Sustainable and Resilient Infrastructure

ISSN: 2378-9689 (Print) 2378-9697 (Online) Journal homepage: http://www.tandfonline.com/loi/tsri20

Embracing human noise as resilience indicator: twitter as power grid correlate

Nicolas LaLone, Andrea Tapia, Christopher Zobel, Cornelia Caraega, Venkata Kishore Neppalli & Shane Halse

SUSTAINABLE AND RESILIENT INFRASTRUCTURE, 2017
VOL. 2, NO. 4, 169–178
https://doi.org/10.1080/23789689.2017.1328920

Taylor & Francis
Taylor & Francis Group

Embracing human noise as resilience indicator: twitter as power grid correlate

Nicolas LaLone[a], Andrea Tapia[a], Christopher Zobel[b], Cornelia Caraega[c], Venkata Kishore Neppalli[c] and Shane Halse[a]

[a]Information Sciences and Technology, Pennsylvania State University, University Park, PA, USA; [b]College of Business, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA; [c]Computer Science and Engineering, University of North Texas, Denton, TX, USA

ABSTRACT

There are typically two approaches for measuring disaster resilience: technically dynamic measures produced by sensors attached to physical objects and socially static metrics that engage demographic indicators within a given geographic location. Although these approaches allow resilience to be represented before and after disruption, it can be difficult to measure resilient behavior *during* an event. We propose that social media data can be used to nowcast the ongoing state of critical infrastructure during a disaster. Through an analysis of tweets made during Hurricane Sandy and power outage data obtained after the event, we find that tweets that mention power, utility, or electricity were correlated with loss of power. We conclude with a discussion of barriers to realizing this concept.

ARTICLE HISTORY
Received 1 August 2016
Accepted 13 March 2017

KEYWORDS
Critical infrastructure; resilience; social media; power grid

Introduction

In this paper, we discuss the concept of disaster resilience as it relates to the noise of humanity – the data we create. Resilience itself is a somewhat noisy concept in that it has been a topic of interest for many decades for a variety of disciplines. Resilience gained a special significance in the United States when the 9/11 Commission issued a homeland security oriented goal to create a 'stronger, safer, and more resilient' United States after the World Trade Center attacks (Kahan, 2015; Napolitano, 2011). Resilience became a buzzword within development spaces as different groups worked on both defining and measuring the concept (Béné, Headey, Haddad, & von Grebmer, 2015; Grünewald & Warner, 2012; Winderl, 2015). The term was even said to have surpassed the popularity of sustainability (Editor, 2012). Since rising to buzzword status, the definition and parameters of resilience has quickly become difficult to parse yet work on understanding resilient behavior continues unabated and is becoming increasingly interdisciplinary in nature.

The most basic definition of resilience is simply a system's ability to absorb change (Holling, 1973). However, definitions built from this starting point are extremely varied. Resilience has been measured as a property (Gunderson, 2000), a capacity (Norris, Stevens, Pfefferbaum, Wyche, & Pfefferbaum, 2008; Olsson, 2003), an ability (Bruneau & Reinhorn, 2006; Council, 2012), and a process (Jacelon, 1997), yet no single definition has become universally accepted. Most of these perspectives can be applied in a disaster-related context to the human-designed, human-created cultural systems like neighborhoods, buildings, roadways, electrical grids, and supply chains that exhibit resilient behavior. Each of these systems generates significant amounts of information, with much of it often considered to be just noise or chatter between subsystems, between people, between organizations, and between objects. Although noise removal is often seen as important for better understanding underlying system behavior, it is important to recognize that the noise itself may often provide incredible detail about any given scale of measurement.

The ability of systems to rebound, given a disturbance, represents a unique opportunity for research and practice to interact. Through the proliferation of text-and picture-based communication like social media, text messaging, and other messaging services like Snapchat or YikYak, the people who inhabit and interact with these systems have begun to leave permanent, traceable, and mineable points of data that represent a geographic or topic-based location. In this paper, we present a proof-of-concept evaluation of these types of data as they relate to disaster resilience. We use posts on twitter,

CONTACT Nicolas James LaLone nick.lalone@gmail.com

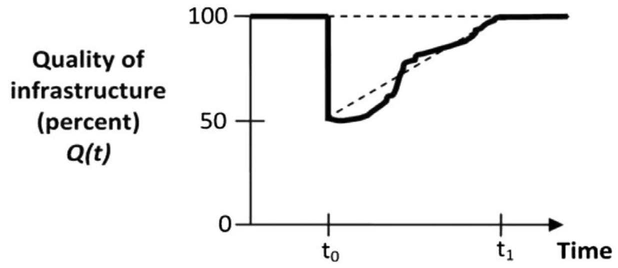

Figure 1. The original resilience triangle (adapted from (Bruneau et al., 2003)).

cross-referenced with the information produced by electrical grid technologies, as a means through which to use human chatter as a stand-in not only for resilience, but also for event detection.

This research is organized in four parts. First, we will discuss how we have used the term resilience in our research. Although there is general agreement on what a resilient system looks like, there is less agreement about the details of resilience, let alone how to measure it. As such, it is important to state what aspects of resilience we are considering. Next, we will discuss several different approaches for measuring resilience. Here, we see quantitative and qualitative, dynamic, and static approaches being deployed by the technical and social sciences, and we will discuss some of the limitations of each approach. Third, we will describe and deploy our methodological toolkit and evaluate the results of a specific case study. Finally, we will conclude with a discussion of the barriers to collecting data about resilient behavior and comment on how data could become more useful in the future.

The dimensions of resilience

In this research, we view resilience not as a singular concept, but as a multi-vocal, multidimensional one. This multidimensional approach was best described by McCreight (2010). The first dimension of resilience in this framework is that of personal, familial, and we would add, communal well-being. This dimension is focused heavily on by the social sciences. Strong personal ties within a region generally indicate the likelihood of resiliency for the people in that area (see Aldrich, 2012). Secondly, given that people are considered, organizations and institutions like schools or neighborhood volunteer organizations are also a dimension of resiliency (see Powley, 2009; Somers, 2009). The third dimension of resilience is that of commerce and production (Hill, Wial, & Wolman, 2008), followed by the fourth dimension: infrastructure (see McDaniels, Chang, Cole, Mikawoz, & Longstaff, 2008). These two dimensions encompass the economic viability of an area, since if the infrastructure of an area is damaged, the ability to conduct business is also impacted and both sales and production may suffer. They therefore fall under the purview of both economists and technically oriented researchers, although these researchers rarely work together. The last dimension of resilience is made up of public safety and government (see Swanstrom, 2008), and researchers in this domain typically examine the role that strong and weak governments play in the resilience of an area.

Approaches to measuring resilience

Taken together, the different dimensions of resilience provide an interconnected series of systems that each falls under the purview of a multitude of disciplines, industries, and interests. We find, however, that most efforts to measure the resilience of a community during and after a crisis have taken one of two approaches. The first is a focus on the *actual* physical loss – i.e. studying the extent to which infrastructure is initially damaged and the time taken to regain normal functionality of that infrastructure (Brauner et al., 2015; Bruneau & Reinhorn, 2007; Bruneau et al., 2003; Cimellaro, Reinhorn, & Bruneau, 2010b; Simpson, Lasley, Rockaway, & Weigel, 2010; Zobel, 2014). The second approach is a focus on the *potential* for resilient behavior – i.e. the sociological study of the static elements that comprise a community's capacity for resilient behavior, such as a home ownership, crime rate, medical capacity, and employment (Cutter, Burton, & Emrich, 2010; Cutter et al., 2008; Norris et al., 2008).

While the first approach considers the changing behavior of the system over time, its traditional focus on the physical response of the system does not easily lend itself to modeling the human behavior that is critical to a community's ability to remediate, repair, and respond to a disaster. In contrast, although the second approach measures these human qualities, it is static in nature and ignores the dynamic nature of crisis and recovery.

Dynamic measures of resilience

Studying the physical loss of infrastructure due to a disaster event involves choosing which aspect of that infrastructure to measure (roads, buildings, water systems, power systems, communication systems, etc.), and requires an understanding of, or at least an appreciation for, the complexity of the interrelationships between the different elements and the society that they have been created to support. As mentioned above, efforts to measure such physical infrastructure resilience typically include consideration of both the initial impact of a disaster event and the time that is then needed to recover from that event. It is because of this explicit recognition of changes over time that we refer to such approaches as 'dynamic.'

The relative extent to which both loss and recovery time are exhibited in a given system may be visually represented by the disaster resilience triangle (Bruneau et al., 2003), with the disaster event occurring at some time t_0 and recovery occurring at a later time t_1 (see: Figure 1). Resilience can be measured as the relative amount of retained infrastructure quality over time, or the area *beneath* the triangle, considered as a percentage of the total amount available if no disaster had occurred (Cimellaro, Reinhorn, & Bruneau, 2010a; Zobel, 2010). This concept has been extended by a number of different researchers, to incorporate characteristics like multidimensionality (Bruneau & Reinhorn, 2004), uncertainty (Cimellaro et al., 2010a), multiple sub-events (Zobel & Khansa, 2014), slow-onset behavior (Zobel & Khansa, 2012), and nonlinear recovery (Zobel, 2014). The tradeoffs between the two characteristics of dynamic resilience (size of impact and recovery time) also allow one to visualize the concept as a collection of hyperbolic curves (Zobel, 2010), and thus to characterize the tradeoffs more precisely even when faced with multiple disaster events (Zobel & Khansa, 2012).

Although the concept of 'quality of infrastructure' implies a strong focus on physical processes, the link between the physical and the social aspects of resilience also has been well recognized by researchers in this area. In their original work, Bruneau et al. (2003) discussed what they considered to be four interrelated dimensions of the resilience concept: technical, organizational, social, and economic resilience (TOSE). The *technical* dimension of resilience is associated with the ability of physical systems to resist and then recover from a disaster, and the *organizational* dimension refers to the ability of organizations to perform their duties during and after such an event. The *social* dimension of resilience is then specified to be the extent to which social systems are able to protect against and recover from the loss of critical services, and the *economic* dimension is associated with the ability to reduce both direct and indirect economic losses.

Despite this recognition of the multidimensional and interdependent nature of resilience, however, relatively little work has been done on expanding the social dimension, in particular, from a dynamic resilience perspective. At least in part, this is because social systems themselves are extremely complicated and thus present 'significant conceptual and measurement challenges' (Bruneau et al., 2003) that differ from those in purely physical systems.

Sociological resiliency measures

Perhaps due to the extreme complexity of combined social and physical systems, there has yet to be an empirical study that proves that increases in resources available to a given area actually has an impact on the time it takes for an area to recover (Aldrich, 2012). There are a number of aspects of any given community – social capital and the strength of ties – that are extremely difficult to quantify, let alone to track over time. This ties in to the first dimension of resiliency – personal or familial well-being.

When a unit of a particular area is intact, methods from the social sciences can be used to count various facets of that community's worth. Indicators such as average income, incarceration rates, school graduation rates, and other measures are used to indicate the general vulnerability of an area (Cutter et al., 2010; Norris et al., 2008; Rivera & Settembrino, 2013). Such indicators tend to be relatively static, however, in that their values don't change much over time, either because the data is naturally granular, or because it is only collected on an infrequent basis (i.e. if it is drawn from the U.S. Census). Furthermore, even if such data is generated on less than a 10-year time interval, it simply can be difficult or impossible to collect during a disaster event because of the need to focus resources elsewhere. Most socially oriented analyses occur after-the-event, therefore, as any analysis of social structures requires that structure be deployed in order to study it (see Norris, Tracy, & Galea, 2009).

Other aspects of a community, however, such as social capital, cannot easily be captured using simple indicators. This is significant because Aldrich (2012) found that existing social capital positively affected a given area's recovery in that it increased the likelihood that members of a community would volunteer to organize recovery efforts. Additionally, such social capital served as a means through which resilience itself was made manifest in that before a crisis hit, communities with strong social ties often helped each other prepare for the oncoming disruption event. Unfortunately, this also meant that communities with enormous social capital could impede the recovery efforts with lower social capital (Aldrich, 2012; Aldrich & Meyer, 2014).

Despite these difficulties, the work of sociologists provides necessary confounding work for any stable empirical measure that represents a given area. These confounding factors create situations wherein interdisciplinary work is pursued and the barriers between industry and academia are temporarily weakened. Unfortunately, this leaves a large gap of knowledge between resiliency measures before an event, and post-event analyses. The impact of the data sciences in trying to fill this gap has occurred in parallel to the birth and development of social media. The resulting sharing economy has helped create a means through which responders and resilience researchers can get real-time data from the field.

Data science resiliency measures

The iSchool movement has been building since the late 1980s. This movement is essentially the study of information, its structure, and its uses and it has developed concurrent to an ever increasing production of new data (Dillon, 2012). While this study is nothing new, the disposal of bounded discipline within the iSchools is not. As such, the information sciences have a unique ability to move between and among disciplines in order to assemble unlikely representations of data that can be used by any discipline. Within crisis management, there have been considerable efforts made to bridge the gap between pre-existing concepts of resilience and situational awareness during recovery. These efforts further confound the concept of resilience by adding a significant degree of granularity to the resilience triangle. It does so by offering a means through which to measure not only human efforts to bring an affected area back to pre-event status, but also how those events might actually slow or hasten the progress of recovery. Most of this new work has been pursued through the lens of social media.

In particular, there has been significant scholarship on microblogging, or posting information to twitter and other platforms like it. It was believed that due to the proliferation of smart phones and the growing ubiquity of social media that posts could be mined to provide significant amounts of situational awareness generated by those impacted by crisis-level events (De Longueville, Smith, & Luraschi, 2009). From this starting point, research began to examine the use and behaviors of social media users during the recovery and crisis events that followed. For example, Starbird and Palen (2010) examined the behaviors of passing along information to others. They found that re-posting, or re-tweeting posts from sources like the media or response organizations during a crisis had become commonplace for those seeking information about an event. While useful, there are numerous challenges to using social media data during a crisis (Hughes & Palen, 2009).

There are two issues within any microblogging or social media analysis. The first is finding reliable tweets given the lack of contextual clues to help identify rumors within a corpus of tweets (Mendoza, Poblete, & Castillo, 2010). In addition to rumors, truthfulness and deception is also of note (Tapia, Moore, & Johnson, 2013). However, each of the issues surrounding trustworthiness seems to be resolvable as the various ways products allow deception to proliferate become less powerful due to the tenure of a products existence (Palen, Vieweg, & Anderson, 2010; Palen, Vieweg, Liu, & Hughes, 2009; Tapia, LaLone, & Kim, 2014). The second issue is the quantification of performance. Twitter has consistently reduced researchers' ability to scrape tweets in real time. This has led to computational work-arounds for textual problems rather than blending quantitative and qualitative work (MacEachren et al., 2011). This has been the trend since that time. Researchers have applied more and more complex data mining and natural language processing to tweets in order to attempt to make these data useful and more quickly to responders (see: (Caragea et al., 2011; Caragea, Silvescu, & Tapia, 2016; Li et al., 2015; Truong, Caragea, Squicciarini, & Tapia, 2014; Vieweg, 2010)).

While useful, the gap in this knowledge is foundational. It remains unclear if these data actually correlate to situational knowledge on the ground or if the desire for these data to correlate has resulted in hopeful work. Most relevant to our research, Asur and Huberman (2010) worked to provide a different sort of stopgap showing that tweets could be used to indicate problem areas before crisis events occurred. We believe that it is possible for twitter to be used as indicator of resilience before, during, and after a crisis and our method has been constructed to show that twitter data does correlate to an event. This proof of concept is our effort to fill the gap of social media correlation to situational knowledge during a crisis-level event. After our method and analysis, we will discuss the numerous roadblocks to using social media like twitter during response and recovery.

Method

As this research is being conducted as a proof-of-concept, the methodology we chose is a standard correlation test with this research question:

> H_1: Tweets per hour that mention our keywords and Customers affected by Power Outages during Hurricane Sandy are positively correlated.

This research question encapsulates resilience as it relates to social media data by allowing us a simple way to detect power grid outages. These electrical grid data are typically found well after the fact on the Department of Energy's website and are often unknowable by the public unless the power company broadcasts the outage. Using social media, we can detect outages in real time.

Data collection

The data used in our experiment comes from two different data sources. The first set of data is from Twitter during Hurricane Sandy. This hurricane took place in October 2012 and is one of the deadliest and costliest hurricanes in the history of the United States. We collected tweets between 26 October 2012 and 11 November 2012 through the Twitter Streaming API. In total, we gathered 10,042,769 tweets that mention Hurricane Sandy.

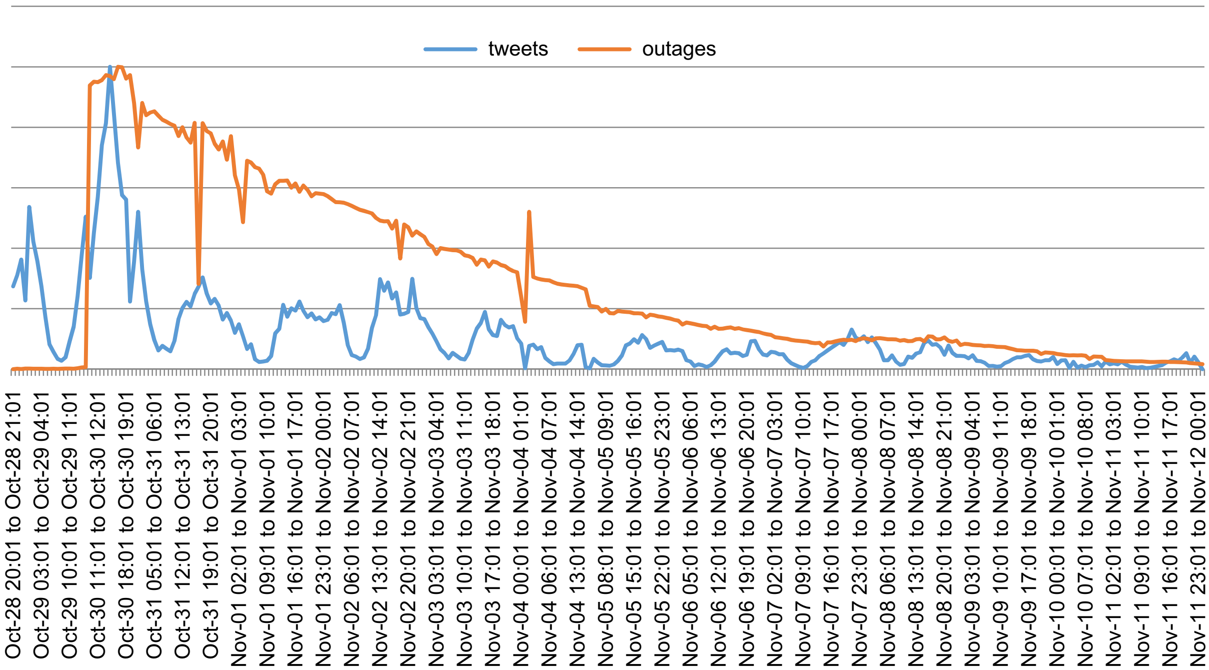

Figure 2. Tweets to outages over time that mention electricity or outages.

For this specific research, 553,689 tweets were selected based on the keywords 'power,' 'outage,' 'electri,' and 'utility.' Our second data-set is outage data from power companies in and around the New York, New Jersey, and Pennsylvania areas. There are 10,702,938 data points that reflect electrical grid readings. These data include county, number of outages, number of affected customers, and a timestamp. In order to correlate the data, the tweets and power data were grouped into one-hour increments. As the outage data-set contained multiple reports for each area within the hour segment, we used the average outage over all locations for each hour.

The following time segments from both the Twitter data and the corresponding outage data are missing:

- Oct-29 16:01 through Oct-30 08:01
- Oct-30 20:01 through Oct-30 23:01
- Nov-04 16:01 through Nov-05 04:01
- Nov-10 15:01 through Nov-11 00:01

It is unclear if the electric grid and Twitter were down during these times or if our crawler failed.

Descriptive statistics

After controlling for one-hour intervals, the Hurricane Sandy data was condensed to 295 time segments. On average, 1928 tweets containing our power-related keywords were sent per hour with a maximum of 17192 during the hour of 1401–1501 on 30 October 2012. This time marks when Hurricane Sandy had just passed over the eastern coast. The outage data contained 295 time segments as well. The average number of outages was 1867 with a maximum of 6516 being reported between 1601 and 1701 on 30 October 2012. There is a one-hour difference between Hurricane Sandy reaching the East Coast and power outage data being reported. One reason for this could be due to the lag between an outage and a report or that there were no outages reported for the first hour of Hurricane Sandy's presence in the affected area. The one-hour difference could also mean that using Twitter to search for power outages could result in faster information flow during a crisis event.

Correlation

The Pearson's correlation test applied to the entire data-set resulted in an *R-value of .57* and a significance value of $\rho < 2.2e\text{-}16$. This finding is both statistically significant and supports a moderate to strong correlation between the two data-sets. This correlation can be seen in Figure 1, which shows the number of tweets (in blue), and the reported outages (in red) normalized to the maximum.

The data was then truncated to include only observations starting on 30 October, when the abnormal power outages began to be reported. Because this eliminated the effect of the pre-impact, relatively steady-state power data, it resulted in the correlation between the two data-sets being increased to .71.

Finally, the cyclical behavior of the Twitter usage data over the course of a day was addressed by further calculating the daily average values of both tweets and power loss levels. When once again the data was restricted to the data range beginning on 30 October, the resulting correlation between the two data-sets increased once more to .86 (Figure 2).

Discussion

The advantages and potential of information about resilience and recovery efforts that can be derived from social media outweighs the difficulty of the work that must be performed in using it. Tapia, Bajpai, Jansen, Yen, and Giles (2011) have already described pockets of use of social media data and illustrated both the frustration and hope of one day being able to use these data effectively. Work like this is becoming more commonplace as the ability to work with textual data at scale increases. The initial correlation analysis provided above is a first step in the direction of being able to assess the resilience of communities in terms of both people and infrastructure.

This current effort involved a *post hoc* analysis of the two data-sets. While this is may not be the real-time analysis that is required by the time-sensitive decisions made by crisis responders, it serves as a proof-of-concept that can provide responders with filtered, useful data. These filtered data provide a snapshot of the community developing around and within an event. This allows a narrative representing an affected zone to emerge organically, and to be observable via social media. By allowing for detection of the state of infrastructure through Twitter, social media data can become much more relevant and valuable within the domain of crisis response and management.

Six barriers to data collection for resilience work

While the analysis above provides an initial indication of the value of social media data in this context, there are barriers to working with the information produced by and about infrastructure systems. In fact, we count six specific barriers to data collection that affect the potential further development of such efforts.

Barrier one – privatized data

Within the United States, power outage data are typically released to the public in one of several ways. First, users can go to their power company's webpage. On this page is a map through which real-time outage data is displayed. Second, many power companies release daily updates on power outages through official press releases or on platforms such as Facebook (Zobel, 2013, 2014). As a third option, interested individuals can go to the United States Energy Information Administration's Electricity Data page. Here, interested individuals can find outage data that is aggregated by month. In the first two cases, the data is updated in real time, or almost real time. This is useful, but only for those who have power or access to data services for telephones. The data in the third case has a relatively significant time lag that makes it unusable for operational decision-making during disaster recovery.

In countries like the United Kingdom, electrical grid data is readily available and all manner of detail can be gleaned from pages like the U.K. National Grid status web page. This citizen-run monitor of energy usage data provides usage by five-minute intervals but falls short of providing concise outage reports. Within the U.K., the proliferation of smart grid technologies has increased the number of annual meter reads from 75 million reads per year to over 120 billion (Raftery, 2013). With meter readings being taken every 30 min, access to this information would provide an enormous potential to pair with tweets about power outage data. However, as in the United States, privatized data due to corporate and not government controlled energy production facilities has restricted these data.

Barrier two – incompatible scales

From the example of the electrical grid, we see a rather robust 30-min interval for electrical grid data, per house, per community, and per municipality. Efforts to quantify resilient behavior, however, are often criticized for not pre-defining what is meant by a term such as 'city,' let alone for not using a more specific analyzable unit (Vale, 2014). Even if an API were available to support combing through such data, the resulting need to find a consistent and appropriate scale provides a unique hurdle to overcome for resilience and recovery personnel. If these data were then paired with tweets, for example, as responders looked for mentions of transformer or pole damage, then the issues of finding and maintaining scale in real time would become even further complicated. Furthermore, it is nearly impossible to institute a standard through which to establish routine procedures due to the variability of the term, 'local.' This falls back to problems of computational speeds and needs superseding the needs of responders and victims of crisis-level events (MacEachren et al., 2011).

One solution to these issues is to feed these data into an application that is then engaged in a similar fashion as a citizen science project (see Tapia, LaLone, MacDonald, Priedhorsky, & Hall, 2014). Here, interested individuals with surplus time and labor can donate their time and efforts to comb through social media data in real time. In much the same way as the Boston Marathon Bombing, these incompatible scales could be overcome with the sheer willingness of the so-called digilantes (Nhan, Huey, & Broll, 2015). While the accuracy and verifiability of these digilantes is often called into question (Tapia et al., 2014), their labor and their accuracy can be mitigated by software affordances, and additional efforts surrounding training and quality standards can improve both situational awareness of responders but also the accuracy of these virtual bystanders (Vieweg, 2012).

Barrier three – incompatible data standards

Another barrier to effective data collection is the issue of incompatible data standards. Working with robust datasets requires an enormous amount of preparation. Data must be scraped, formatted, concatenated or organized as needed, and it must be cleaned, and possibly trained against a pre-existing data-set that has also gone through these processes. Given the nature of the disconnectedness of academia and industry, the likelihood that any one person could do all of these tasks in real time is unlikely or impossible. While this is a reflection of the 'newness' of the information sciences, it represents a massive barrier for the collection and use of data for deriving resilience and recovery measures.

The aging infrastructure of the United States has recently received a boost due to the National Infrastructure Improvement Act (NIIA) (Doyle et al., 2008). Unfortunately, the terms of this act do not include sensor or data standard requirements that are conducive to future work with the information that these structures could provide. Furthermore, although the NIIA was passed in 2006, recent events in Flint, Michigan concerning the quality of drinking water indicate that this bill has not yet achieved the impact that it was intended to provide. It remains to be seen if the predictable lag of technology to law will ever reach a point wherein these structures are constructed with forethought instead of afterthought from 50 years prior.

Barrier four – geolocation, privacy, and the terms of service

The fourth barrier in our discussion also falls into the category of legal structures in that the Terms of Service for social media platforms are not conducive to working with the data that they produce. As internet-based services gain more tools to monetize their user's behaviors (Gerlach, Widjaja, & Buxmann, 2015), there are valid concerns for privacy and transparency of user data. These privacy concerns also present a barrier to using social media to answer vital questions about affected residents within an affected area. In some circumstances, using something like mobile device tracking can let responders know where people are inside of an area. However, the use of technologies that assume the proliferation of devices or that assume that these devices are evenly spread across social strata, may end up hindering response efforts or increasing the vulnerability of already vulnerable populations (Taylor, 2015).

With services like Twitter, geo-locating tweets are an optional service that is not that widespread. Many accounts will not have geo-location, either due to desire for privacy or to lack of literacy about the need to turn it on during a crisis. Without geo-location, reliance on twitter data in and of itself has not produced significant results. In an examination of relevant tweets presented during four different crisis responses, Vieweg (2012) found that it was difficult to identify tweets that provide any sort of awareness of what is occurring within an event at any given time. Twitter itself does not currently offer crisis response practitioners any type of access that it does not already offer regular users, and it is unclear if this will change in the future.

Barrier five – better statistical measures

The information sciences and the iSchool movement have only been around for around 30 years. The ability for scientists of any type to interact with their data is still extremely limited; evidenced by the nearly infinite ways that researchers have tried to use the data they possess. The fifth barrier to using data like this for crisis and resilience work, therefore, is the data itself and the newness of the methods being used to engage those data.

Sentiment analysis, for example, allows researchers to gain an understanding of how a corpus of text 'feels' (Pang & Lee, 2008). Although this technique currently has relatively limited utility, as researchers attempt to do more with the concept it is beginning to expand in usefulness. Filtering data is also gaining power as it becomes easier altering large data-sets given the ever-increasing nature of hardware speeds. Finally, without geo-location, researchers are left to attempt to glean information from the words people use in their tweets. This methodology is still in its infancy (Priedhorsky, Culotta, & Del Valle, 2014; Tasse & Hong, 2014).

Barrier six – polyvocality

The final barrier to using social media data efficiently and effectively is the most obvious – their polyvocal nature. Soden et al. (2015) have attempted to engage the inherent complexity of resilience by associating it with the name polyvocality. Within any city, within any neighborhood, within any area impacted by a disruption, many different voices combine to express the nature of loss and resilience. Each voice uses a different vocabulary while pursuing needs communication through media that might potentially be left out of the mainstream discussion (e.g. Ham radio, snapchat, YikYak, etc.). This is a barrier to using data to represent an area because while social media data can represent a variety of people within and surrounding an area, those who do not produce accessible data will not be represented. This extends also to infrastructure personnel who may not be represented locally.

Harkening back to discussions of social capital, we find that this is doubly confounding. If a group with a significant social capital is not participating in communication about resilience but is using their influence to pursue their own agenda, then it can become difficult to understand or to get support for resilience building or recovery efforts. Additionally, entire aspects of a geographic region could be marked with large gaps through which resources are being consistently sent. As such, a polyvocal means through which to engage the complexities of resilience requires multiple modes of data gathering, data dispersion, and communication. This becomes one of the most important barriers to engage, as every possible stakeholder of a given area should be given a voice.

Conclusions and future work

In the research literature, resilience is often quantified using one of two general approaches: the 'technically dynamic' approach of measuring the loss and recovery of physical system functionality over time, and the 'socially static' approach of capturing demographic indicators of the social capacity for resilient behavior in a given area. We believe that using social media as a means for collecting data about system behavior during an event can help to bridge the gap between these two approaches and help to capture information about other aspects of the socio-technical system, such as the strength of social capital. As a proof of concept to demonstrate this potential, we showed that social media usage correlated highly with electrical grid outages during Hurricane Sandy, as users in the affected areas began to discuss the electrical grid disruptions on Twitter when they occurred.

This correlation between the social and the physical is important because it can be used to connect the various stakeholders interested in response efforts. Responders can measure many of the physical aspects of disasters like earthquakes, hurricanes, or blizzards, but they typically cannot know in real time the effects that a disaster is having on the population. They may rely on eyewitness accounts after the fact from survivors, or on eyewitness information offered in real time, by reporters in the field or by those who are able to make contact by phone. Our case study demonstrates the potential for social media data to augment and strengthen situational awareness. It does this by showing that outage data posted well after the outage and social media posted in real time are correlated. If harnessed and deployed at the municipal level, it could help response organizations be more effective in their operations.

In order for a system that incorporates such data to be useful, however, several barriers to its potential effectiveness must be overcome, as discussed above. We urge those working on incorporating sensors into critical infrastructure (CI) to integrate their work with local response teams who will eventually need the data that will be produced. We urge those working on upgrading the electrical grid of the United States to consider the utility of their data to media outlets and first responders trying to gather information on the ground. We urge those information scientists working with data to continue to develop their methodologies in such a way as to allow anyone to use them. Most of all, we urge those working on defining and characterizing resilience to bridge their definitions and their measures with researchers from other disciplinary backgrounds and perspectives. Only by actively working to share information, and then applying the knowledge that we gain from access to that information, can we improve the ability to help people resist against and recover from the impacts of disasters.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the National Science Foundation CRISP Type 2/Collaborative Research Award (#1541155) titled: *Resilience Analytics: A Data-Driven Approach for Enhanced Interdependent Network Resilience*.

References

Aldrich, D. P. (2012). *Building resilience: Social capital in post-disaster recovery*. Chicago, IL: University of Chicago Press.

Aldrich, D. P., & Meyer, M. A. (2014). Social capital and community resilience. *American Behavioral Scientist, 59.2*, 254–269. 0002764214550299.

Asur, S., & Huberman, B. A. (2010). *Predicting the future with social media*. Paper presented at the Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on, Los Alamitos, CA.

Béné, C., Headey, D., Haddad, L., & von Grebmer, K. (2015). Is resilience a useful concept in the context of food security and nutrition programmes? Some conceptual and practical considerations. *Food Security,* 1–16.

Brauner, F., Münzberg, T., Wiens, M., Fiedrich, F., Lechleuthner, A., & Schultmann, F. (2015, May). *Critical infrastructure resilience: A framework for considering micro and macro observation levels*. Paper presented at the Proceedings of the international conference on information systems for crisis response and management – kristiansand, Norway.

Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O'Rourke, T. D., Reinhorn, A. M., … von Winterfeldt, D. (2003). A framework to quantitatively assess and enhance the seismic resilience of communities. *Earthquake Spectra, 19*, 733–752.

Bruneau, M., & Reinhorn, A. (2004, June 28–July). *Seismic resilience of communities-conceptualization and operationalization.* Paper presented at the Proceedings, International Workshop on Performance-based Seismic-Design, Bled, Slovenia.

Bruneau, M., & Reinhorn, A. (2006). *Overview of the resilience concept*. Paper presented at the Proceedings of the 8th US National Conference on Earthquake Engineering, San Francisco, CA.

Bruneau, M., & Reinhorn, A. (2007). Exploring the concept of seismic resilience for acute care facilities. *Earthquake Spectra, 23*, 41–62.

Caragea, C., McNeese, N., Jaiswal, A., Traylor, G., Kim, H. W., Mitra, P., … Jansen, B. (2011). *Classifying text messages for the haiti earthquake*. Paper presented at the Information Systems for Crisis Response and Management, ISCRAM, Lisbon.

Caragea, C., Silvescu, A., & Tapia, A. (2016). *Identifying informative messages in disasters using convolutional neural networks.* Paper presented at the The 13th International Conference on Information Systems for Crisis Response and Management, Kristiansand, Norway.

Cimellaro, G. P., Reinhorn, A. M., & Bruneau, M. (2010a). Framework for analytical quantification of disaster resilience. *Engineering Structures, 32*, 3639–3649.

Cimellaro, G. P., Reinhorn, A. M., & Bruneau, M. (2010b). Seismic resilience of a hospital system. *Structure and Infrastructure Engineering, 6*, 127–144.

Council, N. R. (2012). *Disaster resilience: A national imperative*. Washington, DC: The National Academies Press.

Cutter, S. L., Barnes, L., Berry, M., Burton, C., Evans, E., Tate, E., & Webb, J. (2008). A place-based model for understanding community resilience to natural disasters. *Global environmental change, 18*, 598–606.

Cutter, S. L., Burton, C. G., & Emrich, C. T. (2010). Disaster resilience indicators for benchmarking baseline conditions. *Journal of homeland Security and emergency management, 7*(1).

De Longueville, B., Smith, R. S., & Luraschi, G. (2009). *Omg, from here, i can see the flames!: A use case of mining location based social networks to acquire spatio-temporal data on forest fires.* Paper presented at the Proceedings of the 2009 international workshop on location based social networks, Seattle, WA.

Dillon, A. (2012). What it means to be an iSchool. *Journal of Education for Library and Information Science, 53*, 267.

Doyle, M. W., Stanley, E. H., Havlick, D. G., Kaiser, M. J., Steinbach, G., Graf, W. L., … Riggsbee, J. A. (2008). Aging infrastructure and ecosystem restoration. *Science, 319*, 286–287.

Editor, D. (2012). The international development buzzword of 2012. *Devex*. Retrieved from Devex website https://www.devex.com/news/the-international-development-buzzword-of-2012-80011

Gerlach, J., Widjaja, T., & Buxmann, P. (2015). Handle with care: How online social network providers' privacy policies impact users' information sharing behavior. *The Journal of Strategic Information Systems, 24*, 33–43.

Grünewald, F., & Warner, J. (2012, October). Resilience: buzz word or critical strategic concept? *Humanitarian Aid on the move*, 12

Gunderson, L. H. (2000). Ecological resilience-in theory and application. *Annual Review of Ecology and Systematics, 31*, 425–439.

Hill, E., Wial, H., & Wolman, H. (2008). *Exploring regional economic resilience: Working Paper*. Berkeley, CA: Institute of Urban and Regional Development.

Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics, 4*, 1–23.

Hughes, A. L. & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management, 6*, 248–260.

Jacelon, C. S. (1997). The trait and process of resilience. *Journal of Advanced Nursing, 25*, 123–129.

Kahan, J. H. (2015). Resilience Redux: Buzzword or Basis for Homeland Security. *Homeland Security Affairs*, 11.

MacEachren, A. M., Robinson, A. C., Jaiswal, A., Pezanowski, S., Savelyev, A., Blanford, J., & Mitra, P. (2011). *Geo-twitter analytics: Applications in crisis management.* Paper presented at the Proceedings, 25th International Cartographic Conference, Paris, France.

McCreight, R. (2010). Resilience as a goal and standard in emergency management. *Journal of homeland Security and emergency management, 7*(1), 1–7.

McDaniels, T., Chang, S., Cole, D., Mikawoz, J., & Longstaff, H. (2008). Fostering resilience to extreme events within infrastructure systems: Characterizing decision contexts for mitigation and adaptation. *Global Environmental Change, 18*, 310–318.

Mendoza, M., Poblete, B., & Castillo, C. (2010). *Twitter Under Crisis: Can we trust what we RT?* Paper presented at the Proceedings of the first workshop on social media analytics, New York, NY.

Napolitano, J. (2011). Progress Toward a More Secure and Resilient Nation. *Homeland Security Affairs 7, 10 Years After: The 9/11 Essays.* Retrieved February 10, 2016, from https://www.hsaj.org/articles/585

Nhan, J., Huey, L., & Broll, R. (2015). Digilantism: An Analysis of Crowdsourcing and the Boston Marathon Bombings. *British Journal of Criminology*, azv118.

Norris, F. H., Stevens, S. P., Pfefferbaum, B., Wyche, K. F., & Pfefferbaum, R. L. (2008). Community resilience as a metaphor, theory, set of capacities, and strategy for disaster readiness. *American Journal of Community Psychology, 41*, 127–150.

Norris, F. H., Tracy, M., & Galea, S. (2009). Looking for resilience: Understanding the longitudinal trajectories of responses to stress. *Social Science & Medicine, 68*, 2190–2198.

Olsson, P. (2003). *Building capacity for resilience in social-ecological systems*. Stockholm: Stockholm University

Palen, L., Vieweg, S., & Anderson, K. M. (2010). Supporting "everyday analysts" in safety- and time-critical situations. *The Information Society, 27*, 52–62.

Palen, L., Vieweg, S., Liu, S. B., & Hughes, A. L. (2009). Crisis in a networked world: features of computer-mediated communication in the April 16, 2007, Virginia tech event. *Social Science Computer Review, 27*, 467–480.

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and trends in information retrieval, 2*(1–2), 1–135.

Powley, E. H. (2009). Reclaiming resilience and safety: Resilience activation in the critical period of crisis. *Human Relations, 62*, 1289–1326.

Priedhorsky, R., Culotta, A., & Del Valle, S. Y. (2014). *Inferring the origin locations of tweets with quantitative confidence.* Paper presented at the Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing, Baltimore, MD.

Raftery, T. (2013). Utilities should open up API's to their smart meter data. http://greenmonk.net/2013/10/10/utilities-should-open-up-apis-to-their-smart-meter-data/

Rivera, F. I., & Settembrino, M. R. (2013). Sociological insights on the role of social capital in disaster resilience. *Disaster resiliency: Interdisciplinary perspectives*, 48–60.

Simpson, D. M., Lasley, C. B., Rockaway, T. D., & Weigel, T. A. (2010). Understanding critical infrastructure failure: examining the experience of Biloxi and Gulfport, Mississippi after Hurricane Katrina. *International Journal of Critical Infrastructures, 6*, 246–276.

Soden, R., Palen, L., Chase, C., Deniz, D., Arneson, E., Sprain, L., … Dashti, S. (2015). The polyvocality of resilience: Discovering a research agenda through interdisciplinary investigation & community engagement. *Under Review*.

Somers, S. (2009). Measuring resilience potential: An adaptive strategy for organizational crisis planning. *Journal of Contingencies and Crisis Management, 17*, 12–23.

Starbird, K., & Palen, L. (2010). *Pass It On?: Retweeting in Mass Emergency*. In Proceedings of the 7th International ISCRAM Conference–Seattle (Vol. 1).

Swanstrom, T. (2008). *Regional resilience: A critical examination of the ecological framework*. Baltimore, MD: Institute of Urban & Regional Development.

Tapia, A. H., Bajpai, K., Jansen, B. J., Yen, J., & Giles, L. (2011). *Seeking the trustworthy tweet: Can microblogged data fit the information needs of disaster response and humanitarian relief organizations.* Paper presented at the Proceedings of the 8th International ISCRAM Conference, Lisbon.

Tapia, A. H., LaLone, N., & Kim, H.-W. (2014). Run amok: Group crowd participation in identifying the bomb and bomber from the Boston marathon bombing. *Proceedings of the 11th ISCRAM, University Park, PA*.

Tapia, A. H., LaLone, N., MacDonald, E., Priedhorsky, R., & Hall, M. (2014). Crowdsourcing rare events: Using curiosity to draw participants into science and early warning systems. *Proceedings of the 11th ISCRAM, University Park, PA*.

Tapia, A. H., Moore, K. A., & Johnson, N. J. (2013). *Beyond the trustworthy tweet: A deeper understanding of microblogged data use by disaster response and humanitarian relief organizations.* Paper presented at the Proceedings of the 10th International ISCRAM Conference, Baden-Baden.

Tasse, D., & Hong, J. I. (2014). *Using social media data to understand cities*. Proceedings of NSF Workshop on Big Data and Urban Informatics, Chicago, IL.

Taylor, L. (2015). *Environment and Planning D: Society and Space,* 0263775815608851.

Truong, B., Caragea, C., Squicciarini, A., & Tapia, A. H. (2014). Identifying valuable information from twitter during natural disasters. *Proceedings of the American Society for Information Science and Technology, 51*(1), 1–4.

Vale, L. J. (2014). The politics of resilient cities: whose resilience and whose city? *Building Research & Information, 42*, 191–201.

Vieweg, S. (2010). Microblogged contributions to the emergency arena: Discovery, interpretation and implications. *Computer Supported Collaborative Work,* 515–516.

Vieweg, S. E. (2012). Situational awareness in mass emergency: A behavioral and linguistic analysis of microblogged communications. ProQuest LLC.

Winderl, T. (2015). Disaster resilience measurements: stocktaking of ongoing efforts in developing systems for measuring resilience. *UNDP*, 25.

Zobel, C. W. (2010). *Comparative visualization of predicted disaster resilience.* Paper presented at the Proceedings of the 7th International ISCRAM Conference, Seattle, WA.

Zobel, C. W. (2013). *Analytically comparing disaster recovery following the 2012 derecho.* Paper presented at the Proceedings of the 10th International ISCRAM Conference, Baden-Baden.

Zobel, C. W. (2014). Quantitatively representing nonlinear disaster recovery. *Decision Sciences, 45*, 1053–1082.

Zobel, C. W., & Khansa, L. (2012). Quantifying cyberinfrastructure resilience against multi-event attacks. *Decision Sciences, 43*, 687–710.

Zobel, C. W., & Khansa, L. (2014). Characterizing multi-event disaster resilience. *Computers & Operations Research, 42*, 83–94.