Title
The Spacetime Continuum: Using Spatio-temporal Filtering Techniques to Revisit Rumor Theory

Author
Fitzhugh, Sean

Publication Date
2015

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

The Spacetime Continuum: Using Spatio-temporal Filtering Techniques to Revisit Rumor Theory

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Sociology

by

Sean Marcus Fitzhugh

Dissertation Committee:
Carter T. Butts, Chair
Katherine Faust
John R. Hipp

2015

DEDICATION

To Wendy
Thank you for all your support and your patience.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Professor Carter Butts has my eternal gratitude. His wisdom and tenacity have made me a better scientist.

I would also like to thank Professor Katie Faust for her continual guidance and support.

I thank the final member of my dissertation committee, Professor John Hipp, and the two additional members of my dissertation proposal committee, Professor Athina Markopoulou and Professor Padhraic Smyth, for their helpful suggestions and advice.

I appreciate the support and camaraderie of both the 2009 Sociology PhD program incoming cohort and the members of the NCASD Lab.

I thank my parents for their encouragement and unconditional love. I am thankful for all the opportunities they have given me. I also thank my grandfather for instilling in his family the importance of education

Finally, I must again thank my loving wife Wendy.

This material is based on research supported by NSF awards CMMI-1031779, OIA-1028394, and IIS-1251267 and the Office of Naval Research under award N00014-08-1-1015

CURRICULUM VITAE

Sean Marcus Fitzhugh

EDUCATION

Doctor of Philosophy in Sociology **2015**
University of California, Irvine *Irvine, CA*

Master of Arts in Mathematical Behavioral Sciences **2012**
University of California, Irvine *Irvine, CA*

Bachelor of Arts in Sociology **2009**
University of California, Berkeley *Berkeley, CA*

Bachelor of Arts in Political Science **2009**
University of California, Berkeley *Berkeley, CA*

ABSTRACT OF THE DISSERTATION

The Spacetime Continuum: Using Spatio-temporal Filtering Techniques to Revisit Rumor Theory

By

Sean Marcus Fitzhugh

Doctor of Philosophy in Sociology

University of California, Irvine, 2015

Carter T. Butts, Chair

A characteristic feature of disasters is the disruption of routine activities in a society. This disruption spurs novel activity from a variety of social units, from organizations to public officials to ordinary citizens. A major component of this response is informal communication among the public and this public communication increasingly occurs in online settings. Much of this online, informal communication falls under the formal definitions of rumoring and I harness the precision of Twitter's timestamped, geolocated messages to probe the spatial and temporal features of rumoring. A challenge of relying on streams of online data is finding signal amid the overwhelming volume of activity. To overcome this obstacle I use what is known about the spatial and temporal characteristics of both rumoring and disaster-related information transmission to develop a spatio-temporal filtering approach for measuring rumoring. I demonstrate in the first chapter that this approach offers strong increases in signal of hazard-related rumoring activity across a variety of events, both natural and anthropogenic. The results shed light onto the distribution of rumoring activity across large spatial scales. In the second chapter I use the same spatio-temporal filtering approach to identify surges in signal of content of hazard-related rumoring as I look for evidence of topical convergence and content evolution throughout the rumoring process. In the third chapter I test a variety of rumor theories by measuring which features of each U.S. county—its demographics, history of tornado events, and volume interpersonal ties to the tornado-affected county—determine its propensity to rumor about severe tornado events. Using data whose spatio-temporal precision and scope (both in geography and variety of events) have historically been infeasible, this dissertation makes several valuable contributions to rumor theories.

Chapter 1

Introduction

Over the past century, the development of broadcast media relayed through radio and television has greatly increased the speed of information dissemination to the public. During disaster events, however, the dominant channel through which individuals share information has long been informal. In disaster settings, official sources of information, such as emergency management organizations and elected officials, typically fail to outpace the flow of informal information spread among the public (Caplow, 1947; Greenberg, 1964; Kreps, 1984; Richardson et al., 1979). When populations receive insufficiently comprehensive or timely information from official sources, they turn to informal exchange of information to acquire new information and validate existing information (Bordia and DiFonzo, 2004; Caplow, 1947; Perry et al., 1981; Quarantelli, 1954). This informal exchange among individuals falls under the formal definition of rumoring: unsubstantiated (i.e., unofficial) discussion about current, newsworthy issues of interest (Allport and Postman, 1947; Bordia and DiFonzo, 2004; Caplow, 1947; Kapferer, 2013; Rosnow and Kimmel, 2000; Shibutani, 1966). Although social scientists have studied rumoring for decades, the challenges of measuring rumoring with any degree of precision outside of laboratory environments have limited our understanding of the *mechanisms* by which rumoring operates. Little is known about the spatial characteristics of rumoring, especially across large geographic scales. This dissertation uses online, informal communication data with historically unmatched precision to characterize hazard-driven informal communication at large spatial scales.

While the phenomenon of informal exchange of information has continued unabated from the radio age through the internet age, members of the public now have more media through which to exchange information amongst themselves. Many have turned to the Internet – and particularly social media sites like Twitter – to share information. With nearly 300 million members posting 500 million messages per day, Twitter has become an increasingly useful platform for event-specific communication during disasters (Guan and Chen, 2014; Sutton et al., 2008, 2013a,b; Vieweg et al., 2010, 2008). Brief, topically relevant statements that are well suited for retransmission, disaster-related Twitter messages closely adhere to the type of communication historically analyzed in rumoring studies (Allport and Postman, 1947; Caplow, 1947). Additionally, the types of communicative activity observed in disaster contexts on Twitter mirror traditional rumoring behaviors: information diffusion, collective sensemaking, and message retransmission. Accordingly, I treat online informal communication on Twitter as synonymous with rumoring and in this dissertation I will use the terms interchangeably. Because online populations respond in systematic, consistent, and measurable ways to disaster events, Twitter is an ideal online context for studying rumoring. Furthermore, the Twitter platform allows us to observe the timing, content, and location of messages, which enables us to observe rumoring behavior with unprecedented precision and across large spatial scales. The precision offered by the online environment allows novel investigations of rumoring activity and gives us the opportunity to revisit classical theories of rumoring behavior.

1.1 Rumoring and Disrupted Environments

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 et al., 1981; Leik et al., 1981). This action takes the form of emergent and often novel behavior. For example, emergency management organizations conduct response-related activities such as triage, search and rescue, and mitigation; affected individuals take protective action and seek shelter, aid, and assistance; and throughout the immediate aftermath of an event, individuals, families, and organizations seek, collect, and disseminate information pertaining to ongoing circumstances. Disasters routinely spur action from a wide variety of social units.

Much of the social action generated by disasters is informal. While most disasters have some element of emergency response activity by formal organizations (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 (i.e. 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; Leik et al., 1981; Quarantelli, 1954, 1980). 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. 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; Shibutani, 1966). A key component of this definition of rumor is that such statements are not confirmed by official sources. 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). While the general principles of the rumoring phenomenon are understood in this context, measuring informal communication with any degree of precision has been an ongoing challenge.

1.2 The Challenge of Observing Informal Communication

Although the rumor literature has a rich history, studying rumoring has posed numerous methodological challenges. Measuring person-to-person communication has been challenging in both natural, observational settings and controlled, laboratory settings. Allport and Postman (1947) describe how the laboratory setting poses several artificial restrictions on the flow and content of rumor, an argument also made by Caplow (1947). In such settings the narrative color of rumor (e.g., exaggeration, humor, excitement) is frequently suppressed as subjects focus on maintaining the precision and accuracy of their statements. This phenomenon is frequently attributed to the experimenters' instructions and the accuracy and fidelity implied by the university environment in which such experiments typically occur.

Further confounding matters, such experiments typically are short-lived and cannot capture the evolution of rumor over the span of hours or days, nor do they typically permit cross-examination of the rumorer from the rumor recipient. Although Allport and Postman concede that "indoor" rumors typically capture the "basic phenomena" of rumor spreading, such rumors lack the spontaneity of "outdoor" rumors (Allport and Postman, 1947, p. 65). Caplow, however, finds that both the form and content of rumors outside the laboratory environment have "alien" configurations that bear little semblance to indoor rumors (Caplow, 1947, p. 300). To the extent that we can monitor rumoring in a natural setting, we limit our exposure to these constraints on the behavior.

The laboratory environment's superior control, however, allows researchers to avoid some of the challenges of studying rumoring in a natural setting. Scanlon (2007) describes many of the challenges faced by Carleton University's Emergency Communications Research Unit in its decades of studying rumoring and disasters. Sampling issues are common, as defining the study population is difficult and finding the population may prove challenging during disaster contexts because individuals frequently relocate in order to follow evacuation orders, seek medical assistance, or flee an imminent threat. Informant accuracy issues also complicate matters. 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 misattributed the source of 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.

In light of these many challenges, social scientists have called for principled approaches to studying emergent phenomena during disaster environments, including rumoring (Drabek and McEntire, 2003). Such approaches ought to capture the complete process of rumoring—from inception to the spread of rumors to the elimination of ambiguity—in a methodologically sound, generalizable manner (Bordia and Rosnow, 1998). Understanding the complete process of rumoring is key to understanding the phenomenon (Rosnow, 1974). Recent studies have begun to incorporate precise temporal (Blanford and MacEachren, 2014) or spatial (Starbird and Palen, 2010) components of informal communication in response to disaster (and on rare occasion, both simultaneously (Guan and Chen, 2014)). However, the aforementioned papers only cover single case studies and the authors do not attempt to generalize the findings to rumor theories.

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, OpenSteetMap, 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. This allows us to capture the phenomena of rumoring without becoming bogged down with extraneous, unrelated communication.

1.3 Determinants of 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. It does not propagate endlessly, which is what separates rumors from urban legends and folk tales (Miller, 1992). Some have described rumoring as a task of collective problem solving (Bordia and DiFonzo, 2004; Shibutani, 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 to the 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 or interesting 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 will only persist 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). Informants may misremember the timing of when they heard or spread rumors and they may mis-attribute the sources or recipients of rumors (Scanlon, 2007). 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. Developing a sampling frame for such during these periods of population displacement is infeasible, as is measuring rumoring activity by interviewing populations across large spatial scales. Merely measuring rumoring has historically been quite challenging; measuring rumoring with accurate spatial and temporal precision has been infeasible with human informants.

In this dissertation I take advantage of the opportunities afforded by the online environment in order to characterize the spatio-temporal characteristics of rumoring. Salience effects and collective problem-solving goals limit the spread of rumors across time and space; these complement the spatio-temporal characteristics of disasters and the disruption and anxiety induced by such events. Knowing the general spatio-temporal characteristics of rumoring in disaster, I can harness online data to identify and characterize rumoring behavior. Using metadata with precise measures of time, location, and content, I can simultaneously measure rumoring activity with historically unmatched precision and across large spatial scales. This dissertation develops spatio-temporal filtering techniques to characterize the rumoring activity at large scale across a variety of disasters. The remainder of this introductory chapter provides an overview of the ongoing data set collected under Project HEROIC. I then demonstrate a spatio-temporal filtering approach that reliably measures activity levels of hazard-related rumoring across a wide variety of disaster events. Next I utilize this spatio-temporal filtering approach to determine if we can measure the *content* of rumors across time and space during disasters. The final substantive chapter examines predictors of county-level, tornado-related rumoring in the aftermath of severe tornado events. Finally, I conclude the dissertation with a reflection on its contributions and speculation on how to develop further this research.

1.4 Data: Project HEROIC

To obtain streams of online, informal communication containing a variety of disaster-related keywords, I 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 that account. These messages are timestamped and some users have opted to include geographic metadata with each message. The limited scope 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 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 9th-most visited website in the world [2]. Although users may opt to restrict access to their tweets, many follow Twitter's recommendation to contribute to the public stream by sharing their tweets publicly. Twitter has a robust search engine that allows visitors (both registered and unregistered users alike) to perform keyword-based searches in the public stream of tweets. This allows users to discover who else is talking about specific topics and promotes contribution to the public discussion. Twitter handles an estimated 2.1 billion search queries per day [3]. Twitter also supports the use of hashtags, metadata tags that link messages to informal, user-generated topic channels. For example, someone posting a tweet about Hurricane Sandy may use the #HurricaneSandy hashtag to include that message in the stream of tweets about that topic. Similar to the search function, users can click on (or directly search) hashtags to return a list of all recent posts containing that hashtag. These hashtags often serve as channels along which users discuss topical information, frequently with strangers (i.e. accounts they are not following). Hashtags are entirely user-generated and the rise and fall in popularity of hashtags is an organic process. Between the search function and hashtags, Twitter serves as a robust, powerful tool for informal, public communication.

Twitter is an increasingly useful platform for 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). In the online environment individuals 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). Topically relevant informal communication characterized by high levels of retransmission, Twitter messages closely adhere to the type of communi-

[1] https://about.twitter.com/company
[2] http://www.alexa.com/topsites
[3] http://www.statisticbrain.com/twitter-statistics/

cation historically studied in the rumoring literature (Allport and Postman, 1947; Caplow, 1947). These aforementioned 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.

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 80% 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. Reported by the mobile device's GPS, this has high levels of precision and accuracy. To test the accuracy of Twitter's reported coordinates, I enabled Twitter's location-tracking services and sent messages from a variety of locations (both indoors and outdoors) in Southern California and Northeastern Georgia during Fall 2013. I noted my position on Google Maps at the time of sending and compared it to the coordinates reported by Twitter. The location reported by Twitter was consistently within 20 feet, which is consistent with observed accuracy of commercial-grade GPS systems, such as those found in mobile phones (Modsching et al., 2006; Wing et al., 2005; Zandbergen and Barbeau, 2011). While users previously had to opt in to Twitter's geolocation services, VGI is increasingly common as the *default* setting in a variety of Twitter's mobile phone applications (and third party applications). Before the rise of geolocated tweets, reliably inferring message location was a challenge. Inferring geolocation from terse tweet text is possible, but difficult and imprecise (Davis Jr et al., 2011; De Longueville et al., 2009; Twaroch et al., 2008). Such techniques struggle to capture tweets that contain no clear geographic information in the tweet (e.g., "There is a fire a few blocks away from my house"), and may be susceptible to false positives. The reliability of GPS data is much higher and the ease of obtaining precise VGI ensures that I will not have to leverage text-based location inference techniques. Recent estimates suggest about 3.5% of tweets contain this location data (Weidemann and Swift, 2013). My own estimates based on the HEROIC data stream are consistent with these findings. Although this is a small fraction of all tweets, this still represents over 17.5 million geotagged tweets per day.

The Twitter service hosts a fertile environment for rumoring activity. The character-constrained messages are concise and conducive to rapid diffusion. In fact Twitter has its own convention for message retransmission called a *retweet*. Users retransmit one another's messages by typing "RT" (short for retweet) in front of the original message. Without the sterile laboratory environment suppressing the narrative color of rumor (Allport and Postman, 1947; Caplow, 1947), the Twitter environment provides an excellent opportunity for studying "outdoor" rumors. Despite operating outside the laboratory setting, studies on rumoring on Twitter are free from many of the methodological obstacles inherent in studying rumoring in a natural, social environment. With the precise timing and geolocation metadata of Twitter messages, I observe when and where users discuss particular topics. Furthermore, the Twitter API (discussed below) allows me to monitor rumoring *prospectively*. This allows me to observe when and where rumoring activity takes place as well as when and where it *does not* take place. Historically, the absence of rumoring has been exceptionally difficult to monitor outside experimental settings. Finally, I am able to monitor the general phenomenon of *rumoring* rather than limiting monitoring to individual rumors. With a prospective approach to studying rumoring activity measured precisely across space and time, I simultaneously tackle many of rumoring's most difficult methodological challenges.

1.4.1 HEROIC's Data Collection System

The data in this dissertation 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 dissertation uses keyword-driven data collected from Twitter's Streaming API. Under HEROIC Project we continuously collect data from 72 unique, hazard-related keywords such as "wildfire," "earthquake," "collapsed," "flood," and "shooting" which cover 19 distinct hazard types. This keyword-driven data collection continuously gathers all public tweets returned by Twitter containing the specified keywords. For this dissertation I use data from the most recent version (1.1) of the Twitter Streaming API, which allows me 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 unmatched precision despite our large-scale data collection. However, the opportunity afforded by Twitter's abundance of data is balanced by the challenge of finding signal amid an overwhelming mass of information. The first chapter of this dissertation develops an approach to refine our search such that we maximize signal of hazard-related rumoring.

Chapter 2

Spatio-temporal Filtering Techniques for Detection of Disaster-Related Communication

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 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 activity. 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. 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" and the site's collective activity is awash with signals from a nearly endless array of topics. 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 Jr 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 on sites like Twitter enhance the difficulty of identifying signals of social processes (Davis Jr 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 this 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 can harness what is known about 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 my analysis in this chapter. Here my goal is to detect aberrations in human activity in response to disasters. Using timestamped streams of geolocated, informal communication activity, I selectively filter communication streams by time and location in order to identify surges of rumoring activity in response to a disaster. Spatial filtering enhances our ability to detect events by utilizing the "signal" produced by sources that are known (or expected) to produce reliable data, thereby enhancing our ability to detect distinct activity patterns above and beyond typical global activity (i.e. background noise). I begin this chapter by characterizing different types of informal response to disaster, reviewing techniques for event detection, and examining how volunteered geographic information (VGI) enables us to use spatial filtering to identify local surges in rumoring activity. I then illustrate such a technique with a series of activity logs of online, informal communication in the context of several different types of disaster events.

2.1 Stages of Informal Response

A routine pattern of activity characterizes response to disaster events. 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 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). Milling is a classic example of a traditional rumoring behavior. This pre-event activity is not boundless, 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.

The immediate response to a hazard event primarily occurs at the site of the event as the population situates itself in its new environment and begins response activities such as rescue and mitigation. 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). This sharing of news primarily occurs where the impact of the event is salient to the population. Rumors pertaining to a particular topic typically survive only as long as that topic is relevant to the population (Allport and Postman, 1946). That is, distant populations wholly unaffected by an event will have a low propensity to rumor about it. In addition to existing while salient, rumors also thrive in disrupted environments. In disruptive and high-anxiety situations, rumors spread quickly (Caputo, 1999; Stein, 1980; Walker and Beckerle, 1987) as populations attempt to make sense of their circumstances. These salience and disruption effects also extend to spatial contexts. In addition to temporal proximity, one's spatial proximity to a disrupted environment plays a role in the propensity to rumor about the event. Those who witness an event (or its impact) typically communicate about the event with a specific audience, those to whom the event is salient and interesting. 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.

In the aftermath of the event 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 begins to emerge 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 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 may emerge 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.

A variety of social processes—including pre-event milling and generation of warning beliefs, immediate response to the disrupted environment induced by the disaster, and mass convergence following the event—generates signal in response to hazard events. Distinguishing among them and distinguishing them from typical activity can be a challenge, however. That they occur at different points in time and space implores us to use a spatio-temporal filtering approach to identify them. We ought to observe distinct rumoring activity at specific points in time and space in the time leading up to, during, and after an event. Detecting these distinct processes will be an important step towards validating a spatio-temporal filtering approach for measuring online rumoring behavior.

2.2 Anomaly Detection

Informal communication is an essential component of disaster response and different types of event-driven communication occur in particular spans of time and space. While we know that anomalous communication activity occurs in response to disaster, the challenge here is distinguishing that response from ordinary activity. Filtering signal from noise is key to identifying surges of activity, whether said 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). The primary challenge is to balance sensitivity of the alert mechanism while preserving robustness to noise in the signal (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, I treat keyword-driven samples of online informal communication as distinct streams of activity, whose disruption I identify using aggregate counts and spatio-temporal metadata to help identify changes in behavior across time and space.

To ensure proper detection of communication surges in response to disasters, I 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 I use a *profiling method* (Fawcett and Provost, 1999). The profiling method approach establishes a baseline level of normal activity and looks for 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). I follow this same approach by monitoring activity streams over an extended period of time and establishing a baseline level of activity and measures of typical levels of variation. Having established baseline measures I can then begin to look for aberrant surges of activity.

Orthogonal to the profiling method is the procedure for identifying responses to specific hazard events in the stream of activity. The abundance of keywords collected under Project HEROIC allows me to identify responses to a wide variety of disasters, including wildfires, shootings, floods, tornadoes, structure collapses, and more. To account for differences in volume and location of usage across keywords, I separate each keyword into its own distinct stream of activity. 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). I illustrate this in Table 2.1, where I list the average daily posting rate for a variety of keywords within 100 miles of three distinct areas: the small town of Carrington, North Dakota (population 2,065); the large city of Denver, Colorado (population 600,158); and the extremely large population center of New York, New York (population 8,491,079). I also illustrate the total number of geolocated messages for each term on a daily basis.

Table 2.1: 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)

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. Accordingly, I take steps to recognize the difference between an observation of 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, I use individual profiling (Fawcett and Provost, 1999) for each keyword. By characterizing the typical amount of activity across space and independently across keywords, I create multiple information streams to monitor and thereby improve my ability to detect responses to different kinds of events in different locales. We harness our ability to detect change in aggregate activity to identify signals of hazard-related rumoring.

In addition to detecting when and where aberrations in activity occur, I 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), I describe the direction and magnitude of change in activity relative to normal, baseline activity levels. Establishment of this baseline allows me to determine *when* a particular aberration occurs. With our timestamped (at the second) series of hazard-keyword communications, I use differences in the volume of keyword usage to determine the magnitude and direction of change over time.

2.3 Spatial Filtering

The spatial inhomogeneity of activity in response to disaster implores me to use spatial subsampling to detect surges in activity across space. By filtering activity into distinct blocks based on distance from the epicenter of a disaster event, I am 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 volunteered geographic information (VGI) in social media, locating online informal communication is easier than it has ever been. VGI is the voluntary creation, assembly, and dissemination of geographic information by the public (Goodchild, 2007). Historically, geographic information has largely been produced by a small collection of formal entities with access to high precision mapmaking tools. With the spread of tools such as Google Maps and GPS-enabled cell phones, individuals have unprecedented access to exceptionally precise mapmaking tools (Elwood et al., 2012). Websites such as OpenStreetMap, Wikipedia, and Wikimapia are being joined by major services such as Twitter and Facebook as platforms for individuals to provide geoinformation. The range of information shared varies considerably, from reviews of local businesses to photographic travel logs to geotagged Wikipedia entries that become accessible via online maps. Much of the recent development of VGI initiatives has fostered the production of local geoinformation. This is not a new phenomenon, however. Members of the population without expertise in geography have historically contributed to the production of geographic information, albeit less frequently and on a smaller scale. A series of land use surveys in 1930s and 1940s Britain was carried out by schoolteachers and their pupils (hardly official sources in land use), and the Audubon Society's Christmas Bird Count has regularly harnessed volunteer information about geolocated counts of birds for over a century (Elwood et al., 2012). The *widespread* usage of VGI, however, is a novel phenomenon and I harness this location data to determine where individuals are using particular event-specific keywords.

The often informal nature of VGI allows it to be produced and disseminated very quickly. During disaster events citizens function as local experts and use VGI to provide situational reports, frequently before official sources are able to verify and release their reports (Goodchild, 2007). During a series of wildfires in 2009, user-generated maps of the affected areas supplanted official sources, which were slower to incorporate details of the situation (Goodchild and Glennon, 2010). The lack of constraints on informal sources provokes questions about determining the accuracy of VGI data. 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. Geographers have addressed this by invoking 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 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. If we notice a widespread surge of "blizzard" messages during the winter in Denver, for example, we ought to be safe to assume that a response to a blizzard-related event is occurring. 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. Tobler's first law has been demonstrated to be accurate with VGI, as geotagged Wikipedia entries in nearly two dozen languages demonstrate a high level of precision and accuracy (Hecht and Moxley, 2009). Generally, consistent information across geographic space implies reliability and accuracy.

Because geographic facts are typically objective and replicable, theories of informant accuracy help to verify the accuracy of VGI. 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. In the case of VGI-aided information streams, during a disaster event any noise or error ought to be 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. We are more likely to observe an 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). These principles of informant accuracy have 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 the city as either residential, industry-driven ("office towns"), nightlife-driven, or a combination of activities (Wakamiya et al., 2011). Not unlike my purpose in this chapter, timestamped, topical, geolocated activity data enables the identification of specific types of aberrant events, such as traffic events (Ribeiro Jr et al., 2012). Through spatial filtering, geographic metadata greatly enhances anomaly detection by allowing us to focus on activity signal at the site of the event.

2.4 Anomaly Detection on Twitter

I have described a typical set of responses to hazard events and now intend to identify those responses in online data. Response to the warnings/alerts, response to the event, and post-event mass convergence all occur at distinct points in time and space. Accordingly, I ought to be able to measure different patterns in excitation across space and time. Responses to warnings and alerts ought to occur in areas where such warnings were issued and they should occur prior to the event. Reaction to the event itself ought to occur near the event epicenter immediately following the event. Finally, mass convergence of attention will occur following the event but may occur at any location. By sorting the data into distinct bins separated across time and space, I attempt to identify surges in activity consistent with the response we expect to follow a disaster.

To demonstrate how I detect informal communicative responses to disaster events, I 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 1,150 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. In Figure 2.1 I show a map of all geotagged tweets containing the word “tornado” in the United States from May 18th through May 23rd and during a typical six-day period.

We observe a substantial increase in nationwide tornado tweets during the period around Moore’s onset. I counted 14,871 geotagged tornado tweets worldwide during this six-day period, compared to a typical six-day period, in which we would find an average of 1,634 worldwide, geotagged “tornado” tweets; the z-score for such an increase over this span is 5.18. Most of the activity during the Moore event is in the “Tornado Alley” section of the United States, where several milder tornadoes touched down throughout the period. We also see large increases in tornado tweets in unaffected areas, particularly in population centers such as New York, Los Angeles, Chicago, and Atlanta. While we are primarily interested in activity surges in affected areas, this surge in activity in unaffected areas is an interesting byproduct, perhaps related to the severity of the Moore tornado. Many of these tweets outside the affected areas share information, post news stories and photos, encourage aid

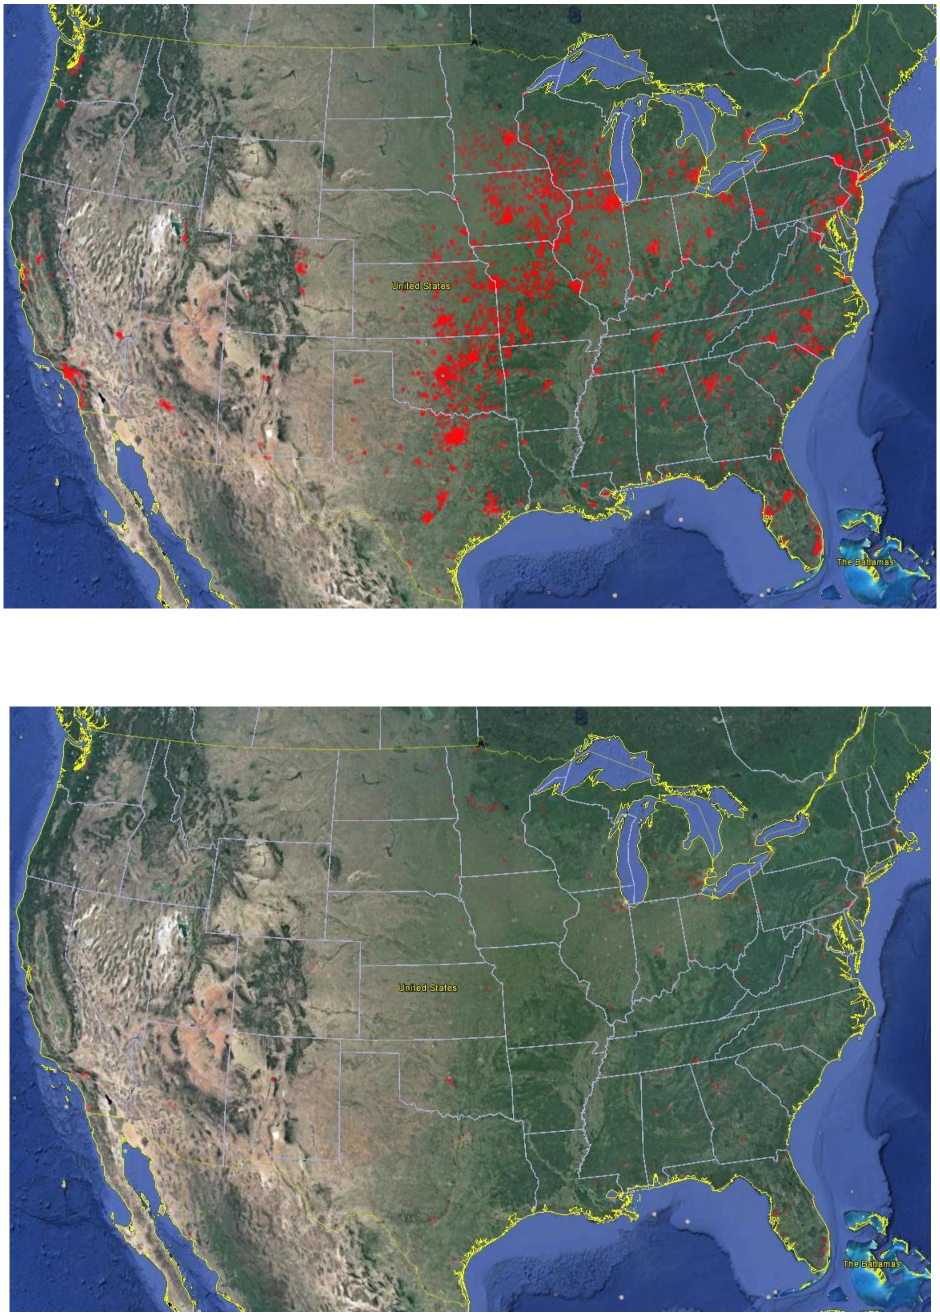

Figure 2.1: I illustrate the location of all geotagged tweets originating from the United States and containing the term "tornado." Concentrations of red indicate higher levels of tornado-related messaging activity. Above depicts the locations of geotagged tweets during the six-day period surrounding the onset of an EF5 tornado in Moore, Oklahoma. Below depicts a typical six-day period beginning on a Friday and ending on a Wednesday.

and donation, and express condolences to those affected by the tornado–all characteristics of mass convergence phenomena. In the typical six-day period I find minimal discussion of tornadoes, as demonstrated by the small, faint sections of red on the map.

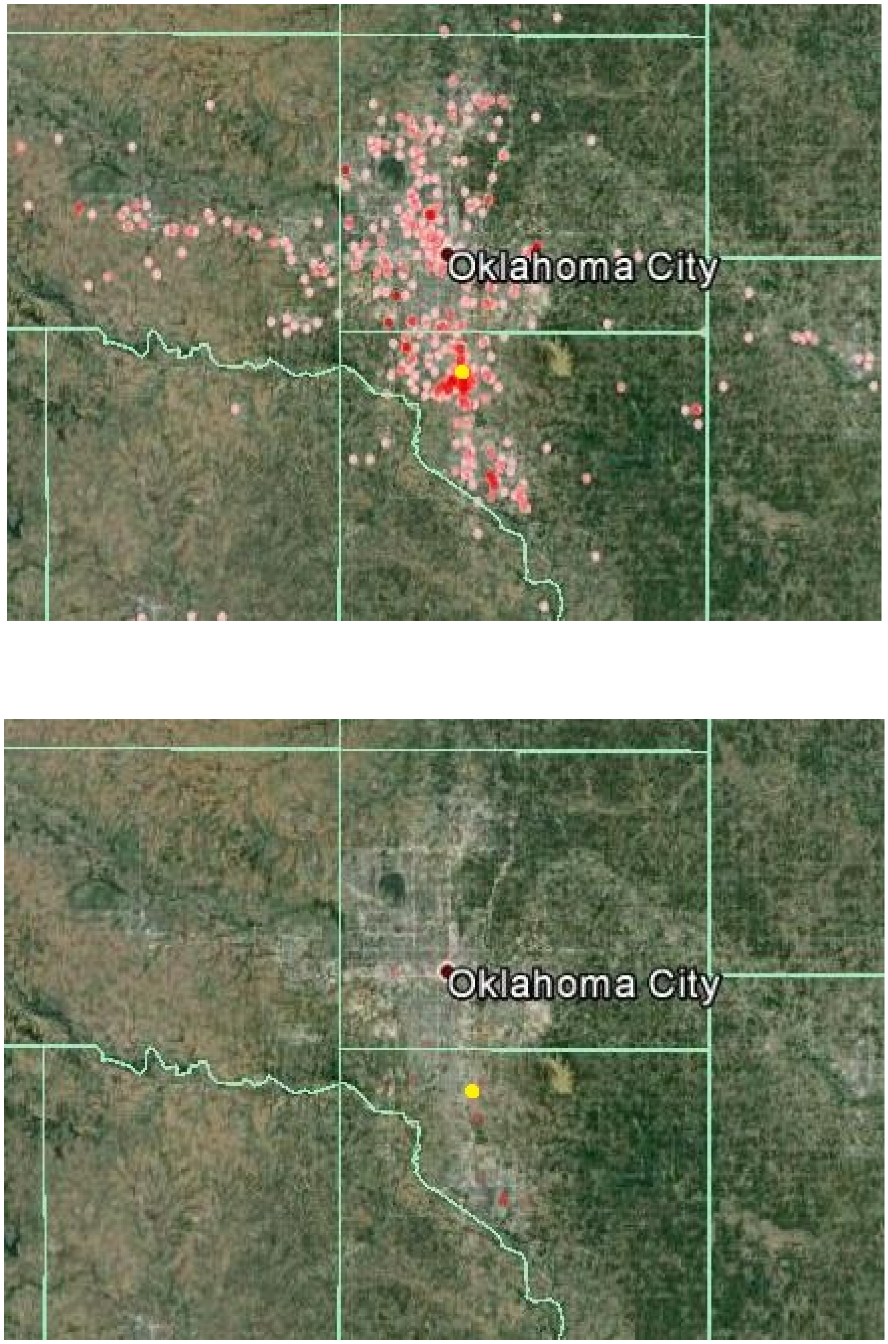

Figure 2.2: Concentrations of red represent usage of the term “tornado” in geolocated tweets in the Oklahoma City metro area, with the Moore city center highlighted in yellow. Above depicts the six-day period surrounding the onset of an EF5 tornado in Moore Oklahoma. Below depicts a typical six-day period.

I illustrate the Oklahoma City metro area in Figure 2.2. We see a substantial increase in “tornado” tweets during the disaster time period and very few “tornado” tweets during a typical period. The "tornado" tweets are widespread throughout the Oklahoma City metro area during the tornado time period, with a particularly large concentration in the city of Moore (highlighted in yellow). Of the 4,935 geotagged "tornado" tweets on May 20th, 716 originated within 100 miles of Moore. This is nearly 100 times more than the daily average of 7.45 "tornado" tweets in the same location (z-score=68.18). This helps demonstrate the utility of spatial filtering techniques for identifying signal of hazard-related communication in response to disaster events. Globally, we find 20 times as many "tornado" tweets as normal on May 20th. However, with spatial filtering we are able to find that this signal peaks around Moore, where we note 96 times as many tweets as normal. I illustrate in Figure 2.3 how filtering across space allows us to isolate elevated signals local to the disaster event, in contrast to lower (but nonetheless elevated) signals at greater distances. Spatial filtering allows us to refine our observations to identify the strongest signal of activity.

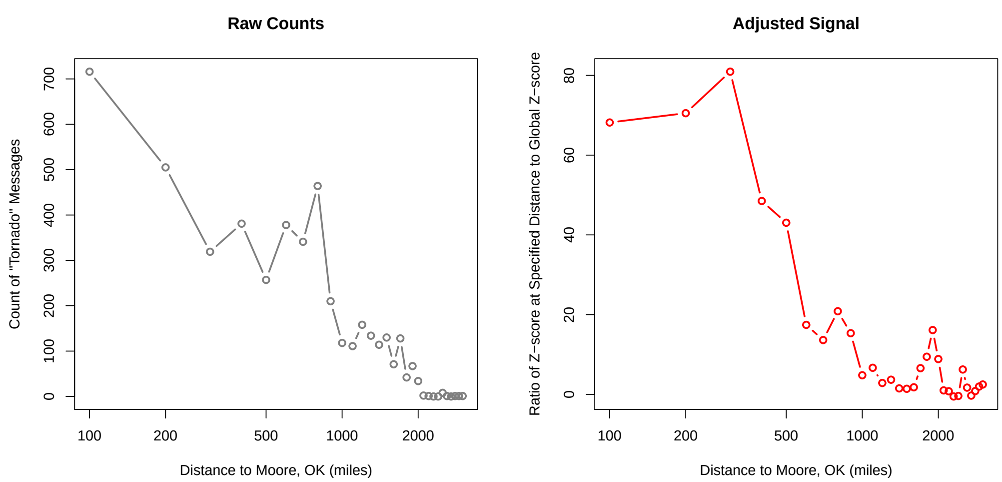

Figure 2.3: Tornado-related message activity on May 20th, day of the Moore tornado. On the left I illustrate raw counts of messages in each 100-mile bin around Moore and on the right I depict the z-scores associated with each count. Both the raw counts and z-scores indicate strong signal of activity 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 of a departure from normal "tornado" message counts as we observe between 100 and 300 miles, due to higher average activity in the 600 and 800-mile bins.

We find a large difference between the volume of tornado-related informal communication during the Moore tornado and a typical period, but I need to represent this as a systematic, quantitative difference from a baseline level of tornado keyword activity. To obtain such a baseline measurement, I compare the observed number of tweets to a measure of "normal" activity, determined through daily averages of activity of geotagged "tornado" keywords. Many "tornado" tweets throughout the year do not refer to ongoing weather events. Users frequently refer to their homes, whose rooms' disorganization and clutter resemble the aftermath of a tornado. Additionally, many weather-related "tornado" tweets refer to past events or fictional events in movies and television. All this background noise is captured in my measure of "normal" activity. Surges in activity related to an actual tornado event ought to produce a distinct signal above the typical "tornado" noise on Twitter, particularly in areas near the site of the event. I use this hypothesis to develop an signal enhancement technique that calculates the signal-to-noise ratio of activity at the local level relative to global signal-to-noise level, which I illustrate in the following section.

2.5 Spatially Filtered Activity Signals

Through spatial filtering of activity data, I calculate the signal-to-noise ratio at sequential distances from the event epicenter. I begin by taking the set of all geolocated tweets containing a keyword related to the hazard, such as "tornado," for example. Using the geographic coordinates of each tweet, I bin tweets in concentric circles around the event epicenter, such that I include all tweets within 100 miles of the epicenter in one bin, all tweets further than 100 miles from the epicenter up to 200 miles in another bin, and so on. I measure the volume of activity in each of these bins. One could employ a wide variety of approaches for measuring activity in each bin. A raw count of messages reflects the total volume of hazard-related rumoring activity. As illustrated in Table 2.1, the counts of messages vary by location based on the populations in those locations. Accordingly, raw counts will be influenced by the population in each bin. One could adjust for this by calculating the ratio of the observed count of messages to the average daily count of messages. This relative rate would indicate the posting rate on any given day relative to the average posting rate. Without recognizing the typical variation in message counts on any given day, this relative rate offers no insight into whether this departure is large or small compared to typical variation in posting rates. 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 this chapter I employ a measure based on z-scores.

In each of these bins I compare the observed number of tweets to an expected number of tweets in that bin, where the expected number is an average calculated from the baseline level of daily activity. This baseline is based on a nine-month set of messages beginning in May of 2013. As demonstrated in Equation 1, I calculate z-scores at both the bin and global levels and report the ratio of these z-scores. I define activity A as the count of messages containing a keyword (or set of keywords) k in geographic bin b during time interval t. I observe tweet volumes in 24-hour intervals, which illustrates daily changes in activity. For message volume A during a given day in a particular geographic bin, I take the difference between the observed number of messages in that geographic bin 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 bin. This follows the standard z-score calculation for how far the observed spatio-temporal activity level is from the mean activity level. I call this the *raw spatially filtered signal*. This tells us the extent to which the bin-level communication volume varies compared to activity in a typical day.

In addition to calculating the z-score for a given keyword in a given spatio-temporally filtered bin, I also calculate a z-score of all geolocated messages containing that z-score on that day.

While a z-score in the spatio-temporally filtered bin indicates variation in activity in that bin, such variation may not necessarily be specific to that bin. If we find that worldwide the count of geolocated "tornado" messages is ten standard deviations above typical activity, then we would be misguided to suggest that an increase of ten standard deviations in any given bin is a feature specific to that bin. Instead this is likely a reflection of the increased global activity. 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 I calculate the *global* z-score, which I call the global signal. The ratio of z-scores gives us the ratio of spatially filtered signal to the global signal, which I call the *adjusted spatially filtered signal*. I illustrate this equation below.

$$\frac{\left(\dfrac{A_{kbt} - \overline{A_{kb}}}{\sigma(A_{kb})}\right)}{\left(\dfrac{A_{kgt} - \overline{A_{kg}}}{\sigma(A_{kg})}\right)} \tag{2.1}$$

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 activity helps us account for spurious, global variation in keyword traffic that may account for the change in activity we see at the local level. I calculate the *adjusted spatially filtered signal* in each of the concentric bins around the event epicenter and observe how the signal varies over time.

2.5.1 Moore Tornado Results

Calculating the adjusted spatially filtered activity across all geographic bins for the days surrounding the Moore tornado, I plot the results below in Figure 2.4. On the y-axis is the adjusted spatially filtered signal and the x-axis illustrates the distance from the epicenter (in 100-mile bins). I use distinct colors to plot each day's adjusted spatially filtered signals as we move further and further from Moore, Oklahoma, the epicenter of the EF5 tornado. I find typical levels of tornado keywords during May 18th, two days before the event. On May 19th, however, we begin to notice anticipatory excitation. I tallied 590 "tornado" messages within 100 miles of Moore, well above the daily average of 7.45. This translates to a 60-fold increase in the adjusted spatially filtered signal, above and beyond the global level of "tornado" messages on that day. This signal persists within several hundred miles of Moore and declines once we get further from 300 miles away. Much of this activity is in response to the severe weather alerts issued throughout the area on the day preceding the EF5. Reflecting this surge of event-specific keyword usage, the adjusted signal for that day rises dramatically. This surge only appears in a handful of bins around Moore; locations not affected by these warnings do not show surges in tornado communication. This is consistent with the kind of activity we expect to observe in response to warnings and alerts, per rumor theory.

Primary excitation on the day of the event is characterized by a strong signal of activity in and around the impacted areas on the day of the event. I note more than a 100-fold increase in signal within 100 miles on May 20th. This surge of activity in an extremely disrupted environment is consistent with theories of rumoring during periods of disruption or anxiety. 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 we begin to notice secondary excitation which is characterized by strong signals of a mass convergence phenomenon. Signal within 100 miles declines on the

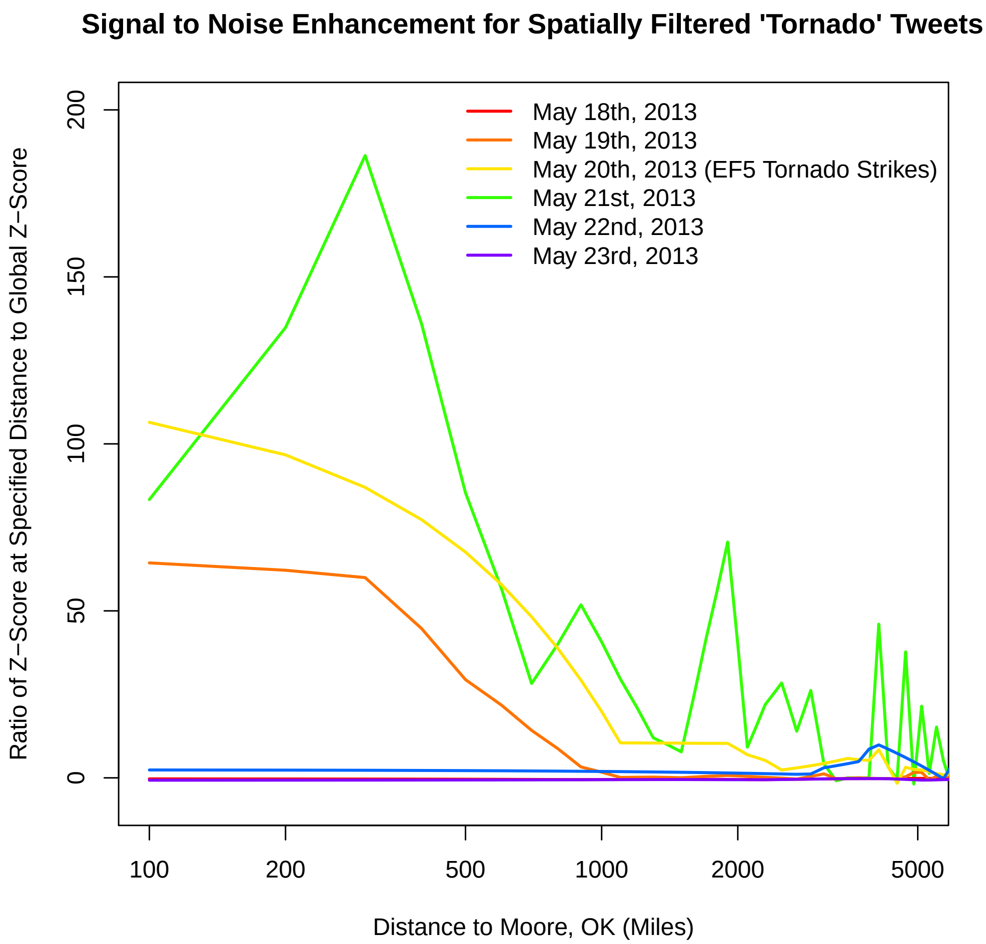

Figure 2.4: Signal to noise ratio for "tornado" tweets binned at 100-mile increments from the epicenter. I have smoothed the lines with a LOWESS routine.

day after the event, but remains elevated with an 80-fold increase in signal. Many bins 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 discussion throughout the United States. We even observe surges in activity abroad, which is reflected in the bins 4,000-5,000 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 event. These are typical elements of mass convergence phenomena (Dynes, 1970; Hughes and Palen, 2009; Sutton, 2010).

Except for some residual signal 4,000-5,000 miles away, activity levels return to normal on May 22nd. By this point, tornado-related rumoring activity has dropped back to baseline levels. The event no longer appears to be a salient topic for rumors, perhaps because most cognitive unclarity surrounding the event has been extinguished. By May 23rd all signals have reverted back to baseline levels.

Using a spatial filtering approach, I am able to identify signals associated with the typical sequence of hazard-related information. Anticipatory excitation occurs locally in response to the series of warnings and alerts. Primary excitation is concentrated 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 5,000 miles away. The Moore tornado was one of North America's costliest disasters in 2013 and this pattern of response may be atypical. To determine whether we can observe similar responses using spatial filtering, I now turn to fourteen additional events.

2.5.2 Cumulative Results

In addition to the Moore tornado, I examined spatially filtered activity signals in response to several disaster events in the United States and Canada. I provide some basic descriptives about each event in Table 2.2, as well as a rating of how strong the adjusted spatially filtered signal was at the event epicenter on the day of the event. The events cover a wide spectrum of disasters experienced in the United States. We have natural disasters and man-made disasters, both accidental (structure collapse) and intentional (shooting). Several of these events occurred without warning while others were preceded by official warnings and alerts. 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 earthquakes in

Spanish Springs, NV and Atka, AK, which had no damage reported. Casualty estimates also vary significantly across events: several tornadoes 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. Despite the diversity of these events in form and impact, I routinely detect a strong, local activity 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 I consistently identify informal responses to these events regardless of measured impact, population size, or type of event.

Table 2.2: 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	1,129	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,

In Figure 2.5 I average the signal across all fifteen events and plot the adjusted activity signal as a function of distance from event epicenter. I find typical levels of activity three days prior to events and two days prior to events, on average. On the day prior to the events, however, we begin to notice anticipatory excitation. I 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. Primary excitation is characterized by a strong signal of activity in and around the impacted areas on the day of the event. I note more than a 20-fold increase in average signal around the event epicenter on the day of the event. This surge of activity in a disrupted environment is consistent with theories of rumoring during periods of disruption or anxiety. Signal declines with distance to event and, except for a small surge around 800 miles, is not distinguishable from noise after 600 miles. Although signal declines within 200 miles on the day after the event, signal of secondary excitation roughly matches the signal of primary excitation at 300 miles and beyond. This sustained signal at great distances is reminiscent of a mass convergence phenomenon. As news spreads about events, we see a surge of attention to hazard-related discussion. In many of the more severe events such as the Moore tornado, Peoria tornado, and Calgary flooding, many of these messages express condolences to those affected, raise awareness of how to help recovery efforts, and spread information about the event. These are typical elements of mass convergence phenomena (Dynes, 1970; Hughes and Palen, 2009; Sutton, 2010).

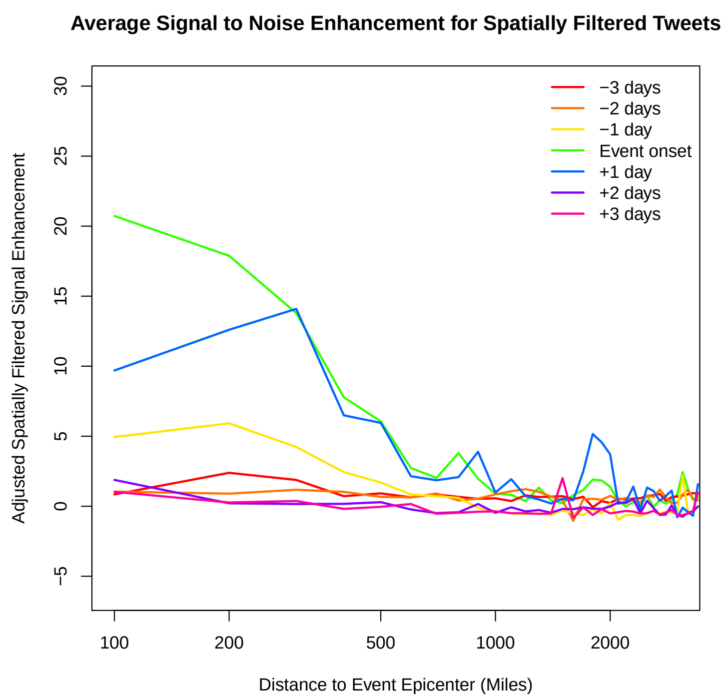

Figure 2.5: Adjusted spatially filtered signal for hazard keywords tweets across fifteen different natural and man-made disaster events.

As this is a mix of events that are foreseeable and events that occur spontaneously, only some are characterized by warnings and alerts. In Figure 2.6 I illustrate two plots: one depicting signal for events with warnings and one depicting signal for events without warnings.

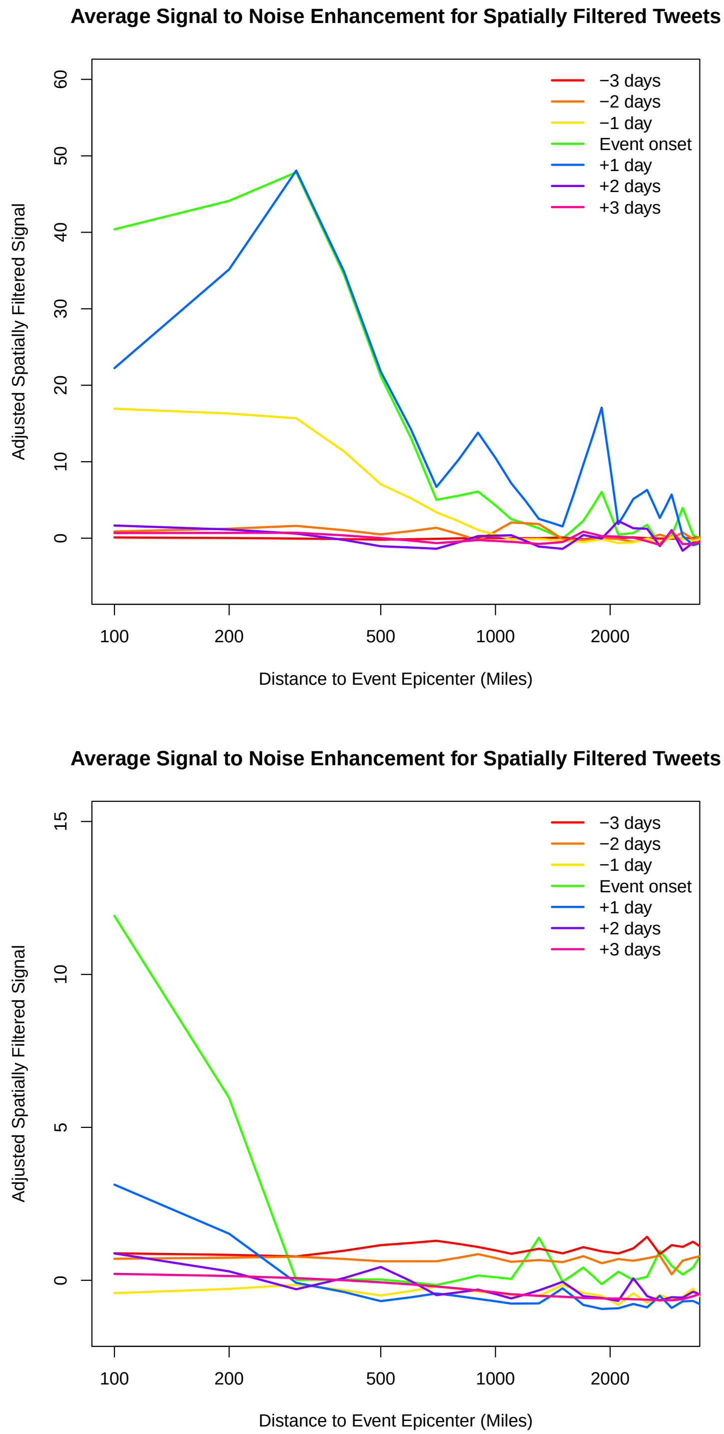

Figure 2.6: Anticipatory excitation is present in events with warnings (top) but disappears when we examine events without warnings (bottom)

As expected, I find evidence of anticipatory excitation in events with warnings while events without warnings have no such excitation. Furthermore, the excitation occurs within about 300 miles on the day prior to the event. As these warnings may span large areas, these results are reasonable. 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, I am not measuring false positives of signal. Furthermore, signal generally returns to baseline levels by the second and third days after the event, after mass convergence of attention has subsided. While the spatial filtering technique is sensitive enough to pick up earthquake events near the tip of the Aleutian Islands in Atka, AK, it is robust enough to noise that I do not falsely report any surges in activity signal several days before or after any event. These results help validate the overall findings.

Although our observations of mass convergence are constrained to events with warnings, I do not expect a direct relationship between the two to exist. Instead this appears to be related to the fact that the most destructive events we observe—events that should be characterized by post-event mass convergence—are events such as floods and tornadoes, events typically characterized by warnings and alerts.

I find clear evidence of three distinct activities occurring during these disasters. Among events with warnings, I find anticipatory excitation occurring near the site of the event impact. This is consistent with salience theories (Allport and Postman, 1946; Caplow, 1947; Larsen, 1954) and Allport and Postman (1947)'s goal gradient hypothesis. Those to whom the alerts/warnings are salient engage in greater levels of hazard-related rumoring. On the day of the event I find primary excitation peaking at the site of the event and declining with distance. Disaster epicenters are frequently characterized by disruption, anxiety, and informal communication (Caputo, 1999; Schachter and Burdick, 1955; Stein, 1980; Walker and Beckerle, 1987) and we find strong signals of this activity. Finally, I observe mass convergence on the day after severe events and the spatial heterogeneity of this response is consistent with typical patterns of mass convergence (Hughes and Palen, 2009; Sutton, 2010). By filtering the data across time and space, I am able to distinguish among these responses to disaster events.

2.6 Discussion

By monitoring streams of geolocated informal messaging, I am consistently able to detect surges of activity in response to—and for certain events, in anticipation of—a wide range of disaster events. Following longstanding rumor theories and using VGI whose accuracy is supported by Tobler's First Law of Geography and informant accuracy theories, I filter these data streams accordingly to enhance signal and eliminate noise from distant locations. I find the strongest signal-to-noise ratio close to the site of disaster, with the activity signal steadily declining with distance to the epicenter. Comparison to baseline levels of activity—at the local level across time and the global level for that particular day—enables detection 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, I am able to harness these deviations in typical activity to locate surges of communication in response to these disasters.

These results demonstrate the utility of using filtering techniques to identify and classify signals in 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. Finally, the results support rumor theories and theories of mass convergence by demonstrating that certain responses have distinct patterns across space and time. The proliferation of VGI data enable us to revisit and refine classical theories with data whose precision was previously unattainable.

Chapter 3

Spatial Variation in Hazard-Related Content in Online Informal Communication

Over the past few decades the volume of text transmitted electronically has grown dramatically. Major newspapers publish volumes of stories on a daily basis, mobile phone users exchange text messages, and the public increasingly posts text through weblogs and sites such as Twitter and Facebook. Identifying topics in these voluminous streams of text is an ongoing research frontier (Hoffman et al., 2010; Ramage et al., 2010; Zheng et al., 2006). One of the primary challenges in this domain is that the population of documents (articles, text messages, tweets) in these streams is too large to examine manually, as reading millions of documents is infeasible. This has spurred the development of automated and semi-automated methods for identifying signals in text, a research problem that has increasingly drawn the interest of social scientists (Hopkins and King, 2010; Lucas et al., 2015; Roberts et al., 2014).

Deciphering topics in length-constrained streams of messages poses unique challenges for language processing. Media such as text messages and Twitter posts impose strict limits on the text length, resulting in linguistic idiosyncrasies such as acronyms, obscure shorthand, and grammatical oddities (Kireyev et al., 2009). When performed without a strong understanding of the underlying text, automated analysis of this kind of text can lead to confounding or utterly incorrect results (Back et al., 2011; Bollen et al., 2011; Lansdall-Welfare et al., 2012). To avoid these pitfalls, I test the viability of an approach that leverages a corpus of human-coded messages to identify and categorize topics in streams of terse text. One of the key advantages is that this relies on the expertise of human coding to develop a corpus of coded messages, which we can use to identify topics in *subsequent* events. This eliminates the need for hand-coding every case study one wishes to analyze.

In the previous chapter I demonstrated that we can harness what is known about the rumoring phenomenon to develop spatio-temporal filtering techniques that identify rumoring activity amid disaster events. This chapter follows the same spatio-temporal filtering approach of the previous, but uses this approach to identify surges in particular message *content* during disaster events. Rather than identify changes in volume in response to disaster, this chapter identifies surges in content resulting from topic coalescence during the rumoring process. By harnessing what is known about when and where topics converge during rumoring (and how the convergence varies across space), this chapter assesses the feasibility of using this spatio-temporal filtering approach in conjunction with coded messages to measure content of hazard-related rumoring.

3.1 Content Evolution in Rumors

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). Rosnow (1974) describes rumoring as a three-step process, beginning with the parturition stage during which rumor statements are generated. While rumoring activity increases during this period, rumor statements converge onto the subject of interest. This period is followed by the diffusion stage when some of those rumor statements gain traction and begin to spread through a population. The final stage is control, when the rumor achieves saturation among the population or facts become disclosed by an official source. After this final stage rumoring activity dissipates as the rumoring process comes to an end.

As the rumoring phenomenon undergoes its typical pattern of evolution, the content of rumor statements likewise undergoes evolution. Allport and Postman (1947) were first to describe three of the primary processes that shape rumor statements: leveling, sharpening, and assimilation. Leveling refers to the process by which a message loses non-essential details and becomes shorter and more precise as it spreads from person to person. Complementing leveling, sharpening is the selective reporting of details from a larger context. Finally, assimilation refers to the distortion of memories, as subjects recall what *ought* to have been observed rather than what was observed (for example, a marquee at a movie theater that mentioned actor "Gene Antry" was usually recalled as "Gene Autry"). These processes simplify messages to make them more suitable for retransmission, but they also help ground the rumor in reality (or, at least, what reality *ought* to be) in order to improve its function as a sensemaking tool. These developments typically catalyze the convergence of rumor statements onto a particular topic by trimming any extraneous information. This development helps to focus rumors onto a specific substantive problem for the population to solve. 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). Upon first hearing a rumor, people generally consult with others to determine the validity of the rumor (Caplow, 1947; Kapferer, 1989). In their observations of rumors among internet discussion groups, Bordia and DiFonzo (2004) find mostly interrogatory messages in the early phase of rumoring, followed by disbelief messages. Sensemaking then peaks and discussion shifts away from rumoring as digressive messages peak at the final time period. With most uncertainty eliminated in this final time period, individuals divert their attention to new, unrelated topics. Bordia and DiFonzo (2004) recognize this systematic evolution of content and argue that this illustrates that rumoring is a task of collective problem solving. The convergence of rumors onto a particular problem and the population's process of understanding that problem are characteristic features of rumoring behavior. Populations purposively engage in rumoring and we can harness known characteristics of the rumoring process to identify topical convergence and evolution over the course of disaster events.

While rumor statements show variation in content across time, they also vary across *space*. As discussed in the previous chapter, rumor diffusion decreases as a function of distance due to salience effects (Caplow, 1947). Among populations that do engage in rumoring during a disaster event, different populations will rumor about distinct aspects of that event that they find to be salient. The content of their rumors will be differentiated accordingly. 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). In a study of rumoring during a forest fire, Larsen (1954) finds that the local "in group" of firefighters and local residents discussed property damage and casualties while the distant "out group" of public officials and reporters focused on identifying leadership during the event and determining who was to blame (Larsen, 1954, p. 115). Although the fire event was salient to both populations, each focused on unique features of the event and the subject of their rumors reflects that. Differential attention to specific aspects of the subject of a rumor can lead populations to converge on distinct rumor topics during the same disaster event.

Although not described in the context of rumor theories, literature on disasters and disrupted environments suggests a convergence and evolution of information discussed informally among those affected by the event. Perry et al. (1981) identify four components that define a post-warning environment, all of which have analogs in the rumoring literature. These components include: confirming the warning, developing a "warning belief" to determine whether a threat exists, determining how threatening the risk is, and concluding that a protective response may or may not be necessary. This evolution is similar to the interrogatory-disbelief-sensemaking-digressive sequence described by Bordia and DiFonzo (2004). Others find a common sequence of events where people engage in milling behavior in order to verify official information, after which they determine whether or not to act on that information (Drabek et al., 1981; Leik et al., 1981; Quarantelli, 1954, 1980). This pattern of post-warning rumoring behavior is reproduced in online, informal communication as the public increasingly turns to Twitter for information sharing, collective sensemaking, and message retransmission (Sutton et al., 2008, 2013a,b; Vieweg et al., 2010, 2008). Time-ordered convergence of message topics has also been reported in online, informal communication as "siren" and "tornado" messages operated on different "life cycles" during the Moore tornado (Blanford and MacEachren, 2014). Message topics also vary across space during disaster events, as individuals who self-identify as locals were more likely to retransmit locally relevant information during a series of Oklahoma fires while non-local individuals shared more abstract information (Starbird and Palen, 2010). Starbird et al. (2010) report similar findings from the Red River floods of 2009: those who posted autobiographical messages or adapted existing knowledge were primarily locals. These recent case studies provide initial evidence that online informal communication follows typical rumoring processes of convergence onto a set of problems and a content evolution as the population strives to eliminate cognitive unclarity.

3.2 Spatio-temporal Content Filtering

To detect topical convergence and content evolution of rumor messages during disaster events, I return to the spatio-temporal filtering approach employed in the last chapter. I use the same timestamped, geolocated streams of messages containing hazard-related keywords and put them into bins based on the timestamp and location. Again, I build bins that span seven days (three days prior, the day of event onset, and the following three days) around the disaster event. Instead of creating 100-mile-wide bins in concentric circles around the event epicenter, I create three spatial bins. The first includes all messages within 100 miles of the event epicenter; I call this the local bin. Next I create the regional bin, which includes all messages between 100 and 500 miles of the event. Finally, I create a bin of distant messages, which are posted between 500 and 2500 miles from the event. Compared to the spatially filtered bins of activity, several of these bins have been consolidated to reduce sparsity of messages in certain bins. Per theories of rumoring and information transmission during disasters, I 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 ought to vary over time, with a convergence on a set of rumors during the onset of the disaster event, an establishment of a set of ground truths to eliminate cognitive unclarity, and a divergence from those topics once the disruption induced by the disaster has passed.

3.2.1 Content Codes

To assess the content of each of these spatio-temporal bins, I compare each bin's content to a corpus of messages that have been human-coded for a variety of hazard-related topics. This corpus of over 12,000 coded tweets comes from five different disaster events: the Boston Marathon bombing of 2013, Hurricane Sandy, Waldo Canyon Fire, Boulder floods of 2013, and Winter Storm Nemo. Drawn from the HEROIC database, these sets of tweets come from a targeted sample of accounts that covers all officials who were actively posting messages during the event. These accounts represent a mix of appointed and elected public officials whose duties include some aspect public safety and those tasked with providing information during disasters, including news media. For each event we collected all tweets posted by those accounts during the period of imminent threat or hazard, which ranged in duration from approximately twenty-four hours during the Waldo Canyon Fire to five days during the aftermath of the Boston Marathon bombing. These messages were independently coded by Jeannette Sutton and Britta Johnson (Sutton et al., 2015, 2013b), two members of Project HEROIC, and cross-checked for a sufficient level of inter-coder reliability. We identified several content themes that were common to the five events and coded each message for its presence of that theme. The content codes represent a mix of directives and information dissemination in addition to more amiable content such as gratitude and encouragement. The former categories represent typical statements from emergency managers and others relaying firsthand information during disasters while the latter categories appeared when organizations would thank others for their assistance, thank the public for cooperation, or promote cohesion and unity in the aftermath of a disruptive event (the "Boston Strong" statement was common in these kinds of cohesion-building statements). None of the content elements are mutually exclusive; messages may belong to multiple categories. I list in table 3.1 the content elements for which each message was coded.

These coded messages are drawn from events representing a variety of disasters, both natural and anthropogenic. Additionally, they represent a variety of impacted populations, from the major metropolitan areas impacted by the Boston Bombing and Hurricane Sandy to moderately sized cities of Boulder and Colorado Springs, CO. The diversity of event types and populations impacted limit the chance that I will mistakenly classify similarity based on hazard type or population characteristics rather than message content.

Table 3.1: Coded content categories

Category	Content
Advisory	Advice that guides people on specific actions they should take in response to the event
Closures/Openings	Closure and re-opening of roads, facilities, public transportation
Corrections	Corrections to previously posted information
Emotive/Judgment	Comments expressing emotion, judgment, evaluative content, or encouragement
Hazard Impact	Descriptions including when and where the event will strike, the impact of the event, and descriptions of damage
Help/Directed Comm.	Direct requests for assistance
Information	Any messages that provide information regarding ongoing response and recovery actions
Re-entry	Information about returning after evacuation
Thank Yous	Messages offering gratitude and thanks
Volunteer/Donate/Help	Information on how to help

3.2.2 Assessing Topic Similarity

I use pairwise comparison between the spatio-temporally filtered bins of observed messages and the sets of coded messages to determine how well the content of each bin aligns with the content in the set of messages belonging to a particular code. By observing the similarities between bins and our coded corpus over time, I observe how the content of observed messages varies across time and space and determine if there is convergence on any particular topic or set of topics. I use a bag of words approach to measure content of the bin and the content of the set of messages belonging to a code, each of which I will refer to as a *document*. The bag of words approach simplifies each document by representing it as a multiset of words. This approach disregards word order (and hence grammar) and represents the document as a vector. Each entry in that vector represents a word and the value refers to the number of times the word appears in the document. While frequency of appearance is the simplest value, other weights based on bigrams, trigrams, or term frequency–inverse document frequency (tf-idf) are also frequently employed. Using this simplified approach we can derive topical information from each document based on the weights associated with each term.

I use cosine similarity to assess the topical similarity between the two documents. Cosine similarity calculates the cosine of an angle between two documents in vector space (Singhal, 2001). To aid exposition, I visualize the concept in Figure 3.1. I illustrate four color-coded documents containing a total of two terms (for the sake of simplicity), which creates a two-dimensional space for these documents to inhabit. We can represent each document with a vector from the origin to its position in space, where position is is based on the weights of terms in the document.

Documents further apart in this topic space are more dissimilar while documents closer together have more similar topics. I quantify similarity based on the cosine of the angle between the vectors representing documents. In the schematic figure, documents 1 and 3 share no terms and the angle between their vectors is 90 degrees. Maximally dissimilar in this space, the cosine similarity between them is 0. A 45-degree angle separates document 2 from document 1 and document 3; the cosine similarity between 2 and 1 and between 1 and 3 is .707, suggesting notable similarity between them. With an angle of 0 between the vectors of documents 3 and 4, their cosine similarity is maximal at 1. Although documents 3 and 4 contain different counts of the term "alert," the cosine similarity approach normalizes counts of terms to prevent measures of (dis)similarity from reflecting differences in magnitudes of word counts. Accordingly, document 2 would sit on the same vector as a document stating "tornado alert" and a document containing the terms "tornado" and "alert" 100 times each. While this two-dimensional topic space is easy to represent through visualization, in practice I will calculate cosines between vectors in k-dimensional space, where k is the total number of terms appearing in either document.

I calculate the cosine similarity between each pair of documents, where one document is the set of messages in a spatio-temporally filtered bin and the other is the set of all messages belonging to a specific code (e.g. information, advisory, emotive, etc.). I represent each

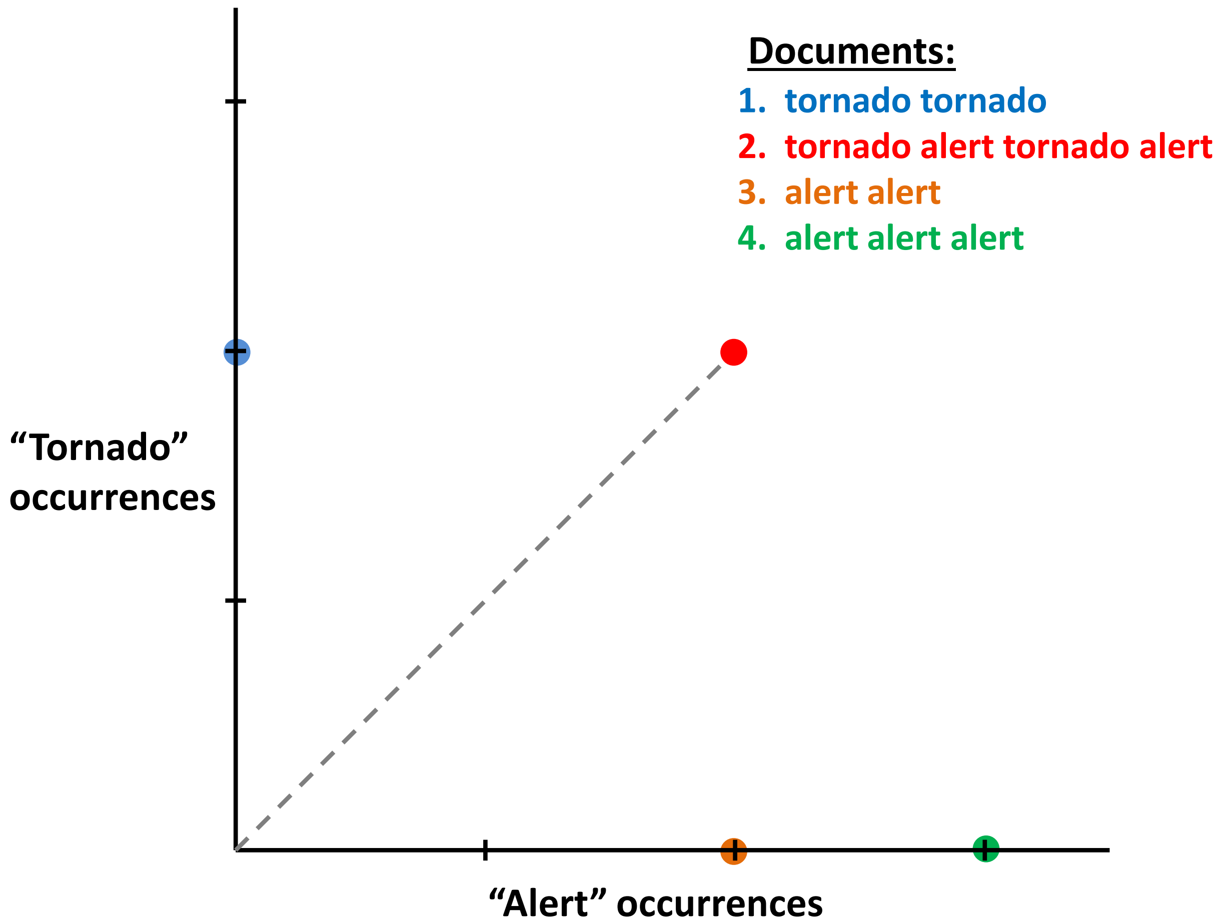

Figure 3.1: Cosine similarity measures the similarity of documents when projected into vector space, where location is defined by the frequency of terms in that document. The size of the angle between documents in this space is proportional to their dissimilarity. Orthogonal to the vectors of documents 3 and 4, document 1 is maximally dissimilar from documents 3 and 4. Because 3 and 4 sit on the same vector, they are maximally similar. Document 2 is moderately similar to document 1 and documents 3 and 4. Because the documents only contain two terms among them, I only need two dimensions to represent the vector space. In practice, vector space often has hundreds or thousands of dimensions.

document by a vector of weights representing the document's bigram distribution. These bigrams are composed of consecutive words in each tweet. For example, a message stating (“I saw a tornado outside”) would have the bigrams “I saw,” “saw a,” “a tornado,” and “tornado outside.” Using bigrams allows me 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 I am accordingly hesitant to discard them. By using bigrams, I preserve these stopwords without allowing them to dominate the vector weights. For each document, the weights are equivalent to the total number of times each bigram appears in the document.

3.3 Results

Following the same adjusted spatio-temporally filtered signal enhancement technique from the previous chapter, I calculate how much signal enhancement we achieve through spatio-temporal filtering of tweets. Instead of identifying enhancement of activity signal, however, I measure enhancement of *topic* signal. Consider an example where I measure the amount of emotive topic enhancement we attain in a local bin of "tornado" messages on the day of a disaster event. To calculate the topic enhancement I measure the cosine similarity between the observed set of local "tornado" messages on the event day and the set of messages coded for the presence of emotive content. I compare the observed local-emotive cosine similarity to the expected local-emotive cosine similarity, calculated from the average daily local-emotive cosine similarity. I calculate a z-score by taking the difference between the observed and average emotive-local cosine similarity and dividing it by the standard deviation of the daily local-emotive cosine similarities. This gives us the raw spatially filtered topic enhancement, a measure of how much more similar the "tornado" messages in the local bin are to the emotive messages on that given day, compared to a typical day. I adjust for global topic similarity by calculating the global z-score, where this represents how much more (or less) emotive content we observe among *all* "tornado" messages on the event day, relative to typical global variation. I demonstrate in Equation 3.1 how I calculate z-scores at both the bin and global levels and report the ratio of these z-scores.

$$\frac{\left(\dfrac{CS_{btc} - \overline{CS_{bc}}}{\sigma(CS_{bc})}\right)}{\left(\dfrac{CS_{gtc} - \overline{CS_{gc}}}{\sigma(CS_{gc})}\right)} \tag{3.1}$$

In the general case I define cosine similarity *CS* as cosine similarity between our observed messages in geographic bin b during time interval t and the set of messages coded for topic c. The numerator gives us the *raw spatially filtered topic signal* while the denominator indicates the global signal. Together the ratio of these two gives us the *adjusted spatially filtered topic signal*, which tells us how similar messages in a bin are to a topic, net of the typical similarity of messages in that bin to that topic *and* net of global similarity to that topic on that particular day.

For each of the seven days surrounding a disaster event, I calculate the adjusted spatially filtered topic enhancement (ASFTE) at the local, regional, and distant levels for all ten of the coded categories. This is a computationally expensive process that requires me to calculate the cosine similarity between each code and each of the four bins (local, regional, geographic, and all global messages) for the entirety of the nine-month reference period. This is necessary in order to establish the mean and standard deviation of similarities for each category in each bin. I illustrate in Table 3.2 the number of messages in each content category and for each keyword (globally) on a daily basis. Due to the expense of parsing the messages, constructing vectors of bigrams, and comparing them, I use a subset of events examined in the previous chapter.

In Table 3.3 I list the events for which I calculate the adjusted spatially filtered topic enhancement. This covers a mix of events preceded by warnings and events occurring spontaneously.

Table 3.2: Counts of messages and bigrams

Document	Messages	Bigrams
Content Codes		
Advisory	395	4,506
Closure/Opening	424	3,959
Correction	20	233
Emotive	119	1,543
Evacuation	115	1,250
Impact	207	2,578
Help/Directed Comm.	50	708
Information	882	9,802
Thanks	155	1,712
Volunteer/Donate	29	344
Daily Keyword Streams		
"Collapsed" avg.	41.7	599.2
"Earthquake" avg.	930.8	5,840.10
"Flood" avg.	385.2	3,546.70
"Shooting" avg.	651.8	7,281.20
"Tornado" avg.	248.8	2,434.20
"Wildfire" avg.	21	231.4

Additionally, these events inflict a wide range of impact, with some having no reported damage and some causing billions of dollars in damage.

Table 3.3: 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	17510	8	151	$40 million
Building collapse	collapsed	Philadelphia, PA	6.05.2013	1,526,006	6	14	< $1 million
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
Flood	flood	Calgary, AB	6.20.2013	1,214,839	4		$1.7 billion
School shooting	shooting	Centennial, CO	12.13.2013	103,743	2	0	0

I illustrate the average spatially filtered topic enhancement in Figure 3.2, with each frame representing a different geographic region. Each coded topic is distinctly colored to show its individual topic enhancement over a seven day-span of time surrounding the disaster event. The y-axis indicates the adjusted spatially filtered topic enhancement, with higher values indicating more similarity with that topic (net of both typical daily similarity and global similarity on that day). The x-axis indicates the day relative to event onset, beginning three days prior to the event onset and ending three days following the event onset. In the local region we observe a distinct increase in topic signal on the day of the event; this activity subsides on the following day and returns to typical levels on the second day after the event. The increase in topic enhancement suggests that messages in that local bin are more closely aligned with each category of coded messages. We find less topic enhancement in the regional bin and no clear signal enhancement in the distant region, suggesting less similarity between topics discussed in those regions and the coded corpus of messages. The Appendix contains tables with significance tests for the plotted values.

Although we observe topic enhancement in the local region, there is notable similarity among the signal enhancement for each individual topic. This suggests that I am picking up some dimension of hazard-related communication that is common to each individual coded category, rather than observing distinct signal from each of the ten categories. The homogeneity of signal across categories may represent an emergency management dimension of discussing disasters that is common to emergency managers, public officials, and news media that compose the accounts in the coded corpus. Language they use may be more germane to topics at ground zero. This surge in similarity of messages in the local bin with the coded messages suggests a topical convergence on the day of the event. This is a distinct contrast from the pattern of signal enhancement in the regional and distant bins. That we observe notable differences between regions is consistent with rumor theories and theories of information diffusion during disasters. Groups in different locations have different informational needs during a disaster and the content of their rumor statements will reflect this. These results offer some support of these theories and, more generally, endorse a spatio-temporal filtering approach for isolating particular rumoring signals.

While the average adjusted spatially filtered topic enhancement shows some increase in local hazard signal in anticipation of the event, not all of these events were preceded by warnings

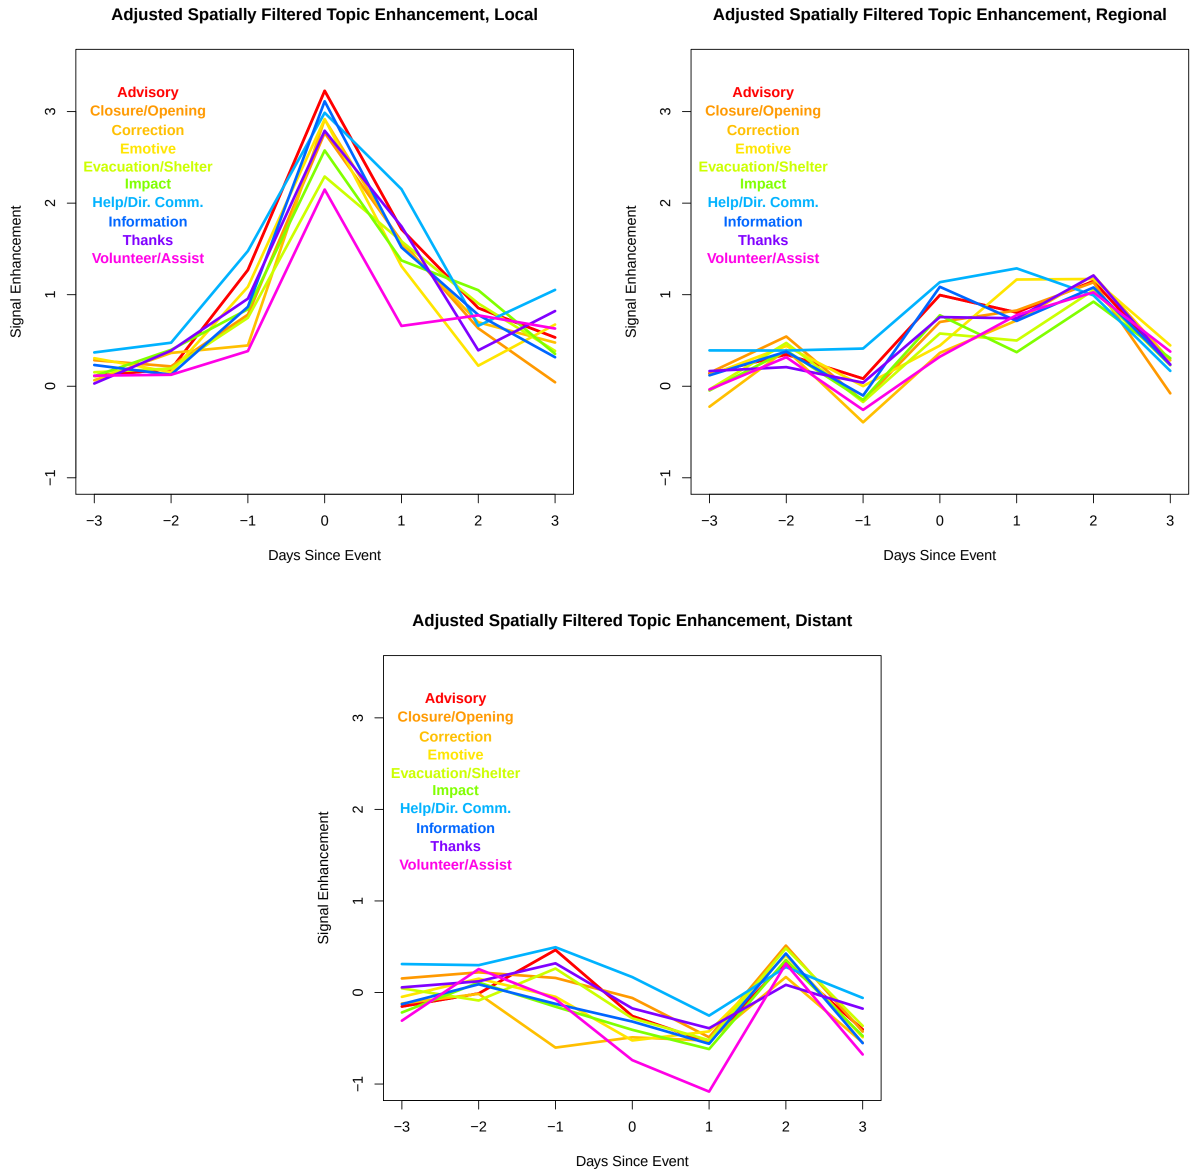

Figure 3.2: Adjusted spatially filtered topic enhancement averaged across all eight events and plotted over a seven-day span surrounding a disaster event. Each color represents a different category from the coded corpus of messages. Higher values indicate the observed messages have elevated similarity with that topic, net of both typical daily topic similarity in that region and global topic similarity on that specific day. The results show a local convergence of hazard-related topics surrounding the disaster, with no clear convergence on these topics in the regional and distant bins. Please refer to the Appendix for significance tests on the plotted values.

and alerts. For further validation of the results, I split the events into those that occurred with advance warning and those that did not. I plot the split results in Figure 3.3.

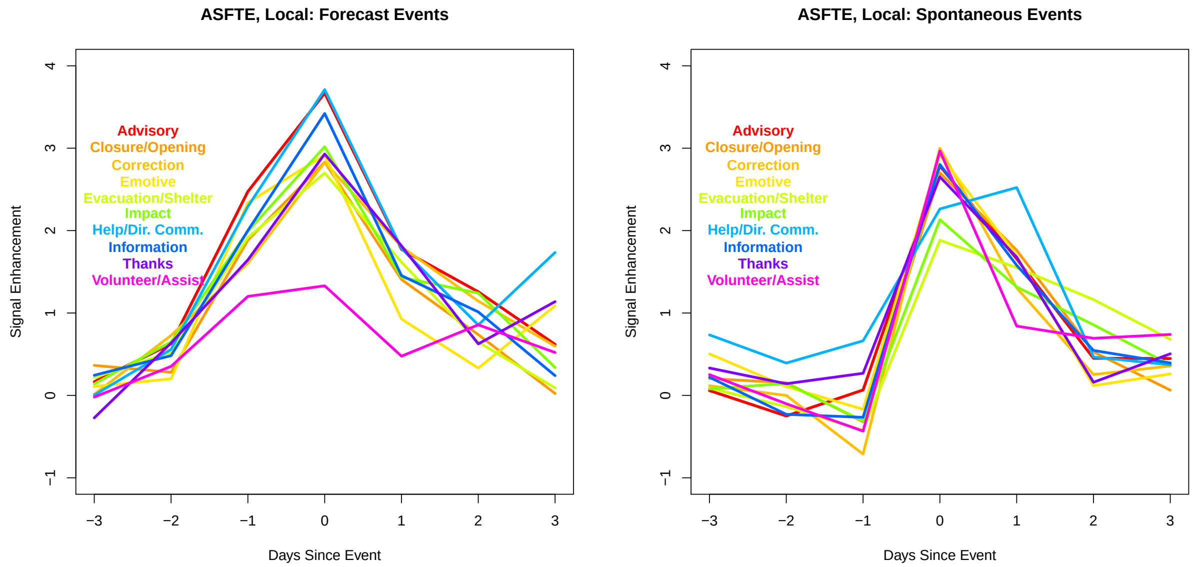

Figure 3.3: I separate local adjusted spatially filtered topic enhancement into events that are preceded by warnings (left) and those that occur without warning (right). The contrast between the two helps to validate that I am indeed picking up signal of hazard-related topics.

The results provide further face validity for the local signal I am picking up through spatial filtering. On the day prior to the event we observe no signal enhancement for hazard topics occurring without warning. By contrast, we observe an increase in topic signal prior to events that are preceded by warnings. We observe elevated signal of the coded topics on the day prior to the event, followed by a peak on the day of the event and a decline towards normality afterwards. In addition to helping validate this technique, this also supports that we can pick up signal of the Allport and Postman (1947) goal-gradient phenomenon, where populations are increasingly likely to rumor about an imminent event. We observe topic coalescence in anticipation of the event, as the content of communication begins to converge on hazard-related topics.

The results indicate that the spatio-temporal filtering approach can distinguish between local, regional, and distant topic signals across a variety rumoring events. The surge of similarity to the coded corpus in the local bin suggests a convergence on emergency management-related topics common to the coded corpus. That we observe no such convergence in the regional and distant bins indicates a contrast between the content of messages observed in the local area and the regional/distant bins. Because we cannot readily distinguish the signals of individual topics, we have no evidence of *evolution* of rumoring topic content over the course of disaster events. To evaluate whether these results are indeed indicative of a lack of content evolution or whether this is an artifact of the corpus of coded messages, I turn to the Moore case study for further investigation of rumoring content.

3.3.1 Moore Case Study

In both the aggregate case and in Moore (as demonstrated in Figure 3.4) I find evidence of local convergence on some rumoring topic (or set of topics), but I cannot readily distinguish among the individual topic signals. Because the coded topics seem to represent a general signal of some emergency management dimension of discussion, I cannot identify any evidence of evolution of content from topic to topic over the course of the event(s). To remedy this, I examine rumoring during the Moore event in greater depth to determine if I can find any evidence of content evolution. This investigation will offer further validation (or perhaps cast doubt) on the patterns of topical convergence we observed during the rumoring process.

As illustrated in Figure 3.4 the adjusted spatially filtered topic enhancement in Moore's local, regional, and distant bins looks qualitatively similar to the adjusted spatially filtered topic enhancement averaged across all eight events. We observe local topic convergence in anticipation of the event, some evidence of delayed regional topic convergence following the event (which is slightly more pronounced than the aggregate regional signal), and no notable signal enhancement in the distant bin. The similarities between the Moore event and the aggregate pattern indicate that it is not wholly inappropriate to suggest that the findings of the Moore case may generalize to other events. The analyses of the Moore case study serve two purposes: 1) further investigation of patterns of topic convergence and 2) identification

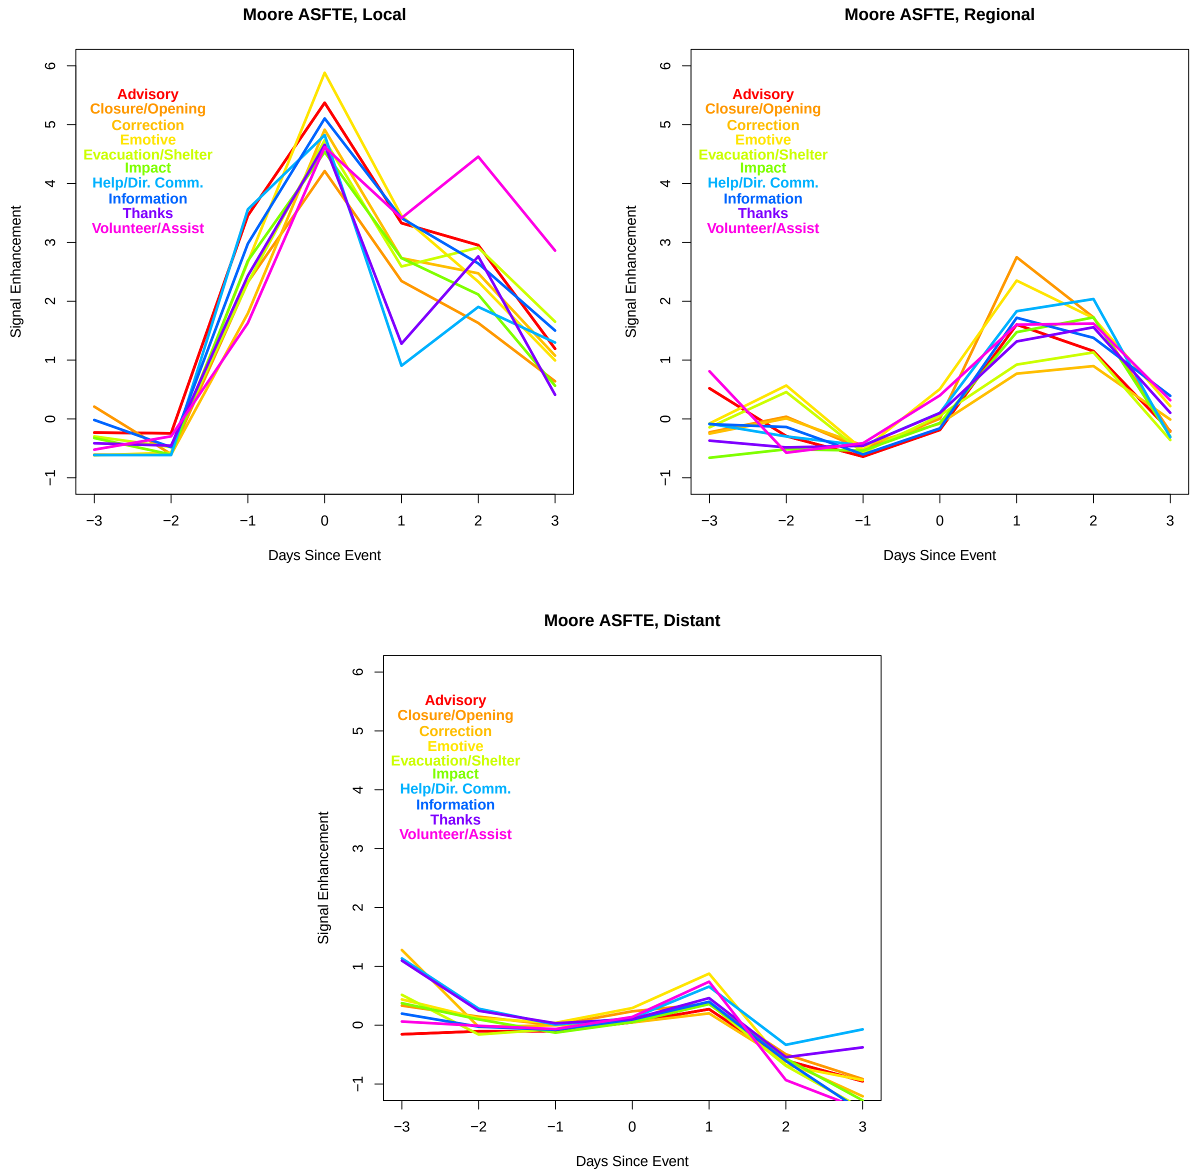

Figure 3.4: Adjusted spatially filtered topic enhancement during the Moore tornado. The pattern bears a strong resemblance to the average ASFTE.

of the presence or absence of topic evolution over the course of the event.

3.3.2 Interregional Similarity

I begin by investigating interregional topic convergence by measuring the extent to which different regions discuss the same topics. For each of the seven days surrounding the Moore tornado I calculated the cosine similarities among the sets of all "tornado" messages originating from each region. We observe little similarity between regions three days prior to the event and two days prior, when most tornado-related messages were a mixture of off-topic and on-topic. On the day before the Moore tornado, watches were issued throughout Oklahoma, Kansas, Nebraska, Iowa, and Missouri. On this day we observe a large increase in cosine similarities, suggesting a coalescence of topics common to all three regions. We observe the strongest overlap in topics between the regional and distant bins. Although elevated above typical levels, the similarity between local topics and regional/distant topics is substantially lower than the similarity between regional and distant bins. This suggests a division of topics between the local bin and the regional/distant bins.

This distinction between local topics and regional/distant topics persists on the day of the Moore tornado. Again, the similarity between regional and distant bins is much higher than the similarity either bin has with the local region. On the day following the Moore tornado the distinction declines somewhat. While the local-regional similarity remains quite low, the local-distant similarity becomes the highest interregional similarity we observe on that particular day. By the second day after the event topics diverge across regions and on the third day they return to their pre-event levels which indicate no notable similarity across the regions.

These results provide further evidence for the finding that hazard-related rumoring does not fully converge on a single topic across all regions. This was evident in the aggregate case because we found convergence on emergency management topics local in the local bins but no such convergence elsewhere. Although we did not observe convergence on emergency

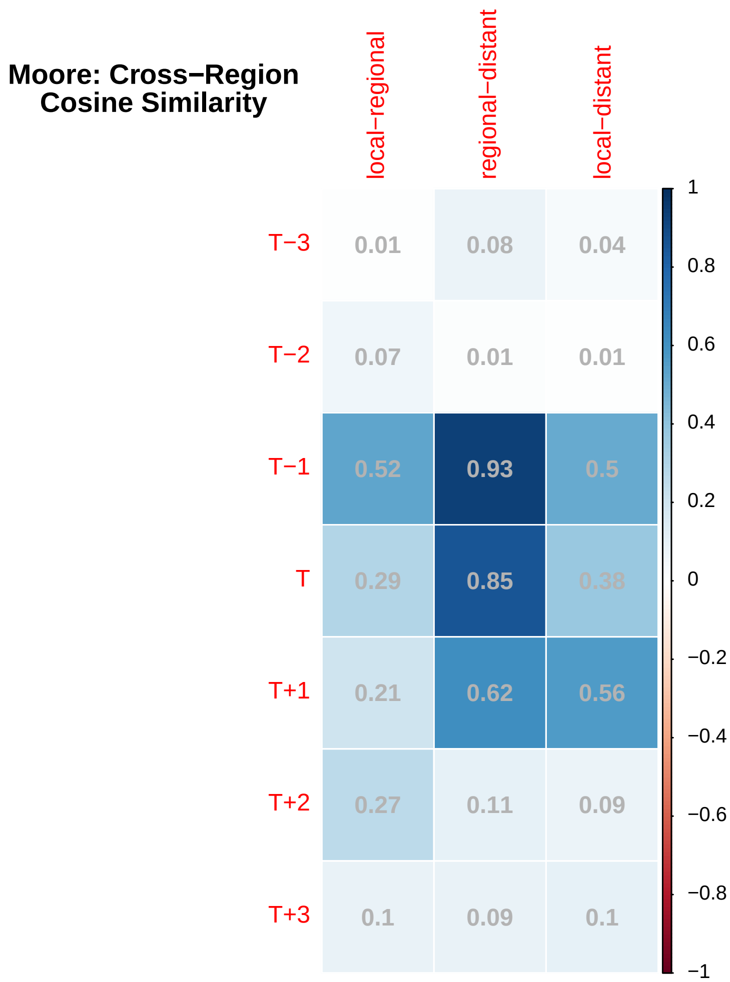

Figure 3.5: Over the seven-day span surrounding the Moore tornado I represent the cosine similarity between each region's set of "tornado" messages. Darker blue areas represent stronger cosine similarities, while lighter areas indicate less similarity. The results highlight an overall trend of topic coalescence on the three-day period surrounding the Moore tornado, but not all locations appear to converge on the same topics.

management topics in the regional or distant bins, the results here suggest that they may instead have converged on other topics, as suggested by their remarkably high cosine similarities. Although we do not find strong evidence of an overall topic convergence, the notable increase in baseline interregional cosine similarity suggests that there is some minor to moderate overlap across all regions. They may not necessarily attend to the same specific subject matter, but they all appear to be discussing a common topic.

3.3.3 Intraregional Similarity

I next compare sets of "tornado" messages within each bin to compare intraregional similarity to interregional similarity. By identifying whether similarity was stronger within or across regions, we can more accurately characterize the convergence phenomenon. I calculated the cosine similarity between all days' "tornado" messages within each region and plotted the results in Figure 3.6. In each bin we observe elevated autocorrelation around the time of the Moore tornado. In fact, all three bins exhibit the highest cosine similarities between the pairing of days T-1 and T and the pairing of T and T+1. This not a clear continuation of topics across the entire time period, however, as in all three cases the similarity between T-1 and T+1 is substantially lower than the aforementioned similarities. Instead this suggests some degree of evolution in topic from day to day. Although the results indicate that there is intraregional topical convergence during the period of time surrounding the Moore tornado, we also have evidence that *what* these populations converge on appears to shift over time.

A comparison between the intraregional and interregional similarities helps solidify a distinction between local topic convergence and regional/distant topic convergence. On T-1 the local topics are more similar to the local messages at T than they are to either regional or distant topics on T-1. Similarly, local topics on T are more similar to local topics on T-1 and T+1 than they are to either the regional or distant bins on that day. That intraregional similarity is stronger in the local bin than *inter*regional similarity suggests further evidence of local topic coalescence. On the day after the event, however, local topics are more similar to distant topics on T+1 than they are to local topics on either T-1 or T. An examination of the most popular bigrams in each region over T-1, T, and T+1 reveals that this may be part of a mass convergence phenomenon.

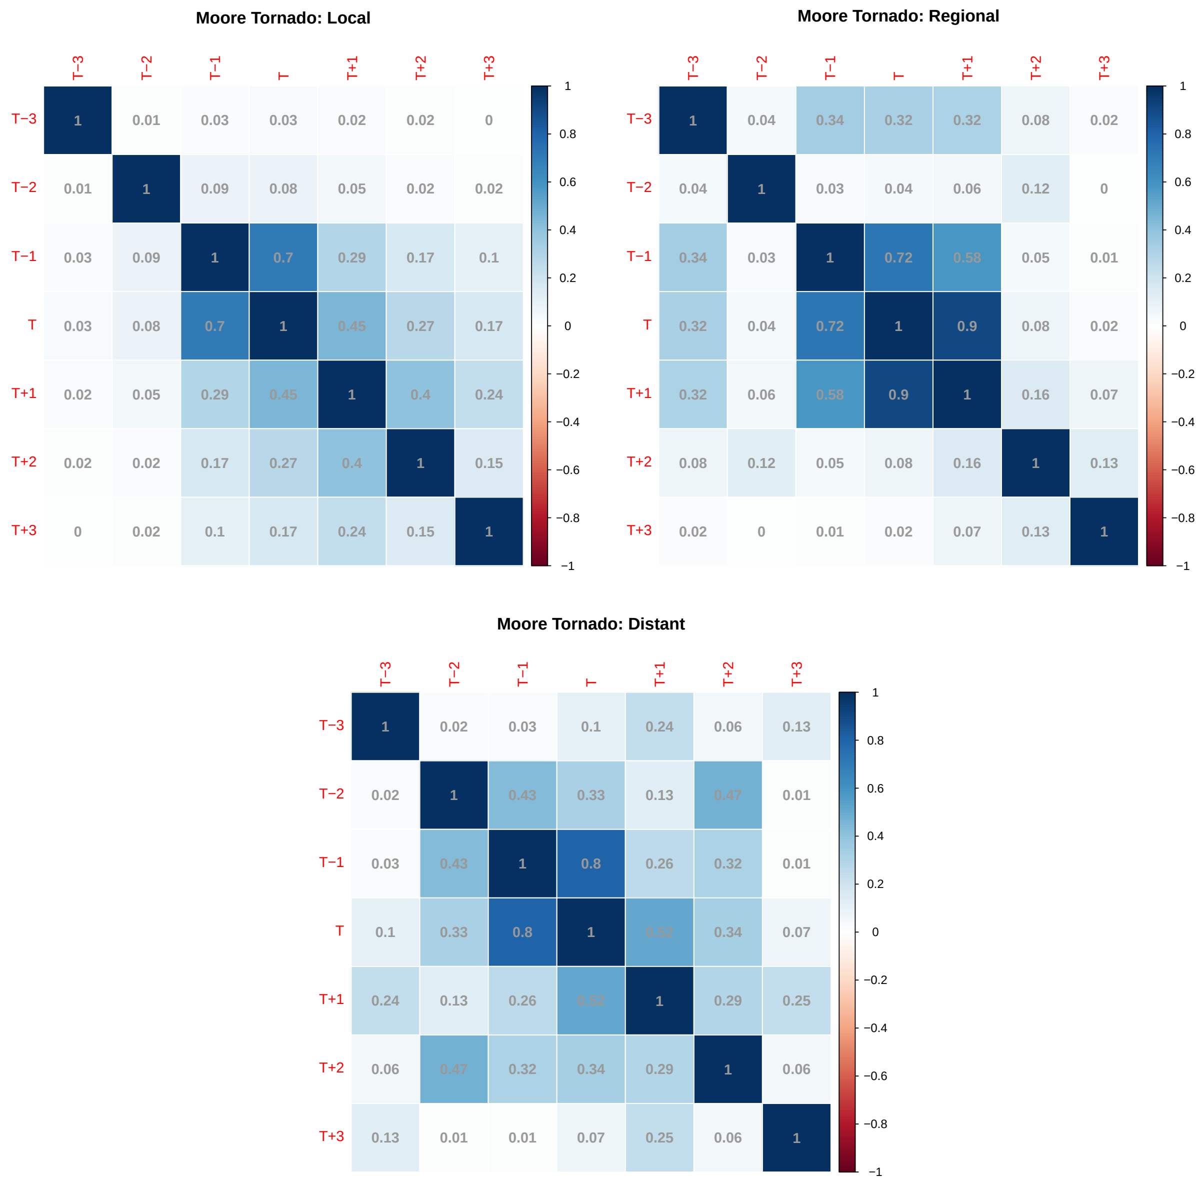

Figure 3.6: For each region I plot the cosine similarity between each day's "tornado" messages for the seven-day period surrounding the Moore tornado. Each region shows moderate to strong autocorrelation during the three-day period before, during, and after the Moore tornado.

3.3.4 Moore Bigrams

To provide further insight into the patterns of similarity we observe, I enumerate the fifteen most common bigrams in each region on the day prior to, day of, and day after the Moore tornado. In the following tables I rank them based on frequency of appearance, beginning with local bigrams in Table 3.4.

Table 3.4: 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)

On the day prior to the Moore tornado the most popular bigrams in the local bin primarily refer 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 one preceding 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 following day we continue to see a variety of watch and warning messages, but several of the most popular bigrams describe tornadoes 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). The prominence of “this tornado” and “in ok [Oklahoma]” provide further evidence that individuals are discussing a specific event that has occurred in a defined location (Oklahoma). On the day after the event 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 we see a transition in content of the most popular bigrams in the local bin. From warnings to a realized event to recovery, these messages display a notable shift in content. This is consistent with the transformation in content we observe with the cosine similarities between local bins. Although the underlying topic of tornado events remains constant, the local population converges on different aspects of that topic over time (warnings, event, recovery).

Table 3.5: 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)

Within the regional bigram we observed the highest levels of topic autocorrelation, as the similarities between T-1/T and T/T+1 were .72 and .90, respectively (compared to .7 and .45 in the local bin and .8 and .52 in the distant bin). The consistency of bigrams from day to day reflects this elevated autocorrelation. During all three days the most frequent bigrams 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 following four bigrams are identical across all three dates and 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 Midwest and into the Mississippi River Valley.

Table 3.6: 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 1000pm (188)	the tornado (260)	until may (94)

The inter-regional cosine similarities suggest strong convergence between the regional and distant bins on the day prior to and day of the Moore tornado and the prevalence of watch-related bigrams during those days is consistent with those results. On the day after the Moore tornado, however, the distant topics are almost equally similar to regional and local message content. The top bigrams on that day offer insight into how the distant bin resembles both the local and regional bins on that day. The top bigrams refer specifically to a tornado event in Oklahoma (almost certainly 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.

3.4 Discussion

By relating the coded corpus of messages to spatio-temporally filtered bins of messages I find strong and consistent evidence of local topical convergence on the day of hazard events. Events preceded by warnings and alerts show local topical convergence in anticipation of the events and all event types show, on average, a reversion to baseline similarity levels with the coded corpus two days after the event. These results support theories of local topic coalescence during periods of elevated rumoring. We find little evidence of coalescence onto the coded corpus in regional and distant bins, suggesting that all populations do not converge on a single topic. Instead, populations outside the impacted communities may converge on distinct topics, as observed in the Moore tornado. While the Moore results suggest that across all locations there is a heightened level of similarity, intra-group similarity tends to be higher than intergroup similarity in local and non-local (i.e., regional/distant) populations.

Local rumoring activity showed similarity with all types of coded official messages, which prevented us from identifying in the aggregate case whether rumoring is an evolutionary process. An examination of intraregional and interregional similarity in Moore suggested that there is some degree of evolution among messages in all bins. Although the results may not necessarily generalize to all events, I do find some initial support that rumoring is a process marked by topical evolution over time.

3.5 Conclusions

The pattern of the results bears semblance to rumor theories and this offers support of the spatio-temporal filtering approach for topic enhancement during rumors. We clearly identified types of content that were specific to ground zero of a wide variety of disaster events. The contrast against the signal (or lack thereof) observed in regional and distant bins suggests that this spatio-temporal filtering approach can be tuned to recognize distinct topical streams of rumoring. The case study results refine our understanding of rumor across large spatial scales by suggesting that it is an evolutionary process that retains location-specific topical foci. Given the speed at which information flows and develops in the online environment, spatio-temporal filtering at the daily scale is a rather blunt approach towards identifying topical signals in rumoring. That makes it more encouraging that I could identify topical signals with this approach. A more refined spatio-temporal filtering approach utilizing much thinner time slices to analyze a variety of events could shed further light onto the rough outline of the evolutionary processes observed in the Moore case.

Chapter 4

Spatial Excitation: Testing a Network Activation Theory of Rumor Transmission

Interpersonal ties play an important role in the diffusion of information during disasters. During disaster events, immediately available alters such as coworkers, neighbors, friends, and family members (Erickson et al., 1978; Richardson et al., 1979; Scanlon, 2007) become likely candidates for information exchange (Al-Makaty et al., 1994; Dodd, 1958; Miller, 1992). I broadly refer to the notion that individuals impacted by or learning of a disaster event primarily respond by activating existing ties in their personal networks as conduits for information exchange (rather than e.g. forming new ties, or relying on formal communication channels) as the *network activation theory* of information transmission in disasters. In this chapter I test the network activation theory and several other theories of rumoring to determine what accounts for county-level rumoring activity in the aftermath of severe tornado events.

It is widely recognized that the likelihood of forming an interpersonal relationship such as friendship or marriage declines with distance (Bossard, 1932; Brakman et al., 1999; Festinger et al., 1950; Freeman et al., 1988). It follows, therefore, that interpersonal relationships tend to be geographically constrained and, accordingly, information spread through those ties will likewise be spatially constrained (Caplow, 1947; Dodd, 1953). To illustrate this concept, I demonstrate in Figure 4.1 the relationship between rumoring activity and interpersonal tie volumes in the aftermath of a severe tornado event. I illustrate county-level counts of messages on Twitter containing the term "tornado" during the 24 hours following the touchdown of the Moore tornado, an EF5 tornado that struck on May 20th, 2013 and killed 24, injured 377, and caused $2 billion in damage. Using a spatial interaction function to simulate networks across geographic space, I estimate the number of interpersonal ties from Cleveland County (highlighted red), site of the Moore tornado, to the rest of the state. I superimpose that network on a map of tornado-related rumoring activity to illustrate the relationship between tie volumes and rumoring activity. The figure shows a strong relationship between tie volumes to Cleveland County and county-level tornado-related rumoring activity in the aftermath of the Moore tornado.

In addition to constraints on rumoring activity imposed by the spatial distribution of ties (i.e. potential rumor pathways), salience also plays an important role in constraining the spread of rumoring activity. Rumoring is not a ubiquitous response to disrupted environments; it typically occurs in a bounded slice of space and time. A rumor does not propagate endlessly, which is one factor separating rumors from urban legends and folk tales (Miller, 1992). 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; Caplow, 1947). However, topics salient to one population may not necessarily be 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 salient to a population (Caplow, 1947). Like the phenomenon of rumoring, salience is a concept that is easily understood yet remains very

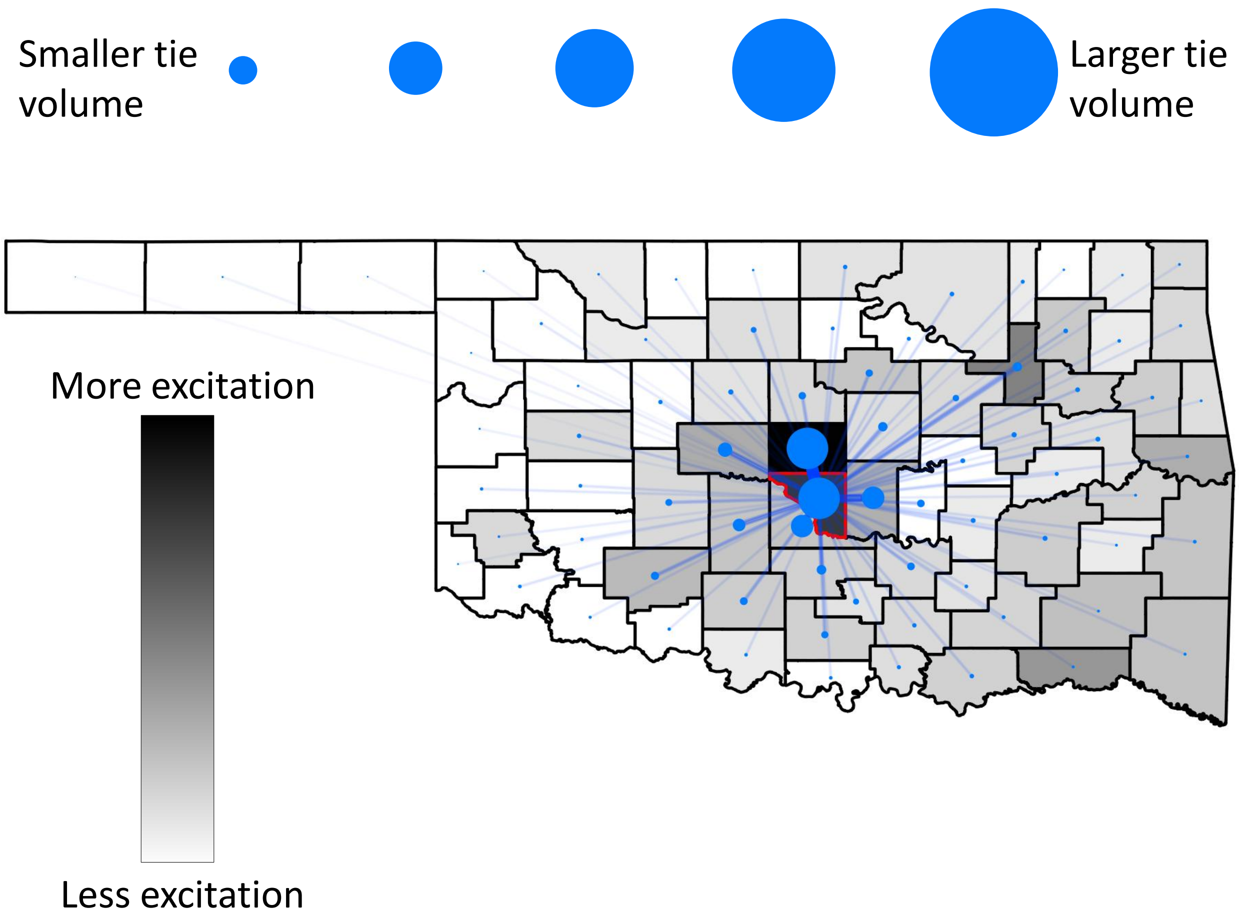

Figure 4.1: Volumes of interpersonal ties between Cleveland County, Oklahoma and all other counties in Oklahoma are projected on a map in which counties are shaded according to the level of tornado-related online rumoring activity in response to the Moore tornado. Counties depicted with larger nodes have greater volumes of ties to Cleveland County. I find that counties with more ties to Cleveland County tend to have greater levels of tornado-related rumoring activity in response to the Moore tornado. The overlap between the spatial distributions of tie volumes and rumoring activity is consistent with the network activation theory of rumoring.

difficult to measure reliably outside laboratory environments.

4.1 Overcoming Traditional Challenges to Measuring Rumoring

The online environment allows us to employ novel approaches for observing rumoring activity. Historically, studies of hazard-related rumoring have relied on a post-hoc approach where researchers interview informants about their rumoring activity *after* the disaster event has occurred. Accuracy issues emerge when asking informants to report on their own behavior (Back et al., 1950; Romney et al., 1986; Romney and Weller, 1984; Sudman et al., 1996), particularly in hazard contexts. Informants may, for example, mis-attribute sources of rumors or misremember when or where they heard a rumor (Scanlon, 2007). Additionally, defining sample populations for observational rumor studies is a challenge, as is interviewing anything larger than a modest population. These sampling issues are exacerbated by a disaster's propensity to displace populations that have left their homes to reach safety, find shelter, seek medical attention, or assist with disaster response (Scanlon, 2007). Furthermore, many studies outside the laboratory environment 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 such, we know less about situations under which rumors fail to emerge or diffuse. Instead of employing a post hoc approach I opt to observe rumoring activity continuously, which enables us to monitor activity (or lack thereof) before, during, and after a disaster event. Computer-mediated interactions provide a valuable opportunity for unobtrusively monitoring rumoring (Bordia and DiFonzo, 2004) and we have accordingly seen a shift in attention towards rumoring behaviors in online, informal communication (Maddock et al., 2015; Mendoza et al., 2010; Sutton et al., 2013a,b; Vieweg et al., 2008). With precise data on the timing, location, and content of rumors, the online environment bypasses many of these sampling and informant accuracy challenges that have plagued rumor studies for decades.

I turn to Twitter to capture rumoring behaviors during disaster contexts. Studies of Twitter usage during disasters indicate that online populations respond in systematic, consistent, measurable ways to disaster events. This is demonstrated by increased volume of event-specific keywords, changes in message structure, and changes in information retransmission behavior (Sutton et al., 2008, 2013a,b; Vieweg et al., 2010). This regular behavior and relevance to classical definitions of rumoring suggest the Twitter platform as a tool for understanding rumoring in the context of hazards.

Using Twitter's Streaming API, I monitor rumoring activity based on the posting of messages containing the term "tornado." Metadata accompanying each message indicates its exact timestamp and the poster's latitude and longitude coordinates if the message comes from a GPS-enabled device, such as a mobile phone or tablet. These metadata enable me to harness the spatial and temporal characteristics of rumoring activity to measure the phenomenon with precision historically unmatched in studies of rumoring. Although the spatio-temporal variation in rumor quantity and content has long been of interest to scholars, collecting data that measures both temporal *and* spatial characteristics of the phenomenon has been extraordinarily difficult to do with any degree of precision. Some have been able to capture temporal (Bordia and Rosnow, 1998; Danzig, 1958; Greenberg, 1964) or spatial (Larsen, 1954) aspects of rumoring, but bridging the two has been difficult. Recent studies have begun to incorporate more precise temporal (Blanford and MacEachren, 2014) or spatial (Starbird and Palen, 2010) measures of informal communication in response to disaster (and on rare occasion, both simultaneously (Guan and Chen, 2014)). I build on this prior work by employing such spatio-temporal measurement of communication activity to test competing theories regarding rumoring in disaster.

4.2 Results

To measure rumoring activity I collected geolocated, timestamped tornado-related messages during three major tornado events in the United States during 2013. These events represent all fatal tornadoes with an Enhanced Fujita (EF) score of 4 or 5 that occurred after May 17th, 2013 (when we began collecting geolocated messages through the new version of Twitter's Streaming API). The severity of these events ensures that I will observe substantial rumoring activity in response to the event, including rumoring across a large spatial scale. I observe activity during two stages of the response, which I call the *primary excitation* and *secondary excitation* phases. I define the primary excitation phase as the period from the moment the tornado touches down until one hour after it has dissipated, during which we typically observe an initial proliferation of rumoring and saturation of awareness in response to the event (Danzig, 1958; Erickson et al., 1978; Greenberg, 1964; Richardson et al., 1979). Secondary excitation covers the 24-hour period following the end of primary excitation, during which we typically observe a mass convergence of attention on the disaster site (Hughes and Palen, 2009; Sutton, 2010). Response during this latter period is often characterized by individuals sharing information about the event, expressing concern about those impacted, and/or attempting to increase awareness about how to help those affected. While primary excitation is typically described as a local phenomenon, mass convergence may occur on a nationwide scale (Sutton, 2010). To measure geolocated rumoring activity during each time period, I use the geospatial metadata from each "tornado" message to identify the U.S. county from which the message originated. These county-level counts demonstrate the extent to which each county engages in rumoring activity in response to the event and illustrate the spatial distribution of rumoring activity across the entire nation.

I use negative binomial regression to model counts of county-level "tornado" messages as a function of covariates related to salience and network activation. I also incorporate a number of controls (including population size, urbanicity, and lagged effects), as discussed below.

These models demonstrate how features of each individual county influence its propensity to engage in rumoring activity in response to each tornado event. I illustrate the model coefficients for primary excitation and secondary excitation in Figure 4.2 and Tables 4.1 and 4.2.

Because we observe county-level counts over three events, we observe each county three times in the model. Accordingly, the total N in our model is 9,411.

Table 4.1: Negative binomial regression of county-level “tornado” message counts during primary excitation

Primary Excitation	Coef	exp(Coef)	SE	p-value
Intercept	-12.1723	0.0000	0.6031	0.0000
Responding county: number of warnings	0.0044	1.0044	0.0019	0.0207
Responding county: number of events	-0.0271	0.9733	0.0151	0.0729
Responding county: impact	-0.0006	0.9994	0.0208	0.9773
Total migration: responding-event counties	-0.0004	0.9996	0.0004	0.3479
Event county: whiter population	0.0250	1.0253	0.0033	0.0000
Event county: poorer population	0.0142	1.0143	0.0066	0.0312
Responding county: rural index	-0.0268	0.9735	0.0356	0.4515
Log (base 2) tie volume	0.3480	1.4163	0.0126	0.0000
Responding county: log (base 2) population	0.5110	1.6669	0.0302	0.0000
N=9,411	AIC: 8,941			

Table 4.2: Negative binomial regression of county-level “tornado” message counts during secondary excitation

Secondary Excitation	Coef	exp(Coef)	SE	p-value
Intercept	-10.3653	0.0000	0.4563	0.0000
Responding county: number of warnings	0.0061	1.0062	0.0014	0.0207
Responding county: number of events	0.0077	1.0078	0.0115	0.0729
Responding county: impact	-0.0068	0.9932	0.0168	0.9773
Total migration: responding-event counties	-0.0006	.9994	0.0004	0.3479
Event county: whiter population	0.0148	1.0149	0.0023	0.0000
Event county: poorer population	0.0268	1.0263	0.0049	0.0312
Responding county: rural index	0.0396	1.0403	0.0271	0.4515
Log (base 2) tie volume	0.2021	1.2240	0.0108	0.0000
Responding county: log (base 2) population	0.5002	1.6491	0.0231	0.0000
Responding county: log (base 2) primary excitation	0.3481	1.4164	0.0274	0.0000
N=9,411	AIC: 14,136			

These models capture three categories of variables that may influence rumoring activity:

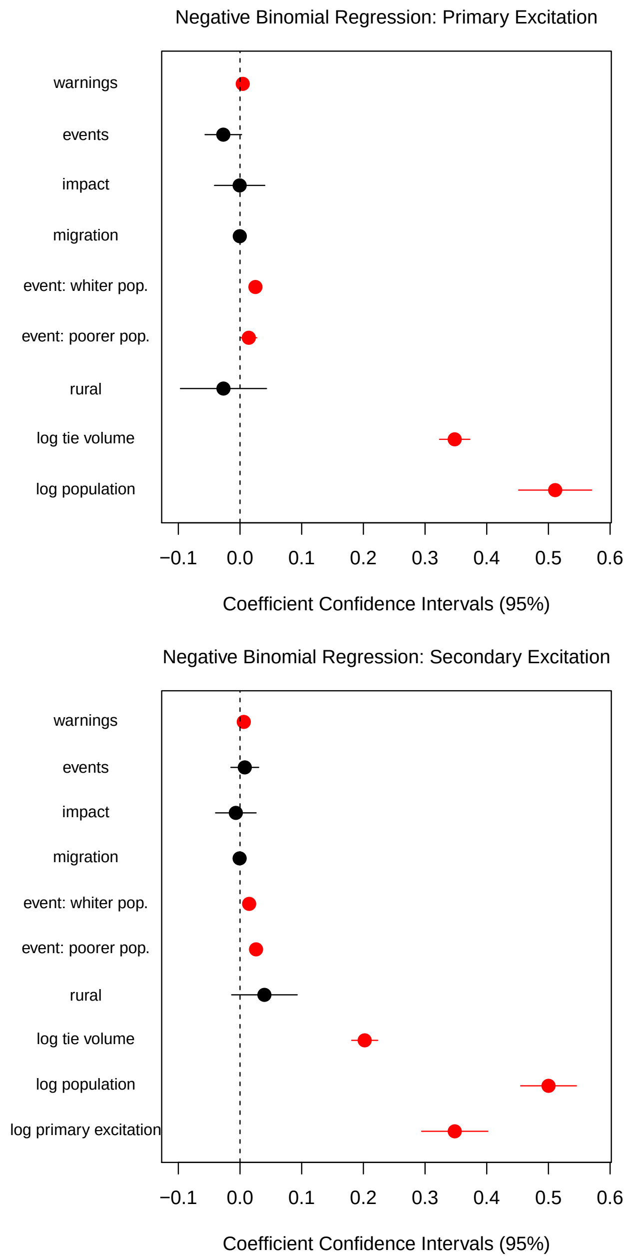

Figure 4.2: Negative binomal regression coefficients for county-level counts of "tornado" messages during periods of primary excitation (left), within 1 hour of the tornado event, and secondary excitation (right), the 24-hour period following primary excitation. Positive, statistically significant ($p < 0.05$) terms are red while non-significant terms are black. Line width for each coefficient represents its 95% confidence interval. I measure effects for event salience, demographic salience, and network activation on county-level, tornado-related rumoring activity.

salience of the event itself, salience of the population affected by the event, and interpersonal tie volumes between the rumoring and event-stricken populations. Table 4.3 lists the mean, median, and standard deviation for county-level excitation and a number of other predictors of excitation. I have broken these measures down to statistics that are general county measures and event-specific statistics. The county measures represent descriptive statistics for each individual county, such as number of warnings over a five-year period or population. The event-specific county measures refer to statistics in each county that may change across the three events under study, including volumes of messages during primary and secondary excitation and the number of ties to the tornado-affected county.

Table 4.3: County and event-level statistics

	Mean	Median	SD
County Measures			
Warnings	30.82	28.00	22.70
Warnings*	33.44	30.00	21.71
Events	1.87	1.00	2.54
Events*	2.91	2.00	2.65
Impact	0.00	-0.16	1.53
Pct. White Residents	17.23	16.30	6.60
Pct. Poverty Residents	82.91	89.11	16.85
Population	98,398.08	25,893.00	313,176.30
Event-specific County Measures			
Primary Excitation	0.43	0.00	3.26
Secondary Excitation	0.87	0.00	4.96
Avg. Individual's Ties to Tornado County	3.10	0.00	172.69

*Counties with 1 or more warning/event

In the following sections I will describe each of the model terms in much greater detail. I begin with the salience category.

4.2.1 Salience

Outside laboratory experiments, salience is very difficult to measure and manipulate in rumoring studies, as it is a mental state that is not directly observable. Relying on informant reports is infeasible, as informants would have great difficulty recalling what was and was not salient, particularly if the event was traumatic (Haber and Haber, 2000) or if they were unaware of the event. Instead of relying on informant reports of salience, I directly observe factors that are known to affect salience.

4.2.1.1 Event Salience

I measure two components of salience: salience of the tornado event itself and salience of the population affected by that event. The first three model terms in Figure 4.2—warnings, events, and impact— represent the salience of past tornado events. These terms measure the extent to which each county's exposure to past tornado events (potential or realized) influences its level of observed rumoring activity in response to severe tornado events affecting the United States. Where disaster events are made more salient by past exposure to events and warnings, subsequent events and warnings enhance public response and communication (Mileti and O'Brien, 1992). Those regularly impacted by tornado events and warnings, for example, may be more likely to discuss a severe tornado event, even if that event strikes a distant community. Likewise, those lightly affected or unaffected by a series of disaster events tend to have a normalization bias, which suppresses their likelihood to respond to future warnings and events (Mileti and O'Brien, 1992). To measure factors contributing to tornado event salience at the county level, I use a five-year measure of several tornado-related statistics. I collected these data from three agencies within the National Oceanic and Atmospheric Administration (NOAA): the Storm Prediction Center, National Climatic Data Center, and National Weather Service. From these agencies I collected data on numbers of tornado events, numbers of tornado warnings, and measures of impact from tornado events including injuries, fatalities, crop damage, and property damage. I matched those warning and event statistics to each U.S. county and used these county-level measures from the five-year period of 2010-2014 to measure tornado event salience. These measures indicate whether any given county's past experience with tornado events and warnings impact is propensity to engage in tornado-related rumoring during our EF4 and EF5 events under study.

In the model I preserve the measures of warning counts and event counts, but I consolidate the impact measures. To generate the "impact" measure I conducted a principal components analysis on the county-level measures of injuries, fatalities, crop damage, and property damage. I take the first principal component (which captures 58.3% of the variance) and assign a score to each county. (I note that the consolidated impact measure produced similar results in the model when compared to individual impact measures, whether modeled jointly or separately.) The results indicate that during neither primary excitation nor secondary excitation does the number of events or the impact of past events in any given county make that county more likely to engage in tornado-related rumoring during severe events in 2013. During both periods, however, I find a positive, statistically significant ($p < 0.05$) relationship between tornado-related rumoring and the number of warnings in a county. That is, counties with more warnings were more likely to generate increased counts of "tornado" messages in response to our major 2013 tornado events. However, the effects are rather weak in both models. An increase of 10 tornado warnings (the mean is approximately 30 over a 5-year period) in any given county is associated with a 4.5% increase in "tornado" messages during primary excitation and a 6.3% increase during secondary excitation. The weakness of these effects coupled with the lack of significance of our other terms suggests that event-related salience does not have a major effect on rumoring activity across large geographic scales.

4.2.1.2 Demographic Salience

In addition to event-related salience, I also examine social dimensions of salience. The characteristics of the population *affected* by a disaster influence the amount of attention paid to the event by other populations. However, past studies do not converge on one single explanation for how the social dimensions of salience operate. Some claim that events are more salient to a population when that event affects a demographically similar population (Rogers, 2000). By contrast, social status theories posit that disasters affecting wealthier communities receive disproportionate attention (Chomsky, 1998; Fothergill et al., 1999), even when several communities with varying economic backgrounds are impacted by the same event (Rovai and Rodrigue, 1998). Vulnerability theories argue the opposite and suggest that greater attention is paid to dramatic, exceptional circumstances in disaster, which is frequently its impact on vulnerable and fragile communities (Cutter et al., 2003; Tierney et al., 2006). I evaluate all of these theories by relating county-level rumoring activity to the differences between demographics of the tornado-affected county and the responding counties.

I use several demographic measures to determine whether salience is driven by characteristics of the population affected by a tornado event. To measure the demographics of each U.S. county, I use racial demographic data from the U.S. Census (Almquist, 2010) and poverty data from the U.S. Census' Model-based Small Area Income and Poverty Estimates. I match these data to each U.S. county and compare the demographic data of the county where the tornado struck to all other U.S. counties. I calculate the signed differences between the percentage of White residents and the percentage of residents in poverty of the event county and all other counties. Together these demographic measures help to indicate whether the salience of tornado events is based on demographic similarity between the affected and responding county, or whether attention is disproportionately paid to communities that differ demographically from the activated county.

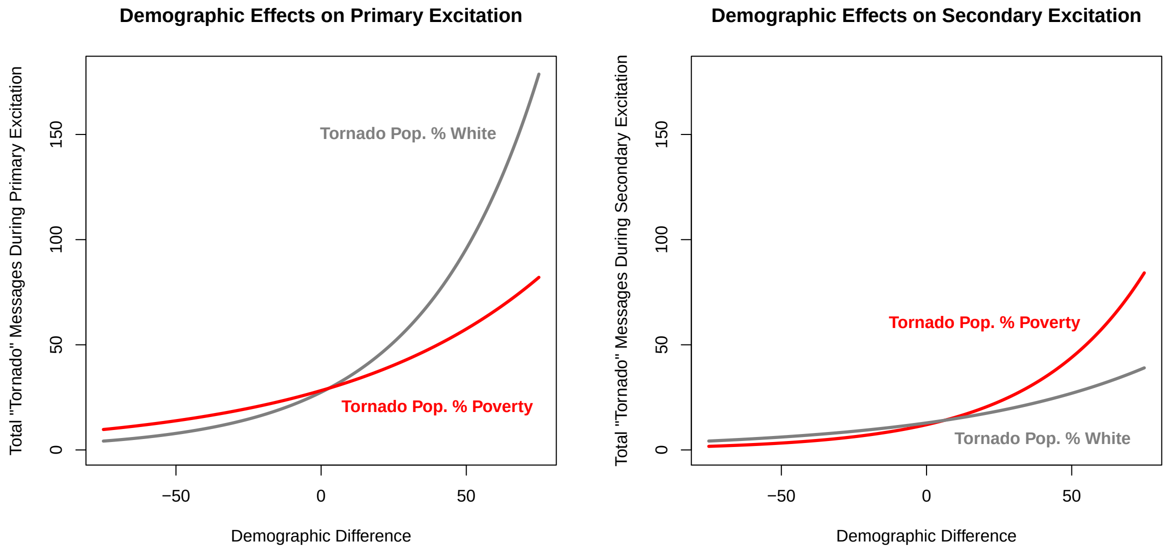

Figure 4.3: Projected activation based on the demographic differences between the tornado event county and the activated county. Event counties with larger absolute percentages of white residents and residents in poverty than activated counties generate greater activation during both primary and secondary excitation. The plotted projections are based on a county with average event salience measures, average population size, and an average of five ties per person to the tornado event county. While the effect for Whiter populations is stronger than the effect for poorer populations during primary excitation, the strength of the results is flipped during secondary excitation.

As I demonstrate in Figure 4.3 the findings for both primary excitation and secondary excitation indicate that counties have a positive, significant propensity to rumor about tornado events that impact a population that is *Whiter* than the activated county or a population that is *poorer* than the activated county. These findings offer mixed support for both social status and vulnerability theories, as each theory is supported by one effect and contradicted by another. The evidence clearly disagrees with theories that salience is driven by demographic *similarity* between the event population and the activated population. A variety of experiments demonstrate that sympathy and assistance in the aftermath of disasters are mediated by perceptions of whether the affected population is deserving or worthy of assistance (Batson et al., 2005; Cheung and Chan, 2000; Zagefka et al., 2012). One possible explanation for these mixed findings is that poorer, whiter communities are perceived as "worthy victims" who reside at the intersection between positive attention and a position conducive to sympathy (Herman and Chomsky, 1988).

I also examine the role of inter-county migration on salience. One may expect that higher migration flows between counties would increase the attention paid to an event striking one of those counties. To measure the role of migration, I collected county-to-county migration data from the American Community Study. I use the sum of counts of inbound and outbound migration for each event county to and from all other U.S. counties. In neither model do I find a significant effect for migration, which suggests that salience vis a vis migration patterns does not influence rumoring activity.

4.2.2 Network Activation

Finally, I test the network activation theory of information diffusion. If rumors spread through interpersonal ties, then rumoring activity in any given county ought to be proportional to the volume of ties it has to the tornado-affected county. Because I cannot feasibly measure tie volumes of all residents in all U.S. counties, I instead follow the Butts et al. (2012) approach of using simulated tie volumes. I employ a spatial interaction function (SIF) to govern tie formation across space, where the marginal tie probability between two individuals declines with distance, a common feature of social networks (Bossard, 1932; Hägerstrand, 1966; Latané et al., 1995; Zipf, 1949). Spatial structure is sufficient to account for most network structure at large geographic scales (Butts, 2003) and network models based on SIFs have been shown to be predictive of social phenomena such as neighborhood crime rates (Hipp et al., 2013), regional identification (Almquist and Butts, 2014), and interorganizational collaboration Butts and Acton (2011). An SIF $F_d(d, \Psi)$ defines the marginal probability of a tie between two individuals i and j as a function of some real parameter vector Ψ and the physical distance d between i and j. In most social networks, the marginal tie probability declines with distance. Beyond the individual case, I can employ SIFs to estimate tie volumes between arbitrary geographic units such as Census tracts, cities, or countries. In the present case I use U.S. counties as our geographical unit. Holding distances constant, I expect to observe greater tie volumes between more populated geographic units. Given two point locations i and j the expected tie volume between would then be $P_i P_j F_d(d_{ij}, \Psi)$, where P_i and P_j are the populations of the regions in question, and d_{ij} the distance between them. To model tie volumes over extended areas, I employ the expectation $\mathbf{E}\mathcal{V}(A, A') = \int_A \int_A' \frac{P_A}{|A|} \frac{P_{A'}}{|A'|} F_d(D(v, v'), \Psi) dv dv'$ where $\mathcal{V}$ is the interregional tie volume, A and A' are the areas in question, v and v' are coordinates, P is population, $|A|$ is land area, and D is the geodesic distance; this is halved for ties from A to itself. I employ an SIF that approximates the spatial distribution of "technologically mediated communication" (Butts and Carley, 2002; Hägerstrand, 1966). This is an attenuated power law distribution with a slowly decaying distance function that declines at $d^{-2.96}$. Based on this SIF I use the `networkSpatial` package in the R statistical computing environment (Butts and Almquist, 2013) to estimate tie volumes by numerical quadrature.

For each tornado event, I use the SIF to estimate the county-to-county tie volumes between the county affected by the tornado and all U.S. counties (including reflexive tie volumes within the affected county). To reduce computational complexity I assume a uniform spatial distribution of population within each county. I improve the accuracy of our estimation of reflexive ties within the tornado-affected county by taking the sum of tie volumes between all tracts within the county. With a higher spatial resolution, tract-level tie volumes provide a more precise estimate at an acceptable level of added computational complexity.

For the tie volume coefficients I find strong, positive, significant results ($p < 0.001$) for both primary excitation and secondary excitation. Net of other factors, doubling of the tie volume between a county with a tornado event and any other county is associated with a 42% increase in the count of "tornado" messages during primary excitation and a 22% increase during secondary excitation. I illustate projected tie volume effects in Figure 4.4.

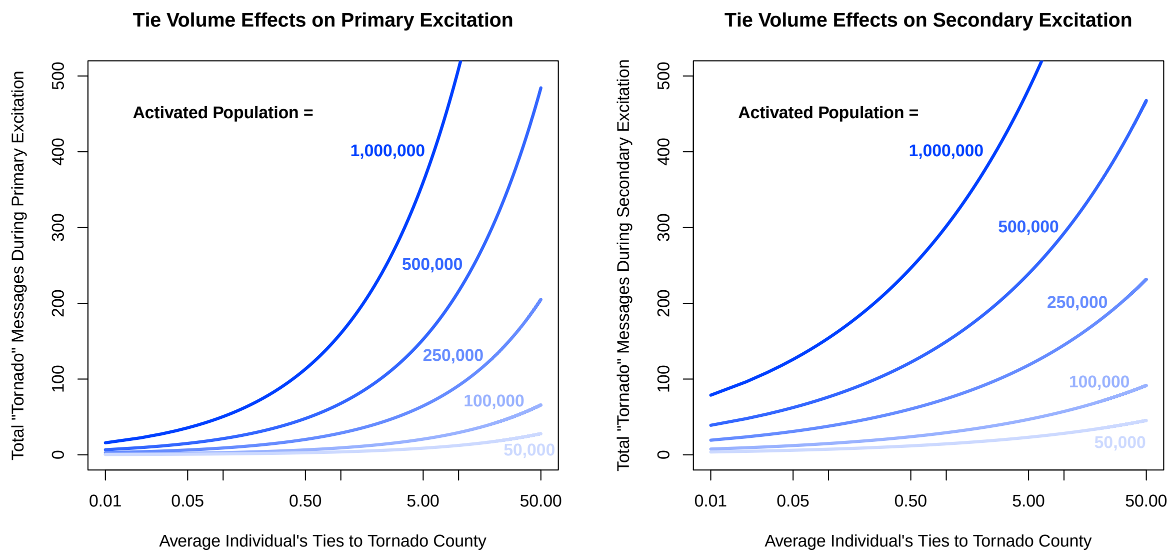

Figure 4.4: Projected activation as a function of tie volume and population size. Each curve represents a different population size. We find that as each individual's average number of ties to the tornado event county increases, the activated county increases its estimated count of tornado-related messages. The projections are based on a county with average event salience measures and average demographic salience measures.

These results provide strong support for the network activation theory. Where we have large volumes of ties between counties, I find increased activity in response to a tornado event striking one of those counties.

In addition to the model terms for event salience, demographic salience, and tie volume, I include terms to help control for other exogenous drivers of rumoring activity. Users of Twitter are more likely to live in urban environments (Duggan and Brenner, 2013) and I account for this with a term that measures each county on a rural-urban index. The Centers for Disease Control and Prevention's National Center for Health Statistics developed a six-point scale to identify where each U.S. county lies on the urban-rural spectrum. I use this index in the model to account for overrepresentation of urban communities on Twitter. In the full models the results indicate no evidence that rural counties are less likely to engage in tornado-related rumoring in the aftermath of an event. I do, however, find *marginal* effects for urban status on the volume of tweets; urban counties have higher marginal counts of "tornado" messages during both primary and secondary excitation. I also control for the population size of each county, as locations with large populations are much more likely to engage in rumoring activity in the aftermath of a disaster (Buckner, 1965; Richardson et al., 1979). I find strong effects for population during periods of primary excitation and secondary excitation. Doubling the population of the activated county produces an estimated 66.7% increase in primary excitation and 64.9% increase during secondary excitation. Finally, I also include a term for primary excitation in our secondary excitation model. Although the mass convergence phenomenon is well documented (Hughes and Palen, 2009; Sutton, 2010), we know little about the mechanisms driving it. The primary excitation term is strong, positive, and significant, which suggests that secondary excitation (during which we typically observe mass convergence) is carried by the momentum of primary excitation. Locations where we observe primary excitation tend to continue generating excitation during the secondary excitation period. I also note that the significant effects are quite similar in both models, which suggests that secondary excitation—and the mass convergence phenomena that encompasses it—is a continuation of primary excitation rather than a separate phenomenon.

4.3 Conclusions

The results build on classic studies of rumor by scaling rumor theory from small group and small town settings to a spatially heterogeneous population of hundreds of millions. At this spatial scale, however, I find limited evidence that rumoring activity is driven by salience of past tornado events. Instead I find that demographic salience plays a role, although the implications are mixed. I find that activated counties have increased response to events impacting Whiter counties or poorer counties. Separately, each of these findings supports and contests social status theory and vulnerability theory. I find strong and consistent support for network activation theory, as we observe greater activity in counties that have many ties to the event county. This suggests that across large spatial scales, interpersonal ties play an important role in rumoring, either as conduits of information transmission or as channels for attention. As the first paper of its kind to test network activation theory at scale, this chapter sheds new insight into information diffusion processes across large, spatially heterogeneous populations.

This approach uses online informal communication to measure county-level, tornado-related rumoring activity in response to severe tornado events. A novel source for measuring rumoring activity, the online environment allows us to measure the phenomenon with a combination of scale and precision that have historically been infeasible. I study this post-disaster rumoring phenomenon at a fairly high spatial and temporal resolution with events that likewise impact a very narrow slice of time and space. Accordingly, the findings may not necessarily represent the pattern of communication during disasters that unfold across larger spatial and temporal scales. I expect that future work will expand this approach to examine spatio-temporal characteristics of rumoring during events that develop more slowly, such as hurricanes, blizzards, and wildfires, and events that impact larger areas such as earthquakes. During such events the pace of rumoring may change, however it may be driven by the same mechanisms we observe in this paper. These results are a first step towards extending classical rumor theories to large geographic scales and at this scale I find strong support for network activation theories of rumoring activity.

Chapter 5

Concluding Matters

In this final chapter of the dissertation I review the key findings across all three chapters, recognize limitations of my analyses, and look ahead to future research that builds off these findings.

5.1 Key Findings

- **Spatio-temporal filtering works!** Even though geolocated messages represent a small fraction (3-4%) of Twitter's stream, data sparsity is not major a problem for detecting and analyzing rumoring activity in a spatially and temporally resolved manner. I am 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 I observed with a *single* keyword. Despite data sparsity, I did not have to develop techniques to identify joint signal across multiple keywords in order to detect rumoring activity 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 I was regularly able to detect rumoring activity *during* events. Although disaster sites may feature power outages, congested cellular networks, population displacement due to evacuation, and individuals' preoccupation with preserving personal safety over posting online, I consistently observed hazard-related communication during a wide variety of events. This is a promising finding for the future of online, informal communication analyses, especially for those requiring spatial metadata.

- **Distance is not dead:** Although digital media have transformed communication and commerce by connecting far-reaching segments of the globe over the last two decades, distance still appears to play an important role in online responses to disaster events. The largest surge in signal of informal communicative activity following a disaster occurs at the site of that disaster. As we become further removed from the event, signal of communicative activity declines. While severe events are followed by a mass convergence of attention at great distances (thousands of miles), that mass convergence does not exceed the strongest local signal. Furthermore, I find evidence that the topics discussed around ground zero differ from those further away, suggesting different informational needs based on a population's distance from the event epicenter.

- **Warnings matter:** Events preceded by warnings show anticipatory excitation, as indicated by a surge in signal of local communicative activity and local, topical convergence on emergency management-related topics. Where warnings are issued ahead of events we observe a notable increase in hazard-related communication. While this does not necessarily imply that individuals take preventive action taken to protect themselves from the event, this finding suggests that warnings can spur social action as far as a day ahead of a disaster event. This has important policy implications for the efficacy of warnings and alerts.

- **Network activation theory:** I develop a network activation theory of information exchange which suggests that during disaster events individuals activate existing ties in their personal networks (rather than form new ones) as conduits for information exchange. I find a strong effect for tie volume between the affected county and the activated county on the volume of messages coming from the activated county. This effect holds, net of features such as population size of the activated county, which suggests strong support for network activation theory. Furthermore, the wide geographic scope of analysis suggests that network activation theory accounts for rumoring activity at national scales.

5.2 Limitations

- **Twitter idiosyncrasies:** 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. The results of this dissertation may overrepresent rumoring activity from minority, urban, or young populations. It is not immediately clear, however, how this would bias the results, as the 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.

- **Content evolution?** Unable to distinguish among the coded content categories in the topic filtering chapter, I could not reliably demonstrate clear evolution of content across all cases. The Moore case study provides some evidence of content evolution in each bin according to intraregional cosine similarities. Additionally, changes in the top bigrams suggest content evolution. However these results are limited to a single case. While the results suggest some initial evidence of content evolution, further analysis is necessary to obtain a stronger understanding of the evolutionary process.

- **Network activation:** Although I test a network activation theory of information transmission, I do not conduct a true test of the network activation process. I did not measure transmission of information across ties. Instead, I related county-level rumoring activity to county-level tie volumes. The findings indicate that the spatial distribution of rumoring is consistent with what we would *expect* communication to look like under a network activation theory, but we did not directly observe that network activation process. It is possible that tie volumes at the county level may be proxies for attention rather than conduits for information transmission.

5.3 Future Directions

- **Content evolution:** Measuring topical convergence and content evolution at the daily level proved to be feasible, although the timespan may have been too wide to capture the phenomenon adequately. Two modifications to the approach could help better measure the phenomena. First, coding tweets from a sample from the public stream may allow for better matching between observed data and coded messages. The results suggested that measures of similarity picked up on an emergency management aspect that was common across the set of messages from officials rather than specific content within those messages. A coded corpus based on public messages may prevent this issue. Secondly, an analysis conducted at the hourly level will allow us to keep pace with the rapid flow of information. Some events may be characterized by an interrogatory-disbelief-sensemaking-digressive process (Bordia and DiFonzo, 2004) that occurs within the span of a single a day, thus preventing me from detecting it with daily time slices. Analyses with finer temporal filtering would pose additional challenges, such as message sparsity and the need to account for hourly periodicity. With these finer intervals, however, I could obtain stronger evidence in favor of or against the topical convergence and content evolution findings. While preliminary results suggested that there is topical convergence and some evidence of content evolution, we still have much to learn about *which* topics the population converges on and how message content changes during this evolutionary process.

- **Network activation cascades:** I observed at the county level during primary excitation and secondary excitation that hazard-related communication is elevated where the population has a large number of ties to the event epicenter. It is worth revisiting the network activation problem in order to obtain a stronger understanding information diffusion along these ties. I looked at first order tie volumes between activated counties and the event county, but hazard-related rumoring often follows paths with lengths longer than one. To trace rumoring across paths of length two, three, four, and so on, I will need to examine information diffusion with higher spatial resolution and with finer temporal resolution. Cellular phone data would be ideal in this case. Such records indicate person-to-person information transmission (which I could not directly observe) with very precise temporal metadata. Using cell tower data from sender and recipient, I could geolocate both parties. Examining this information transmission during periods of primary and secondary excitation would yield information on direct information flows to and from the event epicenter, an improvement over my analysis which could not observe the direction of information flows. Furthermore, I could then trace the flow of information from tower to tower to observe secondary and tertiary message transmission This could build on these initial results to provide a much richer picture of information diffusion during and after disaster events.

- **Allocation of attention to warnings:** In both of the spatial filtering chapters I found that warnings and alerts trigger anticipatory excitation ahead of disaster events. Not much is known about allocation of attention to such messages, however. By observing activity following alerts and warnings, I have the opportunity to answer a wide variety of questions related to attentional dynamics in rumoring. What determines the allocation of attention of the public to warnings and alerts? Is it driven by past experience with such events, salience of the event/warnings, risk perception, or other factors? How does sentiment vary in response to these events? Using Project HEROIC's comprehensive data on wireless alerts sent to mobile phones and issued through traditional National Weather Service/Storm Prediction Center channels, I can compare the responses to warnings in a wide variety of contexts: false positives (alerts where no event occurs), realized events, and successive sets of false positives, realized events, or a mixture thereof. For example, does a series of false positives dampen rumoring activity for a subsequent warning? If so, how long does this dampening effect last? Furthermore, I can distinguish attentional effects based on the type of alert: advisory, watch, warning, or warning transmitted directly to cell phones. In addition to contributing to the literature on attentional dynamics and salience, this has practical applications for policy makers and emergency practitioners.

Bibliography

Al-Makaty, Safran S, Douglas A Boyd, and G Norman Van Tubergen. 1994. "Source credibility during the Gulf War: A Q-study of rural and urban Saudi Arabian citizens." *Journalism & Mass Communication Quarterly* 71:55–63.

Allport, Gordon W and Leo Postman. 1946. "An analysis of rumor." *Public Opinion Quarterly* 10:501–517.

Allport, Gordon W and Leo Postman. 1947. "The psychology of rumor." .

Almquist, Zack W. 2010. "US Census Spatial and Demographic Data in R: The UScensus2000 Suite of Packages." *Journal of Statistical Software* 37:1–31.

Almquist, Zack W and Carter T Butts. 2014. "Predicting Regional Self-Identification from Spatial Network Models." *Geographical Analysis* .

Andrade, E.L., S. Blunsden, and R.B. Fisher. 2006. "Modelling Crowd Scenes for Event Detection." In *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, volume 1, pp. 175–178.

Back, Kurt, Leon Festinger, Bernard Hymovitch, Harold Kelley, Stanley Schacter, and John Thibaut. 1950. "The methodology of studying rumor transmission." *Human Relations* .

Back, Mitja D, Albrecht CP Küfner, and Boris Egloff. 2011. "Automatic or the people? Anger on September 11, 2001, and lessons learned for the analysis of large digital data sets." *Psychological Science* 22:837–838.

Baron, Robert S, James P David, Bethany M Brunsman, and Mary Inman. 1997. "Why listeners hear less than they are told: Attentional load and the teller–listener extremity effect." *Journal of personality and social psychology* 72:826.

Batson, C Daniel, David A Lishner, Jennifer Cook, and Stacey Sawyer. 2005. "Similarity and nurturance: Two possible sources of empathy for strangers." *Basic and applied social psychology* 27:15–25.

Blanford, Justine I.; Bernhardt, Jase; Savelyev Alexander; Wong-Parodi Gabrielle; Carleton Andrew M.; Titley David W.; and Alan M. MacEachren. 2014. "Tweeting and Tornadoes." In *Proceedings of the 11th International ISCRAM Conference*, pp. 317–321.

Bollen, Johan, Huina Mao, and Xiaojun Zeng. 2011. "Twitter mood predicts the stock market." *Journal of Computational Science* 2:1–8.

Bordia, Prashant and Nicholas DiFonzo. 2004. "Problem solving in social interactions on the Internet: Rumor as social cognition." *Social Psychology Quarterly* 67:33–49.

Bordia, Prashant and Ralph L Rosnow. 1998. "Rumor Rest Stops on the Information Highway Transmission Patterns in a Computer-Mediated Rumor Chain." *Human Communication Research* 25:163–179.

Bossard, James HS. 1932. "Residential propinquity as a factor in marriage selection." *American Journal of Sociology* pp. 219–224.

Boyd, Danah and Kate Crawford. 2012. "Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon." *Information, Communication & Society* 15:662–679.

Brakman, Steven, Harry Garretsen, Charles Van Marrewijk, and Marianne Van Den Berg. 1999. "The Return of Zipf: Towards a Further Understanding of the Rank-Size Distribution." *Journal of Regional Science* 39:183–213.

Buckner, H Taylor. 1965. "A theory of rumor transmission." *Public Opinion Quarterly* 29:54–70.

Butts, Carter T. 2003. *Predictability of large-scale spatially embedded networks*. na.

Butts, Carter T. and Ryan M. Acton. 2011. "Spatial Modeling of Social Networks." In *The SAGE Handbook of GIS and Society Research*, edited by Timothy Nyerges, Helen Couclelis, and Robert McMaster, chapter 12, pp. 222–250. SAGE Publications.

Butts, Carter T, Ryan M Acton, John R Hipp, and Nicholas N Nagle. 2012. "Geographical variability and network structure." *Social Networks* 34:82–100.

Butts, Carter T. and Zack W. Almquist. 2013. *networkSpatial: Tools for the Generation and Analysis of Spatially-embedded Networks*. R package version 1.00.

Butts, Carter Tribley and Kathleen Carley. 2002. *Spatial models of large-scale interpersonal networks*. Ph.D. thesis, Carnegie Mellon University Pittsburgh.

Caplow, Theodore. 1947. "Rumors in war." *Social Forces* pp. 298–302.

Caputo, Philip. 1999. *A rumor of war*. Random House.

Cheung, C-K and C-M Chan. 2000. "Social-cognitive factors of donating money to charity, with special attention to an international relief organization." *Evaluation and Program Planning* 23:241–253.

Chomsky, Noam. 1998. *The Common Good*. Odonian.

Cutter, Susan L, Bryan J Boruff, and W Lynn Shirley. 2003. "Social vulnerability to environmental hazards*." *Social science quarterly* 84:242–261.

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, Gisele L Pappa, Diogo Rennó Rocha de Oliveira, and Filipe de L Arcanjo. 2011. "Inferring the Location of Twitter Messages Based on User Relationships." *Transactions in GIS* 15:735–751.

De Longueville, Bertrand, Robin S. Smith, and Gianluca Luraschi. 2009. ""OMG, from Here, I Can See the Flames!": A Use Case of Mining Location Based Social Networks to Acquire Spatio-temporal Data on Forest Fires." In *Proceedings of the 2009 International Workshop on Location Based Social Networks*, LBSN '09, pp. 73–80, New York, NY, USA. ACM.

Dodd, Stuart Carter. 1953. "Testing message diffusion in controlled experiments: charting the distance and time factors in the interactance hypothesis." *American Sociological Review* pp. 410–416.

Dodd, Stuart Carter. 1958. "Formulas for spreading opinions." *Public Opinion Quarterly* 22:537–554.

Drabek, Thomas E and David A McEntire. 2003. "Emergent phenomena and the sociology of disaster: lessons, trends and opportunities from the research literature." *Disaster prevention and Management* 12:97–112.

Drabek, Thomas E, Harriet L Tamminga, Thomas S Kilijanek, and Christopher R Adams. 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 and Joanna Brenner. 2013. *The demographics of social media users, 2012*, volume 14. Pew Research Center's Internet & American Life Project.

Dynes, Russell. 1970. "Organizational Involvement and Changes in Community Structure in Disaster." *American Behavioral Scientist* 13:430–439.

Elwood, Sarah, Michael F Goodchild, and Daniel Z Sui. 2012. "Researching Volunteered Geographic Information: Spatial Data, Geographic Research, and New Social Practice." *Annals of the Association of American Geographers* 102:571–590.

Erickson, Bonnie H, TA Nosanchuk, Liviana Mostacci, and Christina Ford Dalrymple. 1978. "The flow of crisis information as a probe of work relations." *Canadian Journal of Sociology* 3.

Fawcett, Tom and Foster Provost. 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, pp. 53–62, New York, NY, USA. ACM.

Ferrari, Laura, Alberto Rosi, Marco Mamei, and Franco Zambonelli. 2011. "Extracting Urban Patterns from Location-based Social Networks." In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*, pp. 9–16. ACM.

Festinger, Leon, Kurt W Back, and Stanley Schachter. 1950. *Social pressures in informal groups: A study of human factors in housing*. Number 3. Stanford University Press.

Fothergill, Alice, Enrique GM Maestas, and JoAnne DeRouen Darlington. 1999. "Race, ethnicity and disasters in the United States: A review of the literature." *Disasters* 23:156–173.

Freeman, Linton C, Sue C Freeman, and Alaina G Michaelson. 1988. "On human social intelligence." *Journal of Social and Biological Structures* 11:415–425.

Fritz, Charles E. 1961. "Disasters." In *Contemporary Social Problems*, edited by R. K. Merton and R. A. Nishet, pp. 97–122. Riverside, CA: University of California Press.

Go, Alec, Richa Bhayani, and Lei Huang. 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 and J Alan Glennon. 2010. "Crowdsourcing Geographic Information for Disaster Response: A Research Frontier." *International Journal of Digital Earth* 3:231–241.

Goodchild, Michael F and Linna Li. 2012. "Assuring the quality of volunteered geographic information." *Spatial statistics* 1:110–120.

Greenberg, Bradley S. 1964. "Diffusion of news of the Kennedy assassination." *Public opinion quarterly* 28:225–232.

Guan, Xiangyang and Cynthia Chen. 2014. "Using social media data to understand and assess disasters." *Natural Hazards* pp. 1–14.

Haber, Ralph Norman and Lyn Haber. 2000. "Experiencing, remembering and reporting events." *Psychology, Public Policy, and Law* 6:1057.

Hägerstrand, Torsten. 1966. "Aspects of the spatial structure of social communication and the diffusion of information." *Papers in Regional Science* 16:27–42.

Hamid, Raffay, A. Johnson, S. Batta, A. Bobick, C. Isbell, and G. Coleman. 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.

Hecht, Brent and Emily Moxley. 2009. "Terabytes of Tobler: Evaluating the first law in a massive, domain-neutral representation of world knowledge." In *Spatial information theory*, pp. 88–105. Springer.

Herman, Edward S and Noam Chomsky. 1988. *Manufacturing consent: The political economy of the mass media*. Random House.

Hipp, John R, Carter T Butts, Ryan Acton, Nicholas N Nagle, and Adam Boessen. 2013. "Extrapolative simulation of neighborhood networks based on population spatial distribution: Do they predict crime?" *Social Networks* 35:614–625.

Hoffman, Matthew, Francis R Bach, and David M Blei. 2010. "Online learning for latent dirichlet allocation." In *advances in neural information processing systems*, pp. 856–864.

Hollenstein, Livia and Ross Purves. 2013. "Exploring place through user-generated content: Using Flickr tags to describe city cores." *Journal of Spatial Information Science* pp. 21–48.

Hopkins, Daniel J and Gary King. 2010. "A method of automated nonparametric content analysis for social science." *American Journal of Political Science* 54:229–247.

Hughes, Amanda Lee and Leysia Palen. 2009. "Twitter adoption and use in mass convergence and emergency events." *International Journal of Emergency Management* 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 Opinion Quarterly* 53:467–481.

Kapferer, Jean-Noel. 2013. *Rumors: Uses, interpretations, and images*. Transaction Publishers.

Killian, Lewis M and Ralph H Turner. 1972. "Collective behavior."

King, Gary. 2011. "Ensuring the data-rich future of the social sciences." *Science* 331:719–721.

Kireyev, Kirill, Leysia Palen, and K Anderson. 2009. "Applications of topics models to analysis of disaster-related twitter data." In *NIPS Workshop on Applications for Topic Models: Text and Beyond*, volume 1.

Kouloumpis, Efthymios, Theresa Wilson, and Johanna Moore. 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." *International Journal of Mass Emergencies and Disasters* .

Kreps, Gary A. 1984. "Sociological Inquiry and Disaster Research." *Annual Review of Sociology* pp. 309–330.

Lansdall-Welfare, Thomas, Vasileios Lampos, and Nello Cristianini. 2012. "Nowcasting the mood of the nation." *Significance* 9:26–28.

Larsen, Otto N. 1954. "Rumors in a Disaster." *Journal of Communication* 4:111–123.

Latané, Bibb, James H Liu, Andrzej Nowak, Michael Bonevento, and Long Zheng. 1995. "Distance matters: Physical space and social impact." *Personality and Social Psychology Bulletin* 21:795–805.

Latonero, Mark and Irina Shklovski. 2010. "Respectfully Yours in Safety and Service: Emergency Management & Social Media Evangelism." In *Proceedings of the 7th International ISCRAM Conference–Seattle*, volume 1.

Lazer, David M, Ryan Kennedy, Gary King, and Alessandro Vespignani. 2014. "The Parable of Google Flu: Traps in Big Data Analysis." *Science* 343:1203–1205.

Leik, Robert K, T Michael Carter, John P Clark, Stephanie D Kendall, and Gregory A Gifford. 1981. "Community Response to Natural Hazard Warnings." Technical report, DTIC Document.

Leinweber, David J. 2007. "Stupid data miner tricks: overfitting the S&P 500." *The Journal of Investing* 16:15–22.

Leung, Daniel and Shawn Newsam. 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., P. Mausel, E. Brondzio, and E. Moran. 2004. "Change detection techniques." *International Journal of Remote Sensing* 25:2365–2401.

Lucas, Christopher, Richard A. Nielsen, Margaret E. Roberts, Brandon M. Stewart, Alex Storer, and Dustin Tingley. 2015. "Computer-Assisted Text Analysis for Comparative Politics." *Political Analysis* 23:254–277.

Macleod, Robb D and Russell G Congalton. 1998. "A quantitative comparison of change-detection algorithms for monitoring eelgrass from remotely sensed data." *Photogrammetric engineering and remote sensing* 64:207–216.

Maddock, Jim, Kate Starbird, Haneen Al-Hassani, Daniel E Sandoval, Mania Orand, and Robert M Mason. 2015. "Characterizing Online Rumoring Behavior Using Multi-Dimensional Signatures." In *ACM Conference on Computer-Supported Cooperative Work and Social Computing*.

Mamei, Marco, Alberto Rosi, and Franco Zambonelli. 2010. "Automatic analysis of geo-tagged photos for intelligent tourist services." In *Intelligent Environments (IE), 2010 Sixth International Conference on*, pp. 146–151. IEEE.

Mendoza, Marcelo, Barbara Poblete, and Carlos Castillo. 2010. "Twitter under crisis: Can we trust what we RT?" In *Proceedings of the first workshop on social media analytics*, pp. 71–79. ACM.

Mileti, Dennis S and Paul W O'Brien. 1992. "Warnings during disaster: Normalizing communicated risk." *Social Problems* 39:40–57.

Miller, Dan E. 1992. "Snakes in the Greens and Rumor in the Innercity." *The Social Science Journal* 29:381–393.

Modsching, Marko, Ronny Kramer, and Klaus ten Hagen. 2006. "Field trial on GPS Accuracy in a medium size city: The influence of built-up." In *3rd Workshop on Positioning, Navigation and Communication*, pp. 209–218.

Pak, Alexander and Patrick Paroubek. 2010. "Twitter as a Corpus for Sentiment Analysis and Opinion Mining." In *LREC*.

Perry, Ronald W, Michael K Lindell, and Marjorie Greene. 1981. *Evacuation planning in emergency management*. Lexington Books.

Pozdnoukhov, Alexei and Christian Kaiser. 2011. "Space-time dynamics of topics in streaming text." In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*, pp. 1–8. ACM.

Quarantelli, Enrico L. 1954. "The Nature and Conditions of Panic." *American Journal of Sociology* pp. 267–275.

Quarantelli, Enrico Louis. 1980. "Evacuation Behavior and Problems: Findings and Implications from the Research Literature." Technical report, DTIC Document.

Quarantelli, Enrico L and Russell R Dynes. 1977. "Response to Social Crisis and Disaster." *Annual Review of Sociology* 3:23–49.

Ramage, Daniel, Susan T Dumais, and Daniel J Liebling. 2010. "Characterizing Microblogs with Topic Models." *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media* 10:130–137.

Ribeiro Jr, Sílvio S, Clodoveu A Davis Jr, Diogo Rennó R Oliveira, Wagner Meira Jr, Tatiana S Gonçalves, and Gisele L Pappa. 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*, pp. 5–11. ACM.

Richardson, RJ, Bonnie H Erickson, and TA Nosanchuk. 1979. "Community size, network structure, and the flow of information." *Canadian Journal of Sociology/Cahiers canadiens de sociologie* pp. 379–392.

Roberts, Margaret E, Brandon M Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G Rand. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58:1064–1082.

Rogers, Everett M. 2000. "Reflections on news event diffusion research." *Journalism & Mass Communication Quarterly* 77:561–576.

Romney, A Kimball, Susan C Weller, and William H Batchelder. 1986. "Culture as Consensus: A Theory of Culture and Informant Accuracy." *American anthropologist* 88:313–338.

Romney, Kimball A. and Susan C. Weller. 1984. "Predicting Informant Accuracy from Patterns of Recall Among Individuals." *Social Networks* 6:59–77.

Rosnow, Ralph L. 1974. "On Rumor." *Journal of Communication* 24:26–38.

Rosnow, Ralph L. 1988. "Rumor as communication: A contextualist approach." *Journal of Communication* 38:12–28.

Rosnow, Ralph L. and Alan J. Kimmel. 2000. "Rumor." In *Encyclopedia of Psychology, Volume 7*, edited by Alan E. Kazdin, pp. 122–23. Oxford, England: Oxford University Press.

Rovai, Eugenie and Christine M Rodrigue. 1998. *The Northridge and Ferndale earthquakes: Spatial inequities in media attention and recovery*. National Emergency Training Center.

Scanlon, Joseph. 2007. "Sampling an unknown universe: problems of researching mass casualty incidents (a history of ECRU's field research)." *Statistics in medicine* 26:1812–1823.

Scanlon, T Joseph. 1977. "Post-disaster rumor chains: A case study." *Mass Emergencies* 2:22–27.

Schachter, Stanley and Harvey Burdick. 1955. "A field experiment on rumor transmission and distortion." *The Journal of Abnormal and Social Psychology* 50:363.

Shibutani, Tamotsu. 1966. *Improvised news*. Ardent Media.

Shklovski, Irina, Leysia Palen, and Jeannette Sutton. 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." *International Journal of Remote Sensing* 10:989–1003.

Singhal, Amit. 2001. "Modern information retrieval: A brief overview." *IEEE Data Eng. Bull.* 24:35–43.

Sinha, Durganand. 1952. "Behaviour in a catastrophic situation: A psychological study of reports and rumours." *British Journal of Psychology. General Section* 43:200–209.

Starbird, Kate, Grace Muzny, and Leysia Palen. 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*.

Starbird, Kate and Leysia Palen. 2010. *Pass it on?: Retweeting in mass emergency*. International Community on Information Systems for Crisis Response and Management.

Starbird, Kate, Leysia Palen, Amanda L Hughes, and Sarah Vieweg. 2010. "Chatter on the red: what hazards threat reveals about the social life of microblogged information." In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, pp. 241–250. ACM.

Stauffer, Chris and W.E.L. Grimson. 2000. "Learning patterns of activity using real-time tracking." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 22:747–757.

Stein, Howard F. 1980. "Wars and rumors of wars: A psychohistorical study of a medical culture." *Journal of Psychohistory* 7:379–401.

Sudman, Seymour, Norman M Bradburn, and Norbert Schwarz. 1996. *Thinking about answers: The application of cognitive processes to survey methodology.* Jossey-Bass.

Sutton, Jeannette, Cedar League, Timothy L Sellnow, and Deanna D Sellnow. 2015. "Terse messaging and public health in the midst of natural disasters: the case of the Boulder floods." *Health communication* 30:135–143.

Sutton, Jeannette, Leysia Palen, and Irina Shklovski. 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, Emma Spiro, Carter Butts, Sean Fitzhugh, Britta Johnson, and Matt Greczek. 2013a. "Tweeting the Spill: Online Informal Communications, Social Networks, and Conversational Microstructures during the Deepwater Horizon Oilspill." *International Journal of Information Systems for Crisis Response and Management (IJISCRAM)* 5:58–76.

Sutton, Jeannette, Emma S Spiro, Britta Johnson, Sean Fitzhugh, Ben Gibson, and Carter T Butts. 2013b. "Warning tweets: serial transmission of messages during the warning phase of a disaster event." *Information, Communication & Society* pp. 1–23.

Sutton, Jeannette N. 2010. *Twittering Tennessee: distributed networks and collaboration following a technological disaster*. ISCRAM.

Tierney, Kathleen, Christine Bevc, and Erica Kuligowski. 2006. "Metaphors matter: Disaster myths, media frames, and their consequences in Hurricane Katrina." *The Annals of the American Academy of Political and Social Science* 604:57–81.

Tobler, Waldo R. 1970. "A computer movie simulating urban growth in the Detroit region." *Economic geography* 46:234–240.

Twaroch, Florian A, Philip D Smart, and Christopher B Jones. 2008. "Mining the web to detect place names." In *Proceedings of the 2nd international workshop on Geographic information retrieval*, pp. 43–44. ACM.

Vieweg, Sarah, Amanda L Hughes, Kate Starbird, and Leysia Palen. 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*, pp. 1079–1088. ACM.

Vieweg, Sarah, Leysia Palen, Sophia B Liu, Amanda L Hughes, and Jeannette Sutton. 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)*.

Wakamiya, Shoko, Ryong Lee, and Kazutoshi Sumiya. 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*, pp. 77–84. ACM.

Walker, Charles J and Carol A Beckerle. 1987. "The effect of state anxiety on rumor transmission." *Journal of Social Behavior & Personality* .

Weidemann, C and J Swift. 2013. "Social Media Location Intelligence: The Next Privacy Battle-An ArcGIS add-in and Analysis of Geospatial Data Collected from Twitter. com." *International Journal of Geoinformatics* 9.

Weller, Susan C and A Kimball Romney. 1988. *Systematic data collection*, volume 10. Sage.

Wing, Michael G, Aaron Eklund, and Loren D Kellogg. 2005. "Consumer-grade global positioning system (GPS) accuracy and reliability." *Journal of Forestry* 103:169–173.

Yang, Shuang-Hong, Alek Kolcz, Andy Schlaikjer, and Pankaj Gupta. 2014. "Large-scale high-precision topic modeling on twitter." In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1907–1916. ACM.

Zagefka, Hanna, Masi Noor, Rupert Brown, Tim Hopthrow, and Georgina Randsley Moura. 2012. "Eliciting donations to disaster victims: Psychological considerations." *Asian Journal of Social Psychology* 15:221–230.

Zandbergen, Paul A and Sean J Barbeau. 2011. "Positional accuracy of assisted gps data from high-sensitivity gps-enabled mobile phones." *Journal of Navigation* 64:381–399.

Zheng, Rong, Jiexun Li, Hsinchun Chen, and Zan Huang. 2006. "A framework for authorship identification of online messages: Writing-style features and classification techniques." *Journal of the American Society for Information Science and Technology* 57:378–393.

Zipf, George K. 1949. *Human Behavior and the Principle of Least Effort*. Addison Wesley.

Appendix A

Appendix

In this appendix I provide the list of tables illustrating the result of t-tests on the adjusted spatio-temporally filtered topic signals. Each test compares the cosine similarities between messages in a spatio-temporal bin and one of the coded categories. For each test I compare the observed cosine similarities on a given day relative to a disaster event across all eight events with the cosine similarities observed between messages in that bin and messages in the coded corpus across the entire period of observation. This is effectively a comparison between observed messages during the period of time surrounding a disaster and a baseline measure of similarity, where the baseline is a measure of *all* cosine similarities during the nine-month reference period.

T-tests, three days prior to disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	-0.113	7	-0.414	.691
Closure/Opening	-0.008	7	-0.037	.972
Correction	-0.207	7	-0.492	.638
Emotive	-0.033	7	-0.125	.904
Evacuation	0.034	7	0.142	.891
Impact	-0.158	7	-0.481	.645
Help/Directed Comm.	-0.022	7	-0.082	.937
Information	-0.031	7	-0.122	.907
Thanks	-0.315	7	-0.856	.420
Volunteer/Donate	-0.202	7	-0.588	.575
Regional				
Advisory	-0.267	7	-0.819	.440
Closure/Opening	-0.268	7	-1.084	.314
Correction	-0.557	7	-1.634	.146
Emotive	-0.280	7	-1.104	.306
Evacuation	-0.358	7	-1.986	.087
Impact	-0.453	7	-1.384	.209
Help/Directed Comm.	-0.010	7	-0.038	.971
Information	-0.268	7	-1.065	.322
Thanks	-0.239	7	-0.624	.552
Volunteer/Donate	-0.411	7	-1.418	.199
Distant				
Advisory	-0.352	7	-0.906	.395
Closure/Opening	-0.064	7	-0.149	.886
Correction	-0.247	7	-0.506	.628
Emotive	-0.245	7	-0.857	.420
Evacuation	-0.115	7	-0.380	.715
Impact	-0.438	7	-1.082	.315
Help/Directed Comm.	0.028	7	0.095	.927
Information	-0.302	7	-1.160	.284
Thanks	-0.148	7	-0.430	.680
Volunteer/Donate	-0.530	7	-1.707	.132

T-tests, two days prior to disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	-0.069	7	-0.169	.870
Closure/Opening	-0.384	7	-0.384	.712
Correction	0.264	7	1.099	.308
Emotive	-0.181	7	-0.819	.440
Evacuation	0.009	7	0.031	.976
Impact	0.132	7	0.445	.670
Help/Directed Comm.	0.053	7	0.163	.875
Information	-0.187	7	-0.544	.604
Thanks	0.131	7	0.446	.669
Volunteer/Donate	-0.087	7	-0.395	.705
Regional				
Advisory	0.175	7	0.632	.547
Closure/Opening	0.344	7	1.294	.237
Correction	0.363	7	1.298	.236
Emotive	0.278	7	0.864	.416
Evacuation	0.367	7	1.300	.235
Impact	0.112	7	0.401	.700
Help/Directed Comm.	-0.030	7	-0.143	.890
Information	0.152	7	0.425	.684
Thanks	0.001	7	0.008	.994
Volunteer/Donate	0.200	7	0.499	.633
Distant				
Advisory	-0.102	7	-0.439	.674
Closure/Opening	0.049	7	0.205	.843
Correction	0.000	7	-0.002	.999
Emotive	-0.059	7	-0.188	.856
Evacuation	-0.111	7	-0.604	.565
Impact	0.008	7	0.034	.974
Help/Directed Comm.	-0.010	7	-0.032	.976
Information	-0.079	7	-0.275	.792
Thanks	-0.003	7	-0.012	.991
Volunteer/Donate	0.225	7	0.888	.404

T-tests, one day prior to disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	0.868	7	1.053	.327
Closure/Opening	0.436	7	0.568	.588
Correction	0.089	7	0.098	.925
Emotive	0.751	7	1.161	.284
Evacuation	0.393	7	0.458	.661
Impact	0.430	7	0.516	.622
Help/Directed Comm.	1.150	7	1.844	.108
Information	0.474	7	0.548	.601
Thanks	0.602	7	1.088	.313
Volunteer/Donate	0.064	7	0.120	.908
Regional				
Advisory	-0.302	7	-0.618	.556
Closure/Opening	-0.506	7	-0.910	.393
Correction	-0.629	7	-0.826	.436
Emotive	-0.348	7	-0.777	.463
Evacuation	-0.480	7	-0.691	.512
Impact	-0.572	7	-1.083	.315
Help/Directed Comm.	0.116	7	0.412	.693
Information	-0.452	7	-0.866	.415
Thanks	-0.327	7	-0.964	.367
Volunteer/Donate	-0.590	7	-1.508	.175
Distant				
Advisory	0.278	7	0.378	.717
Closure/Opening	-0.012	7	-0.017	.987
Correction	-0.820	7	-0.989	.356
Emotive	-0.267	7	-0.565	.590
Evacuation	0.064	7	0.071	.945
Impact	-0.412	7	-0.874	.411
Help/Directed Comm.	0.266	7	0.513	.624
Information	-0.319	7	-0.629	.549
Thanks	0.100	7	0.164	.875
Volunteer/Donate	-0.243	7	-0.578	.581

T-tests, day of disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	3.075	7	5.078	.001
Closure/Opening	2.578	7	6.687	.000
Correction	2.808	7	3.417	.011
Emotive	3.005	7	6.764	.000
Evacuation	2.281	7	4.171	.004
Impact	2.316	7	4.107	.005
Help/Directed Comm.	2.826	7	4.866	.002
Information	2.882	7	5.022	.002
Thanks	2.744	7	6.124	.000
Volunteer/Donate	2.267	7	3.982	.005
Regional				
Advisory	0.618	7	1.134	.294
Closure/Opening	0.316	7	0.953	.372
Correction	0.029	7	0.046	.965
Emotive	0.230	7	0.604	.565
Evacuation	0.133	7	0.238	.819
Impact	0.391	7	0.679	.519
Help/Directed Comm.	0.815	7	1.235	.257
Information	0.704	7	1.095	.310
Thanks	0.441	7	0.561	.592
Volunteer/Donate	0.407	7	0.609	.562
Distant				
Advisory	-0.564	7	-1.522	.172
Closure/Opening	-0.503	7	-1.359	.216
Correction	-0.671	7	-2.131	.071
Emotive	-0.812	7	-1.271	.244
Evacuation	-0.480	7	-1.33	.227
Impact	-0.724	7	-2.497	.041
Help/Directed Comm.	-0.094	7	-0.234	.821
Information	-0.684	7	-1.821	.112
Thanks	-0.379	7	-1.201	.269
Volunteer/Donate	-0.833	7	-1.247	.252

T-tests, one day after to disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	1.957	7	3.512	.010
Closure/Opening	1.810	7	3.610	.009
Correction	1.605	7	3.594	.009
Emotive	1.725	7	2.550	.038
Evacuation	1.915	7	3.133	.017
Impact	1.601	7	2.736	.029
Help/Directed Comm.	2.483	7	2.960	.021
Information	1.701	7	3.117	.017
Thanks	2.116	7	3.185	.015
Volunteer/Donate	0.823	7	1.062	.324
Regional				
Advisory	0.801	7	2.360	.050
Closure/Opening	0.745	7	1.625	.148
Correction	0.882	7	1.818	.112
Emotive	1.254	7	1.569	.161
Evacuation	0.403	7	1.218	.263
Impact	0.418	7	0.912	.392
Help/Directed Comm.	1.309	7	2.36	.050
Information	0.747	7	1.853	.106
Thanks	0.937	7	1.366	.214
Volunteer/Donate	0.444	7	0.703	.505
Distant				
Advisory	-0.803	7	-2.407	.047
Closure/Opening	-0.720	7	-2.322	.053
Correction	-0.767	7	-2.992	.020
Emotive	-0.620	7	-1.090	.312
Evacuation	-0.746	7	-2.266	.058
Impact	-0.818	7	-2.717	.029
Help/Directed Comm.	-0.436	7	-1.044	.331
Information	-0.838	7	-2.428	.046
Thanks	-0.598	7	-1.784	.118
Volunteer/Donate	-1.511	7	-3.143	.016

T-tests, two days after disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	0.854	7	1.902	.099
Closure/Opening	0.425	7	1.854	.106
Correction	0.808	7	1.623	.149
Emotive	0.254	7	0.632	.548
Evacuation	0.878	7	2.097	.074
Impact	0.857	7	1.903	.099
Help/Directed Comm.	0.714	7	1.783	.118
Information	0.754	7	1.834	.109
Thanks	0.262	7	0.795	.453
Volunteer/Donate	0.498	7	1.154	.286
Regional				
Advisory	0.696	7	1.209	.266
Closure/Opening	0.688	7	1.032	.336
Correction	0.719	7	1.203	.268
Emotive	0.779	7	1.215	.264
Evacuation	0.642	7	1.236	.257
Impact	0.483	7	0.985	.358
Help/Directed Comm.	0.700	7	1.268	.245
Information	0.608	7	1.014	.344
Thanks	0.839	7	1.241	.255
Volunteer/Donate	0.356	7	0.521	.618
Distant				
Advisory	0.050	7	0.119	.908
Closure/Opening	0.236	7	0.434	.677
Correction	-0.120	7	-0.376	.718
Emotive	0.204	7	0.416	.690
Evacuation	0.195	7	0.453	.664
Impact	0.062	7	0.147	.887
Help/Directed Comm.	0.157	7	0.412	.693
Information	0.126	7	0.267	.798
Thanks	-0.090	7	-0.222	.831
Volunteer/Donate	-0.120	7	-0.248	.811

T-tests, three days after disaster

	Mean of Differences	DF	t-value	p-value
Local				
Advisory	0.343	7	0.987	.357
Closure/Opening	-0.111	7	-0.404	.698
Correction	0.455	7	1.305	.233
Emotive	0.437	7	1.099	.308
Evacuation	0.320	7	1.003	.349
Impact	0.161	7	0.683	.517
Help/Directed Comm.	1.011	7	1.674	.138
Information	0.078	7	0.308	.767
Thanks	0.756	7	1.793	.116
Volunteer/Donate	0.410	7	0.964	.367
Regional				
Advisory	-0.200	7	-0.723	.490
Closure/Opening	-0.386	7	-1.642	.145
Correction	-0.148	7	-0.453	.664
Emotive	0.072	7	0.197	.850
Evacuation	-0.207	7	-0.657	.532
Impact	-0.105	7	-0.339	.744
Help/Directed Comm.	-0.302	7	-0.691	.512
Information	-0.182	7	-0.695	.509
Thanks	-0.208	7	-0.463	.657
Volunteer/Donate	-0.151	7	-0.495	.634
Distant				
Advisory	-0.751	7	-3.650	.008
Closure/Opening	-0.752	7	-3.370	.012
Correction	-0.743	7	-2.519	.040
Emotive	-0.905	7	-3.800	.007
Evacuation	-0.622	7	-2.684	.031
Impact	-0.722	7	-3.404	.011
Help/Directed Comm.	-0.465	7	-1.374	.212
Information	-0.896	7	-3.302	.013
Thanks	-0.545	7	-2.358	.050
Volunteer/Donate	-1.078	7	-2.576	.037