

Emergent Forms of Online Sociality in Disasters Arising from Natural Hazards

by

MARINA KOGAN

M.S., University of Colorado Boulder, 2015

M.A., University of Illinois, 2010

B.S./B.A., City University of New York, 2005

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Doctor of Philosophy
Computer Science Department

2017

This dissertation entitled:

Emergent Forms of Online Sociality
in Disasters Arising from Natural Hazards

written by Marina Kogan
has been approved for the Department of Computer Science

Leysia Palen

Aaron Clauset

Clayton Lewis

Jed Brubaker

Mor Naaman

Date _____

The final copy of this thesis has been examined by the signatories, and we
Find that both the content and the form meet acceptable presentation standards
Of scholarly work in the above mentioned discipline.

ABSTRACT OF THE DISSERTATION

Kogan, Marina (Ph.D., Computer Science)

Emergent Forms of Online Sociality in Disasters Arising from Natural Hazards

Thesis directed by Professor Leysia Palen

Disasters arising from natural hazards are associated with breakdown of existing structures, but they also result in creation of new social ties in the process of self-organization and problem solving by those affected. This dissertation focuses on emergent forms of sociality that arise in the context of crisis. Specifically, it considers collaborative work practices, social network structures, and organizational forms that emerge on social media during disasters arising from natural hazards. Social media platforms support highly-distributed social environments, and the forms of sociality that emerge in these contexts are affected by the affordances of their technical features, especially those that more or less successfully facilitate the creation of a shared information space. Thus, this dissertation is organized around two important aspects of social media spaces: the availability of an explicitly-shared site of work and the availability of a visible, legible record of activity.

This dissertation investigates the forms of sociality that emerge during disasters in three social media activities: retweeting, crisis mapping in OpenStreetMap (OSM), and Twitter reply conversations. These three social media activities highlight various availability of an explicitly-shared site of work and visible record of activity. The studies of retweeting and reply conversations investigate the Twitter activity in response to the 2012 Hurricane Sandy—the second costliest hurricane in US history and the most tweeted about event to date at the time. Analysis of crisis mapping in OpenStreetMap—an open, editable, volunteer-based map of the world—focuses on the OSM activity after the 2010 Haiti earthquake, which was the first major disaster event supported by OpenStreetMap. For these investigations, the dissertation elaborates and develops human-centered data science methods—a set of methodological approaches that both harness the power of computational techniques and account for the highly-situated nature of the social activity in crisis. Finally, the dissertation positions the findings from the three studies

within the larger context of high-tempo, high-volume social media activity and highlights how the framework of the two intersecting dimensions of the shared information space reveals larger patterns within the emergent forms of sociality across contexts.

ACKNOWLEDGMENTS

First and foremost, I'd like to thank my incredible advisor Leysia Palen. She always knows how to push me to become a better, more reflexive scholar, while also being the most supportive mentor. I only hope to be such an inspiring role model to my future students.

I'd like to thank all my collaborators for all the support and productive discussions that made this work possible: Ken Anderson, Julie Demuth, Jennifer Henderson, Heather Lazrus, Rebecca Morss, Sid Redner, and Olga Wilhelmi. I'm especially grateful to my committee members—Aaron Clauset, Clayton Lewis, Jed Brubaker, and Mor Naaman—who graciously invested their time and energy into improving this dissertation and through that process into challenging me to be a better scholar.

A special thank you goes to my friends and colleagues Melissa Bica, Jennings Anderson, Lise St. Dennis, Nicole Beckage, Robert Soden, and Abigale Stangl, who in addition to academic discussions and collaborations offered invaluable emotional support throughout this process. Thank you for your patience with me.

Finally, I'd like to thank my husband, Koji Otsuka, for helping me stay grounded and sane through this incredible roller coaster ride and for lovingly nudging me to take it as a learning experience both professionally and personally.

Table of Contents

CHAPTER 1 : INTRODUCTION	1
1.1 SOCIAL MEDIA IN DISASTER.....	1
1.2 EMERGENT FORMS OF SOCIALITY IN DISASTER	1
1.3 SOCIAL MEDIA CONTEXT	2
1.4 STRUCTURE VS. AGENCY DEBATE: STRUCTURATION AND COMPLEXITY SCIENCE AS AN ANSWER	4
1.5 OVERVIEW OF THE STUDIES.....	6
CHAPTER 2 : BACKGROUND LITERATURE	8
2.1 DISASTER SOCIOLOGY.....	8
2.1.1 <i>The Public as First Responders</i>	8
2.1.2 <i>Convergent Behavior</i>	8
2.1.3 <i>Temporality of Disaster</i>	9
2.1.4 <i>Organizational Structure in Disaster</i>	10
2.2 CRISIS INFORMATICS.....	13
2.2.1 <i>Active Participation in Early CI Research</i>	13
2.2.2 <i>Social Media and Microblogging in Crisis</i>	15
Social Media and Situational Awareness	17
Credibility of Information.....	18
Collective Intelligence.....	20
2.2.3 <i>Emergent Organizations in Crisis</i>	21
2.3 EMERGENT FORMS OF SOCIALITY	23
2.3.1 <i>Organizational Forms</i>	24
2.3.2 <i>Network Structures</i>	27
2.3.3 <i>Collaborative Work Practices</i>	29
2.4 SOCIAL MEDIA CONTEXT	31
CHAPTER 3 : CONCEPTUAL FRAMING AND METHODS	35
3.1 RESEARCH QUESTIONS	36
3.2 CONCEPTUAL FRAMING.....	36
3.3 INCLUSION OF PUBLISHED WORK.....	39
3.4 METHODOLOGICAL CONSIDERATIONS: HUMAN-CENTERED DATA SCIENCE.....	39
CHAPTER 4 : RETWEETING BY THE GEOGRAPHICALLY-VULNERABLE DURING HURRICANE SANDY	43
4.1 STUDY OVERVIEW	43
4.2 STUDY INTRODUCTION	44
4.2.1 <i>Study Objective & Background</i>	45
4.3 THE 2012 HURRICANE SANDY EVENT (US LANDFALL).....	48
4.4 METHODS.....	49
4.4.1 <i>Data Collection Steps</i>	49
4.4.2 <i>Creating the Data Sets</i>	50
Three Retweet Counts.....	52
4.5 ANALYSES & FINDINGS.....	52
4.5.1 <i>Geo-Vulnerable Retweet Networks for the 4 Time Slices</i>	53
Time Evolution of Retweet Networks	53
Network Size & Density.....	53

Degree Distributions.....	55
Network Mixing Patterns.....	57
Transitivity	59
Important Users.....	60
4.5.2 <i>Geographic Patterns of Retweet Activity</i>	60
Proportion of Geo-Vulnerable Sources within the Four Time Sliced Networks.....	60
Proportion of the Geo-Vulnerable Source Tweets.....	61
4.5.3 <i>Retweet Count Distributions of Various Populations</i>	63
Retweet Distributions: The Geo-Vulnerable Retweeting Other Geo-Vulnerable Tweets vs Global Users Retweeting Global Tweets	63
Retweet Count Distributions: The Geo-Vulnerable Retweeting Global Tweets vs Global Users Retweeting Global Tweets	64
4.5.4 <i>Content</i>	65
Overrepresented tweets in the Global/Geo-Vulnerable count	65
4.6 STUDY SUMMARY.....	67
4.7 EMERGENT FORMS OF SOCIALITY AND TWO ASPECTS OF SHARED INFORMATION SPACE	68
4.7.1 <i>Collaborative Work Practices</i>	69
4.7.2 <i>Network Structures</i>	69
4.7.3 <i>Organizational Forms</i>	70
CHAPTER 5 : OPENSTREETMAP CRISIS MAPPING IN THE WAKE OF THE 2010 HAITI EARTHQUAKE	71
5.1 STUDY OVERVIEW	71
5.2 STUDY INTRODUCTION	72
5.2.1 <i>The Cooperative Work of Digital Volunteerism</i>	72
5.2.2 <i>OpenStreetMap</i>	74
5.3 RESEARCH MOTIVATION	74
5.3.1 <i>Prior OpenStreetMap Studies</i>	75
5.3.2 <i>An Expanded View of Map Collaboration</i>	77
5.4 RESEARCH PLAN	77
5.4.1 <i>Collection of Geospatial OSM Data</i>	78
5.4.2 <i>Initial Network Construction to Generate the Dataset for Qualitative Analysis</i>	78
Initial Method 1: Changeset-Level Networks.....	79
Initial Method 2: Road-Level Network	81
5.4.3 <i>Network-Based Data Scoping and Qualitative Analysis</i>	81
5.5 MAPPING PRACTICES & THE NETWORK TECHNIQUES FOR ISOLATING THEM.....	82
5.5.1 <i>Temporal Patterns of Practice</i>	82
Temporal Practice 1: Shift Work	82
Temporal Practice 2: Sustained Convergence	83
5.5.2 <i>Spatial Patterns of Practice</i>	84
Spatial Practice 1: Mapping of Externally Salient Features	84
Spatial Practice 2: Mapping Continuous Structures.....	86
Spatial Practice 3: Map Filling	87
5.5.3 <i>Interpersonal Patterns of Practice</i>	89
Interpersonal Practice 1: Parallel Mapping.....	89
Interpersonal Practice 2: Tag Correction	91
5.6 STUDY SUMMARY & DISCUSSION.....	93
5.6.1 <i>Expanded View of Collaboration</i>	95
5.7 ELUCIDATING NETWORK STRUCTURES AND ORGANIZATION FORMS	96
5.7.1 <i>Introduction to Analysis of Network Structure and Organizational Forms</i>	96
5.7.2 <i>Network Structures</i>	96
Time Slice Construction for Network Analysis	96
Network Statistics under Configuration Model	97
Network Interconnectedness	97
Degree Distribution.....	98

Pro-Social Network Structures: Reciprocity and Transitivity	99
Assortativity: Who Maps with Whom	101
5.7.3 <i>Organizational Forms: Social Roles of Expert and Novice Mappers</i>	102
Assortativity with respect to Mapper Experience	102
Community Structure.....	103
5.8 FORMS OF SOCIALITY AND TWO ASPECTS OF SHARED INFORMATION SPACE.....	104
CHAPTER 6 : TWITTER CONVERSATIONS AMONG LOCALS DURING HURRICANE SANDY	106
.....	
6.1 STUDY OVERVIEW	106
6.2 STUDY INTRODUCTION	107
6.3 THE 2012 HURRICANE SANDY EVENT (US LANDFALL).....	111
6.4 DATA SETS AND METHODS.....	111
6.4.1 <i>Data Collection Steps</i>	111
6.4.2 <i>Data Sets</i>	112
6.4.3 <i>Conversation Construction</i>	112
6.4.4 <i>Temporal Analysis Methods</i>	113
6.4.5 <i>Network Analysis Methods</i>	113
6.4.6 <i>Conversationalist-Based Analysis Methods</i>	114
6.4.7 <i>Content Analysis Methods</i>	114
6.5 ANALYSES AND FINDINGS	115
6.5.1 <i>Temporality in Replying Activity</i>	115
6.5.2 <i>Inter-tweet Time</i>	115
Time Spans between Outgoing Replies	116
Inter-Reply Time	116
Conversational Tempo & Duration across Time Periods	117
Interpretation of Temporal Behavior Patterns.....	118
6.5.3 <i>Relationality in Replying Activity</i>	119
Size and Density of Time Period Networks	119
Degree and Strength of the Time Period Networks	120
Reciprocity & Degree Assortativity in Time Period Networks.....	121
Summary of Relationality in Replying Activity	122
6.5.4 <i>Conversationalist Composition</i>	122
By Number of Participants.....	123
By Popularity of Conversationalists.....	124
6.5.5 <i>Conversation Content</i>	126
On-Topic Conversations and Frequent Conversationalists	127
6.6 STUDY SUMMARY & DISCUSSION.....	128
6.6.1 <i>Temporality</i>	128
6.6.2 <i>Relationality</i>	129
6.6.3 <i>Conversationalist Composition & Content</i>	130
6.6.4 <i>Implications for Design</i>	130
6.7 BEYOND THE TWEET	131
6.8 FORMS OF SOCIALITY AND TWO ASPECTS OF SHARED INFORMATION SPACE.....	131
6.8.1 <i>Collaborative Work Practices</i>	132
6.8.2 <i>Network Structures</i>	132
6.8.3 <i>Organizational Forms</i>	133
CHAPTER 7 : DISCUSSION	134
7.1 INTRODUCTION: EMERGENT FORMS OF SOCIALITY IN SOCIAL MEDIA	134
7.2 EMERGENT FORMS OF SOCIALITY AND TWO DIMENSIONS OF SHARED INFORMATION SPACE	135
7.2.1 <i>Collaborative Work Practices</i>	135
7.2.2 <i>Network Structures</i>	137
7.2.3 <i>Organizational Forms</i>	141

7.3 CONTEXTUAL CONSIDERATIONS: LIMITATIONS AND OPPORTUNITIES	142
7.4 BEYOND THE THREE STUDIES: A BROADER VISION	145
REFERENCES.....	149

LIST OF TABLES

Table 2.1: Disaster Stages according to Powell (1954).	10
Table 2.2: Types of Organizations in Disaster according to Dynes (1970).	11
Table 4.1: Dataset Names and Descriptions.....	47
Table 4.2: Retweet volume for in-common Twitterers across all time slices.	50
Table 4.3: Glossary of Retweet Counts.....	53
Table 4.4: Network size and density.	55
Table 4.5: Nodes with the highest PageRank.....	60
Table 4.6: Source tweet author details.	61
Table 4.7: Proportion of tweets with locally useful content in Global/Geo-Vulnerable retweet samples.....	66
Table 5.1: Average Network Statistics for the Two Methods.....	81
Table 5.2: Basic Network Statistics for Three Time Slices	98
Table 5.3: Reciprocity, Transitivity, and Assortativity.....	100
Table 6.1: Conversations Across Time	113
Table 6.2: Median Conversation Tempos and Lengths.....	117
Table 6.3: Network Statistics for Time Period Networks	119
Table 6.4: Types of Conversations by Number of Users	123

LIST OF FIGURES

Figure 3.1: Twitter and OSM along the two axes	37
Figure 3.2: Retweeting, Twitter Replying, and Crisis Mapping along the two axes	38
Figure 4.1: Path of Hurricane Sandy provided by the US National Weather Service.	46
Figure 4.2: Bounding Box for Data Filtering.....	49
Figure 4.3: Loglog of out-degree distribution with linear fit to the tail.....	56
Figure 4.4: Hub percentage for each network.....	57
Figure 4.5: Average neighbor in-degree vs. node out-degree.....	58
Figure 4.6: Percentage of source tweets in geo area of interest.....	62
Figure 4.7: Retweet count distributions, A.....	63
Figure 4.8: Retweet count distributions, B.....	64
Figure 4.9: Fraction of tweets with 1-10, 11-80, and >80 retweets	65
Figure 5.1: Port-Au-Prince, Haiti in OpenStreetMap. Before and 4 days after the earthquake.....	74
Figure 5.2: Four users with overlapping changesets mapping distinct objects.....	80
Figure 5.3: Three mappers extending each other's roads.....	80
Figure 5.4: Mapping around Salient Geographic Features	85
Figure 5.5: Filling in the Map	88
Figure 5.6: The Parallel Mapping of Two Mappers.....	90
Figure 5.7: During Degree Distribution	99
Figure 5.8: After and Long After Degree Distributions.....	99
Figure 6.1: (a) Inter-tweet times; (b) Times between outgoing replies; (c) Inter-reply times	117
Figure 6.2: Relationship of in-strength to in-degree in <i>During</i>	121

CHAPTER 1: Introduction

1.1 Social Media in Disaster

Social media have been increasingly used by those affected by disasters arising from natural hazards, as well as global onlookers, for a variety of purposes. These include information seeking and sharing (Qu et al., 2011; Starbird & Palen, 2010), collective sense-making (Palen et al., 2009; Starbird, 2013), self-organizing around a task deemed important in the disaster situation, such as pet rescue coordination (White et al., 2014; White & Palen, 2015), critical information collation (Starbird & Palen, 2013), event reporting (Keegan et al., 2012; Keegan et al., 2013; Perng et al., 2013), and mapping the affected areas (Dittus et al., 2016; Soden & Palen, 2014). In the high-tempo (Keegan, 2012), high-volume context of disaster, this social media activity produces an immense volume of digital traces, documenting the experiences and decision-making of the public. This dissertation engages with these digital traces and, through the use of novel methods and theoretical framing, investigates the complexities of sociality that emerge across a variety of social media activities in disaster.

1.2 Emergent Forms of Sociality in Disaster

Large-scale disruptions, such as disasters arising from natural hazards, often lead to infrastructural failures and destruction of the physical environment, as well as putting immense pressure on the social fabric. Thus, disruption events are often associated with the breakdown of existing social structures (Mileti, 1999, Fritz, 1961). However, they are also ripe for the creation of new social structures through the formation of new relationships among those affected. Sociologists who study disaster have long pointed out that people directly affected by the crisis events are the true first responders who actively participate in the rescue and recovery efforts (Fritz & Mathewson, 1957; Dynes, 1970; Kendra & Wachtendorf, 2003). This active participation, together with the phenomenon of convergence (Fritz & Mathewson, 1957), where people coalesce around the physical or digital site of disruption, result in new relationships being established, new sources of information found, and new communities being formed.

Thus, in disruption the social structures are remade anew (Comfort, 1999; Petrescu-Prahova & Butts, 2005).

This makes disruption events not only important objects of study in their own right, which has been well-established in the field of crisis informatics (Hagar & Haythornthwaite, 2005; Palen et al., 2009), but also a powerful lens into formation and reconfiguration of social structures. Sociologists of disaster point out that disaster events can be seen as catalysts that bring into relief the ongoing processes of organizational structures emerging out of social activity (Kreps & Bosworth, 1994). They suggest that crisis events highlight social “structure in process,” as it emerges from activities and practices of those affected, and thus can shed light on how people self-organize and problem-solve more generally, and what organizational forms emerge in the process.

While this theoretical perspective resulted from the analysis of the offline post-disaster social activity (Kreps & Bosworth, 1994), social media platforms provide their own sets of tools, resources, and specific affordances around communication, information seeking, and collaboration as part of sense-making process. Thus, what social structures emerge in the post-disaster social media activity is an open question that is ripe for investigation. This dissertation, therefore, focuses on the forms of sociality—organizational forms, collaborative work practices, and social network structures—that arise through online activity in disaster.

1.3 Social Media Context

Social media sites provide a highly distributed and mostly asynchronous setting for the flurry of the sense-making activity that takes place in disasters arising from natural hazards. Thus, social activity in such environment faces classical issues of Computer-Supported Cooperative Work (CSCW), especially with respect to how the large numbers of participants manage to coordinate remotely. CSCW research has shown that achieving a shared information space is crucial for distributed work (Schmidt & Bannon, 1992). It is defined as achieving a common understanding of meaning among a group of people working in distributed fashion. Bertelsen and Bødker show that, in massively distributed setting, the common

information space is achieved not through collocation but through position and mobility, suggesting that the site of work matters in cooperative work (Bertelsen & Bødker, 2001).

Thus, two aspects of shared information spaces are especially important in a massively distributed setting of social media. The first relates to how well the environment facilitates peripheral awareness. This is passively collected information on the activity of others, which provides a context for one's own activity (Dourish & Bellotti, 1992). One could argue that the first step towards peripheral awareness is an understanding that others are working in the same collaborative space and thus an awareness of this shared site of work. Therefore, one way to operationalize peripheral awareness for the distributed social media context is through availability of an explicitly-shared site of work. This relates to where the collaborative work takes place and how clear it is to the participants that they are working in the same digital space: the more clearly and explicitly-shared the site of work is, the easier it is to find others who are working on the same issue. With respect to this dimension, Twitterers who are making sense of a hurricane warning are working in a decentralized environment with no single point of entry or visible common site of work. On the other hand, OpenStreetMap in the time of crisis provides a clear, shared site of work—a map of the affected area.

Moreover, the social media platforms are complex interactional environments, allowing for multiple types of interaction and coordination. Thus, various types of work might emerge within each platform. A way to differentiate the activities relates to the second prominent aspect of the shared information space—how well the environment facilitates articulation work. This is the meta work of coordination and division of labor that allows the main activity to take place (Schmidt & Bannon, 1992). Arguably, the first step towards successful articulation work is being able to see what others have done and what still needs to be accomplished. Therefore, one way to operationalize articulation work for the social media context is through an availability of the visible record of activity: the more clear and explicit the record of activity is, the more it facilitates articulation work, without requiring interpersonal communication.

These two intersecting dimensions of the shared information space—availability of explicitly-shared site of work and human-readable record of activity—then allows us to organize the vast space of social media platforms and activities within them. Moreover, they also potentially offer insight into how to analyze these digital phenomena and find points of entry into the associated data.

Thus, I position the social media activities addressed in this dissertation within this framework. Twitter does not provide a well-defined site of work, and retweeting, specifically, is an outward-facing information sharing activity. It provides the users with essentially no record of the many steps tweets may take to propagate. On the other hand, in reply activity, the back-and-forth of the collective sensemaking is visibly recorded through the @reply convention. In OpenStreetMap, crisis mapping provides a well-defined site of work (map of the affected area), but it does not produce a human-readable record of the mapping activity, since all the map changes are stored in a large, behind-the-scene geospatial database, and only the newest version of the map is visible.

Having positioned the work of this dissertation within this theoretical framework, let me briefly turn to disciplinary commitments, tensions, and syntheses that underpin this dissertation and propel it forward.

1.4 Structure vs. Agency Debate: Structuration and Complexity Science as an Answer

Historically, sociology committed to a dichotomy of one of two stances: the structural perspective where the macroscopic structures determine and constrain people's behavior, or the interactionist perspective that succeeds at understanding small group interaction while largely ignoring the structural constraints. This dichotomy resulted in the so-called “structure vs. agency” debate in social science, which is still ongoing. However, clearly neither perspective is fully satisfying, and they are not necessarily mutually exclusive.

This false dichotomy of structure vs. agency is relevant to questions in crisis informatics, because those affected are the active first responders in the disaster situation (Dynes, 1970), who actively navigate the situation based on their specific context in what Suchman calls *situated action* (Suchman, 1987). On

the other hand, they are responding within the context of macroscopic structural forces such as the class hierarchy, gender roles and expectations, and organizational structure of the formal emergency response. Moreover, through their active participation, geographically-vulnerable populations affect existing meso- and large-scale structures, such as community organizations, disaster response policies, and societal attitudes towards various kinds of hazards. Moreover, they even participate in the production of new forms of sociality through their activity. From the methodological standpoint, the large-scale quantitative analyses produce broad-stroke insights into these general dynamics, but the highly situated nature of social activity in crisis also requires close attention to context traditionally exhibited in quantitative analyses associated with the interactionist tradition (Geertz, 1973). Thus, to truly understand social behavior in crisis, there is a clear need to move beyond the structure vs. agency dichotomy.

Social theorists grappled with the structure vs. agency dilemma and attempted to bridge the two with various degree of success, including Bourdieu with his concept of the habitus, Foucault with his decentered post-structuralist view of power, and Latour with Actor Network Theory. In the end, I find Giddens' theory of structuration especially helpful as a theoretical framing, as it explicitly addresses the dichotomy and attempts to account for both structure and agency without giving primacy to either, while illuminating how they constitute each other through practice. It is then not surprising that pivotal strands of both human computer interaction and disaster sociology build on the theory of structuration: from Orlikowski's view of society and technology as mutually constitutive (Orlikowski, 1992) to organizational theorists adapting structuration to understand how emergent groups in disaster come together or fall apart (Kreps and Bosworth, 1994; Weick, 1993).

Though a useful theoretical perspective, the theory of structuration does not necessarily offer measurable and well-operationalized constructs, especially those suitable to the high-volume context of social media. Promisingly, computer science offers the complex systems perspective, which, not dissimilarly from the functionalist perspective, views societal systems as constituted by interconnected and interdependent parts, where small-scale processes can "add up" to system-level dynamics. Yet, unlike

functionalism, the complex systems perspective provides specific methods for elucidating and modeling these interdependencies and emphasizes the emergent and dynamic properties of the social systems.

This focus on emergent properties is especially suitable to the social media context in disaster, with its high-tempo, high-volume convergent activity and new social connections, practices, and organizational forms arising from the activity. This brings me to the overarching research question at the center of this dissertation:

What forms of sociality emerge in the social media activity during disasters arising from natural hazards?

1.5 Overview of the Studies

There are three studies included in the dissertation that address the main research question across various social media contexts. Study One (Chapter 4) largely consists of a re-print (Kogan et al., 2015) that explores collaborative practices, network structures, and organizational forms that emerge in the process of information propagation through retweeting by the geographically-vulnerable Twitter users and those global to the event in response to the 2012 Hurricane Sandy. Thus, this study analyzes the forms of sociality that emerge in the context of retweeting, where Twitter provides no explicitly-shared site of work and its information propagation infrastructure does not produce a human-readable record of the many steps tweets may take to propagate. Additional discussion, included at the end of the chapter, reflects on how the findings of this analysis relate to these affordances and how they can be viewed with respect to the larger notion of visibility afforded by the social media platforms to various types of activities in the high-tempo, high-volume convergent situations.

Study Two (Chapter 5) is also partially comprised of a re-print (Kogan et al., 2016). The original analysis focuses on the mapping practices of contributors to OpenStreetMap (OSM) during the response to the 2010 Haiti Earthquake. Additional analyses further explore the organizational forms that emerge in this coordination, including the social roles taken on by the contributors, such as experienced and novice mappers, as well as the network structures arising from this distributed crisis mapping. This study

analyzes the forms of sociality that emerge in the environment with a well-defined site of work (map of the affected area), but lack of the human-readable record of the activity. Additional discussion at the end of the chapter also positions the finding within the proposed theoretical framework and relates them to the affordances of OSM as a system and crisis mapping as a particular activity.

Finally, Study Three (Chapter 6) is primarily comprised by a pre-print of the work currently under review. It explores organizational forms, collaborative practices, and network structures that arise in Twitter @reply conversations of the geographically vulnerable in response to the 2012 Hurricane Sandy. Therefore, it focuses on the emergent forms of sociality that arise in online activity in the context of Twitter reply conversations, where Twitter still provides no explicitly-shared site of work, but the conversational back-and-forth of the activity is reflected in the social media record that is readily available to the participants. Additional discussion, included at the end of the chapter, highlights how the findings from this study can be positioned within the proposed framework of the two intersecting dimensions of the shared information space, as well as in the larger context of high-tempo, high-volume social media activity across affordances of various platforms and types of work.

CHAPTER 2: Background Literature

2.1 Disaster Sociology

Disasters has been historically analyzed from the variety of perspectives mostly focusing and the physical disruption and engineering challenges, but social scientists have seriously engaged with disasters arising from the natural hazards starting in the second half of the twentieth century. This interest in social aspects of disaster coincided in the US with the intensification of the Cold War and consequent concern with the psychology and social dynamics of potential mass disruption and mass casualty events (Alexander, 2002).

2.1.1 The Public as First Responders

Early sociologists of disaster surveyed accounts of responses to a wide range of natural hazards, and were surprised not to find much evidence for the stereotypical view of the panicked, erratic mob that flees the scene or responds irrationally (Fritz & Mathewson, 1957). Instead they found rational, sometimes methodical, and often even organized response by those immediately affected by the disasters arising from the natural hazards. Nevertheless, sociologists of disaster found that the pervasive stereotype of helpless or erratic victims has persisted, often perpetuated by the media coverage of the disaster events that tends to emphasize sensationalist view of crowd dynamics such as looting and other crime (Fischer III, 1998). Survey of the scientific accounts, on the other hand, highlight the fact that those affected as well as bystanders become in most cases the true first responders in disaster (Dynes, 1970). Instead of playing the role of victim, they organize search and rescue efforts and coordinating relief work (Fritz & Mathewson, 1957; Kendra & Wachtendorf, 2003).

2.1.2 Convergent Behavior

Moreover, people are often drawn to the disaster area in the phenomenon of convergence first outlined by Frit & Mathewson (Fritz & Mathewson, 1957). In this phenomenon, people, resources and messages spontaneously converge towards the affected area, driven by a variety of motivations. Fritz and

Matthewson teased out five main roles associated with these various motivations for human convergence: returnees, anxious, helpers, curious onlookers, and exploiters. The returnees are the residents of the affected area who attempting to return to their homes and re-establish their routines. The anxious are the people who indirectly affiliated with the affected area through family, friendship, or community ties. The helpers are people eager to volunteer their time, efforts, or resources to the response effort. The curious onlookers are drawn to the affected area by the sensationalist nature of disasters arising from natural hazards. And finally, the exploiters are attracted by the selfish motivations of potential personal or monetary gain. Later research added two more types of convergent behavior: supporters who gather to encourage and express gratitude to emergency workers and mourners who coalesce around lighting candles, creating memorials, and mourning the dead (Kendra & Wachtendorf, 2003).

The influx of convergent behavior increases the number of information, resources, and people to coordinate, and makes it more difficult to manage for the official response efforts. Thus, traditionally the convergent behavior has been seen as a problem to be managed by the agencies of the official disaster response and preparedness. In this context of potentially disruptive increase in flows of people, resources, and information, accurate, fast, and specific information is crucial to the successful management of the situation. Moreover, the information needs to be adapted to the needs of various groups, as people in the five main convergence roles are likely to have drastically different information needs.

2.1.3 Temporality of Disaster

Fritz elaborated on how disasters—and their social impacts—are concentrated in time and space (Fritz, 1961). Dynes further developed a theory of how specifically disasters are spatialized in terms of their impact (Dynes, 1970). He conceptualized this a series of concentric circles or zones representing different level of impact and resulting social dynamics. At the very center is an area that sustained most severe impact, surrounded by a fringe area, which also sustained a significant level of damage and disruption. These impact zones receive aid from the more distant community and regional aid zones. On its way, the aid passes through the filter zones that immediately surround the impact zones. As secondary

disaster impacts—such as fires or hazardous materials spills triggered by earthquakes and environmental pollution resulting from flooding—often propagate outward from the impact zones, they may cause a geo-spatial ‘ripple effect’ into the surrounding zones. The spatialized nature of the disasters arising from the natural hazards then requires the information to also be localized and locally relevant to be useful.

Similarly, Fritz emphasized that the disasters are also concentrated in time. Building on this insight and work of other social scientists of disaster, Powell (1954) developed his theory of disaster stages. These are macroscopic temporal constructs that describe socio-behavioral phenomena for large-scale disasters arising from the natural hazards. The theory defines seven main disaster stages, as delineated in the Table 2.1. Boundaries across phases are often fuzzy, and so the stage-related behavior can be concurrent. Nevertheless, disasters are highly temporal and high-tempo events, and so the relevant information must be timely.

2.1.4 Organizational Structure in Disaster

The active participation of those affected and the bystanders in the response effort also led to a new conceptualization of the citizenship role in disaster. Dynes (Dynes, 1970) combined the strands of

Stage 0: Pre-Disaster Social structures before impact
Stage 1: Warning Precautionary activity, consultation with own social network
Stage 2: Threat Perception of change in conditions prompting survival action
Stage 3: Impact “Holding on,” shift in recognition from individual to community affect and involvement
Stage 4: Inventory Individual taking stock, moving to collective inventory
Stage 5: Rescue Local, spontaneous, unorganized rescue efforts and first aid
Stage 6: Remedy Arrival of official organized relief
Stage 7: Recovery Individual rehabilitation, community restoration of property, organization of preventative measures

Table 2.1: Disaster Stages according to Powell (1954).

previous work on the altruistic motivations in disaster and proposed an expanded view of citizenship in crisis, suggesting that after the impact citizens are motivated to take on roles and responsibilities needed by their affected communities. This increased focus on community also led Dynes to focus on the organizational structure and how it changes to accommodate the disruption of disaster. Instead of relying on previous definitions of formal complex organizations, which are drawn from the more stable and static contexts, his theory of organizations in disaster emphasized the adaptations and dynamic nature of organizational structure. He proposed a typology consisting of four types of organizations that differ along the axes of familiarity of tasks and routines to the members of the organization and permanence of the existing organizational structure (see Table 2.2).

	Regular Tasks	Non-regular Tasks
Old Structure	Type I: Established	Type III: Extending
New Structure	Type II: Expanding	Type IV: Emergent

Table 2.2: Types of Organizations in Disaster according to Dynes (1970).

Type I—established organizations—is the most formalized type that maintains the old structure and relies of regular and familiar tasks. It represents formal complex organizations, which have established hierarchies and bureaucratic structures, and adhere to predetermined policies and procedures. These are typically formal response organizations such as fire and police departments. Type II—expanding organizations still has familiar regular tasks but expands and adapts its organizational structure to the emergency situation. These are typically volunteer organizations trained in disaster response, which rely on the core personnel and expand the ranks with volunteers in crisis. Type III—extending organizations—maintain their social structure since they exist before the emergency, but take on new tasks to accommodate the needs of the situation, such as feeding of those affected by an existing food truck company. And finally, type IV is emergent organizations, which form in response to the emergency, thus producing new organization structure and orienting towards new tasks. They tend to form when some parts of the affected population are underserved by the existing formal response or when response is not well coordinated.

Additionally, disaster sociologists focused on the emergent organizational structures and shifts in social roles that arise in the wake of disasters that result from natural hazards. Drawing from Giddens' theory of structuration (Giddens, 1984), Kreps and Bosworth see organizations as continuous enactment of organizational structure and individual social roles (role-making and role-playing), which are mutually constitutive (Kreps and Bosworth, 1994). They suggest that disruption events can be seen as catalysts that bring into relief the ongoing processes of organizational structures emerging out of social activity—role-making and role-playing. To reflect their conception of structure as process (p.78), the authors develop a structural code that includes four elements: domains, tasks, resources (human and material), and activities. Domains and tasks are the structural goals, and resources and activities are the structural means. Domains is the “collective representations of bounded units and their reasons for being”, while the tasks are “collective representations of a division of labor for the enactment of human activities.” On the other hand, resources are the “individual capacities and collective technologies of human populations,” and activities are defined as “conjoined actions of individuals and social units.” Depending on whether the group's goals or the means of activity are first to crystalize, a variety of organizational forms emerges: from formal organizations that have been extensively studied in the emergency response to the much more informal and loosely-organized forms of collective action. Moreover, four dimensions of role enactment intersect and reinforce the progression through the structural codes (all four elements are required to be manifested, in whatever order, for a full organization; less than four elements constitute social structures in process).

Weick also draws heavily on Giddens' theory of structuration (Giddens, 1984). However, he focuses on sensemaking in organizations, including the emergent and minimally structured organizations such as a troop of smokejumpers (Weick, 1993). He sees social roles and meaning as mutually constitutive in disruption. The social roles provide the structural aspect of stability and continuity, while meaning and sensemaking take on the more dynamic and emergent aspect associated with agency in social science (Weick, 1995). Weick suggests that meaning and structure—in this case social roles structure—constitute each other, but they can also destroy each other as in his examples of collapse of

sense-making in the Mann Gulch fire (Weick, 1993). Since they are so directly related, Weick suggests that it would be helpful to insert the inverse relationship: less meaning would focus people's attention on anchoring effects of structure (formal roles), and less structure (collapsing social ties and roles) would reorient people towards the individual sense-making and improvisation.

2.2 Crisis Informatics

Crisis informatics is a relatively young field of study that focuses on the use of Information and Communication Technologies (ICTs) and sociotechnical system and their adaptations in relation to mass emergencies (Hagar & Haythornthwaite, 2005; Palen et al., 2009). Many research strands in crisis informatics revolve around the issues of social computing in disaster—specifically, how people navigate and leverage the affordances of various sociotechnical systems in order to share information, collectively problem-solve and self-organize in crisis.

Similarly to early disaster sociology, crisis informatics literature also highlights the active participations of those affected in the disaster-related information sharing and collective problem solving, but now through the use and appropriation of ICTs. Moreover, researchers in crisis informatics devoted early years of that field convincing crisis response practitioners and academics alike that the information produced and shared through all this active participation, now on social media, is a valuable source of information (Hughes, Palen, Sutton, Liu, & Vieweg, 2008; Liu, Palen, Sutton, Hughes, & Vieweg, 2008; Palen, Vieweg, Liu, & Hughes, 2009; Sutton, Palen, & Shklovski, 2008; Vieweg, Palen, Liu, Hughes, & Sutton, 2008). The next section discusses some of the early crisis informatics research that highlights various aspects of this active participation.

2.2.1 Active Participation in Early CI Research

Some of the earliest crisis informatics research looked at online forums as gathering places and informational resources for information sharing and self-organization around the disaster situations. In their study of Pentalk—a community networking initiative in the 2001 UK Foot and Mouth Disease crisis—Hagar and Haythornthwaite (2005) find that this forum provided a needed channel of

communication and information exchange for the community that has been largely geographically isolated by the crisis. Foot and colleagues examined how people expressed themselves online in response to the tragedy of September 11, 2001 (Foot et al., 2005; Foot & Schneider, 2004). Palen and Liu (2007) analyzed both online and offline forms of persistent citizen communications for multiple crisis events. They found elevated levels of public participation, connected to the three new information pathways facilitated by the ICTs: communications within the public affected by the crisis, communications between those affected and others outside of the area of impact, and finally communications between the public information officer roles and the members of the public. They point out that the members of the public always played an active role in disaster response, but the new affordances of the ICTs facilitated communication and self-organizing both at a larger scale and with more visibility. For example, open source disaster management software developed by the Sahana Software Foundation, which emerged in the wake of the 2004 Indian Ocean Tsunami, facilitated active and broad public participation and self-organization as part of the response (Careem et al., 2006; Currion et al., 2007).

In their analysis of online forums dedicated to various disaster events, Palen, Hiltz, and Liu (2007) found that they create a possibility for effective gathering places online. They also found the evidence of not only active participation by the members of the public, but their readiness to take on the role of true first responders, in an ongoing attempt to meet the needs of their disrupted communities. Similarly, in the analysis of discussions on a popular online forum in response to the 2008 Sichuan earthquake, Qu et al (2008) teased four important roles played by the forum: information exchange, opinion-related communication, action-related communications, and exchange of emotional support. All these aspects of the forum reflected the active engagement of the affected public in the response efforts, as well as serving as feedback mechanism for the public and the official response. Online forums also played important role in the information exchange (Procopio & Procopio, 2007) and disaster relief such as donation coordination (Torrey, 2007) in the wake of Hurricane Katrina, which was the first major disaster arising from a natural hazard to highlight the potential of peer-to-peer communication in response to a crisis event (Macias et al., 2009; Robinson, 2009).

During the 2007 Southern California wildfires, residents of mountain communities faced a dearth of localized and locally-relevant information, as the traditional broadcast sources were heavily biased towards covering the metropolitan areas and even the local news struggled to keep up with rapid changes in situation in the mountain communities (Shklovski, Palen & Sutton, 2008; Sutton, Palen & Shklovski, 2008). This gap in the traditional information sources facilitated emergence of online forums as new gathering places: first motivated by information seeking, yet eventually evolving into a community based on the shared identity and geographic allegiance. Moreover, residents contributing to this community expressed a sense of pride and ownership in their active participation, as they saw themselves as producers of information in addition to consumers. Similarly, in their analysis of disaster-related photo sharing on Flickr Liu et al. (2008) find that the members of this emergent community also view themselves as becoming producers and not mere consumers of content through actively contributing to the visual narratives of disaster events.

Finally, in the broad survey of disaster-related websites, Hughes et al. (2008) analyze the online activity from the perspective of seven types of convergent social behavior prominent in the social science of disaster (Fritz and Matthewson, 1957; Kendra and Wachtendorf, 2003). They teased out seven main types of convergent behavior as they manifest in online environment and offered examples of online activities for each, showing that many of them are congruent with the classical sociological understanding of convergent behavior.

2.2.2 Social Media and Microblogging in Crisis

Twitter—a microblogging platform that limits posts to 140 characters (280 since the recent change)—has been one of the most widely used ICTs for sharing information in disaster. It is suited for rapid transmission of information in crisis, as its decentralized structure accommodates the convergent behavior documented above. As Twitter is geared towards the latest information, it accommodates the need for the timely information. Moreover, the Twitter social network consists of the follower and followee relationships, which can easily and fluidly shift to adjust to changing information needs. Finally,

the Twitter hashtag convention allows users to tune into the discussions around specific topics described by the hashtag text (Messina, 2007), while the retweet convention allows to effortlessly pass on existing information and express some level of trust in that information.

The affordances of Twitter are then especially attuned to the needs of information exchange in disaster, and so it is not surprising that the disaster-related posts are more often focused on information broadcast and information brokerage, in comparison to the general Twitter traffic (Hughes & Palen, 2009). Next I discuss a range of studies that have focused on how information seeking and sharing on Twitter contribute to shaping the collective discussion and information space in the disasters resulting from the natural hazards.

In their analysis of the microblogging activity in response to the 2009 Red River Valley flood, Starbird, Palen, Hughes, & Vieweg (2010) focused on identifying the mechanisms of information production, distribution, and organization. They found three main types of microblogging behavior in response to crisis: generative, synthetic, and derivative. Generative activity is the creation of the original content, which can serve as a source material for the later activity. While only 10% of the tweets in the overall sample were generative, interestingly, twitterers local to the event produced over 80% of the generative tweets. Synthetic information production takes place when information from other sources is incorporated into the tweets, synthesizing existing sources. About a quarter of all activity in the sample was synthetic in nature, which Palen later corrected to the term “syncretic” (Palen, 2014). The most common form of information production was derivative such as retweets, recommendations, and re-sourcing. Retweeting propagates information to broader audiences, while also signifying an informal endorsement of the content. On the other hand, the @mention convention engages with and provides endorsement of a Twitter user. Finally, re-sourcing—such as copying information from other sources or pointing to them—directs users to new information sources. All these mechanisms serve as informal and distributed means of information curation in the crisis environment where seeking and sharing relevant information is one of the pivotal activities that aid people in protective decision making (Morss et al., 2015) and enhance situational awareness (Cameron et al., 2012; Ireson, 2009; Johnson et al., 2011).

Social Media and Situational Awareness

Situational awareness is the human perception of the multifaceted and context-specific circumstances surrounding a crisis event, which facilitates complex information interpretation, decision-making, and possible outcome prediction. Attaining and improving situational awareness is pivotal for people affected by disaster, as crisis situations are especially complex and can have real high-stakes consequences (Johnson et al., 2011).

In their study of the 2009 Red River Valley floods and the 2009 Oklahoma Fires, Vieweg and colleagues qualitatively analyzed vast amounts of tweets to locate those containing information facilitating situational awareness, such as fire locations and flood levels (Vieweg et al., 2010). Other research relied on natural language processing to automatically classify tweets as contributing to situational awareness (Corvey et al., 2012; Verma et al., 2011), although purely machine-based approaches to this task are not always reliable. Similarly, despite the varying quality of posts and their conversational quality, Ireson (2009) found that the public forum posts in response the 2007 Sheffield floods contained relevant information contributing to situational awareness.

Researchers apply a variety of techniques to extract situational awareness from the social media posts for specific types of crises-related tasks. Some research strands combine natural language processing, machines learning, and crowdsourcing to detect and track epidemics through social media posts (Brennan et al., 2013; Chen et al., 2016; Munro, 2011; Olteanu et al., 2015). Other researchers have used the geospatial data provided within the social media posts to predict earthquake impact and resulting damage (Avvenuti et al., 2014; Earle et al., 2012; Sakaki et al., 2012). Yang and colleagues (2012) used online reports to provide early estimates of the death toll in the Japan Earthquake of 2011, which were more accurate than previously available static models. The visual data contained in the social media posts was also found helpful for guiding civil engineers in their surveying of the often-large affected areas (Dashti et al., 2014).

Another study of microblogging activity around the Red River Valley flood (and the 2009 Oklahoma fires) focuses specifically on the retweeting behavior in crisis (Starbird & Palen, 2010). They

found that despite their derivative nature, retweets can be considered an informal recommendation system for the content. From this perspective, the Twitter users local to the event and the onlookers tended to recommend different types of content: while locals more frequently retweeted specific, disaster-relevant information, the global audience more often propagated the high-level, abstract view of the event. In addition, a later study (Starbird & Palen, 2012) has shown that the patterns of information propagation through retweeting in the context of political unrest can help identify users local to the event. Another study that looked at the mechanisms of information propagation on Sina-Weibo in response to the 2010 Yushu earthquake (Qu et al., 2011) found that while users used this microblogging platform for four major purposes—situation update, opinion expression, emotional support, and calling for action—action-related messages dominated the information that was propagated by the users.

While research has shown that social media can also play a role in offering sense of normalcy and providing channels for emotional support in long-term conflicts (DeChoudhury et al., 2014; Mark et al., 2009; Mark & Semaan, 2008), the focus on actionable information outlined above is understandable in the information exchange during the high-tempo disaster events arising from the natural hazards. However, as disaster situations tend to change quickly, even the timely information crowd-sourced updates on Twitter are often necessarily incomplete. Palen, Vieweg, and Anderson (2010) point out that in this type of situation, the confirmed accuracy is not as key as the helpfulness of the information. They suggest that Twitter users are proficient at evaluating the credibility of sources and piecing together multiple sources of information. Moreover, in this study, twitterers relied on the information from both people they knew and trusted and completely new sources or perfect strangers, as “information management under emergency conditions ... become increasingly socially distributed” (Palen, Vieweg, & Anderson, 2010).

Credibility of Information

In this highly-distributed environment full of incomplete information, dissemination of intentional and unintentional rumors and misinformation is a legitimate concern. Luckily, the path of

rumor propagation can be easily tracked, and a growing strand of crisis informatics research has focused on understanding this process and how it relates to the credibility of information shared on social media. Mendoza, Poblete, and Castillo (2010) found that in the aftermath of 2010 Chile earthquake, unconfirmed rumors were questioned significantly more than confirmed information, introducing the idea of self-correcting online crowd in disasters arising from the natural hazards. In a later study (Castillo, Mendoza, & Poblete, 2011), they proposed using machine learning and natural language processing techniques to evaluate credibility of a tweet. They found that trustworthy tweets often included URLs, which is consistent with an earlier finding by Hughes and Palen (2009) that disaster-related tweets were more likely to include URLs and be retweets from media and government sources, spreading information from the confirmed credible sources. Tapia and colleagues (2011) also find that focusing on the trusted participants within a controlled online environment assuaged trust and credibility issues faced by the humanitarian and relief NGOs in disaster. Moreover, a recent study shows that the social media behavior of the “official” accounts, such as formal emergency responders, tends to slow down the flow of misinformation on social media (Andrews et al., 2016). Similarly, a study by Maddock and colleagues (2015) finds some evidence for the self-correcting nature of the online activity in crisis, showing that misinformation and correction rise and fall in tandem, though they exhibit different magnitudes and a lag between the onset of misinformation and the correction.

More computational methods have been developed to assess the credibility of social media posts. Unsupervised machine learning techniques (Xia et al., 2012) and supervised methods with human in the loop (Gupta & Kumaraguru, 2012) have shown some promise. Since there is some evidence that the social media users located closer to the physical epicenter of disaster tend to post more accurate information (Thomson & Ito, 2012), several strands of research also focus on using computational methods to determine the locality of the Twitterers based on their attributes, such as profile information and social network connections (Schlieder & Yanenko, 2010; Starbird et al., 2012).

Collective Intelligence

Such a view of self-correcting crowd in disaster is taken to the next level in the studies of collective intelligence in disaster. Social media has been shown to facilitate the phenomenon of collective intelligence, which is the process of collective problem solving by large, distributed groups (Palen et al., 2007; Vivacqua & Borges, 2010). Palen et al. (2007) point out that while those affected have always been the true first responders and based their decisions on many sources (official and personal), the logs available in ICTs make this type of distributed decision-making traceable and fully visible. They attempt to capture the distributed problem-solving that took place in the aftermath of the 2007 Virginia Tech shooting. In the intervening time between the announcement of number of casualties and the names of the victims, Virginia Tech students exchanged safety and welfare checks and other information in the process that eventually evolved into a distributed attempt to identify all the victims. In this process of pooling and verifying information in the distributed environment across multiple sites, the community has developed new norms and standards, such as fact-checking and identifying the information sources, in order to reach the consensus on who should be included in the final list of victims.

Collective intelligence in crisis has been studied across a variety of platforms. While analyzing the structure and dynamics of Wikipedia editing in crisis and other breaking news events, Keegan and colleagues demonstrate that this open online encyclopedia facilitates distributed problem-solving through information exchange as the high-tempo events develop (Keegan et al., 2013; Keegan et al., 2015). Reddit contributors in crisis participate in collective decision-making through promoting (or demoting) some information and making it more (or less) visible (Leavitt & Clark, 2014; Leavitt & Robinson, 2017). As these communities coalesced around particular tasks that they collectively deemed important in disaster, it provides an excellent transition to the topic of emergent communities and organizations that often arise in disaster context.

2.2.3 Emergent Organizations in Crisis

Self-organized groups emerge around a variety of tasks members deem important in the process of disaster response and recovery. They may emerge among those affected by the disaster in the push to rebuild the impacted community. They may also emerge among strangers not affected by the event who feel strongly about certain aspects of the relief effort. While ICTs make self-organization easier and more visible, the extent of technology use varies widely across the self-organized groups. For example, in the wake of September 11 attacks, the self-organized watercraft evacuation coalesced in response to only a radio call by the Coast Guard. Many vessels operators converged based on the radio call, while others joined in after observing ongoing evacuation. In the end, this self-organized effort resulted in orderly and well-managed evacuation of those stranded in Manhattan after the attack (Wachtendorf & Quarantelli, 2003).

In another example, people who organize around a specific cause use both online and offline tools to coordinate and accomplish their goals. White and Palen (2015) describe the coordination activity that started with a Facebook post about a stranded herd of horses and eventually evolved into gathering of expertise for their rescue both online and offline. The social media both facilitated some of the coordination efforts and provided coverage and fundraising tools for this emergent community. In other instances, individuals and groups on the ground gain visibility and signal their expertise through the use of social media (Sarcevic et al., 2012).

In other instances, communication and coordination among the members of self-organizing groups are primarily mediated by technology. White, Palen, and Anderson (2014) describe another pet-centered effort, where a Facebook group was dedicated to matching lost and found pets in the wake of the 2012 Hurricane Sandy. This community organized its pet-matching efforts around digitally produced content, such as the online flyers for each pet, delineating its distinguishing features and noting its current lost-and-found status.

In the midst of the intense humanitarian crisis following the 2010 Haiti Earthquake, crisis informatics researchers proposed a technological intervention that was intended to make the social media

reports flowing out of Haiti more actionable—Tweak the Tweet (TtT). It’s a micro-syntax for disaster tweeting that uses the hashtag symbol to highlight key pieces of information such as location, making the tweets machine-readable and thus facilitating the response effort. Several other initiatives later developed methods for extracting the location information from social media posts (Intagorn & Lerman, 2011; Sultanik & Fink, 2012). While the TtT syntax was intended to be used by the affected populations in their reports of damage and emergency needs, instead Twitter users outside of the earthquake zone appropriated the convention and started serving as translators for the impacted communities (Starbird & Palen, 2011). As the translation of actionable information into machine-readable format gained momentum, these remote translators started to recognize other participants and coordinate their efforts. This ad-hoc emergent organization developed two primary roles: translators and organizers who performed complex organizational tasks to coordinate the effort. The community grew increasingly and interconnected and eventually developed a common identity—they referred to themselves collectively as “voluntweeters.” After the Haiti response, voluntweeters formalized their efforts into a non-profit virtual volunteer-based organization called Humanity Road, which supports communities affected by disasters arising from the natural hazards through screening the social media activity for actionable information (Starbird & Palen, 2013). In that sense, Humanity Road is similar to some of the earlier digital volunteer groups, such as the Random Hacks of Kindness “barcamps” and the CrisisCommons, which were composed of “technology volunteers” with technical expertise and emergency management experience who donated their time to provide services and applications for those affected by crisis (Boehmer, 2010). Standby Task Force is another crisis-centered digital volunteer organization that focuses on crisis mapping.

The Humanitarian OpenStreetMap Team (HOT) is another emergent organization that arose in response to the Haiti earthquake. OpenStreetMap (OSM) is an open, editable, and volunteer-produced map of the world that is often called Wikipedia of maps, making it a form of volunteer geographic information (Goodchild, 2007), which can be a valuable crowd-based source of information in disaster response (Goodchild & Glennon, 2010; Heipke, 2010; Meier, 2015; Norheim-Hagtun & Meier, 2010;

Zook et al., 2010). OSM is also a global community of volunteers who contribute their time and (sometimes) expertise in order to create a completely open-source world-wide geo-spatial data. Since the OSM data often provides the most accurate and accessible maps of the disaster-affected areas, OSM rapidly rose in prominence as a computer-mediated means of contributing to the humanitarian response by improving the quality of the map. In weeks after the Haiti earthquake, a small team on the ground and hundreds of remote volunteer “crisis mappers” (Shanley et al., 2013) who contributed by tracing satellite imagery collectively built a base map that documented the affected area in great detail. They provided data extracts for response efforts and many humanitarian organizations relied on the base map in their activities. Soden and Palen (2014) delineate how HOT formalized its activities after the immediate Haiti response and established an on-the-ground citizen-based effort there. They also worked on making the OSM user interface more collaboration- and coordination-friendly, as the Haiti crisis mapping highlighted the highly convergent nature of the volunteer OSM efforts in disaster (Soden & Palen, 2014). Moreover, after the Haiti response, HOT began to promote community-based mapping efforts in other seismically active regions, such as Nepal (Soden, Budhatoki, & Palen, 2014) and other disaster-affected areas (Dittus et al., 2016; Palen et al., 2015; Soden & Palen, 2016).

2.3 Emergent Forms of Sociality

All the active participation first outlined in disaster sociology studies and further documented in increasing detail and diversity in the crisis informatics literature leads to new social arrangements in disaster. People seek out new sources of information that are helpful in a crisis situation, expanding and reconfiguring their information networks. Both those affected and distant volunteers who are motivated to help coalesce into new self-organized communities and groups in an attempt to accomplish a goal they feel strongly about in the disaster context: from delivering food to the impacted areas to remotely mapping the details of the landscape from the satellite imagery. In these emergent communities, new norms of behavior are being established to aid in coordination efforts, such as the norms and practices of

fact-checking and source-tracking for information that crystalized in the process of distributed problem-solving after the Virginia Tech shooting.

With new relationships being established, new sources of information being found, and new communities being formed, **social structures are remade anew in disaster** (Comfort, 1999; Petrescu-Prahova & Butts, 2005). Moreover, Kreps and Bosworth (1994) suggest that the social disruption caused by disasters arising from the natural hazards illuminate and bring into relief the ongoing processes of organizational structures emerging out of social activity. Thus, disasters force us to examine social structure in process. In a similar vein, crisis can be viewed as a catalyst for social processes of self-organizing (Hagar & Haythornthwaite, 2005) and establishment of new norms of behavior that facilitate and legitimate new communities (Liu et al., 2008).

In light of this view of disaster as a setting and a catalyst for the remaking of existing social structures, it befits the researchers in the crisis informatics space to examine the emergent forms of sociality that arise online in disasters arising from natural hazards. Forms of sociality can be defined in a variety of ways, but the purpose of this research investigation I limit this concept to:

- Organizational forms
- Social network structures
- Collaborative practices

2.3.1 Organizational Forms

As I've shown in the Disaster Sociology section, the organizational forms can be quite diverse, ranging from the formal organization with established hierarchies and bureaucratic procedures to the emergent organizations with ad-hoc, often egalitarian structure, and loosely organized sets of activities and practices. From the point of view of emergent forms of sociality, we are most interested here in the emergent organizations, which are defined by the "structural means" of activities and resources emerging first and the "structural ends" of domains and tasks following afterwards (Kreps & Bosworth, 1994). Weick (1993) offers a notion of a minimal organization, following Mintzberg's definition of simple

organization structure as having coordination by direct supervision, organic structure, little formalized behavior, strategy being planned at the top, and person in charge tending to formulate plans intuitively or congruently with his personality. In addition, Weick also adds series of interlocking routines to his definition of minimal organization, suggesting that routines and practices are important to group cohesion and coordination. Finally, Weick emphasizes what he calls “generic subjectivity,” where roles and rules exist to enable individuals to be interchanged with little disruption to the ongoing pattern of interaction (Weick, 1993).

The organizational theory perspective can be well-complemented by the classical social theory view of groups. In his essay “Quantitative Aspects of the Group,” Georg Simmel considers of the most abstract characteristics of a group—it’s mere size. He explores how the number of participants affects the groups processes and structural arrangements. A dyad is different from all the other types of groups because only two participants means that withdrawal of either would destroy the entire group: “A dyad depends on each of its two elements alone—in its death though not in its life: for its life it needs both, but for its death, only one.” There is no superpersonal structure in the dyad, meaning that the group does not dominate the individual with only two people involved. That also leads to an intense absorption of participants into relationship, which relies on immediate reciprocity. Moreover, with only two participants, each one is directly responsible for the collective action and cannot shift responsibility to the group as a whole. Moving to a triad makes a major qualitative difference in the dynamics: here like in any group with more than two members any participant now can be outvoted, with two members forming a coalition over one. Thus, a triad is the simplest social form where the group can dominate over each member, producing a superpersonal structure and a framework for constraining of individual participants for collective purposes. Therefore, a triad is the simplest form reflecting the tension of social life between freedom and individuality and constraints of being ruled and conformity. In addition, Simmel notes that various interpersonal dynamics could emerge in a triad: the third person serving as a mediator, or seeking advantages by pitting the other two against each other, or even dividing instilling a disagreement on purpose to gain dominance.

Simmel also emphasizes the structural differences between small and large groups. In small groups there is a direct interaction, while in larger one it has to be mediated by formal arrangements and organizational structures (or as I will soon discuss, by affordances of the systems and algorithms in social media). Because of this direct interaction, the smaller the group the greater the member involvement, even if only due to the frequency of interaction. However, this close-knit structure comes with more scrutiny and responsibility for each member. The larger groups, on the other hand, need organizational arrangements to mediate the interaction, necessarily creating differentiation of status and delegation of responsibilities. A larger group "gains its unity, which finds expression in the group organs and political notions and ideals, only at the price of a great distance between all of these structures and the individual" (Simmel, 1950).

In his seminal work "Web of Group Affiliations," Simmel explores how traditional and modern societies differ in terms of group membership (Simmel, 1955). In traditional societies, people were members of only a few groups, all of which were overlapping and even concentric: immediate family, extended family, and village in which they live, for example. In the modern societies, with the division of labor and increased travel, we become members of many different groups that do not necessarily have much overlap. And Simmel suggests that it is this position at the intersection of many social groups that facilitates our unique individuation, or what Stuart Hall calls identity differentiation (Hall & du Gay, 1996). But this very position at the intersection of many distinct groups also creates characteristic web-like structure of the social networks, with their community structure and importance of long-range links that are so prominent in modern society. And while Simmel has been criticized for not attending as much to the content of those links, his focus on the web-like structure of modern group affiliations rightfully brings our attention to the network structure of communities emerging in crisis. And as Wellman and Hampton (1999) suggest, communities are always networks of social ties.

2.3.2 Network Structures

It is impossible to discuss the web-like networks of social relationships and their structure without addressing the concept of the social capital. It was first elucidated by Bourdieu, who defined it as “the actual or potential resources which are linked to a durable network of more or less institutionalized relationships of mutual acquaintance or recognition” (Jenkins, 2014). Thus, it is advantages and benefits derived from one’s position in the position in a social network, including the types of ties maintained and the resources attached to those ties (Wellman & Wortley, 1990). Social capital is often described as two constructs: bonding and bridging. Bonding social capital is usually conceived as the social ties (and their associated resources) connecting people who are similar on some social dimensions and often are quite close. These are the strong ties; they often leverage resources of emotional support. On the other hand, the bridging capital is the social ties that connect people across groups or communities. These are weak ties, which leverage resources of new information and exposure to diverse communities.

Granovetter famously described the benefits of weak ties in acquiring new informational resources, such as finding out about available jobs (Granovetter, 1983). He describes the tie strength as “a (probably linear) combination of the amount of time, the emotional intensity, the intimacy (mutual confiding), and the reciprocal service which characterize that tie.” The stronger the tie, the more friends two people have in common; the friends of friends are also likely to become friends themselves (triadic closure). In such a tight-knit group, people benefit from each other’s emotional support through the bonding social capital. Information is also evenly distributed across the group, so it is harder to learn from your friends something you did not already know. That kind of resource is much more likely to be obtained through bridges—links that connect otherwise disparate communities and therefore confer bridging social capital. Granovetter shows that by definition all bridges must be weak ties, otherwise there would have been other ties between the two groups. Such bridges sit in the path of a lot of information (have high betweenness centrality in network science terms), and therefore are pivotal in information diffusion. Moreover, echoing Simmel’s notion of the web of group affiliation, Granovetter

points out that for bridges' existence people need to connect across multiple dimension (or groups for Simmel), such as family, work, volunteer organizations, and so on.

Other social scientists shown in empirical studies that social capital (Wellman & Hampton, 1999) is connected to the breadth of individual's weak ties. Researchers in community informatics and social computing domains embrace this conceptualization of social capital as operationalized with tie strength in computer-mediated communication. For example, many studies focused on the importance of tie strength to the outcomes of online social interactions, such as requests and question answering on Facebook (Grey et al., 2013; Ellison et al., 2013), and overall satisfaction and well-being (Burke, Marlow, & Lento, 2010; Burke, Kraut, & Marlow, 2011; Fingerman, 2009). Finally, in her study of media use and tie strength among co-located and distanced participants, Haythornthwaite finds that media use differs with the tie strength of the communicators. Specifically, pairs in strong ties maintained a greater number of relations and communicated more frequently than others. Moreover, people use different media for different purposes, but stronger ties relate to the more diverse means of communication. That is, what is communicated does not differ by media, but rather by the type (strength) of tie.

Another aspect of the network structure is the notion of reciprocity—whether social relationship or interaction is reciprocated across the social tie. For example, Facebook friendship relationship is reciprocal by definition: you cannot add someone as a friend until they confirm that they consider you a (Facebook) friend as well. On the other hand, Twitter relationship of following has directionality: you can follow someone without them following you back. In this latter case, lack of reciprocity may signify difference in social status (Ball & Newman, 2013) and roles. Social status and roles are especially important consideration in the emergent communities that arise in crisis. The minimal organizations that form in crisis (see Section 2.1.4) rely on the role structure as a stable, structural element (Weick, 1993). Weick builds on Giddens' theory of structuration, where structure and agency are mutually constitutive (see deeper discussion of structuration in Disciplinary Commitment chapter). In the context of social disruption of crisis, Weick proposes that the social roles play a part of structure, while meaning-making can be seen as a form of agency. In the tradition of structuration, then, roles and meaning are mutually

constitutive and collapse in role structure leads in collapse of sense-making, which is imposing of order and rationality on the environment or “interaction patterns that stabilize meaning by creating shared interpretive schemes” (Weick, 1993). Thus, social roles are important for anchoring and facilitating sense-making in the context of disruption and collapse of organizational order. Conversely, as they are mutually reinforcing, the processes of sense-making can also produce new role structures.

Moreover, in this perspective, frameworks or structures that facilitate meanings include not only roles, but also “rules, procedures, and configured activities” (Weick, 1993). Specifically, roles are reflected and rooted in what Weick calls interlocking routines. This brings us to collaborative practices as an important consideration as a form of sociality in emergent organizations.

2.3.3 Collaborative Work Practices

The notion of practice has become increasingly important in social science, including Giddens’ theory of structuration, according to which “[t]he basic domain of study of the social sciences... is neither the experience of the individual actor, nor the existence of any form of societal totality, but social practices ordered across space and time” (Giddens, 1984, p. 2). According to the theory of structuration, the duality of structure lies in the fact that it is recreated and sometimes newly reconstituted through social practices of the agents (individuals and groups), while at the same time affecting and reconstituting the social practices.

In her classic article, *Office Procedures as Practical Action: Models of Work and Systems Design*, Suchman juxtaposes the abstraction of office procedures and the situatedness of actual practical actions of work (Suchman, 1983). She approaches the topic from the point of view of designing office productivity software systems that would support the actual work activity, which is relevant from the point of view of practices that emerge in online socio-technical systems like the microblogging and social networking sites in disaster. Suchman notes that most of these productivity software systems are constructed around the office procedures, which are assumed to reflect the underlying structure of daily office work. She argues

that procedures do not just lie beneath the work and simply have to be “unearthed” from the messiness of everyday activity. Instead, the smooth flow of office procedures is an outcome and product of work.

These socially constructed and abstracted procedures are then used to conceptualize the design of computer systems. The systems with such built-in procedural constraints inevitably limit the possible actions available to the users, often complicating and constraining their activities and practices. The procedures are thus imposed on the users, as opposed to being the underlying structure of the actual practical activity. Moreover, such overlaying of procedures is necessarily restrictive, since according to Suchman, human activity is irreducible to the algorithmic specification (Suchman, 1983). This is because practical action is necessarily contextual and situated: “every course of action depends in essential ways upon its material and social circumstances” (Suchman, 1987, p. 50). This is especially the case in the volatile and time- and safety-critical environment of crisis, where the material and social circumstances often change quickly and dramatically, and the resulting practical actions may have life-saving or dire consequences.

Thus, instead of abstracting action away from the circumstances of practical action, Suchman suggests that truly understanding it would imply “to study how people use their circumstances to achieve intelligent action” (Suchman, 1987, p. 50). Similarly, Hollan, Hutchins, and Kirsh see “cognition in the wild” as situated in its material and social context. We use social cues and materiality of the world to supplement, direct, and “distribute” our cognition. Furthermore, cognition is not simply extended or given more memory in this distributed form. Instead, “[t]he material world also provides opportunities to reorganize the distributed cognitive system to make use of a different set of internal and external processes” (Hollan, Hutchins, & Kirsh, 2000).

In their agenda-setting article for the field of Computer-Supported Cooperative Work (CSCW), Schmidt and Bannon laid out the main features of the collaborative activities in the distributed setting of computer-mediated communication (Schmidt & Bannon, 1992). They suggest that cooperative work implies mutual interdependence of actors, and thus requires cooperation in accomplishing the tasks: “being mutually dependent *in work* means that ‘A’ relies positively on the quality and timeliness of ‘B’s

work and vice versa and should primarily be conceived of as a positive, though by no means necessarily harmonious, interdependence” (Schmidt & Bannon, 1992, p.8). In order to coordinate their activities, “cooperating workers have to *articulate* (divide, allocate, coordinate, schedule, mesh, interrelate, etc.) their distributed individual activities” (p. 8). This work of articulation, then, implies an additional overhead cost in terms of time, resources, and labor, in comparison to the individual activities. In this way, cooperative work is distinctively different from individual activity, because “cooperating workers have to articulate their distributed individual activities and, thus, must engage in activities that are extraneous to the activities that contribute directly to fashioning the product or service and meeting requirements” (p.8). Moreover, to articulate the distributed activities participants need appropriate means of communication. In order to diminish the complexity of the communication and interaction needed to successfully articulate, those means of interaction often evolve associated procedures, plans, schemes and organizational structures. And so as they evolve into mechanisms of interaction, they themselves require articulation work.

2.4 Social Media Context

Building on Giddens’ theory of structuration, Orlikowski argues that social structure and technology are mutually constitutive and reinforcing (Orlikowski, 1992). She points out that within a variety of socio-technical milieus, including computer-mediated communities of social media, there is room for agency, structure, and especially their mutually constitutive interplay (Orlikowski, 1992; Orlikowski, 2002). But if technology and social structure mutually constitute each other, then we should consider how the affordances of these socio-technical systems interrelate with the forms of sociality that emerge as result of active participation, including in disaster.

The online environment, and especially the social networking sites, provide certain affordances and impose certain constraints on their users. Social media sites and even specific activities within a site come with different physical and social affordances. To highlight this, danah boyd uses term *architecture* as both the software architecture but also the structural force constraining the users (boyd, 2010).

According to her, “[n]etworked publics are not just publics networked together, but they are publics that have been transformed by networked media, its properties, and its potential” (boyd, 2010, p. 4).

Moreover, she expands that "networked publics' affordances do not dictate participants' behavior, but they do configure the environment in a way that shapes participants' engagement" (boyd, 2010, p. 1).

Similarly, Langlois et al. view technical infrastructure as specific cultural context, which restricts and creates new opportunities. They also highlight the algorithmic black box, as a common situation where social media users are directly affected by the software algorithms that make somewhat arbitrary choices and are not transparent to the users. Such algorithms, like Facebook news feed filtering, both constrain and open new opportunities for social interaction.

Another aspect of social media affordances that is important for social interaction is its “me-centricity” (Langlois et al., 2009). The fact that most social networking sites and microblogging platforms are centered around users’ profiles, newly emphasize the importance of presentation of the self in the new context (boyd, 2010; Goffman, 2012). Moreover, the profiles and personal “walls” serve as spaces of interaction: places where people congregate, gather, and share information, opinions, and emotions. The news feed gives users a sense of a public constructed by those with whom they connect. The comments that often follow can be viewed as performance of social connection before an audience: checking in as social grooming of existing connections (boyd, 2010, Walther et al., 2011). Additionally, the persistence of the social media record is another important structural affordance of the medium, suggesting that persistent, visible, and often human-readable record of activities may be available to the users, depending on the system’s settings.

In a related vein, the scalability of the social media is another structural affordance that may affect the social activity on social networking sites. The system affords the user and his/her content potential of a tremendous visibility, but it does not guarantee it (Wellman & Hampton, 1999). The social voting mechanisms of liking and retweeting do not necessarily propagate what people want scaled, but rather what the collective chooses to amplify (boyd, 2010). Thus, while some activities, like @reply Twitter conversations may produce visible and shared record of activity, other social media activities that

rely on up-voting and amplification—like retweeting—do not necessarily produce an explicit and visible record of what has been done.

Social media sites also provide a highly distributed and mostly asynchronous setting for the flurry of activity that takes place in disaster. Thus, such activity in a highly-distributed setting faces the classical issues of CSCW. Schmidt and Bannon establish the importance of common information space in CSCW, which they define as achieving a common understanding of meaning among a group of people working in distributed fashion (Schmidt & Bannon, 1992). In addition, Bertelsen and Bødker show that, in a massively distributed setting, the common information space is achieved not through collocation but through position and mobility, suggesting that the site of work matters in cooperative work (Bertelsen & Bødker, 2001). Thus, the degree to which there is an explicitly shared site of work must affect how well the common information space is achieved and negotiated. For example, in OpenStreetMap (OSM) crisis mappers organize around an explicitly shared site of work—the map of the affected area, whereas Twitterers who are making sense of a hurricane warning are working in a decentralized environment with no single point of entry or visible common site of work, though still constituting the implicit ad-hoc public through hashtag use (Bruns, & Burgess, 2011) and network connections (Langlois et al., 2009).

Moreover, the social media platforms are complex interactional environments, allowing for multiple types of interaction and coordination. Thus, various types of work might emerge within each platform, producing various levels of peripheral awareness (Dourish & Bellotti, 1992), which Bertelsen and Bødker show is important in massively distributed information spaces (Bertelsen & Bødker, 2001). As I alluded to above, the visible record of social media activity is one such mechanism of peripheral awareness, as it allows users to keep track of what has been done and what still needs doing, who has done what, and in response to whom. The explicit record of activity, such as Wikipedia talk pages, visibly encodes the mutual dependence and division of labor. It thus facilitates the articulation work—the meta-work of coordination and division of labor that allows the primary work to take place (Schmidt & Bannon, 1992)—without requiring personal communication.

While these classical works of CSCW offer some direction, they do not explicitly operationalize a more or less successful shared information space, especially as it manifests in the highly distributed, high-volume context of social media. In the next chapter, I discuss the two dimensions of the shared information space, which are largely based on classical CSCW literature, but are more explicitly defined and operationalized for the high-tempo, high-volume activity on social media in disaster. These two intersecting dimensions allow us to characterize the affordances of specific platforms and the particular activities within those platforms in terms of how well they facilitate creation of the shared information space. These two dimensions allow us to place various social media activities within a single framework with respect to the shared information space, which then allows us to ask more insightful research questions about how these dimensions are related to the emergent social dynamics and forms of sociality that arise in conjunction to those online activities in disaster.

CHAPTER 3: Conceptual Framing and Methods

As I highlighted in Chapter 2, classical CSCW literature shows that arriving at and maintaining shared information space is paramount to the success of distributed cooperative work (Schmidt & Bannon, 1992). However, the specific dimensions of the successful shared information space are not well operationalized in the CSCW literature, especially as they manifest in the massively distributed, high-volume environment of social media. In their focus on mobility as opposed to collocation (Bertelsen & Bødker, 2001) and peripheral awareness of others working in the same environment (Dourish & Bellotti, 1992), CSCW researchers point to the importance of the site of work and how well it accommodates these needs. Thus, one salient dimension of the shared information space can be operationalized, including for the highly distributed context of social media, in terms of how explicitly shared the site of work is—how well the affordances of the system and the particular activity support users’ peripheral awareness through highlighting the shared digital environment and others working in it.

In addition, the centrality of articulation work to achieving shared information space (Schmidt & Bannon, 1992) points to the importance of affordances that facilitate informed and explicit division of labor. Thus, another dimension of the shared information space that can be defined and operationalized for the social media context is how visible the record of activity is—how well the specific activities within a platform generate an explicit, human-readable log of who has done what, when, and in response to whom. Having seen the importance of these two measurable aspects of the shared information space repeatedly in my empirical work, in this dissertation I offer them as a partial operationalization of the multi-faceted construct of the shared information space and then rely on them to organize the space of various platforms and activities, in terms of how well they support the successful creation of the shared information space.

Moreover, these two dimensions—how clear it is to the users that they are sharing the online space with others and how explicit, legible, and human-readable the record is for the particular type of activity—is likely to influence the social dynamics and cooperative arrangements that emerge in that

activity, as these affordances more or less effectively facilitate the creation of the shared information space.

3.1 Research Questions

This brings us to the overarching research question of this dissertation:

What are the emergent forms of sociality that arise through social media activity in disasters arising from natural hazards?

As discussed above, I focus on the presence of two dimensions of the shared information space in the social media context, and how that supports sociality:

- Availability of explicitly-shared site of work, and
- Availability of visible, human-readable record of work

As these two dimensions delineate some of the affordances under which the emergent forms of sociality arise, I use them as organizing themes and primary axes for this dissertation. Thus, in this dissertation I explore how these two intersecting dimensions of activities on social media platforms translate into the emergent forms of sociality that arise in disasters arising from the natural hazards. To this end, I ask the following more targeted research questions:

- What forms of sociality arise in the presence of more or less explicitly shared site of work?
- What forms of sociality arise in the presence of more or less explicit, human-readable record of work?
- What kinds of forms of sociality are produced at the different intersections of these two organizing principles?

3.2 Conceptual Framing

To answer these questions, in this dissertation I engage with two online platforms, affordances of which produce the shared site of work to varying degrees: in OpenStreetMap crisis mappers find an explicitly shared site of work—the map of the affected area, while Twitterers who are making sense of a disaster situation operate in a decentralized environment with no single point of entry or visible common

site of work. Thus, on the axis of how explicitly shared the site of work is, Twitter disaster-related activity should be placed closer to the origin, whereas OpenStreetMap crisis mapping belongs further along this axis. This conceptual mapping is illustrated in Figure 3.1.

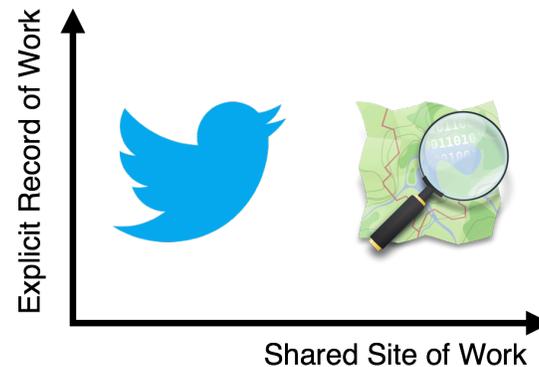


Figure 3.1: Twitter and OSM along the two axes

Moreover, as discussed in the Chapter 2, social media platforms are complex interactional environments, where multiple types of activities and interactions are possible. Thus, various types of disaster-related work might emerge within each platform, producing various levels of peripheral awareness (Dourish & Bellotti, 1992), which is important in massively distributed information spaces (Bertelsen & Bødker, 2001). Thus, in addition to the two platforms, I focus on three distinct online activities: retweeting in Twitter, @reply activity in Twitter, and crisis mapping in OpenStreetMap, which are represented in the three empirical studies that comprise this dissertation. These three activities fit into my conceptual framework of intersecting influence of the shared site of works and explicit record of work in the following way (see the visual representation in Figure 3.2).

In addition to Twitter not providing an explicit shared site of work, retweeting is an outward-facing information sharing activity, with pretty much no record immediately available to the users of the many steps tweets take to propagate. Therefore, retweeting falls into the lower left corner of the coordinate system comprised by the two axes of interest. On the other hand, in @reply activity, the back and forth of the collective sense-making is visibly recorded, placing this activity into the upper left corner of the coordinate system. While OpenStreetMap in crisis constitutes a clearly-defined shared site work—map of the affected area, crisis mapping does not produce a human-readable record of work, because all

the map changes are stored in a large, behind-the-scene geospatial database, and only the newest version of the map is visible. Thus, crisis mapping in OpenStreetMap can be placed into the lower right corner of our coordinate system. Certain activities in OSM might produce slightly more visible traces of work accomplished—such as perhaps when OSM mappers correct each other’s tags—textual labels for the map features that are crucial for its usability. Such work of correction might be more visible on the map, since the changes in the textual elements are often easier to notice than that of geospatial features. However, no OSM activity produced a human-readable log of changes that mappers can coordinate around and use for explicit division of labor, since, once again, all the changes are stored in the complex geo-spatial database. Thus, OSM does not have a type activity that would rightfully fit into the upper right corner of the coordinate system I propose.

To summarize, Study One that investigates retweeting answers the main research question of what emergent forms of sociality arise in disaster with respect to an online activity that does not provide a clearly shared site of work or the explicit record of work. Study Two, which investigates crisis mapping in OpenStreetMap, seeks to answer the primary research question in the context of online environment that provides a clear site of work, but no human-readable record of what has been accomplished. And finally, Study Three, which focuses on the Twitter @reply conversations in disaster, investigates what emergent forms of sociality arise in crisis through the activity in the highly-distributed environment with

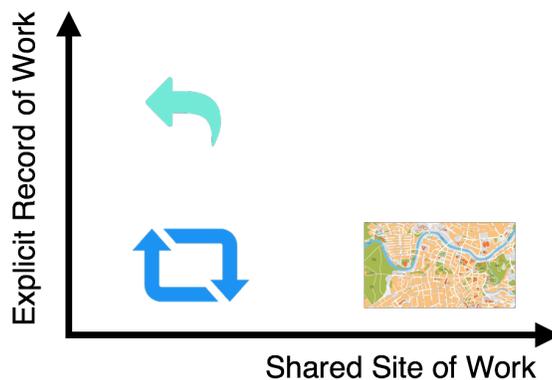


Figure 3.2: Retweeting, Twitter Replying, and Crisis Mapping along the two axes

no clear site of work, but the back-and-forth of which produces a persistent, human-readable record of the sense-making process.

3.3 Inclusion of published work

This dissertation consists of three empirical studies, two of which have been published in the human-computer interaction venues, and the third is currently under review. While I was a lead author on all the studies, they were conducted in collaboration with multiple co-authors:

- Study One, which appears as Chapter 4, is a reprint of Kogan, M., Palen, L. & Anderson, K.. Tweet Local, Retweet Global: Retweeting by the Geographically-Vulnerable during Hurricane Sandy. *Proc of CSCW*, 2015.
- Study Two, which appears as Chapter 5, is a reprint of Kogan, M., Anderson, J., Palen, L., Anderson, K., & Soden, R. Finding the Way to OSM Mapping Practices: Bounding Large Crisis Datasets for Qualitative Investigation. *Proc of CHI*, 2016.
- Study Three, which appears as Chapter 6, is a pre-print of Kogan, M. & Palen, L. Twitter Conversations Among Locals During a Hurricane: At-Scale Features of Conversational Structure in a High-Tempo, High-Stakes Microblogging Environment. Under review for CHI 2018.

All reprinted, previously published research appears in this dissertation with the permission of my co-authors, Leysia Palen, Jennings Anderson, Ken Anderson, and Robert Soden.

3.4 Methodological Considerations: Human-Centered Data Science

Since this dissertation is focused on emergent forms of sociality, theoretically it draws from Giddens' theory of structuration (Giddens, 1984), in that it views social structure and agency as mutually constitutive, where social structures arise from practices which they in turn affect. Specifically, the main thrust of this work is around how social dynamics arise from individual interaction and practices. Thus, this line of inquiry requires a methodology that is sensitive to emergent phenomena. The complexity science provides a perfect methodological lens for this investigation, since it offers formalized approaches for exploring how individual activity "adds up" to social dynamics. Specifically, this dissertation in the

large part relies on the methods of network science—a subfield of complexity science that investigates large complex networks. This is one of my methods of choice because it is well suited for analyzing social media data, which is often relational as it represents online social interactions.

Giddens’ theory of structuration, with its emphasis on structure and agency being mutually constitutive, has also exerted significant influence on the pivotal strands of HCI (Orlikowski, 1992) and disaster sociology research (Kreps and Bosworth, 1994; Weick, 1993). Specifically, organizational theorists have adopted it to understand how emergent groups in disasters arising from the natural hazards come together or fall apart, while also highlighting the centrality of social roles in this process (Kreps and Bosworth, 1994; Weick, 1993). Thus, in this dissertation I consider the effect of the social roles that are salient to the social media activity in particular contexts.

The high-tempo (Keegan, 2012; Keegan et al., 2012), high-volume convergent nature of crisis events produces vast amounts of social media data, necessitating the use of the data science methods to find the signal in the noise. However, the socio-behavioral nature of social media communication makes it context-dependent and situated (Suchman, 1987). Thus, to glean meaningful insight from the crisis-related social media activity, it is necessary to use methods that account for the complex context of the user activity. Historically, such contextual understanding has been associated with the use of “thick description” (Geertz, 1973) in qualitative methods.

Algorithmic study of online disaster-related activity is further complicated by the informal nature of social media communication. For example, from our study of Twitter activity of the geographically-vulnerable in the 2012 Hurricane Sandy, we know that the residents who were faced with the decision of whether to evacuate rarely used the word “evacuate” when describing their actions. Often, casual and even idiomatic expressions were used instead, like a resident who tweeted that he is “getting out of dodge.” Thus, searching social media only on a predefined list of keywords ignores other disaster-related user activity. For instance, less than 20% of Sandy-relevant tweets by Far Rockaway, Queens residents in our dataset were found through a keyword search, and the rest were discovered by reading through all the residents’ tweets for the period (Anderson et al., 2016).

Moreover, we know that Twitter users perceive their tweeting as a conversational, discursive activity (Palen & Anderson, 2016). Thus, it is often possible to infer the meaning of tweets only in the context of adjacent posts, as the important pieces of information are frequently spread across multiple messages. This is especially the case with the complex protective-decision making. For example, in our Hurricane Sandy research we find that we often can determine whether resident had evacuated or sheltered-in-place only retroactively, when they express regret or satisfaction with their decision after the fact (Anderson et al., 2016). Hence, careful analysis of social media activity requires treating the data conversationally, as opposed to a string of isolated posts.

Thus, in order to gain meaningful insights about emergent forms of sociality that arise in social media activity in crisis, it is necessary to develop and thoughtfully adapt methodological approaches that both harness the power of computational techniques and account for the highly situated and contextual nature of the social activity in crisis. Following Aragon, I refer to this class of methods as human-centered data science (Aragon et al., 2016). In this dissertation, human-centered data science is represented by three main approaches:

- Collecting contextual data
- Applying context-sensitive methods
- Iterating between the micro scale of individual activity and the macro scale of social dynamics

In this dissertation, I rely on datasets that retain the social context of the behavior of interest. For example, because Twitter users think of their tweeting as discursive activity, the data science that explores it should also treat these communications as such. Instead of studying isolated tweets resulting from a keyword search, we advocate to studying the entire “contextual streams” of users of interest—all their tweets within a certain time frame, including those seemingly unrelated to the keyword search, as they often contextualize and elaborate keyword-found “on-topic” tweets (Anderson et al., 2016). The methodological challenge of determining the appropriate temporal resolution for analysis of such streams could be solved with the work sessions approach pioneered for Wikipedia (Geiger & Halfaker, 2013). In the OpenStreetMap study (Chapter 5), I collect more context-imbued data by expanding the definition of

collaboration beyond co-editing the same object to broader operationalizations, which are more aligned with the computer-supported cooperative work notions of collaboration.

In addition to studying the contextual data, I apply context-sensitive computational methods as another step towards the human-centered data science suited for crisis-related social media communication. As mentioned above, one such method is network science. It allows researchers to look at the user activity from a structural perspective, based on their interactions, communications, and other cooperative activities. It enables us to tease out the kinds of structures that emerge from these activities. Simultaneously, the relational nature of the network data allows us to retain the individual connections that comprise some of the social context for the individual user activity. Consequently, Granovetter sees social network analysis as a potential way of bridging the situatedness of individual agency and the macro scale of the social structures (Granovetter, 1983).

While attending to the contextual data and utilizing context-sensitive methods are important steps towards human-centered data science, many aspects of social media activity in crisis are so complex that they also require a varied mixture of methods: quantitative for finding structural trends at the macroscopic scale of social dynamics and qualitative for understanding the rich detail of individual and small-group activity at the microscopic scale, often crucial for interpreting the big picture. Thus, in this dissertation I also contextualize the analyses by iterating between these two scales and employing mixed-methods approaches.

CHAPTER 4: Retweeting by the Geographically-Vulnerable during Hurricane Sandy

4.1 Study Overview

The content in the Sections 4.2 through 4.6 of this chapter is a reprint of Kogan, M., Palen, L. & Anderson, K. Tweet Local, Retweet Global: Retweeting by the Geographically-Vulnerable during Hurricane Sandy. In the Proceedings of the *ACM 2015 Conference Computer-Supported Cooperative Work and Social Computing (CSCW 2015)*. It is included in the dissertation with the permission of my co-authors.

This article was published in 2015 and mainly focused on the differences in the retweeting behavior between the Twitter users who were geographically-vulnerable to the Hurricane Sandy and the globally-convergent Twitter crowd. This focus on how those who are local to the disaster spread information differently from the global public was partly motivated by the earlier findings that tweets with local utility may be more important for those affected (Starbird & Palen, 2010). This focus on locality is especially important in understanding how people coordinate and collaborate online in crisis and what aspects of the social media environment support them in these activities. Specifically, since Twitter does not provide a well-defined site of work, local coordination and the activity of the global onlookers may produce two distinct, salient sites of work in disaster. So while the published paper did not explicitly feature the proposed framework of the two intersecting aspects of the shared information space, it inherently relied on the concept of the site of work in its focus. Moreover, lack of the record for the retweeting path largely informed the methods of the published paper, where network analysis was used to partially reconstruct those paths of information propagation among the local users.

The published work explores collaborative practices, network structure, and organizational forms that emerge in the process of information propagation through retweeting by the geo-vulnerable Twitter users and those global to the event in response to the 2012 Hurricane Sandy. Hurricane Sandy wrought \$6 billion in damage, took 162 lives, and displaced 776,000 people after hitting the US Eastern seaboard on October 29, 2012. Because of its massive impact, the hurricane also spurred a flurry of social media

activity, both by the population immediately affected and by the globally convergent crowd. In the published paper I explore how retweeting activity of the geographically vulnerable differs from that of the general Twitter population. I investigate whether they spread information differently, including what and whose content they chose to propagate. I investigate whether the Twitter-based relationships are preexisting or if they are newly formed because of the disaster, and if so if they persist. I find that the people in the path of the disaster favor in their retweeting locally-created tweets and those with locally-actionable information. Thus, while the importance of the site of work was not explicitly articulated in the published article, its focus on the differences between the two emergent sites—those of local and global activity—proved to be a fruitful approach for understanding social behavior in crisis.

Section 4.7 further focuses on how the two aspects of shared information space that constitute the conceptual framing of the dissertation—lack of explicitly-shared, well-defined site of work in Twitter and lack of easily traceable record of activity in retweeting—are related to the emergent forms of sociality that arise in retweeting. This section also positions the findings of this study with respect to the larger notion of visibility afforded by the social media platforms to various types of activities in the high-tempo, high-volume convergent situations.

4.2 Study Introduction

Social media are platforms for one-to-many communication that are being used by the public during disaster response for a range of purposes. A flourishing body of research on the topic has developed in the last 5 years, with a community of researchers working apace with the socio-behavioral phenomena that are advancing in kind and number with each disaster. Some research in the *crisis informatics* space examines how formal responders use or do not use social media (Blake et al., 2013; Gibbs & Holloway, 2013; Hughes & Palen, 2009; Hughes et al., 2014). Another line considers how online communication affects on-the-ground action before, during, and after disasters (Palen et al., 2009; Soden & Palen, 2014; White et al., 2014; Wulf et al., 2011).

However, the largest body of work is on the internal behaviors of the “Twittersphere” and other social media environments. Here again the research directions branch, with much work devoted to deriving data opportunistically from the social media logs using supervised and unsupervised machine learning (for example, (Imran et al., 2014; Verma et al., 2011)) as well as other techniques (Abel et al., 2012; Boersma et al., 2009; Dashti et al., 2014). The second branch of research that focuses on the social media posts as the objects of study is that which attempts to understand socio-behavioral phenomena. This includes sentiment analysis across the population represented in the streams of social media communication (Caragea et al., 2014; De Choudhury et al., 2014). It also includes analysis of self-organizing behaviors of groups that come together through social media to accomplish some task such as event reporting (Keegan, 2012; Keegan et al., 2011; Keegan et al., 2012; Perng et al., 2013) or a task they articulate as a particular need during that event (Starbird & Palen, 2011; White et al., 2014). A third branch considers how information diffuses across social networks (Bakshy et al., 2012; Coviello et al., 2014; Romero et al., 2011), particularly as the special conditions of disaster affect such diffusion (Qu et al., 2011; Starbird & Palen, 2010; Starbird et al., 2014).

4.2.1 Study Objective & Background

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

First, we investigate whether those who are geographically vulnerable spread information differently than the general public—a public that was highly active on Twitter around the world because of Hurricane Sandy’s catastrophic potential and intense media attention. In addition, when the

geographically vulnerable population retweets, what information do they spread, and what and whose tweets do they propagate if they do so at all? We investigate whether the Twitter-based relationships are preexisting or if they are newly formed during the disaster, and if those new relationships persist after the storm.

The objective of this line of inquiry is multi-fold. First, we want to continue to validate and deepen earlier findings by Starbird and Palen from the 2009 Red River Flood event that “locals” are more likely to retweet content that has “local utility” and that non-locals are more likely to retweet the “abstract” of the event (Starbird & Palen, 2010). In 2009, hardly any emergency management groups were on Twitter—but things had radically changed between 2009 and 2012 as emergency management groups tried to develop an online presence, including increasingly more police and fire departments (Hughes et al., 2014). In addition, relatively few media outlets were on Twitter in 2009, certainly as reflected in the Red River Flood data. In addition, the 2012 Sandy event had far more global attention as well as more extreme immediate consequences to a larger population than the US/Canadian Red River Flood threat of 2009—a critical event for Red River Valley locals, but incomparable to Sandy in terms of media reach. Furthermore, the number of active twitterers between 2009 and 2012 increased about seven fold (30M to 200M).¹ Finally, the act of retweeting in 2009 was only a user innovation—it was not built into the software (and therefore the metadata of tweets). We want to confirm, as well as elaborate, the earlier local utility versus

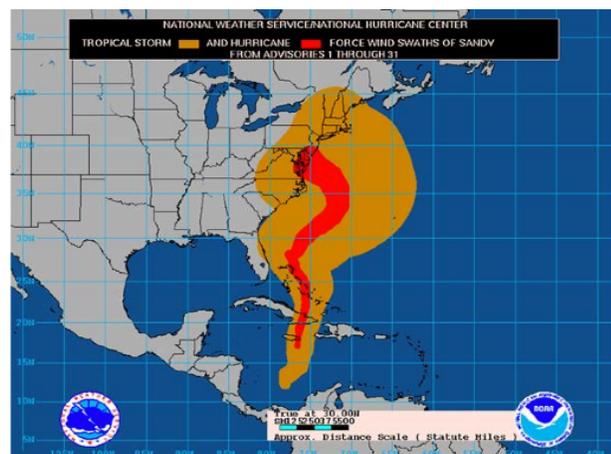


Figure 4.1: Path of Hurricane Sandy provided by the US National Weather Service via (FEMA, 2013).

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

global interest framing under these changing and far more expansive conditions.

Second, the Sandy-affected population is a desirable one to study because it is high density, where people are more likely to have enough relationships via Twitter within that geographical space to truly see what

technology-abetted socio-behavioral phenomena—and particularly information diffusion—might be. The affected population also represents high variance in socio-economic status (SES) (FEMA, 2013), though current research suggests that the affected populations in less central (and less affluent) areas might be underrepresented in the Twitter feed (Shelton et al., 2014).

Such conditions, however, bring us closer to seeing if the empirical findings of geographical behavior that social science has previously revealed are echoed in the digital world. Specifically, do new relationships between people form during disasters in the online world? We expect that they do, but we do not know to what extent those relationships persist. We know from the disaster literature that in areas that experience seasonal hazards, the affiliations and connections made in prior events often (but not always) extend to subsequent events, and to the extent to which these connections can be used as a kind of “organizational memory” (Ackerman, 1998) is exactly what those who are increasingly attached to ideas of “community resilience” (Goldstein, 2012) are banking on. This research aims to provide baseline information—as well as some longitudinal information—about an important event that will undoubtedly precede similar events in the same region.

Furthermore to this point, if we find that there are retweet behaviors that are particular to the affected population, then we may be able to use those features to derive data in future disaster events more quickly, for both scientific and applied purposes, that could aid in emergency management. As Twitter volume increases, we need more sophisticated techniques beyond keyword filtering to find the

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

Table 4.1: Dataset Names and Descriptions.

people of interest—specifically those who are needing or supplying information about the impending hazard and its effects on the social and built environment (Palen, 2014).

Third, this research is the foundation for an important new area of research in crisis informatics: that of understanding how people who are geographically vulnerable during hazards make “protective decisions”—decisions to evacuate or to “shelter in place,” which may include collecting sufficient provisions. These decisions are made as part of a “web” of sources that are available to them; they do not rely on official sources or mandates alone (Gibbs & Holloway, 2013; Lazo et al., 2009; Mileti, 1999; Morss & Hayden, 2010; Zhang et al., 2007). Because people are increasingly turning toward online sources, including social media, we must understand social media behavior as much as possible. This research is necessarily scoped to information diffusion behaviors, but it is part of a line of inquiry about “protective decision-making” by affected populations, especially as those activities vary across SES and other demographics (Zhang et al., 2007).

4.3 The 2012 Hurricane sandy event (US Landfall)

Hurricane Sandy made landfall on October 29, 2012 in southern New Jersey, affecting one of the most populated regions in the US that include New Jersey, New York, and Connecticut, with its impacts felt over a total of 24 States (FEMA, 2013). It dissipated by November 1. Prior to US landfall, it had passed through the Caribbean Sea, causing much damage to island nations before putting the entire US eastern seaboard under threat from the Gulf of Mexico and north into Canada (see Figure 4.1).

In the US, the number of deaths directly attributed to the storm totaled 162 (FEMA, 2013). It was the second costliest hurricane to hit the US; damage was estimated to be US \$6 billion (Blake et al., 2013). Approximately 776,000 people were displaced (Yonetani & Morris, 2013) and 650,000 homes were damaged or destroyed. 8.5 million people lost power as a result of the storm (Blake et al., 2013), and many were without power for weeks following.

4.4 Methods

4.4.1 Data Collection Steps

Data collection is part of our ongoing and committed effort to study disaster-related Twitter data. Using a four-node Cassandra cluster, our research group collects Twitter data 24/7 in specialized software designed for high-volume Twitter data collection (Anderson & Schram, 2011). For the Sandy event, we started collecting data using the Streaming API on October 24 2012, using the following keywords for the first round of data collection: `frankenstorm`, `hurricane`, `hurricanesandy`, `perfectstorm`, `sandy`, `sandycam`, `stormporn`, `superstorm`. This produced the *keyword data set*. In our disaster-related data collection procedures, we mitigate against potential bias in the Streaming API sampling (Morstatter, 2013) by using carefully chosen keywords and then focusing on specific subsets of users to gather a complete set of their tweets via the REST API. In particular here, we determined which users contributed to our keyword set, and then filtered again to those who had at least one geolocated tweet that fell within the geographical area of interest (see Figure 4.2). We then pulled the user streams—or what we call the *contextual streams*—for each of these *geolocated users* using the Twitter REST API. The reason we collect contextual streams is because we seek the fuller semantic context before and after a tweet that contains a found keyword. A “found” tweet based on a keyword search can have a different or enhanced meaning when one examines the surrounding tweets by the same user. Even when doing non-linguistic analysis as we do here, the full contextual streams remain important because the twitterers *themselves* presume contextual continuity across their tweets as they write them such that they do not necessarily invoke a disaster-relevant keyword each time they tweet. Therefore, our thinking is that the unit of analysis should not depend on isolated keyword-found tweets (Palen, 2014). Note that



Figure 4.2: Bounding Box for Data Filtering.

the REST API returns up to 3200 of a user’s tweets from most recent to least; this usually allows examination of a user’s behavior before the hurricane as well.

4.4.2 Creating the Data Sets

We next describe all the data sets used in this paper, with a summary of main sets in Tables 1 and 2. The original raw keyword search spans a long period of time. For this study we bounded the set to Oct 24-Nov 30 to constrain it to the first weeks of the recovery. We refer to this as the *Global Keyword* dataset (see Table 4.1). Data before October 24 preceded any reasonable predictions of where the hurricane would likely make US landfall.

Next we worked in collaboration with meteorologists, social scientists, and GIS researchers at the National Center for Atmospheric Research to define the geographic bounding box of the region most affected by Sandy. The bounded area spans a great portion of the eastern seaboard (Figure 4.2), and intentionally covers inland locations, to which people closer to the coastline were likely to evacuate.

Users who produced at least one geolocated tweet within the bounding box during the *global keyword* search time window (Oct 24-Nov 30) were designated as *geographically* (or “*geo*”) *vulnerable users*. Their tweets that were part of the keyword collection yield the *Geo-Vulnerable Keyword* dataset. All tweets in the *Global Keyword* dataset produced by the geographically vulnerable users—including their non-geolocated tweets—are part of this set. With the set of local users known, we collected their contextual streams as described above. This is called the *Geographically (“Geo”) Vulnerable Contextual* dataset.

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

Table 4.2: Retweet volume for in-common Twitterers across all time slices.

We further isolate the *Geo-Vulnerable Contextual* dataset into four distinct 5-day time slices to capture the activity of the most geo-vulnerable users *Geo-Before* (Oct 15-19), *Geo-During* (Oct 27-31), *Geo-Short-After* (Nov 8-12), and *Geo-Long-After* the event (Oct 22-26) (see Table 4.2). We choose these time frames based on our knowledge of socio-behavioral phenomena with respect to different phases of disaster events (drawing primarily from Powell (1954)). In sum, *Geo-Before* captures the time before people in the geographical area of interest could know they would be under threat but not so far beforehand that their social networks would be very different simply as a function of time. This gives us a before-disaster view of their behavior. The *Geo-During* captures the intense *high warning, evacuation and storm impact periods*, when people are making “protective decisions” and then living through the storm. The *Geo-Short-After* period captures our population as they move into the second week of *recovery*: some here could return home, many others could not, and some were able to start formulating plans for repair work. In other words, they have been able to take stock and have made initial post-disaster plans. The *Geo-Long-After* represents almost a full month after the event. Even more people have stabilized at this point. Many will have returned back to normal routines, but those who lost their homes are making other arrangements.

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

contributed retweets to the contextual dataset in all four time periods. Limiting the time slice sets to only those overlapping users resulted in smaller sets as reported in Table 4.2.

Three Retweet Counts

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

Since the technique for acquiring the retweet count requires going through all the original tweets in a dataset, these counts are necessarily timeframe- and dataset-dependent. For the retweet count distribution analysis, we were interested in the activity for the *Geo-During* time frame as it definitively includes the period when the evacuation notices were issued and the storm made landfall. We computed *three* types of retweet counts within this time frame based on the *Global Keyword* and *Geo-Vulnerable Keyword* datasets, described in Table 4.3. *Retweet counts* are named based on the relationship between *where the original tweet was sourced* and the *population it was retweeted by*.

4.5 Analyses & Findings

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

4.5.1 Geo-Vulnerable Retweet Networks for the 4 Time Slices

Time Evolution of Retweet Networks

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

Network Size & Density

The first difference in the four time sliced networks is the sheer size. Table 4.4 illustrates the size and density measures for the four networks. The *Geo-During* network is considerably larger, with more nodes and edges, than the prior and later slices. The larger size of the network corresponds to the higher volume of retweeting activity we

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

Table 4.3: Glossary of Retweet Counts.

observed during disaster, which is not so surprising considering that the underlying dataset was constructed by collecting the contextual streams of geo-vulnerable users who were found by searching on the hurricane-related keywords.

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

The relative size of the largest connected component can serve as a proxy for how densely interconnected the network is. The largest weakly connected component of the *Geo-During* network encapsulates 92.88% of all its nodes and an impressive 96.34% of the edges—a considerably larger fraction than for the other three time periods (see Table 4.4). Moreover, if we construct completely comparable “internal” versions of the four networks by excluding the original tweet authors who are not also among the 7,988 overlapping users and thus retaining only the retweets between the core users, they show a similar trend. The “internal” version of *Geo-During* has the largest weakly connected component, with the fractional size considerably larger than the internal networks for other time slices. Therefore, we see higher density in the *During* period even in the internal networks, which allow the analysis to focus solely on the retweet dynamics among the overlapping users to avoid the varying number of external original tweet authors (though this method does not control for the underlying dynamic of varying retweet volume). The higher density suggests that the retweet activity of geo-vulnerable twitterers during the disaster connects multiple subnetworks, thereby constructing a more interconnected, dense social network.

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

Degree Distributions

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

All the distributions in Figure 4.3 are long-tailed, as is common in degree distributions of real-world complex networks (Clauset et al., 2009; Lazo et al., 2009). However, some of the distributions are prone to rarer events than others. For example, the *Geo-Before* network has the highest out-degree of 181. Thus, an in-common retwitterer in this network (@BKdotNet) has retweeted tweets from 181 authors in the 5 days of interest before the event (in-degree=0). Similarly, and somewhat more impressively, a user (@LaborIrishDem) in the *Geo-During* network retweeted tweets from 192 authors during the 5 days of the event (in-degree=0). In contrast, the *Geo-Short-After* and *Geo-Long-After* networks have maximum out-degrees of 134 and 97, respectively (in-degree=0 for both). Moreover, the out-degree distributions for all the time slices are power-law-like only in the tail (out-degree>10), but not the head of the distribution. This suggests that the core users who remained active retwitterers through all the stages of the event tended to retweet tweets of many authors. This is especially true for the *Geo-During* network, where the linear fit to the loglog distribution is the

least steep of all the time slices, decreasing less rapidly and thus pushing more of the distribution's density away from the origin. Hence, the core in-common users, in aggregate, tended to retweet information from a greater variety of sources during the disaster than before or after.

The in-degree here represents how popular various original authors were in the

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

Table 4.4: Network size and density.

(Fraction of nodes stat. sig. with Chi-square=3304.55, $p < 0.0001$, fraction of edges stat. sig. with Chi-square=3664.98, $p < 0.0001$).

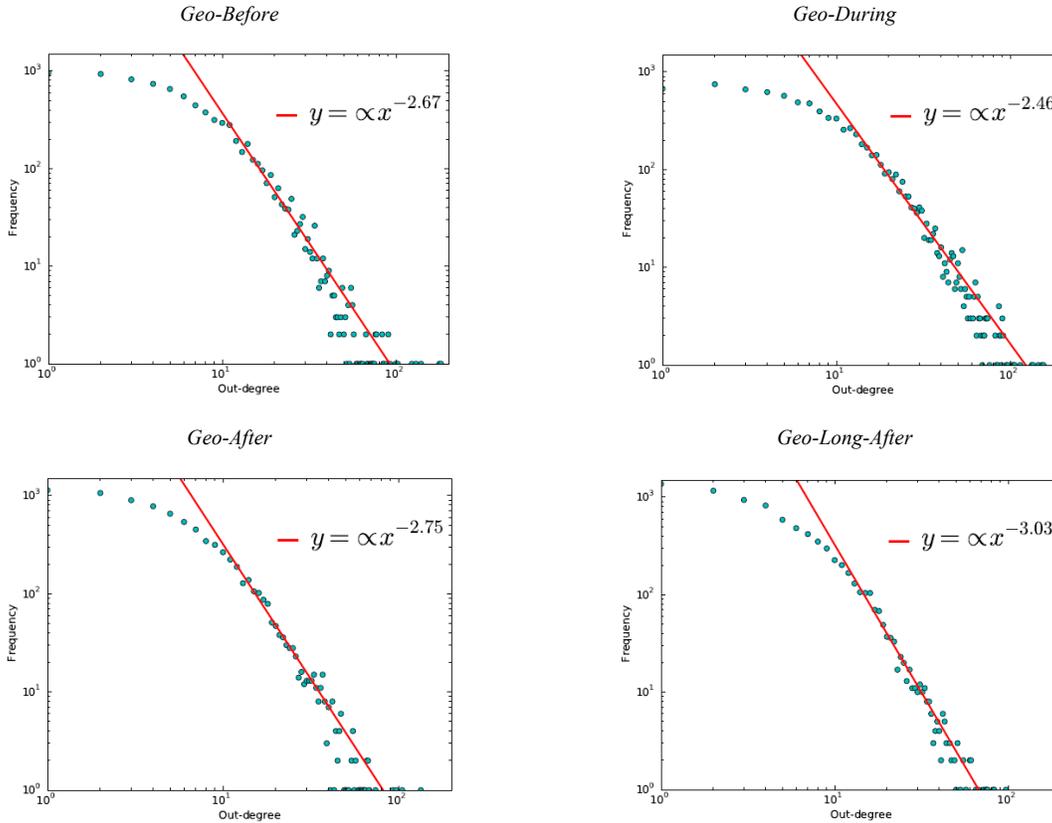


Figure 4.3: Loglog of out-degree distribution with linear fit to the tail ($x > 10$).

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

Users with high degree—many links—are often called “hubs” and are an important feature of social networks. The *Geo-During* network has more hubs based on in-degree, while the *Geo-Before*, *Geo-Short-After*, and *Geo-Long-After* networks have fewer of them. Figure 4.4 illustrates the proportion of

nodes for each network that have 100 links or more, for in- and out-degree. The differences for the out-degree are less dramatic, but the *Geo-Before*, and especially *Geo-During* have more hubs than the later time networks. This suggests that among the users who appear across all datasets, there was more hub-like activity during the disaster period—retweeting from multiple sources and being retweeted by multiple retwitterers—and thus connecting many subnetworks.

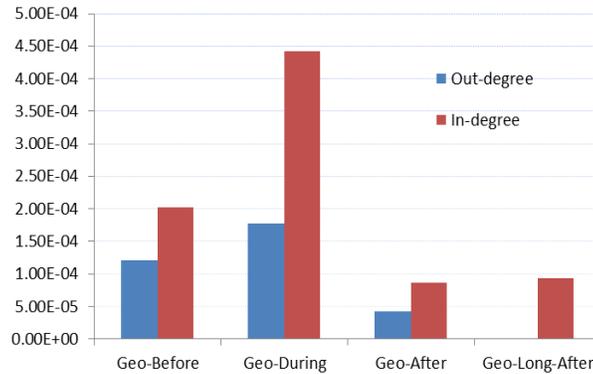


Figure 4.4: Hub percentage for each network.
 (Stat. sig. for in-degree with $\chi^2=19.53$, $p<0.001$).

Network Mixing Patterns

More insight into understanding and comparing network structure can be gleaned from “mixing patterns,” which refers to what type of nodes tend to connect to each other. “Degree assortativity” is measured by a network-level coefficient quantifying the strength of the relationship between the degree of the network’s nodes and degree of their neighbors. For our networks, which do not have data on the retweet activity of original authors unless they are also in-common retwitterers, we are most interested in out-in degree assortativity. This metric represents the correlation between the out-degree of a node and the average in-degree of its neighbors—tweets of how many authors the user retweeted and how many users, on average, retweeted the tweets of those original authors. The out-in degree assortativity for all the time sliced networks is low, signifying, on average, a weak relationship between each node’s degree and degree of its neighbors. The out-in degree assortativity of all the networks is also weakly negative, suggesting that nodes with high out-degree are on average slightly more likely to be connected to the nodes with low in-degree. This implies that for all four networks, when users retweet many sources, those sources are slightly more likely to be the ones that have not been retweeted very often by the in-common retwitterers. All the patterns above also hold if we remove all the nodes with an out-degree of one; we do

this to verify that one-off retwitterers of popular memes (high in-degree) do not produce the observed disassortativity. This disassortative relationship is weakest for the *Geo-During* network, suggesting that during disaster, the geo-vulnerable users who retweet many different sources are slightly more likely to retweet more popular sources than before or after an event. The fact that popularity plays a role in retweet activity in disaster is somewhat intuitive, as retweeting can be seen to be about trust (in either content or here in the authority of the author as confirmed by her/his popularity), and rarely is trustworthiness more important than in disaster. We will attempt to separate the part played by content utility in this process from that of author popularity later in the paper.

We can disaggregate the concept of out-in degree assortativity by looking at the average neighbor connectivity for each node (see Figure 4.5). The sub-figures illustrate the average in-degree of neighbors (y-axis) for the node with out-degree k (x-axis).

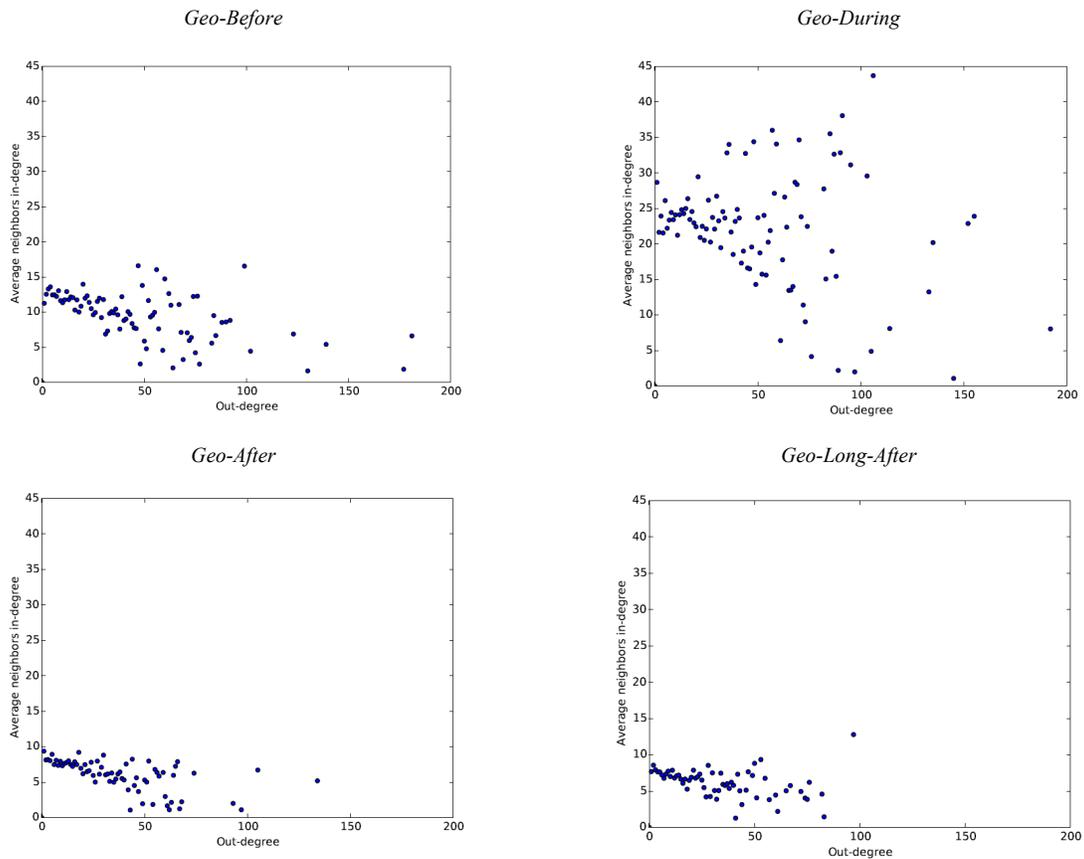


Figure 4.5: Average neighbor in-degree vs. node out-degree.

Figure 4.5 shows a less pronounced negative relationship between node out-degrees and average in-degrees of their neighbors for the *Geo-During* network, supporting our assertion of slightly more frequent hub-to-hub retweeting (out-degree hub retweeting in-degree hubs). For example, @AthenasAika—the user represented by the highest point on the *Geo-During* graph—has retweeted 106 sources who on average have been retweeted 43.70 times. This user retweeted such high-degree hubs as @BarackObama, @MikeBloomberg, and other national and local authorities and media outlets.

Transitivity

A relation is transitive if when A relates to B, and B relates to C, A also relates to C. Thus transitivity captures the idea of “a friend of a friend is a friend,” and specific to retweeting information, “a source of my source is also a source.” In the directed case, the most meaningful network-level measure of transitivity is the proportion of transitive triples derived from the triadic census as defined by Holland and Leinhardt (Holland & Leinhardt, 1970; Wasserman & Faust, 1994). We corrected this measure by excluding the few transitive triples in which a user retweets two authoritative sources that retweet each other. Thus, the corrected metric excludes the configurations that relate to authority-based broadcasting and retains only the transitive triple configurations associated with community-building retweeting, which is a typical understanding of transitivity. This corrected proportion, while rather low for all four time sliced networks, is the highest for the *Geo-During* network. The *Geo-After* has slightly higher corrected proportion of transitive triples than *Geo-Before* and *Geo-Long-After*, but still lower than *Geo-During*. Therefore, the higher proportion of transitive triples in *Geo-During*, suggests that when retweeting communities are likely to form, they are most likely to form during the event and persist to some degree immediately after (the difference is statistically significant with $\chi^2=83.29$, $p<0.0001$).

The retweet behavior displays low transitivity because retweeting need not be limited by the boundaries of users’ following networks, which are themselves low-transitivity (Lerman & Ghosh, 2010) due to their directional nature and low rates of following reciprocity (Kwak et al., 2010).

Important Users

We further explored the differences between the four time sliced retweet networks by establishing each network’s important users, with the specific focus on the influential sources (See Table 4.5).

PageRank—a variant of eigenvector centrality—has at its core a notion that the importance of a node can be judged by looking at the importance of nodes that link to it. Thus, this metric is especially well suited for finding the influential nodes based on their incoming edges, like the important original tweet authors. Based on this metric, the majority of the most important nodes of the *Geo-During* network tend to be local government authorities and the media, while those accounts are much less represented in the “most important nodes” list of the other time slices.

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

4.5.2 Geographic Patterns of Retweet Activity

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

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

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

Table 4.5: Nodes with the highest PageRank.

Table 4.6 indicates that for each of the time sliced networks, there is an overlap between the authors of the original tweets and the retweet authors. Thus, a small portion of the core retwitterers play both roles: authority figures whose tweets are propagated by others, and propagators of information.

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

This is especially the case in the time of disaster. According to Table 4.6, the percentage of geo-vulnerable source authors is rather consistent among the *Geo-Before*, *Geo-Short-After*, and *Geo-Long-After* networks (61-62%). However, in the *Geo-During* network this percentage of geo-vulnerable source authorship goes up to 68%.

Proportion of the Geo-Vulnerable Source Tweets

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

On the other hand, recall that we consider users to be geo-vulnerable if they produce at least one geotagged tweet within

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

Table 4.6: Source tweet author details.

(Fraction of geo-vulnerable sources stat. sig. with $\chi^2=23.49$, $p<0.0001$).

the boundary in the time frame of interest. This procedure produces many more tweets whose authors are considered geo-vulnerable than the actual tweets with geo-vulnerable geographical coordinates—and this larger set now includes retweets. Thus, in the rest of this analysis we consider tweets to be in the geographical area of interest if they were authored by a geo-vulnerable twitterer.

The geo-vulnerable twitterers in both the *Geo-Vulnerable Contextual* and *Geo-Vulnerable Keyword* dataset retweeted more tweets from geo-vulnerable authors than from non-geo-vulnerable. This is true for both overall time period of interest—Oct 13-Nov 30—and the *Geo-During* timeframe. For the *Geo-Vulnerable Keyword* dataset, the percentage remained essentially unchanged regardless of the time period (Figure 4.6). For *Geo-Vulnerable Contextual* dataset, however, the percentage of original tweets by geo-vulnerable authors increased noticeably during the disaster time period. This suggests that during the disaster, the Twitter conversations of geo-vulnerable authors tend to favor the local sources more strongly.

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

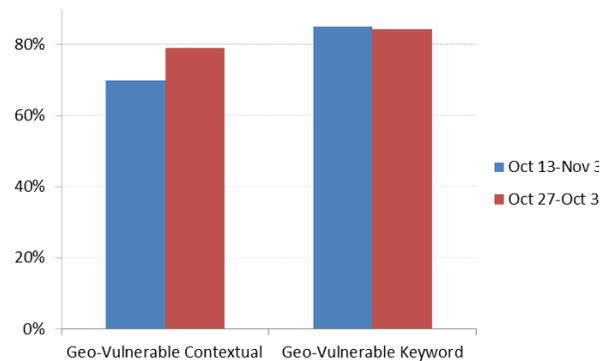


Figure 4.6: Percentage of source tweets in geo area of interest.
(Stat. sig. with $\chi^2=827.11$, $p<0.0001$)

4.5.3 Retweet Count Distributions of Various Populations

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

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

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

The distribution of *Geo-Vulnerable/Geo-Vulnerable* retweet counts on loglog scale looks significantly different from the distribution of *Global/Global* retweet count (Figure 4.7). The negative slope of the linear fit to the retweet frequencies of the former is less steep than that of the latter, producing a more heavy-tailed distribution with more density dispersing from the origin. The histogram inset in Figure 4.7 makes this more visually apparent: the *Geo-Vulnerable/Geo-Vulnerable* distribution shows more density at retweet counts greater than 10, suggesting that the geo-vulnerable tweets get retweeted 10 or more times by the geo-vulnerable users more frequently than the global tweets get retweeted 10 or more times by all users. Specifically, we can see that the tweets with 10-80 retweets are visibly over-represented in *Geo-Vulnerable/Geo-Vulnerable* distribution, compared to the *Global/Global* histogram.

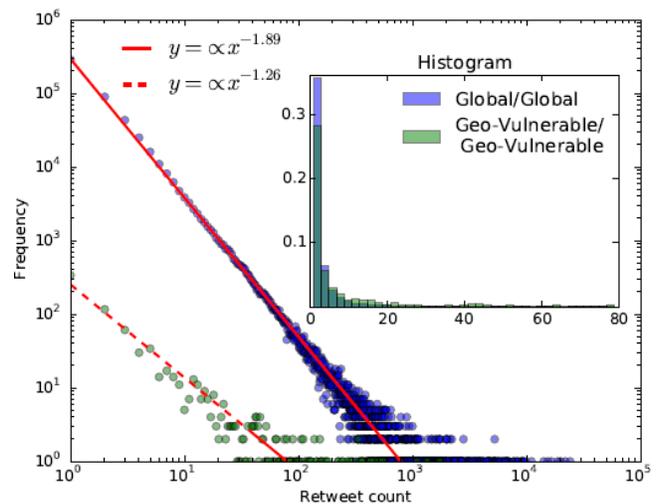


Figure 4.7: Retweet count distributions.

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

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

To move away from using location as an identifying characteristic, we can look at the distributions of geo-vulnerable users retweeting tweets from the global keyword data set, not just the geo-vulnerable tweets. Comparing the loglog distribution of *Global/Geo-Vulnerable* retweet counts to the *Global/Global* distribution, we again observe a less steep linear fit for the former, suggesting that this distribution decreases more slowly and thus produces more rare events—tweets with larger retweet counts (Figure 4.8). The fact that this distribution is more heavy-tailed is more visually obvious from the inset of Figure 4.8, where we see the *Global/Geo-Vulnerable* histogram diffuse density away from the origin and over-represent tweets with higher retweet counts in comparison to the *Global/Global*

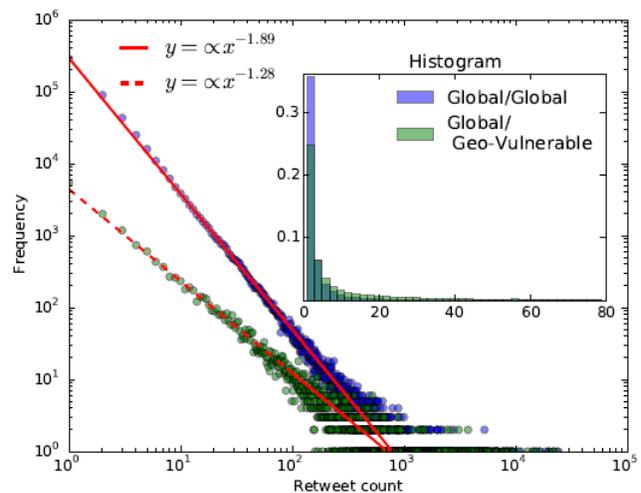


Figure 4.8: Retweet count distributions.

distribution. Specifically, the tweets with 10-80 retweets are again visibly over-represented, even to the higher degree than we observed for geo-vulnerable users propagate

In summary, for all the retweet counts, the retweet patterns of the geo-vulnerable users seem to be quite different from the global users (whether the geo-vulnerable users are retweeting the geo-vulnerable or global sources). Figure 4.9 provides the overview of these differences.

Thus, we can conclude that the tweets that end up with the higher retweet counts and hence over-represent those counts in the distributions are propagated by the geo-vulnerable users more, not necessarily because of their locality but because of some other aspect of the tweets. There are many aspects of the tweet that might motivate geo-vulnerable users to retweet it—informational utility and social conformity are two obvious contenders, especially from the perspective of retweet as an informal recommendation system (either for content or its author) (Starbird et al., 2012). Though we cannot fully disentangle these motivations in this analysis, earlier research (Starbird & Palen, 2010) suggests that local utility of the content is likely to play a role in these tweets being retweeted more by geo-vulnerable users than we would expect from the retweeting patterns of Twitter’s general population of users.

4.5.4 Content

Overrepresented tweets in the Global/Geo-Vulnerable count

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

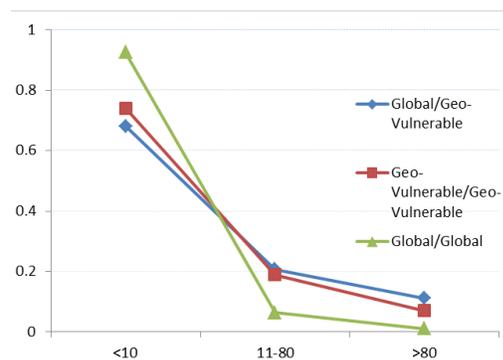


Figure 4.9: Fraction of tweets with 1-10, 11-80, and >80 retweets

(Stat. sig. with $\chi^2 = 18876.38$, $p < 0.00001$).

local utility. We focused on this retweet count because the global source of the tweets ensures that over-representation of certain tweets is not due to their geographic origin.

We divided the sample tweets into three groups: 10 tweets and below, 11-80 tweets, and 81 and up. We qualitatively coded the three samples with a simple binary flag indicating whether or not the tweet contains locally useful information. We considered the content to be locally useful if it provided practical information on the state of affairs, such as exact weather and path of the hurricane prediction, notification of road and school closures, public transportation announcements, declaration of state of emergency, concrete opportunities to help in the recovery and so on. We distinguished those specific and locally-applicable tweets from the general expression of fear, awe, and disbelief and from text or images that provide a large-scale overview of the event. Table 4.7 summarizes the findings from this content analysis.

Tweets in all three samples ranged in their content from the local utility that aids in situational awareness to the broad appeal of the bird’s eye view ‘abstract’ of the event (Starbird & Palen, 2010), including jokes and other memes. However, the 11-80 retweet subsample included much higher proportion of the tweets with locally-useful, actionable information compared to the other two samples. Clearly, the boundary cutoffs between subsamples are somewhat arbitrary, as all the samples had a varied mixture of content with local utility and broad appeal, and would be better represented by a continuum rather than discrete thresholds. However, these thresholds we empirically obtained from the retweet count distributions offer us a reasonable starting point for finding locally-useful tweets where they are most highly concentrated.

We did not find a one-to-one direct relationship between the local utility and retweet count, as evidenced by numerous locally-useful tweets we found in 1-10 retweet category. Our basic content analysis suggests that there might

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

Table 4.7: Proportion of tweets with locally useful content in Global/Geo-Vulnerable retweet samples

(Fraction with local utility stat. sig. with $\chi^2=30.11$, $p<0.0001$).

be some non-textual features that impede or promote a tweet's retweetability (Suh et al., 2010). For example, many locally-useful tweets in the 1-10 retweet sample were hard to read due to overabundance of mentions. On the other hand, the fact that retweet count distributions have considerable densities above the retweet count of 80 while the local utility decreases at this point, suggest that other factors, such as author popularity, social conformity, and imitation might be in play. The data concur: 93.58% of the original authors of tweets with more than 80 retweets are popular twitterers, as operationalized by having a thousand followers or more (a statistically significantly higher proportion than for tweets with 11-80 retweets). Qualitative analysis of over-80-retweet tweets shows that celebrities and the media have a very strong presence among authors. In future work, we plan to explore how various non-content features affect the retweet potential for the tweets with local utility, and conversely what structural features characterize the well-retweeted locally-useful and non-useful tweets.

4.6 Study Summary

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

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

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

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

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

4.7 Emergent Forms of Sociality and Two Aspects of Shared Information Space

This section offers an additional discussion of this study’s findings, as they fit into the larger context of high-tempo, high-volume social media activity. It also explores how the two intersecting aspects of shared information space—in this case lack of explicitly shared site of work in Twitter and

absence of visible record of the path of retweeting—result in particular emergent forms of sociality documented in Kogan et al. (2015).

4.7.1 Collaborative Work Practices

As discussed in Kogan et al. (2015), information propagation practices of the Twitter users local to Hurricane Sandy led to the more densely interconnected retweet network that linked previously disconnected communities in the high emergency period than before or after. In the height of the disaster the retweeting activity included a broader range of participants, as the geo-vulnerable users tended to retweet information from a greater variety of sources. The information sharing activity broadened to accommodate the geo-vulnerable users' need for new sources and channels of information. Moreover, among the users who appear across all datasets, there was more hub-like activity during the disaster period—retweeting from multiple sources and being retweeted by multiple retwitterers—and thus connecting many subnetworks. Since Twitter does not provide a well-defined and well-bounded site of work, search for new information sources is both complicated but also unbound by the affordances of the platform, often relying on the ad hoc publics of hashtags and other search mechanisms (Bruns and Burgess, 2011). Thus, this environment of loosely-defined, open-ended, and often fractured site of work likely facilitate new collaborative practices, such as the engagement of the geo-vulnerable Twitterers with broader range of information sources, producing more interconnected networks with more hubs in the height of disaster.

4.7.2 Network Structures

Under the conditions of no clear site of work and no easily traceable record of the retweeting path, information sharing activities also produce a particular type of network structure. All the four networks display low reciprocity because retweeting is not limited by the boundaries of one's following networks, and in the absence of clear record of retweeting path this directed activity does not facilitate the bi-directional visibility that leads to reciprocal ties. Similarly, low transitivity—meaning that a source of a source is rarely a source—suggests that the affordances of retweeting do not facilitate the types of

visibility in the shared information space that would lead to distinctly pro-social behavior represented by triadic closure.

4.7.3 Organizational Forms

Organizational forms, as most clearly reflected in the social roles embodied by the participants are also important forms of sociality that emerge in the process of high-tempo crisis-related activity. As shown in Kogan et al. (2015), the social role that became most prominent in retweeting activity in Hurricane Sandy had to do with the locality of the participants—whether they were geo-vulnerable to the event or comprised the global Twitter audience. These two groups shared information quite differently in the wake of Sandy. The geo-vulnerable twitterers propagated more posts from other geo-vulnerable authors than from non-geo-vulnerable. This is true for both overall time period of interest, and especially for the high-emergency period. Moreover, further analysis showed that irrespective of locality of the source, the geo-vulnerable users tend to favor in their retweeting tweets with locally actionable information, such as exact weather and hurricane path prediction, notification of road and school closures, public transportation announcements, and declaration of state of emergency. The global users, on the other hand, were more likely to propagate tweets representing the abstract of the event, such as an image of the storm system from space. In the decentralized environment of Twitter, with no clear entry point into the site of work and no easily traceable record of who propagated a post, those affected and the globally-attending Twitter public gravitated towards certain kinds of content—locally actionable and an abstract of the event, respectively—as opposed to focusing their retweeting on specific trusted users (such as their friends), to which their retweeting would likely have been largely limited in the more well-defined, but also circumscribed site of work.

CHAPTER 5: OpenStreetMap Crisis Mapping in the Wake of the 2010 Haiti Earthquake

5.1 Study Overview

The content in the Sections 5.2 through 5.6 is a reprint of Kogan, M., Anderson, J., Palen, L., Anderson, K., & Soden, R. Finding the Way to OSM Mapping Practices: Bounding Large Crisis Datasets for Qualitative Investigation. In the Proceedings of the *ACM 2016 Conference on Human Factors in Computing Systems (CHI 2016)*. It is included in the dissertation with the permission of my co-authors.

The original focus of the published paper was on understanding what collaboration looked like in OpenStreetMap. Previous network analysis research contended that mapping in OSM did not show any evidence of collaboration, but from the point of view of CSCW, with its broader conceptions of collaboration, it seemed implausible that large numbers of contributors working on the map at the same time did so in complete isolation. Thus, the original paper was focused on how the mappers collaborate and specifically how these large numbers of participants accomplish remote work when the map does not immediately show the activity of others. While the availability of the record of work did not explicitly figure as an organizing principle in the published paper, an implicit understanding that the lack of such a record in OpenStreetMap complicated collaborative crisis mapping was central to the premise of this article. Moreover, the paper's focus on the map of the affected area as an object around which the remote contributors organized relates to the importance of the well-defined site of work to the creation of the shared information space. Thus, while the proposed framework of the two intersecting aspects of the shared information space did not explicitly organize the published article, the concerns with these aspects nevertheless permeate this work.

OpenStreetMap (OSM) is the most widely used volunteer geographic information system. Although it is increasingly relied upon during humanitarian response as the most up-to-date, accurate, or accessible map of affected areas, the behavior of the mappers who contribute to it is not well understood. In the original paper, I explored the collaborative work practices of volunteer mappers operating in the high-tempo, high-volume context of disasters. To do this, I built upon and expanded prior network

analysis techniques to select high-value portions of the vast OSM data for further qualitative analysis. I then performed detailed content analysis of the identified activity and, where possible, conducted interviews with the participants. This research allowed the identification of seven distinct mapping practices that can be classified according to dimensions of time, space, and interpersonal interaction.

Section 5.7 in addition focuses on the network structures arising from the distributed crisis mapping, highlighting whether the resulting network densifies or sparsifies as the crisis mapping proceeds in response to the natural hazard with no advanced warning such as an earthquake. The additional analyses also emphasize the organizational forms that emerge in crisis mapping, including the social roles taken on by the contributors, such as experienced and novice mappers. The findings show that while the lack of human-readable record of work complicates the collaboration in OpenStreetMap, the presence of well-defined, spatially-bounded site of work—map of the affected area—still allows for the emergence of such forms of sociality as explicitly pro-social network structures. The section 4.8 places the findings of this study into a larger context of the high-tempo, high-volume distributed social media activity and highlights how the two intersecting aspects of shared information space relate to the forms of sociality that emerge within those affordances.

5.2 Study Introduction

The field of crisis informatics studies how people make use of information and communication technology in disaster. As is well documented in the disaster studies literature, people converge around the physical location of disaster (Dynes, 1970; Kendra & Wachtendorf, 2003), and now also converge on-line (Hughes et al., 2008; Meier, 2015; Palen et al., 2009).

5.2.1 The Cooperative Work of Digital Volunteerism

A prominent branch of research in crisis informatics focuses on understanding socio-behavioral phenomena in the on-line behavior of digital volunteers, examples of which include lost-and-found pet matching (White et al., 2014), critical information collation (Starbird & Palen, 2013), event reporting (Keegan, 2012; Keegan et al., 2012; Perng et al., 2013), and mapping on open source platforms (Dittus et

al., 2016; Soden & Palen, 2014). The methods by which these phenomena are studied usually end in deep qualitative investigation, and begin either by extended ethnographic study of a disaster event to identify interesting behaviors, or via a combination of quantitative analyses and content analysis to find leads in large crisis data sets.

In this paper, we examine how network analysis techniques, which are based on the relational nature of social interactions, can be used to more rapidly constrain large socio-behavioral data sets to bound inquiry (Kogan et al., 2015) and to identify sites for meaningful qualitative investigation.

We apply network analysis techniques here to examine the socio-behavioral phenomena that occur during the rapid production of *volunteered geographic information* (Elwood, 2008; Goodchild, 2007), i.e. cooperative map-making, specifically in the context of crisis. We do this for two reasons. First, “high-tempo events” (Keegan, 2012) such as disasters arising from natural hazards, create a high-speed, high-volume convergence of the digital crowd that yields a well-scoped opportunity for studying social interaction that occurs during cooperative map-making (Figure 5.1). Second, the volunteer-built maps of OpenStreetMap (OSM), our object of study, are often used by humanitarian response teams as the most up-to-date, accurate, or accessible maps of affected areas (Palen et al., 2015; Soden & Palen, 2014), and thus examining how those maps are produced has bearing on the physical sites of disaster those maps represent.

Our aim is to use network analysis techniques to scope the questioning around mapping practices in the aftermath of the 2010 Haiti earthquake, the first major disaster event that OpenStreetMap supported, where mapping practices were not yet affected by implicit and explicit institutional and technological imperatives that came to gradually govern crisis mapping (Soden & Palen, 2014). While there are a growing number of crisis events that the OSM community is responding to, we believe using this first crisis mapping dataset allows the research community to further compare mapping practices in subsequent events to origins, using the techniques we provide to help scope investigation appropriately.

5.2.2 OpenStreetMap

Since its creation in 2004, OpenStreetMap has become the most widely used example of cooperative map-making (Goodchild, 2007). Simultaneously a wiki and an open geographic data initiative, OSM provides an interactive map and open dataset

that anyone can contribute to and use. The OSM platform and the data it contains are developed and maintained by a global network of contributors that includes individual enthusiasts as well as employees of government, private sector companies, and NGOs. In 2010, a group concerned with the use of OSM in humanitarian and development contexts formed an organization called the Humanitarian OpenStreetMap Team (HOT) (Soden & Palen, 2014). HOT has since played a leading role in coordinating OSM's activities related to disaster response and preparedness.

The data underlying the map consists of *nodes*, *ways*, and *relations*, which are geographic representations of elements in the physical world (such as a water fountain, a road, or a country border, respectively) that a contributor adds to the map through one of the available map editing interfaces. These objects are interdependent upon one another such that ways are ordered series of nodes, and relations are collections of nodes and/or ways. Each object can then have associated metadata as a set of key-value tags, such as “amenity=café.” A log of each contribution to the database is recorded as a *changeset*. With such a detailed record of every edit to the map, the database contains rich information on how users create and edit these basic map elements. This edit history is our research site, as we reconstruct the social interactions and work practices that supported the construction of the map.

5.3 Research Motivation

Revealing and describing socio-behavioral phenomena in OSM is challenging. First, when viewing the map on-line, the image shown represents the accumulation of all individual edits. However,



Figure 5.1: Port-Au-Prince, Haiti in OpenStreetMap. Before the earthquake (left) and 4 days after (right), showing the rapid mapping activity. Credit: Mikel Maron (Maron, 2010).

the history of *how* the particular map came to be cannot be easily viewed. While this edit history is well documented, it is recorded at the object level. An edit to a node does not propagate to the way that contains it. If a contributor changes the shape of a building footprint by moving a node that represents the building corner, or moves a node along a road to more accurately reflect the curvature of the highway, the parent way (the building or the road), is unaware of the update. Since OSM is often called the “Wikipedia of maps” (Palen et al., 2015), the closest comparison for research on map construction is the practice of co-editing in Wikipedia, where researchers have access to “talk pages” and rich edit histories in human- and machine-readable form; unfortunately, the social interaction behind OSM is significantly less visible. Furthermore, the interdependent relationship of the OSM data structure just described makes for a much more complex edit history than Wikipedia because it must be reconstructed piece-by-piece, rather than existing as a single textual change log. In addition to the computational complexity, the reasons for modifying a map feature are not always discussed or as apparent as an addition or correction to a Wikipedia article.

Second, the OSM community supports an active wiki, mailing lists, in-person meetups, and conferences, but these discussions are not easily traceable to the immediate work of map creation. In understanding the process of map creation, therefore, it is essential to look beyond the visual representation of the map to the massive database that encodes details of who edited what and when.

Moreover, to understand *collaborative work* here we need to examine mapping practice *qualitatively*. However, that requires scoping the vast OSM edit data *quantitatively* first. Thus to understand mapping practice and collaborative mapping work, a researcher is confronted by an empirically nested problem; a problem we attempt to unravel in this paper with the use of network analysis techniques and in-depth qualitative analysis.

5.3.1 Prior OpenStreetMap Studies

As the largest open database of worldwide geographic information (Alarabi et al., 2014), OSM has been the subject of research in fields from behavioral and social sciences (Budhathoki &

Haythornthwaite, 2012) to Geographical Information Systems (GIS) research (Alarabi et al., 2014) and the GISciences (Mooney & Corcoran, 2014), however, perhaps because of the challenges mentioned above, OSM has received much less scholarly attention in HCI. Much of the existing OSM literature has focused on data quality (Haklay, 2010). Other analyses have looked at global trends of membership growth and contribution (Neis & Zipf, 2012). With some exceptions (see (McConchie, 2013; Mooney & Corcoran, 2012; Mooney & Corcoran, 2014; Palen et al., 2015)), efforts to study social interaction in OSM have not included study of the underlying database. They have instead looked at meetups (Hristova et al., 2013) and conferences (Lave & Wenger, 1991) that take place away from the map itself—the site of work.

Earlier work by Mooney and Corcoran (Mooney & Corcoran, 2012) presents the first network analysis study of OSM user interaction. They looked at co-edit networks where vertices represent users, and edges indicate that two users edited the same OSM object (any node, way, or relation) (Mooney & Corcoran, 2012). They carried out this research on OSM edits that span the entire history of their study regions. Their initial study finds a “lack of significant collaboration amongst contributors” (Mooney & Corcoran, 2012). A later study expands this research to more cities and their respective top contributors, but still uses overall aggregate OSM edit data, and thus potentially obscures interesting behaviors specific to different time periods (Mooney & Corcoran, 2014). In their “quantitative only approach,” the authors provide “statistical confirmation” of some form of recurring co-edit collaboration among top contributors in a region, but leave deeper understanding and categorization of these editing behaviors to future research (Mooney & Corcoran, 2014).

The prior studies that do categorize the work styles of OSM contributors examine them in the overall OSM population. In 2012, Neis and Zipf divided OSM contributors into Senior (≥ 1000 new nodes): 5%, Junior (10-1000): 14%, Non-Recurring mappers (< 10) 19%, and Non-Editors (0 nodes): 62% (Neis & Zipf, 2012). This quantitative definition of mapping styles was partly validated by later qualitative research by Budhatoki and Haythornthwaite (Budhatoki & Haythornthwaite, 2012), who broadly classified mappers as either “serious” or “casual.” The uneven distribution of work found in these

studies is due to most edits being made by a small portion of mappers in the general OSM population. We find that the distribution of mapper contributions is markedly less skewed in the crisis context where volunteers are highly motivated to help (Starbird & Palen, 2011).

5.3.2 An Expanded View of Map Collaboration

The prior OSM studies focus on patterns of mapper behavior in the general OSM population; we, however, are specifically interested in this socio-behavioral phenomena in the context of high-tempo crisis events. This brings to the foreground the temporal dimension of interaction, largely omitted in prior research. Moreover, the convergent nature of crisis mapping suggests that mappers participating in a response are likely to have different patterns of contribution from the general population. Thus, we expect to find different patterns of work and interaction in these events.

Additionally, we contend that the definition of collaboration is highly constrained in prior OSM research (Mooney & Corcoran, 2012; Mooney & Corcoran, 2014), and not reflective of the rich notions of collaboration developed in the CSCW and HCI communities. We propose that there are nuanced notions of map-based social interaction that are helpful to identify, and therefore focus on the range of relational activity at the site of work.

5.4 Research plan

To appropriately scope the vast map data in support of detailed qualitative analysis, we have created two basic network representations of OSM user interactions that differ by scope and attention to temporality from Mooney and Corcoran's co-editing objects (Mooney & Corcoran, 2012; Mooney & Corcoran, 2014). We constructed these networks for the first two weeks after the 2010 Haiti earthquake—the disaster event that was the first major humanitarian event for the OSM community. We think of this period as “mapping without training wheels” and, as such, it allows for examination of emergent, natural crisis mapping. This initial two-week period corresponds to what is known as the *response* phase when humanitarian aid converges to save life and limb (Powell, 1954). We used standard network analysis metrics (node degree and edge weight) on these networks to scope the dataset for qualitative inquiry.

We then employed detailed qualitative analysis of the scoped OSM data to derive the most empirically salient behaviors and types of interaction. Next, we interviewed fifteen mappers to elaborate these found mapping practices, to understand how those practices set them up relationally to other mappers, and to ensure that we understood the practices correctly. The descriptions of these mapping practices constitute the main contribution of this paper.

Simultaneously, by studying these mapping practices in detail, we built upon the initial techniques to identify a set of network metrics for faster scoping of OSM data to more readily find these behaviors—which by their complexity still need subsequent human analysis—in other situations.

5.4.1 Collection of Geospatial OSM Data

Using epic-osm, an OSM data collection and analysis tool (Anderson et al., 2015), we collected the first two weeks of edits to OSM in the Haiti region. Our data set contains edits by 429 mappers active in this time period, 85 of whom were new to OSM. Storing each individual OSM edit in a NoSQL database, epic-osm exposes an extensible query language that allows easy slicing and exporting of specific segments of data for rapid, iterative analysis with visualization and network analysis tools such as NetworkX (Hagberg et al., 2008) and d3.js (Bostock et al., 2011).

5.4.2 Initial Network Construction to Generate the Dataset for Qualitative Analysis

In all the network techniques discussed in this paper, vertices represent distinct OSM mappers; an edge denotes a specific measure of social interaction between two mappers, with a weight representing the frequency of the interaction. The interactions of the mappers over time constitute a dynamic social network, with different social relations comprising the network at different moments in time. Therefore, we use one of the canonical techniques for working with dynamic networks: discretizing the temporal network into static snapshots (Snijders, 2006). Since the mapping activity of worldwide remote volunteers overrides daily periodicity, we cannot apply existing methods to determine the appropriate time step by inferring the period (Budhathoki & Haythornthwaite, 2012). Examining the timing of contributions inductively reveals that, cumulatively, 86% of all edits occurred between 10AM and 2AM UTC, with two

distinct peaks of activity during this time window, potentially corresponding to the convergence of the European mappers (most numerous among the interviewees) and later the US and Haitian mappers. Thus, we divided the data into eight-hour time slices—2AM-10AM, 10AM-6PM, and 6PM-2AM UTC—corresponding to the period of least activity and the two peaks, respectively. Indeed, we found that an 8-hour time interval was the most empirically illuminating, as the resulting networks accumulated enough interactions to avoid the sparsity problem of smaller intervals (2, 3, 4 and 6 hours) but still retain daily patterns of interaction unlike larger intervals (12 and 24 hours).

Initial Method 1: Changeset-Level Networks

The geo-spatial bounding box associated with each changeset—the set of contributions a user makes in one editing session—can act as a geographic summary of a mapper’s contribution. Therefore, our first method of defining mapping interaction involves looking at mappers whose changeset bounding boxes overlap geographically. This network is comprised of directed edges: pointing from the mapper who contributed later to the mapper whose changeset was overlapped. An overlap in changeset bounding boxes implies at least one of the following:

1. Users are editing the same object (Same as (Mooney & Corcoran, 2012; Mooney & Corcoran, 2014)).
2. Users are editing objects that are connected, such as a bridge on a road.
3. Users are editing objects that are geographically intertwined, like roads and houses in a neighborhood.

It is important to note that the geographical footprint of changesets varies dramatically. Since a changeset is a collection of edits, the bounding box depends entirely on the number of edits and their geographic size. If a mapper edits a large way that represents a country’s border, then that changeset’s bounding box will include that entire country. Furthermore, automated bots performing basic map maintenance can create large changesets with many overlaps. To account for this, we only included changesets with an area less than 45km² in our analysis. We identified this size as the 75th percentile, and a clear break in the changeset size distribution in our data.

This particular metric for on-the-map interaction casts a wider net than just mappers editing the same object as offered by Mooney and Corcoran (Mooney & Corcoran, 2012; Mooney & Corcoran, 2014) since we posit that users editing near each other is also a form of social interaction. For example, Figure 5.2 shows map edits that were discovered through overlapping changesets that have edits to entirely distinct objects. The mappers colorized in brown and green were mapping buildings while the mappers in blue and pink were editing roads, representing a clear division of labor. These edits all happened within 5 hours, with the closest overlaps at just 25 seconds apart (the green building mapper and the pink road mapper). The overlapping changeset network for these edits is considerably more dense than co-editing objects graph produced by Mooney and Corcoran’s technique (Mooney & Corcoran, 2012), which contains many disconnected nodes. They represent mappers who were editing near each other, but not the same object.

The changeset-level definition is an all-inclusive approach to identifying interaction of any sort, including the type of division of labor seen in Figure 5.2. However, in casting such a wide net, this approach may also find instances of collisions where users are mapping over one-another (Palen et al., 2015). These collisions were especially prevalent in early crisis-mapping events (Palen et al., 2015) and add noise to the resulting interaction networks. We therefore provide a technique for representing more



Figure 5.2: Four users with overlapping changesets mapping distinct objects—example of interactions picked up by Method 1.

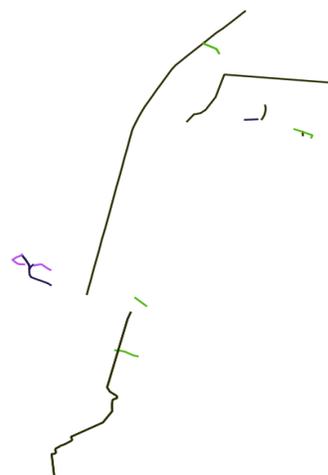


Figure 5.3: Three mappers extending each other’s roads—example of interactions picked up by Method 2

2

specific social interaction by looking deeper into the data and focusing on specific objects.

Initial Method 2: Road-Level Network

When considering objects in the OSM database that represent direct interaction and exclude collisions, we must start with a map object which is often edited and built upon: *a road*. With such a rich dataset, we can define a very particular form of co-editing that may only be interpreted as direct interaction, and not any form of collision.

To do this, we constructed a network in which vertices represent mappers who created (not edited) at least one road. An edge then exists between mappers if their roads share one and only one node. This edge is directed from the second mapper to the first, implying that the second user created a road which intersects with the road that the first user created. Figure 5.3 shows the case of intersecting roads, colored by mapper. The resulting interaction network is considerably less dense than the overlapping changeset network based on Figure 5.2. As Table 5.1 shows, on average Method 2 networks were less dense and connected than Method 1 networks, due to the more selective type of interaction depicted.

5.4.3 Network-Based Data Scoping and Qualitative Analysis

We used the two network methods above to find ten mappers with the highest degree—that is, many interactions of the specific kind with other mappers—and ten mappers with lowest degree—that is, a few interactions with other mappers (but enough edits to have an opportunity for interaction) for a total of forty mappers. We also look for frequent interactions between mappers (high weight edges—top 20%) and interactions between two mappers that persisted over time (more than three time slices); this identified an additional eight mappers for study.

We then investigated each case by hand qualitatively in detail, to gather as much information about the mappers and patterns of practice as possible. This includes reading the

Network Statistic	Method 1	Method 2
Density	0.0510	0.0073
Fractional Largest Component Size	0.5738	0.1360
Clustering	0.3207	0.0242

Table 5.1: Average Network Statistics for the Two Methods

OSM profiles and personal pages of the mappers when available. We sent out interview invites to these forty-eight OSM contributors: the mappers with the top highest and lowest degrees and the mappers with whom they frequently interacted on the map. We conducted semi-structured email and Skype interviews with the fifteen mappers who responded, four of whom were new mappers in the Haiti response, to further elaborate and validate our findings. Through this analysis, we found prominent patterns of work, and then refined techniques to further isolate them.

5.5 Mapping practices & the network techniques for isolating them

We distilled mapping practices to three main themes: *temporal*, *spatial*, and *interpersonal patterns of practice*.

5.5.1 Temporal Patterns of Practice

Temporality is a salient feature in crisis work, as the speed of data production has direct consequences on the disaster response. We describe two temporally-based practices.

Temporal Practice 1: Shift Work

Through qualitative analysis, we discovered that some mappers show a clear affinity to “shift work”— they more or less consistently participated in the Haiti mapping at regular times of the day, suggesting that they were fitting in their mapping activities to accommodate existing obligations in their schedules.

To further isolate this behavior, we calculated the frequency with which each of the 48 mappers appeared in *low-activity* (2AM-10AM UTC), *peak 1* (10AM-6PM), and *peak 2 slice* (6PM-2AM) networks. For example, if a user appeared in ten time slices, and six of those were low-activity slices, and the remaining four were peak 1, then this mapper’s low-activity frequency would be 60% and the peak 1 40% (with peak 2 at 0%). Then we selected the users with at least 75% of their activity falling into the same part-of-day time slices.

We found many instances of shift-work mapping. For example, a German mapper contributed to the map exclusively during the same 8-hour part of the day. Moreover, all of his mapping activity is concentrated between 9PM and midnight in his local time zone. There was no consistency in what he mapped: roads, major highways, rivers, and stretches of coast. Similarly, we found no clear pattern in the users near whom the German mapper edited. Nevertheless, out of the five time slices in which he was active, the user mapped near three other mappers in two instances each. None of the three nearby mappers were shift mappers.

A Canadian mapper described why he often mapped in the same hours of the day—he is a hobbyist mapper who maps after work while watching his kids. Thus, in this case, the temporal patterns of work are not only related to work schedules, but also are determined by other demands.

In yet another example of shift work, a Swedish mapper had previously established a bi-weekly schedule for mapping his local area in the months before the earthquake. After the disaster, he re-routed his efforts and time towards mapping Haiti, while still maintaining his established mapping schedule. Moreover, this mapper consistently contributed to the map in the same 8-hour part of the day, similar to the German shift worker.

Temporal Practice 2: Sustained Convergence

Another characteristic behavior is when mappers converge onto the scene and have no apparent shift work beyond perhaps sleep cycles. These are the classic post-disaster convergers (Fechner et al., 2015) who, for a period, will alter their schedules to accommodate the response. This sustained convergence activity explains the behavior of most of the other Haiti mappers whose work did not fit the practice of shift work.

To further isolate this convergence behavior, we simply reversed the filtering described for the shift work temporal practice to select mappers whose activity fell into the same time of day less than 75% of the time. A mapper from the US is a good example of sustained convergence. As a consultant, his working hours were flexible, especially when business was slow. This allowed him to map whenever he

had time across the span of the day. This flexibility allows mappers to adapt and react to the needs of the response, while also playing to their areas of expertise (which has been shown to be useful in other geowikis (Priedhorsky & Terveen, 2008)) and interest, as we see explained in the Spatial and Interpersonal patterns below. One mapper explains:

“I find that at different stages of disaster response, the responding [relief & humanitarian] organizations ask for different kinds of mapping data. Some of those requests are more down my alley and some aren't; I think this is key to making crowd data gathering work: allowing each person to find requests for the particular kind of work they're interested in doing.”

5.5.2 Spatial Patterns of Practice

The spatial characteristics of a mapper's contributions are another major organizing theme. They explain how mapping practices relate to the geography of the area being mapped or to the various types of spatial features.

Spatial Practice 1: Mapping of Externally Salient Features

Some mappers focus their work around major geographical elements, such as major roads and highways, neighborhood boundaries, coastlines, and the airport in Port-au-Prince that would appear, without any external information, to be important to the disaster response. Whether they are important sometimes emerges later, but in the early days of a disaster, there can be little direction about what is needed, and so mappers make judgments about what seems salient.

Though the network methods do not directly account for geography, these important spatial features tend to be mapped by many different mappers. Thus, to isolate this mapping pattern we use the concentration of contributors as a proxy for finding these prominent spatial objects. We look for maximal cliques—groups where every user is connected to one another—larger than a certain threshold (in this case, five) in the overlapping changeset network, which would signify many users mapping the same geographic area within the same 8-hour time slice. This identifies the areas of mass convergence appropriate for detailed hand analysis.

We found many instances of this behavior, with Figure 5.4 offering a specific example. Nine mappers contributed to the edits represented in this map. Five of them mapped three different segments comprising the major coastal road represented by the bell-like shape near the top of the map, colored in purple, green, and blue, going from left to right. The coastal outline shown above this road was mapped by three of the five mappers. The more-or-less horizontal line in the center of the map represents a major highway and it includes contributions by four mappers. Three of these mappers also mapped the coastal road above and two contributed to the coastline. Yet others mapped local streets, including major north-south thoroughfares. Three of the users mapped a major river, with one of these focusing on mapping all the waterways in the area, which relates to a preference for continuous objects discussed below.

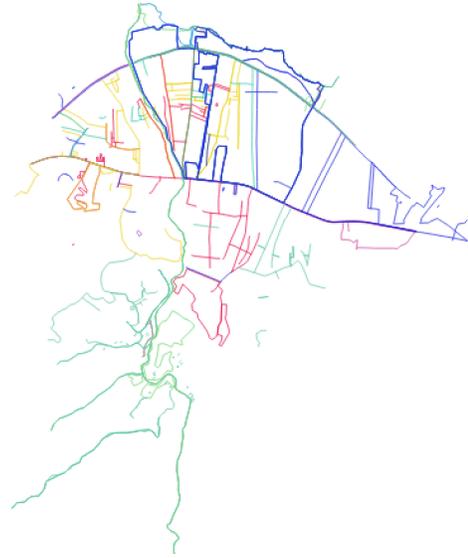


Figure 5.4: Mapping around Salient Geographic Features

The version numbers for most ways represented in Figure 5.4 are greater than one, suggesting that the participating mappers were not adding new elements, but instead editing existing ones, both cleaning up the geometry and settling on the correct tags. For example, the comments associated with all the changesets for one mapper refer to the practice of increasing data quality: “close landuse polygon,” “fixing crossing ways,” “fix unconnected ways,” and “repair duplicate nodes, ways.” We further discuss this *corrective* or *supervisory* behavior later with respect to tag correction, but it is especially salient in the context of mapping prominent geographical features, like major highways that might be important in relief efforts. A Swedish mapper explains his practice:

“Of course I looked at [the] illustration of the epicent[er] for the quake, and did as much as possible at a close distance to that. No need to map the other side of Haiti in detail.”

A US mapper mentioned that he quickly came to specialize in mapping hospitals and clinics throughout the earthquake aftermath, after beginning with a clinic in Christianville where he knew someone doing relief work. This

suggests that preference for specific objects might be a salient spatial pattern, but we did not find many other cases of that specialization otherwise. Yet, we know from other kinds of crowdworking tasks in disaster response that sometimes people will specialize: in the case of lost-and-found pet matching, for example, some volunteers specialized by breed (White et al., 2014), and our observations across many events suggest that others will, e.g., focus on creating lists of shelters before moving on to other tasks. We suspect that mapping specialization around special use locations does occur, but that the Haiti event as the first major OSM event did not allow this specialization to crystallize. It is something to which future analysts of mapping behavior should attend as humanitarian mapping practices formalize.

Spatial Practice 2: Mapping Continuous Structures

A behavior we consistently found through qualitative analysis is expanding out the map by adding roads, rivers and coastlines from one mapper to the next—what we refer to as *continuous structures*.

Once qualitatively found, we isolated this pattern of work by finding the time slices where the weakly connected component of the intersecting road network was larger than 50%—a natural threshold of many mappers (more than half) engaging in interaction. A weakly connected component implies that there is an undirected path that exists between each mapper in the component. In our networks, the weakly connected component represents a group of mappers who are building roads off of each others' work, regardless of who mapped first. The large component indicates relatively many mappers connected in this way (more than 50% of the active mappers). We used the weakly connected component here to retain as many interactions as possible, because the intersecting roads network is rather sparse.

We isolated several instances of this behavior with respect to roads. For example, we observed one instance of 9 mappers extending each other's roads over a geographical area of 220 km². These contributions are much less dense than in Figure 5.4, since they only include newly built roads, but we found it interesting because it represents the road building activity for the first 8-hour time slice on January 13—one day after the earthquake—which constitute some of the first map contributions.

Our interviewees often talked about the need to expand the map to stay out of the way of other mappers. For example, a Japanese mapper at “an early stage ... experienced some conflict in Port-au-Prince. Then [he] mapped a provincial town.” Similarly, a Canadian mapper was contributing roads mostly in the rural areas away from the crowds. He suggested that he strongly preferred mapping roads to buildings. He found mapping roads to be relaxing and even liberating, since one can just follow the road and incrementally expand the map further into space. Corcoran et al. (2013) refer to this behavior as “exploration” of the map, borrowing the term from urbanization literature where it is originally defined as an “elementary governing process” of spatial network growth (Strano et al., 2012). Roads are easy to map even when a mapper is interrupted: One can always wrap up mapping at the next intersection. Buildings on the other hand require more devoted focus because their precise and closed nature requires completion. A German mapper expressed a similar sentiment about a joy of mapping roads:

“to ride along a track with your computer mouse is fascinating ... When you can do something useful with this; it's ... double fun.”

Spatial Practice 3: Map Filling

In the data we observed many contributors who focus on filling in the map as much as possible. We find this behavior akin to completing a puzzle: users build upon each other’s work mostly by mapping adjacent areas, but intersect on major roads and neighborhood boundaries.

Once qualitatively found, we isolated this behavior by finding the time slices where the strongly connected component of the overlapping changeset network is larger than 50%. Differing from the weakly connected component above, the strongly connected component contains a directed path between every node in the component. In our network, the strongly connected component represents a group of mappers who are sequentially making edits near or overlapping each other. Large components indicate many mappers connected in this way. We used the strongly connected component here because the overlapping changeset network is very dense and can therefore afford losing some edges in return for identifying a group of mappers in the order they mapped.

The map in Figure 5.5 provides an illustrative example of filling in activity. Forty-seven users contributed to this map within a single 8-hour time slice. The area represented spans 17 km from Carrefour in the southwest to beyond Cite Soleil in the northeast with Port-au-Prince in between. Despite the large area, swaths of the map are dense with detail, which include local and major streets, highways, buildings, rivers, coastline, and public spaces such as parks, soccer fields, and city squares. Corcoran et al. (2013) refer to this behavior as “densification”, another “elementary governing process” of spatial network growth (Strano et al., 2012).

This spatial behavior also intersects with temporal patterns of mapping. All the instances of such filling in the map occurred in the first five days of the volunteer mapping effort (January 13th-17th). This suggests that this is a spatial mapping practice that is prominent early in the response, as the mappers simply strive to complete a map before other priorities are known. A British mapper explains:

“I...looked at areas that we had source data to work off and looked for areas where the apparent OSM feature ‘density’ looked to be lower compared to what was present on the source data. Then I would dive into that area and work along it. A bit like mining a coal seam.”



Figure 5.5: Filling in the Map

Many interviewees said that they selected where to map based on lack of other mapper activity. A US user mapped “further from downtown, in the shanties.” Another US mapper “zoomed in on sparse looking area[s] of [the] satellite overlay map,” and yet another “picked empty-looking areas of the map that had some imaging.” They did this to look for map sparsity that they could fill in, but some wanted to stay out of others’ way, suggesting that though it is not directly interactive, they were aware of others’ work, which leads to the final dimension: interpersonal patterns.

5.5.3 Interpersonal Patterns of Practice

Interpersonal interactions can be viewed as collaborative work that depends on the products of others’ work.

Interpersonal Practice 1: Parallel Mapping

We observed patterns of work where pairs of mappers were consistently mapping next to each other and even going back and forth on editing the map data in a contained area. We might also think of this as awareness of others’ activity in order to coordinate one’s own (Dourish & Bellotti, 1992; Heath & Luff, 1992).

We isolated this behavior by finding edge pairs with high weight and high reciprocity (top 20%) (Wasserman & Faust, 1994)—two mappers overlapping each other’s work more or less equally often—in the overlapping changeset network. Combined, these measures locate pairs of mappers who edited in the same area frequently, took turns mapping, and with roughly the same degree of reciprocity (editing each others’ work).

We found a large number of such interactions, with Figure 5.6 illustrating one of these. The entire exchange represented by this map occurred over four hours. The first user was located in Great Britain and the second was mapping from Germany. Since their time zones differed by an hour, we present their activity in UTC, which is also the local time for the first user. The mapping activity of Mapper1 is represented in solid lines and that of Mapper2 is in dotted lines. The color signifies the passing of time, with darker colors representing later time.

First, at around 22:30 UTC on January 14, 2010, Mapper1 began by contributing the solid ways in light blue on the very right of this map. Then he added some light blue solid ways at the center of right cluster around 22:45. At about 23:50, Mapper2 started adding the dotted ways in light blue at the very left edge of the map. Meanwhile, at about 20 minutes past midnight, Mapper1 added more ways—in slightly darker solid blue—in the middle and above his previous edits, creating more or less a horizontal line in the middle of the right cluster. Twenty minutes later, Mapper2 made edits represented by the large blue dotted closed loop on the right of the map, between two groups of Mapper1 edits. In another 30 minutes, Mapper2 also added the darker blue dotted vertical line prominent on the right side of the map, connecting to a number of horizontal ways mapped by Mapper1 less than an hour earlier. Finally, twenty minutes on, Mapper1 mapped the dark blue solid way that connects the solid edits with the dotted cluster on the left, and in another 15 minutes mapped the dark blue solid buildings in the middle of the dotted cluster on the left.

We also isolated multiple instances where this parallel back and forth behavior repeated over multiple time slices, potentially signifying a semi-persistent interaction style. Most of these occur in the second week of the response, suggesting that disaster phasing is important to understanding what is happening: Editing each others' work, perhaps after people have become known to each other, can begin after the initial filling-in activity.



Figure 5.6: The Parallel Mapping of Two Mappers.
Mapper 1 is solid line, Mapper 2 is dotted line, opacity denotes passing of time.

Finally, in interviews, though some mappers wanted to deliberately stay out of the way of other mappers, perhaps to reduce editing conflicts and to spread the work around, others emphasized the importance of edits staying connected to other mappers to enable “good junction from person to person,” to ensure the continuity of the map.

Interpersonal Practice 2: Tag Correction

Correct metadata is important to many mappers. Getting the geometry of the physical space is just the beginning of getting a full map representation; getting the right metadata is the crucial next step in producing usable data that can be ingested into other platforms and applications downstream.

Since the tagging conventions are fluid in the OSM community—the prolific British mapper talked about a “tagging scheme made up on the spot”—mappers can often disagree on how to tag a certain map object, leading to prolific corrective behavior around the metadata specifically with respect to tagging.

To isolate the tag correction behavior, we constructed another type of network that is a subset of the overlapping changeset network. It is based on mappers editing the same object, but is even further restricted to having an edge between two mappers only when the object’s tags change between two consecutive edits. The directed edges between mappers in this network are weighted by the frequency of them overwriting each other’s tags. We then isolated the edges with high weights to find pairs of mappers where one corrected the other’s tags on many occasions.

We found many instances of one user correcting another’s tags. For example, in one such instance Mapper3 changed tags from “highway: road” to “highway: residential” for 41 of 43 objects newly added by Mapper4. All these corrections happened within the span of 3 hours.

We did not find any definitive instances of a tagging war, where two mappers consistently change each other’s tags for an object. This is likely because many less experienced mappers (including all the interviewees with whom we discussed this issue) do not have strong opinions about how an object should be tagged, and the prominence of correction is likely rather a function of less experienced mappers

needing to learn the open-ended and evolving tagging scheme. Tagging wars also might require longer time to develop than our 8-hour time slices.

A US mapper who is a GIS professional recalled his difficulties in learning this aspect of OSM mapping when he started to contribute during the Haiti earthquake. Though the GIS aspects of mapping were very familiar to him, the open-ended and community-resolved nature of OSM tags proved quite difficult. He somewhat jokingly suggested that there are “15 different ways of tagging something” and that learning these conventions for each context was the steepest part of his learning curve. He said that none of what he mapped was “a complete data point” from a GIS perspective, as OSM data is inherently open-ended.

Similarly, learning the correct use of the tagging scheme was difficult for the Canadian mapper, who spent a lot of time in an earlier event clarifying tags with experienced mappers on IRC and watching how experienced and especially local mappers tagged certain features. For example, it took some extensive consultation to establish that the dots he saw on the satellite map were actually water wells, and so those nodes should be tagged as such.

Our interviewees talked frequently about corrections in humanitarian mapping in OSM. The US mapper put it simply: “[If you put something on the map] someone will change it.” He suggested that the corrections, which are common, are less useful for new mappers without formal feedback. Other interviewees also brought up the need for training and mentorship of new mappers.

They explained that such mentoring interactions often happen at face-to-face mapping events. The Canadian mapper viewed this training and supervision as essential to the credibility of his local OSM community, which depended on the quality of the data they produced. Similarly, the British mapper regularly participates in local mapping events to “show the ropes” to new mappers.

However, it is not clear what the mentorship interactions would look like in the map itself. We explored several conceptualizations of how such supervision might present itself in the map, but the interactions that we isolated looked like close-up parallel mapping, and did not, on their own, suggest a mentoring or supervising interaction. However, *OSM Notes* is a post-Haiti feature that is now in place for

this kind of training support; Notes may be able to substitute for OSM edit analytics in finding mentorship behavior in the map. Indeed, Palen et al. (2015) report that the community wants to create opportunities to mentor. Here an experienced mapper explains to a novice mapper during Typhoon Yolanda that *Notes* should be used to:

“report an error in the data or to give some additional information, for instance the name of a street or an address etc. When available, an OSM contributor will attempt to resolve it.”

5.6 Study Summary & Discussion

The massive editing history of OSM is a rich dataset for exploring crowdwork. Specifically, in the context of crisis, these high-tempo collaborative mapping efforts produce large datasets that need to be spatiotemporally scoped for analysis. We use network techniques to help scope these datasets and isolate distinctive mapping behaviors that can then, and only then, be investigated qualitatively.

Through in-depth qualitative analysis of our scoped dataset, we refined the initial network techniques to introduce four succinct quantitative methods that isolate mapping practices and social interaction in OSM. We use frequency of participation during the same time of day to find mappers doing shift work or participating in sustained convergence. We used large clique sizes to find instances where many mappers were working in the same area, which in turn led to identifying mappers who converged around features that seemed salient to the external, on-looking world. We used the relative size of the largest connected component to isolate those who map continuous structures or fill in the map. Our methods can also filter by edge weights and reciprocity to find instances of parallel mapping. Finally, to identify instances of tag correction, we used subsets of overlapping changeset data to construct a new type of network where mappers were connected if they edited the same object, and the tags on that object changed from one version to another.

Our clique and component-based network methods were inspired by Keegan et al.’s (2012) article trajectories relating to Wikipedia’s breaking news collaborations. Our methods do not look for the cycles or the long, directed chains employed by Keegan et al. because our interest is in areas and not map

objects, but we still built on the kernel of their method: That network motifs represent particular collaborative practices in high-tempo events. This is especially the case in cooperative environments—like Wikipedia and OpenStreetMap—where users are co-constructing the final product through the interaction with the site of work and the results of each other’s labor.

Related research on Wikipedia’s editing roles and practices focuses on the distinction between novices, who engage through legitimate peripheral participation (Bryant et al., 2005; Lave & Wenger, 1991), and “Wikipedians”—the full members of the community who see their role as central to building Wikipedia (Bryant et al., 2005). Similarly, studies of other geo-wikis focus on the lifecycle of the users (Panciera, 2011). We do not find such a clear distinction in OSM: the work performed by the new and experienced mappers was not distinct, though quality varied. This might be due to the absence of clearly demarcated peripheral tasks in OSM. Many novice users did, however, apprentice (Bryant et al., 2005) by emulating mapping practices of experienced users. Other Wikipedia research shows that early on a few prolific users create most of the edits (Kittur et al., 2007) and value (Priedhorsky et al., 2007). However, in the disaster context in OSM, we do not find this long-tailed distribution, making these distinctions less salient.

Some spatial behaviors that we found qualitatively and then further isolated through network analysis confirm and elaborate on an earlier statistical study (Corcoran et al., 2013). Specifically, the spatial pattern of mapping continuous structures agrees with the pattern of map growth Corcoran et al. borrow from urbanization literature called “exploration” (Corcoran et al., 2013; Strano et al., 2012). Similarly, our found spatial pattern of filling in the map is consistent with their borrowed notion of “densification,” another mechanism of map growth (Corcoran et al., 2013; Strano et al., 2012). However, our qualitative approach further describes mapper editing patterns, including as understood by the mappers themselves. In addition, our temporal scoping of activity by the two weeks following the Haiti earthquake allowed identification of other patterns of work—temporal, spatial, and interpersonal—that are likely to arise from the convergent nature of volunteer mapping.

Refining Mooney and Corcoran’s co-editing network (Mooney & Corcoran, 2014) with further restrictions on tags and temporality proved to be a good tool for detecting tag correction, an interpersonal pattern of behavior. With this method, we found one-directional correction: one user changing tags initially made by another. Though we did not find back-and-forth tagging wars, the tag changing practice still relates to the form of collaboration represented by Wikipedia’s edit wars (Kittur et al., 2007).

Our work has the potential to provide design implications for the OSM community: The first is how our found behaviors highlight the need for collaborative features in OSM map editors; “mapping continuous structures” is a behavior that could be supported by highlighting map areas where work has recently been performed on a road or river by someone with whom one has worked in the past. A second is the need for real-time tagging support and discussion within the map editors for regions with many conflicting tags.

Finally, in this work we explored the mapper practices at the origins of large-scale volunteer crisis mapping. Since the Haiti response, geo-wiki tools have evolved, including creation of the HOT Tasking manager in OSM (Palen et al., 2015) and Ethermap for realtime editing (Fechner et al., 2015). We look forward to applying our network techniques in future research to understand the impact of these tools on mapping practices.

5.6.1 Expanded View of Collaboration

The CSCW literature has expanded understandings of *collaboration*. As a first-order concept, collaboration is constitutive of many things—mutual awareness of work, supervision and mentorship, division of labor, conflict and disagreement—and not simply, for example, mapping simultaneously on one feature. But these are challenging features of sociality to disentangle in the vast edit data of OpenStreetMap. So, to two concurrent ends, this paper offers 1) quantitative, network analysis techniques that were used to bound data for the qualitative analyses that were ultimately required to explain 2) the atomic practices that are the basis for highly distributed shared mapping. By scoping inquiry to the 2010 Haiti earthquake, the first major OSM crisis-mapping event, the conditions for various forms of mapping

interaction were made most visible because of the sudden convergence of new and experienced mappers onto the virtual space of that country.

5.7 Elucidating Network Structures and Organization Forms

5.7.1 Introduction to Analysis of Network Structure and Organizational Forms

The analyses in Kogan et al. (2016) above largely focuses on the collaborative work practices that emerged in the crisis-mapping activity after the 2010 Haiti earthquake. However, it did not include an investigation of the other two forms of sociality around which I organize this dissertation: organizational forms and social network structures. Thus, in the remainder of this chapter I discuss further analyses that focus on these two emergent forms of sociality. Moreover, at the end of the chapter I take the opportunity to reflect on how framing the analyses around the two intersecting aspects of the shared information space—explicitly shared site of work and human-readable record of work—helps in interpreting the findings and positioning them within the larger context of highly-distributed collaborative work in the high-tempo environment.

5.7.2 Network Structures

The analysis of collaborative practices of OSM crisis mappers in Kogan et al. (2016) above successfully demonstrated that the broadest conceptualization of collaboration in OpenStreetMap—changesets of two mappers overlapping within a certain time interval, indicating that they are mapping the same geographic area—effectively captures the various mapper practices and thus can serve as a productive operationalization of collaboration in OpenStreetMap. Thus, I focus on the overlapping changeset networks in order investigate the network structures that emerged in crisis mapping after the Haiti earthquake.

Time Slice Construction for Network Analysis

While the eight-hour time slices were a useful tool in finding the daily patterns of activity and related mapper practices, this temporal resolution is too finite for the emergence of network structure, on

which I would like to focus in this section. Thus, I divided the two weeks of the activity described in Kogan et al. (2016) into three time stages: *During*—January 12th through January 16th, immediately *After* (simply ‘*After*’ from here on)—January 16th to January 21th, and *Long After*—January 21st to January 25th. Unlike the analysis of retweet activity in response to Hurricane Sandy in Chapter 4, this analysis does not include a *Before* period. This is because no mapping of Haiti was taking place in the days before the earthquake, since it is a type of natural hazard with essentially no warning. Thus, no activity could be analyzed for that period. The three stages on which I base the analyses loosely coincide with the initial phases of disaster response: threat and impact correspond to the *During* period, inventory and rescue correspond to the *After* phase, and remedy and recovery—to the *Long After* (see the Chapter 2: Background for more details). For each period—*During*, *After*, and *Long After*—I constructed an accumulated network of activity, where an edge exists between two mappers if their changesets overlapped during this period, and the edge weights signify how many of their changesets overlapped.

Network Statistics under Configuration Model

I used the empirical in- and out-degree distributions for the three time slice networks to generate 100 synthetic directed networks for each period using configuration model—a random graph model that rewires the edges to simulate an empirical heterogeneous degree distribution. I calculated various network statistics for the synthetic networks, providing a distribution against which to compare empirical values. When network statistics for the synthetic networks were not normally distributed, I performed nonparametric sign test to test how likely the empirical values for the *During*, *After*, and *Long After* are under the configuration model. Otherwise, I performed Student t-test for the same purpose.

Network Interconnectedness

As these are directed networks—the directionality signifies who “worked on top” of whom—the interconnectedness of the networks can be measured with the fractional size of the strongly connected components, which accounts for the directionality of the edges in determining if there is a path between the mappers.

85% of all the contributors in the *During* time period comprised the largest strongly connected component (Table 5.2). This suggests that a vast majority of mappers in this period consecutively built on each other’s work, whether intentionally or not. For the two later periods—*After* and *Long After*—the fractional size of the strongly connected component went down to 0.67. As the fractional size of the largest weakly connected component for all the time periods is close to one, this suggests that while most contributors in those periods overlapped their crisis mapping spatially, only two thirds of them sequentially built on each other’s work.

The largest strongly connected component for the *During* network is statistically significantly larger than would be expected from configuration model (one-sided sign test, p-value= 6.939×10^{-18}). On the other hand, largest strongly connected component for the *After* network is significantly smaller than we would expect from the model (4.206×10^{-8}). The differences are not statistically significant for the *Long After* network.

	<i>During</i>	<i>After</i>	<i>Long After</i>
# of nodes	176	172	126
# of edges	1938	1044	669
Frac. size of strongly LCC	0.85	0.67	0.67
Median degree	14	7	6

Table 5.2: Basic Network Statistics for Three Time Slices

Degree Distribution

While the interconnected nature of the *During* network is not explained by the degree distribution alone, it is certainly telling. The median overall degree—number of other mappers with whose changesets the mapper’s contribution overlapped—is 14. Moreover, while most contributors overlap their mapping with only a few others, a sizable minority worked in the same part of the map as many other mappers. This heterogeneous degree distribution reflects a few mappers whose work overlapped with that of others very often—such as *cetest* with the degree of 106—likely connecting otherwise disconnected parts of the network (Figure 5.7).

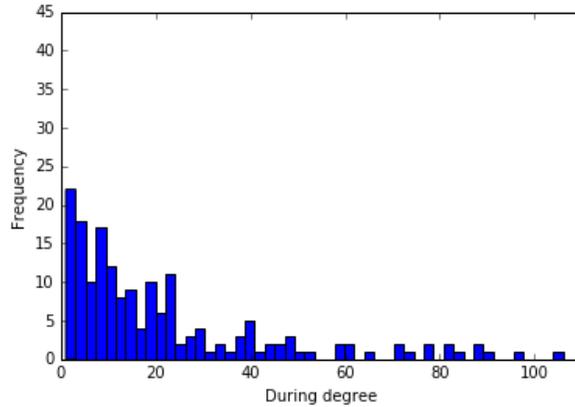


Figure 5.7: During Degree Distribution

The median degree for the *After* period is 7—half of that in *During* (Figure 5.8). The distribution below shows that more of the mass is pushed towards the origin, with the vast majority of the mappers only overlapping with few others. Similarly, in the *Long After* period, the median overall degree is only 6. Accounting for the edge weights magnified the difference in how mappers overlapped with the work of others between the *During* period and *After* and *Long After*.

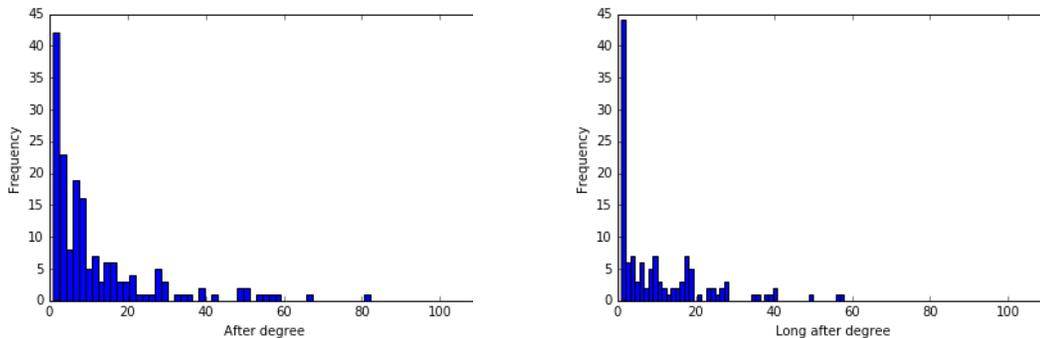


Figure 5.8: After and Long After Degree Distributions

Pro-Social Network Structures: Reciprocity and Transitivity

In these spatially-motivated networks, reciprocity means the proportion of all pairs of contributors where the spatial footprint of A’s mapping overlapped B’s work, in which B also overlapped A’s work. The reciprocity for all three networks is not very high, but considerable at 30%-31%. For all three networks, the empirical reciprocity is statistically significantly higher than would be expected from

the configuration models (one-sided t-test, *During*: p-value= 4.195×10^{-125} , *After*: p-value= 9.829×10^{-120} , *Long After*: p-value= 3.447×10^{-106}).

Since OpenStreetMap does not provide a human-readable record of activity where the contributors could keep track of the work that has been accomplished so far and by whom, the interface does not actively facilitate reciprocity—it is not easy to ‘respond’ to another mapper who edited your work by editing theirs. On the other hand, in crisis mapping the site of work is very well defined and is explicitly shared—the part of the map that represents the affected area. Thus, such collocation of mappers within the map likely contributes to the considerable reciprocity we observe, since the mappers are editing the same limited area, increasing the chances of overlapping someone’s work after they overlapped yours.

Transitivity measures the concept of triadic closure, which is especially common in social networks and relates to the notion that “a friend of a friend is also a friend.” In a directed network, the most meaningful measure of transitivity is the proportion of transitive triples derived from the triadic census (Holland and Leinhardt, 1970). However, in the case of overlapping changeset networks, where the directionality of edges is based on timing, the directed metric does not necessarily reflect any meaningful social relationship. Therefore, I focus on the undirected measure of transitivity here, which simply denotes how often when A’s and B’s work overlaps, and B’s and C’s work overlaps, A’s and C’s work also overlaps on the map. As Table 5.3 illustrates, the *During* period has the largest transitivity of the three periods—0.41, while *After* and *Long After* show transitivity of 0.29 and 0.36, respectively. This suggests that in the *During* period, neighbors of neighbors were also neighbors themselves (in terms of their mapping) more often than in other time periods. In the *During* and *After* periods, the observed

	<i>During</i>	<i>After</i>	<i>Long After</i>
Reciprocity	0.31	0.33	0.33
Undirected transitivity	0.41	0.29	0.36
Degree assortativity	-0.16	-0.11	-0.05
Experience assortativity	-0.01	0.06	-0.01

Table 5.3: Reciprocity, Transitivity, and Assortativity

transitivity is statistically significantly higher than that expected from configuration model (one-sided t-test; during: $p\text{-value}=6.489*10^{-141}$; after: $p\text{-value}=9.060*10^{-102}$). The difference is not significant for the *Long After* period.

For this type of spatially-produced network, however, it is impossible to disentangle to social and spatial dynamics: it is not clear whether the higher transitivity in the *During* period implies a more pro-social behavior or simply reflects the spatial density of the high-tempo crisis mapping in the immediate aftermath of the earthquake. This well-defined, and spatially constrained, site of work likely contributes to the considerable levels of transitivity in all three periods, especially considering that OpenStreetMap does provide a human-readable record of work that would actively facilitate the triadic closure by explicitly highlighting who maps with whom (or near) to the contributors.

Assortativity: Who Maps with Whom

As you might remember, network assortativity metrics quantify what types of nodes tend to connect to each other. Degree assortativity is a network-level coefficient quantifying the strength of the relationship between the degree of the network's nodes and degree of their neighbors. The *During* and *After* networks are weakly degree-disassortative, as shown in the Table 5.3. These networks are more disassortative with respect to the node degree than would be expected from configuration model (on-sided t-test; after: $p\text{-value}=6.679*10^{-73}$; after: $p\text{-value}=5.324*10^{-21}$). The *Long After* network is also degree-disassortative, but the relationship is very weak. This network is less disassortative than would be expected from configuration model (one-sided t-test; $p\text{-value}=2.540*10^{-6}$). This suggests that in *During* and *After*, for the contributors whose mapping overlaps with the work of many others, those others tend to overlap with only a few. It is not clear what practices are at play here, but the fact that in *During* and *After* the mappers who were less active in overlapping their work with others more frequently worked with the more prolific overlappers suggests the potential for some type of symbiotic or possibly even a mentoring relationship, like when the novice contributors map alongside the established mappers in order to learn the norms and practices of the community.

5.7.3 Organizational Forms: Social Roles of Expert and Novice Mappers

Social roles taken on by the participants are an important aspect of the organizational forms that emerge in disaster-related online activity. As already alluded above, the social roles that are especially prominent in crisis mapping on OpenStreetMap are those of experienced and novice mappers, since as a content-producing organization OSM is concerned with successful incorporation of new mappers into the community and instilling its norms and practices in the newcomers. Therefore, to explore the practices of experienced and novice mappers, I investigate how these practices were reflected in the collaboration network.

Assortativity with respect to Mapper Experience

For this purpose, I focus on network assortativity with respect to experienced/novice node attribute. However, in the two weeks under consideration, most of the activity was generated by the experienced OpenStreetMap contributors, with only three or four novice user in each time period. Thus, the assortativity with respect to experience shows a very weak relationship, as reflected in Table 5.3. While these values are really small for all the networks, comparing them to those expected under the configuration model for each time period still reveals an interesting pattern. In *During*, the network is statistically significantly more disassortative on this attribute than expected under the model (one-sided sign test; $p\text{-value} = 7.889 \times 10^{-31}$). This suggests that at the start of crisis mapping, experienced users and novices mapped together (or near each other) more than expected from just how many others they overlapped with (encoded in configuration model). Ostensibly, the novice users likely tended to map around the more experienced contributors, while they were learning the ropes of crisis mapping in OpenStreetMap. The other two networks are statistically significantly less disassortative (more assortative) with respect to experience attribute than expected from the configuration model (one-sided sign test; *After*: $p\text{-value} = 1.002 \times 10^{-21}$, *Long After*: $p\text{-value} = 2.756 \times 10^{-8}$). This suggests that as the time progressed, the novice users were somewhat less likely to map around the veterans and more likely to map amongst themselves, compared to what is expected solely based on their number of overlapping neighbors.

Community Structure

Community detection is another way to glean some aspects of emergent organizational structure and how experienced and novice users fit into it. Since these are directed networks with edge weights, I rely on the Weighted Stochastic Block Model (WSBM) for this analysis (Aicher, Jacobs, & Clauset, 2013). This is a probabilistic model that effectively coarse-grains the network by assigning nodes into bundles of structural equivalence—meaning that they perform the same function in the network and are interchangeable from the point of view of network structure. In order to select the WSBM that best fits the empirical data, I performed model comparison across the number of communities ranging from 2 to 20, two types of weight distribution (Exponential and Poisson), and two types of edge distribution—Bernoulli and degree-corrected. This resulted in the comparison of 76 models for each time period.

For the *During* period, the best-fit WSBM that has 2 communities, exponentially-distributed edge weights, and a Bernoulli-distributed edge probability. Within this structure, the novice mappers were split across the two communities, with one—Jesús Gómez—a Spanish mapper who has since accumulated a vast record of contributions—in Community A and two other mappers in Community B (a Brazilian and a German mapper). In the *After* period, all the novice mappers were concentrated in Community A: Jesús Gómez, an American mapper with username bogande, and three contributors with clearly German usernames. Finally, in the *Long After* period, three of the novice mappers were in Community A—two of the German mappers active in the *After* and another contributor with a German username, while mapper HB9DTX was the sole novice in the Community B. Either the majority of the novice mappers or the most productive of them (such as Jesús Gómez) consistently gravitated to the Community A across the three time periods. These transitions in the composition of novel users also suggest that there might be a geographic pattern to onboarding of new OpenStreetMap contributors: mappers from across the world, including Europe and South America were at the forefront of the crisis mapping, while the Europeans and American mappers, who have historically had stronger representation within the community, were more represented among the contributors who onboarded later, as the crisis mapping effort gathered momentum across the established channels of international OpenStreetMap community.

5.8 Forms of Sociality and Two Aspects of Shared Information Space

The analyses in Kogan et al. (2016) above largely focused on the mapper practices and found seven prominent mapper practices, including the broad categories of temporal, spatial, and interpersonal. Spatial practices were more prevalent in the initial stages of the response, while the interpersonal practices more prominently emerged later in the crisis mapping. The clearly defined site of work—map of the affected area—likely enabled the spatialized collaborative practices early on, by placing the mappers into the same circumscribed digital space where the map urgently needed to be updated. The interpersonal practices, on the other hand, were not prominently displayed until later, since no affordances of the OpenStreetMap interface enabled mappers to easily keep track of who has done what and whose work they were modifying. Only with time and the spatial co-situation within the map did they eventually develop enough passive awareness and possibly familiarity with each other to build up to the interpersonal collaborative practices.

The network structures arising from crisis mapping in response to the Haiti earthquake show a rather consistent temporal pattern. The *During* period—at the height of initial response activity such as search and rescue, where accurate maps of the area are vital to the success of the response efforts—produced the most interconnected, dense network with somewhat high reciprocity, transitivity, and degree-disassortative relationship among the mappers. These patterns cannot be explained by the frequency of mapper interaction alone (represented as the degree distribution in the configuration model). As the response progressed, most of these metrics (reciprocity excepted) showed lower values in subsequent networks, suggesting that the mapper activity produced less interconnected, dense networks, with less signs of classically pro-social behavior such as the neighbor of the neighbor becoming a neighbor (triadic closure), and fewer connections between those overlapping with many and mappers overlapping with a few, which may reflect instances of legitimate peripheral participation and possibly even mentoring.

The somewhat high levels of reciprocity and transitivity are notable considering that the OpenStreetMap interface does not provide easy opportunities for the mappers to keep track of who has

done what. However, the spatially concentrated nature of the crisis mapping ensures that the mappers often overlap each other, and even the triads of mappers tends to be connected through their overlapping changesets.

In addition to the network structure, organizational forms are also important emergent forms of sociality in crisis. As discussed in the Background chapter (Chapter 2), social roles embodied by the participants are often central to the organizational forms that emerge in the process of high-tempo crisis-related activity. In the case of crisis mapping in OpenStreetMap, a set of the prominent social roles has to do with the length of contribution to the platform, since familiarity with its norms and practices, as well as the technical mapping expertise, all take a certain time to develop. As I've shown above, all three networks are weakly disassortative with respect to the mapper experience, suggesting that novice mappers tend to map slightly more frequently around experienced contributors. This dynamic around social roles is more prominent in the height of the emergency period, with later networks showing a less disassortative relationship, as the novice mappers acquire more experience in the norms and policies of the community. Additionally, either the majority of or the most productive novice mappers tended to contribute to the same community from one period to the next. Even though the human-readable log of the work accomplished was not available, the spatial proximity of their contributions in this well-defined explicitly-shared site of work facilitated the novice crisis mappers' ability to participate in an ad hoc community of mappers.

CHAPTER 6: Twitter Conversations Among locals during Hurricane Sandy

6.1 Study Overview

The content in the Sections 6.2 through 6.8 of this chapter is a pre-reprint of Kogan, M. & Palen, L. Twitter Conversations Among Locals During a Hurricane: At-Scale Features of Conversational Structure in a High-Tempo, High-Stakes Microblogging Environment. Under review for the *ACM 2018 Conference on Human Factors in Computing Systems (CHI 2018)*. It is included in the dissertation with the permission of my co-author.

This is the latest of the three studies, and while the conceptual framing of the two aspects of the shared information space does not figure in this article, it certainly informed the thinking behind this work in more explicit ways than in the earlier work. Conversational nature of Twitter interactions has long been documented, and Palen and Anderson's work (2017) articulates the need to go beyond the single post and look at the content conversationally especially well. Thus, my interest in the Twitter conversations as a unit of analysis has been long-standing, but the specific approaches to foregrounding the conversations in this work have been more recently informed by the focus on the record of activity as important for the creation of the shared information space. Specifically, the focus on the back-and-forth of the reply conversations—whether represented through a directed graph or a temporality of the turntakes—zeroes in on the availability of such a record in replying as both a digital trace log for the researchers and as an articulation work resource for the participants. And as the findings show below, the availability of such a record allows the users to also treat it as a persistent log and thus enable them to return to the shared information space of disaster-related sensemaking as their circumstances permit—despite the fact that Twitter as a whole does not provide a well-defined site of work.

The original work explores organizational forms, collaborative practices, and network structure that arise in the reply conversations of the geographically vulnerable in response to the 2012 Hurricane Sandy. It propels social media research beyond the single post as the unit of analysis toward fuller treatment of interaction by making the construct of the *conversation* analytically available. I offer a

method for constructing @reply conversations in Twitter to apprehend social media conversational features at scale. I apply this analysis to the high-tempo, high-stakes environment of 2012's Hurricane Sandy, with its high volume of online talk by affected locals and distinct disaster-stage phasing by which to consider interactional difference. I investigate the temporality of conversations; the relationality of who speaks to whom; the number and kind of conversationalists; and how content affects temporal features. The analysis reveals that, during the height of the emergency, people expand conversations both in number and kind of conversational partners—just as their information search intensifies. It also shows that this expansion contributes to slower-paced and longer conversations in the high-emergency period, suggesting reliance on online informational relationships during times of greatest uncertainty.

Section 6.9 highlights how the two aspects of shared information space on which I focus in organizing this dissertation—human-readable record of who responded to whom and the lack of explicitly-shared, well-defined site of work in Twitter—help us understand and position the findings from this study in the larger context of high-tempo, high-volume social media activity across affordances of various platforms and types of work.

6.2 Study Introduction

Social media platforms are increasingly being used by the public during disaster response for a range of purposes. A flourishing body of crisis informatics research has been focused on understanding the socio-behavioral phenomena emerging in crisis. This includes areas such as sentiment analysis of the social media stream for various populations (Celli & Rossi, 2012; De Choudhury et al., 2014); information diffusion (Bakshy et al., 2012; Kogan et al., 2015; Qu et al., 2011; Romero et al., 2011; Starbird & Palen, 2010); collective sensemaking (Palen et al., 2009; Starbird, 2013); and self-organizing behaviors of people who come together through social media to accomplish disaster-related tasks, such as lost-and-found pet matching and pet rescue coordination (White et al., 2014; White & Palen, 2015) critical information collation (Starbird & Palen, 2013), event reporting (Keegan et al., 2012; Keegan et

al., 2013; Sung-Yueh Perng et al., 2013) and mapping using an open source platform (Dittus et al., 2016; Soden & Palen, 2014).

Furthermore, sociologists of disaster have described how people directly affected by crisis events are far more than passive victims or uncontrollable crowds—historically pervasive stereotypes even in academic circles. Rather members of the public participate in rescue, recovery and rebuilding efforts (Dynes, 1970; Fritz & Mathewson, 1957; Kendra & Wachtendorf, 2003). This active participation may result in new relationships being established, new sources of information found, and new communities formed (Starbird & Palen, 2011). Researchers in crisis informatics investigated how on-line content produced during disasters events can be used to understand the nature of participation by members of the public, and how it could be leveraged for other disaster warning, rescue and recovery purposes (Hughes et al., 2008; Liu, et al., 2008; Palen et al., 2009; Sutton et al., 2008; Vieweg et al., 2008).

However, in the case of analyzing microblog content—specifically, Twitter—an inherent limit of such research is that the tweet is the unit of analysis, one that is the outcome of built-in technical parameters that reify the single speech act as a dominant discursive element. This is what Palen and Anderson refer to as the “tyranny of the tweet” (Palen & Anderson, 2016). That is to say, the manner in which data are provided, collected, stored, and sampled make the tweet as a unit of analysis far more convenient than other treatments. Single-tweet treatment favors natural language processing methods (Kim et al., 2012; Ritter et al., 2010) and even manual coding (Celli & Rossi, 2012) to annotate tweets with categorical characteristics, even in the conversational context. Other studies constructed social networks from reply activity Bak et al., 2012; Gonçalves et al., 2011; Hannak et al., 2014), still making use of the individual tweet without considering conversations as a whole. Indeed, other computational treatments—at least over very large quantities of data that are produced in a mass disaster event—are quite difficult to architect.

This research presses the social computing work of crisis informatics toward fuller treatment of the social media interactivity between affected locals in a disaster event— Hurricane Sandy in 2012—by extending network analytic methods to produce the “conversation” as the unit of analysis, using the

“micro-scale communicative exchanges” (Bruns & Burgess, 2011) of the @reply feature. It makes monologues, dialogues, and group discussions analytically available for apprehending longer form online communication in emergency situations, and at scale. Who is speaking to whom, when during the disaster, and over what time frames? Does the composition of actors shift upon the introduction of new information sources? Do conversations about the disaster have different characteristics than those not disaster-related, but during the same time period?

Consider this Twitter conversation generated and found through the @reply convention to signal turntakes. It begins with @susanruben, post-landfall, notifying contacts of a blog post about managing damaged belongings:

2012-11-04 17:06:26, susanruben:
Please Take Nothing Not Ready. Blog Post. @groovydude @phillyinquirer
<http://t.co/ENe8LXRO> #sandynj

Three hours later, one contact makes an observation about what is happening to damaged appliances left curbside:

2012-11-04 20:13:30, MoeThatch:
@susanruben @groovydude @PhillyInquirer I was helping in Vent Heights and the scrap trucks were scavaging peoples appliances w/o asking
2012-11-04 20:14:52, MoeThatch:
@susanruben @groovydude @PhillyInquirer and the house we were at, chased one down and said u cant take anything, we need for insurance

Three hours later, another declares a rule about appropriate versus inappropriate salvaging of damaged property:

2012-11-04 23:00:59, SillyRaven:
@MoeThatch @susanruben @groovydude unspoken rule: on the curb fair game. Driveway, Lawn, alley - don't take #scrapping #sandy etiquette

After a long interval of about 19 hours, the conversation continues, with some contention about the implied “etiquette,” and the unfortunate consequence of such ideas not being accepted social norms:

2012-11-05 18:13:23, MoeThatch:
@SillyRaven @susanruben @groovydude im not sure everyone knows the scrapping etiquette, appliances were in driveway.
2012-11-05 20:38:27, SillyRaven:
@MoeThatch hmmm - near the ,...street or midland on the driveway? that is really poor form if it was midland.
2012-11-05 20:44:07, MoeThatch:
@SillyRaven on property, beyond sidewalk area, so i would consider that far enough from curb to be off limits.

2012-11-05 20:45:05, SillyRaven:
@MoeThatch bastards! i have my 'good' trash = stuff adjusters need to see
set far back near my garage door...
2012-11-05 20:48:15, MoeThatch:
@SillyRaven lesson learned!

In this conversation, we see a number of features that we will examine in this paper, at scale. There is branching that occurs between conversationalists—in this case right at the first turntake. We see long durations between rapid bursts of interaction, with important events related to the conversation seemingly happening during the long in-between. Finally, we see prompting and discussion about what to do in a situation that may not have much precedent for these people.

The research presented here offers a method for constructing @reply conversations as units of analysis for then apprehending conversational features at scale in the high-tempo, high-stakes environment of disaster. We investigate the **temporality** of conversations; the **relational properties** of who speaks to whom; the number and kind of **conversationalists** in discussions; and how **content** affects temporal features.

We examine these structures at scale over four periods of hurricane event, derived for at-scale analytical purposes from the sociological conceptualizations of disaster phases (Powell, 1954), which describe socio-behavioral phenomena for mass emergencies from natural hazards. Social media and other on-line activity allow examination of temporality at even finer time scales because of the timestamped posts, as well as how temporality factors in relation to other conversational features.

Social scientists have shown that those affected often make protective decisions based not only on the official sources but also considering the social cues from personal ties such as family and friends (Lazo et al., 2015). Moreover, importance of weak ties (Granovetter, 1983) in information seeking (Morss et al., 2015) would suggest that those affected may engage in conversations with twitterers outside their immediate social circle. Thus, relationality of who speaks with whom and the characteristics of the conversation participants may also illuminate socio-behavioral phenomena in natural hazards emergencies.

6.3 The 2012 Hurricane sandy event (US Landfall)

Hurricane Sandy made US landfall late on October 29, 2012 in Brigantine, New Jersey. It affected one of the most populated regions of the country, including the New York, New Jersey, and Connecticut tri-state area. Overall it impacted a total of 24 states (FEMA, 2015). In the US, Sandy claimed 162 lives and, with estimated US \$6 billion of damage, it was the second costliest hurricane (Blake et al., 2013). Approximately 776,000 people were displaced (Yonetani & Morris, 2013) and 650,000 homes destroyed or damaged. More than 8 million people lost power (Blake et al., 2013), and for many, power was not restored for weeks.

6.4 Data Sets and Methods

6.4.1 Data Collection Steps

We started data collection on October 24, 2012 using Twitter’s streaming API to collect on a range of broad and highly localized terms. We isolated twitterers who had at least one geo-located tweet within a bounding box (Figure 4.2), which we defined in collaboration with meteorologists, social scientists, and GIS researchers at the National Center for Atmospheric Research. Though only about 1.5% of tweets are geo-located, prior research shows that this sparse geo-spatial data is sufficient to accurately detect population location and mobility patterns (Hawelka et al., 2014). Moreover, our focus on Twitterers who have only one geo-located tweet within the bounding box of interest enables us to better navigate the limitations of geo-tag-based data collection by producing a less restrictive notion of locality. The users isolated in this manner are considered “geographically vulnerable” to the impending hurricane. We then collected all of these twitterers’ tweets starting from October 13, 2012 for a dataset of 5.6M tweets created by 28.5K twitterers. Note that the REST API returns up to 3200 of a user’s tweets from most recent to least, which usually allows examination of a user’s behavior before the hurricane.

6.4.2 Data Sets

From this, we then constructed a subset dataset for analysis by selecting @reply tweets, and the tweets to which those replies were made. We further restricted the dataset to only include the tweet-reply pairs where the original tweet is also in the geo-vulnerable dataset. This achieves retaining only the conversations between geo-vulnerable users who were all found through the keyword search, meaning they tweeted something Sandy-related at least once to be found (though we work with their full contextual streams that also include off-topic conversations). This *conversational geo-vulnerable set* spans October 15 to November 23, 2012.

We further divided this data set into four 10-day periods that roughly correspond to disaster stages (Powell, 1954). The *Before* period spans October 15-24 and captures the time before most people in the affected area knew they would be under threat, but not so far back that their social networks would be dramatically different just as a function of time. The *During* period spans October 25-November 3 and includes high warning, evacuation, and hurricane's landfall, when people were engaging in protective decision-making (Morss et al., 2015). This period also includes people taking stock of the damage immediately after and grappling with whether the NYC marathon was an appropriate use of city resources. The *After* period of November 4-13 includes initial days of recovery when some of the geo-vulnerable are still displaced and/or without power and heat. Adding to this difficult situation, a second storm—a *nor'easter*—happened in this time period. Finally, *Long After* spans November 14-23, and includes the start of long-term recovery, when situations have stabilized for most.

6.4.3 Conversation Construction

We used network methods to analyze the complicated conversation structure. We constructed a Directed Acyclic Graph (DAG) for each conversation from replies pointing to each other. Each conversation as a DAG constitutes a separate component in the tweet-reply network.

In some branching conversations, the newer branches reply to much earlier tweets, complicating the conversational order. We store such the branching conversations by the topographical order of the

DAG—going from tweet to its reply and following the tweet tree, which sometimes means suspending tweets’ temporal order. This process produced 5350 conversations composed of 14310 tweets involving 2765 unique twitterers. 5281 conversations fall

	Before	During	After	Long After
# of tweets	2215	5732	3740	2399
# of unique twitterers	791	1772	1216	863
# of conversations	823	2110	1418	930

Table 6.1: Conversations Across Time

into the *Before*, *During*, *After*, and *Long After* periods (see Table 6.1 for details). The 69 conversations that span more than one period were excluded from the analysis.

6.4.4 Temporal Analysis Methods

Considering Twitter activity conversationally and discursively, even beyond @reply convention, presents a methodological challenge of determining the appropriate length of Twitter streams that represent coherent streaks of discursive activity. For this, we turn to the technique used to determine work session length in the Wikipedia research (Geiger & Halfaker, 2013), by identifying characteristic lengths of the punctuated bursts of activity—here periods of twitter engagement. We constructed a frequency distribution of log-scaled time spans between events—log scaled because these time spans can be very long-tailed. If there are several distinct user behaviors, this process produces multi-modal mixtures of lognormal distributions. We use medians as measures of centrality as many distributions are skewed. We used non-parametric rank-based tests such as Kruskal-Wallis and post hoc Dunn test with Holm-Šidák correction for multiple comparisons. We used Chi-square tests to check for independence of categorical variables.

6.4.5 Network Analysis Methods

We then constructed a directed graph where a twitterer points to another twitterer, to whose tweet(s) they replied. The number of replies between twitterers is stored as the edge weights. We then used the 10-day periods described above to construct four time period networks for *Before*, *During*, *After*, and *Long After*. We settled on these four cumulative networks to match all the temporal analyses in this

paper, and because the distribution of inter-reply time is long-tailed (as shown in the findings), making it difficult to empirically select a fine-grained temporal resolution.

Since we are interested in social dynamics of conversations, we distinguish between twitterers who communicate broadly but do not necessarily engage with their contacts frequently from those who converse often across their contacts. For this we calculated each twitterer's strength to degree ratio, which is high when the twitterer communicates frequently across his or her contacts. We computed this ratio for both in- and out-measures; both directed ratios have a median of 1. We consider a twitterer to be a frequent conversationalist if both their directed strength-to-degree ratios are above 1.

6.4.6 Conversationalist-Based Analysis Methods

For this analysis, we considered a twitterer to have a high follower count—a proxy for popularity, social status, and influence (Cha et al., 2010; Kwak et al., 2010)—if the twitterer had more than 500 followers. For the purposes of classifying twitterers as having high follower counts, we chose to use the earliest follower count in our data set for each twitterer, to indicate what kind of user they were going into the event.

6.4.7 Content Analysis Methods

We could determine which conversations were on topic as those containing tweets collected based on the search terms. However, though our search terms were carefully chosen, they insufficiently represent the range of ways people might be talking about the hurricane without ever using a search terms (Anderson et al., 2016). In addition, having a rich contextual stream for users enabled use of qualitative analysis to find the on-topic tweets and conversations that were missed by the automated keyword search. We hand-coded all conversations as on-topic/off-topic; analysis was done by one author and checked randomly by a second for consistency. Twitter handles that belong to individuals are anonymized for reporting purposes in this paper.

6.5 Analyses and Findings

We now present our analyses in the following order. First, we consider the temporal dynamics, where we found, unexpectedly, that the conversations *slow* down in the high emergency period. We then analyze the relational structure, the user composition, and the topic of the Twitter conversations to discover what social aspects contributed to the slower conversations in the *During* period. We intersperse interpretation to animate statistical reporting.

6.5.1 Temporality in Replying Activity

Following the technique used by Geiger and Halfaker to determine the length of a characteristic *work session* in Wikipedia (Geiger & Halfaker, 2013), we analyzed the temporal patterns of tweeting behavior by calculating the span of time elapsed between each pair of a twitterer's consecutive tweets. Since this time span ranges widely between and even among users, analyzing it on a log scale elucidates the patterns more clearly. To contextualize temporal patterns of replying, we first need to engage with inter-tweet time and times between outgoing replies

6.5.2 Inter-tweet Time

Figure 6.1a shows that the tweeting activity of the geo-vulnerable twitterers has a multi-modal frequency distribution that displays three distinct peaks. This suggests that the time spans between tweets have three characteristic lengths: one peak has a center of around 3 minutes, another around 2 hours, and another around 16 hours. The first corresponds to rapid-fire tweeting, when users are actively engaged with the Twitter content and participate very rapidly. The second peak corresponds to the practice of the periodic check-ins throughout the day, when breaks between tweets can take hours. And the last peak corresponds to the diurnal patterns of tweeting, when some users check on their feed in the beginning and end of day.

Time Spans between Outgoing Replies

The temporal practices of outgoing reply activity display similar patterns. The peaks are positioned closely to the inter-tweet time distribution, but the mass is shifted towards the right-most mode (Figure 6.1b). This suggests that, as a temporal practice, outgoing replying activity has the same temporal signature as tweeting overall and fits into tweeting practice. However, in this case, many more outgoing replies happen within the diurnal pattern—almost as many as in the rapid-fire mode. That is, when all the tweeting activity is considered, we find that geo-vulnerable twitterers have three likely types of behavior—*rapid*, *periodic*, and *diurnal*. The majority of the time they tweet rapidly, with periodic tweeting being less frequent, and diurnal less frequent still. However, if only users' outgoing replies are considered, the three frequent patterns of behavior remain unchanged, and the majority still reply in rapidity, but more twitterers reply as part of their diurnal routine than periodic check-ins.

Hence, many users take longer between replying to tweets than just tweeting, potentially lingering for a variety of reasons. We explore these reasons further in the rest of the paper. Here is an example of an unrushed conversation that occurs in the high emergency *During* portion of the event between 3 people. Note that @tcbraider took a total of 17 hours between two of her/his turntakes because the intervening turntake by @BethHowe takes 12 hours.

```
2012-10-30 19:05:10, tcbraider:
@BethHowe @pkrobert8 Haven't connected all day. How you making out? Not
exactly how'd I want to show you my beloved city
2012-10-30 23:08:04, pkrobert8:
@tcbraider @BethHowe Doing fine.'Being treated very well by your city and
its people
2012-10-31 06:55:13, BethHowe:
@tcbraider Your city is wonderful and it will be soon be again. Amazing
work by u, @danatele and all of the volunteers at @RedCrossPhilly.
2012-10-31 12:35:03, tcbraider:
@BethHowe Thanks @BethHowe. Great work by you and @pkrobert8 as well
```

Inter-Reply Time

The notion that users might linger with the response is also supported, if somewhat counterfactually, by frequency distribution of the log-scaled time elapsed between a tweet and a reply directed towards it. Figure 6.1c shows a much more long-tailed unimodal distribution. The single mode

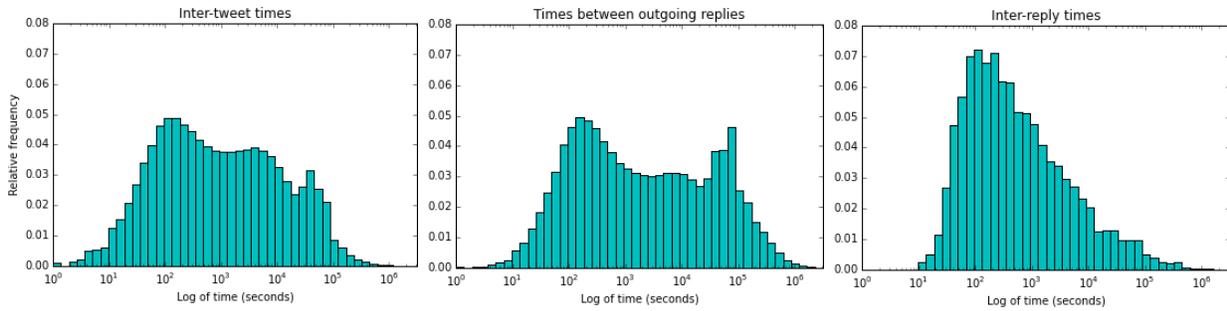


Figure 6.1: (a) Inter-tweet times; (b) Times between outgoing replies; (c) Inter-reply times

suggests a single prevalent behavior, but the long tail indicates that to the right of the mode this behavior is extremely variable. Thus, for tweet-reply spans longer than 3 minutes, the timing of the replies is better explained by users' own tweeting habits rather than by when the original tweet to which they are replying has arrived.

Conversational Tempo & Duration across Time Periods

The changes of the conversational tempo (inter-reply time) across the time periods representing disaster stages are especially telling. The median inter-reply times for *During* and *Long After* are longer than for the *Before* period (Table 6.2). Thus, the social contexts of high-emergency period and post-disaster recovery drive more people into taking slightly longer to reply to tweets, thus slowing down the exchanges between twitterers. Looking at the conversation-level patterns, we find that the median conversation durations also get longer as time progresses across the four periods, though only the

Median	Before	During	After	Long After
Inter-reply time (Kruskal-Wallis H=12.76, p=0.005)	323	383* p=0.039	363 p=0.156	404** p=0.002
Conversation length (Kruskal-Wallis H=11.42, p=0.010)	656	693 p=0.502	665 p=0.502	900** p=0.006

Table 6.2: Median Conversation Tempos and Lengths

Post hoc Dunn test comparing to Before period with $\alpha=0.05$

difference between *Before* and *Long After* is statistically significant (Table 6.2). Since, on average, the number of tweets in the conversations are the same across the four periods, the lengthening of the conversations in the disaster and recovery stages must be largely due to longer periods between conversational turntakes above (inter-reply

time).

Interpretation of Temporal Behavior Patterns

The longer and slower conversation in crisis are notable, as we might rather expect turn-taking to go faster as most everything else does in a high-tempo environment of high emergency, a time when people deal with safety-critical concerns and gather vital information for situational awareness. Furthermore, we know that conversations slow down in the *During* period not because of infrastructural failures like power outages, since we found that the overall tweeting pace picks up in this period (Kruskal-Wallis $H=1592.77$, $p \ll 0.001$, post hoc Dunn $p \ll 0.001$). Therefore, the socio-behavioral aspects of replying in crisis must be responsible for instead *slowing down the conversations*, as opposed to speeding up.

We propose that longer inter-reply times suggest that Twitter posts are referred to as an activity log or record that people return back to when possible. This is in contrast to stereotypically perceived use of Twitter as a real-time capture of things happening at the moment. It seems that in disaster, more attention is paid to the record of past tweets.

To understand this further, we analyzed at which times of day users tweet across the time periods, and found that in *During*, more people tweeted throughout the day than in other periods (Kruskal-Wallis $H=2338.54$, $p \ll 0.001$, post hoc Dunn $p \ll 0.001$). Whether due to more time to reply on the hurricane-sanctioned day off (from work and school), or because they chose to be more available for crisis-related replies during the day, increasing number of twitterers were active throughout. If those tweets were normally answered in the morning or evening when they first arrived, through-the-day replying would contribute to slower conversations.

The duration and tempo of conversations may also depend on many aspects of the conversational environment, such as how many users are communicating and what type of socio-behavioral characteristics they contribute. But before we focus on the individual users in the conversations, let us

explore the structure of the environment that partly comprises their social context—the @reply social network.

6.5.3 Relationality in Replying Activity

Intuitively, we would expect that

who is replying to whom affects the replying behavior. This is especially important in the context of disaster where these @reply conversations could be a resource for situational awareness and support.

Since we are attempting to understand what factors contributed to slowing down the conversation in *During*, we are especially interested in how relational structures of replying activity change over time.

Size and Density of Time Period Networks

Table 6.3 shows number of nodes, number of edges, and the size of the largest weakly connected component for the time period networks. More users are active in the high emergency *During* period than at any other time. It might be that users join Twitter in disaster since it is a valuable information source, but it is also likely that users who were inactive for some time came back to the platform for its informational resources. More edges—that is, the replying relationships—are also present in the *During* network.

The relative size of the largest connected component (LCC) can serve as a proxy for how densely interconnected the network is. While most of the networks are very sparse and do not have large weakly connected components, the one for *During* is quite a bit larger than those for other times. It encapsulates 29.89% of all *During*'s nodes, which is low compared to other social networks (Newman et al., 2002), but still impressive since it means that about a third of all the geo-vulnerable twitterers who were active in *During* were interconnected through their conversations alone. Such interconnectedness is usually facilitated by “bridges”—people who connect otherwise disconnected parts of the network because they talk to many disparate communities. The much more interconnected *During* network suggests that the

	Before	During	After	Long After
Number of nodes	773	1753	1212	869
Number of edges	770	1981	1261	862
Size of LCC	6.60%	29.89%	10.64%	2.76%

Table 6.3: Network Statistics for Time Period Networks

sense-making work of twitterers in disaster connects the parts of the network that do not normally talk with each other.

The *After* largest weakly connected component is also slightly larger than for *Before* or *Long After*. This suggests that some of that denser interconnectedness from the high-emergency period persists into the recovery. This speaks to the sociology of disaster concern for community resilience (Goldstein, 2012): that is, some of the interconnections between normally separate parts of the network remain after the high emergency period, carrying over into recovery (Kogan et al., 2015).

Degree and Strength of the Time Period Networks

The node degree distributions for all the four time period networks have the canonical long-tail shape and the same median degree of 2. This suggests that while some users have many reply conversation partners, at least half of the geo-vulnerable twitterers in each time period network only communicate with two other users. On average, that does not change for the *During* period, implying that the higher interconnectedness of that period comes not from users talking to more contacts, but rather from the users who become active in the high emergency period and who through their (few) contacts help to connect the network.

Though overall degree tells us about overall number of contacts, we might want to know how many others a user replies to (out-degree) and gets replies from (in-degree). These measures allow us to account for the role of *status* in online interactions, where celebrity accounts, for example, reply to a few, but are barraged by replies from fans.

Figure 6.2 shows the relationship between the nodes' in-strength and in-degrees for the *During* period. Some of the low-in-degree users have many reply interactions coming their way (in-strength) suggesting that they get replies from conversation partners very often. For example, @fraggle only talks to @JJJoe replying to him 50 times; @JJJoe responds 52 times. They are young adults who mostly talk about movies, TV shows, hockey, and feeling bored at work. But in the *During* period these exchanges turn to how Sandy has disrupted their daily routines.

On the other hand, the outlier on the right of Figure 3 is @NYCMayorsOffice, which received 87 replies from 65 twitterers in the *During* period. Tweets from this account tend to be about subway closures, evacuation notices, and other city information.

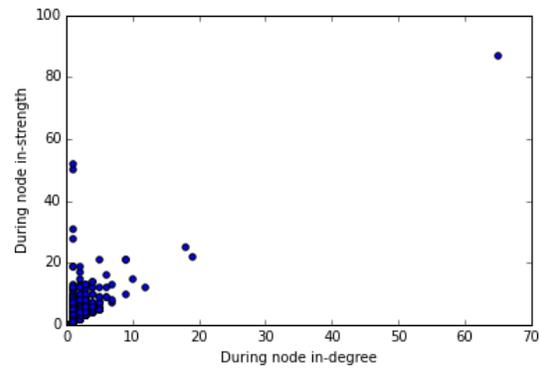


Figure 6.2: Relationship of in-strength to in-degree in *During*

Reciprocity & Degree Assortativity in Time Period

Networks

The portion of outgoing connections that are reciprocated—reciprocity (Wasserman & Faust, 1994)—is high for all the time period networks. However, while the other time periods show reciprocity of 71-75%, the *During* network has reciprocity of only 64%. It is likely that *During* reciprocity rate is dampened by the prevalence of replying to state and city official accounts, which do not tend to reciprocate. For example, all those replies sent towards to @NYCMayorsOffice in the *During* period did not receive a single response back.

More insight into the network structure can be gleaned from the mixing patterns—what types of nodes tend to connect to each other (Bliss et al., 2012; Newman, 2003). Degree assortativity quantifies the strength of the relationship between the degree of the network’s nodes and that of their neighbors. Interestingly, it changes drastically over the four time periods.

The trend is especially clear in the in-in directed assortativity. After somewhat high positive assortativity in the *Before* period (0.24), *During* and *After* show weak disassortative mixing (-0.07 and -0.02). This suggests that in the *Before* period, the geo-vulnerable twitterers who received replies from many contacts were more likely to get them from others like them in that respect. But in the *During* and less strongly in *After*, this dynamic shifted, with users who got replies from many contacts more likely to get them from those who are replied to by a few. This suggests that during the event, people were less

likely to talk to others with a similar number of conversation partners, and their communication patterns became less homogeneous. In *Long After*, assortativity stabilized back to positive (0.19).

The event-related dip in in-in assortativity can again be explained by geo-vulnerable users actively responding to the official accounts such as @NYCMayorsOffice, which has high in-degree, while people replying to it do not. Accounting for the edge weights fortifies this trend.

Summary of Relationality in Replying Activity

In summary, in the *During* period we find a much more interconnected network that is likely brought together by newly participating users who connect different communities. The popular government and news accounts, with whom many geo-vulnerable users are communicating, are also likely contributing to connecting the network. In general, in the high emergency *During* period, twitterers talk more outside their normal conversational patterns, also likely making the network more tightly interconnected. Some of this interconnectedness persists into *After*. The interactions more outside of the normal conversation patterns, including increased communication with government and news accounts, is one of the factors contributing to the slower conversation pace in *During*.

These results provide some structural insight into how conversations happen over the course of the disaster. However, these time period networks are based on the instantaneous reply events and their frequencies. We should incorporate these insights into the conversation-level analyses that are better suited to represent the highly temporal and longitudinal nature of social media interactions. Thus, the next sections of the findings rely on Twitter conversations—sequences of tweets replying to each other—as the primary unit of analysis.

6.5.4 Conversationalist Composition

We return to the longitudinal nature of conversations, and analyze them with respect to the twitterers involved.

By Number of Participants

When we break down the conversations by the number of twitterers who participated in them, we find some interesting insight into @reply conversations. First, the surprising 25% of all the conversations consist of one user replying to his or her own tweets. These are mostly two-tweet sequences where the second is the correction for a small mistake or a typo in the first, not a factual correction, which are rare on Twitter according to Arif et al. (2017). Most self-corrections happen very rapidly, but in some cases twitterers reply to their own tweets to elaborate or expand a point; those monologues may take any amount of time, as this example illustrates:

```
2012-11-02 20:56:24, annasuenyc:
Power is restored, traffic downtown looks like #LA right now. Not moving at
all, absolute stand-still. #NYC http://t.co/erI7IJ68'
2012-11-03 22:02:10, annasuenyc:
@annasuenyc Just realized, all the cars are sitting waiting to get to gas
from station on East Houston that was out of power all week. Wow.
```

The two-person exchanges that we typically associate with a conversation comprise 73% of conversations among the geo-vulnerable twitterers. And only 2% are made of reply exchanges involving more than two people (Table 6.4).

Table 6.4 also shows that the conversations twitterers have with themselves are quite a bit shorter than the other types of conversations. Moreover, the median conversation turntakes are also much shorter for monologues than conversations with more participants. This makes sense, as these twitterers need not wait for another party for a reply. The two-person conversations (dialogues) are not only longer than monologues, but also on average consist of more tweets. The longest dialogue is @fraggie and @JJJoe’s 43-step discussion about hockey and tv shows, with concerns about Sandy surfacing at the end.

	One person	Two people	More than two people
% of conversations	25%	73%	2%
Med conversation length (Kruskal-Wallis H=261.30, p<<0.001)	234	964*** p<<0.001	1287*** p<<0.001
Med (90th %tile) number of tweets	2 (2)	2 (4)	4 (6)
Med time span between replies (Kruskal-Wallis H=116.87, p<<0.001)	226	515*** p<<0.001	770*** p<<0.001

Table 6.4: Types of Conversations by Number of Users
Post hoc Dunn test comparing to monologues with $\alpha=0.05$

The group discussions—conversations with more than two participants—tend to be longer than conversations with fewer users. The number of tweets in these conversations is on average higher than in others. Moreover, turntakes for group discussions and dialogues are longer than in one-person conversations (Table 6.4). This points to the differences between individual and social behavior. In this interaction, @mandyOdoe says in the mid- afternoon:

```
2012-11-04 14:43:22, mandyOdoe:  
Hey. East Village. Did we lose service again? Using Verizon and I only have  
1 bar @villageeye @VZWSupport #SandyKeepsOnGiving.
```

But by the time @marthasandrews responds to her and an intervening tweet by @gardenlily, it is almost midnight.

```
2012-11-04 20:53:13, gardenlily:  
@mandyOdoe @villageeye @VZWSupport I'm having the same problem. Dipped down  
to 1 or 2 bars, can't text or make a call. #sandy #eastvillage  
2012-11-04 22:34:33,marthasandrews:  
@mandyOdoe @villageeye @VZWSupport Sandy doesn't want to leave us. At least  
I got out of an annoyingly long phone call
```

Since we are interested in factors that contribute to slower conversations in the *During* period, we next consider how these conversation types are distributed across time periods. The monologues are less frequent in *During*, comprising only 19% of all *During* conversations, compared to 27-31% in other periods. The dialogues are represented fairly evenly across time periods. And out of all the group discussions, a surprising 57% fall into the *During* period, overrepresented compared to 39% of all conversations that happen in that time period ($\chi^2=68.13$, $df=6$, $p\text{-value}\ll 0.001$). It seems that coordination needs of sense-making in disaster prompt more group activity, although some of these are several users replying to @NYCMayorsOffice. Furthermore, since group discussions tend to be slower and longer, their high presence in *During* likely contributes to both longer conversations and turntakes in this period.

By Popularity of Conversationalists

Because conversations with popular users such as celebrity accounts, government accounts, and news accounts seem to be a factor in the slower *During* conversations, next we analyze the conversation composition with respect to number of followers, which can serve as a reliable proxy for popularity (Cha

et al., 2010; Kwak et al., 2010). From here on, twitterers with a high number of followers are referred to as “popular.”

Conversations with no popular users are strongly overrepresented among the monologues, where they comprise 77% of the conversations, and are underrepresented among the dialogues ($\chi^2=1114.81$, $df=2$, $p \ll 0.001$). Here, the conversations that include popular twitterers are overrepresented, comprising 74% of the dialogues. Group discussions are almost exclusively comprised of the conversations with popular users (92%). So it appears that popular news and government accounts tend to avoid the self-correction that generates most monologues and engage in two-way or group interactions.

The conversations with popular twitterers also tend to be quite a bit longer than those without (med=856 seconds vs. 507, Mann-Whitney $U=2948800$, $p\text{-value} \ll 0.001$), while both types have the same distribution of number of turntakes. Thus, the conversations with the popular users go slower, with median inter-reply time higher than in the conversations without popular users (514 vs. 321 seconds, Mann-Whitney $U=2989563$, $p\text{-value} \ll 0.001$). An example of @OccupySandy, which emerged in response to the hurricane, illustrates this by taking 10 hours to respond to @HumanityRoad, likely due to its volunteers being busy with more pressing tasks.

```
2012-11-13 15:54:58, HumanityRoad:  
MT @OccupySandy Jacobi Church and @880FultonOS need more drivers over the  
next couple days to move supplies. #sandy #sandyvolunteers #hmrdr'  
2012-11-14 02:26:27: OSWarehouse:  
@HumanityRoad Thank you for helping get word out about the need for  
drivers. We need 6+ vans an hour but settle for bikes, #peoplepower!'
```

For conversations with no popular twitterers, distributions of conversation length and inter-reply time show clear work sessions peaks. And for conversations including popular twitterers, these distributions are more like Figure 6.1c, with a long tail. Thus we propose that popular users slow down the conversation pace since, after a certain turntake length, the timing of their replies depends less on the timing of the original than on their own reply patterns. Since conversations with these popular users, like the official and news accounts, are overrepresented in *During* ($\chi^2=28.66$, $df=3$, $p\text{-value} \ll 0.001$) when geo-vulnerable users engage with the relevant information they provide, they may contribute to slower and longer conversations in *During*.

6.5.5 Conversation Content

In addition to who participates in the conversation, the matter of what they discuss is likely to affect the practices and timing of the conversational back-and-forth. Thus, we coded all conversations as off- or on-topic: discussing the experience, consequences, or recovery efforts for Sandy.

Overall, out of 5350 conversations between the geo-vulnerable twitterers, 1895 (35%) are about Sandy. There are no on-topic conversations in the *Before* period, as most people were not aware of the impending storm until the beginning of the *During* period. This is when on-topic conversations are overrepresented (72%, vs. 40% of all conversations, $\chi^2=1511.62$, $df=3$, $p\text{-value}\ll 0.001$).

The conversations about Sandy are longer than off-topic ones, with median duration of 858s. compared to 671s. (Mann-Whitney $U=3040391$, $p\text{-value}\ll 0.001$), while the number of turntakes is the same on average. This means that the turntakes are longer in on-topic conversations than in off-topic (med=499 vs. 402 seconds, Mann-Whitney $U=3066680$, $p\text{-value}\ll 0.001$). And since on-topic conversations are mostly present in the *During* period, the discussions of the hurricane are contributing to slowing down of the conversation tempo in *During*. In the example bellow, @hgcellar starts a warning-period conversation about Sandy preparation in the early evening:

```
2012-10-26 17:12:41, hgcellar:  
@FreddieJensen Well, looks like I live across the street from the  
evacuation zone. Should I be relieved? #hurricanenewbie #nywx #sandy
```

As they both work on preparing for the worst, the long first turntake makes the conversation go into the late evening.

```
2012-10-26 20:01:03, FreddieJensen:  
@hgcellar Biggest threat for burbs: power outage. Turn fridge/freezer low,  
store lots of water in jugs+bathtub, full tank gas, lots of MREs
```

```
2012-10-26 20:04:52, hgcellar:  
@FreddieJensen Check, check and check! :) Got the last D batteries at  
Target, folks who waited until tomorrow to prep will be out of luck.
```

On-topic conversations are also slower when twitterers take their time to answer a question to the best of their ability:

```
2012-11-02 01:52:03, gansta0718:  
@mtracey635: @TheBoken still can't find my mom no shelters answer phones  
looking for Meredith Tracey (mimi) 530johnson st#Hoboken
```

2012-11-02 03:03:48, TheBoken:
@gansta0718 @mtracey635 I sent your request for Mimi to a City Official
with a propensity to follow up.

On-Topic Conversations and Frequent Conversationalists

Frequent conversationalists are twitterers who both reply and are replied to frequently across their contacts. The technical definition of this concept can be found in the Methods section. We find that conversations with two frequent conversationalists are much more common (43% of all conversations) than those with one (29%) or zero (29%). This is intuitive since this measure is accounting for a pairwise replying activity. Thus, we divided conversations into those with two or more frequent conversationalists and those with less than two.

The conversations with two or more frequent conversationalists are slightly less often about Sandy (33%) than the conversations with less than two frequent conversationalists (37%, $\chi^2=12.22$, $df=1$, $p\text{-value}<0.001$). This is surprising, as we would expect twitterers to rely on their established Twitter relationships to navigate in crisis. However, the official and news accounts that supply a lot of relevant on-topic information tend to not be frequent conversationalists, placing conversations involving these accounts into less-than-two-conversationalists category and increasing the proportion of on-topic conversations there.

But even excluding conversations with popular users, we still find that their prevalence alone does not explain why more-than-two-conversationalist discussions are less often about Sandy. This is especially clear across the time periods. In the *During* period, when most on-topic conversations happen, only 25% of the conversations with two or more conversationalists are about Sandy, compared to 40% for fewer than two ($\chi^2=8.31$, $df=1$, $p\text{-value}=0.004$).

Analyzing conversations qualitatively suggests that, in crisis, people tend to check-in more broadly across their local network, even explicitly asking for or offering help in broadcast fashion—tweeting to their entire network:

2012-11-03 02:48:24, Ken_TY:
Does anyone have info on Coney Island or Brighton Beach volunteer locations
for tomorrow? #SandyHelp #brooklyn'

2012-11-03 11:25:22, TSPayForward:
@Ken_TY check out <http://t.co/AGGp2qDL>
2012-11-03 14:21:03, TSPayForward:
@genevivelarue @Ken_TY thanks for helping get the word out!!
2012-11-03 14:29:46, Ken_TY:
@genevivelarue no problem at all

We saw broader conversational reach in the network findings; we further confirmed it by examining how consistently users talk to each other from one time period to another. We find that there were 603 unique groups (mostly pairs) who had a Twitter conversation in the *Before* period. This number grew to 1516 unique groups in *During*, only 206 of which were the same as in *Before*. Thus, with more users being active in the high emergency period, there were more new user pairings, 75% of which included less than two frequent conversationalists. Hence, the broader reach in this period means that frequent conversationalists were paired off with users who did not have the history of the conversational back-and-forth, likely also contributing to the slower tempo of discussions in the *During* period.

6.6 Study Summary & Discussion

We have analyzed the Twitter reply behavior conversationally, promoting the geo-vulnerable twitterer's monologues, dialogues, and group discussions to the analytical center. Further, we focused on the structure and temporality of these Twitter conversations, elucidating user behaviors around the reply exchanges that constitute them across four time periods. A focus on long form discussion allows analysis of how intensive information activities interactionally unfold among those affected by the event.

6.6.1 Temporality

We found that twitterers' behavior in the *During* period of high emergency differed considerably from the other periods. In *During*, conversations were longer and slower-paced than in other periods, despite the high-tempo, high-pressure nature of crisis. Moreover, conversations slowed down not because users tweeted less frequently due to infrastructural constraints, but rather due to socio-behavioral factors. In contrast, Wikipedia research documents accelerated and bursty qualities of collaborative editing for high-tempo crisis topics as those events unfold (Keegan et al., 2013). Sociologists also note that activities and information flows speed up in crisis contexts (Weick, 1995), especially as related to the phenomenon

of convergence (Dynes, 1970), where people, resources, and information coalesce around physical or digital sites of disruption. Here we have found that conversations in Twitter are different: they are slower during the high-emergency period. However, Weick reminds us that, “a growing issue for sensemaking is the disparity between the speed and complexity of information technology and the ability of humans to comprehend the outputs of the technology” (Weick, 1995), which might explain this result. Additionally, creativity and improvisation, which are central sensemaking processes in response to collapsing organizational order (Weick, 1993) also take time to emerge, making the slower Twitter conversations during high emergency less surprising.

6.6.2 Relationality

In addition, the conversations of the geo-vulnerable twitterers were broader in terms of who spoke to whom in the *During* period. Twitterers conversed with a wider variety of contacts, creating a more interconnected reply network that joined previously disparate communities. This interconnectedness persisted somewhat into *After*, indicating perhaps mechanisms of “community resilience” at work, where social ties acquired in one crisis might provide social capital for weathering the next (Goldstein, 2012), and extends prior sociological findings to an online context.

Additionally, in other periods the geo-vulnerable twitterers tended to converse with others similar to them in terms of number of contacts, but in *During* users with a few conversational partners talked more often with popular users, including news and government accounts. Thus, in crisis, twitterers engaged with types of contacts that are outside of their normal patterns of interaction, validating Granovetter’s assertion of importance of weak ties for the crisis context (Granovetter, 1983).

This broader engagement includes popular news and government accounts, which may not be viewed by some as sources of “organic” collaborative sensemaking. However, sociologists of disaster point out that well-defined and formalized social roles—such as government officials—provide needed structures for sensemaking activity in the crisis context (Weick, 1993). This broader engagement, with both these popular and new other (non-celebrity) conversational partners, also contributed to the slower

conversation pace the in *During*. This suggests that popular users may be better able to “hold the floor”—that is, draw attention that encourages others to revisit their content over a longer period of time. For further verification, when we removed popular users’ replies, we confirmed that slower turntakes in the *During* period are *not* explained simply by a slower response of these high-status accounts, indicating that other factors are at a work in the temporalities of these exchange.

6.6.3 Conversationalist Composition & Content

The *During* conversations were also broader in terms of how many users participated. We know that social configurations of participants, including simply their number, affect the dynamics of the social interaction (Simmel, 2010). Monologues were less common, and the group discussions were much more prevalent, likely due to the coordination needs of the high emergency period. As conversations broadened in this way, this contributed to the slowing and lengthening of conversations in *During*. Conversations about Sandy, which were slower than conversations about other topics during this time, also contributed to this slower tempo in *During*, as geo-vulnerable Twitterers took the time to answer questions as accurately as they could.

6.6.4 Implications for Design

The longer inter-reply times suggest that in the high-emergency period, tweets are referred to as a record perhaps for longer term reference, challenging notions of ephemerality of Twitter communications as a real-time sensor. In disaster conversations, it seems that more attention is paid to the record of past tweets. Design of user interfaces that account for such longer term engagement with the Twitter record would better support sensemaking activities during the information convergence of high-tempo events, where people must attend to a variety of tasks and concerns that extend the pacing of their participation in conversations. For example, an interface that facilitates the ability to come back easily to certain parts or junctures in past conversations—perhaps enabling a conversational branch curation rather than of the tweet stream—would be helpful in the high emergency context.

The current Twitter interface connects tweets that reply to each other with a vertical line, visually representing some of the back-and-forth. Yet these conversational snippets are placed in the same linear layout, not visually accounting for the branching nature of many conversations. Representing the discussion as a branching tree would be more intuitive and useful, especially if accompanied by the ability to curate certain branches. Moreover, a mechanism for suggesting external resources in response to the questions in the conversation would facilitate the question-answering behavior that we found often slows down the sensemaking.

Our findings suggest that those affected require interfaces that support the broader reach of their conversations. For example, such interfaces could suggest new conversational partners who could be helpful in sensemaking activity. Weak tie connections of those affected could be offered as candidates for the conversations based on their locality, tempo of their conversations, or perhaps specific topics they have recently discussed. Local government and news accounts that are likely to respond to those affected could be recommended as potential information sources.

6.7 Beyond the Tweet

This work expands social media research from analyzing single Twitter posts to considering whole conversations by placing monologues, dialogues, and group discussions of Twitterers who are geographically vulnerable at the analytical center. In offering approaches for making conversations analytically available, this work contributes to social computing research for crisis informatics that can now investigate more nuanced questions about the social roles that twitterers play, how they draw on social capital in protective decision making, and how online conversations interweave with the activities on the ground.

6.8 Forms of Sociality and Two Aspects of Shared Information Space

In this section, I take the opportunity to reflect on how framing the analyses around the two intersecting aspects of the shared information space—explicitly shared site of work and human-readable

record of work—helps in interpreting this study’s findings and positioning them within the larger context of highly-distributed collaborative work in the high-tempo environment.

6.8.1 Collaborative Work Practices

As discussed in the pre-print above, collaborative practices of the Twitter users affected by the hurricane Sandy led to broader conversations in terms of who spoke to whom in the high emergency period. Twitterers conversed with a wider variety of contacts, creating a more interconnected reply network, joining previously disparate communities. This broader reach of conversations in high-emergency period includes both popular news and government accounts and new other (non-celebrity) conversational partners. Somewhat similarly to the broadening of the information sharing networks in the Retweet activity, geo-vulnerable Twitter users also expand their conversational networks in search of new information, access to resources, and emotional support. This suggests that in the time of crisis, affected twitterers do not rely solely on their established ties (such as their friends and follower networks), but find new connections in the Twitter deluge despite the absence of clear site of work within the platform, likely through reliance on the ad hoc publics of hashtags (Langlois et al., 2009; Bruns & Burgess, 2011; Ellison & boyd, 2013).

6.8.2 Network Structures

The clear record of the back-and-forth of a conversation produced through @reply convention facilitates distinctive network structures as a result of this activity. The reply networks for all time periods have high rates of reciprocity (60-70%), suggesting that the back and forth is easy to track in the Twitter @reply record and the geo-vulnerable users actively took advantage of this feature. Reciprocity is the lowest in the *During* period, when those affected engaged most with the popular government and media accounts, which were less likely to reciprocate with a reply.

Interestingly, the highly pro-social behavior in dyadic relationships (reciprocity) does not translate into a pro-social behavior within triads (transitivity). Since replying is by definition a directed activity, replies of the users to whom one replies are not necessarily visible, unless one also follows those

accounts. As mentioned above, in the high-tempo high-volume setting of disasters arising from natural hazards, the geo-vulnerable Twitter users expand their conversations beyond their normal patterns of interaction and likely seek information, resources, and support from additional sources outside of their following network. This broadening of the conversations, in combination with Twitter not providing a well-defined, explicitly shared site of work produce and environment where it is rare to respond to an account that is replied to by someone to whom you also reply—due to their lack of visibility in this very loosely defined and highly distributed site of work.

6.8.3 Organizational Forms

Organizational forms as reflected in number of participants and the social roles they take on in the sensemaking activity are also important forms of sociality that emerge in the process of high-tempo crisis-related activity. As discussed in the pre-print, a social role that became prominent in the analysis of replying activity is that of the official government or media account. Considering the issues of visibility in this highly-distributed site of work mentioned above, the popular government and media accounts hold a special advantage because of their prominence and reach in this loosely organized online environment. Thus, it is not surprising that the geo-vulnerable twitterers engage with these popular accounts more in the height of the disaster than at other times, since these sources are highly visible despite the lack of clearly-defined site of work even in the chaotic high-emergency period. Such prominence and visibility allow the official government and media accounts to even “hold the floor” longer, slowing down the conversations in the height of disaster.

CHAPTER 7: Discussion

7.1 Introduction: Emergent Forms of Sociality in Social Media

As described in the introductory chapters, the active participation of those affected in recovery efforts (Fritz and Mathewson, 1957; Dynes, 1970; Kendra and Wachtendorf, 2003), together with the phenomenon of convergence (Fritz & Matthewson, 1957), result in new relationships being established, new sources of information found, and new communities being formed. Furthermore, Weick suggests that sense making is a critical process in which people engage after a major disruption (Weick, 1993). Information seeking, communication, and coordination as part of sense-making process also lead to new connections and emergent community structures.

Thus, this thesis has focused on the emergent properties of social media in disaster and how individual behaviors “add up” to social dynamics. Specifically, the thesis revolves around the emergent forms of sociality that arise in social media activity during the disasters arising from natural hazards. The forms of sociality that I focused on, based on the social science and community informatics literature, are collaborative work practices, organizational forms, and social network structures.

Additionally, various social networking sites provide different affordances for the social interaction (boyd, 2010), and thus for emergence of new forms of sociality. Moreover, various platforms and activities within a single medium might fulfill different communication needs, and so people use them for different purposes, relating to the concept of media multiplexity (Haythornthwaite, 2005; White & Palen, 2015). Thus, various social media platforms and activities within them provide different affordances for the emergence of new forms of sociality. Therefore, to organize this multifaceted space, in Chapter 3 I offered an operationalization of two intersecting aspects of the shared information space salient to the high-tempo, high-volume context of social media activity in crisis. I then organized the social media platforms and specific activities within them in this dissertation with respect to how well they facilitate creation of the shared information space based on how explicitly shared the site of work is and how visible and human-readable record of activity is.

In this chapter, I synthesize the findings from the three empirical studies (Chapters 4-6), focusing especially on how the framework of the two intersecting dimensions of the shared information space reveals larger patterns within the emergent forms of sociality across contexts, and what other analytical advantages it offers. I highlight collaborative practices, organizational forms, and social network structures that emerged in each social media context and discuss how they relate to the affordances of the social media platforms and specific activities within them, and specifically to the two intersecting dimensions of the shared information space. In the final part of this chapter, I take the opportunity to generalize from these findings and discuss how the framework developed in this thesis could be helpful in other high-volume, high-tempo social media contexts and how it relates to the CSCW concepts more broadly.

7.2 Emergent Forms of Sociality and Two Dimensions of Shared Information Space

7.2.1 Collaborative Work Practices

Different affordances of the social computing platforms that facilitate creation of shared information spaces produce an interesting range of collaborative practices across the three studies. In crisis mapping with OpenStreetMap after the Haiti earthquake (Chapter 5), the clearly defined and spatially bounded site of work facilitates a spatially-defined shared information space that successfully attracts contributors from all over the world. This international participation in the same digital map space then relates to temporal practices of shift work, where mappers from different continents incorporated the crisis mapping of Haiti into their existing daily routines and often edited the affected area in shifts based on their respective time zones.

The spatial practices found in the post-earthquake crisis mapping (Chapter 5)—mapping externally salient features, mapping continuous structures, and filling in the map—largely emerged early on also as a result of the highly concentrated and spatialized nature of crisis mapping. On the other hand, the interpersonal practices—parallel mapping and tag correction—took longer to develop after the earthquake, since the lack of the human-readable record of work in OSM does not actively facilitate

passive awareness required for the back-and-forth inherent in these practices. Nevertheless, in the qualitative analysis of mapper contributions and especially in the interviews with the mappers, I found many examples of less experienced mappers learning the norms and policies of the OSM community through legitimate peripheral participation (Lave and Wenger, 1991) and more experienced contributors mentoring the newcomers on and off the map. These types of interpersonal practices simply took longer to develop, since they were not actively facilitated due to the lack of clear record of who has done what on the map.

While I found some similar exchanges about the community norms—specifically the etiquette of managing damaged belongings—in the @reply conversations of those affected in Hurricane Sandy (Chapter 6), my structural analysis of these exchanges is especially revealing of the temporal practices of their participants. While Twitter does not provide a clearly-defined site of work, the @reply record generated by these conversations and user interface notifications accompanying them allow the geo-vulnerable to treat it as a log, to which they can return when time permits. This use of Twitter conversations as less of a real-time sensor and more of a persistent log is especially prominent in the height of the emergency, when pressing on-the-ground concerns and answering questions to the best of their ability necessitate longer turntakes in the Twitter conversations. Thus, the explicit, human-readable record of @reply back-and-forth allowed the geo-vulnerable Twitter users to develop new temporal practices with respect to replying, which were more congruent with their disaster context. It allowed the affected Twitter users to return to the shared information space of crisis-related sensemaking as their circumstances allowed, despite the lack of clearly-defined shared site of work.

On the more macroscopic scale of temporality, all three analyses are based around several time stages, which were selected on the basis of the theory of the disaster phases (Powell, 1954). The two Twitter studies—of the Retweet and Reply activity (Chapters 4 and 6, respectively)—show that in the period of high emergency those affected by Hurricane Sandy seek out new sources of information and new types and configurations of conversation partners, opening up and broadening their respective networks to new types of participants. Specifically, Twitter users change their sensemaking practices in

disaster to include the types and configurations of partners, with which they do not normally engage as much—in search of new information, new resources, new pathways through the networks, and new ways to leverage their social capital, because old reliable social connections are often unavailable or less useful in the shifting context of disaster. This might suggest that the lack of the well-defined site of work in Twitter might also offer an advantage in disaster, since not being explicitly limited by it, users easily change their collaborative practices to more actively embrace new information sources and types of conversation partners useful in navigating the high-tempo environment of crisis.

7.2.2 Network Structures

In terms of the network structure that arises through these new patterns of participation, all three studies provide evidence that the resulting social networks—retweet network, overlapping changeset network in OSM, and Twitter @reply network—densify in the high emergency period, albeit to a different degree. Specifically, all three networks show higher fractional size of the largest strongly connected component—a proxy of how interconnected the network is that denotes what portion of the network provides a (directed) path from one participant to another—in the height of the disaster than before or after.

While some of this higher interconnectedness is undoubtedly the result of broader participation and the high-tempo convergence around the digital site of disruption at the height of the disaster, I employed a variety of methods across the studies to ensure that the densification we observe is not purely a function of the convergent behavior. In the retweet study (Chapter 4), I analyzed the retweet activity of only the Twitter users who were active across all four time periods, ensuring that the identity and the number of core retwitterers was consistent across all time slices and only their information sources (number and identity) varied over time. In this case, a broader range of and more numerous information sources, as well more retweeting activity, led to the more interconnected network in the high-emergency period. In the OSM crisis mapping and Twitter conversational activity studies (Chapter 5 and 6, respectively), I controlled for the varying sizes of time slice networks by comparing the network statistics

to their distributions under the configuration model and testing how statistically different the observed network metrics were from those expected under the model.

Since we observe the densification in all three networks—across the range of affordances produced by the different levels of explicitly shared site of work and human-readable record of activity,—it suggests that the active participation in disaster manifests similarly (though to various degree) across the platforms and activities with variously successful shared information spaces. On the other hand, the fact that some of this interconnectedness persists somewhat after the high-emergency period for the Twitter reply networks (Chapter 6), but not the overlapping changeset network in OSM (Chapter 5) or the Retweet network (Chapter 4), suggests that the visible, human-readable record of @reply back-and-forth is especially useful in maintaining the connections within the network after the immediate emergency period, contributing to the “organizational memory” (Ackerman, 1998) of the affected populations and the related sociological notions of “community resilience” (Goldstein, 2012).

While all three networks show more interconnectedness in the height of the emergency, the specific configurations and network structures that emerge in this more interconnected activity are somewhat different for the three sets of affordances surrounding each activity. The retweet networks exhibit low reciprocity and low transitivity across all time periods. In the absence of clear record of the path each tweet takes to propagate through the network (Twitter only displays the original source tweet), the directed nature of retweeting is not conducive to producing the two-way visibility that would facilitate reciprocity in social interaction. Similarly, the low transitivity of the retweeting networks indicates that lack of shared site of work and visible record of activity do not produce the shared information space that would facilitate mutual visibility needed for a source of a source to become a source. Instead, this combination of affordances is geared much more towards information propagation, which is also important in the high-uncertainty context of disaster, but is less conducive to the mutual visibility that allows for more distinctly pro-social engagement (reciprocal or triadic) of the Twitter users.

All of the overlapping changeset networks for OpenStreetMap crisis mapping after the Haiti earthquake exhibit medium levels of reciprocity and transitivity. Here, the well-defined site of work—the

map of the affected area—placed contributors into a spatially-constrained digital environment where they were bound to overlap with each others’ work, simply because they were working in the same space. On the other hand, considering the lack of clear human-readable record of their contributions, it is likely that many of these overlaps are not necessarily purposeful—in the sense of replying to someone’s edits to an area by adding your own, —but likely rather have to do with the spatiality of the map, although some of the more purposeful, interpersonal practices of interaction do emerge later in the analysis period (as mentioned in the Collaborative Work Practices section 7.2.1 above). However, due to the lack of the human-readable record of the contributions, these types of more purposeful engagements are still rare enough, even in the later stages of the response, to produce decidedly pro-social network structure patterns such as high reciprocity and transitivity.

On the other hand, the reply networks of the geo-vulnerable Twitter users in the wake of the Hurricane Sandy exhibit the patterns of high reciprocity but low transitivity. The high reciprocity suggests that the @reply convention serves a clear, accessible record of the conversational back-and-forth in the sensemaking activities, making it easy for those affected to reciprocate in their conversations. Interestingly, the highly pro-social dyadic behavior, which translates into high rates of reciprocity, did not coincide with the highly pro-social behavior among triads of users—transitivity. The lack of well-defined site of work in Twitter, besides the ad hoc publics defined through hashtags (Bruns & Burgess, 2011) and follower networks (Langlois et al., 2009), together with the directional nature of replying that precludes visibility of conversational partners of one’s conversational partners unless one also follows them, produce a particular type of environment. In this context, responding to an account that is replied to by someone to whom you also reply is rare because of their lack of visibility in this loosely defined and highly distributed site of work. Thus, one aspect of the shared information space—visible record of work—is strongly present in the context of Twitter replying, facilitating a highly pro-social dyadic engagement (reciprocity). However, this aspect alone is not sufficient to produce a highly pro-social group engagement (transitivity), as directionality of replies and lack of visibility beyond one’s friends

network precludes the development of the other aspect of shared information space—well-defined shared site of work.

OpenStreetMap interface does not provide a human-readable record of the history of contributions, as all the edits are instead stored in a behind-the-scenes geo-spatial database. Nevertheless, based on the patterns observed across the existing intersections of the two aspects of shared information space, we could imagine what such environment may look like in terms of social interaction. For example, when a contributor is editing a map object, the interface could highlight the last 5 edits to this object, who produced them, and when. Another option may be highlighting some history of edits in the area adjacent to the object currently being edited. As the site of work is well defined in OSM crisis mapping, including in the OSM user interface a human-readable and easily-accessible log of all, or more tractably some, prior contributions would further facilitate the distinctly pro-social aspects of interaction.

While in the analysis of activity after the Haiti earthquake I found that the more purposeful, interpersonal practices of interaction do eventually emerge within the current interface, adding a human-readable record of activity to the interface would likely allow such interactions to emerge much faster and easier. I imagine that mapper-to-mapper mentoring, learning both technical and interpersonal norms of the OSM community, and various forms of legitimate peripheral participation would be actively facilitated by such interface. Novice mappers could find role models more easily within the map activity itself, instead of reaching out to experienced mappers, as they do currently, on the OSM wiki and other off-the-map channels. In addition to the spatially-motivated interactions I found in post-Haiti OSM activity, more purposeful, intentional co-working interactions would emerge sooner in the process of crisis mapping. The recent addition of Notes feature, where contributors can annotate the map with comments, which in turn can be answered, combined with this imagined human-readable log of activity would be especially beneficial for training the novice mappers and thus improving the quality of contributions in the high-stakes environment of crisis. These types of more purposeful interactions would produce higher levels of explicitly pro-social markers in the social network structure, such as reciprocity and transitivity.

Moreover, from the researcher perspective, having a human readable record of OSM contributions would make gathering and analyzing the traces of computer-supported cooperative work in OSM crisis mapping much more efficient, as it would eliminate the need to rebuild the entire history of OSM edits from the geo-spatial database in order to select the activity of interest.

7.2.3 Organizational Forms

In all three studies, social roles featured prominently as an important aspect of organizational forms. While the specific roles that were central to each environment differed, some similarities were clear. For both Twitter activities—Retweet and Reply—the government and media accounts were important players at the height of the disaster. In the highly-distributed environment of Twitter that does not provide a well-defined site of work, well-known and popular sources like the government agencies and the media stood out in the uncertainty of crisis, partly because of their outsized visibility, making them easier to turn to in this highly-fractured shared information space. For the OpenStreetMap, which is both a map and a content-producing organization, incorporating new contributors is a central task, especially in the periods of intense growth following disasters. Thus, the roles of experienced and novice mappers emerged as important organizational features in the social practices that arise.

The organizational typology based on the structural codes, as denoted by Kreps and Bosworth (Kreps and Bosworth, 1994) is also relevant here. As a reminder, depending on the order in which the four structural codes become present—Resources, Domains, Activities, and Tasks—organizations in disaster are shaped into a variety of forms, broadly outlined by four types: established, expanding, extending, and emergent. For example, formal organizations manifest structural goals before means, with the Domains and Tasks appearing before Activities and Resources. On the other hand, emergent organizations arise when structural means appear first, such as Resources and Activities predating Tasks and Domains. Then communities that emerge around Retweet and Reply activities in the wake of disaster are clearly emergent organizations—specifically in that structural means such as Activities and Resources surface first, before the structural ends of Domains and Tasks. With the crisis mapping in

OpenStreetMap, however, the crisis-related activity could be viewed as extending OSM into a new Domain of crisis mapping—and thus necessitating the creation of the Humanitarian OpenStreetMap Team (HOT)—as well as expanding existing organization of OSM with new members. Thus, incorporating, mentoring, and communicating the community norms and technical specifications to the novice contributors is one of the central tasks for OSM in the crisis context, when the convergence of new volunteers is both a valuable Resource, but also an organizational challenge to this distributed community. As mentioned in Section 7.2.1, the well-defined site of work in OSM—map of the affected area—was somewhat helpful, eventually, in facilitating some purposeful exchanges needed for such incorporation of the novice mappers. However, the lack of the human-readable record of contributions made this process quite slow, as this aspect of the shared information space was not there to provide mutual mapper visibility that is essential in purposeful pro-social engagement at the heart of mentoring and incorporation of novices into an expanding organization.

7.3 Contextual Considerations: Limitations and Opportunities

The analyses that comprise this dissertation focused on social media activity within particular platforms at the specific moments in time. As the technological aspects and affordances of platforms and activities evolve over time, in this section I would like to highlight how this particular technological context affected the analyses by both limiting them in some ways and also providing opportunities for observing specific kinds of situated behavior.

This is especially relevant for Study 2, which engages with OpenStreetMap crisis mapping after the 2010 Haiti earthquake. In this study, I specifically focused on post-earthquake OSM contributions in order to investigate what practices and organizational forms arose in what could be considered “mapping in the wild.” The 2010 Haiti earthquake was the first major humanitarian event supported by OpenStreetMap, and the contributor behaviors we observed in this study were not yet constricted and formalized by the variety of technical interventions since undertaken by the OSM community. Therefore,

focusing on this historical moment allowed us to examine what can be argued are natural and emergent crisis-mapping behaviors in OSM.

As the Study 2 highlights, due to the convergent nature of crisis response some of this emergent crisis mapping involved contributors overriding each other's work, and thus not mapping the affected area in the most efficient fashion. To combat this inefficiency and potential for chaos associated with the convergence of contributors, Humanitarian OpenStreetMap Team (HOT) since introduced a number of technical interventions, the most prominent of which is the Tasking Manager. It essentially divides the area that requires mapping into orderly rectangular sections, intended to be checked out for editing by one contributor at a time to prevent collisions. This technological intervention has arguably reduced the number of unintended editing overwrites and the wasted contributor effort. On the other hand, this pre-defined orderly division of labor potentially offers less opportunity for serendipitous mapping encounters and organically-emergent collaborations. As a content-producing organization, OpenStreetMap has been grappling with this dilemma of potentially divergent priorities: focus on data quality, which is especially important in crisis context when map accuracy can be safety-critical, vs. the organic growth and development of the OSM mapping community. Another content-producing community—Wikipedia—historically prioritized content quality through bureaucratization, which often poses barriers to entry, community growth, and diversity. OpenStreetMap as a community is still in the process of grappling with this tension, and the shifting priorities are reflected in the socio-technical interventions introduced over time and thus in the affordances of OSM as a platform.

While the types of mapping interactions were likely different after the introduction of the Tasking Manager and other interventions such as the Notes feature, the methods I used to collect and analyze the crisis mapping data would be just as effective for analysis of these subsequent behaviors. On the whole, the main obstacle for understanding collaborative work in OSM still persists: the meta discussions around how the work should be accomplished with respect to community standards, practices, and conventions is still largely separated from the map itself (the primary site of work), and instead resides in the OSM wiki's and chat rooms. Thus, even after the numerous technological interventions, researchers of OSM

are still faced with the same difficulties in understanding the crisis mapping activity, where asking contributors about their mapping practices still remains the only way to ensure that we truly understand the mapper interactions reflected in the digital traces of the map.

Somewhat similarly, the affordances of the Twitter platform have changed since the 2012 Hurricane Sandy: the character limit for tweets has doubled to 280 characters, from the traditional 140. At the first glance it may be tempting to assume that the longer format might have an effect of producing more of a community gathering place, like that researchers often observe in crisis-related groups on Facebook, potentially producing more pro-social organizational structures and practices. However, despite the change in the character limit, the central affordances of Twitter are still oriented toward information propagation, with the focus on the most recent updates and a view of the Twitter feed as an up-to-date sensor. This “now-centric,” fleeting nature of the feed relates to the lack of the well-defined site of work—one of the central concepts of this dissertation. Therefore, this suggests that the character of Twitter is not likely to change based solely on a longer format: users will simply have more space to continue with the same kind of now-centric interactions.

Finally, the methods I chose to utilize in this work produced a particular view of these technologically-enabled social interactions. This is certainly one particular view among many. If I performed a deep qualitative analysis, I would have been able to arrive at a more in-depth, nuanced understanding of very specific activities within particular environments. However, utilizing only qualitative approaches it would be more difficult to lift out patterns and commonalities and draw conclusions across the range of affordances. On the other hand, if I only utilized large-scale quantitative methods, I would likely have been able to tease out more large-scale generalizable trends and possibly make some causal inferences in an attempt to disentangle social phenomena from the technological ones (affordances of the platforms). While I am interested in engaging with the causal questions prominent in this context, I nevertheless fully embrace the sociotechnical lens with respect to these phenomena. I view these platforms and activities as sociotechnical systems where the sociality and its technological context are thoroughly intertwined. Methodologically, the quantitative aspects of my

analyses enable me to glean larger patterns across the affordances in the studies, and thus engage with the broader landscape of technical affordances in social media. Simultaneously, the qualitative analyses enable me to contextualize and interpret the inferred larger patterns, glean the decision-making and in some cases motivations of the participants, thus bringing this research closer to understanding social behavior in all its situatedness.

7.4 Beyond the Three Studies: A Broader Vision

The synthesis in section 7.2 highlights how the conceptual framing relying on the two axes of the explicitly shared site of work and human-readable record of activity help us bring into relief aspects of social media platforms and activities within them that are more or less conducive to the creation of a shared information space, and thus to emergence of various forms of sociality. I believe these two intersecting aspects of the shared information space provide a successful framework, because, to some extent, they define and operationalize classic CSCW concepts relating to the creation of shared information space for the highly distributed, decentralized, and high-tempo, high-volume environment of social media. The shared site of work relates to the issues of peripheral awareness (Dourish & Belotti, 1992), and the accessible record of work reflects and facilitates articulation work and resulting division of labor—the concepts at the heart of CSCW (Schmidt and Bannon, 1992). Thus, defining and fleshing out these concepts for the environment of social media—especially with the data deluge that accompanies a high-tempo, high-volume crisis event—provides concrete ways of positioning the affordances of various platforms and types of activities within them in the larger CSCW context, bringing commonalities and tensions into relief, and doing it at scale.

This type of conceptual framing, which segments the vast social media landscape into easily recognizable quadrants based on detectable features, at scale, can be a valuable tool for better understanding the affordances and types of shared information spaces created in many platforms and many activities, and the forms of social interaction that emerge in conjunction with such environments. It can help us place the phenomena we observe across a range of social media platforms into a common

framework that orients around a coherent set of operationalizable CSCW concepts. For example, it allows us to understand the strong community/subreddit orientation of Reddit (well-defined site of work) within the same framework as the appeal of Snapchat's ephemerality (lack of the record of activity removes some of the visibility and accountability). The same framework would be helpful for understanding location-based and geographically-scoped communities like now-defunct Yik Yak and still existing dating apps like Tinder, where geographically-defined site work is counterweighted by easily-accessible, but anonymized record of activity.

The two axes of my conceptual framework, and especially the explicitly shared site of work, which can be delineated by geography or by the distinct social spaces within the platforms, relate the concept of co-situation—inhabiting the same online space and situation while being in different physical spaces (Blackwell, Birnholtz, & Abbot, 2015). The idea of intersecting spaces—physical and digital—is closely related to various mechanisms of visibility, again building on classical CSCW concerns with peripheral awareness within the shared information space. Visibility is important for the actors in these social media environments, who want to see and be seen: represent their views, identity, and preferences and find suitable partners for their online conversations and other activities. Or in the case of those affected in disasters arising from a natural hazards, to find reliable information sources and relate their on-the-ground experiences. Moreover, visibility is also important for the researchers who strive to find meaningful signal in the social media data deluge, especially in the high-tempo, high-volume, noisy and complex context of disaster.

This focus on whose voices get heard, what messages dominate the narrative has been long standing in crisis informatics, where the researchers strive to privilege the lived experience and voices of those affected in an attempt to understand their decision-making processes and behavior (Hughes et al., 2008; Liu et al., 2008; Sutton, Palen, & Shklovski, 2008; Anderson et al., 2016). In the context of social media, where often the loudest voices get amplified, the issue of who gets heard and what messages spread through the network—issues of visibility—are closely intertwined with the rise of attention economy, the main challenge of which has been articulated by Herbert Simon long before its apogee:

“...in an information-rich world, the wealth of information means a dearth of something else: a scarcity of ... the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it”

(Simon, 1971).

In the current climate of misinformation and fake news, the issues of attention economy are becoming even more prominent. Thus, a conceptual framework that can highlight various mechanisms of visibility would be useful for understanding how the affordances of the platforms and their activities relate to the rise in prominence of particular types of messages or actors. Moreover, drawing the connections between the aspects of affordances and the forms of sociality that tend to arise in those environments may help us detect and describe the behavioral patterns of malicious actors: we could potentially detect their behavioral signatures or network motifs suggestive of their activity.

The tools I have developed for studying the online social activity in aftermath of disasters arising from the natural hazards, described in this dissertation, can be used to understand the social dynamics in other types of large-scale disruptions. While we find misinformation in the context of disasters arising from natural hazards, most of it is unintentional, as in propagating unverified information or false rumors, and gets corrected by the convergent crowd quite rapidly (Mendoza, Poblete, and Castillo, 2010). However, the situation is often much more dire in other high-tempo, high-volume, and high-stakes contexts such as political unrest, where we often find much more blatant and purposeful misinformation. Thus, the conceptual framing developed in this dissertation, as it focuses on aspects of the shared information space that relate to mechanisms of visibility, would also be helpful for understanding the dynamics of visibility and attention in the context of political disruption. As I plan to study this domain in my future research, I expect to find different temporality of the events, with more prolonged, less clear stages. I also expect to find less opening up to new sources we saw in response to the natural hazards, and instead likely will find closing in of the networks and reliance on the in-group ties, similar to the “turtling up” found in the hedge fund communication network after a price shock (Romero, Uzzi, Kleinberg, 2016). Nevertheless, the conceptual and methodological tools developed in this thesis would be helpful in

bringing into relief and understanding social “structures in process” (Kreps & Bosworth, 1994) that emerge in this environment that is both different and similar to the context of disasters arising from natural hazards.

Overall, I see large-scale disruptions, such as disasters arising from natural hazards and political crises, as both fruitful objects of study in their own right and also as an opportunity to observe how new forms of sociality emerge out of high-tempo, high-volume social media activity. Such large-scale disruptions provide a lens into formation and reconfiguration of social forms and practices, allowing us to observe how people come together, how online communities form, and how some of them succeed and why—the social dynamics at the heart of my research interests. I am ultimately motivated by these types of questions, such as why some communities are better than others at accomplishing their goals or just at being good, supportive communities. What forms of sociality play a role here? Is their network structure (core/periphery, bridges and bottlenecks) related to their success and how? And how do the social roles taken on by participants factor in? What about norms and policies within the communities, and the role of founders in establishing them? I believe the framework developed in this dissertation, which connects classical concepts of CSCW pertaining to the creation of shared information space in the distributed context of social media to these forms of sociality, would be a valuable tool for beginning to answer some of these important questions.

References

- Abel, F., Hauff, C., Houben, G. J., Stronkman, R., & Tao, K. Semantics+ filtering+ search= twitcident. exploring information in social web streams. In *Proc. of the 23rd ACM Conference on Hypertext and social media*. ACM (2012), 285-294.
- Ackerman, M. S. Augmenting organizational memory: a field study of answer garden. *ACM Transactions on Information Systems (TOIS)*, 16(3), (1998), 203-224.
- Alarabi, Ahmed Eldawy, Rami Alghamdi, and Mohamed F. Mokbel. 2014. TAREEG: a MapReduce-based system for extracting spatial data from OpenStreetMap. In *Proceedings of the SIGSPATIAL Conference on Advances in Geographic Information Systems (SIGSPATIAL '14)*, 83-92. DOI=10.1145/2666310.2666403
- Alexander, D. (2002) "From civil defence to civil protection – and back again", *Disaster Prevention and Management: An International Journal*, Vol. 11 Issue: 3, pp.209-213, <https://doi.org/10.1108/09653560210435803>
- Anderson, B. (2006). *Imagined communities: Reflections on the origin and spread of nationalism*. Verso Books.
- Anderson, J., Kogan, M., Bica, M., Palen, L., Anderson, K., Morss, R., Demuth, J., Lazrus, H., Wilhelmi, O., & Henderson, J. Far Far Away in Far Rockaway: Responses to Risks and Impacts during Hurricane Sandy through First-Person Social Media Narratives. To appear at ISCRAM, 2016.
- Anderson, J., Soden, R., Anderson, K., Kogan, M. & Palen, L.. 2015. EPIC-OSM: A Software Framework for OpenStreetMap Data Analytics. To Appear in *Proceedings of the Hawaii international conference on system sciences (HICSS-49)*.
- Anderson, K. M., & Schram, A. Design and implementation of a data analytics infrastructure in support of crisis informatics research (NIER track). In *Proc. of the 33rd International Conference on Software Engineering*. ACM (2011), 844-847.
- Andrews, C., Fichet, E., Ding, Y., Spiro, E. S., & Starbird, K. (2016). Keeping Up with the Tweet-dashians: The Impact of “Official” Accounts on Online Rumoring. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (pp. 452–465). New York, NY, USA: ACM. <https://doi.org/10.1145/2818048.2819986>
- Aragon, C., Bayer, J., Echenique, A., Fiore-Gartland, B., Huang, Y., Hutto, C., Kim, J., Neff, G., Xing, W. Developing a Research Agenda for Human-Centered Data Science, Workshop, CSCW 2016.
- Arif, A., Robinson, J., Stanek, S., Fichet, E., Townsend, P., Worku, Z. & Starbird, K. 2017. A Closer Look at the Self-Correcting Crowd: Examining Corrections in Online Rumors. In *Proc. of CSCW* (pp. 155-168). ACM.
- Avvenuti, M., Cresci, S., Marchetti, A., Meletti, C., & Tesconi, M. (2014). EARS (Earthquake Alert and Report System): A Real Time Decision Support System for Earthquake Crisis Management. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1749–1758). New York, NY, USA: ACM. <https://doi.org/10.1145/2623330.2623358>
- B. Ball and M. E. J. Newman, “Friendship networks and social status,” *Network Science*, vol. 1, no. 01, 16–30, 2013.
- Bak, JY., Kim, S., & Oh, A. 2012. Self-disclosure and relationship strength in Twitter conversations. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2* (pp. 60-64). Association for Computational Linguistics.
- Bakshy, E., Hofman, J., Mason, W., and Watts, D. 2011. Everyone's an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining (WSDM '11)*. ACM, New York, NY, USA, 65-74.
- Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. The role of social networks in information diffusion. In *Proc. of the 21st international conference on World Wide Web*, (2012), 519-528.
- Beggs, J. J., Haines, V. A., & Hurlbert, J. S. (1996). Revisiting the rural–urban contrast: Personal networks in nonmetropolitan and metropolitan settings. *Rural sociology*, 61(2), 306-325.

- Bertelsen, O. W., & Bødker, S. (2001). Cooperation in massively distributed information spaces. In *ECSCW 2001* (pp. 1-17). Springer Netherlands.
- Blackwell, C., Birnholtz, J., & Abbott, C. (2015). Seeing and being seen: Co-situation and impression formation using Grindr, a location-aware gay dating app. *New media & society*, *17*(7), 1117-1136.
- Blake, E. S., Kimberlain, T. B., Berg, R. J., Cangialosi, J. P., & Beven II, J. L. Tropical Cyclone Report: Hurricane Sandy. *National Hurricane Center*, *12*, (2013).
- Bliss, C.A., Kloumann, I. M. , Harris, K. D., Danforth, C. M. & Dodds, P.S. (2012). Twitter reciprocal reply networks exhibit assortativity with respect to happiness. *Journal of Computational Science*, *3*(5), 388-397.
- Boehmer, E. (2010, July 22). Coordinating Efforts by Volunteer and Technical Communities for Disaster Preparedness, Response, and Relief. Science and Technology Innovation Program - Woodrow Wilson International Center for Scholars.
- Boersma, F. K., Groenewegen, P., & Wagenaar, P. Emergency Response Rooms in Action: an ethnographic case-study in Amsterdam. In *Proc. of ISCRAM* (2009), 1-8.
- Bostock, M., Ogievetsky, V. & Heer, J.. 2011. D3: Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* *17*, 12: 2301-2309. <http://dx.doi.org/10.1109/TVCG.2011.185>
- boyd, d.m. (2010). Social network sites as networked publics: Affordances, dynamics and implications. In Z. Papacharissi (Ed.), *Networked self: Identity, community, and culture on social network sites* (pp. 39-58). New York, NY: Routledge.
- Brennan, S., Sadilek, A., & Kautz, H. (2013). Towards Understanding Global Spread of Disease from Everyday Interpersonal Interactions. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence* (pp. 2783–2789). Beijing, China: AAAI Press. <http://dl.acm.org/citation.cfm?id=2540128.2540530>
- Bruns, A., and Burgess J. E. (2011). The use of Twitter hashtags in the formation of ad hoc publics. In *6th European Consortium for Political Research General Conference*, 25-27 August 2011, University of Iceland, Reykjavik.
- Bryant, S. L., Forte, A., & Bruckman, A. (2005). Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work* (pp. 1-10). ACM.
- Budhathoki, N.R. & Haythornthwaite, C. 2012. Motivation for Open Collaboration: Crowd and Community Models and the Case of OpenStreetMap. In *American Behavioral Scientist*, *57*, 5: 548–575. DOI=10.1177/0002764212469364.
- Burke, M., Kraut, R., & Marlow, C. (2011). Social capital on Facebook: Differentiating uses and users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 571-580). ACM.
- Burke, M., Marlow, C., & Lento, T. (2010). Social network activity and social well-being. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1909-1912). ACM.
- Cameron, M. A., Power, R., Robinson, B., & Yin, J. (2012). Emergency Situation Awareness from Twitter for Crisis Management. In *Proceedings of the 21st International Conference Companion on World Wide Web* (pp. 695–698). New York, NY: ACM Press. <https://doi.org/10.1145/2187980.2188183>
- Caragea, C., Squicciarini, A., Stehle, S., Neppalli, K., & Tapia, A. Mapping Moods: Geo-Mapped Sentiment Analysis During Hurricane Sandy. In *Proc. of ISCRAM* (2014).
- Careem, M., De Silva, C., De Silva, R., Raschid, L., & Weerawarana, S. (2006). Sahana: Overview of a Disaster Management System. In *Proceedings of the International Conference on Information and Automation* (pp. 361–366). IEEE Computer Society.
- Celli, F. & Rossi, L. 2012. The role of emotional stability in Twitter conversations. In *Proceedings of the workshop on semantic analysis in social media* (pp. 10-17). Association for Computational Linguistics.
- Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American journal of Sociology*, *113*(3), 702-734.

- Cha, M., Haddadi, H., Benevenuto, F. & Gummadi, K. P. Measuring User Influence in Twitter: The Million Follower Fallacy. In *Proceedings of fourth international conference of Weblogs and Social Media (ICWSM '10)*. pp 10-17 (2010): 30.
- Chen, L., Hossain, K. S. M. T., Butler, P., Ramakrishnan, N., & Prakash, B. A. (2016). Syndromic Surveillance of Flu on Twitter Using Weakly Supervised Temporal Topic Models. *Data Mining and Knowledge Discovery*, 30(3), 681–710. <https://doi.org/10.1007/s10618-015-0434-x>
- Clauset, A., & Eagle, N. 2012. Persistence and periodicity in a dynamic proximity network. arXiv:1211.7343.
- Clauset, A., Shalizi, C. R., & Newman, M. E. Power-law distributions in empirical data. *SIAM review*, 51(4), (2009), 661-703.
- Comfort, L. K. (1999). *Shared risk: Complex systems in seismic response*. Pergamon.
- Corcoran, P., Mooney, P. & Bertolotto, M. 2013. Analysing the growth of OpenStreetMap networks. In *Spatial Statistics* 3: 21–32. DOI=10.1016/j.spasta.2013.01.002
- Corvey, W. J., Verma, S., Vieweg, S., Palmer, M., & Martin, J. H. (2012). Foundations of a Multilayer Annotation Framework for Twitter Communications during Crisis Events. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*. Istanbul, Turkey. http://www.lrec-conf.org/proceedings/lrec2012/pdf/1008_Paper.pdf
- Coviello, L., Sohn, Y., Kramer, A. D., Marlow, C., Franceschetti, M., Christakis, N. A., & Fowler, J. H. Detecting Emotional Contagion in Massive Social Networks. *PloS one*, (2014), 9(3), e90315.
- Curry, P., De Silva, C., & Van de Walle, B. (2007). Open Source Software for Disaster Management. *Communications of the ACM*, 50(3), 61. <https://doi.org/10.1145/1226736.1226768>
- Dashti, S., Palen, L., Heris, M. P., Anderson, K. M., Anderson, J., & Anderson, S. (2014). Supporting Disaster Reconnaissance with Social Media Data: A Design-Oriented Case Study of the 2013 Colorado Floods. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM 2014)*. University Park, PA. http://idl.iscram.org/files/dashti/2014/423_Dashti_et al2014.pdf
- De Choudhury, M., Mason, W. A., Hofman, J. M., & Watts, D. J. Inferring relevant social networks from interpersonal communication. In *Proc. of WWW*, (2010), 301-310.
- De Choudhury, M., Monroy-Hernández, A. & Mark. G. 2014. "Narco" emotions: affect and desensitization in social media during the mexican drug war. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3563-3572. <http://dx.doi.org/10.1145/2556288.2557197>
- Dittus, M., Quattrone, G., & Capra, L. (2016). Analysing Volunteer Engagement in Humanitarian Mapping: Building Contributor Communities at Large Scale. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (pp. 108–118). New York, NY, USA: ACM. <https://doi.org/10.1145/2818048.2819939>
- Dourish, P., & Bellotti, V. (1992). Awareness and coordination in shared workspaces. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work* (pp. 107-114). ACM.
- Dynes, R. R. (1970). *Organized behavior in disaster*. Heath LexingtonBooks.
- Earle, P. S., Bowden, D. C., & Guy, M. (2012). Twitter Earthquake Detection: Earthquake Monitoring in a Social World. *Annals of Geophysics*, 54(6). <https://doi.org/10.4401/ag-5364>
- Ellison, N. B., & boyd, D. M. (2013). Sociality through social network sites. *The Oxford handbook of internet studies*, 151.
- Elwood, S. 2008. Volunteered Geographic Information: Future Research Directions Motivated by Critical, Participatory, and Feminist GIS. *GeoJournal* 72, 3&4: 173–183. DOI=10.1007/s10708-008-9186-0
- Fechner, T., Wilhelm, D. & Kray, C. 2015. Ethermap: Real-time Collaborative Map Editing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '15)*, 3583-3592. <http://doi.acm.org/10.1145/2702123.2702536>

- FEMA (July 1, 2013). Hurricane Sandy FEMA After Action Report, http://www.fema.gov/media-library-data/20130726-1923-25045-7442/sandy_fema_aar.pdf
- Fingerman, K. L. (2009). Consequential strangers and peripheral ties: The importance of unimportant relationships. *Journal of Family Theory & Review* 1.2:69-86.
- Fischer III, H. W. (1998). Response to disaster: Fact versus fiction & its perpetuation: The sociology of disaster. University Press of America.
- Foot, K., & Schneider, S. M. (2004). Online Structures for Citizen Engagement in the September 11th Web Sphere. *Electronic Journal of Communication*, 14(3 & 4).
- Foot, K., Warnick, B., & Schneider, S. M. (2005). Web-Based Memorializing After September 11: Toward a Conceptual Framework. *Journal of Computer-Mediated Communication*, 11(1), 72–96. <https://doi.org/10.1111/j.1083-6101.2006.tb00304.x>
- Fritz, C. (1961). Disaster. In Merton, R. & Nisbet, R. (eds). *Contemporary Social Problems*, 651-694. New York: Harcourt.
- Fritz and J. H. Mathewson. 1957. Convergence Behavior in Disasters: A Problem in Social Control. Committee on Disaster Studies, Disaster Study No. 9. Washington, DC, National Academy of Sciences, National Research Council.
- Geertz, C. (1973) Thick Description: Toward an Interpretive Theory of Culture, In *The Interpretation of Cultures: Selected Essays*, New York: Basic Books, 3-30.
- Geiger, R. S., & Halfaker, A. (2013). Using edit sessions to measure participation in Wikipedia. In *Proceedings of the 2013 conference on Computer supported cooperative work* (pp. 861-870). ACM.
- Gibbs, L., & Holloway, C. Hurricane Sandy After Action: Report and Recommendations to Mayor Michael R. Bloomberg. *Hurricane Sandy After Action: Report and Recommendations to Mayor Michael R. Bloomberg, The City of New York, New York, NY*, 36, (2013).
- Giddens, A. (1984). The constitution of society: Outline of the theory of structuration. Univ of California Press.
- Goffman, E. (2002). The presentation of self in everyday life. 1959. *Garden City, NY*.
- Goldstein. (Ed.). Collaborative resilience: moving through crisis to opportunity. MIT Press. (2012).
- Gonçalves, B., Perra, N. & Vespignani, A. 2011. Modeling users' activity on twitter networks: Validation of dunbar's number. *PloS one*, 6(8), e22656.
- Goodchild, M. F. (2007). Citizens as Sensors: The World of Volunteered Geography. *GeoJournal*, 69(4), 211–221. <https://doi.org/10.1007/s10708-007-9111-y>
- Goodchild, M. F., & Glennon, J. A. (2010). Crowdsourcing Geographic Information for Disaster Response: A Research Frontier. *International Journal of Digital Earth*, 3, 231–241. <https://doi.org/10.1080/17538941003759255>
- Granovetter, M. (1983). The strength of weak ties: A network theory revisited. *Sociological theory* 1.1:201-233.
- Gray, R., Ellison, N., Vitak, J., and Lampe, C. (2013). Who Wants to Know? Question-Asking and Answering Practices Among Facebook Users. In *Proc. CSCW, 1213-1224. New York: ACM*.
- Gupta, A., & Kumaraguru, P. (2012). Credibility Ranking of Tweets during High Impact Events. In *Proceedings of the 1st Workshop on Privacy and Security in Online Social Media* (p. 2:2–2:8). New York, NY: ACM Press. <https://doi.org/10.1145/2185354.2185356>
- Hagar, C. and C. Haythornthwaite (2005). Crisis, Farming & Community. *The Journal of Community Informatics* 1, 3:41–52.
- Hagberg, A., Swart, P., and Schult, D. 2008. Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the Python in Science Conference (SciPy 2008)*, 11-15.
- Haklay, M. 2010. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and Planning B: Planning and Design* 37: 682-703. DOI=10.1068/b35097
- Hall, S., & du Gay, P. (Eds.). (1996). *Questions of Cultural Identity: SAGE Publications*. Sage.

- Hannak, A., Margolin, D., Keegan, B., & Weber, I. 2014. Get Back! You Don't Know Me Like That: The Social Mediation of Fact Checking Interventions in Twitter Conversations. In *ICWSM*.
- Hawelka, B., Sitko, I., Beinat, E., Sobolevsky, S., Kazakopoulos, P. & Ratti, C. 2014. Geo-located Twitter as proxy for global mobility patterns. *Cartography and Geographic Information Science*, 41(3), 260-271.
- Haythornthwaite, C. (2005). Social networks and Internet connectivity effects. *Information, Community & Society*, 8(2), 125-147.
- Heath, C. & Luff, P. 1992. Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *Journal of Computer Supported Cooperative Work*, 1,1: 24-48. DOI=10.1007/BF00752451
- Heipke, C. (2010). Crowdsourcing Geospatial Data. *Journal of Photogrammetry and Remote Sensing*, 65(6), 550–557. <https://doi.org/10.1016/j.isprsjprs.2010.06.005>
- Hiltz, S. R., Kushma, J., & Plotnick, L. Use of Social Media by US Public Sector Emergency Managers: Barriers and Wish Lists. In *Proceedings of ISCRAM* (2014).
- Hollan, J., E. Hutchins, and D. Kirsh (2000). Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. *ACM Transactions on Computer Human Interaction*, Vol. 7, No. 2, June 2000, p. 174-196.
- Holland, P. W., and Leinhardt, S.. "A method for detecting structure in sociometric data." *American Journal of Sociology* (1970): 492-513.
- Hristova, D., Quattrone, G., Mashhadi, A. & Capra, L. 2013. The Life of the Party: Impact of Social Mapping in OpenStreetMap. In *Proceedings of AAAI Conference on Weblogs and Social Media (ICWSM 2013)*, 234-243.
- Hughes, A. L., & Palen, L. Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, 6(3), (2009), 248-260.
- Hughes, A. L., Denis, L. A. S., Palen, L., & Anderson, K. M. Online Public Communications by Police & Fire Services during the 2012 Hurricane Sandy. In *Proc. of CHI* (2014), 1505-1514.
- Hughes, A. L., Palen, L., Sutton, J., Liu, S. B., & Vieweg, S. 2008. "Site-Seeing" in Disaster: An Examination of On-Line Social Convergence. In *Proceedings of Conference on Information Systems for Crisis Response and Management (ISCRAM '08)*, 324-333.
- Imran, M., Castillo, C., Lucas, J., Patrick, M., & Rogstadius, J. Coordinating human and machine intelligence to classify microblog communications in crises. *Proc. of ISCRAM* (2014).
- Intagorn, S., & Lerman, K. (2011). Mining Geospatial Knowledge on the Social Web. *International Journal of Information Systems for Crisis Response and Management*, 3(2), 33–47. <https://doi.org/10.4018/jiscrm.2011040103>
- Ireson, N. (2009). Local Community Situational Awareness during an Emergency. In *Proceedings of the 3rd IEEE International Conference on Digital Ecosystems and Technologies (DEST 2009)* (pp. 49–54). Washington, DC: IEEE Computer Society. <https://doi.org/10.1109/DEST.2009.5276763>
- Jenkins, R. (2014). *Pierre Bourdieu*. Routledge.
- Johnson, D., Zagorecki, A., Gelman, J. M., & Comfort, L. K. (2011). Improved Situational Awareness in Emergency Management through Automated Data Analysis and Modeling. *Journal of Homeland Security and Emergency Management*, 8(1). <https://doi.org/10.2202/1547-7355.1873>
- Keegan, B. (2015). Emergent Social Roles in Wikipedia's Breaking News Collaborations. In E. Bertino & S. A. Matei (Eds.), *Roles, Trust, and Reputation in Social Media Knowledge Markets* (pp. 57–79). Springer International Publishing. https://doi.org/10.1007/978-3-319-05467-4_4
- Keegan, B. Breaking news on wikipedia: dynamics, structures, and roles in high-tempo collaboration. In *Proc. of CSCW Companion*, (2012), 315-318.
- Keegan, B., Gergle, D., & Contractor, N. (2013). Hot Off the Wiki: Structures and Dynamics of Wikipedia's Coverage of Breaking News Events. *American Behavioral Scientist*, 57(5), 595–622. <https://doi.org/10.1177/0002764212469367>

- Keegan, B., Gergle, D., & Contractor, N. 2011. Hot off the wiki: dynamics, practices, and structures in Wikipedia's coverage of the Tōhoku catastrophes. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11)*, 105-113. <http://doi.acm.org/10.1145/2038558.2038577>
- Keegan, B., Gergle, D. & Contractor, N. 2012. Staying in the loop: structure and dynamics of Wikipedia's breaking news collaborations. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration (WikiSym '12)*. Article 1, 10 pages. <http://doi.acm.org/10.1145/2462932.2462934>
- Kendra, J. M., & Wachtendorf, T. (2003). Elements of resilience after the world trade center disaster: reconstituting New York City's Emergency Operations Centre. *Disasters*, 27(1), 37-53.
- Kendra, J.M. & Wachtendorf, T. 2003. Reconsidering Convergence and Converger Legitimacy in Response to the World Trade Center Disaster. *Terrorism and Disaster: New Threats, New Ideas (Research in Social Problems and Public Policy)* 11: 97-122. DOI=10.1016/S0196-1152(03)11007-1
- Kim, S., Bak, J., & Oh, A. H. 2012, May. Do You Feel What I Feel? Social Aspects of Emotions in Twitter Conversations. In *ICWSM*.
- Kittur, A., Chi, E.H., Pendleton, B.A., Suh, B. & Mytkowicz, T. 2007. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. In *Proceedings of alt.CHI 2007*
- Kittur, A., Suh, B., Pendleton, B.A. & Chi, E.H. 2007. He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, 453-462. <http://doi.acm.org/10.1145/1240624.1240698>
- Kogan, M., Anderson, J., Palen, L., Anderson, K., & Soden, R. Finding the Way to OSM Mapping Practices: Bounding Large Crisis Datasets for Qualitative Investigation. To appear at CHI, 2016.
- Kogan, M., Palen, L. & Anderson, K. 2015. Think Local, Retweet Global: Retweeting by the Geographically-Vulnerable during Hurricane Sandy. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*, 981-993. <http://doi.acm.org/10.1145/2675133.2675218>
- Kraut, R. E., & Fiore, A. T. (2014). The role of founders in building online groups CSCW'14: Proceedings of the ACM Conference on Computer Supported Cooperative Work (pp. 722-732). NY: ACM.
- Kreps, G. A., & Bosworth, S. L. (1994). Organizing, role enactment, and disaster: A structural theory. University of Delaware Press.
- Kwak, H., Lee, C., Park, H., & Moon, S. What is Twitter, a social network or a news media?. In *Proc. of WWW*, (2010), 591-600.
- Langlois, G., Elmer, G., McKelvey, F., & Devereaux, Z. (2009). Networked Publics: The Double Articulation of Code and Politics on Facebook. *Canadian Journal of Communication*, 34(3).
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge university press.
- Lazo, J. K., Bostrom, A., Morss, R. E., Demuth, J. L. & Lazrus, H. 2015. Factors affecting hurricane evacuation intentions. *Risk analysis*, 35(10), 1837-1857.
- Lazo, J. K., Morss, R. E., & Demuth, J. L. 300 BILLION SERVED. *Bulletin of the American Meteorological Society*, 90(6), (2009).
- Lazrus, H., Morrow, B.H., Morss, R.E., Lazo, J.K. (2012) Vulnerability beyond Stereotypes: Context and Agency in Hurricane Risk Communication. *Weather, Climate, and Society*, 4, 103–109.
- Leavitt, A., & Clark, J. A. (2014). Upvoting Hurricane Sandy: Event-based News Production Processes on a Social News Site. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1495–1504). New York, NY, USA: ACM. <https://doi.org/10.1145/2556288.2557140>
- Leavitt, A., & Robinson, J. J. (2017). The Role of Information Visibility in Network Gatekeeping: Information Aggregation on Reddit During Crisis Events. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 1246–1261). New York, NY, USA: ACM. <https://doi.org/10.1145/2998181.2998299>

- Lerman, K., & Ghosh, R. Information Contagion: An Empirical Study of the Spread of News on Digg and Twitter Social Networks. In *Proc. of ICWSM, 10*, (2010), 90-97.
- Lin, Y. R., Margolin, D., Keegan, B., Baronchelli, A., & Lazer, D. (2013). # Bigbirds Never Die: Understanding Social Dynamics of Emergent Hashtag. *arXiv preprint arXiv:1303.7144*.
- Lin, Y.W. 2011. A qualitative enquiry into OpenStreetMap making. *New Review of Hypermedia and Multimedia* 17, 1: 53-71. <http://dx.doi.org/10.1080/13614568.2011.552647>
- Liu, S., Palen, L., Sutton, J., Hughes, A. & Vieweg, S. 2008. In Search of the Bigger Picture: The Emergent Role of On-Line Photo-Sharing in Times of Disaster. Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)
- Macias, W., Hilyard, K., & Freimuth, V. (2009). Blog Functions as Risk and Crisis Communication During Hurricane Katrina. *Journal of Computer-Mediated Communication*, 15(1), 1–31. <https://doi.org/10.1111/j.1083-6101.2009.01490.x>
- Mark, G., M. Bagdouri, L. Palen, J. Martin, B. Al-An, and K. Anderson (2012). Blogs as a Collective War Diary. *Proceedings of CSCW 2012, ACM Press* (2012), 37–46.
- Maron, M. 2010. Haiti OpenStreetMap Response. Image. (14 January 2010). Retrieved May 20, 2015 from <http://brainoff.com/weblog/2010/01/14/1518>
- McConchie, A. 2013. Introducing “Map-Gardening”. Retrieved September 15, 2015 from <http://mappingmashups.net/2013/05/25/introducing-map-gardening/>..
- Meier, P. 2015. *Digital Humanitarians: How Big Data is changing the Face of Humanitarian Response*. Taylor & Francis Press.
- Mendonca, D., Beroggi, G. E., & Wallace, W. A. (2001). Decision support for improvisation during emergency response operations. *International journal of emergency management*, 1(1), 30-38.
- Messina, C. (2007). Twitter Hashtags for Emergency Coordination and Disaster Relief. Oct 22, 2007. *Blog in: FactoryCity*.
- Mileti, D. (1999). *Disasters by Design: A Reassessment of Natural Hazards in the United States*. Joseph Henry Press.
- Mintzberg, H. (1983). Structure in fives. *Designing effective organizations*. Englewood Cliffs, NJ: Prentice Hall.
- Mooney, P. & Corcoran, P. 2012. How social is OpenStreetMap? 24–27. In *Proceedings of International Conference on Geographic Information Science* (AGILE 2012), 282-287.
- Mooney, P. & Corcoran, P. 2014. Analysis of Interaction and Co-editing Patterns amongst OpenStreetMap Contributors. *Transactions in GIS* 18, 5: 633–659. DOI= 10.1111/tgis.12051
- Morss, R. E., & Hayden, M. H. Storm Surge and "Certain Death": Interviews with Texas Coastal Residents following Hurricane Ike. *Weather, Climate & Society*, 2(3), (2010), 174-189.
- Morss, R. E., Demuth, J. L., Lazo, J. K., Dickinson, K., Lazrus, H., and Morrow, B. H. (2015) Understanding public hurricane evacuation decisions and responses to forecast and warning messages, submitted to *Weather and Forecasting*.
- Morstatter, F., et al. "Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose." *ICWSM*. 2013.
- Munro, R. (2011). Tracking Epidemics Through Crowdsourcing and Natural Language Processing. Presented at the International Conference of Crisis Mappers (ICCM), Geneva.
- Neis, P. & Zipf, A. 2012. Analyzing the Contributor Activity of a Volunteered Geographic Information Project — The Case of OpenStreetMap. *ISPRS International Journal of Geo-Information* 1, 2: 146-165. DOI=10.3390/ijgi1020146
- Newman, M. E .J., Watts, D. & Strogatz, S. 2002. Random graph models of social networks. *Proceedings of the National Academy of Sciences*, 99(suppl 1), 2566-2572.

- Newman, M. E. J. (2003). Mixing patterns in networks. *Physical Review E*, 67(2), 026126.
- Nguyen, V. A., Boyd-Graber, J., & Resnik, P. (2012). SITS: A hierarchical nonparametric model using speaker identity for topic segmentation in multiparty conversations. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1* (pp. 78-87). Association for Computational Linguistics.
- Norheim-Hagtun, I., & Meier, P. (2010). Crowdsourcing for Crisis Mapping in Haiti. *Innovations: Technology, Governance, Globalization*, 5, 81–89. https://doi.org/10.1162/INOV_a_00046
- Olteanu, A., Vieweg, S., & Castillo, C. (2015). What to Expect When the Unexpected Happens: Social Media Communications Across Crises. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 994–1009). New York, NY, USA: ACM. <https://doi.org/10.1145/2675133.2675242>
- Orlikowski, W. J. (1992). Learning from Notes: Organizational Issues in Groupware Implementation.
- Orlikowski, W. J. (2002). Knowing in practice: Enacting a collective capability in distributed organizing. *Organization science* 13.3:249-273.
- Palen, L. & Anderson, K. (2016). Crisis Informatics—New Data for Extraordinary Times. *Science*, Vol 353, Issue 6296, 15 Jul 2016, pp. 224-225. [10.1126/science.aag2579](https://doi.org/10.1126/science.aag2579).
- Palen, L. and S. Vieweg, S.B. Liu, and A.L. Hughes (2009). Crisis in a Networked World. *Social Science Computing Review* 27, no. 4: 467-480.
- Palen, L. Frontiers in Crisis Informatics. Keynote Address at *ISCRAM*, (2014).
- Palen, L., & Liu, S. B. (2007). Citizen communications in crisis: Anticipating a future of ICT-supported public participation. In *Proceedings of CHI*, 727-736.
- Palen, L., Soden, R., Anderson, T. J., & Barrenechea, M. (2015). Success & Scale in a Data-Producing Organization: The Socio-Technical Evolution of OpenStreetMap in Response to Humanitarian Events. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 4113–4122). New York, NY, USA: ACM. <https://doi.org/10.1145/2702123.2702294>
- Palen, L., Vieweg, S., & Anderson, K. (2010) Supporting “Everyday Analysts” in Safety- and Time-Critical Situations, *The Information Society*, 27:1, 52-62, DOI: 10.1080/01972243.2011.534370
- Pancier, K. 2011. User lifecycles in cyclopath: a survey of users. In *Proceedings of iConference (iConference '11)*, 741-742. <http://doi.acm.org/10.1145/1940761.1940892>
- Paul, M. J. (2012). Mixed membership Markov models for unsupervised conversation modeling. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning* (pp. 94-104). Association for Computational Linguistics.
- Perng, S-Y., Buscher, M., Wood, L., Halvorsrud, R., Stiso, M., Ramirez, L. & Al-Akkad, A. Peripheral response: microblogging during the 22/7/2011 Norway attacks, *International Journal of Information Systems for Crisis Response and Management*, (2013), 41-57
- Petrescu-Prahova, M., & Butts, C. T. (2005). Emergent coordination in the World Trade Center disaster. *Institute for mathematical behavioral sciences*.
- Powell, J. W. (1954). *An introduction to the natural history of disaster*. Univ. of Maryland: Disaster Research Project.
- Priedhorsky, R. & Terveen, L. 2008. The computational geowiki: what, why, and how. In *Proceedings of the ACM conference on Computer supported cooperative work (CSCW '08)*, 267-276. <http://doi.acm.org/10.1145/1460563.1460606>
- Priedhorsky, R., Chen, J., Lam, S. K., Panciera, K., Terveen, L. & Riedl, J. 2007. Creating, destroying, and restoring value in wikipedia. In *Proceedings of the ACM conference on Supporting group work (GROUP '07)*, 259-268. <http://doi.acm.org/10.1145/1316624.1316663>

- Procopio, C., & Procopio, S. (2007). Do You Know What It Means to Miss New Orleans? Internet Communication, Geographic Community, and Social Capital in Crisis. *Journal of Applied Communication Research*, 35(1), 67–87. <https://doi.org/10.1080/00909880601065722>
- Qu, Y., Huang, C., Zhang, P., & Zhang, J. Microblogging after a major disaster in China: a case study of the 2010 Yushu earthquake. In *Proc. of CSCW*, (2011), 25-34.
- Qu, Y., Wu, P. F., & Wang, X. (2009, January). Online community response to major disaster: A study of Tianya forum in the 2008 Sichuan earthquake. In *System Sciences, 2009. HICSS'09. 42nd Hawaii*
- Ritter, A., Cherry, C. & Dolan, B. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics* (pp. 172-180). Association for Computational Linguistics.
- Robinson, S. (2009). “If You Had Been with Us”: Mainstream Press and Citizen Journalists Jockey for Authority over the Collective Memory of Hurricane Katrina. *New Media & Society*, 11(5), 795–814. <https://doi.org/10.1177/1461444809105353>
- Rogers, E. M. *Diffusion of innovations*. Simon and Schuster, (1962).
- Romero, D. M., Meeder, B., & Kleinberg, J. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc. of WWW*, (2011), 695-704.
- Sakaki, T., Okazaki, M., & Matsuo, Y. (2012). Tweet Analysis for Real-Time Event Detection and Earthquake Reporting System Development. *IEEE Transactions on Knowledge and Data Engineering*, 25(4), 919–931. <https://doi.org/10.1109/TKDE.2012.29>
- Sarcevic, A., Palen, L., White, J., Starbird, K., Bagdouri, M., & Anderson, K. (2012). Beacons of hope in decentralized coordination: Learning from on-the-ground medical twitterers during the 2010 Haiti earthquake. In *Proceedings of the ACM 2012 conference on computer supported cooperative work* (pp. 47-56). ACM.
- Schlieder, C., & Yanenko, O. (2010). Spatio-Temporal Proximity and Social Distance: A Confirmation Framework for Social Reporting. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Location Based Social Networks* (pp. 60–67). New York, NY, USA: ACM. <https://doi.org/10.1145/1867699.1867711>
- Schmidt, K., & Bannon, L. (1992). Taking CSCW seriously. *Computer Supported Cooperative Work (CSCW)*, 1(1-2), 7-40.
- Shanley, L. A., Burns, R., Bastian, Z., & Robson, E. S. (2013). Tweeting Up a Storm: The Promise and Perils of Crisis Mapping. *Photogrammetric Engineering & Remote Sensing*, 79(10), 865–879.
- Shelton, T., Poorthuis, A., Graham, M., & Zook, M. "Mapping the data shadows of Hurricane Sandy: Uncovering the sociospatial dimensions of 'big data'." *Geoforum* 52 (2014): 167-179.
- Simmel, G. (1950). *The sociology of Georg Simmel* (Vol. 92892). Simon and Schuster.
- Simmel, G. (1955). Conflict and the web of group affiliations. *New York: The Free Press*.
- Simon, H. A. (1971). Designing organizations for an information-rich world.
- Snijders. 2006. Statistical Methods for Network Dynamics. In *Proceedings of the XLIII Scientific Meeting of the Italian Statistical Society*, 281-296.
- Soden, R., & Palen, L. (2016). Infrastructure in the Wild: What Mapping in Post-Earthquake Nepal Reveals About Infrastructural Emergence. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 2796–2807). New York, NY, USA: ACM. <https://doi.org/10.1145/2858036.2858545>
- Soden, R., & Palen, L. From Crowdsourced Mapping to Community Mapping: The Post-Earthquake Work of OpenStreetMap Haiti. In *11th International Conference on the Design of Cooperative Systems (COOP)*, (2014).
- Starbird, K., Palen, L., Hughes, A. L., & Vieweg, S. (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.

- Starbird, K. & Palen, L. 2013. Working and sustaining the virtual "Disaster Desk". In *Proceedings of the ACM conference on Computer supported cooperative work (CSCW '13)*, 491-502. <http://doi.acm.org/10.1145/2441776.2441832>
- Starbird, K. and L. Palen (2011). "Voluntweeters:?" Self-Organizing by Digital Volunteers in Times of Crisis. In *Proceedings of the 2011 Conference on Human Factors in Computing Systems (CHI 2011)*, 1071-1080. Vancouver, BC: ACM.
- Starbird, K. and L. Palen (2012). (How) Will the Revolution be Retweeted?: Information Propagation in the 2011 Egyptian Uprising. In *Proceedings of the 2012 Conference on Computer Supported Cooperative Work (CSCW 2012)*, TBA. Bellevue, WA: ACM.
- Starbird, K., & Palen, L. (2010). *Pass it on?: Retweeting in mass emergency*. International Community on Information Systems for Crisis Response and Management.
- Starbird, K., Maddock, J., Orand, M., & Mason, R. M. Rumors, False Flags, and Digital Vigilantes: Misinformation on Twitter after the 2013 Boston Marathon Bombing. In *Proc. of iConference*, (2014).
- Starbird, K., Muzny, G., & Palen, L. Learning from the crowd: Collaborative filtering techniques for identifying on-the-ground Twitterers during mass disruptions. In *Proc. of ISCRAM*, (2012), 1-10.
- Starbird, K. 2013. Delivering patients to sacré coeur: collective intelligence in digital volunteer communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 801-810. <http://dx.doi.org/10.1145/2470654.2470769>
- Strano, E., Nicosia, V., Latora, V., Porta, S., & Barthélemy, M. (2012). Elementary processes governing the evolution of road networks. *Scientific reports*, 2.
- Suchman, L. A. (1983). Office procedure as practical action: models of work and system design. *ACM Transactions on Information Systems (TOIS)*, 1(4), 320-328.
- Suchman, L. A. (1987). *Plans and situated actions: the problem of human-machine communication*. Cambridge university press.
- Suh, B., Hong, L., Pirolli, P., & Chi, E. H. Want to be retweeted? Large scale analytics on factors impacting retweet in twitter network. In *IEEE second international Conference on Social computing*, (2010), 177-184.
- Sultanik, E. A., & Fink, C. (2012). Rapid Geotagging and Disambiguation of Social Media Text via an Indexed Gazetteer. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM 2012)*. Vancouver, BC. http://idl.iscram.org/files/sultanik/2012/212_Sultanik+Fink2012.pdf
- Sutton, J, Palen, L. & Shklovski, I. 2008. Backchannels on the Front Lines: Emergent Use of Social Media in the 2007 Southern California Fires. *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)*
- Tapia, A. H., & Moore, K. (2014). Good enough is good enough: Overcoming disaster response organizations' slow social media data adoption. *Computer Supported Cooperative Work (CSCW)*, 23(4-6), 483-512.
- Tapia, A. H., Bajpai, K., Jansen, B. J., & Yen, J. (2011). Seeking the Trustworthy Tweet: Can Microblogged Data Fit the Information Needs of Disaster Response and Humanitarian Relief Organizations. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM 2011)*. Lisbon, Portugal. http://idl.iscram.org/files/tapia/2011/991_Tapia_etal2011.pdf
- Thomson, R., & Ito, N. (2012). Social Responsibility and Sharing Behaviors Online: The Twitter-Sphere's Response to the Fukushima Disaster. *International Journal of Cyber Society and Education*, 5(1), 55-74.
- Torrey, C., Burke, M., Lee, M., Dey, A., Fussell, S., & Kiesler, S. (2007). Connected Giving: Ordinary People Coordinating Disaster Relief on the Internet. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences* (p. 179a). Washington, DC: IEEE Computer Society. <https://doi.org/10.1109/HICSS.2007.144>
- Verma, S., Vieweg, S., Corvey, W., Palen, L., Martin, J. H., Palmer, M., & Anderson, K. M. NLP to the Rescue?: Extracting "Situational Awareness" Tweets During Mass Emergency. In *Proc. of ICWSM*, (2011), 385-392.

- Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010). Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness. In *Proceedings of the ACM 2010 Conference on Computer Human Interaction* (pp. 1079–1088). New York, NY: ACM Press. <https://doi.org/10.1145/1753326.1753486>
- Vieweg, S., Palen, L., Liu S., Hughes, A.L. & Sutton, J. 2008. Collective Intelligence in Disaster: An Examination of the Phenomenon in the Aftermath of the 2007 Virginia Tech Shootings. Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM).
- Vivacqua, A. S., & Borges, M. R. S. (2010). Collective Intelligence for the Design of Emergency Response. In *Proceedings from the 2010 International Conference on Computer Supported Cooperative Work in Design (CSCWD)* (pp. 623–628). <https://doi.org/10.1109/CSCWD.2010.5471897>
- Wachtendorf, K. T., & Quarantelli, E. L. (2003). The Evacuation of Lower Manhattan by Water Transport on September 11: An Unplanned ‘Success.’ *Joint Commission Journal on Quality and Patient Safety*, 29(6), 316-318.
- Walther, J. B., C. T. Carr, S. S. W. Choi, D. C. DeAndrea, J. Kim, S. T. Tong, and B. Van Der Heide (2011). Interaction of Interpersonal, Peer, and Media Influence Sources Online: A Research Agenda for Technology Convergence. In Z. Papacharissi (Ed.), *Networked self: Identity, community, and culture on social network sites* (pp. 39-58). New York, NY: Routledge.
- Wasserman, S., and Faust, K.. "Social Network Analysis: Methods and Applications." *Cambridge University Press*, (1994).
- Weick. 1993. The collapse of sensemaking in organizations: The Mann Gulch disaster. In *Administrative science quarterly*, 628-652.
- Weick. 1995. *Sensemaking in organizations* (Vol. 3). Sage.
- Wellman, B., & Wortley, S. (1990). Different strokes from different folks: Community ties and social support.
- Wellman, B. and K. Hampton (1999). Living Networked On and Offline. *Contemporary Sociology*, 648-654.
- White, J. I., & Palen, L. (2015). Expertise in the wired wild West. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 662-675). ACM.
- White, J. I., Palen, L., & Anderson, K. M. (2014). Digital mobilization in disaster response: the work & self-organization of on-line pet advocates in response to hurricane sandy. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*(pp. 866-876). ACM.
- Wulf, V., Rohde, M., Pipek, V., & Stevens, G. Engaging with practices: design case studies as a research framework in CSCW. In *Proc. of CSCW*, (2011), 505-512.
- Xia, X., Yang, X., Wu, C., Li, S., & Bao, L. (2012). Information Credibility on Twitter in Emergency Situation. In *Proceedings of the 2012 Pacific Asia Conference on Intelligence and Security Informatics* (pp. 45–59). Berlin, Heidelberg: Springer-Verlag. https://doi.org/10.1007/978-3-642-30428-6_4
- Yang, J., & Leskovec, J. (2010). Modeling information diffusion in implicit networks. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on* (pp. 599-608). IEEE.
- Yang, X., Wu, Z., & Li, Y. (2012). Using Internet Reports for Early Estimates of the Final Death Toll of Earthquake-Generated Tsunami: The March 11, 2011, Tohoku, Japan, Earthquake. *Annals of Geophysics*, 54. <https://doi.org/10.4401/ag-5169>
- Yonetani, M., & Morris, T. (2013). Global Estimates 2012: people displaced by disasters. IDMC.
- Zhang, F., Morss, R. E., Sippel, J. A., Beckman, T. K., Clements, N. C., Hampshire, N. L., & Winkley, S. D. An In-Person Survey Investigating Public Perceptions of and Responses to Hurricane Rita Forecasts along the Texas Coast. *Weather & Forecasting*, 22(6), (2007), 1177-1190.
- Zook, M., Graham, M., Shelton, T., & Gorman, S. (2010). Volunteered Geographic Information and Crowdsourcing Disaster Relief: A Case Study of the Haitian Earthquake. *World Medical & Health Policy*, 2(2), 7–33. <https://doi.org/10.2202/1948-4682.1069>