Expanding Awareness: Comparing Location, Keyword, and Network Filtering Methods to Collect Hyperlocal Social Media Data

Rob Grace, Shane Halse, William Aurite & Andrea Tapia
College of Information Sciences and Technology
Pennsylvania State University

wrg123@ist.psu.edu, seh297@ist.psu.edu, wza5064@ist.psu.edu, atapia@ist.psu.edu

Abstract

Opportunities to collect real-time social media data during a crisis remain limited to location and keyword filtering despite the sparsity of geographic metadata and the tendency of keyword-based methods to capture information posted by remote rather than local users. Here we introduce a third, network filtering method that uses social network ties to infer the location of social media users in a geographic community and collect data from networks of these users during a crisis. In this paper we compare all three methods by analysing the distribution of situational reports of infrastructure damage and service disruption across location, keyword, and network-filtered social media data during a weather emergency. We find that network filtering doubles the number of situational reports collected in real-time compared to location and keyword filtering alone, but that all three methods collect unique reports that can support situational awareness of incidents occurring across a community.

1. Introduction

To effectively collect social media data that can support situational awareness among crisis responders and affected citizens during a crisis has long motivated researchers and systems designers [34]. In the case of Twitter, efforts have been made to collect tweets providing situational reports of events “on the ground” in order to assess damage caused by earthquakes [3], gauge flood levels [1], detect power outages [4, 17], and support the work of crisis responders and digital volunteers [7, 14].

However, existing methods to collect situational reports provide only a partial view of all crisis-related information posted on social media. In the case of Twitter, typical data collection methods rely on sparse geographic metadata and crisis-related keywords that return a fraction of all potentially-relevant tweets [21, 28, 29]. Consequently, “data sets must get bigger... before they can be sampled or filtered accordingly,”

Palen and Anderson [27] explain, “the bounds of observation must be done through decisions—which may have acknowledged limitations—to scope the data.” To widen observation of disruptive events occurring on the ground, crisis responders require new methods to collect more data than now available and, at the same time, better understanding of the limitations of each method so that multiple methods can be combined in ways that expand awareness during a crisis.

This study contributes to the critical examination of big crisis data [6, 23, 27] by comparing existing location and keyword filtering methods with a new data collection method- network filtering- to show how each conditions particular opportunities for situational awareness during a hyperlocal weather emergency. Our findings offer two primary contributions.

First, we introduce a novel data collection method that uses social network ties to infer Twitter users living in a geographic community and collect tweets they post during a crisis. We deploy and compare network filtering with existing methods during a hyperlocal weather emergency to find that over half (52%) of all situational reports are ignored when using only location and keyword-based methods to collect social media data during a crisis.

Second, we show that each of the three methods identify unique incidents of infrastructure damage and service disruption reported on Twitter, but network filtering alone identifies nearly three quarters (73%) of all incidents reported during the emergency. These findings suggest that combining multiple data collection methods is necessary when using Twitter to support situational awareness during a crisis.

2. Collecting Social Media Data During Crisis

Collecting real-time Twitter data during a crisis typically involves two primary methods [25]. The first, location filtering, uses Twitter’s Streaming API to return a sample of tweets ($\leq 1\text{-}3\%$ all tweets worldwide) including geographic metadata,

latitude/longitude coordinates associated with a GPS-enabled device (e.g. smartphone) or user-tagged “place,” that fall within a geographic bounding box. The second, keyword filtering, uses Twitter’s Streaming API to return any tweets that include selected crisis and place-related keywords (including hashtags). Based on the affordances of Twitter’s Streaming API, these two methods have become de facto standards for collecting social media data, however, other methods are possible. A third and hitherto untried method, network filtering, infers the location of users via social network ties associated with a geographic area to collect tweets from networks of these users located near a crisis. Importantly, each method introduces limitations for data collection that, in turn, shape opportunities for situational awareness during crises.

2.1. Bias of Geographic Metadata

To collect information from people in crisis-affected areas, crisis informatics researchers often first filter tweets by location, and then apply subsequent filters to identify situational reports [1]. However, location filtering identifies only tweets including geographic metadata, a mere fraction- 1-3%- of all tweets posted [22]. Location filtering thus excludes up to 97-99% of tweets posted during a crisis.

Moreover, studies show that geotagged tweets provide a biased representation of Twitter user activity [6, 11, 19], to include the types of information users post in a geographic area [29]. Per capita, more users post geotagged tweets in cities than rural areas, and tend to be younger than the general population [11, 18, 19]. Uneven tweeting activity during a crisis can, in turn, bias representations of events occurring on the ground [6, 23, 29]. Separate studies of Twitter activity in and around New York City during Hurricane Sandy, for instance, observe increased geotagged tweeting in urban centers damaged by the storm, but relatively sparse Twitter activity in neighboring urban areas that were, in some cases, more adversely affected [11, 29]. Shelton et al. [29] conclude “that places on the spatial periphery of the metropolitan area, e.g., Staten Island or the Bronx, are more likely to be marginalized within data shadows than more central locations, e.g., Manhattan and Brooklyn” (p. 173).

Using linguistic features to identify non-geotagged tweets posted within the New York City metropolitan area, Hecht and Stephens [11] discover reports of flooding in the neighboring city of Hoboken, New Jersey that are missing from geotagged tweets posted in that area. Among geotagged tweets posted in Hoboken, however, the authors find reports of flooding in Manhattan (e.g. flooding of New York Times

building). By revealing the sparsity of geotagged tweets and a reporting bias favoring incidents in urban centers over peripheral locations, these studies suggest that location filtering alone likely fails to identify the breadth and local diversity of situational reports posted on Twitter.

2.2. Bias of Keyword Filtering

Researchers also commonly employ keyword filtering methods to gather tweets by constructing queries that seek to match select crisis and place-related keywords with words people are likely to include in tweets during a crisis [25, 27]. Consequently, keywords must be selected that are common among crisis-related tweets and relatively unique compared to all tweets posted globally to comprehensively gather relevant data while preventing rate limiting and levels of noise that can quickly become prohibitive when filtering the global Twitter stream.

The necessary balance between recall and precision, however, often introduces bias towards course-grained geographic information (e.g. keywords matching city rather than street names) and information posted by geographically-remote users or oriented to them [34]. For this reason, using combinations of crisis-related words, hashtags, and globally-distinct place names [25, 34], keyword filtering collects “the most visible tweets relating to the event in question, since it is the purpose of topical hashtags to aid the visibility and discoverability of Twitter messages” [5]. As a result, the use of keywords aids the discovery of information about a crisis, but often that posted and consumed by remote crowds lacking direct ties with people located in crisis-affected areas [16]. Examining multiple keyword-filtered crisis datasets, Olteanu et al. find that eyewitness reports account for approximately 9% of all crisis-related tweets [26].

Conversely, Vieweg et al. [34] observe that tweets posted in crisis-affected areas often lack visible keywords associated with the event as people living in a geographic area often assume a shared context:

“...certain places, landmarks or items become taken-for-granted and expected when referred to in more general terms. The... [dataset] was collected based on search terms “red river” and “redriver”, and within this data set, if someone mentioned “the river” or “the flood level” it was commonly understood to be about the Red River, which makes the Red River “unmarked”— no detail is necessary when referring to it.” (p. 1086)

Tweets about the “flood level,” for instance, would never be collected unless a user also included at least one of the two selected keywords. As a result, keyword filtering often excludes situational reports that lack the

globally-visible, course-grained toponyms that tend to be assumed among Twitter users in a geographic area. Analyzing Twitter activity across three crisis events- a tornado, flood, and school shooting- Saleem, Xu, and Ruths [11] find that “the first tweets carrying situational information tended to lack the kind of identifying keywords and hashtags that would make them easy to discover in a full Twitter stream.”

2.3. Deploying Network Filtering

After comparing location and keyword filtering methods, Carley et al. conclude that “they miss most of the user population, and hence may miss critical information about who needs what help. Improved procedures for inferring location based on the user ties... are needed” [6]. Despite established research on geolocation inferencing [11, 15, 35], methods that use social network ties to infer the location of social media users, crisis informatics research has not adopted this approach to collect data from networks of Twitter users inferred near a crisis. We refer to this third method as network filtering.

Applied to Twitter, geolocation inference methods have been used to predict a users’ home city (associated with a geographic area) by comparing social network relationships among users whose locations are known (e.g. users who post geotagged tweets) and unknown [15, 35]. Someone who follows Twitter accounts followed by many people known to be living in the same geographic area, for example, may be inferred to also live in that area [20].

The limitations of location and keyword filtering recommend new methods of data collection that can capture some of the 97-99% of tweets lacking geotags, as well as compensate for the urban and global biases associated with each method, respectively. In an approach we refer to as network filtering, geolocation inferencing methods can be adopted to identify and collect social media data from networks of users associated with a geographic area. Unlike location and keyword filtering, network filtering relies on neither geographic metadata or the content of tweets to geolocate information posted on Twitter and might be deployed to collect more and more diverse geolocated Twitter data than now possible using location and keyword-based methods. However, lacking empirical evidence, the relative utility of network filtering in this respect remains unknown.

2.4. Research Questions

In this study we deploy and compare all three data collection methods- location, keyword, and network filtering- to analyze the relative opportunities they

afford when constructing situational awareness during a crisis. In the context of a severe storm in the Centre County, United States, we consider the following questions: (RQ1) How are situational reports distributed across location, keyword, and network-filtered Twitter data during an emergency? (RQ2) How are location, keyword, and network-filtered situational reports distributed across incidents observed by Twitter users during an emergency?

3. Methods

Below we describe the three data collection methods we employed to collect tweets during the storm of May 1st, our qualitative coding process, and our analysis of situational awareness information.

3.1. Location and Keyword Filtering

Location filtering involved the use of Twitter’s Streaming API to collect tweets within a bounding box encompassing Centre County, Pennsylvania during a twelve-hour period (12pm-12am) before, during, and after a severe storm and tornado that struck the area on May 1st, 2017. This produced the *Location Dataset* totaling 17,849 original tweets including either lat/long coordinates (i.e. geotag) or user-tagged places located within the county.

Keyword filtering was also performed to filter tweets that include 48 place names, including “Centre County” and the names of its 47 municipalities, boroughs, and census-designated places.¹ Data was collected for the 12-hour period to produce the *Keyword Dataset* totaling 9455 tweets.

3.2. Network Filtering

To infer and collect tweets from networks of users in Centre County we deployed a simple geolocation inferencing method that we introduce as a novel network filtering technique to collect Twitter data posted within a geographic community [9]. Typical geolocation inference methods attempt to infer n-locations for a set of Twitter users, and require “(1) a definition of what constitutes a relationship in Twitter to create the social network, and (2) a source of ground truth location data to use in inference” [15]. Most approaches utilize following or mention ties among Twitter users [11, 35] and geographic metadata, geographic references in tweet content, or profile location information as the source of ground truth for inferring the locations of users lacking geographic

¹ <http://centrecountypa.gov>

information [11, 20]. Lacking external information sources to seed the network with ground truth user locations, these approaches rely on these sparse sources of ground truth data (e.g. geotagged tweets) because they are the only sources available and suitable for automated extraction using the Twitter API.

As we seek to infer n-users for only a single geographic location (e.g. Centre County) we can approach the geolocation inferencing problem differently. Exploiting the tendency of local people to follow local organizations [20], we ascertained ground truth data by manually cataloging 195 Twitter accounts belonging to categories of organizations located in the county (bars, civic and emergency services, citizens' associations, entertainment, media, schools, restaurants) in order to identify and extract their networks of followers. These procedures are described in detail in [10]. We extracted account IDs for 185,176 users and, as our approach initially prioritizes recall over precision, we began continually collecting all tweets posted by this network of users via Twitter's Streaming API beginning in March 2016. Using this network filtering technique, data was collected during the 12-hour period of the storm on May 1st, producing the *Network Dataset* totaling 17351 tweets.

We evaluated the accuracy of this broad but potentially coarse-grained inferencing approach in two ways. First, to evaluate if most users were located in the area of interest, we used Google Fusion Tables to geocode and compare the profile locations of approximately 80k users who self-entered an identifiable location on their profiles with our geographic area of interest. Among these users, 68% entered a profile location within the county, and over 90% within the state. These results indicated that the network of users significantly overlaps with users located in the county, and that tweets posted by this network would be likely to provide situational information during the storm. Second, during our qualitative coding process, we manually investigated every tweet providing a situational report from all three datasets to determine if the post provided local information. We discuss this process in detail below.

3.3. Qualitative Content Analysis

We manually coded each tweet of the Location, Keyword, and Network datasets to understand the types of information posted on Twitter during the 12-hour period of the storm and identify tweets-situational reports- that might support situational awareness during the emergency. Qualitative content analysis provides a grounded and systematic approach for understanding the diversity of information people

report on social media during a crisis, including those that support situational awareness [13, 34]. This analysis involved three stages analyzing, in turn, tweet relevance, situational information, and location information.

First, we coded tweets as “on-topic” if any part of the tweet content referred to weather or its consequences (e.g. damage caused by high winds), and “off-topic” if the content of the tweet did not. In this initial coding process, we attempted to distinguish between emergency-related, on-topic tweets and the diversity of off-topic posts that accompany disruptive events [26, 34]. To ensure coding accuracy a random set of 1000 tweets were first given to all three coders and a Cronbach alpha test was run yielding $\alpha = 0.92$. Coding differences were deliberated and reconciled, and then the entire dataset of 44655 tweets was then subdivided and coded for relevance, resulting in 3113 (7%) on-topic tweets and 41542 (93%) off-topic tweets.

Second, on-topic tweets were coded for a second time to understand the types of information reported. Together, the authors engaged in a grounded, iterative process of open coding that involved assigning meanings, in the form of emergent code categories, to all on-topic tweets in all three datasets [13]. As an iterative process, we refined our code categories through a process of constant comparison by re-analyzing assigned codes when new themes emerged throughout the coding process [8]. This process involved the grouping and refinement of categories and sub-categories created during open coding (e.g. axial coding) [32]. We eventually arrived at 19 code categories accounting for the diversity of all on-topic information.

During this process we consulted categories developed in prior content analyses of crisis-related social media [26, 31]. While coding we noticed a diversity of information reporting forms of infrastructure damage prior studies suggest can support situational awareness during a crisis, including tweets reporting damage to buildings [26], roadways [12, 33], and electrical infrastructure [4, 17]. While this work informed our grounded analysis, the data we encountered revealed types of information that unpacked categories developed in prior research. For what Olteanu et al. categorize as “Infrastructure & Utilities” [26], for example, we develop five distinct categories: property, road, and power line damage, Internet outage, and power outage. Given the potential utility of this situational information [4], we focused subsequent analysis on these six categories (Table 1).

Third, we assessed if the tweets describing infrastructure damage and service disruption provided local information, here understood as a description of a

physical event occurring in the geographic area of interest. We recorded these events to establish a catalogue of incidents reported by Twitter users across the three datasets. To do so, we adapted criteria for determining local information utilized in prior studies [26, 30]. We determined a tweet provided local information if it a) made a geographic reference to a place(s) within Centre County, was posted by a user who b) posted a geotagged tweet(s) within the county on May 1st, or c) self-entered a profile location and, in their extant tweet stream, made a geographic reference(s) within the county. During this process we encountered many tweets we determined to be non-local although they provided information about locations in nearby, adjacent counties. The tweets ultimately assessed to provide local situational information constitute the *Situational Report Dataset*, totaling 352 tweets reporting 44 incidents across the county.

Table 1. Coding categories for infrastructure damage and service disruption

Code Category	Description Example
Power Line Damage	Tweets reporting downed or damaged power lines "our neighbors reported line dwn across [road] & [road] and that was 4 hrs ago..."
Property Damage	Tweets reporting damage to building and property "storm passes. no problems for us but two neighbors had trees hit their homes"
Road Damage	Tweets reporting damage to and obstruction of roadways "Tree down across [road] near the Meridian. Police have it blocked"
Storm Damage	Tweets reporting unspecified damage caused by high winds and rain "Major tree damage and flooding around the county. Please drive carefully!"
Internet Outage	Tweets reporting loss of internet connectivity "Either the storm is knocking out wifi in [building name] or this place is haunted"
Electricity Outage	Tweets reporting the loss or interruption of electricity "Lights out workout at East Coast Health & Fitness in [place]. literally! #blackout"

4. Analysis

At approximately 2pm on May 1st, 2017, the National Weather Service (NWS) issued Tornado Watch Number 185: "A fast-moving line of storms is expected to progress across parts of New York and Pennsylvania into this evening. Damaging wind will be the primary hazard, with a few tornadoes also possible" [24]. Over the next few hours, Twitter activity marked the eastward progress of the storm as it approached and

then struck communities in Centre County (Figure 1). Tweets warning of the storm and possible tornadoes spike after the 2pm NWS notice, followed by a flurry of weather forecasts at 5pm anticipating the impact of the storm.

The peak of the storm occurs approximately 20 minutes after 6pm with a sudden downpour of rain and wind gusts reaching over 60 mph. In a small community in the east of the county, an EF1 tornado touched down damaging several buildings, severing power lines, and uprooting trees over a one mile path [24]. For other communities, severe winds downed trees blocking roads and damaging buildings, while heavy rains caused flooding throughout the county. Immediately following the impact of the storm, reports of damage as well as power and Internet outages instantly spiked. Near 7pm the skies cleared rapidly to reveal a suddenly calm and beautiful sky.

Over the course of the storm, Centre County 9-1-1 would process over 500 calls. Most callers reported damage sustained from downed trees, including fires started from trees fallen on power lines. More than 12,000 people lost power, causing the local power company to call in utilities crews from neighboring areas, and the activation of the emergency operations center to notify electricity and telecommunications repair crews of areas reporting outages [2].

Fig 1. Total on-topic tweets (blue) and situational report tweets (red) in the Location, Keyword, and Network Datasets collected on May 1st

4.1. Distribution of Situational Reports

If public safety officials were monitoring Twitter on May 1st using existing data collection methods, they would identify less than half of all situational

reports of infrastructure damage and service disruption. Location and keyword filtering identify 28% (n=97) and 20% (n=72), respectively, of situational reports collected across the three methods (n=352). In contrast, network filtering identifies over half, 52% (n=183), of all situational reports posted during the storm (Table 2).

After removing overlaps, tweets collected by more than one method, network filtering identifies 56% (n=119) of unique situational reports. Among the three methods, keyword filtering returns the least unique data, responsible for 38% (n=27) of all situational reports. As location and keyword filtering remain de facto standards for real-time data collection, our empirical findings suggest that existing methods used to establish situational awareness during a crisis remain severely limited.

We performed a one-way ANOVA to assess if the types of situational reports (e.g. property damage, electricity outage, etc.) vary in frequency across the three data collection methods. We found no statistically significant difference between the types of situational reports collected by each method ($F(2,15) = 1.311$, $p > 0.1$). Furthermore, a Tukey post hoc test revealed no significant differences, with each grouping following the same ratio (Location-Keyword, $p=0.935$; Network-Keyword, $p=0.299$; Network-Location, $p=0.474$). Thus, while location, keyword, and network filtering collect different data, and in different volumes, each method tends to collect the same types of situational reports.

Table 2. Situational reports collected in the Location, Keyword, and Network Datasets

	Location	Keyword	Network	Total
Power Line	2	6	6	14
Property	17	12	51	80
Road	11	18	26	55
Storm	14	13	14	41
Internet	4	-	4	8
Electricity	49	23	82	154
Total	97 (28%)	72 (20%)	183 (52%)	352 (100%)

Importantly, the three methods each collect different data. While this might be expected when submitting three different queries to Twitter's Streaming API, the large number of unique tweets returned by each method demonstrates the diversity, and volume, of information users post on social media during a crisis. Overlaps occur across all three datasets, but in very small numbers (Figure 2). For example, while location and network filtering collected a total of 17,295 and 16,733 unique tweets, respectively, only 530 tweets were collected by both methods. Only 10 tweets of the 44,655 total tweets were collected by all three methods.

Interestingly, however, overlaps are much more likely to provide relevant information than unique data returned by a single collection method. For example, 5% of tweets overlapping the Location and Network Datasets, and 28% of tweets overlapping the Keyword and Location Datasets, provide situational reports. Relatedly, we find that users who posted situational reports provided more geographic information than other users, including those discussing the storm (i.e. users posting on-topic tweets). In comparison to approaches that combine real-time and post-hoc data collection, such as methods collecting the entire tweet stream of users first identified using real-time location and keyword-filtering methods [27], we find that the vast majority of all tweets in the Location (95%), Keyword (98%), and Network (99%) datasets are posted by users not otherwise identifiable in another dataset. However, among situational reports the proportions of tweets posted by unique users to the Location (58%), Keyword (32%), and Network (41%) datasets drastically decrease. This means that users who posted situational reports more often included multiple types of geographic information- for instance, by geotagging their tweets and following organizations located in the county- in tweet(s) posted during the storm than users who posted other types of information.

Fig 2. Unique and overlapping tweets in the Location, Keyword, and Network Datasets

4.2. Mapping Situational Reports

Observing that location, keyword, and network filtering collect different data in different volumes, we also analyzed the distribution of situational reports identified by each method across the 44 extant incidents- physical events of damage or disruption- Twitter users collectively reported during the May 1st storm. Mapping situational reports collected by each method to the geographic location of the incidents they

report demonstrates the bounds of observation scoped by each data collection method and suggests how each affords different opportunities for situational awareness during a crisis.

Each circle in Figure 3 represents an incident during the storm, with the circle's radius indicating the number of tweets reporting the incident, and color indicating the incident type (Figure 3). Over half of the incidents ($n=25$) are reported by multiple tweets, with the flooding of a high school football stadium reported most often among tweets in the Situational Reports Dataset ($n=21$).

Of 44 total incidents reported by Twitter users, network filtering identifies 73% ($n=32$), while location and keyword filtering both identify 43% ($n=19$) each. However, all three methods identify incidents not reported in another dataset, with unique incidents identified by location ($n=5$), keyword ($n=6$), and network filtering ($n=10$) collectively accounting for nearly half, or 48%, of all incidents reported on Twitter.

Fig 3. Incidents identified by situational reports in the Location, Keyword, and Network Datasets: power line damage (teal), property damage (yellow), road damage (pink), storm (orange), internet outage (red), and power outage (blue)

Users posting situational reports during the storm describe the location of an incident using either a street name or local landmark, but rarely both. For example, all 21 reports of flooding at the high school football stadium refer to specific (i.e. Memorial Stadium) or general landmarks (e.g. football field). Of the 44 total incidents, 45% ($n=20$) were described using only street

information (occasionally including street numbers), while 43% ($n=19$) were described using only local landmarks (e.g. names of buildings, businesses, neighborhoods, etc.). In contrast, users provided both street and landmark information for only five of the 44 incidents. Geographically, incidents identified through each method demonstrate a similar pattern: situational reports are concentrated within the largest city in the county, State College, with only scattered reports in less populous communities. This distribution can be expected given the different populations of communities across the county, but also recalls prior studies that find concentrated reporting around visible incidents in urban centers while less reporting in peripheral areas that may potentially experience more damage.

In this respect, comparing the most reported incidents- flooding of the high school football stadium ($n=21$), downed trees in a busy intersection ($n=13$), and downed power lines across a major roadway ($n=10$)- with the least reported- unique incidents identified by a single tweet- provides insight into the information behaviors of social media users reporting events during an emergency. While the three most reported incidents occur in highly frequented areas in the largest city in the county, they also are first reported by an influential social media account and subsequently reported by other, less influential users who may be providing derivative or non-eyewitness accounts. The stadium flooding was first reported by the popular Penn State University news site, Onward State, while the incidents of roadway obstruction were first reported by a local news reporter and meteorologist, respectively. These tweets mentioned other influential accounts (e.g. the local news station) and included established, highly-visible hashtags (e.g. pawx). In contrast, unique incidents were often reported by personal accounts lacking mentions and hashtags (e.g. “No Power on 700 Block of Bishop St. Damn this Sucks”). These patterns point to generative and derivative information behaviors among social media users during a crisis: some users post first-hand accounts of events while other, geographically remote users share, modify, or discuss this information [31]. We further discuss the methodological and theoretical implications of these findings in the next section.

4. Discussion

This study introduces a novel network filtering method to collect Twitter data during a crisis and, in the context of a hyperlocal weather emergency, compares and assesses how location, keyword, and network filtering methods can enhance situational awareness. In doing so we present two contributions to crisis

informatics research examining the relationship between data collection methods and opportunities for situational awareness during a crisis.

First, we introduce network filtering as a novel data collection method and empirically demonstrate how network filtering can dramatically increase the ability to collect data supporting situational awareness during a crisis. In the case of a severe storm we find that network filtering doubles the amount of on-topic, weather-related information, as well as situational reports of infrastructure damage and service disruption, collected from Twitter compared to location and keyword filtering methods. Conversely, these findings suggest that situational awareness technologies employing typical location and keyword-based data collection methods overlook a significant amount of relevant information during a crisis.

This study also reveals that location, keyword, and network filtering all provide unique opportunities for situational awareness. That is, each method collects different data in different amounts, including unique data providing unique insights into incidents occurring across a geographic community. We find that nearly half (48%) of all incidents reported on Twitter during the May 1st storm can be identified only by combining all three data collection methods. Furthermore, tweets collected by multiple collection methods are more likely to provide situational information than tweets collected by a single method alone, suggesting potential filtering strategies that can reduce dataset size and noise. By introducing network filtering as an effective data collection technique and recommending the pairing of multiple data collection methods—location, keyword, and network filtering—to expand and scope data collection during a crisis, this study makes an important methodological contribution to the design of situational awareness tools that can expand awareness during times of crisis.

Second, this study contributes to our understanding of crisis information behavior by suggesting that the types of situational information users report on social media are shaped by highly-visible information posted by influential social media accounts. Recalling the distinction between generative, eyewitness reports and derivative reports posted by social media users [31], our analysis illustrates how influential social media accounts can distribute reports of events during an emergency, that, in turn, become topics of discussion among other social media users in a geographic community. Importantly, this finding provides insight into urban reporting biases observed among social media users in prior studies [11, 19, 29]. As observed here, the most reported incidents on Twitter were those early reported by influential accounts in the community and subsequently reported by others. This finding

suggests that urban reporting biases result both from the demographics of social media users posting geotagged situational information (e.g. younger, more urban, etc.), as well as derivative information behaviors shaped by popular social media accounts that influence what information becomes visible and discussed among those social media users.

In addition, our findings provide further evidence that local social media users often omit the types of course-grained geographic information (e.g. city names) that makes tweets visible to remote Twitter users and more easily collected using keyword-based methods [28, 34]. However, we find that users posting situational reports of infrastructure damage and service disruption do include geographic information by naming local streets and landmarks in their tweets. That users often include one or the other, even when multiple users report the same incident, suggests social media users in a geographic area tend to share local knowledge, including standard place names, and communicate with others possessing the same local knowledge. Importantly, the tendency of social media users to include the names of local streets and landmarks when posting situational information recommends the creation and use of local gazetteers when designing situational awareness tools to geolocate situational reports posted during a crisis.

Lastly, we acknowledge possible limitations of this study that may arise from our analysis of a hyperlocal emergency and its difference in scale from disasters affecting larger populations and geographic areas [26]. While we find similar types of information reported by social media users, the method of network filtering we have introduced would likely require the incorporation of automated techniques to infer and collect data from the more expansive networks of users in areas affected by a disaster.

5. Conclusion

In this study we introduce and assess a new method for real-time social media data collection during crises. We review the respective biases attending existing location and keyword filtering methods and introduce network filtering as an alternative method that uses social network ties to infer the location of social media users in a geographic area and collect data from networks of these users during a crisis. Comparing the distribution of situational reports of infrastructure damage and service disruption collected by all three methods during a hyperlocal weather emergency, we find that network filtering doubles the number of situational reports collected in real-time compared to location and keyword filtering alone. However, we also find that all three methods collect unique reports, and

therefore can be deployed together to expand awareness of incidents occurring across a community.

6. Acknowledgement

We would like to thank the anonymous reviewers for their helpful comments. This research was supported by the National Science Foundation under grant Nos. 541155 and No. 741370.

7. References

- [1] de Albuquerque, J.P., B. Herfort, A. Brenning, and A. Zipf, "A geographic approach for combining social media and authoritative data towards identifying useful information for disaster management", *International Journal of Geographical Information Science* 29(4), 2015, pp. 667–689.
- [2] Annarelli, S., and J. Hartley, "Power out, schools closed after Tuesday storm", *Centre Daily Times*, 2017.
<http://www.centredaily.com/news/local/community/article148010634.html>
- [3] Avvenuti, M., S. Cresci, A. Marchetti, C. Meletti, and M. Tesconi, "EARS (Earthquake Alert and Report System): A Real Time Decision Support System for Earthquake Crisis Management", *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, pp. 1749–1758.
- [4] Bauman, K., A. Tuzhilin, and R. Zaczynski, "Using Social Sensors for Detecting Emergency Events", *ACM Transactions on Management Information Systems* 8(2–3), 2017, pp. 1–20.
- [5] Bruns, A., and Y.E. Liang, "Tools and methods for capturing Twitter data during natural disasters", *First Monday* 17(4), 2012.
- [6] Carley, K.M., M. Malik, P.M. Landwehr, J. Pfeffer, and M. Kowalchuck, "Crowd sourcing disaster management: The complex nature of Twitter usage in Padang Indonesia", *Safety Science* 90, 2016, pp. 48–61.
- [7] Cobb, C., T. McCarthy, A. Perkins, et al., "Designing for the deluge", *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing - CSCW '14*, ACM Press (2014), 888–899.
- [8] Glaser, B., and A. Strauss, *The discovery of grounded theory: Strategies for qualitative research*, Aldine de Gruyter, Hawthorne, NY, 1967.
- [9] Grace, R., J. Kropczynski, S. Pezanowski, S. Halse, P. Umar, and A. Tapia, "Social Triangulation: A new method to identify local citizens using social media and their local information curation behaviors", *Proceedings of the 14th International Conference on Information Systems for Crisis Response and Management*(May 2017), 2017, pp. 902–915.
- [10] Grace, R., J. Kropczynski, S. Pezanowski, S. Halse, P. Umar, and A. Tapia, "Social triangulation: A new method to identify local citizens using social media and their local information curation behaviors", *Proceedings of the 14th ISCRAM Conference 2017–May*, 2017, pp. 902–915.
- [11] Hecht, B., and M. Stephens, "A Tale of Cities: Urban Biases in Volunteered Geographic Information", *Eighth International AAAI Conference on Weblogs and Social Media (ICWSM)*, (2014), 197–205.
- [12] Hong, L., C. Fu, P. Torrens, and V. Frias-Martinez, "Understanding Citizens' and Local Governments' Digital Communications During Natural Disasters", *Proceedings of the 2017 ACM on Web Science Conference - WebSci '17*, ACM Press (2017), 141–150.
- [13] Hsieh, H.-F., and S.E. Shannon, "Three Approaches to Qualitative Content Analysis", *Qualitative Health Research* 15(9), 2005, pp. 1277–1288.
- [14] Hughes, A., and R. Shah, "Designing an Application for Social Media Needs in Emergency Public Information Work", *Proceedings of the 19th International Conference on Supporting Group Work - GROUP '16*, ACM Press (2016), 399–408.
- [15] Jurgens, D., T. Finethy, J. Mccorriston, Y.T. Xu, and D. Ruths, "Geolocation Prediction in Twitter Using Social Networks: A Critical Analysis and Review of Current Practice", *Proceedings of the Ninth International AAAI Conference on Web and Social Media*, (2015), 188–197.
- [16] Kogan, M., L. Palen, and K.M. Anderson, "Think Local, Retweet Global", *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15*, ACM Press (2015), 981–993.
- [17] LaLone, N., A. Tapia, C. Zobel, C. Caraega, V.K.

- Neppalli, and S. Halse, “Embracing human noise as resilience indicator: twitter as power grid correlate”, *Sustainable and Resilient Infrastructure*, 2017, pp. 1–10.
- [18] Lin, J., and R.G. Cromley, “Evaluating geo-located Twitter data as a control layer for areal interpolation of population”, *Applied Geography* 58, 2015, pp. 41–47.
- [19] Malik, M.M., H. Lamba, C. Nakos, and U. Pfeffer, “Population Bias in Geotagged Tweets”, 2015.
- [20] McGee, J., J. Caverlee, and Z. Cheng, “Location prediction in social media based on tie strength”, *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management - CIKM '13*, ACM Press (2013), 459–468.
- [21] Morstatter, F., N. Lubold, H. Pon-Barry, J. Pfeffer, and H. Liu, “Finding Eyewitness Tweets During Crises”, 2014.
- [22] Morstatter, F., J. Pfeffer, H. Liu, and K.M. Carley, “Is the Sample Good Enough? Comparing Data from Twitter’s Streaming API with Twitter’s Firehose”, 2013, pp. 400–408.
- [23] Mulder, F., J. Ferguson, P. Groenewegen, K. Boersma, and J. Wolbers, “Questioning Big Data: Crowdsourcing crisis data towards an inclusive humanitarian response”, *Big Data & Society* 3(2), 2016, pp. 205395171666205.
- [24] NWS, *Tornado Watch 185*, NOAA, 2017.
- [25] Olteanu, A., C. Castillo, F. Diaz, and S. Vieweg, “CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises”, *Proc. of the 8th International Conference on Weblogs and Social Media*, 2014, pp. 376.
- [26] Olteanu, A., S. Vieweg, and C. Castillo, “What to Expect When the Unexpected Happens”, *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15*, ACM Press (2015), 994–1009.
- [27] Palen, L., and K.M. Anderson, “Crisis informatics: New data for extraordinary times”, *Science* 353(6296), 2016, pp. 224–225.
- [28] Saleem, H.M., Y. Xu, and D. Ruths, “Novel Situational Information in Mass Emergencies: What does Twitter Provide?”, *Procedia Engineering* 78, 2014, pp. 155–164.
- [29] Shelton, T., A. Poorthuis, M. Graham, and M. Zook, “Mapping the data shadows of Hurricane Sandy: Uncovering the sociospatial dimensions of ‘big data’”, *Geoforum* 52, 2014, pp. 167–179.
- [30] Starbird, K., and L. Palen, “(How) will the revolution be retweeted?”, *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work - CSCW '12*, ACM Press (2012), 7.
- [31] Starbird, K., L. Palen, A.L. Hughes, and S. Vieweg, “Chatter on the red”, *Proceedings of the 2010 ACM conference on Computer supported cooperative work - CSCW '10*, ACM Press (2010), 241.
- [32] Strauss, A., and J. Corbin, *Basics of qualitative research: Grounded theory procedures and techniques*, Sage, Thousand Oaks, CA, 1990.
- [33] Tien, I., M. Aibek, D. Benas, and C. Pu, “Detection of damage and failure events of critical public infrastructure using social sensor big data”, *Proceedings of International Conference on Internet of Things and Big Data*, (2016).
- [34] Vieweg, S., A.L. Hughes, K. Starbird, and L. Palen, “Microblogging during two natural hazards events”, *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, ACM Press (2010), 1079.
- [35] Zheng, X., J. Han, and A. Sun, “A Survey of Location Prediction on Twitter”, *IEEE Transactions on Knowledge and Data Engineering*, 2018, pp. 1–1.