

Contents lists available at [ScienceDirect](#)

Social Science Research

journal homepage: www.elsevier.com/locate/ssresearch

Spatio-temporal filtering techniques for the detection of disaster-related communication

Sean M. Fitzhugh^{a,*}, C. Ben Gibson^a, Emma S. Spiro^b, Carter T. Butts^{a,c,d}^a Department of Sociology, University of California, 3151 Social Science Plaza A, Irvine, CA 92697, USA^b Information School, Mary Gates Hall, 370, University of Washington, Seattle, WA 98195, USA^c Institute for Mathematical Behavioral Sciences, Department of Statistics, Bren Hall 2019, University of California, Irvine, CA 92697, USA^d Department of Electrical Engineering and Computer Sciences, 2200 Engineering Hall, University of California, Irvine, CA 92697, USA

ARTICLE INFO

Article history:

Received 30 July 2015

Received in revised form 15 April 2016

Accepted 26 April 2016

Available online 20 May 2016

Keywords:

Disasters

Rumoring

Communication

Geography

Event detection

Big data

ABSTRACT

Individuals predominantly exchange information with one another through informal, interpersonal channels. During disasters and other disrupted settings, information spread through informal channels regularly outpaces official information provided by public officials and the press. Social scientists have long examined this kind of informal communication in the rumoring literature, but studying rumoring in disrupted settings has posed numerous methodological challenges. Measuring features of informal communication—timing, content, location—with any degree of precision has historically been extremely challenging in small studies and infeasible at large scales. We address this challenge by using online, informal communication from a popular microblogging website and for which we have precise spatial and temporal metadata. While the online environment provides a new means for observing rumoring, the abundance of data poses challenges for parsing hazard-related rumoring from countless other topics in numerous streams of communication. Rumoring about disaster events is typically temporally and spatially constrained to places where that event is salient. Accordingly, we use spatio-temporal subsampling to increase the resolution of our detection techniques. By filtering out data from known sources of error (per rumor theories), we greatly enhance the signal of disaster-related rumoring activity. We use these spatio-temporal filtering techniques to detect rumoring during a variety of disaster events, from high-casualty events in major population centers to minimally destructive events in remote areas. We consistently find three phases of response: anticipatory excitation where warnings and alerts are issued ahead of an event, primary excitation in and around the impacted area, and secondary excitation which frequently brings a convergence of attention from distant locales onto locations impacted by the event. Our results demonstrate the promise of spatio-temporal filtering techniques for “tuning” measurement of hazard-related rumoring to enable observation of rumoring at scales that have long been infeasible.

© 2016 Elsevier Inc. All rights reserved.

The increasing prevalence and accessibility of online data heralds a new regime for data collection and analysis. We enjoy greater access to large-scale, observational social science data, much of which has a textual component. This new data regime

* Corresponding author.

E-mail addresses: sean.fitzhugh@uci.edu (S.M. Fitzhugh), cbgibson@uci.edu (C. Ben Gibson), espiro@uw.edu (E.S. Spiro), butts@uci.edu (C.T. Butts).

brings great opportunities but also poses several challenges (Boyd and Crawford, 2012; King, 2011). While many new data sources offer opportunities to analyze and understand social phenomena with precision and at scales that have long been infeasible, the scope of such data can be overwhelming. These kinds of large-scale data are often too large to examine manually and are frequently noisy due to a conglomeration of different competing signals of social processes. This challenge creates an opportunity for the development of filtering techniques to refine signal from noise in order to highlight a particular activity or set of activities.

Scholars from a variety of fields have recently turned attention towards identifying signals of social processes in large-scale, user-generated online text data from platforms such as Twitter and Facebook. At any given moment individuals use online platforms to discuss a wide variety of topics such as news, sports, work/school, weather, and what they ate for lunch. Each of these topics has its own “signal” of user activity (reflecting the text and information generated by users) and the site’s collective activity is awash with signals from a nearly endless array of topics. Identifying particular signals of user behavior, such as discussing specific topics or frequenting specific locations, is an ongoing challenge that is complicated by the bevy of signals of other activity. The challenge of parsing large volumes of online text to identify signal has drawn much attention in recent years (Ferrari et al., 2011; Hollenstein and Purves, 2013; Mamei et al., 2010; Pozdnoukhov and Kaiser, 2011). While the application in the online environment is novel, the general problem is not. Interest in automated signal processing in noisy environments (Fawcett and Provost, 1999; Hamid et al., 2005; Macleod and Congalton, 1998; Ribeiro et al., 2012; Singh, 1989; Stauffer and Grimson, 2000) predates the proliferation of user-generated online activity, and we can apply the lessons learned in those contexts to the online context.

Although the signal identification problem is not new, detecting signal in the online environment poses new challenges that require novel solutions. Short message length and esoteric language and grammar in media like text messages and Twitter enhance the difficulty of identifying signals of social processes (Davis et al., 2011; Go et al., 2009; Kireyev et al., 2009; Kouloumpis et al., 2011; Pak and Paroubek, 2010; Yang et al., 2014). The challenges of detecting signals of social phenomena in the online environment implore us to develop a fundamental understanding of the social phenomena we intend to detect. Failure to understand the social processes underlying activity observed at large scale is dangerous and may lead to misleading or spurious results (Back et al., 2011; Boyd and Crawford, 2012; Johnson, 2014; Lazer et al., 2014; Leinweber, 2007), such as misclassifying failure-to-connect error messages as “anger” messages in a stream of pager messages or overestimating flu incidence based on search term activity on Google. To prevent such outcomes we harness what is known about the social phenomenon—in this case, rumoring behavior in the context of disaster events—in order to filter activity selectively as we search for signal of hazard-related rumoring.

This principle of systematic filtering motivates the techniques we develop in this paper. Here our primary goal is to develop an approach for measuring signal of human activity that varies markedly across time and space. Understanding the basic principles of a phenomenon of interest is key to developing effective strategies for measuring that phenomenon. We introduce a spatio-temporal filtering approach for measuring signal and demonstrate its utility by illustrating an example of detecting signal of disaster-related rumoring activity. Using what is known about social responses to disaster events (impending or realized), we selectively filter timestamped streams of geolocated, informal communication activity by time and location in order to identify surges of rumoring activity in response to a disaster. Spatio-temporal filtering enhances our ability to detect events by utilizing the signal produced by sources that are known (or expected) to produce reliable information, thereby enhancing our ability to detect distinct activity patterns above and beyond typical global signal (i.e. background noise representing the array of signals irrelevant to our focus). We proceed from here by characterizing rumoring in disaster contexts, reviewing techniques for event detection, and illustrating strategies for how we can filter our data across time and space to identify local surges of disaster-related rumoring activity. We then describe various approaches one can take for filtering data and illustrate our technique with online, informal communication in the context of several different types of disaster events.

1. Informal response to disasters

The disruption of social structures and routines is a characteristic feature of disasters. Kreps (1984) defines a disaster as an event inducing damages, losses, and/or disruption, where the impact affects social structures and/or societies. Key to this definition is that a disaster compels a response from social units, often in response to the disruption of essential functions of society (Fritz, 1961). This response translates into action from individuals, households, organizations, and other social units (Drabek and McEntire, 2003; Drabek et al., 1981; Hughes and Palen, 2009; Leik et al., 1981; Sutton, 2010), often in the form of emergent and frequently novel behavior. While most disasters have some element of emergency response activity by formal organizations (e.g., search and rescue, firefighting, aid provision, triage), informal actors such as individuals and families engage in a variety of activities. They take protective action, verify official information (e.g., emergency alerts), supplement official information with information from friends, family, and neighbors, check the status of others, and engage in mitigation activities such as search and rescue, evacuation, and cleanup (Drabek et al., 1981; Hughes and Palen, 2009; Leik et al., 1981; Quarantelli, 1954, 1980; Sutton et al., 2013a,b). Citizen response routinely serves as the first emergency response during disasters, followed shortly thereafter by official emergency responders (Kreps, 1983). Informal response to emergency situations is a critical aspect of the hazard environment.

Informal communication is an essential component of disaster response. Informal communication includes, for example, face-to-face conversations between neighbors, phone conversations between friends, and emails exchanged among students.

By contrast, formal communication typically occurs through channels such as press conferences, official mailings, press releases, and emergency alerts. During disaster events informal channels are heavily utilized as citizens regularly provide situational reports from the disaster zone—information on road closures, online maps of affected areas, reports of damaged/destroyed structures—before such information is released by official sources (Goodchild, 2007; Goodchild and Glennon, 2010; Sutton et al., 2008). Informal communication in this context falls under formal definitions of *rumoring*: informal, person-to-person communication pertaining to current, newsworthy topics of interest among a population (Allport and Postman, 1947; Bordia and DiFonzo, 2004; Caplow, 1947; Rosnow and Kimmel, 2000; Shibusani, 1966). A key component of this definition of rumor is that such statements are not confirmed by official sources. Although rumoring has historically been a notoriously difficult phenomenon to measure, the field has a developed basic understanding of the general principles of rumoring behaviors in disaster contexts. These common behavioral patterns allow us to refine our approach for measuring signal of disaster-related rumoring activity.

1.1. Constraints on disaster-related rumoring

Rumoring is known to have several properties that constrain it to a particular point in space and time. Rumoring is not a ubiquitous response to disrupted environments. Unlike urban legends and folk tales, it does not propagate endlessly (Miller, 1992). Some have described rumoring as a task of collective problem solving (Bordia and DiFonzo, 2004; Shibusani, 1966) with reasonably well defined start and end points. A process helping individuals make sense of their environments and cope with uncertainty, rumoring tends to dissipate once cognitive unclarity has been eliminated (Caplow, 1947). Salience also plays an important role in determining the lifespan of a rumor. As an event/entity becomes more and more likely to affect individuals, those individuals become more and more likely to rumor about it (Allport and Postman, 1947). This process also works in reverse; populations further removed in time from an event become less likely to rumor about it. In addition to temporal proximity, one's spatial proximity to the subject of rumors plays a role one's propensity to engage in rumoring. Rumors pertaining to a particular topic typically survive only as long as that topic is relevant and interesting to the rumoring population (Allport and Postman, 1946), but not all topics are salient to all populations. Although capable of traveling great distances in short periods of time, rumors fail to gain traction and quickly perish if they are not relevant to a population (Caplow, 1947). These mechanisms driving rumoring activity place limits on its spatio-temporal prevalence.

Characteristics of disasters further constrain the spread of rumoring activity across time and space. Rumors thrive in disrupted settings, and disasters have long provided contexts for studying rumoring. In disruptive and high-anxiety situations, rumor viscosity decreases and rumors spread at relatively high rates (Caputo, 1999; Stein, 1980; Walker and Beckerle, 1987). Periods of disruption or anxiety are finite in time and space, however, and rumoring activity typically persists as long as these states are active. Although the spatio-temporal variation in rumor quantity and content has long been of interest to the field, collecting data that accounts for temporal *and* spatial characteristics of rumoring has been extraordinarily difficult to do with any degree of precision. Some have been able to capture rumoring data with some degree of temporal precision (Bordia and Rosnow, 1998; Danzig, 1958; Greenberg, 1964) or with some spatial precision (Larsen, 1954), but bridging the two has been difficult. Synthesizing temporal and spatial rumoring data across a wide variety of events had long been beyond the capabilities of researchers. Simply gathering reliable data on rumoring was already fraught with challenges.

Measuring rumoring has historically been challenging in hazard contexts. The field has traditionally relied on *post hoc* informant reports of rumoring, which are subject to a variety of informant errors (Back et al., 1950; Romney et al., 1986; Romney and Weller, 1984; Sudman et al., 1996). In the aftermath of a highly publicized kidnapping of the British Trade Commissioner, Scanlon (2007) finds that people frequently forgot from whom they heard rumors and periodically mis-attributed the source of the rumor. Sampling issues often emerge as disasters frequently displace portions of the population that evacuate to safety, seek medical attention, or assist with emergency response. Furthermore, those who have left the area may have been more likely to have heard the rumor and acted on that information, thus leaving the remaining population with a lower likelihood of having heard or spread the rumor. To maintain a manageable scope of analysis, many studies outside the laboratory ask respondents where they heard a *particular* rumor (Erickson et al., 1978; Greenberg, 1964; Miller, 1992; Richardson et al., 1979; Scanlon, 1977; Walker and Beckerle, 1987). This is in contrast to the classic Allport and Postman (1947) and Caplow (1947) analyses on *rumoring*, which capture multiple threads of rumor. With a few exceptions (Kapferer, 1989; Schachter and Burdick, 1955) studies of rumor typically operate retrospectively and suffer from a success bias. As a result, we know less about situations under which rumors *fail* to emerge or diffuse. Merely measuring rumoring has historically been quite challenging; measuring rumoring with accurate spatial and temporal precision has long been infeasible with human informants. Until the recent rise of computer-mediated communication mediated communication, advances in the measurement of rumoring had stagnated for decades (Bordia and DiFonzo, 2004; Bordia and Rosnow, 1998). This paper builds on recent advances by harnessing the opportunities of the new big data regime in order to overcome many of these historic obstacles to measuring rumoring. Using what is known about when and where rumoring flourishes, we develop techniques to filter data such that we maximize the signal of particular, hazard-related rumoring activity.

1.2. Stages of hazard-related rumoring

Just as rumoring is bound to specific slices of space and time, so too is informal response to disaster events. We define three distinct phases of rumoring activity in response to disaster events: anticipatory excitation, primary excitation, and secondary

excitation. Anticipatory excitation encompasses rumoring that precedes events forecast by warnings and alerts. Many, but certainly not all, disasters are preceded by a series of warnings or alerts. Organizations such as the National Weather Service generate these messages and issue them to be distributed by local offices in areas that may be affected by the impending hazard event. Salience effects play an important role in determining the level of informal communicative activity in response to these warnings. Allport and Postman (1947) use the goal-gradient phenomenon to link salience and rumoring behaviors. As an event becomes increasingly likely to affect a population, the “more fertile is the soil for anticipatory rumors” (Allport and Postman, 1947, p. 63). In anticipation of the disaster event, individuals to whom the warning is salient go through processes of confirming the warning, developing a “warning belief” to assess whether the event poses a threat, determining how threatening the risk is, and concluding that taking protective action is or is not necessary (Perry et al., 1981). Part of this process involves milling behaviors, in which individuals collectively determine whether to pursue action in response to these warning messages (Drabek et al., 1981; Fritz, 1961; Killian and Turner, 1972; Quarantelli and Dynes, 1977). Constrained by typical bounds on rumoring activity, this pre-event rumoring is not pervasive, however. Warnings and alerts are issued to specific locations and persist for a finite period of time (ranging from an hour to a few days). Accordingly, rumoring activity in response to warnings ought to occur when and where warnings are active.

During primary excitation the population situates itself in its new environment and begins to respond to the impact of the hazard event. During and immediately following the event, people engage in informal discussion of events and the passing of improvised news (Shibutani, 1966). Such news frequently is more timely than news from official sources and therefore becomes the dominant mode of information transmission (Caplow, 1947; Kreps, 1984). As disaster-related rumoring primarily occurs where the disaster is salient and where anxiety thrives (Caputo, 1999; Stein, 1980; Walker and Beckerle, 1987), this sharing of news primarily occurs among those impacted (directly or indirectly) by the disaster. That is, distant populations wholly unaffected by an event will have a low propensity to rumor about it. It follows then that while the disaster is ongoing and in its immediate aftermath we ought to observe rumoring activity in and around locations directly impacted by the event.

Secondary excitation occurs in the aftermath of the disaster as we see additional communicative activities in response to the event. In the areas directly impacted by the event we observe a continuation of response and recovery efforts. Activity also occurs in areas outside the disaster zone. The event aftermath is frequently characterized by a series of distinct responses across different distances to the epicenter (Dynes, 1970). Those near the disaster site frequently discuss firsthand accounts of the event and spread information of these firsthand accounts to others. Populations further away may hear about the impact of the event from news media. Some severe events may be followed by a mass convergence phenomenon, in which individuals share information about the event, express concern about those impacted, and/or attempt to increase awareness about how to help those affected (Sutton, 2010). This mass convergence increasingly occurs in an online context and may emerge from populations in a variety of locations, from local individuals to those situated hundreds or thousands of miles away (Hughes and Palen, 2009; Sutton, 2010). Eventually the response will taper off as the event becomes less and less salient. Additionally, a return to normal routines accompanies an elimination of the ambiguity under which rumoring thrives (Caputo, 1999; Schachter and Burdick, 1955; Stein, 1980). Unlike anticipatory and primary excitation in response to the event, post-event secondary excitation may take place in locations distant to the affected area in addition to the event epicenter.

In order to validate that we are measuring signals of hazard-related rumoring, we ought to detect anticipatory excitation in events preceded by warnings and alerts, primary excitation, and secondary excitation (including mass convergence in some events). Per theories of rumor and crisis informatics, these responses occur in reasonably well defined points of space and time. If we can capture these signals across large spatial scales, this would validate our approach and represent a great step forward for the rumoring subfield, which has long been limited by the difficulty of measuring the phenomenon.

2. Measuring hazard-related, online, informal communication

The online environment offers an opportunity to overcome many of the historic challenges of studying informal communication in disaster contexts. As populations increasingly turn to platforms such as Twitter, OpenStreetMap, Facebook, and Wikimapia to send and receive information during disaster settings (Goodchild, 2007; Goodchild and Glennon, 2010; Sutton et al., 2008), scholars now have opportunities to monitor online, informal communication at global spatial scales and collect metadata on the timing, location, and content of such communication. While the online environment affords new opportunities for observing hazard-related informal communication, the scale of the data can be overwhelming. Taking advantage of this new data regime without succumbing to its heft requires a careful approach. By leveraging what we know about hazard-related informal communication, we can selectively filter communication data to capture and observe the rumoring phenomenon.

To obtain streams of online, informal communication related to disaster events, we turn to Twitter. Twitter is an online microblogging service where users post short messages (called *tweets*) up to 140 characters in length. Like a traditional RSS system, users subscribe to (or “follow”) accounts to receive all tweets posted by those accounts. These messages are time-stamped and some users have opted to include geographic metadata with each message. The limited size of these messages creates an environment conducive to traditional rumoring behaviors. Past studies find that rumor statements tend to be brief, structurally simple, and easily retransmitted (Allport and Postman, 1947; Baron et al., 1997; Caplow, 1947). Further aligning its content with that of traditional rumoring, Twitter’s activity is largely topical and the company has institutionalized features such as robust search capabilities and topic channels called “hashtags” to promote and take advantage of the topicality of its

content. With over 300 million active, registered users posting 500 million tweets per day,¹ Twitter is an extremely popular service and the 8th-most visited website in the world.² As Twitter has become accessed increasingly via applications on GPS-enabled mobile phones and tablets, more and more users opt to include geographic metadata (VGI) with each tweet. Twitter reports that 83% of its active users access the site via mobile devices.¹ When users utilize Twitter's VGI, each message they post contains data about the location from which they send each message. Such data include the city, state, and country from which the message was sent, along with the latitude and longitude coordinates. This gives us an opportunity to observe when and where people discuss particular topics.

Twitter is an increasingly useful platform for observing event-specific communication during disasters. Case study work has highlighted the use of social media for information sharing and collective disaster-related sensemaking (Vieweg et al., 2008). Mimicking behaviors long observed in the rumoring literature, individuals in the online environment employ informal communication to cope with information shortages and relieve feelings of helplessness (Sutton et al., 2008), share eyewitness observations of events in real time (Vieweg et al., 2010), and engage in information diffusion by retransmitting others' messages (Sutton et al., 2013a,b). These studies on Twitter usage during disasters indicate that online populations respond in systematic, consistent, measurable ways to disaster events. Such responses typically include increased volume of event-specific keywords, changes in message structure, and changes in information retransmission behavior. This regular behavior and adherence to definitions of rumor endorse the Twitter platform as a tool for understanding rumoring in the context of hazards.

The data in this paper come from Project HEROIC (Hazards and Emergency Response in Online Informal Communication), a multi-year project dedicated to the collection and analysis of online, informal communication in response to hazard events. This paper uses keyword-driven data collected from Twitter's Streaming API. Under Project HEROIC we continuously collect data from 72 unique, hazard-related keywords such as "wildfire," "earthquake," "collapsed," "flood," and "shooting" which cover 19 distinct natural and anthropogenic hazard types. This keyword-driven data collection continuously gathers all public tweets returned by Twitter containing the specified keywords. For this paper we use data from the most recent version (1.1) of the Twitter Streaming API, which allows us to examine hazard-related messages from May 2013 onward. For each of these terms we continuously observe its stream of activity over time. This long-term, prospective approach allows us to observe when and where rumoring activity occurs on a given keyword, as well as how long it lasts. Likewise, we also observe the *absence* of activity, which sheds light into times and places where rumoring *does not* occur. With a prospective, cross-hazard measure of rumoring activity across a variety of topics, we are well equipped to observe and analyze rumoring activity at scale. The timing, location, and text data allow us to observe the phenomenon with historically unmatched precision and across large spatial scales. However, the opportunity afforded by Twitter's abundance of data is balanced by the challenge of finding signal amid an overwhelming mass of information. In the following section we develop an approach to refine our search such that we maximize signal of hazard-related rumoring.

3. Spatially enhanced signal detection

The spatial inhomogeneity of activity in response to disaster implores us to use spatial subsampling to detect variable surges in activity across space. By filtering activity into distinct blocks based on distance from the epicenter of a disaster event, we are able to harness (what ought to be, per rumor theories) the strongest signal of communicative activity in response to the event. Geographic metadata in our activity streams enables the use of these spatial filtering techniques. Thanks to the rise of geolocated communication such as mobile phone calls and messages triangulated by cell towers and volunteered geographic information (VGI) in social media, geolocating informal communication is easier than it has ever been.

While geographic information has historically followed a set of formal standards and been produced and verified by a handful of experts, no such standardization or verification exists to ensure that individuals do not produce false geographic information online. To combat possible noise in our signal, we use aggregation-based approaches to identify signal of disaster-related rumoring. This approach follows Tobler's First Law of Geography and a variety of techniques from the field of informant accuracy (Goodchild and Li, 2012; Latonero and Shklovski, 2010). Tobler's First Law of Geography states that "All things are related, but nearby things are more related than distant things" (Tobler, 1970). Barring sudden, drastic changes to the landscape, new information from a specific locale should be consistent with what we already know about that location as well as what we know about nearby locations. Similarly, theories of informant accuracy posit that those with accurate domain knowledge provide more reliable responses with less error than those without such knowledge (Romney et al., 1986; Romney and Weller, 1984; Sudman et al., 1996; Weller and Romney, 1988). Their observations will cluster around a single "truth" while inaccurate observations (i.e. error) will be randomly scattered around the truth; that is, error is inhomogeneous and does not typically converge around a small number of data points. Both approaches suggest that if we notice a widespread surge of "blizzard" messages during the winter in Denver, for example, we ought to be safe to assume that we are observing a response to a blizzard-related event. If we identify a small number "brush fire" messages sprinkled throughout the information stream during the same event, we can usually dismiss such messages as noise. During a disaster event any noise or error ought to be

¹ <https://about.twitter.com/company>.

² <http://www.alexa.com/topsites>.

outweighed by the signal produced by accurate observers. Should no such disaster be present, the indiscriminate noise should ensure that we see no such signal in any direction.

This aggregation-based approach for identifying signal of social phenomena has been successfully employed in numerous online contexts. We are more likely to observe a disaster event where we find a series of similar, spatially clustered reports than we are if we only observe a single report (Latonero and Shklovski, 2010). Additionally, this aggregation-based approach has been used to infer characteristics of places from VGI: activity logs have been used to characterize activity “hot spots” in a variety of urban environments (Ferrari et al., 2011; Pozdnoukhov and Kaiser, 2011); geotagged photographs point to bastions of activity and popular sites among tourists (Hollenstein and Purves, 2013; Leung and Newsam, 2010; Mamei et al., 2010); and changes in spatio-temporal activity patterns throughout the day enable the classification of different parts of a city as either residential, industry-driven (“office towns”), nightlife-driven, or a combination of activities (Wakamiya et al., 2011). In multiple contexts, consistent information from multiple sources across geographic space implies a reliable and accurate signal. Space is a powerful tool both for verifying the validity of our data and as a dimension along which we can identify variation in that data. In order to identify variation, however, we need to characterize what normal, baseline activity looks like.

3.1. Establishing baseline signal

While aggregation of geolocated, disaster-related messages provides a signal of event-driven rumoring activity, we need to develop a baseline signal to contrast against our observed signal. We know that anomalous communication activity occurs in response to disaster, but the challenge here is distinguishing that response from ordinary activity. Filtering signal from noise is key to identifying surges of activity, whether that activity is social, man-made (e.g. traffic), or natural (e.g. wind events). The “activity monitoring” class of problems involves observation of streams of activity, followed by an alert if an aberration has occurred in the stream (Fawcett and Provost, 1999). One of the primary challenges is to balance the sensitivity of the alert mechanism (Stauffer and Grimson, 2000). Too many false positives ensure that the alarm loses its effectiveness while excessive false negatives may lead to a failure to correct or respond to a disruptive event, such as a the closure of an emergency exit (Andrade et al., 2006), a traffic jam (Stauffer and Grimson, 2000), or a disruption of ordinary routines in a loading dock (Hamid et al., 2005). In this case, we treat our keyword-driven samples of online informal communication as distinct streams of activity, whose disruption we identify using a measure based on aggregate counts.

To ensure proper detection of communication surges in response to disasters, we need to strike a balance between mistakenly classifying spurious changes of local, hazard-related communication activity and ensuring that hazard-related communication activity in response to disaster events is not overlooked. To detect aberrations in communication we use a *profiling method* (Fawcett and Provost, 1999). The profiling method approach establishes a baseline level of normal activity and identifies deviations from typical activity patterns or sequences; these deviations are then classified as atypical events (Fawcett and Provost, 1999; Hamid et al., 2005; Stauffer and Grimson, 2000). We follow this same approach by monitoring keyword streams over an extended period of time and establishing a baseline level of activity and measures of typical levels of variation for each keyword. Each keyword-driven activity stream ought to have a different spatio-temporal activity profile: when and where people use a technical term such as “aftershock” is very different from usage of a more general term such as “fire.” Accordingly, each keyword has a distinct pattern of variation in volume across regular intervals (hourly, daily, monthly). We illustrate this in Table 1, where we list the average daily posting rate and standard deviation for that rate for a variety of keywords within 100 miles of three distinct areas: the small town of Carrington, North Dakota (population 2065 according to the 2010 U.S. Census); the large city of Denver, Colorado (population 600,158); and the extremely large population center of New York, New York (population 8,491,079). We also illustrate the global number of geolocated messages for each term on a daily basis.

Across each of these three areas posting rates change notably, with New York frequently having an order of magnitude more daily hazard tweets than Denver, whose posting rates are roughly two orders of magnitude higher than those around Carrington. To account for this spatial variation in message activity, we take steps to distinguish between an observation of

Table 1
Daily posting rates for hazard-related keywords.

	Carrington, ND	Denver, CO	New York, NY	Global
Aftershock	0.000 (SD = 0.000)	0.037 (SD = 0.189)	0.211 (SD = 0.802)	12.437 (SD = 14.815)
Alert	0.178 (SD = 0.446)	27.819 (SD = 21.371)	171.948 (SD = 108.915)	2228.119 (SD = 1201.739)
Collapsed	0.000 (SD = 0.000)	0.181 (SD = 0.488)	2.811 (SD = 6.681)	41.044 (SD = 33.833)
Earthquake	0.004 (SD = 0.061)	0.556 (SD = 1.880)	5.393 (SD = 13.227)	934.219 (SD = 638.979)
Fire	0.439 (SD = 1.044)	23.922 (SD = 21.821)	215.442 (SD = 180.595)	3930.097 (SD = 3422.849)
Flood	0.007 (SD = 0.086)	8.126 (SD = 32.929)	13.159 (SD = 17.083)	385.252 (SD = 464.900)
Rain	0.409 (SD = 0.884)	21.475 (SD = 31.239)	265.404 (SD = 378.567)	10720.110 (SD = 5529.984)
Shooting	0.063 (SD = 0.243)	4.830 (SD = 5.812)	38.337 (SD = 40.214)	651.844 (SD = 652.418)
Tornado	0.061 (SD = 0.434)	3.343 (SD = 14.248)	10.600 (SD = 40.636)	272.257 (SD = 609.788)
Warning	0.654 (SD = 2.653)	9.621 (SD = 21.176)	25.561 (SD = 18.832)	812.208 (SD = 707.917)
Wildfire	0.000 (SD = 0.000)	0.980 (SD = 2.242)	0.693 (SD = 1.395)	22.207 (SD = 33.424)

100 more “fire” keywords than normal during a day in Carrington, North Dakota and an equivalent increase in New York City. Given the small population of Carrington, such an increase is likely indicative of a fire event, while the latter is well within the expected daily variation in New York City’s population. Additionally, posting rates vary across keywords. In a city of Denver’s size, an increase of 20 messages more than average containing the term “flood” is distinct from an increase in 20 messages with the keyword “wildfire.” The rarity of the latter term suggests that such an increase may be due to a wildfire event. The former term’s increase is likely part of typical variation in activity, however. To account for this variation across keyword streams and across space, we use individual profiling (Fawcett and Provost, 1999) for each keyword. By characterizing the typical amount of activity across space and independently across keywords, we create multiple information streams to monitor and thereby improve our ability to detect responses to different kinds of events in different locales. We harness our ability to detect change in aggregate activity for each keyword to identify signals of hazard-related rumoring.

In addition to detecting when and where aberrations in activity occur, we need to identify the direction and magnitude of activity changes. Determining the *extent* of the change in activity is a key component of anomaly detection (Lu et al., 2004; Macleod and Congalton, 1998). Emulating the concept of a *spectral change vector* (Singh, 1989), we describe the direction and magnitude of change in activity relative to normal, baseline activity levels. With our timestamped (at the second) series of hazard-keyword communications, we use differences in the volume of keyword usage to determine the magnitude and direction of change over time.

4. Anomaly detection

To illustrate how we identify rumoring activity during disaster events, we use the case of the Moore tornado. On Monday, May 20th, 2013 an EF5 Tornado struck Moore, a city of 55,000 in the Oklahoma City metro area. The Oklahoma Department of Emergency Management reports that the tornado killed 25, injured 377, destroyed approximately 1150 homes, and caused an estimated \$2 billion in damages. A hallmark disaster, this was an event that severely disrupted social routines and spurred a surge of activity in anticipation of and response to the tornado. As with many such incidents, activity on Twitter reflected this disruption. This activity was not uniformly distributed across space, however. To observe this signal at a higher resolution, we refine the signal of rumoring activity across space.

One of the most elementary approaches we can take to filtering our data is by aggregating counts of hazard-specific messages across space. We begin by creating a set of bins, represented by concentric circles around the epicenter of the Moore tornado. We take the set of all geolocated “tornado” tweets on May 20th and put tweets within 100 miles of the epicenter in one bin, tweets between 100 and 200 miles from the epicenter in another bin, tweets between 200 and 300 miles in another bin, and so on. We could employ a wide variety of approaches for measuring signal of rumoring activity in each bin. A raw count of messages reflects the total volume of hazard-related rumoring activity. In the left panel of Fig. 1, we illustrate raw counts of “tornado” messages as a function of distance from the epicenter of the Moore tornado. We observe the largest raw count in the 0–100 mile bin, followed by the second-largest count in the 100–200 mile bin, and the third-largest in the 700–800 mile bin. While the counts of messages in the 100–200 and 700–800 mile bins are nearly identical, the latter bin includes the Atlanta metro area and much of the Chicago metro area, two of the top-ten largest metro areas in the country. As illustrated in Table 1, the counts of messages vary by location based on the populations in those locations. Accordingly, raw counts of “tornado” messages will be influenced by the population in each bin. In many contexts one would want to account for spatial heterogeneity in population by establishing a baseline measure of activity in each bin.

We could utilize a number of normalization techniques to isolate signal of event-driven rumoring activity from normal activity. One approach is to calculate a relative change of signal by measuring the ratio of the observed count of messages to the average daily count of messages in each bin. These bin-specific ratios indicate the magnitude of change compared to typical counts of “tornado” messages. While the ratio-based approach allows us to establish a baseline that accounts for typical quantities of “tornado” messages, this approach offers no insight into whether the observed departure is large or small compared to typical *variation* in posting rates. Beyond measuring how far our observed signal deviates from baseline signal, in some circumstances we may want to know how large this deviation is compared to typical deviations. To measure change in message count relative to typical variations in message counts, one may employ a measure based on a z-score. By reporting how many standard deviations above or below the mean our message count is, this measure tells us activity change net of typical variation. In the right panel of Fig. 1 we illustrate the z-score associated with each bin.

The contrast between the two panels in Fig. 1 demonstrates the different types of signals highlighted by different measurement techniques. As noted previously, the raw counts showed that the 700–800 mile bin had a large amount of tornado-related rumoring activity, perhaps due to the large population within that bin. Once we use the z-score approach to account for typical volume and volatility of “tornado” message activity in each bin, we find that the signal from the 700–800 mile bin is substantially lower than it appears when using raw counts (though it nonetheless remains elevated beyond typical signal). Under the z-score metric the three bins closest to the epicenter of the Moore tornado have the strongest signal of tornado-related rumoring activity. This signal drops off substantially as we get further from the event. These normalization procedures enhance our measurement of the signal on the day of the Moore tornado, but measuring signal across the span of multiple days requires further filtering of the data.

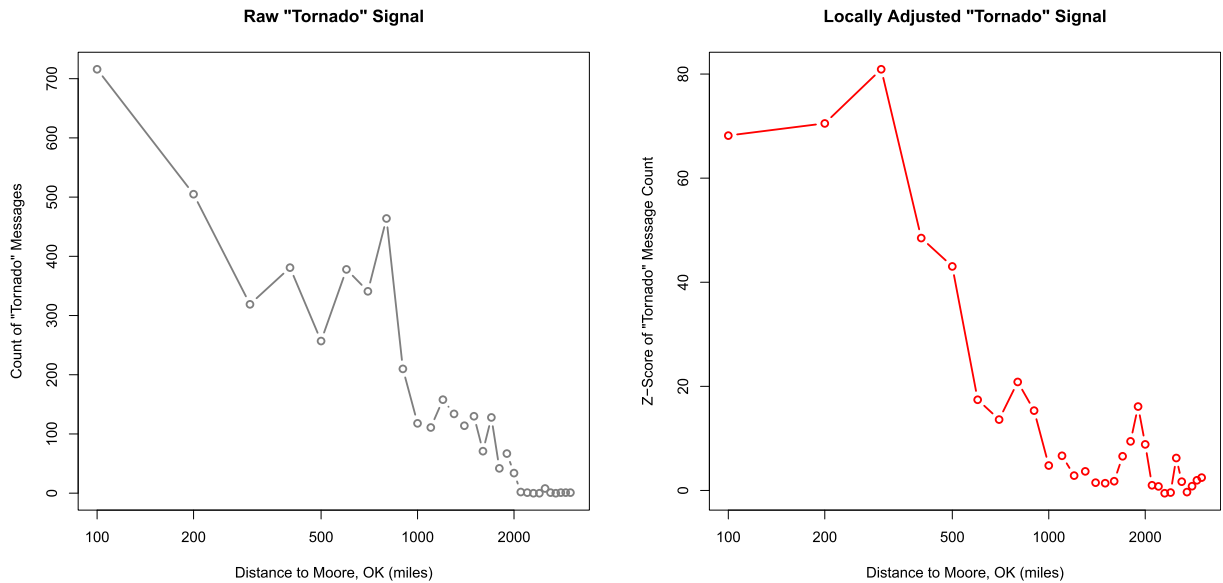


Fig. 1. Tornado-related message activity on May 20th, day of the Moore tornado. On the left we illustrate raw counts of messages in each 100-mile bin around Moore and on the right we depict the z-scores associated with each count. Both the raw counts and z-scores indicate strong signal of activity around Moore. Beyond 300 miles, the signal drops off notably. Although we observe large counts of messages at 600 miles (which includes Denver and New Orleans) and 800 miles (which includes Chicago and Atlanta) from the event, the z-score indicates that this is not as great a departure from normal "tornado" message counts as the departure we observe between 100 and 300 miles, due to higher average activity in the 600 and 800-mile bins.

4.1. Spatio-temporal filtering

Although we previously observed signal of rumoring activity on a single day, in many contexts we may want to measure signal over multiple time points (hours, days, weeks, etc.). To illustrate this approach, we will examine tornado-related rumoring activity over a six-day span surrounding the Moore tornado. When looking for signal across multiple days and multiple locations, we now must turn to *temporal* filtering in addition to our spatial filtering. Let us define activity A_{kbt} as the count of messages containing a keyword (or set of keywords) k in geographic bin b during time interval t . As discussed in the illustrative example above, we may use any number of keywords depending on the type of event: "tornado," "wildfire," "flood," "shooting," etc. Additionally, one may take an equally flexible approach to defining geographic bins. While we employ 100-mile-wide concentric circles to measure activity as a function of distance to the event epicenter, alternative approaches may use different sized concentric bins or examine signal in standardized geographic regions such as U.S. Census blocks, tracts, or counties. Instead of filtering tornado-related rumoring across space by measuring signal as a function of distance to the event, we may instead be interested in filtering rumoring according to a county's history of tornado events (such that we compare signal of activity in locations with many past alerts against those with few alerts). One may employ endless possible strategies for filtering signal, and how one filters signal from noise largely depends on the question(s) one wishes to answer with that signal. Likewise, such strategies also inform our approach to filtering across time. In this case we observe message volumes in 24-h intervals, which illustrates daily changes in activity. For cases with especially rich data (or rapidly changing signals), filtering data at an hourly resolution may be more appropriate, while filtering at a weekly or monthly resolution would be more appropriate for very sparse or slowly changing activity signals. As we expect our signal to vary as a function of spatial distance and temporal distance from the event (per rumor theories), we accordingly use a spatio-temporal filtering approach.

Because we are measuring tornado-related rumoring across multiple days, we must calculate a z-score for each day within each spatial bin. For each day we calculate the bin-specific z-score from the numerator of Equation (1). For message volume A_{kbt} during a given day in a particular geographic bin, we take the difference between the observed number of messages in that geographic bin (e.g. all "tornado" messages) and the daily average number of messages in that geographic bin, and divide the difference by the standard deviation of number of daily messages in that spatial bin. We calculate the average daily message volume and standard deviation for any keyword in any bin from a reference period spanning nine months, beginning in May of 2013. This approach follows the standard z-score calculation for how far the observed activity signal has deviated from typical activity signal. We call this the *raw spatially filtered signal*.

$$\frac{\left(\frac{A_{kbt} - \bar{A}_{kb}}{\sigma(A_{kb})} \right)}{\left(\frac{A_{kgt} - \bar{A}_{kg}}{\sigma(A_{kg})} \right)} \quad (1)$$

We previously discussed strategies for accounting for bin-level activity patterns, and we again have to consider these strategies at the global level if we're examining signal over multiple time points, as global signal varies over time. In some applications we may want to adjust for how large of an increase in *global* signal we observe during that time point. While the z-score in the spatio-temporally filtered bin indicates variation in signal in that spatial bin, such variation may not necessarily be exclusive to that bin. Instead global variation in signal (i.e. many locations worldwide have some shared change in signal) may account for the local variation we observe. If we find a ten-fold increase of tornado-related rumoring in Moore during a given time point and find that during that same time point worldwide the count of geolocated "tornado" messages is ten standard deviations above typical activity, then we would be remiss to conclude that our observed increase in signal is specific to Moore. Instead this is likely a reflection of the increased global signal. Likewise, if the global count of "tornado" messages is ten standard deviations below normal, then a typical count of tornado messages in any given bin may be noteworthy. To account for this we calculate the *global* z-score, which we call the *global signal*, in the denominator of Equation (1). We subtract the average global count of messages (e.g. all messages worldwide containing the term "tornado") from the observed count of messages on that day and divide the difference by the standard deviation of global messages (observed over the nine-month reference period previously mentioned). The ratio of z-scores gives us the ratio of spatially filtered signal to the global signal, which we call the *adjusted spatially filtered signal*. The numerator represents the magnitude of the spatially filtered keyword signal while the denominator accounts for variation in global traffic of the keyword. Adjusting for global variation in signal helps us account for spurious, time-sensitive, global variation in signal that may account for the change in signal we observe at the local level. We calculate the *adjusted spatially filtered signal* in each of the concentric bins around the event epicenter on each of our days of observation and note how the signal varies over time. We illustrate the adjusted spatially filtered signal over six days in Moore in the following section.

4.2. Moore tornado results

Calculating the adjusted spatially filtered activity signal across all geographic bins for the days surrounding the Moore tornado, we plot the results below in Fig. 2. On the y-axis we plot the adjusted spatially filtered signal and on the x-axis we illustrate the distance from the epicenter (in 100-mile bins). We use shading and line breaks to distinguish each day's adjusted spatially filtered signal. An adjusted spatially filtered signal of 0 indicates that the signal of tornado-related rumoring

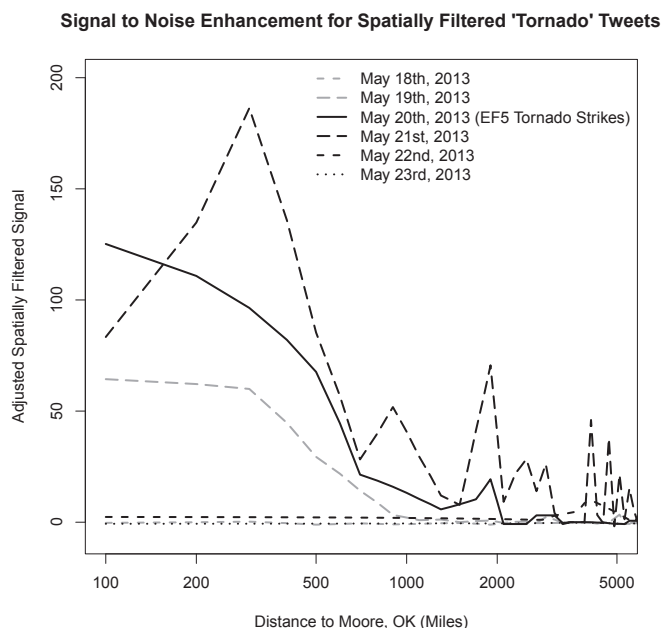


Fig. 2. Signal to noise ratio for "tornado" tweets binned at 100-mile increments from the epicenter. We have smoothed the lines with a LOWESS routine.

activity is at typical levels while a positive value indicates greater signal than normal, compared to typical activity in that spatial bin and relative to global activity on that day. Likewise, a negative score indicates reduced tornado-related rumoring activity.

On May 18th, two days prior to the Moore tornado, we observe an adjusted spatially filtered signal around 0 across all locations, indicating typical levels of tornado-related rumoring. On May 19th, however, we begin to notice a surge in signal within 1000 miles of Moore. The signal is strongest within 300 miles of Moore and drops notably at further distances. Much of this activity is in response to the warnings and alerts issued throughout much of the Great Plains region on that day, including Oklahoma, Kansas, Iowa, Illinois, and Missouri. Locations not affected by these warnings do not show surges in tornado communication. Signal of rumoring activity peaks on the day of the Moore tornado in the areas directly impacted by the tornado. We note more than a 100-fold increase in signal within 100 miles of Moore on May 20th. Sustaining approximately \$2 billion in damage, Moore was devastated following the tornado. During this disruption we observe a tremendous increase in signal of “tornado” messages. This signal declines with distance from the event epicenter, but remains strong elsewhere. From 500 to 600 miles away a 50-fold increase in signal remains, with signal declining as a function of distance to Moore.

On the day after the event, signal within 100 miles of Moore declines from the previous day, but remains nonetheless elevated with an 80-fold increase in signal. Many locations exhibit their highest signal on this day. We see great increases in activity within several hundred miles of Moore, around 1000 miles, and around 2000 miles. As news spreads about the devastation in Moore, we see a surge of attention to tornado-related topics throughout the United States. We even observe surges in activity abroad, which is reflected in the bins several thousands of miles away. Throughout the US and elsewhere, many of these messages express condolences to those affected, raise awareness of how to help recovery efforts, and spread information about the disaster. Except for some residual signal 4000–5000 miles away, activity levels return to normal on May 22nd. By this point, tornado-related rumoring activity has dropped back to baseline levels.

By selectively filtering messages across time and space, we are able to highlight the strongest sources of tornado-related rumoring. Our approach is informed by rumor theory and theories of crisis informatics, and the geo-spatial and temporal patterns of signal in response to the Moore tornado are consistent with those theories. We identify signals associated with the typical sequence of hazard-related information. Anticipatory excitation occurs locally in response to the series of warnings and alerts on the day prior to the event. Primary excitation peaks in Moore on the day of the event and declines with distance to the event epicenter. Finally, many locations’ peak activity occurs during the phase of secondary excitation, in which mass convergence of attention contributes to strong increases in signal as far as 5000 miles away. The Moore tornado was one of North America’s costliest disasters in 2013 and this pattern of response may not be representative of responses to other disaster events. To determine whether we can use our spatio-temporal filtering technique to observe similar patterns of responses, we now turn to fourteen additional events.

4.3. Cumulative results

In addition to the Moore tornado, we examined spatio-temporally filtered signals of rumoring activity in response to several disaster events in the United States and Canada. We provide some basic descriptives about each event in Table 2, as well as a rating of how strong the adjusted spatially filtered signal was at the event epicenter on the day of the event. We deliberately chose a diverse sample of events that cover a wide spectrum of disasters experienced in North America. We have natural disasters that were preceded by official warnings and alerts, including tornadoes and floods, and natural disasters that occurred spontaneously, including earthquakes and wildfires. Additionally, we have anthropogenic disasters, including those that are accidental (structure collapses) or intentional (mass shooting). The impact of the events varies substantially, from events such as the Moore tornado, Calgary floods, and Peoria tornado—whose damage estimates exceeded \$1 billion—to

Table 2
Event descriptions.

Event	Keyword	Location	Date	Population	Deaths	Injuries	Damage
EF5 tornado***	Tornado	Moore, OK	5.20.2013	55,081	24	377	\$2 billion
EF4 tornado**	Tornado	Peoria, IL	11.17.2013	119,698	2	125	\$1.6 billion
EF4 tornado**	Tornado	El Reno, OK	5.31.2013	17,510	8	151	\$40 million
EF3 tornado	Tornado	Alexander, IA	6.12.2013	175	0	0	< \$ 1 million
Bridge collapse**	Collapsed	Skagit River, WA	5.23.2013	31,743	0	3	\$18 million
Building collapse**	Collapsed	Philadelphia, PA	6.05.2013	1,526,006	6	14	< \$1 million
M7.0 earthquake*	earthquake	Atka, AK	8.30.2013	61	0	0	0
M6.5 earthquake*	Earthquake	Atka, AK	9.03.2013	61	0	0	0
M5.7 earthquake**	Earthquake	Greenville, CA	5.24.2013	1129	0	0	< \$100k
M4.2 earthquake*	Earthquake	Spanish Springs, NV	8.26.2013	15,064	0	0	0
Wildfire**	Wildfire	Yarnell, AZ	6.28.2013	649	19	23	\$1.8 million
Wildfire**	Wildfire	Black Forest, CO	6.11.2013	13,116	2	0	\$90 million
Flood***	Flood	Calgary, AB	6.20.2013	1,214,839	4		\$1.7 billion
Flood*	Flood	Boulder, CO	9.09.2013	97,385	8	0	\$1 billion
School shooting*	Shooting	Centennial, CO	12.13.2013	103,743	2	0	0

Adjusted activity signal within 100 miles of event epicenter: * >5, ** >10, *** >50.

earthquakes in Spanish Springs, Nevada and Atka, Alaska, which had no damage reported. Casualty estimates also vary significantly across events: several of the tornado events killed dozens and injured hundreds while the earthquake events had no casualties. We also have considerable variation in the number of people affected, with the Calgary floods affecting a city of over 1 million and the Atka earthquakes impacting a town of only 61. We chose this diverse array of events to determine if this spatio-temporal filtering approach is robust enough to detect signal of rumoring activity in the aftermath of all kinds of disasters.

Despite the diversity of these events in form and impact, we routinely detect a strong, local signal on the day of the disaster (with the lone exception being the Alexander, IA tornado). These findings demonstrate the robustness of this approach, as we consistently identify informal responses to these events regardless of measured impact, population size, or type of event. In Fig. 3 we average the signal across all fifteen events and plot the adjusted activity signal as a function of distance from event epicenter. We find typical levels of rumoring activity three days prior to events and two days prior to events, on average. On the day prior to the events, however, we observe a pattern consistent with anticipatory excitement. We find, on average, a roughly 5-fold increase in signal, which persists within a couple hundred miles of the event epicenter. Much of this activity is in response to warnings and alerts issued in advance of some events. On the day of the event we find a pattern consistent with primary excitation, as we observe a strong signal of rumoring activity in and around the impacted areas on the day of the event. We note more than a 20-fold increase in average signal around the event epicenter on the day of the event. During primary excitation signal declines with distance to event and, except for a small surge around 800 miles, is not distinguishable from noise after 600 miles. We observe a pattern of secondary excitation characterized by rumoring activity that matches the signal of primary excitation at 300 miles and beyond. Relative to the previous day, signal declines in the bins immediately surrounding the disaster event. Two days after the event rumoring activity returns to typical levels and stays there on the third day after the event. When and where we observe increased signal of rumoring activity is consistent with what we would expect from rumoring and crisis informatics theories.

We observed a local uptick in rumoring signal on the day before events, on average, despite our mix of anticipated and spontaneous events. For further validation of our results, we separate the results into two categories: those preceded by warnings and alerts and those without advance notice. We plot those separated results in Fig. 4.

Further validating our spatio-temporal filtering technique, we find evidence of anticipatory excitation in events *with* warnings while events without warnings have no such excitation. Furthermore, the anticipatory excitation only occurs within about 300 miles on the day prior to the event. The National Weather Service usually issues warnings over a large area spanning multiple counties, which is roughly consistent with the spatial spread of anticipatory excitation that we observe. In neither of these cases, however, do we observe an increase in signal two or three days prior to the event, before such warnings are issued. This demonstrates that, on average, we are not measuring false positives of signal. Similarly, signal generally returns to baseline levels by the second and third days after the event, after mass convergence of attention has subsided. This is further reassurance that we are not measuring false positives. This is a strong endorsement of the spatio-temporal filtering approach to measuring rumoring activity. While this approach is sensitive enough to pick up earthquake events in the

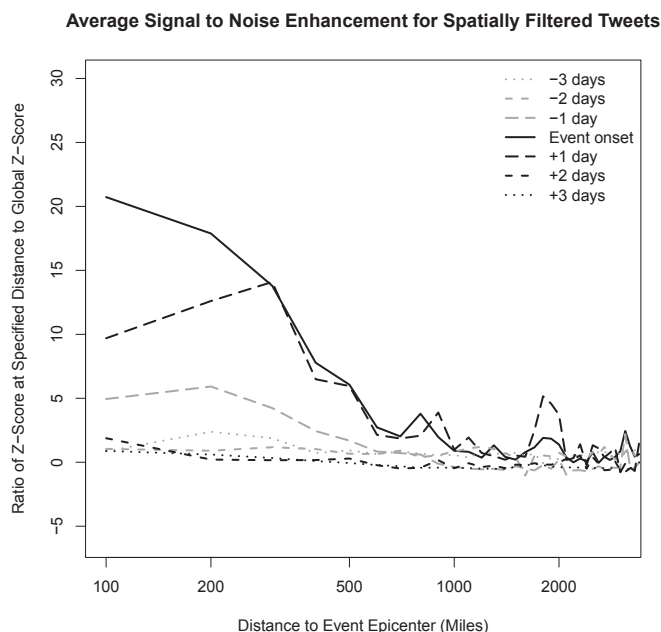


Fig. 3. Adjusted spatially filtered signal for hazard keywords tweets across fifteen different natural and man-made disaster events.

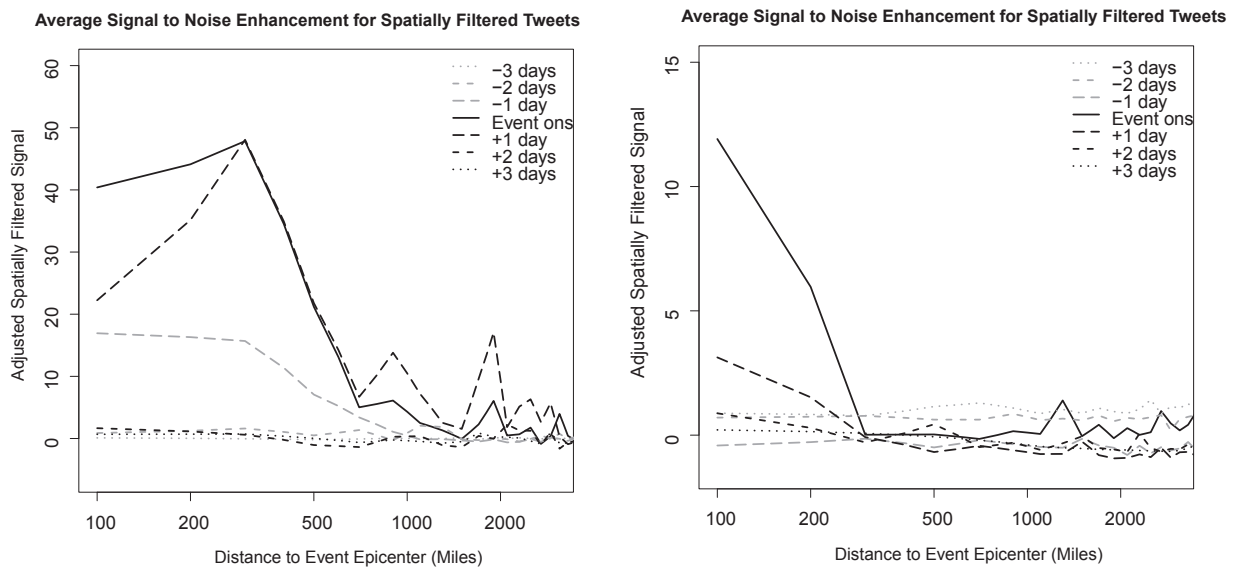


Fig. 4. Anticipatory excitation is present in events with warnings (top) but disappears when we examine events without warnings (bottom).

Aleutian Islands in Atka, Alaska, it is robust enough to noise that we do not falsely report any surges in activity signal several days before or after any event, nor do we spuriously observe anticipatory excitation ahead of disasters that are not foretold by warnings and alerts. These results help to validate the overall findings.

5. Further applications of spatio-temporal filtering

The approach we take above is one instantiation in a broader family of techniques for filtering signal from noise. As we mentioned previously, one can employ a number of different dimensions along which to filter the data. We chose time and space, as those impose strong constraints on the proliferation of disaster-related rumoring. Beyond measuring rumoring, this general approach can be used to detect aberrations in other communicative behaviors such as cell phone calls and text messages. Such communication is timestamped and easily geolocated through cell towers, which makes spatio-temporal filtering trivial. Approaches harnessing mobile phone data could build on interest in using individual communicative activity to classify locales in urban environments (Ferrari et al., 2011; Hollenstein and Purves, 2013; Leung and Newsam, 2010; Mamei et al., 2010; Pozdnoukhov and Kaiser, 2011). For example, we could identify signal of communicative activity about a particular concert, sporting event, or political rally. One would have to change the resolution of the filtering approach (likely using smaller resolution units such as hours and U.S. Census tracts), but it should nonetheless be suitable to measure signals of such activity in these and many other scenarios.

These kinds of spatio-temporal filtering approaches have grown increasingly common in other fields, such as disease surveillance. Recognizing the importance of outbreak-detection methods that harness data with spatial and temporal components, scholars have developed tools such as HealthMap, which uses spatial and temporal data to identify signals of anomalous disease activity (Brownstein et al., 2008; Freifeld et al., 2008; Robertson et al., 2010). Due to its combination of temporal, geographic, and text data Twitter has also been recognized as a valuable tool for monitoring disease outbreaks (Lee et al., 2013). Beyond using spatio-temporal filtering to measure counts of activities such as messages, phone calls, or flu cases, we can use this spatio-temporal filtering approach to refine our signal of more meaningful features of social processes, such as message content.

5.1. Spatio-temporally filtered topic signals

To provide another example of how one could implement a spatio-temporal filtering approach and to provide further validation of the Moore results, we examine how message *content* varies across time and space. While volume of rumors varies across time and space, past studies suggest the content of rumors also varies across time and space. Many describe rumoring as a *process* that helps populations engage in collective problem solving in order to explain how the world (or a small part of it) works (Rosnow, 1988; Shibutani, 1966; Sinha, 1952; Walker and Beckerle, 1987). The problem-solving process frequently follows a common pattern beginning with uncertainty, followed by speculation, and concluding with a reduction of cognitive unclarity through the development of a common understanding or belief (Bordia and DiFonzo, 2004). Populations purposively engage in rumoring, and we can harness known characteristics of the rumoring process to identify topical convergence and topical evolution over the course of disaster events. While rumor statements show variation in

content across time, they also vary across *space*. Populations across space may have different informational needs and they will attend to and disseminate information that pertains to those needs (Shklovski et al., 2008; Starbird et al., 2012). Differential attention to specific aspects of the subject of a rumor can lead populations to converge on distinct rumor subtopics during the same disaster event. Accordingly, we ought to see some variation in topics across both time and space.

To detect topical convergence and content evolution of rumor messages during disaster events, we employ the same basic strategy of filtering data across time and space. To avoid data sparsity in this example, we examine the day before, day of, and day after the Moore tornado and condense our spatial bins into three bins. The first bin includes all messages within 100 miles of the event epicenter; we call this the local bin. Next we create the regional bin, which includes all messages between 100 and 500 miles of the event. Finally, we create a bin of distant messages, which are posted between 500 and 2500 miles from the event. If rumor content varies across time and space, then we ought to be able to differentiate the topics discussed in these bins. Topics discussed at ground zero should be distinct from those discussed hundreds or thousands of miles away. Likewise, topics discussed prior to the event at ground zero should differ from topics discussed in the same location on the day after the tornado.

We use a bag of words approach to measure content across all “tornado” messages in each of the nine bins. This approach simplifies the set of messages by representing it as a multiset of words. We disregard word order (and hence grammar) and instead indicate the number of times each word appears in each bin. Instead of counting single words, however, we count bigrams. These bigrams are composed of consecutive words in each tweet. For example, a message stating (“I saw the tornado outside”) would have the bigrams “I saw,” “saw the,” “the tornado,” and “tornado outside.” Using bigrams allows us to preserve stopwords, which are typically discarded in language processing in order to prevent excess weight from being placed on terms such as “a,” “that,” or “and.” Prepositions, conjunctions, articles, and other stopwords often have subtle yet meaningful implications and we are accordingly hesitant to discard them.

In the following tables we enumerate the fifteen most common bigrams in each region on the day prior to, day of, and day after the Moore tornado. We order the bigrams based on frequency of appearance (in parentheses), beginning with local bigrams in Table 3.

On the day prior to the Moore tornado the most popular bigrams in the local bin primarily refer to tornado warnings and use language that frequently appears in such warnings. “Tornado warning” is the most common bigram and one of several related to warnings and watches. The bigrams “may 19,” “19 at,” “by nws [National Weather Service],” “cdt [Central Daylight Time] by,” “cdt until,” “issued may,” and “until may” all contain snippets of language that frequently accompanies watches issued by the National Weather Service, such as dates and times. Ahead of major storms such as the Moore tornado Twitter users frequently retransmit alerts and warnings, or parts thereof, and this bin reflects a surge of this typical retransmission activity on the day prior to the Moore tornado.

On the following day we continue to see a variety of watch and warning messages, but several of the most popular bigrams describe a tornado event in a more tangible manner. The top bigram on the day of the event is “the tornado,” which is particularly noteworthy in that “tornado” is preceded by a definite article. This language specifically refers to a tornado event that has been realized, rather than a potential event that is typically discussed in the context of warnings and alerts (which is always preceded by an indefinite article such as “a”). The prominence of “this tornado” and “in ok [Oklahoma]” provide further evidence that individuals are discussing a specific event that has occurred in a well defined location (Oklahoma). The subject of these tornado messages appears to have transitioned from discussion of warnings of a potential event to discussion of an actual event.

The topic of messages immediately surrounding Moore again changes on the day after the event. On May 20th we continue observing the prominence of definite articles, pronouns, and location-specifying prepositions in the top bigrams. Additionally, we begin to see words such as “victims” and “donate,” which reflect increased attention towards recovering from the tornado event. Over the course of this three-day period around Moore we see a topical transition in the most popular bigrams

Table 3
Moore: top-15 local bigrams.

T - 1	T	T + 1
tornado warning (90)	the tornado (93)	the tornado (55)
a tornado (81)	tornado warning (80)	tornado victims (21)
may 19 (55)	a tornado (75)	in the (17)
19 at (54)	warning for (50)	of the (17)
the tornado (53)	this tornado (46)	oklahoma tornado (16)
on the (50)	in ok (42)	in moore (15)
warning for (43)	ok until (39)	a tornado (14)
tornado on (38)	of the (38)	to the (14)
the ground (35)	19 at (36)	this tornado (13)
tornado watch (28)	may 19 (36)	to help (13)
by nws (27)	in the (34)	tornado in (13)
cdt by (27)	tornado is (28)	in oklahoma (12)
cdt until (27)	on the (27)	moore tornado (12)
issued may (27)	tornado watch (24)	this is (12)
until may (27)	counties in (23)	to donate (11)

from the local bin. From warnings to a realized event to recovery, these messages display a notable shift in content within the topic of tornadoes. Next we move from local content to the regional content illustrated in Table 4.

During all three days the most frequent bigrams in the regional bin are dominated by language typical of tornado watches and warnings. Except for the date, the top-five bigrams are identical across these three days. Although the order of the rankings varies, the next four bigrams are identical across all three dates and include snippets of text that refer to when tornado watches were issued. Before, during, and after the Moore tornado the regional bins are dominated by messages related to tornado watches and warnings as the larger storm system passes over the Great Plains region and into the Mississippi River Valley. While the content of the messages overlaps with Moore's local messages on the day prior to the event (likely because both areas received several warnings and alerts), the topics in each area diverge once the tornado strikes Moore. Finally, in Table 5 we illustrate message content in the distant geographic bin.

On the day prior to and day of the Moore tornado we observe mostly watch-related bigrams in the distant region. On the day after the Moore tornado, however, the top bigrams refer specifically to a tornado event in Oklahoma (the Moore tornado), while the remainder resemble the usual warning-related messages. The prevalence of Oklahoma-related messages suggests evidence of a mass convergence of attention from distant locales onto Oklahoma and likely onto those impacted by the Moore tornado. As news of the Moore tornado's devastation spreads to areas hundreds and thousands of miles beyond the Oklahoma City metro area, populations distant to Moore increasingly paid attention to the event.

Across the local, regional, and distant bins, these bigrams show how the content of "tornado" messages varies during the Moore tornado. Locally we observe a shift in content from warning-related language to language pertaining to a realized event to recovery-related content. Regionally we observe a continuation of warning-related language across all three days, while at larger distances we see a transition from warning-related topics to Moore-specific topics. As suggested by rumor theory, the rumor topics vary across time and space. Through spatio-temporal filtering of messages we are able to observe evolution of topic signal that is consistent with rumor theory's predictions. In the area immediately surrounding Moore we find the strongest evidence that rumoring is an evolutionary process, as the topic and tone of messages shifts notably over this three-day span. Although these results may not necessarily generalize to all events, we do find some initial support that rumoring is a process marked by topical evolution over time. Once again, the spatio-temporal filtering approach proves to be robust tool for measuring signal of hazard-related rumoring.

6. Discussion

Using spatio-temporal filtering of rumoring activity, we observed three phases of rumoring and each is consistent with theories of rumoring and crisis informatics. Prior to forecast disaster events we observe Allport and Postman (1947)'s "goal-gradient phenomenon" during a period of anticipatory excitation. This anticipatory excitation manifests as increased disaster-related rumoring where that disaster event is most salient (Allport and Postman, 1946; Caplow, 1947; Larsen, 1954). Notably, we find no such anticipatory excitation when we examine unanticipated events such as earthquakes and structure collapses. Once a disaster strikes, event epicenters are frequently characterized by disruption and anxiety, both of which contribute to increased rumoring activity (Caputo, 1999; Schachter and Burdick, 1955; Stein, 1980; Walker and Beckerle, 1987). We observe such activity during primary excitation, as signal of rumoring is consistently concentrated in locations closest to the event epicenter. As information about the event spreads to distant communities, the event receives a mass convergence of attention (Hughes and Palen, 2009; Sutton, 2010). Our results consistently show that while local rumoring begins to taper off during secondary excitation, distant rumoring matches or exceeds its signal compared to the day of the disaster event. In several of the more severe events such as the Moore tornado, Peoria tornado, and Calgary flooding, many of these messages express support of the victims, share news about the event, and discuss how to make charitable contributions to those impacted by the events, all features commonly observed during periods of mass convergence (Dynes, 1970; Hughes and Palen, 2009;

Table 4
Moore: top-15 regional bigrams.

T - 1	T	T + 1
19 at (741)	20 at (744)	20 at (254)
may 19 (741)	may 20 (744)	may 20 (254)
by nws (441)	by nws (479)	by nws (177)
cdt by (441)	cdt by (479)	cdt by (177)
cdt until (441)	cdt until (479)	cdt until (177)
tornado watch (435)	issued may (479)	issued may (177)
issued may (429)	until may (479)	until may (177)
until may (429)	tornado watch (458)	tornado watch (156)
watch issued (392)	watch issued (403)	watch issued (132)
nws storm (303)	a tornado (210)	the tornado (107)
prediction center (303)	1000pm cdt (203)	21 at (100)
storm prediction (303)	at 1000pm (203)	may 21 (100)
1000pm cdt (233)	19 at (190)	tornado warning (96)
at 1000pm (233)	may 19 (190)	a tornado (89)
a tornado (161)	tornado warning (179)	300am cdt (63)

Table 5
Moore: top-15 distant bigrams.

T – 1	T	T + 1
19 at (1268)	tornado watch (684)	the tornado (315)
may 19 (1202)	19 at (640)	in oklahoma (281)
by nws (662)	by nws (625)	oklahoma tornado (238)
tornado watch (618)	may 19 (622)	a tornado (223)
cdt by (615)	watch issued (619)	tornado in (161)
cdt until (615)	cdt by (595)	tornado watch (136)
issued may (615)	cdt until (595)	by nws (111)
until may (615)	issued may (588)	at least (110)
watch issued (575)	until may (588)	of the (110)
at 900pm (343)	may 20 (532)	watch issued (108)
900pm cdt (319)	20 at (527)	by the (103)
nws storm (307)	300am cdt (412)	21 at (100)
prediction center (307)	at 300am (412)	may 21 (100)
storm prediction (307)	a tornado (309)	issued may (94)
at 1000 pm (188)	the tornado (260)	until may (94)

Sutton, 2010). That we detect signals of all three phases of disaster-related rumoring activity—anticipatory, primary, and secondary excitation—helps to validate that we are indeed detecting signals of hazard-related rumoring.

In addition to validating the spatio-temporal filtering approach, these findings allow us to extend traditional rumor theories to large spatial scales. Notoriously difficult to measure throughout the field's history, rumoring had historically been constrained to experimental, laboratory-based (Allport and Postman, 1947; Anthony, 1973; Caplow, 1947; Schachter and Burdick, 1955; Walker and Beckerle, 1987) or small-scale observational studies (Back et al., 1950; Danzig, 1958; Kapferer, 1989; Larsen, 1954). The rise of the big data regime provides new opportunities for measuring rumoring at scales that have long been infeasible, although effectively harnessing big data can be challenging due to its overwhelming size and scope (Boyd and Crawford, 2012; King, 2011). By filtering across time and space, both of which impose strong constraints on the spread of rumor, we demonstrate that we can measure signals of rumoring activity and observe how they vary across time and large spatial scales. Measuring rumoring across these large scales gives us an opportunity to build on past rumor theories by testing their viability at previously infeasible scales. Understanding the determinants of rumoring at these large scales will transform our understanding of the phenomenon.

Using Twitter to detect rumoring activity allows us to measure rumoring at large spatial and temporal scales, but it places limits on the generalizability of our results. Activity on Twitter does not represent all rumoring activity or all online, informal communication. Pew Research Center's Internet and American Life Project provides an in-depth examination of how Twitter's demographics compare to those of the United States (Duggan and Brenner, 2013). Twitter demographics skew (independently) young, urban, and minority, while there is no significant difference in participation across gender, educational attainment, and income. These results may overrepresent rumoring activity from minority, urban, or young populations. It is not immediately clear, however, how this would bias the results, as the existing literature does not clearly demonstrate that these particular groups engage in rumoring activity that differs in volume, content, or form from other populations. Nonetheless, it is worthwhile to recognize the biased population from which these data are drawn, as a number of electronically mediated communication media will have similar biases that must be recognized when applying filtering techniques.

7. Conclusions

We have developed and validated an approach for detecting rumoring activity in an extremely noisy online environment. By selectively filtering streams of geolocated informal messaging, we are consistently able to detect surges of signal in response to—and for certain events, in anticipation of—a wide range of disaster events. We rely on keyword-based aggregation, whose accuracy is supported by Tobler's First Law of Geography and informant accuracy theories, to detect signals of rumoring activity. Comparison of message volume to baseline levels of activity—at the local level across time and the global level for that particular day—enables identification of signal among the noise in areas as remote as the Aleutian Islands or areas as noisy as Philadelphia, Pennsylvania. Due to the disruptive nature of disasters and the public's response to that disruption, we are able to identify these deviations from typical activity to locate surges of communication in response to these disasters. Following rumor theory to harness what is known about informal communication in these contexts, this spatio-temporal filtering approach has both great precision and robustness to false positives.

These results demonstrate the utility of using filtering techniques to identify and classify signals in large data streams. By filtering out data from locations where we do not a priori expect surges in activity, we greatly enhance detectable signal. This improves the effectiveness of event detection techniques, particularly in cases where we have spatial and temporal metadata associated with activity logs. Beyond illustrating changes in signal strength, these changes across distance may reflect different social processes, such as direct experience and word of mouth in the directly affected areas and reaction to news reports elsewhere. This approach is one instantiation of a broad family of filtering approaches one may take to detect signals of specific human activity in noisy environments. We demonstrate the versatility of this approach by moving beyond

rumoring volume to measure distinct rumoring *topics* that occur at various points in time and space during the Moore tornado. Beyond distance-based and daily dimensions along which to filter data, one could employ different time intervals (hours, weeks, etc.) and different spatial units (Census units such as counties, tracts, and blocks). How one should filter the data largely depends on the question one wishes to answer.

Our approach offers plenty of encouragement for studying geolocated communication activity, whether through Twitter, Facebook, mobile phone messages, or any other communication medium with geolocation. Even though geolocated messages represent a small fraction (3–4%) of Twitter's stream (Weidemann and Swift, 2013), data sparsity is not major a problem for detecting and analyzing signal in a spatially and temporally resolved manner. We are able to use this limited data to observe and analyze rumoring activity on a consistent basis. Furthermore, all these analyses were based on events that we observed with a *single* keyword. Despite data sparsity, we did not have to develop techniques to identify joint signal across multiple keywords in order to detect signal from a single event. Such techniques could certainly help boost signal of rumoring, but they were not essential prerequisites for these analyses. Finally, it is noteworthy that we were regularly able to detect signal *during* disaster events. Although disaster events may lead to power outages, congested cellular networks, population displacement due to evacuation, and individuals' preoccupation with preserving personal safety over posting messages online, we consistently observed hazard-related communication during a wide variety of events. These are promising findings for the future of rumor analyses and, more broadly, for future research using spatial and temporal information to characterize human behaviors at large scale.

Acknowledgements

The authors would like to thank Katie Faust, John Hipp, Padhraic Smyth, Athina Markopoulou, and two anonymous reviewers for their helpful comments on earlier drafts. This research was supported by NSF awards CMMI-1031853 and OIA-1028394 and Office of Naval Research award N00014-08-1-1015.

References

- Allport, Gordon W., Postman, Leo, 1946. An analysis of rumor. *Public Opin. Q.* 10, 501–517.
- Allport, Gordon W., Postman, Leo, 1947. *The Psychology of Rumor*.
- Andrade, E.L., Blunsden, S., Fisher, R.B., 2006. Modelling crowd scenes for event detection. In: *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on, Volume 1*, pp. 175–178.
- Anthony, Susan, 1973. Anxiety and rumor. *J. Soc. Psychol.* 89, 91–98.
- Back, Kurt, Festinger, Leon, Hymovitch, Bernard, Kelley, Harold, Schacter, Stanley, Thibaut, John, 1950. The methodology of studying rumor transmission. *Hum. Relat.* 3, 307–312.
- Back, Mitja D., Küfner, Albrecht CP., Egloff, Boris, 2011. Automatic or the people? Anger on September 11, 2001, and lessons learned for the analysis of large digital data sets. *Psychol. Sci.* 22, 837–838.
- Baron, Robert S., David, James P., Brunzman, Bethany M., Inman, Mary, 1997. Why listeners hear less than they are told: attentional load and the teller–listener extremity effect. *J. Personal. Soc. Psychol.* 72, 826.
- Bordia, Prashant, DiFonzo, Nicholas, 2004. Problem solving in social interactions on the Internet: rumor as social cognition. *Soc. Psychol. Q.* 67, 33–49.
- Bordia, Prashant, Rosnow, Ralph L., 1998. Rumor rest stops on the information highway transmission patterns in a computer-mediated rumor chain. *Hum. Commun. Res.* 25, 163–179.
- Boyd, Danah, Crawford, Kate, 2012. Critical questions for big data: provocations for a cultural, technological, and scholarly phenomenon. *Inf. Commun. Soc.* 15, 662–679.
- Brownstein, John S., Freifeld, Clark C., Reis, Ben Y., Mandl, Kenneth D., 2008. Surveillance Sans frontiers: Internet-based emerging infectious disease intelligence and the HealthMap project. *PLoS Med.* 5, e151.
- Caplow, Theodore, 1947. Rumors in war. *Soc. Forces* 298–302.
- Caputo, Philip, 1999. *A Rumor of War*. Random House.
- Danzig, Elliott R., 1958. The Effects of a Threatening Rumor on a Disaster-stricken Community. National Academy of Sciences-National Research Council.
- Davis Jr., Clodoveu, A., Pappa, Gisele L., Rennó Rocha de Oliveira, Diogo, Arcanjo, Filipe de L., 2011. Inferring the location of Twitter messages based on user relationships. *Trans. GIS* 15, 735–751.
- Drabek, Thomas E., McEntire, David A., 2003. Emergent phenomena and the sociology of disaster: lessons, trends and opportunities from the research literature. *Disaster Prev. Manag.* 12, 97–112.
- Drabek, Thomas E., Tamminga, Harriet L., Kilijaneck, Thomas S., Adams, Christopher R., 1981. *Managing Multiorganizational Emergency Responses: Emergent Search and Rescue Networks in Natural Disaster and Remote Area Settings*. Institute of Behavioral Science, University of Colorado Boulder, CO.
- Duggan, Maeve, Brenner, Joanna, 2013. The Demographics of Social Media Users, 2012, vol. 14. Pew Research Center's Internet & American Life Project.
- Dynes, Russell, 1970. Organizational involvement and changes in community structure in disaster. *Am. Behav. Sci.* 13, 430–439.
- Erickson, Bonnie H., Nosanchuk, T.A., Mostacci, Liviana, Dalrymple, Christina Ford, 1978. The flow of crisis information as a probe of work relations. *Can. J. Sociol.* 3.
- Fawcett, Tom, Provost, Foster, 1999. Activity monitoring: noticing interesting changes in behavior. In: *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '99*. ACM, New York, NY, USA, pp. 53–62.
- Ferrari, Laura, Rosi, Alberto, Mamei, Marco, Zambonelli, Franco, 2011. Extracting urban patterns from location-based social networks. In: *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-based Social Networks*. ACM, pp. 9–16.
- Freifeld, Clark C., Mandl, Kenneth D., Reis, Ben Y., Brownstein, John S., 2008. HealthMap: global infectious disease monitoring through automated classification and visualization of Internet media reports. *J. Am. Med. Inf. Assoc.* 15, 150–157.
- Fritz, Charles E., 1961. Disasters. In: Merton, R.K., Nisbet, R.A. (Eds.), *Contemporary Social Problems*. University of California Press, Riverside, CA, pp. 97–122.
- Go, Alec, Bhayani, Richa, Huang, Lei, 2009. Twitter Sentiment Classification Using Distant Supervision. CS224N Project Report, Stanford, pp. 1–12.
- Goodchild, Michael F., 2007. Citizens as sensors: the world of volunteered geography. *GeoJournal* 69, 211–221.
- Goodchild, Michael F., Glennon, J Alan, 2010. Crowdsourcing geographic information for disaster response: a research frontier. *Int. J. Digit. Earth* 3, 231–241.
- Goodchild, Michael F., Li, Linna, 2012. Assuring the quality of volunteered geographic information. *Spat. Stat.* 1, 110–120.
- Greenberg, Bradley S., 1964. Diffusion of news of the Kennedy assassination. *Public Opin. Q.* 28, 225–232.

- Hamid, Raffay, Johnson, A., Batta, S., Bobick, A., Isbell, C., Coleman, G., 2005. Detection and explanation of anomalous activities: representing activities as bags of event n-grams. In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, Volume 1, pp. 1031–1038, vol. 1.
- Hollenstein, Livia, Purves, Ross, 2013. Exploring place through user-generated content: using Flickr tags to describe city cores. *J. Spatial Inf. Sci.* 21–48.
- Hughes, Amanda Lee, Palen, Leysia, 2009. Twitter adoption and use in mass convergence and emergency events. *Int. J. Emerg. Manag.* 6, 248–260.
- Johnson, Jeffrey, 2014. *Knowing Your Social Network Data and Measures*. Sunbelt XXXIV.
- Kapferer, Jean Noel, 1989. A mass poisoning rumor in Europe. *Public Opin. Q.* 53, 467–481.
- Killian, Lewis M., Turner, Ralph H., 1972. *Collective Behavior*.
- King, Gary, 2011. Ensuring the data-rich future of the social sciences. *Science* 331, 719–721.
- Kireyev, Kirill, Palen, Leysia, Anderson, K., 2009. Applications of topics models to analysis of disaster-related twitter data. In: *NIPS Workshop on Applications for Topic Models: Text and beyond*, vol. 1.
- Kouloumpis, Efthymios, Wilson, Theresa, Moore, Johanna, 2011. Twitter sentiment analysis: the good the bad and the omg! *ICWSM 11*, 538–541.
- Kreps, Gary A., 1983. The organization of disaster response: core concepts and processes. *Int. J. Mass Emerg. Disasters* 1, 439–465.
- Kreps, Gary A., 1984. Sociological inquiry and disaster research. *Annu. Rev. Sociol.* 309–330.
- Larsen, Otto N., 1954. Rumors in a disaster. *J. Commun.* 4, 111–123.
- Latonero, Mark, Shklovski, Irina, 2010. Respectfully yours in safety and service: emergency management & social media evangelism. In: *Proceedings of the 7th International ISCRAM Conference—Seattle*, vol. 1.
- Lazer, David M., Kennedy, Ryan, King, Gary, Vespignani, Alessandro, 2014. The parable of Google flu: traps in big data analysis. *Science* 343, 1203–1205.
- Lee, Kathy, Agrawal, Ankit, Choudhary, Alok, 2013. Real-time disease surveillance using Twitter data: demonstration on flu and cancer. In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 1474–1477.
- Leik, Robert K., Michael Carter, T., Clark, John P., Kendall, Stephanie D., Gifford, Gregory A., 1981. *Community Response to Natural Hazard Warnings*. Technical report, DTIC Document.
- Leinweber, David J., 2007. Stupid data miner tricks: overfitting the S&P 500. *J. Invest.* 16, 15–22.
- Leung, Daniel, Newsam, Shawn, 2010. Proximate sensing: inferring what-is-where from georeferenced photo collections. In: *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pp. 2955–2962. IEEE.
- Lu, D., Mause, P., Brondizio, E., Moran, E., 2004. Change detection techniques. *Int. J. Remote Sens.* 25, 2365–2401.
- Macleod, Robb D., Congalton, Russell G., 1998. A quantitative comparison of change-detection algorithms for monitoring eelgrass from remotely sensed data. *Photogramm. Eng. Remote Sens.* 64, 207–216.
- Mamei, Marco, Rosi, Alberto, Zambonelli, Franco, 2010. Automatic analysis of geotagged photos for intelligent tourist services. In: *Intelligent Environments (IE), 2010 Sixth International Conference on*, pp. 146–151. IEEE.
- Miller, Dan E., 1992. Snakes in the greens and rumor in the innercity. *Soc. Sci. J.* 29, 381–393.
- Pak, Alexander, Paroubek, Patrick, 2010. *Twitter as a Corpus for Sentiment Analysis and Opinion Mining*. LREC.
- Perry, Ronald W., Lindell, Michael K., Greene, Marjorie, 1981. *Evacuation Planning in Emergency Management*. Lexington Books.
- Pozdnoukhov, Alexei, Kaiser, Christian, 2011. Space-time dynamics of topics in streaming text. In: *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-based Social Networks*. ACM, pp. 1–8.
- Quarantelli, Enrico L., 1954. The nature and conditions of panic. *Am. J. Sociol.* 267–275.
- Quarantelli, Enrico Louis, 1980. *Evacuation Behavior and Problems: Findings and Implications from the Research Literature*. Technical report, DTIC Document.
- Quarantelli, Enrico L., Dynes, Russell R., 1977. Response to social crisis and disaster. *Annu. Rev. Sociol.* 3, 23–49.
- Ribeiro Jr., Sílvio S., Davis Jr., Clodoveu A., Oliveira, Diogo Rennó R., Meira Jr., Wagner, Gonçalves, Tatiana S., Pappa, Gisele L., 2012. Traffic observatory: a system to detect and locate traffic events and conditions using Twitter. In: *Proceedings of the 5th International Workshop on Location-based Social Networks*. ACM, pp. 5–11.
- Richardson, R.J., Erickson, Bonnie H., Nosanchuk, T.A., 1979. Community size, network structure, and the flow of information. *Can. J. Sociol./Cahiers Can. Sociol.* 379–392.
- Robertson, Colin, Nelson, Trisalyn A., MacNab, Ying C., Lawson, Andrew B., 2010. Review of methods for space-time disease surveillance. *Spatial Spatio-temporal Epidemiol.* 1, 105–116.
- Romney, Kimball A., Weller, Susan C., 1984. Predicting informant accuracy from patterns of recall among individuals. *Soc. Netw.* 6, 59–77.
- Romney, A Kimball, Weller, Susan C., Batchelder, William H., 1986. Culture as consensus: a theory of culture and informant accuracy. *Am. Anthropol.* 88, 313–338.
- Rosnow, Ralph L., 1988. Rumor as communication: a contextualist approach. *J. Commun.* 38, 12–28.
- Rosnow, Ralph L., Kimmel, Alan J., 2000. Rumor. In: Kazdin, Alan E. (Ed.), *Encyclopedia of Psychology*, vol. 7. Oxford University Press, Oxford, England, pp. 122–123.
- Scanlon, T Joseph, 1977. Post-disaster rumor chains: a case study. *Mass Emerg.* 2, 22–27.
- Scanlon, Joseph, 2007. Sampling an unknown universe: problems of researching mass casualty incidents (a history of ECRU's field research). *Stat. Med.* 26, 1812–1823.
- Schachter, Stanley, Burdick, Harvey, 1955. A field experiment on rumor transmission and distortion. *J. Abnorm. Soc. Psychol.* 50, 363.
- Shibutani, Tamotsu, 1966. *Improvised News*. Ardent Media.
- Shklovski, Irina, Palen, Leysia, Sutton, Jeannette, 2008. Finding community through information and communication technology in disaster events. In: *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work (CSCW 2008)*. Citeseer.
- Singh, Ashbindu, 1989. Review article digital change detection techniques using remotely-sensed data. *Int. J. Remote Sens.* 10, 989–1003.
- Sinha, Durganand, 1952. Behaviour in a catastrophic situation: a psychological study of reports and rumours. *Br. J. Psychol. General Sect.* 43, 200–209.
- Starbird, Kate, Muzny, Grace, Palen, Leysia, 2012. Learning from the crowd: collaborative filtering techniques for identifying on-the-ground Twitterers during mass disruptions. In: *Proceedings of the 9th International ISCRAM Conference*.
- Stauffer, Chris, Grimson, W.E.L., 2000. Learning patterns of activity using real-time tracking. *Pattern Anal. Mach. Intell. IEEE Trans.* 22, 747–757.
- Stein, Howard F., 1980. Wars and rumors of wars: a psychohistorical study of a medical culture. *J. Psychohist.* 7, 379–401.
- Sudman, Seymour, Bradburn, Norman M., Schwarz, Norbert, 1996. *Thinking about Answers: The Application of Cognitive Processes to Survey Methodology*. Jossey-Bass.
- Sutton, Jeannette N., 2010. *Twittering Tennessee: Distributed Networks and Collaboration Following a Technological Disaster*. ISCRAM.
- Sutton, Jeannette, Palen, Leysia, Shklovski, Irina, 2008. Backchannels on the front lines: emergent uses of social media in the 2007 southern California wildfires. In: *Proceedings of the 5th International ISCRAM Conference*, pp. 624–632. Washington, DC.
- Sutton, Jeannette, Spiro, Emma, Butts, Carter, Fitzhugh, Sean, Johnson, Britta, Greczek, Matt, 2013a. Tweeting the spill: online informal communications, social networks, and conversational microstructures during the deepwater horizon oilspill. *Int. J. Inf. Syst. Crisis Response Manag. (IJISCRAM)* 5, 58–76.
- Sutton, Jeannette, Spiro, Emma S., Johnson, Britta, Fitzhugh, Sean, Gibson, Ben, Butts, Carter T., 2013b. Warning tweets: serial transmission of messages during the warning phase of a disaster event. *Inf. Commun. Soc.* 1–23.
- Tobler, Waldo R., 1970. A computer movie simulating urban growth in the Detroit region. *Econ. Geogr.* 46, 234–240.
- Vieweg, Sarah, Palen, Leysia, Liu, Sophia B., Hughes, Amanda L., Sutton, Jeannette, 2008. Collective intelligence in disaster: an examination of the phenomenon in the aftermath of the 2007 Virginia Tech shootings. In: *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)*.

- Vieweg, Sarah, Hughes, Amanda L., Starbird, Kate, Palen, Leysia, 2010. Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, pp. 1079–1088.
- Wakamiya, Shoko, Lee, Ryong, Sumiya, Kazutoshi, 2011. Crowd-based urban characterization: extracting crowd behavioral patterns in urban areas from twitter. In: *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-based Social Networks*. ACM, pp. 77–84.
- Walker, Charles J., Beckerle, Carol A., 1987. The effect of state anxiety on rumor transmission. *J. Soc. Behav. Personal.* 2, 353–360.
- Weidemann, C., Swift, J., 2013. Social media location intelligence: the next privacy battle—an ArcGIS add-in and analysis of geospatial data collected from Twitter. *com. Int. J. Geoinform.* 9.
- Weller, Susan C., Romney, A Kimball, 1988. In: *Systematic Data Collection*, vol. 10. Sage.
- Yang, Shuang-Hong, Kolcz, Alek, Schlaikjer, Andy, Gupta, Pankaj, 2014. Large-scale high-precision topic modeling on twitter. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 1907–1916.