Think Local, Retweet Global: Retweeting by the Geographically-Vulnerable during Hurricane Sandy

Marina Kogan, Leysia Palen & Kenneth M. Anderson
Department of Computer Science
University of Colorado Boulder
{marina.kogan, leysia.palen, ken.anderson}@colorado.edu

ABSTRACT
Hurricane Sandy wrought $6 billion in damage, took 162 lives, and displaced 776,000 people after hitting the US Eastern seaboard on October 29, 2012. Because of its massive impact, the hurricane also spurred a flurry of social media activity, both by the population immediately affected and by the globally convergent crowd. In this paper we explore how retweeting activity by the geographically vulnerable differs (if at all) from that of the general Twitter population. We investigate whether they spread information differently, including what and whose content they chose to propagate. We investigate whether the Twitter-based relationships are preexisting or if they are newly formed because of the disaster, and if so if they persist. We find that the people in the path of the disaster favor in their retweeting locally-created tweets and those with locally-actionable information. They also form denser networks of information propagation during disaster than before or after.

Author Keywords
Crisis informatics; Disasters; Protective Decision-Making; Social Computing; Social Media; Twitter

ACM Classification Keywords
H.5.3. Groups & Organization Interfaces—collaborative computing, computer-supported cooperative work; K.4.2. Social Issues

General Terms
Human Factors

INTRODUCTION
Social media are platforms for one-to-many communication that are being used by the public during disaster response for a range of purposes. A flourishing body of research on the topic has developed in the last 5 years, with a community of researchers working apace with the socio-behavioral phenomena that are advancing in kind and number with each disaster. Some research in the *crisis informatics* space examines how formal responders use or do not use social media [6, 15, 17, 18]. Another line considers how online communication affects on-the-ground action before, during, and after disasters [30, 36, 44, 45].

However, the largest body of work is on the internal behaviors of the “Twittersphere” and other social media environments. Here again the research directions branch, with much work devoted to deriving data from existing social media logs using supervised and unsupervised machine learning (for example, [19, 42]) as well as other techniques [1, 6, 10]. The second branch of research that focuses on the social media posts as the objects of study is that which attempts to understand socio-behavioral phenomena. This includes sentiment analysis across the population represented in the streams of social media communication [7, 11]. It also includes analysis of self-organizing behaviors of groups that come together through social media to accomplish some task such as event reporting [20, 21, 22, 31] or a task they articulate as a particular need during that event [38, 44]. A third branch considers how information diffuses across social networks [4, 9, 34], particularly as the special conditions of disaster affect such diffusion [33, 37, 39].

Paper Objective & Background
This paper sits as a contribution in the space that examines online socio-behavioral phenomena, and in particular how information is diffused across a population. However, we target a line of inquiry that includes the population of people who are known to have been in the geographical area of effect of a major disaster during its height. In this case, we examine those who were under the most serious threat before and during the 2012 Hurricane Sandy, and who used Twitter to post during that time. Specifically, this is an analysis of how people who are at risk from a natural hazard with an advanced warning period—a hurricane—retweet information before, during, and after the event.

First, we investigate whether those who are geographically vulnerable spread information differently than the general public—a public that was highly active on Twitter around

CSCW '15, March 14 - 18 2015, Vancouver, BC, Canada

ACM 978-1-4503-2922-4/15/03…$15.00
http://dx.doi.org/10.1145/2675133.2675218

the world because of Hurricane Sandy's catastrophic potential and intense media attention. In addition, when the geographically vulnerable population retweets, what information do they spread, and what and whose tweets do they propagate if they do so at all? We investigate whether the Twitter-based relationships are preexisting or if they are newly formed during the disaster, and if those new relationships persist after the storm.

The objective of this line of inquiry is multi-fold. First, we want to continue to validate and deepen earlier findings by Starbird and Palen from the 2009 Red River Flood event that "locals" are more likely to retweet content that has "local utility," and that non-locals are more likely to retweet the "abstract" of the event [37]. In 2009, hardly any emergency management groups were on Twitter—but things had radically changed between 2009 and 2012 as emergency management groups tried to develop an online presence, including more police and fire departments [18]. In addition, relatively few media outlets were on Twitter in 2009, certainly as reflected in the Red River Flood data. In addition, the 2012 Sandy event had far more global attention as well as more extreme immediate consequences to a larger population than the US/Canadian Red River Flood threat of 2009—a critical event for Red River Valley locals, but incomparable to Sandy in terms of media reach. Furthermore, the number of active twitterers between 2009 and 2012 increased about seven fold (30M to 200M).[1] Finally, the act of retweeting in 2009 was only a user innovation—it was not built into the software (and therefore not into the metadata of tweets). We want to confirm, as well as elaborate, the earlier local utility versus global interest framing under these changing and far more expansive conditions.

Second, the Sandy-affected population is a desirable one to study because it is high density, where people are more likely to have enough relationships via Twitter within that geographical space to truly see what technology-abetted socio-behavioral phenomena—and particularly information diffusion—might be. The affected population also represents high variance in socio-economic status (SES) [12], though current research suggests that the affected populations in less central (and less affluent) areas might be underrepresented in the Twitter feed [35].

Such conditions, however, bring us closer to seeing if the empirical findings of geographical behavior that social science has previously revealed are echoed in the digital world. Specifically, do new relationships between people form during disasters in the online world? We expect that they do, but we do not know to what extent those relationships persist. We know from the disaster literature that in areas that experience seasonal hazards, the affiliations and connections made in prior events often (but not always) extend to subsequent events, and to the extent to which these connections can be used as a kind of "organizational memory" [2] is exactly what those who are increasingly attached to ideas of "community resilience" [14] are banking on. This research aims to provide baseline information—as well as some longitudinal information—about an important event that will undoubtedly precede similar events in the same region.

Furthermore to this point, if we find that there are retweet behaviors that are particular to the affected population, then we may be able to use those features to derive data in future disaster events more quickly, for both scientific and applied purposes, that could aid in emergency management. As Twitter volume increases, we need more sophisticated techniques beyond keyword filtering to find the people of interest—specifically those who are needing or supplying information about the impending hazard and its effects on the social and built environment [29].

Third, this research is the foundation for an important new area of research in crisis informatics: that of understanding how people who are geographically vulnerable during hazards make "protective decisions"—decisions to evacuate or to "shelter in place," which may include collecting sufficient provisions. These decisions are made as part of a "web" of sources that are available to them; they do not rely on official sources or mandates alone [13, 24, 26, 27, 47]. Because people are increasingly turning toward online sources, including social media, we must understand social media behavior as much as possible. This research is necessarily scoped to information diffusion behaviors, but it is part of a line of inquiry about "protective decision-making" by affected populations, especially as those activities vary across SES and other demographics [47].

THE 2012 HURRICANE SANDY EVENT (US LANDFALL)

Hurricane Sandy made landfall on October 29, 2012 in southern New Jersey, affecting one of the most populated

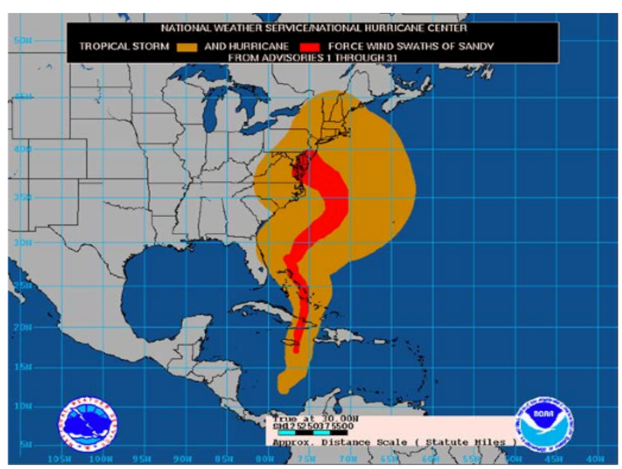

Figure 1. Path of Hurricane Sandy provided by the US National Weather Service via [12].

[1] See <http://bit.ly/1rJXKgA> and <http://bit.ly/1avywUF> for details.

regions in the US that include New Jersey, New York, and Connecticut, with its impacts felt over a total of 24 States [12]. It dissipated by November 1. Prior to US landfall, it had passed through the Caribbean Sea, causing much damage to island nations before putting the entire US eastern seaboard under threat from the Gulf of Mexico and north into Canada (see Figure 1).

In the US, the number of deaths directly attributed to the storm totaled 162 [12]. It was the second costliest hurricane to hit the US; damage was estimated to be US $6 billion [5]. Approximately 776,000 people were displaced [46] and 650,000 homes were damaged or destroyed. 8.5 million people lost power as a result of the storm [5], and many were without power for weeks following.

METHODS

Data Collection Steps

Data collection is part of our ongoing and committed effort to study disaster-related Twitter data. Using a four-node Cassandra cluster, our research group collects Twitter data 24/7 in specialized software designed for high-volume Twitter data collection [3]. For the Sandy event, we started collecting data using the Streaming API on October 24 2012, using the following keywords for the first round of data collection: `frankenstorm`, `hurricane`, `hurricanesandy`, `perfectstorm`, `sandy`, `sandycam`, `stormporn`, `superstorm`. This produced the *keyword data set.* In our disaster-related data collection procedures, we mitigate against potential bias in the Streaming API sampling [28] by using carefully chosen keywords and then focusing on specific subsets of users to gather a complete set of their tweets via the REST API. In particular here, we determined which users contributed to our keyword set, and then filtered again to those who had at least one gelocated tweet that fell within the geographical area of interest (see Figure 2). We then pulled the user streams—or what we call the *contextual streams*—for each of these *geolocated users* using the Twitter REST API. The reason we collect contextual streams is because we seek the fuller semantic context before and after a tweet that contains a found keyword. A "found" tweet based on a keyword search can have a different or enhanced meaning when one examines the surrounding tweets by the same user. Even when doing non-linguistic analysis as we do here, the full contextual streams remain important because the twitterers *themselves* presume contextual continuity across their tweets as they write them such that they do not necessarily invoke a disaster-relevant keyword each time they tweet. Therefore, our thinking is that the unit of analysis should not depend on isolated keyword-found tweets [29]. Note that the REST API returns up to 3200 of a user's tweets from most recent to least; this usually allows examination of a user's behavior before the hurricane as well.

Creating the Data Sets

We next describe all the data sets used in this paper, with a summary of main sets in Tables 1 and 2. The original raw keyword search spans a long period of time. For this study we bounded the set to Oct 24-Nov 30 to constrain it to the first weeks of the recovery. We refer to this as the *Global Keyword* dataset (see Table 1). Data before October 24 preceded any reasonable predictions of where the hurricane would likely make US landfall.

Next we worked in collaboration with meteorologists, social scientists, and GIS researchers at the National Center for Atmospheric Research to define the geographic bounding box of the region most affected by Sandy. The bounded area spans a great portion of the eastern seaboard (Figure 2), and intentionally covers inland locations, to which people closer to the coastline were likely to evacuate.

Users who produced at least one geolocated tweet within the bounding box during the *global keyword* search time window (Oct 24-Nov 30) were designated as *geographically (*or *"geo") vulnerable users.* Their tweets that were part of the keyword collection yield the *Geo-Vulnerable Keyword* dataset. All tweets in the *Global Keyword* dataset produced by the geographically vulnerable users—including their non-geolocated tweets—are part of this set.

With the set of local users known, we collected their contextual streams as described above. This is called the *Geographically ("Geo") Vulnerable Contextual* dataset.

We further isolate the *Geo-Vulnerable Contextual* dataset into four distinct 5-day time slices to capture the activity of the most geo-vulnerable users *Geo-Before* (Oct 15-19), *Geo-During* (Oct 27-31), *Geo-Short-After* (Nov 8-12), and *Geo-Long-After* the event (Oct 22-26) (see Table 2). We choose these time frames based on our knowledge of socio-behavioral phenomena with respect to different phases of disaster events (drawing primarily from Powell [32]). In sum, *Geo-Before* captures the time before people in the geographical area of interest could know they would be

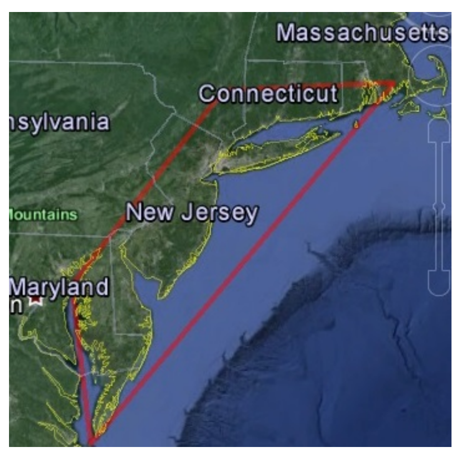

Figure 2. Bounding Box for Data Filtering. The bounding polygon is defined by the coordinates of its four corners: [-76.055416, 36.988536], [-76.416506, 39.084008], [-73.872974, 41.654353] & [-70.874853, 41.732875]).

under threat but not so far beforehand that their social networks would be very different simply as a function of time. This gives us a before-disaster view of their behavior. The *Geo-During* captures the intense *high warning, evacuation and storm impact periods*, when people are making "protective decisions" and then living through the storm. The *Geo-Short-After* period captures our population as they move into the second week of *recovery*: some here could return home, many others could not, and some were able to start formulating plans for repair work. In other words, they have been able to take stock and have made initial post-disaster plans. The *Geo-Long-After* represents almost a full month after the event. Even more people have stabilized at this point. Many will have returned back to normal routines, but those who lost their homes are making other arrangements.

These datasets will be used to internally compare different retweet patterns, and specifically to look at how social networks emerged in relation to the disaster and to what extent they persisted into the recovery. Because we are interested in comparing the retweeting behavior of geo-vulnerable users within these various time periods, it was necessary to *limit these four time slice sets to the unique users common to all the periods*. We did this to be able to establish a baseline comparison from which to derive heuristics about tweeting behavior without additional variables for which to account. We note that we might be losing people in the *Geo-During* period if they did not evacuate and suffered from prolonged power outages (and therefore could not tweet), but this is a limitation of the study that we opted to work around to achieve stability in the analysis elsewhere. (Future work will extend the analyses to account for differences once baseline behaviors are uncovered.) With this approach, we found 7,988 twitterers who contributed retweets to the contextual dataset in all four time periods. Limiting the time slice sets to only those overlapping users resulted in smaller sets as reported in Table 2.

Three Retweet Counts

As we are interested in retweet behavior in this analysis, we rely heavily on the retweet count metric. However, the Twitter retweet count field is not particularly useful, as for each retweet it *indicates only its current turn* in retweeting the original. For example, the tweet with a retweet count of 5 is the 5th retweet of the original source. Thus, the Twitter metric maintains the retweet count of the original at the time of a particular retweet, and stores it associated with that retweet. We instead need to know *how many times the original tweet was propagated overall within the particular time period under study*. To compute that metric, we went through all the original tweets within a particular dataset, found all their retweets, and stored the latest/largest Twitter retweet count among those retweets as a measure of how many times the original tweets have been passed along. In the remainder of this paper, that is what we refer to as *retweet count*.

DATA SET NAME & TIME SPAN	NUMBER OF TWEETS	NUMBER OF USERS
Global Keyword (Oct 24-Nov 30)	16.2M	5.9M
Geographically ("Geo") Vulnerable Keyword (Oct 24-Nov 30)	224.8K	28.5K
Geographically ("Geo") Vulnerable Contextual (Oct 13-Nov 30)	5.6M	28.5K

Table 1. Dataset Names and Descriptions.

DATA SET NAME	DATES (2012)	RETWEET VOLUME FOR IN-COMMON TWITTERS
IN-COMMON TWITTERERS ACROSS ALL TIMES SLICES: 7988		
Geo-Before	Oct 15-Oct 19	64,423
Geo-During	Oct 27-Oct 31	84,597
Geo-Short-After	Nov 8-Nov 12	55,855
Geo-Long-After	Nov 22-Nov 26	50,425

Table 2. Retweet volume for in-common Twitterers across all time slices.

Since the technique for acquiring the retweet count requires going through all the original tweets in a dataset, these counts are necessarily timeframe- and dataset-dependent. For the retweet count distribution analysis, we were interested in the activity for the *Geo-During* time frame as it definitively includes the period when the evacuation notices were issued and the storm made landfall. We computed *three* types of retweet counts within this time frame based on the *Global Keyword* and *Geo-Vulnerable Keyword* datasets, described in Table 3. *Retweet counts* are named based on the relationship between *where the original tweet was sourced* and the *population it was retweeted by*.

ANALYSES & FINDINGS

In this section, we have opted to include the findings uncovered in each progressive step to help the general reader follow the analytic rationale and argumentation.

Geo-Vulnerable Retweet Networks for the 4 Time Slices

Time Evolution of Retweet Networks

For the four time slices with in-common users (*Geo-Before, Geo-During, Geo-Short-After, Geo-Long After)*, we collected the user ids of those who generated the retweets as well as the user ids of the authors of the original tweets that were being retweeted. Retweeting behavior can be seen to signify some kind of loosely-connected social relationship—at a minimum, the retwitterer sees value in the information or in the original twitterer. These relationships can be represented as a directed graph, with the retwitterers as source nodes and original twitterers as target nodes, and the directed edges (from source to target) representing the retweets. Thus, the four time slice datasets produce four distinct directed social networks, with the common core of overlapping retwitterers combined with the original tweet authors who are specific to the time period as nodes. These four networks are discrete time slices in the temporal evolution of the single retweet network. Next, we compare various structural aspects of the four networks to establish how the web of social relations changed during disaster in comparison to pre- and post-disaster activity.

Network Size & Density

The first difference in the four time sliced networks is the sheer size. Table 4 illustrates the size and density measures for the four networks. The *Geo-During* network is considerably larger, with more nodes and edges, than the prior and later slices. The larger size of the network corresponds to the higher volume of retweeting activity we observed during disaster, which is not so surprising considering that the underlying dataset was constructed by collecting the contextual streams of geo-vulnerable users who were found by searching on the hurricane-related keywords.

On the one hand, though the geo-vulnerable users were found through a keyword search, in this analysis we use their entire contextual tweet streams during the four time slices, which provide the most accurate measure of how much they tweeted, even when their tweets do not contain keyword terms. The higher volume of retweet activity during the hurricane suggests that the geo-vulnerable population tends to propagate social media posts more often in disaster (even if the posts do not contain event-related keywords).

The relative size of the largest connected component can serve as a proxy for how densely interconnected the network is. The largest weakly connected component of the *Geo-During* network encapsulates 92.88% of all its nodes and an impressive 96.34% of the edges—a considerably larger fraction than for the other three time periods (see Table 4). Moreover, if we construct completely comparable "internal" versions of the four networks by excluding the original tweet authors who are not also among the 7,988 overlapping users and thus retaining only the retweets between the core users, the networks show a similar trend. The "internal" version of *Geo-During* has the largest weakly connected component, with the fractional size considerably larger than the internal networks for other time slices. Therefore, we see higher density in the *During* period even in the internal networks, which allow the analysis to focus solely on the retweet dynamics among the overlapping users to avoid the varying number of external original tweet authors (though this method does not control for the underlying dynamic of varying retweet volume). The higher density suggests that the retweet activity of geo-vulnerable twitterers during the disaster connects multiple subnetworks, thereby constructing a more interconnected, dense social network.

The reciprocity of directed links does not contribute much to the higher network density of *Geo-During*, as it is consistently low for all four networks—about 0.1%.

Degree Distributions

An even more canonical metric for measuring the difference in structure of different networks is their degree distribution. In this directed case, we are specifically interested in the out-degree, which represents the number of original authors whose tweets the in-common retwitterers have retweeted.

All the distributions in Figure 3 are long-tailed, as is common in degree distributions of real-world complex networks [8, 24]. However, some of the distributions are prone to rarer events than others. For example, the *Geo-Before* network has the highest out-degree of 181. Thus, an in-common retwitterer in this network (@BKdotNet) has retweeted tweets from 181 authors in the 5 days of interest before the event (in-degree=0). Similarly, and somewhat more impressively, a user (@LaborIrishDem) in the *Geo-During* network retweeted tweets from 192 authors during the 5 days of the event (in-degree=0). In contrast, the *Geo-Short-After* and *Geo-Long-After* networks have maximum out-degrees of 134 and 97, respectively (in-degree=0 for both). Moreover, the out-degree distributions for all the time slices are power-law-like only in the tail (out-degree>10), but not the head of the distribution. This

RETWEET COUNT NAME	DESCRIPTION
1. *Global/Geo-Vulnerable*	Originating tweets come from the *Global Keyword* dataset and are retweeted within the *Geo-Vulnerable Keyword* dataset
2. *Geo-Vulnerable/Geo-Vulnerable*	Originating tweets come from the *Geo-Vulnerable Keyword* data set and are retweeted within the *Geo-Vulnerable-Keyword dataset*
3. *Global/Global*	Originating tweets come from the *Global Keyword* data set and are retweeted within the *Global Keyword* data set

Table 3. Glossary of Retweet Counts.

	Geo-Before	Geo-During	Short-After	Geo-Long-After
Network size (nodes, edges)	49,413 64,423	56,527 84,597	46,040 55,855	42,829 50,425
Weakly Connected components	1,465	935	1,728	2,047
Fraction in largest component (nodes, edges)	87.61% 92.76%	92.88% 96.34%	84.10% 89.97%	81.28% 88.15%

Table 4. Network size and density. (Fraction of nodes stat. sig. with Chi-square=3304.55, p<0.0001, fraction of edges stat. sig. with Chi-square=3664.98, p<0.0001).

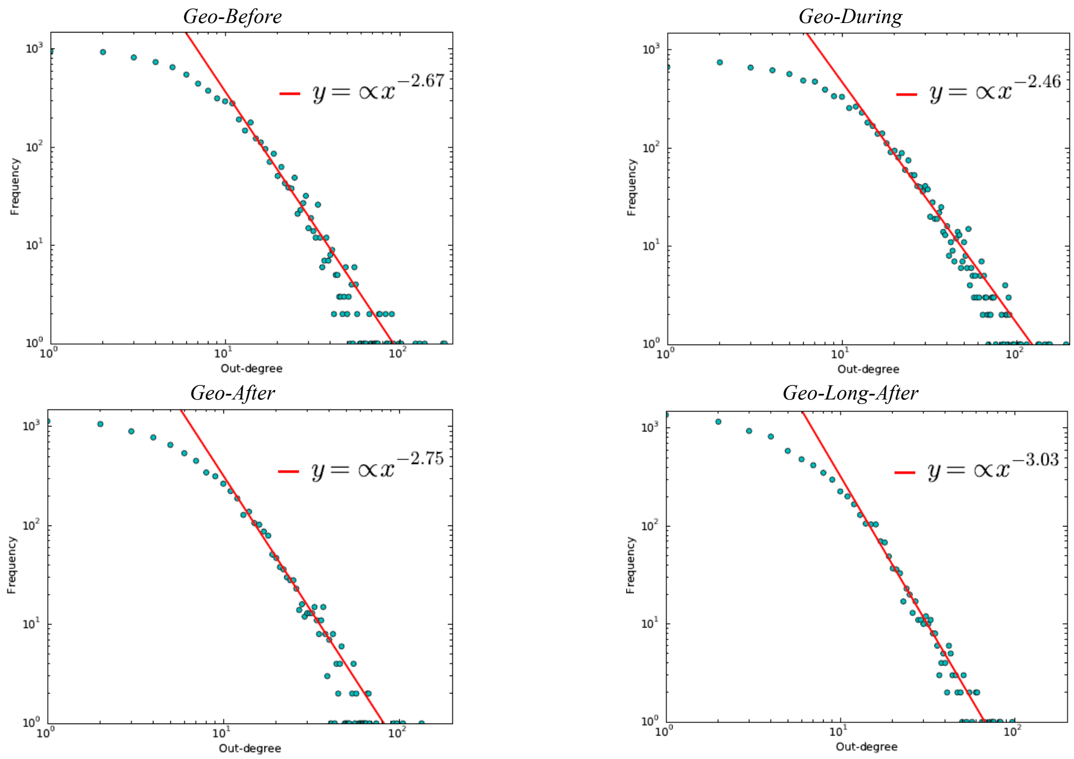

Figure 3. Loglog of out-degree distribution with linear fit to the tail (x>10).

suggests that the core users who remained active retwitterers through all the stages of the event tended to retweet tweets of many authors. This is especially true for the *Geo-During* network, where the linear fit to the loglog distribution is the least steep of all the time slices, decreasing less rapidly and thus pushing more of the distribution's density away from the origin. Hence, the core in-common users, in aggregate, tended to retweet information from a greater variety of sources during the disaster than before or after.

The in-degree here represents how popular various original authors were in the retweet activity of the core overlapping users. For example, the *Geo-Before* network has the highest in-degree of 306. Here @BarackObama was retweeted by 306 different in-common users in the 5 days of interest before the event. This node's out-degree is zero, since we do not have the data on the retweet behavior of the original authors unless they are also one of the in-common retwitterers (see Table 6 for the intersection of the two). Similarly, but even more impressively, @MikeBloomberg —the New York city mayor at the time—was retweeted by 420 different core users during the 5 days of the event (out-degree=0). The author retweeted by the highest number of in-common users in the *Geo-Short-After* and *Geo-Long-After* networks is @XSTROLOGY (in-degree of 167 and 154, respectively; out-degree of 0 for both).

Users with high degree—many links—are often called "hubs" and are an important feature of social networks. The *Geo-During* network has more hubs based on in-degree, while the *Geo-Before, Geo-Short-After,* and *Geo-Long-After* networks have fewer of them. Figure 4 illustrates the proportion of nodes for each network that have 100 links or more, for in- and out-degree. The differences for the out-degree are less dramatic, but the *Geo-Before,* and especially *Geo-During* have more hubs than the later time networks. This suggests that among the users who appear across all datasets, there was more hub-like activity during the disaster period—retweeting from multiple sources and being retweeted by multiple retwitterers—and thus connecting many subnetworks.

Network Mixing Patterns

More insight into understanding and comparing network structure can be gleaned from "mixing patterns," which refers to what type of nodes tend to connect to each other. "Degree assortativity" is measured by a network-level coefficient quantifying the strength of the relationship between the degree of the network's nodes and degree of their neighbors. For our networks, which do not have data on the retweet activity of original authors unless they are also in-common retwitterers, we are most interested in out-in degree assortativity. This metric represents the correlation between the out-degree of a node and the average in-degree of its neighbors—tweets of how many authors the user retweeted and how many users, on average, retweeted the tweets of those original authors. The out-in degree assortativity for all the time sliced networks is low, signifying, on average, a weak relationship between each node's degree and degree of its neighbors. The out-in degree assortativity of all the networks is also weakly

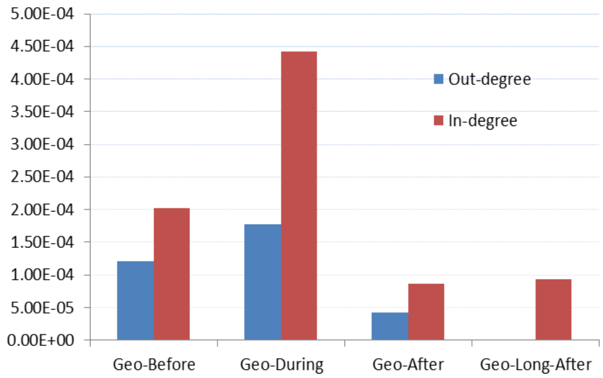

Figure 4. Hub percentage for each network. (Stat. sig. for in-degree with χ^2=19.53, p<0.001).

negative, suggesting that nodes with high out-degree are on average slightly more likely to be connected to the nodes with low in-degree. This implies that for all four networks, when users retweet many sources, those sources are slightly more likely to be the ones that have not been retweeted very often by the in-common retwitterers. All the patterns above also hold if we remove all the nodes with an out-degree of one; we do this to verify that one-off retwitterers of popular memes (high in-degree) do not produce the observed disassortativity. This disassortative relationship is weakest for the *Geo-During* network, suggesting that during disaster, the geo-vulnerable users who retweet many different sources are slightly more likely to retweet more popular sources than before or after an event. The fact that popularity plays a role in retweet activity in disaster is somewhat intuitive, as retweeting can be seen to be about trust (in either content or here in the authority of the author as confirmed by her/his popularity), and rarely is trustworthiness more important than in disaster. We will attempt to separate the part played by content utility in this process from that of author popularity later in the paper.

We can disaggregate the concept of out-in degree assortativity by looking at the average neighbor connectivity for each node (see Figure 5). The sub-figures illustrate the average in-degree of neighbors (y-axis) for the node with out-degree *k* (x-axis).

Figure 5 shows a less pronounced negative relationship between node out-degrees and average in-degrees of their neighbors for the *Geo-During* network, supporting our assertion of slightly more frequent hub-to-hub retweeting (out-degree hub retweeting in-degree hubs). For example, @AthenasAika—the user represented by the highest point on the *Geo-During* graph—has retweeted 106 sources who on average have been retweeted 43.70 times. This user retweeted such high-degree hubs as @BarackObama, @MikeBloomberg, and other national and local authorities and media outlets.

Transitivity

A relation is transitive if when A relates to B, and B relates to C, A also relates to C. Thus transitivity captures the idea of "a friend of a friend is a friend," and specific to retweeting information, "a source of my source is also a source." In the directed case, the most meaningful network-level measure of transitivity is the proportion of transitive triples derived from the triadic census as defined by Holland and Leinhardt [16, 43]. We corrected this measure by excluding the few transitive triples in which a user retweets two authoritative sources that retweet each other. Thus, the corrected metric excludes the configurations that relate to authority-based broadcasting and retains only the transitive triple configurations associated with community-building retweeting, which is a typical understanding of transitivity. This corrected proportion, while rather low for all four time sliced networks, is the highest for the *Geo-During* network. The *Geo-After* has slightly higher corrected proportion of transitive triples than *Geo-Before* and *Geo-Long-After,* but still lower than *Geo-During.* Therefore, the higher proportion of transitive triples in *Geo-During,* suggests that when retweeting communities are likely to form, they are most likely to form during the event and persist to some degree immediately after (the difference is statistically significant with χ^2=83.29, p<0.0001).

The retweet behavior displays low transitivity because retweeting need not be limited by the boundaries of users' following networks, which are themselves low-transitivity [25] due to their directional nature and low rates of following reciprocity [23].

Important Users

We further explored the differences between the four time sliced retweet networks by establishing each network's important users, with the specific focus on the influential sources (See Table 5). PageRank—a variant of eigenvector centrality—has at its core a notion that the importance of a node can be judged by looking at the importance of nodes that link to it. Thus, this metric is especially well-suited for

Geo-Before	Geo-During
BarackObama	MikeBloomberg
MensHumor	GovChristie
XSTROLOGY	LolOhComedy
billmaher	MTAInsider
WomensHumor	NYCMayorsOffice
azizansari	ElBloombito
FillWerrell	NYGovCuomo
Geo-After	**Geo-Long-After**
XSTROLOGY	XSTROLOGY
GovChristie	UberFacts
UberFacts	HuffingtonPost
MensHumor	MensHumor
ElBloombito	BenSavage
WhatTheFFacts	SportsCenter
HuffingtonPost	WhatTheFFacts

Table 5. Nodes with the highest PageRank.

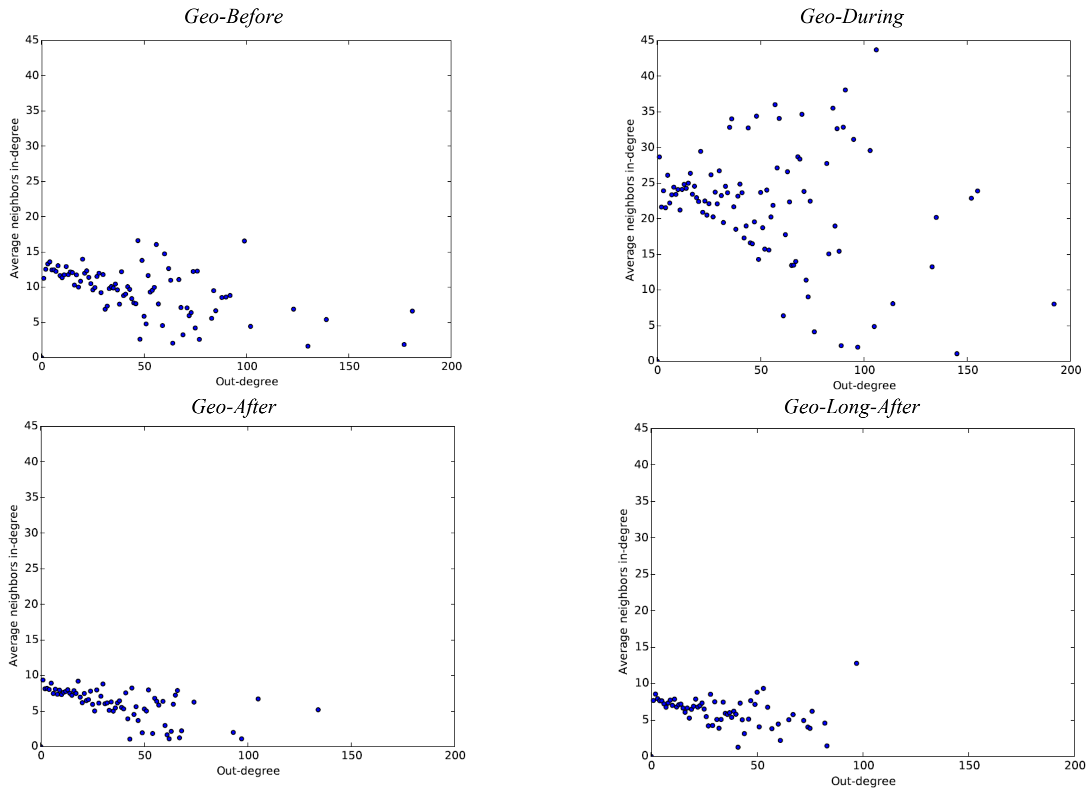

Figure 5. Average neighbor in-degree vs. node out-degree.

finding the influential nodes based on their incoming edges, like the important original tweet authors. Based on this metric, the majority of the most important nodes of the *Geo-During* network tend to be local government authorities and the media, while those accounts are much less represented in the "most important nodes" list of the other time slices.

The strong presence of local government and media sources in the list for the *Geo-During* network is consistent with our earlier observation of more in-degree hubs in this network.

Geographic Patterns of Retweet Activity

Proportion of Geo-Vulnerable Sources within the Four Time Sliced Networks

We are interested in the geographic patterns of the retweet activity. One of the ways to glean the presence of such patterns is to calculate the proportion of geo-vulnerable nodes in the time sliced networks. Remember that the in-common retwitterers around which the time slice networks were constructed are all geo-vulnerable, since we extracted them from the *Geo-Vulnerable Contextual* dataset. Thus, the only nodes with potential for different locations are the authors of the original tweets.

Table 6 indicates that for each of the time sliced networks, there is an overlap between the authors of the original tweets and the retweet authors. Thus, a small portion of the core retwitterers play both roles: authority figures whose tweets are propagated by others, and propagators of information.

All the time sliced networks also display a consistently higher proportion of geo-vulnerable original tweet authors than non-geo-vulnerable sources. Such consistency suggests that the geo-vulnerable retwitterers are more likely to propagate messages created by other geo-vulnerable users.

This is especially the case in the time of disaster. According to Table 6, the percentage of geo-vulnerable source authors is rather consistent among the *Geo-Before*, *Geo-Short-After*, and *Geo-Long-After* networks (61-62%). However, in the *Geo-During* network this percentage of geo-vulnerable source authorship goes up to 68%.

Proportion of the Geo-Vulnerable Source Tweets

Going to the tweet-level analysis of location brings some challenges. First, only a small number of tweets are geolocated (1.1% in *Global Keyword* dataset) and an even smaller portion is located within the bounding box of the event (0.4%). Moreover, our analysis indicates that retweets are never geolocated; that is, the location of the person performing the retweet is not captured. Instead, Twitter passes along the geotag (if present) of the original tweet (the tweet being retweeted). Indeed, if you attempt to load a retweet in a web browser, Twitter simply redirects you to the original tweet. These "features" make location analysis based on the retweet's coordinates impossible.

	Geo-Before	Geo-During	Geo-Short-After	Geo-Long-After
Original Twitterers	41,755	49,017	38,326	35,094
Overlap with in-common Twitterers	330	478	274	253
Geolocated Original Twitterers	1,499	1,838	1,477	1,380
Geo-Vulnerable Original Twitterers	62.58%	68.44%	62.42%	61.16%

Table 6. Source tweet author details. (Fraction of geo-vulnerable sources stat. sig. with χ^2=23.49, p<0.0001).

On the other hand, recall that we consider users to be geo-vulnerable if they produce at least one geotagged tweet within the boundary in the time frame of interest. This procedure produces many more tweets whose authors are considered geo-vulnerable than the actual tweets with geo-vulnerable geographical coordinates—and this larger set now includes retweets. Thus, in the rest of this analysis we consider tweets to be in the geographical area of interest if they were authored by a geo-vulnerable twitterer.

The geo-vulnerable twitterers in both the *Geo-Vulnerable Contextual* and *Geo-Vulnerable Keyword* dataset retweeted more tweets from geo-vulnerable authors than from non-geo-vulnerable. This is true for both overall time period of interest—Oct 13-Nov 30—and the *Geo-During* timeframe. For the *Geo-Vulnerable Keyword* dataset, the percentage remained essentially unchanged regardless of the time period (Figure 6). For *Geo-Vulnerable Contextual* dataset, however, the percentage of original tweets by geo-vulnerable authors increased noticeably during the disaster time period. This suggests that during the disaster, the Twitter conversations of geo-vulnerable authors tend to favor the local sources more strongly.

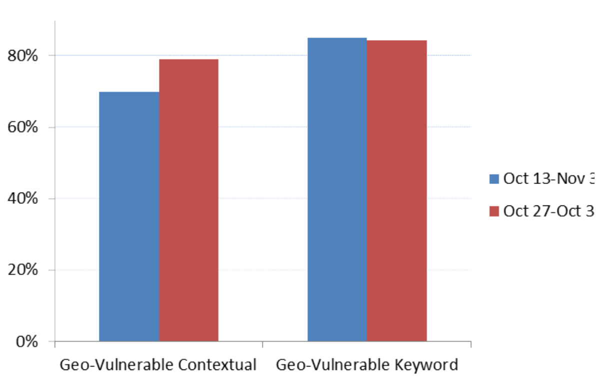

Figure 6. Percentage of source tweets in geo area of interest. (Stat. sig. with χ^2=827.11, p<0.0001)

Since *Geo-Vulnerable Contextual* dataset contains all the tweets of the geo-vulnerable users, some discussions and retweets might be unrelated to Hurricane Sandy, especially on the longer time frame. During the disaster, it is expected that more of the contextual tweets would be focused on Sandy and related local issues, making the percentage of original tweets by geo-vulnerable authors higher. The higher fraction of source tweets by the geo-vulnerable authors in the more on-topic *Geo-Vulnerable Keyword* set is consistent with the above intuition, supporting a hypothesis that the disaster-related geo-vulnerable retweet activity tends to favor the local sources more strongly.

Retweet Count Distributions of Various Populations

Now we turn to the three retweet count metrics we explained earlier, to see in another form how Global and Geo-Vulnerable populations retweet each other.

Retweet Distributions: The Geo-Vulnerable Retweeting Other Geo-Vulnerable Tweets vs Global Users Retweeting Global Tweets

The geographic patterns discussed above suggest that we might glean some insights into the retweeting behavior of the most affected geo-vulnerable twitterers by comparing their retweet count distributions to those of global twitterers. These distributions show how frequently we observe the tweets with a particular retweet count. We do not include tweets with a retweet count of 0 here, since though these tweets might be contentful, they do not contribute to the collective situational awareness because of their very limited audience, which is mostly confined to a twitterer's followers.

The distribution of *Geo-Vulnerable/Geo-Vulnerable* retweet counts on loglog scale looks significantly different from the distribution of *Global/Global* retweet count (Figure 7). The negative slope of the linear fit to the retweet frequencies of the former is less steep than that of the latter, producing a more heavy-tailed distribution with more density dispersing from the origin. The histogram inset in Figure 7 makes this

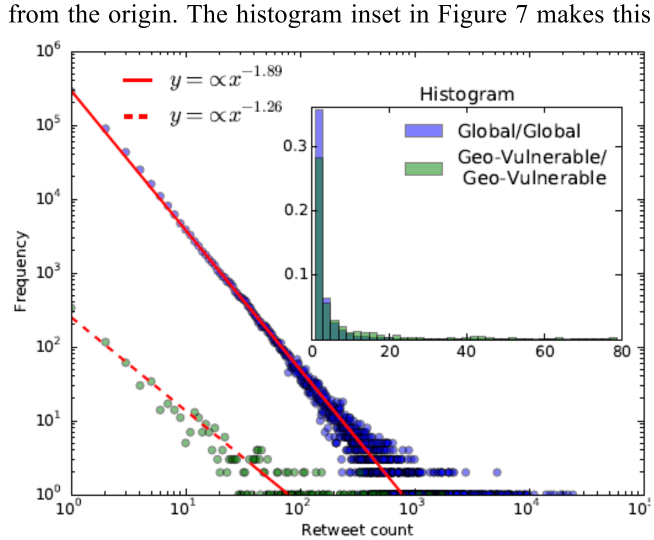

Figure 7. Retweet count distributions.

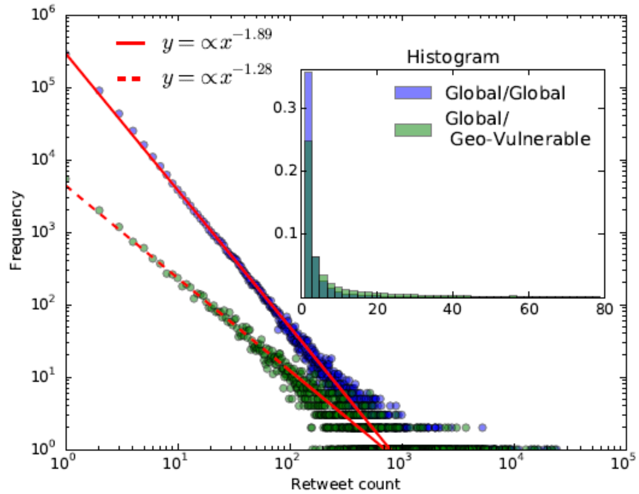

Figure 8. Retweet count distributions.

more visually apparent: the *Geo-Vulnerable/Geo-Vulnerable* distribution shows more density at retweet counts greater than 10, suggesting that the geo-vulnerable tweets get retweeted 10 or more times by the geo-vulnerable users more frequently than the global tweets get retweeted 10 or more times by all users. Specifically, we can see that the tweets with 10-80 retweets are visibly over-represented in *Geo-Vulnerable/Geo-Vulnerable* distribution, compared to the *Global/Global* histogram.

Thus, it seems that certain geo-vulnerable tweets might offer something especially useful to the discussion, making them more appealing to the geo-vulnerable users, and thus more retweeted. This makes the geo-vulnerable tweets with higher retweet counts (at 10-80 retweets) over-represented in the retweet count distributions for the geo-vulnerable users. This supports earlier findings from a smaller disaster much earlier in Twitter's life: tweets from the geo-vulnerable might be more useful for other geo-vulnerable users. Therefore, the geographic similarity might help us derive the most useful tweets. However, the geolocated tweets comprise only a small percentage of all the tweets (about 1.1%) and geo-vulnerable tweets make up an even smaller portion (0.4%). Hence it would be very helpful to identify the locally useful tweets without relying on their locality as an identifying marker, which we discuss in further detail next.

Retweet Count Distributions: The Geo-Vulnerable Retweeting Global Tweets vs Global Users Retweeting Global Tweets

To move away from using location as an identifying characteristic, we can look at the distributions of geo-vulnerable users retweeting tweets from the global keyword data set, not just the geo-vulnerable tweets. Comparing the loglog distribution of *Global/Geo-Vulnerable* retweet counts to the *Global/Global* distribution, we again observe a less steep linear fit for the former, suggesting that this distribution decreases more slowly and thus produces more rare events—tweets with larger retweet counts (Figure 8).

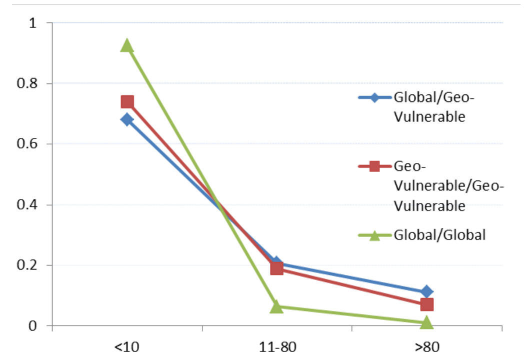

Figure 9. Fraction of tweets with 1-10, 11-80, and >80 retweets (Stat. sig, with $\chi^2 = 18876.38$, p<0.00001).

The fact that this distribution is more heavy-tailed is more visually obvious from the inset of Figure 8, where we see the *Global/Geo-Vulnerable* histogram diffuse density away from the origin and over-represent tweets with higher retweet counts in comparison to the *Global/Global* distribution. Specifically, the tweets with 10-80 retweets are again visibly over-represented, even to the higher degree than we observed for geo-vulnerable users propagating geo-vulnerable tweets.

In summary, for all the retweet counts, the retweet patterns of the geo-vulnerable users seem to be quite different from the global users (whether the geo-vulnerable users are retweeting the geo-vulnerable or global sources). Figure 9 provides the overview of these differences.

Thus, we can conclude that the tweets that end up with the higher retweet counts and hence over-represent those counts in the distributions are propagated by the geo-vulnerable users more, not necessarily because of their locality but because of some other aspect of the tweets. There are many aspects of the tweet that might motivate geo-vulnerable users to retweet it—informational utility and social conformity are two obvious contenders, especially from the perspective of retweet as an informal recommendation system (either for content or its author) [40]**.** Though we cannot fully disentangle these motivations in this analysis, earlier research [37] suggests that local utility of the content is likely to play a role in these tweets being retweeted more by geo-vulnerable users than we would expect from the retweeting patterns of Twitter's general population of users.

Content

Overrepresented tweets in the Global/Geo-Vulnerable count

To test the hypothesis that certain tweets gain more retweets based on usefulness of their content to the geographically-vulnerable, we perform a content analysis of the tweets. We selected a uniform random sample of all the tweets with *Global/Geo-Vulnerable* retweet count larger than zero (17K), since the number of tweets was too large to manually code for the presence of local utility. We focused on this retweet count because the global source of

	1-10 Retweets	11-80 Retweets	>80 Retweets
Tweets	11,518	3,503	1,885
Sample	1,147	353	191
% with Local Utility	38.92%	54.83%	36.84%

Table 7. Proportion of tweets with locally useful content in Global/Geo-Vulnerable retweet samples (Fraction with local utility stat. sig. with χ^2=30.11, p<0.0001).

the tweets ensures that over-representation of certain tweets is not due to their geographic origin.

We divided the sample tweets into three groups: 10 tweets and below, 11-80 tweets, and 81 and up. We qualitatively coded the three samples with a simple binary flag indicating whether or not the tweet contains locally useful information. We considered the content to be locally useful if it provided practical information on the state of affairs, such as exact weather and path of the hurricane prediction, notification of road and school closures, public transportation announcements, declaration of state of emergency, concrete opportunities to help in the recovery and so on. We distinguished those specific and locally-applicable tweets from the general expression of fear, awe, and disbelief and from text or images that provide a large-scale overview of the event. Table 7 summarizes the findings from this content analysis.

Tweets in all three samples ranged in their content from the local utility that aids in situational awareness to the broad appeal of the bird's eye view "abstract" of the event [37], including jokes and other memes. However, the 11-80 retweet subsample included much higher proportion of the tweets with locally-useful, actionable information compared to the other two samples. Clearly, the boundary cutoffs between subsamples are somewhat arbitrary, as all the samples had a varied mixture of content with local utility and broad appeal, and would be better represented by a continuum rather than discrete thresholds. However, these thresholds we empirically obtained from the retweet count distributions offer us a reasonable starting point for finding locally-useful tweets where they are most highly concentrated.

We did not find a one-to-one direct relationship between the local utility and retweet count, as evidenced by numerous locally-useful tweets we found in 1-10 retweet category. Our basic content analysis suggests that there might be some non-textual features that impede or promote a tweet's retweetability [41]. For example, many locally-useful tweets in the 1-10 retweet sample were hard to read due to overabundance of mentions. On the other hand, the fact that retweet count distributions have considerable densities above the retweet count of 80 while the local utility decreases at this point, suggest that other factors, such as author popularity, social conformity, and imitation might be in play. The data concur: 93.58% of the original authors of tweets with more than 80 retweets are popular twitterers, as operationalized by having a thousand followers or more (a statistically significantly higher proportion than for tweets with 11-80 retweets). Qualitative analysis of over-80-retweet tweets shows that celebrities and the media have a very strong presence among authors. In future work, we plan to explore how various non-content features affect the retweet potential for the tweets with local utility, and conversely what structural features characterize the well-retweeted locally-useful and non-useful tweets.

CONCLUSION

The purpose of this research is explain how Twitter activity by those who are geographically affected by a disaster differs or not from the general global reaction. If there are differences, then we know that victims turn to social media for different reasons than the general population. To address these questions, we had to carefully manage a large set of data, comparing retweet behavior across populations and time slices, which can make written explanations difficult, but it makes results more dependable.

In summary, our major findings are that geographically vulnerable twitterers propagate more information during the disaster period than before or after. They can also be both the sources and propagators of information. In doing this, geographically vulnerable twitterers have denser interconnected retweet networks during disasters than before or after. Social network "hubs," especially those based on in-degree, are present in higher numbers during the disaster period than before or after. It also appears that during the disaster period, local government authorities and the media are the most important nodes in comparison to their presence before or after the disaster.

The geographically vulnerable are more likely to propagate tweets from other geographically vulnerable users at any time, but this is especially prominent in disaster.

In addition, the geo-vulnerable retweet quite differently than the global population of twitterers who are interested in the same event. They propagate certain tweets considerably more than the general Twitter population, creating retweet distributions where rare events are more likely. Specifically, tweets that acquire 10-80 retweets make up a higher fraction of the total retweet activity of the geo-vulnerable population, and we see from qualitative analysis of these tweets that the geo-vulnerable select from the global twitterverse and retweet more are more likely to have some kind of local utility.

Though the popularity of social media is hard to deny, some still question its impact and import during disaster response because resource allocation decisions, governance policy, and even life-or-death actions are at issue. This research shows that those who are in the geographic area of effect relate to social media content differently once a disaster strikes, and they relate to it differently from the general population that attends to the event. They tend to propagate information from other geographically-vulnerable people, and focus the bulk of their retweeting activity on the tweets containing locally-useful information.

These findings provide evidence for moving forward on practice- and policy-making initiatives that address the role of social media in disaster emergency response. Technology designers might also be influenced by the needs of geographical neighbors when we think about future social computing innovations. Finally, those who analyze social media data for basic and applied science purposes may employ some of these findings to create sampling techniques to more quickly zero in on the content and propagators in the "big data" of crisis response.

ACKNOWLEDGMENTS

This research is funded through the US NSF grants IIS-0910586 and AGS-1331490. We thank our Project EPIC colleagues for support and critique. We especially thank our reviewers for their insightful comments and assistance.

REFERENCES

1. Abel, F., Hauff, C., Houben, G. J., Stronkman, R., & Tao, K. Semantics+ filtering+ search= twitcident. exploring information in social web streams. In *Proc. of the 23rd ACM Conference on Hypertext and social media.* ACM (2012), 285-294.
2. Ackerman, M. S. Augmenting organizational memory: a field study of answer garden. *ACM Transactions on Information Systems (TOIS)*, *16*(3), (1998), 203-224.
3. Anderson, K. M., & Schram, A. Design and implementation of a data analytics infrastructure in support of crisis informatics research (NIER track). In *Proc. of the 33rd International Conference on Software Engineering.* ACM (2011), 844-847.
4. Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. The role of social networks in information diffusion. In *Proc. of WWW,* (2012), 519-528.
5. Blake, E. S., Kimberlain, T. B., Berg, R. J., Cangialosi, J. P., & Beven II, J. L. Tropical Cyclone Report: Hurricane Sandy. *National Hurricane Center*, *12*, (2013).
6. Boersma, F. K., Groenewegen, P., & Wagenaar, P. Emergency Response Rooms in Action: an ethnographic case-study in Amsterdam. In *Proc. of ISCRAM* (2009), 1-8.
7. Caragea, C., Squicciarini, A., Stehle, S., Neppalli, K., & Tapia, A. Mapping Moods: Geo-Mapped Sentiment Analysis During Hurricane Sandy. In *Proc. of ISCRAM* (2014).
8. Clauset, A., Shalizi, C. R., & Newman, M. E. Power-law distributions in empirical data. *SIAM review*, *51*(4), (2009), 661-703.
9. Coviello, L., Sohn, Y., Kramer, A. D., Marlow, C., Franceschetti, M., Christakis, N. A., & Fowler, J. H. Detecting Emotional Contagion in Massive Social Networks. *PloS one*, (2014), *9*(3), e90315.
10. Dashti, S., Palen, L., Heris, M. P., Anderson, K. M., Anderson, S., & Anderson, S. Supporting Disaster Reconnaissance with Social Media Data: A Design-Oriented Case Study of the 2013 Colorado Floods. In *Proc. of ISCRAM* (2014).
11. De Choudhury, M., Monroy-Hernandez, A., & Mark, G. Narco emotions: affect and desensitization in social media during the mexican drug war. In *Proceedings of the CHI*, (2014), 3563-3572.
12. FEMA (July 1, 2013). Hurricane Sandy FEMA After Action Report, http://www.fema.gov/media-library-data/20130726-1923-25045-7442/sandy_fema_aar.pdf
13. Gibbs, L., & Holloway, C. Hurricane Sandy After Action: Report and Recommendations to Mayor Michael R. Bloomberg. *Hurricane Sandy After Action: Report and Recommendations to Mayor Michael R. Bloomberg, The City of New York, New York, NY*, 36, (2013).
14. Goldstein, B. E. (Ed.). *Collaborative Resilience: Moving through Crisis to Opportunity*. MIT Press. (2012).
15. Hiltz, S. R., Kushma, J., & Plotnick, L. Use of Social Media by US Public Sector Emergency Managers: Barriers and Wish Lists. In *Proceedings of ISCRAM* (2014).
16. Holland, P. W., and Leinhardt, S.. "A method for detecting structure in sociometric data." *American Journal of Sociology* (1970): 492-513.
17. Hughes, A. L., & Palen, L. Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, *6*(3), (2009), 248-260.
18. Hughes, A. L., Denis, L. A. S., Palen, L., & Anderson, K. M. Online Public Communications by Police & Fire Services during the 2012 Hurricane Sandy. In *Proc.* of *CHI* (2014), 1505-1514.
19. Imran, M., Castillo, C., Lucas, J., Patrick, M., & Rogstadius, J. Coordinating human and machine intelligence to classify microblog communications in crises. *Proc. of ISCRAM* (2014).
20. Keegan, B. C. Breaking news on wikipedia: dynamics, structures, and roles in high-tempo collaboration. In *Proc. of CSCW Companion,* (2012), 315-318.
21. Keegan, B., Gergle, D., & Contractor, N. Hot off the wiki: dynamics, practices, and structures in Wikipedia's coverage of the Tōhoku catastrophes. In *Proc. of the 7th International Symposium on Wikis and Open Collaboration,* (2011), 105-113.
22. Keegan, B., Gergle, D., & Contractor, N. Staying in the loop: Structure and dynamics of Wikipedia's breaking news collaborations. In *Proc. of the 8th Annual*

International Symposium on Wikis and Open Collaboration, (2012), 1-10.

23. Kwak, H., Lee, C., Park, H., & Moon, S. What is Twitter, a social network or a news media?. In *Proc. of WWW,* (2010), 591-600.
24. Lazo, J. K., Morss, R. E., & Demuth, J. L. 300 BILLION SERVED. *Bulletin of the American Meteorological Society*, *90*(6), (2009).
25. Lerman, K., & Ghosh, R. Information Contagion: An Empirical Study of the Spread of News on Digg and Twitter Social Networks. In *Proc. of ICWSM*, *10*, (2010), 90-97.
26. Mileti, D. S. *Disasters by design: A reassessment of natural hazards in the United States*. National Academies Press. (1999).
27. Morss, R. E., & Hayden, M. H. Storm Surge and "Certain Death": Interviews with Texas Coastal Residents following Hurricane Ike. *Weather, Climate & Society*, *2*(3), (2010), 174-189.
28. Morstatter, F., et al. "Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose." *ICWSM*. 2013.
29. Palen, L. Frontiers in Crisis Informatics. Keynote Address at *ISCRAM*, (2014).
30. Palen, L., Vieweg, S., Liu, S. B., & Hughes, A. L. Crisis in a networked world features of computer-mediated communication in the April 16, 2007, Virginia Tech Event. *Social Science Computer Review*, *27*(4), (2009), 467-480.
31. Perng, S-Y., Buscher, M., Wood, L., Halvorsrud, R., Stiso, M., Ramirez, L. & Al-Akkad, A. Peripheral response: microblogging during the 22/7/2011 Norway attacks, *International Journal of Information Systems for Crisis Response and Management*, (2013), 41-57
32. Powell, J. W. An introduction to the natural history of disaster. *Univ. of Maryland: Disaster Research Project,* (1954).
33. Qu, Y., Huang, C., Zhang, P., & Zhang, J. Microblogging after a major disaster in China: a case study of the 2010 Yushu earthquake. In *Proc. of CSCW,* (2011), 25-34.
34. Romero, D. M., Meeder, B., & Kleinberg, J. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc. of WWW,* (2011), 695-704.
35. Shelton, T., Poorthuis, A., Graham, M., & Zook, M."Mapping the data shadows of Hurricane Sandy: Uncovering the sociospatial dimensions of 'big data'." *Geoforum* 52 (2014): 167-179.
36. Soden, R., & Palen, L. From Crowdsourced Mapping to Community Mapping: The Post-Earthquake Work of OpenStreetMap Haiti. In *11th International Conference on the Design of Cooperative Systems (COOP),* (2014).
37. Starbird, K., & Palen, L. Pass it on?: Retweeting in mass emergency. In *Proc. of ISCRAM* (2010), 1-10.
38. Starbird, K., & Palen, L. Voluntweeters: Self-organizing by digital volunteers in times of crisis. In *Proc. of CHI,* (2011), 1071-1080.
39. Starbird, K., Maddock, J., Orand, M., & Mason, R. M. Rumors, False Flags, and Digital Vigilantes: Misinformation on Twitter after the 2013 Boston Marathon Bombing. In *Proc. of iConference,* (2014).
40. Starbird, K., Muzny, G., & Palen, L. Learning from the crowd: Collaborative filtering techniques for identifying on-the-ground Twitterers during mass disruptions. In *Proc. of ISCRAM,* (2012), 1-10.
41. Suh, B., Hong, L., Pirolli, P., & Chi, E. H. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *IEEE second international Conference on Social computing,* (2010), 177-184.
42. Verma, S., Vieweg, S., Corvey, W., Palen, L., Martin, J. H., Palmer, M., & Anderson, K. M. NLP to the Rescue?: Extracting "Situational Awareness" Tweets During Mass Emergency. In *Proc. of ICWSM,* (2011), 385-392.
43. Wasserman, S., and Faust, K.. "Social Network Analysis: Methods and Applications." *Cambridge University Pr*ess, (1994).
44. White, J. I., Palen, L., & Anderson, K. M. Digital Mobilization in Disaster Response: The Work & Self-Organization of On-Line Pet Advocates in Response to Hurricane Sandy. In *Proc. of CSCW,* (2014).
45. Wulf, V., Rohde, M., Pipek, V., & Stevens, G. Engaging with practices: design case studies as a research framework in CSCW. In *Proc. of CSCW,* (2011), 505-512.
46. Yonetani, M., & Morris, T. (2013). Global Estimates 2012: people displaced by disasters. IDMC.
47. Zhang, F., Morss, R. E., Sippel, J. A., Beckman, T. K., Clements, N. C., Hampshire, N. L., & Winkley, S. D. An In-Person Survey Investigating Public Perceptions of and Responses to Hurricane Rita Forecasts along the Texas Coast. *Weather & Forecasting*, *22*(6), (2007), 1177-1190.