviors was dependent on the scale of the analysis. Although some effect of geographical scale had been identified in previous studies (Shelton et al. 2014), the magnitude or consistency of the relationship had not been defined. We have defined consistent power law relationships between scale and the observation of extreme levels of Twitter activity. Due to the nature of this relationship, geographical scale has a very strong influence at the smallest scales (less than 5 km²) and produces drastically different data distributions when comparing between scales of very different magnitudes. Additionally, previous research had theorized that increasing scale improved the identification of different, perturbed state behavior (Chen et al. 2013). However, we show that, although increasing scale improves the identification of greater-than-normal Twitter activity, increasing scale decreases the identification of less-than-normal Twitter activity. The inverse relationship between scale and different types of perturbed state behavior can thus strongly affect analytical methods’ ability to identify activity drop-offs caused by catastrophic flooding or energy infrastructure damage. This will be critical to consider as the field continues to try to scale the applications of social media both up and down.




Figure 2: Comparison of the percentages of areas exhibiting non-normal or extremely low Twitter activity behavior between spatial nets, fit to power law distributions.




Figure 3: Comparison of the percentages of areas exhibiting non-normal or extremely high Twitter activity behavior between spatial nets, fit to power law distributions.

In our data, we have identified four strong power law relationships between aggregation scale and the number of identified areas exhibiting non-normal and extreme behavior. Power law distributions have attracted a large amount of attention in almost every field, ranging from microbiology to economics. They are nearly ubiquitous in natural systems, although the statistical confirmation of the reality of each claim has been questioned (Stumpf 2012). It is no small irony to the authors that their identification of the scale-dependence of social media analyses culminates in a distribution defined as scale-independent. We would like to note, then, that the relationship between scale of aggregation and these Twitter behaviors is what is scale-independent; the relationship itself is direct and significant.

The number of identified extremely low events decreases exponentially with increasing scale, and the number of identified extremely high events increases exponentially with increasing scale. The scalar for identification of non-normal (0.602) versus extreme (0.571) events is not very different for the lower thresholds, but the scalar for the identification of non-normal (0.176) versus extreme events (0.107) is higher with similar powers. This increased identification of non-normal behaviors at increased scales suggests the need to apply more stringent thresholds for activity marked as non-normally or extremely high at larger spatial scales, or our analyses cannot triage. Also, the many decreased values identified at the smaller scales confirms the need for more stringent methods of investigation into these areas that fall suddenly silent.

In terms of sudden silence, previous research has identified that drop-offs in Twitter activity are also correlated with damage and theorized that those drop-offs are caused by social vulnerabilities more than behavioral choices (Samuels et al. 2018). Increased scales minimize the potential for a social media analysis to identify these drop-offs as extreme events, a factor that needs to be considered and addressed in social media applications. Especially as many aggregate studies utilize zones such as census tracts and Zip Code Tabulation Areas (ZCTAs), this factor can be a serious problem. ZCTAs vary widely in size and shape; the ones within Houston, for instance, range in size from 0.16 km² to 677.20 km². This variance in size and socially constrained boundaries have been substantially critiqued in the field of GIS (Jelinski and Wu 1996; Saib et al. 2014).

Here, we can confirm the effect of size variance on analysis and warn against the use of varying sizes. Larger tracts tend to be those further from the city, with a higher percentage of populations that do not participate in Twitter. These larger areas with a smaller number of Tweets per person will be far more likely to show no or little increased activity, and thus any localized crises causing activity drop-offs will not register in these analyses. Similarly, at the smallest scales, the majority of areas exhibit extreme behaviors, so these areas' behaviors are more likely to be interpreted as exhibiting a need. On the equity side of that, the majority of small-scaled ZCTAs are in the center of our cities, where both population density and community wealth tend to be higher. We have shown that the scalar dependencies of identifying behavior changes tip the scales in favor of these urban populations.

4.1 Limitations and Future Work

4.1.1 Hurricane-specific Tweeting

Many studies in the field filter for Tweets that are directly related to the hurricane through text analysis. This limits the application of steady state versus perturbed state analysis, as no one was Tweeting about Hurricane Harvey before it formed in the Gulf. As such, using hurricane/disaster-specific Tweets was not possible.

4.1.2 Area exclusion

We excluded from analysis all areas that did not have a single Tweet across the steady state period or a single Tweet across the perturbed state. However, this resulted in a substantial amount of geographic coverage reduction in the smaller scales. We identified a logarithmic increase in excluded area with

decreasing scale. For example, 66% of Houston was excluded in the 0.25 km² net, 53% was excluded from the 1 km² net, and only 20% was excluded from the 20 km² net. Further research is necessary to investigate the effects of Twitter activity thresholds, population thresholds, and the effect of including or excluding areas and populations with a very lower Twitter representation. This information will help inform how applicable Twitter data can be to demographically-different neighborhood distinctions in behavior and will minimize the inclusion of non-participating, unrepresented populations.

4.1.3 Future work

No two disasters are the same, either in terms of the damage caused or in terms of the affected society. The generalizability of the identified power law relationships and the distributions of Twitter activity to other cities and other disasters should be investigated.

5 CONCLUSIONS

Social media analyses use changes in human behavior to identify the location, magnitude, and severity of localized crises within disaster contexts. The marks that these behavior changes leave on crisis data, however, are not consistent or uniform. Previous research in the field has identified both increases and decreases in social media activity in response to severe hurricane damage, and researchers have noticed the influence of scale on crisis informatics analyses, but the two had not been linked. Within this paper, we analyze the influence of scale on the observation of extreme social media behaviors during Hurricane Harvey. We find that increasing scales promote the identification of more areas with extremely high amounts of Twitter activity and the identification of fewer areas with extremely low amounts of Twitter activity. We also find that smaller scales are more likely to identify areas as exhibiting extreme behaviors. In the context of future analyses, researchers need to be aware that using aggregation areas with widely varying sizes can be heavily skewing their results, particularly when that skew favors bursts of activity and those in more populous, potentially more advantaged areas. Ultimately, when social media analyses are a factor influencing resource and aid distribution, we show that defining and understanding the scalar dependencies of our analyses is critical to ensuring equitable emergency response.

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.

A STEADY STATE ANALYSIS

We created CDFs for the steady state Twitter posting activity of every area within each of the spatial nets, resulting in thousands of steady state CDFs. The distributions for five randomly selected areas for two of the spatial nets (50 km² and 1 km²) are depicted in Figure 3 for comparison. Both depict left-leaning log normal distributions, with fairly substantial variation in density within the spatial nets and as compared to each other. The long right tails are also a constant characteristic; however, they are more prominent in the distributions of the 1 km² spatial net, and there is greater variability across and within the 1 km² areas.

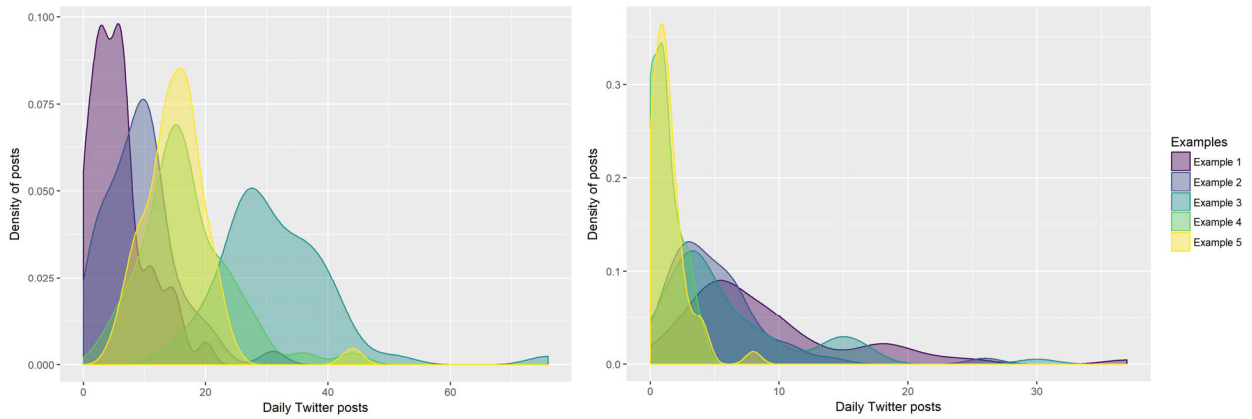


Figure 3: The steady state distributions of 5 random areas within the 50 km² areas spatial net (left) and the 1 km² areas spatial net (right).

B AREAL INCLUSION

As discussed in Section 4, a much larger area of geographic space without a Twitter presence was removed from the analysis for the smaller spatial scales (approximately 65%) than from the larger spatial scales (approximately 20%). The relationship between the geographical scale of analysis and the excluded area is depicted in Figure 4.




Figure 4: The relationship between the percentage of the total study area (the greater metropolitan area of Houston) excluded from the analysis on account of not having sufficient Twitter activity, as defined in Section 2.3, and the geographic scale at which the data was aggregated.

REFERENCES

- Bian, R. and Wilmot, C. G. 2017. "Measuring the Vulnerability of Disadvantaged Populations during Hurricane Evacuation". *Natural Hazards* 85(2):691–707.
- Boyd, D. and Crawford, K. 2012. "Critical Questions for Big Data". *Information, Communication & Society* 15(5):662–679.
- Caragea, C., Silvescu, A., and Tapia, A. 2016. "Identifying Informative Messages in Disaster Events Using Convolutional Neural Networks". In *Proceedings of the 2016 International Conference on Information Systems for Crisis Response and Management*, edited by A. Tapia, P. Antunes, V. A. Bañuls, K.A. Moore and J. Porto de Albuquerque. Rio de Janeiro, Brazil.
- Carr, D., Olsen, A., and White, D. 1992. "Hexagon Mosaic Maps for Display of Univariate and Bivariate Geographical Data". *Cartography and Geographic Information Systems* 19(4):228–236.
- Chen, C., Neal, D., and Zhou, M. 2013. "Understanding the Evolution of a Disaster: A Framework for Assessing Crisis in a System Environment (FACSE)". *Natural Hazards* 65(1):407–422.
- Choe, S., Park, J., Han, S., Park, J., and 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.
- Guan, X. and Chen, C. 2014. "Using Social Media Data To Understand And Assess Disasters". *Natural Hazards* 74(2):837–850.
- Ashktorab, Z.. "Tweedr: Mining Twitter To Inform Disaster Response". In *Proceedings of the 2014 International Conference on Information Systems for Crisis Response and Management*, edited by S.R. Hiltz, M.S. Pfaff, L. Plotnick, and P.C. Shih, 354–358. University Park, Pennsylvania, USA.
- Imran, M., Castillo, C., Diaz, F., and Vieweg, S. 2015. "Processing Social Media Messages in Mass Emergency". *ACM Computing Surveys* 47(4):1–38.
- Imran, M., Castillo, C., Lucas, J., Meier, P., and Vieweg, S. 2014. "AIDR: Artificial Intelligence for Disaster Response". In *Proceedings of the Companion Publication of the 2014 International Conference on World Wide Web Companion*, 159–162. Seoul, Korea.
- Jelinski, D. E. and Wu, J. 1996. "The Modifiable Areal Unit Problem and Implications for Landscape Ecology". *Landscape Ecology* 11(3):129–140.
- Jessop, B., Brenner, N., and Jones, M. S. 2008. "Theorizing Sociospatial Relations". *Environment and Planning D: Society and Space* 26(3):389–401.
- Jiang, B. 2018. "Geospatial Analysis Requires a Different Way of Thinking: The Problem of Spatial Heterogeneity". In *Trends in Spatial Analysis and Modelling Geotechnologies and the Environment*, edited by M. Behnisch & G. Meinel, 23–40. Springer International Publishing.
- Jongman, B., Wagemaker, J., Romero, B., and de Perez, E. 2015. "Early Flood Detection for Rapid Humanitarian Response: Harnessing Near Real-Time Satellite and Twitter Signals". *ISPRS International Journal of Geo-Information* 4(4):2246–2266.
- Knüsel, B., Zumwald, M., Baumberger, C., Hadorn, G. H., Fischer, E. M., Bresch, D. N., Knutti, R. 2019. "Applying Big Data Beyond Small Problems in Climate Research". *Nature Climate Change* 9(3):196–202.
- Kropczynski, J., Grace, R., Coche, J., Halse, S., Obeysekare, E., Montarnal, A., Benaben, F., Tapia, A. 2018. "Identifying Actionable Information on Social Media for Emergency Dispatch". In *Proceedings of the 2018 International Conference on Information Systems for Crisis Response and Management Asia Pacific*. November 5th-7th, Wellington, New Zealand.
- Kryvasheyev, Y., Chen, H., Obradovich, N., Moro, E., Hentenryck, P. V., Fowler, J., Cebrian, M. 2016. "Rapid Assessment Of Disaster Damage Using Social Media Activity". *Science Advances* 2(3):1–11.
- Morstatter, F. and Liu, H. 2017. "Discovering, Assessing, And Mitigating Data Bias in Social Media". *Online Social Networks and Media* 1(1):1–13.
- Polisciuc, E., Maças, C., Assunção, F., and Machado, P. 2016. "Hexagonal Gridded Maps and Information Layers: A Novel Approach for the Exploration and Analysis of Retail Data". In *Proceedings of SIGGRAPH ASIA 2016 Symposium on Visualization*. December 5th-8th, Macau, China.
- Qadir, J., Ali, A., ur Rasool, R., Zwitter, A., Sathiaselan, A., and Crowcroft, J. 2016. "Crisis Analytics: Big Data-Driven Crisis Response". *Journal of International Humanitarian Action* 1(12):1–21.
- Reibel, M. and Agrawal, A. 2007. "Areal Interpolation of Population Counts Using Pre-Classified Land Cover Data". *Population Research and Policy Review* 26(5):619–633.

- Reuter, C. and Kaufhold, M. A. 2018. "Fifteen Years Of Social Media In Emergencies: A Retrospective Review And Future Directions For Crisis Informatics". *Journal of Contingencies and Crisis Management*, 26(1):41–57.
- Saib, M.-S., Caudeville, J., Carre, F., Ganry, O., Trugeon, A., and Cicoella, A. 2014. "Spatial Relationship Quantification between Environmental, Socioeconomic and Health Data at Different Geographic Levels". *International Journal of Environmental Research and Public Health* 11(4):3765–3786.
- Sakaki, T., Okazaki, M., and Matsuo, Y. 2010. "Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors". In *Proceedings of the 2010 International Conference on World Wide Web*, April 26th-30th, Raleigh, North Carolina, 851–860.
- Samuels, R., Taylor, J., and Mohammadi, N. 2018. "The Sound of Silence: Exploring How Decreases in Tweets Contribute to Local Crisis Identification". In *Proceedings of the 2018 International Conference on Information Systems for Crisis Response and Management*, edited by K. Boersma and B. Tomaszewski. May 20th-23rd, Rochester, New York, USA.
- Shelton, T., Poorthuis, A., Graham, M., and Zook, M. 2014. "Mapping the Data Shadows of Hurricane Sandy: Uncovering the Sociospatial Dimensions of 'Big Data'". *Geoforum* 52:167–179.
- Stumpf, M., Porter, M. 2012. "Critical Truths about Power Laws". *Science* 335(6069):665–666.
- Tapia, A. H. and 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):483–512.
- Toepke, S. L. 2018. "Minimum Collection Period For Viable Population Estimation From Social Media". In *Proceedings of the 2018 International Conference on Geographical Information Systems Theory, Applications and Management*, edited by C. Grueau, R. Laurini, and L. Ragia. March 17th-19th, Funchal, Madeira, Portugal, 138–147.
- Wang, Q. and Taylor, J. E. 2016a. "Patterns and Limitations of Urban Human Mobility Resilience under the Influence of Multiple Types of Natural Disaster". *PLoS ONE* 11(1):1–14.
- Wang, Q. and Taylor, J. E. 2016b. "Process Map for Urban-Human Mobility and Civil Infrastructure Data Collection Using Geosocial Networking Platforms". *Journal of Computing in Civil Engineering* 30(2):04015004-1–0401500411.
- Wang, Y. and Taylor, J. E. 2019. "DUET: Data-Driven Approach based on Latent Dirichlet Allocation Topic Modeling". *Journal of Computing in Civil Engineering* 33(3):04019023-1 – 04019023-8.
- World Population Review. 2018. Houston, Texas Population 2018. <http://worldpopulationreview.com/us-cities/houston-population>, accessed 17th April 2018.

AUTHOR BIOGRAPHIES

RACHEL SAMUELS is a civil engineering PhD Student in the Network Dynamics Lab at the Georgia Institute of Technology in Atlanta, GA. She holds a B.S. in Geology and a B.S.c in Business Administration from Washington and Lee. Her research interests include urban resilience, hazard remediation, and data analytics for crisis response, and she has worked on disaster research ranging from the effect of mass flooding on mining waste to the distribution of human-induced seismic activity from wastewater disposal. The main focus of her current research is the identification of intelligent, equitable methods of utilizing big social data for emergency response and hazard mitigation. Her email address is rachel.samuels@gatech.edu.

JOHN E. TAYLOR is the inaugural Frederick Law Olmsted Professor of Civil and Environmental Engineering at the Georgia Institute of Technology. He received his PhD in 2006 from Stanford University on the topic of innovation in A/E/C industry networks. At Georgia Tech, he is the founder and Director of the Network Dynamics Lab (<http://ndl.gatech.edu/>), which specializes in investigating engineering network dynamics of industrial and societal importance. His current research focuses on; (1) achieving sustained energy conservation by coupling energy use with occupant networks and examining inter-building network phenomena in cities, and (2) understanding and improving response times by affected human networks during extreme events in urban areas. His research extends from developing virtual and augmented reality applications to collect and visualize urban scale data, to developing real-time interventions to improve urban sustainability and resilience. He has authored over 200 technical publications. His email address is jet@gatech.edu.